cloudera®

Cloudera Security

Cloudera, Inc.
395 Page Mill Road
Palo Alto, CA 94306
info@cloudera.com
US: 1-888-789-1488
Intl: 1-650-362-0488
www.cloudera.com

Release Information

Version: Cloudera Enterprise 5.14.x
Date: December 18, 2020

Table of Contents

About this Guide

This guide is intended for system administrators who want to secure a cluster using data encryption, user authentication, and authorization techniques. It provides conceptual overviews and how-to information about setting up various Hadoop components for optimal security, including how to setup a gateway to restrict access. This guide assumes that you have basic knowledge of Linux and systems administration practices, in general.

Cloudera Security Overview

As a system designed to support vast amounts and types of data, Cloudera clusters must meet ever-evolving security requirements imposed by regulating agencies, governments, industries, and the general public. Cloudera clusters comprise both Hadoop core and ecosystem components, all of which must be protected from a variety of threats to ensure the confidentiality, integrity, and availability of all the cluster's services and data. This overview provides introductions to:

Security Requirements

Goals for data management systems, such as confidentiality, integrity, and availability, require that the system be secured across several dimensions. These can be characterized in terms of both general operational goals and technical concepts, as shown in the figure below:

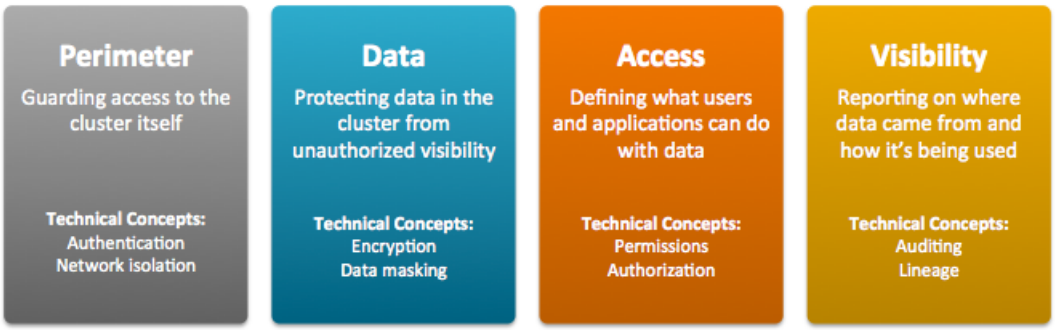

- **Perimeter** Access to the cluster must be protected from a variety of threats coming from internal and external networks and from a variety of actors. Network isolation can be provided by proper configuration of firewalls, routers, subnets, and the proper use of public and private IP addresses, for example. **Authentication** mechanisms ensure that people, processes, and applications properly identify themselves to the cluster and prove they are who they say they are, before gaining access to the cluster.
- **Data** Data in the cluster must always be protected from unauthorized exposure. Similarly, communications between the nodes in the cluster must be protected. **Encryption** mechanisms ensure that even if network packets are intercepted or hard-disk drives are physically removed from the system by bad actors, the contents are not usable.
- **Access** Access to any specific service or item of data within the cluster must be specifically granted. **Authorization** mechanisms ensure that once users have authenticated themselves to the cluster, they can only see the data and use the processes to which they have been granted specific permission.
- **Visibility** Visibility means that the history of data changes is transparent and capable of meeting data governance policies. **Auditing** mechanisms ensure that all actions on data and its lineage—source, changes over time, and so on—are documented as they occur.

Securing the cluster to meet specific organizational goals involves using security features inherent to the Hadoop ecosystem as well as using external security infrastructure. The various security mechanisms can be applied in a range of levels.

Security Levels

The figure below shows the range of security levels that can be implemented for a Cloudera cluster, from non-secure (0) to most secure (3). As the sensitivity and volume of data on the cluster increases, so should the security level you choose for the cluster.

With level 3 security, your Cloudera cluster is ready for full compliance with various industry and regulatory mandates and is ready for audit when necessary. The table below describes the levels in more detail:

Level	Security	Characteristics
0	Non-secure	No security configured. Non-secure clusters should never be used in production environments because they are vulnerable to any and all attacks and exploits.
1	Minimal	Configured for authentication, authorization, and auditing. Authentication is first configured to ensure that users and services can access the cluster only after proving their identities. Next, authorization mechanisms are applied to assign privileges to users and user groups. Auditing procedures keep track of who accesses the cluster (and how).
2	More	Sensitive data is encrypted. Key management systems handle encryption keys. Auditing has been setup for data in metastores. System metadata is reviewed and updated regularly. Ideally, cluster has been setup so that lineage for any data object can be traced (data governance).
3	Most	The secure enterprise data hub (EDH) is one in which all data, both data-at-rest and data-in-transit, is encrypted and the key management system is fault-tolerant. Auditing mechanisms comply with industry, government, and regulatory standards (PCI, HIPAA, NIST, for example), and extend from the EDH to the other systems that integrate with it. Cluster administrators are well-trained, security procedures have been certified by an expert, and the cluster can pass technical review.

Hadoop Security Architecture

The figure below is an example of some of the many components at work in a production Cloudera enterprise cluster. The figure highlights the need to secure clusters that may ingest data from both internal and external data feeds, and across possibly multiple datacenters. Securing the cluster requires applying authentication and access controls throughout these many inter- and intra-connections, as well as to all users who want to query, run jobs, or even view the data held in the cluster.

- External data streams are authenticated by mechanisms in place for Flume and Kafka. Data from legacy databases is ingested using Sqoop. Data scientists and BI analysts can use interfaces such as Hue to work with data on Impala or Hive, for example, to create and submit jobs. Kerberos authentication can be leveraged to protect all these interactions.
- Encryption can be applied to data at-rest using transparent HDFS encryption with an enterprise-grade Key Trustee Server. Cloudera also recommends using Navigator Encrypt to protect data on a cluster associated with the Cloudera Manager, Cloudera Navigator, Hive and HBase metastores, and any log files or spills.
- Authorization policies can be enforced using Sentry (for services such as Hive, Impala, and Search) as well as HDFS Access Control Lists.

- Auditing capabilities can be provided by using Cloudera Navigator.

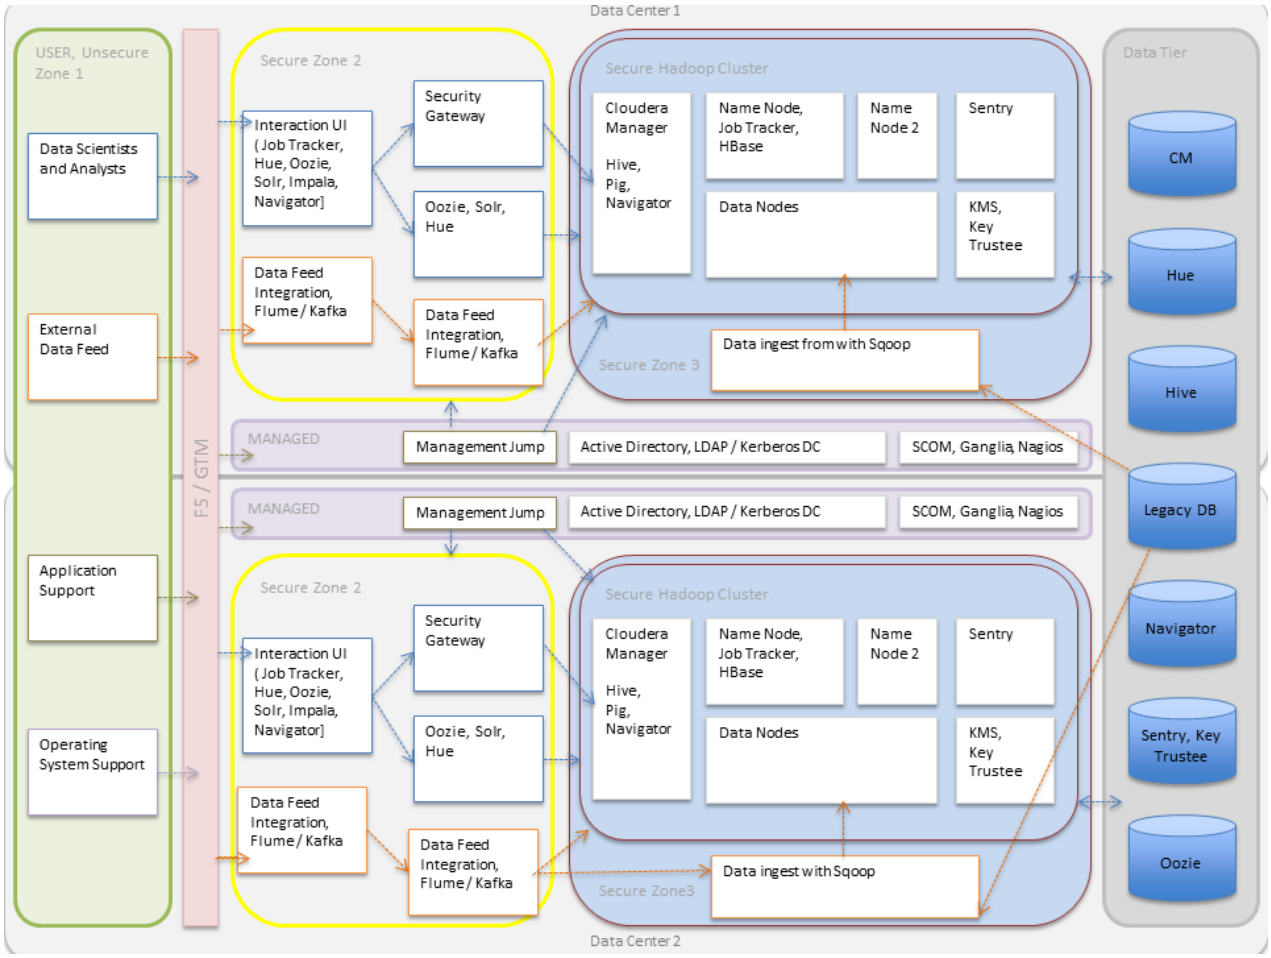

Authentication Overview

Authentication is a basic security requirement for any computing environment. In simple terms, users and services must prove their identity (authenticate) to the system before they can use system features to the degree authorized. Authentication and authorization work hand-in-hand to protect system resources. Authorization is handled in many different ways, from access control lists (ACLs), to HDFS extended ACLs, to role-based access controls (RBAC) using Sentry. See Authorization Overview on page 29 for more information.

Several different mechanisms work together to authenticate users and services in a cluster. These vary depending on the services configured on the cluster. Most CDH components, including Apache Hive, Hue, and Apache Impala can use Kerberos for authentication. Both MIT and Microsoft Active Directory Kerberos implementations can be integrated for use with Cloudera clusters.

In addition, Kerberos credentials can be stored and managed in the LDAP-compliant identity service, such as OpenLDAP and Microsoft Active Directory, a core component of Windows Server.

This section provides a brief overview with special focus on different deployment models available when using Microsoft Active Directory for Kerberos authentication or when integrating MIT Kerberos and Microsoft Active Directory.

Cloudera does not provide a Kerberos implementation. Cloudera clusters can be configured to use Kerberos for authentication, either MIT Kerberos or Microsoft Server Active Directory Kerberos, specifically the Key Distribution Center or KDC. The Kerberos instance must be setup and operational before you can configure the cluster to use it.

Gathering all the configuration details about the KDC—or having the Kerberos administrator available to help during the setup process—is an important preliminary task involved with integrating the cluster and Kerberos regardless of the deployment model.

Kerberos Overview

In simple terms, Kerberos is an authentication protocol that relies on cryptographic mechanisms to handle interactions between a requesting client and server, greatly reducing the risk of impersonation. Passwords are not stored locally nor sent over the network in the clear. The password users enter when logging in to their systems is used to unlock a local mechanism that is then used in a subsequent interaction with a trusted third-party to grant a user a ticket (with a limited lifetime) that is used to authenticate with requested services. After the client and server processes prove their respective identities to each other, communications are encrypted to ensure privacy and data integrity.

The trusted third-party is the Kerberos Key Distribution Center (KDC), the focal point for Kerberos operations which also provides the Authentication Service and the Ticket Granting Service (TGS) for the system. Briefly, the TGS issues a ticket to the requesting user or service which is then presented to the requested service that proves the user (or service) identity for the ticket lifetime (by default, 10 hours). There are many nuances to Kerberos, including defining the principals that identify users and services for the system, ticket renewal, delegated token handling, to name a few. See Kerberos Security Artifacts Overview on page 40.

Furthermore, these processes occur for the most part completely transparently. For example, business users of the cluster simply enter their password when they log in, and the ticket-handling, encryption, and other details take place automatically, behind the scenes. Additionally, users are authenticated not only to a single service target, but to the network as a whole thanks to the tickets and other mechanisms at work in the Kerberos infrastructure.

Kerberos Deployment Models

Credentials needed for Kerberos authentication can be stored and managed in an LDAP-compliant identity/directory service, such as OpenLDAP or Microsoft Active Directory.

At one time a stand-alone service offering from Microsoft, Active Directory services are now packaged as part of the Microsoft Server Domain Services. In the early 2000s, Microsoft replaced its NT LAN Manager authentication mechanism with Kerberos. That means that sites running Microsoft Server can integrate their clusters with Active Directory for Kerberos and have the credentials stored in the LDAP directory on the same server.

This section provides overviews of the different deployment models available for integrating Kerberos authentication with Cloudera clusters, with some of the advantages and disadvantages of the available approaches.

Local MIT KDC

This approach uses an MIT KDC that is local to the cluster. Users and services authenticate to the local KDC before they can interact with the CDH components on the cluster.

Architecture Summary

- An MIT KDC and a separate Kerberos realm is deployed locally to the CDH cluster. The local MIT KDC is typically deployed on a Utility host. Additional replicated MIT KDCs for high-availability are optional.
- All cluster hosts must be configured to use the local MIT Kerberos realm using the `krb5.conf` file.
- All **service and user principals** must be created in the local MIT KDC and Kerberos realm.
- The local MIT KDC will authenticate both the service principals (using keytab files) and user principals (using passwords).
- Cloudera Manager connects to the local MIT KDC to create and manage the principals for the CDH services running on the cluster. To do this Cloudera Manager uses an admin principal and keytab that is created during the setup process. This step has been automated by the Kerberos wizard. See Enabling Kerberos Authentication Using the Wizard for details, or see How to Configure Clusters to Use Kerberos for Authentication for information about creating an admin principal manually.
- The local MIT KDC administrator typically creates all other user principals. However, the Cloudera Manager Kerberos wizard can create the principals and keytab files automatically.

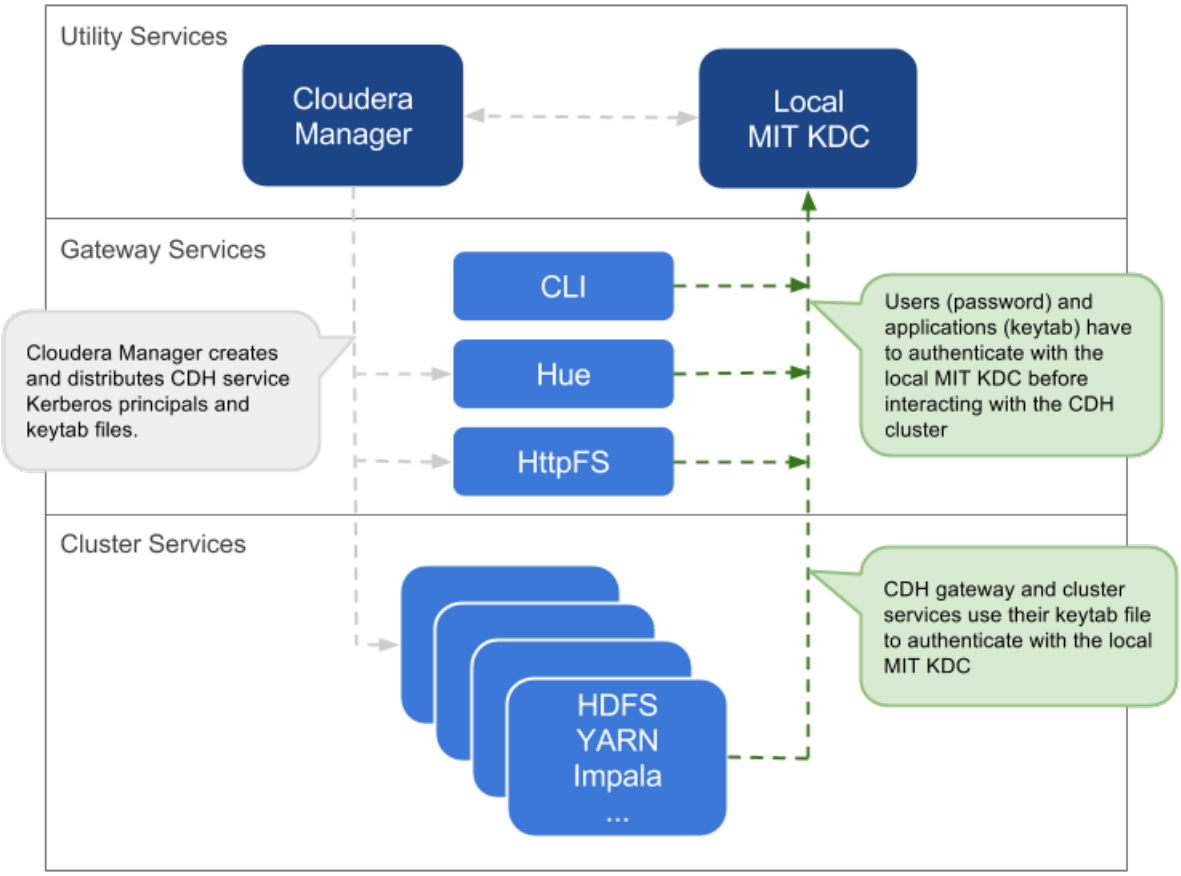

Pros	Cons
The authentication mechanism is isolated from the rest of the enterprise.	This mechanism is not integrated with central authentication system.
This is fairly easy to setup, especially if you use the Cloudera Manager Kerberos wizard that automates creation and distribution of service principals and keytab files.	User and service principals must be created in the local MIT KDC, which can be time-consuming.
	The local MIT KDC can be a single point of failure for the cluster unless replicated KDCs can be configured for high-availability.
	The local MIT KDC is yet another authentication system to manage.

Local MIT KDC with Active Directory Integration

This approach uses an MIT KDC and Kerberos realm that is local to the cluster. However, Active Directory stores the user principals that will access the cluster in a central realm. Users will have to authenticate with this central AD realm to obtain TGTs before they can interact with CDH services on the cluster. Note that CDH service principals reside only in the local KDC realm.

Architecture Summary

- An MIT KDC and a distinct Kerberos realm is deployed locally to the CDH cluster. The local MIT KDC is typically deployed on a Utility host and additional replicated MIT KDCs for high-availability are optional.
- All cluster hosts are configured with both Kerberos realms (local and central AD) using the `krb5.conf` file. The default realm should be the local MIT Kerberos realm.

- **Service principals** should be created in the local MIT KDC and the local Kerberos realm. Cloudera Manager connects to the local MIT KDC to create and manage the principals for the CDH services running on the cluster. To do this, Cloudera Manager uses an admin principal and keytab that is created during the security setup. This step has been automated by the Kerberos wizard.
- A one-way, cross-realm trust must be set up from the local Kerberos realm to the central AD realm containing the **user principals** that require access to the CDH cluster. There is no need to create the service principals in the central AD realm and no need to create user principals in the local realm.

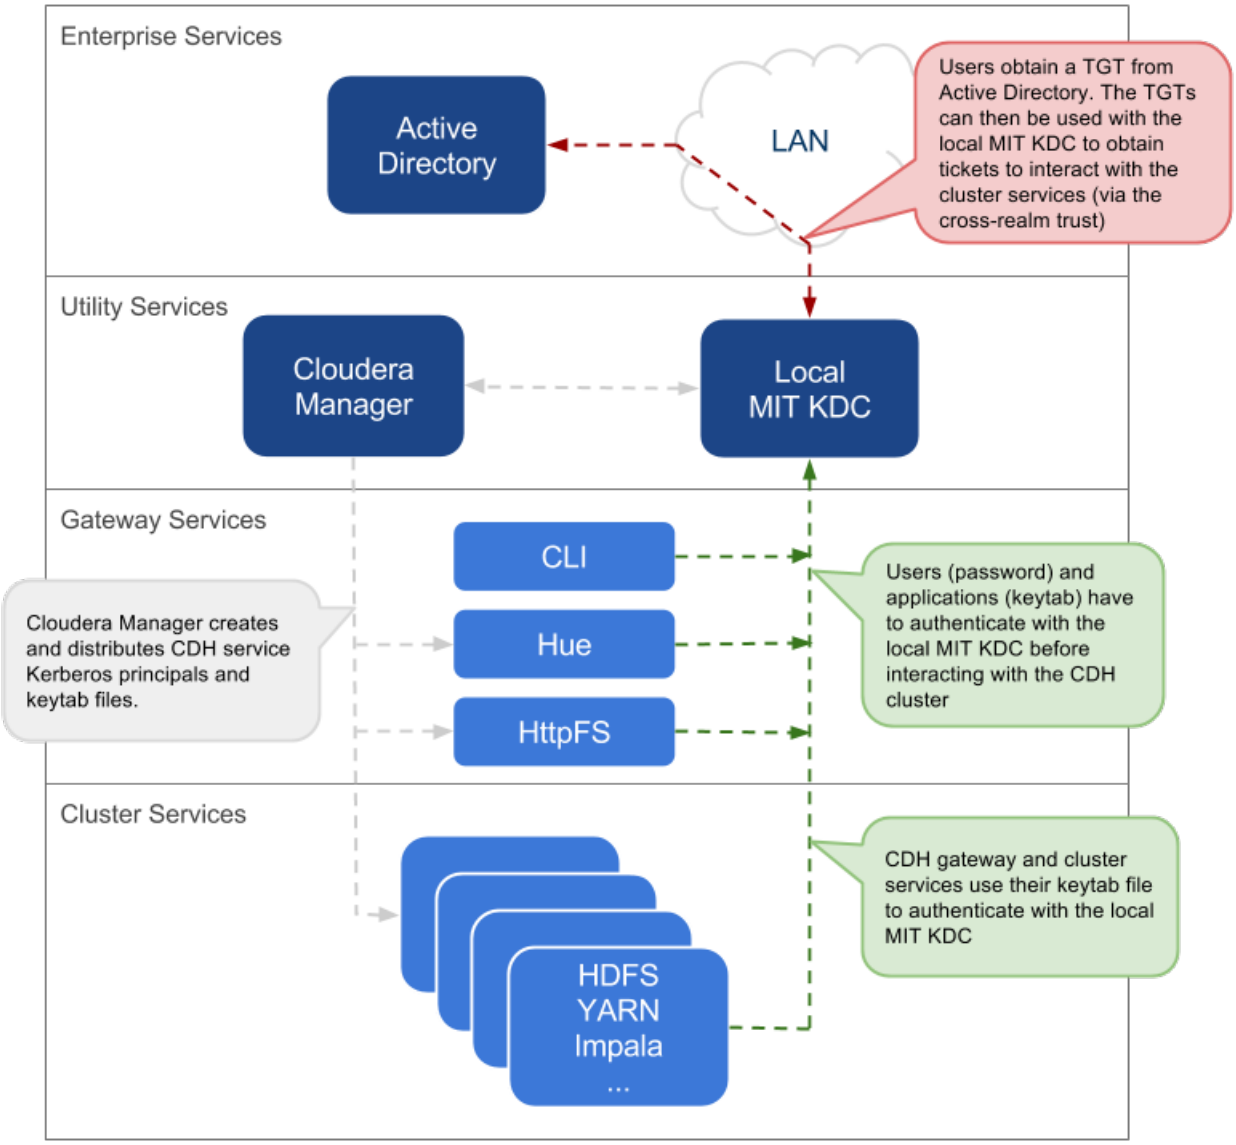

Pros	Cons
The local MIT KDC serves as a shield for the central Active Directory from the many hosts and services in a CDH cluster. Service restarts in a large cluster create many simultaneous authentication requests. If Active Directory is unable to handle the spike in load, then the cluster can effectively cause a distributed denial of service (DDOS) attack.	The local MIT KDC can be a single point of failure (SPOF) for the cluster. Replicated KDCs can be configured for high-availability.

Pros	Cons
This is fairly easy to setup, especially if you use the Cloudera Manager Kerberos wizard that automates creation and distribution of service principals and keytab files. Active Directory administrators will only need to be involved to configure the cross-realm trust during setup.	The local MIT KDC is yet another authentication system to manage.
Integration with central Active Directory for user principal authentication results in a more complete authentication solution.	
Allows for incremental configuration. Hadoop security can be configured and verified using local MIT KDC independently of integrating with Active Directory.	

Using a Centralized Active Directory Service

This approach uses the central Active Directory as the KDC. No local KDC is required. Before you decide upon an AD KDC deployment, make sure you are aware of the following possible ramifications of that decision.

Architecture Summary

- All **service and user principals** are created in the Active Directory KDC.
- All cluster hosts are configured with the central AD Kerberos realm using `krb5.conf`.
- Cloudera Manager connects to the Active Directory KDC to create and manage the principals for the CDH services running on the cluster. To do this, Cloudera Manager uses a principal that has the privileges to create other accounts within the given Organisational Unit (OU). (This step has been automated by the Kerberos wizard.)
- All service and user principals are authenticated by the Active Directory KDC.

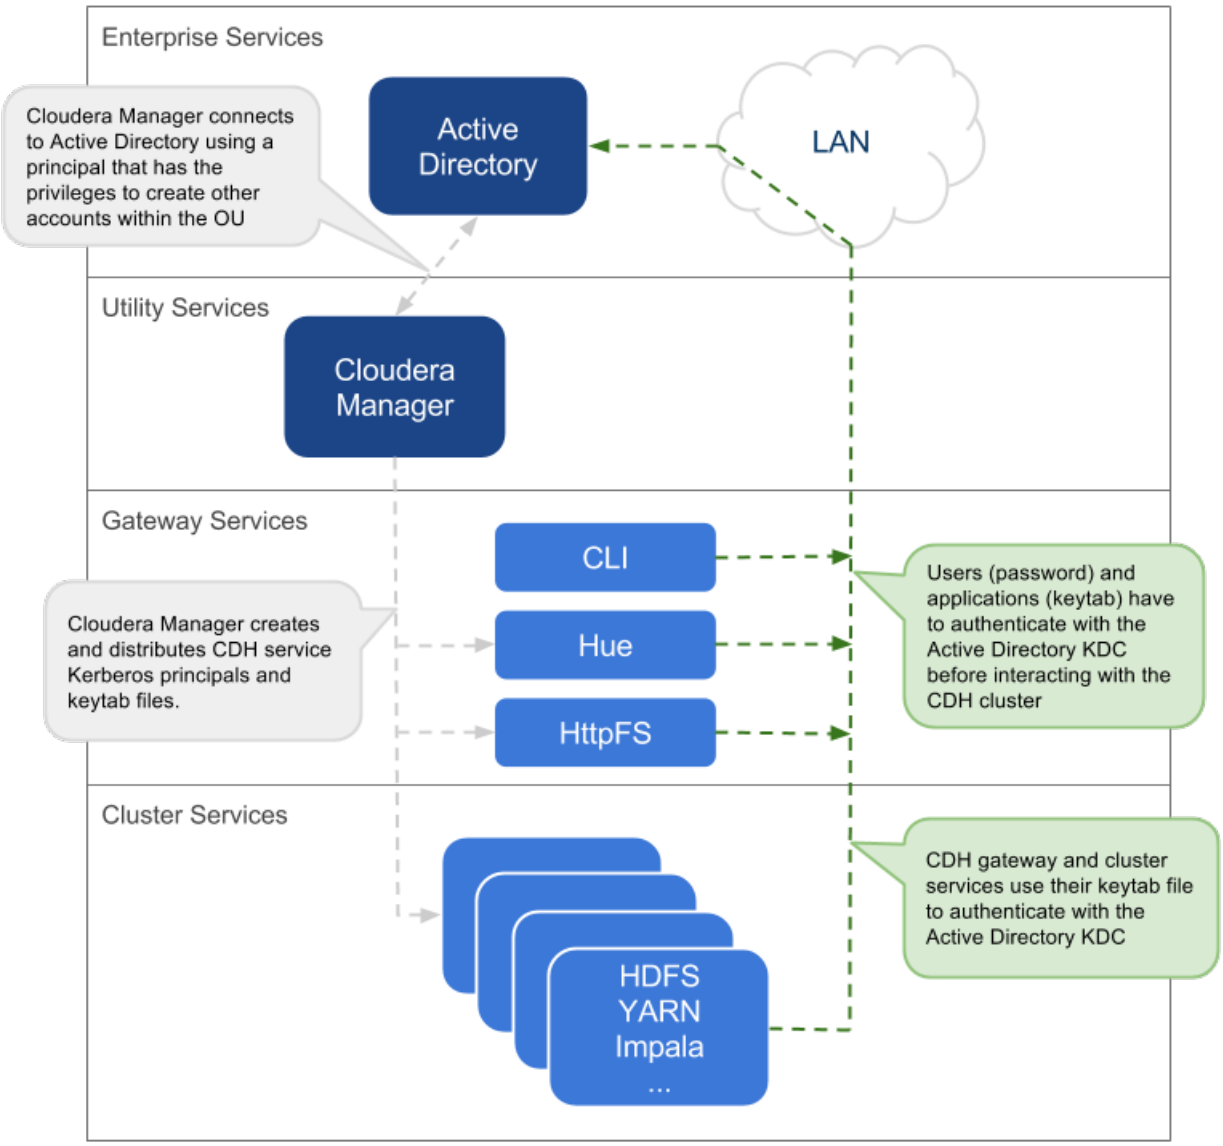

Note: If it is not possible to create the Cloudera Manager admin principal with the required privileges in the Active Directory KDC, then the CDH services principals will need to be created manually. The corresponding keytab files should then be stored securely on the Cloudera Manager Server host. Cloudera Manager's Custom Kerberos Keytab Retrieval script can be used to retrieve the keytab files from the local filesystem.

Recommendations for Active Directory KDC

Several different subsystems are involved in servicing authentication requests, including the Key Distribution Center (KDC), Authentication Service (AS), and Ticket Granting Service (TGS). The more nodes in the cluster and the more services provided, the heavier the traffic between these services and the services running on the cluster.

As a general guideline, Cloudera recommends using a dedicated Active Directory instance (Microsoft Server Domain Services) for every 100 nodes in the cluster. However, this is just a loose guideline. Monitor utilization and deploy additional instances as needed to meet the demand.

By default, Kerberos uses TCP for client/server communication which guarantees delivery but is not as fast at delivering packets as UDP. To override this setting and let Kerberos try UDP before TCP, modify the Kerberos configuration file (`krb5.conf`) as follows:

```
[libdefaults]
udp_preference_limit = 1
...
```

This is especially useful if the domain controllers are not on the same subnet as the cluster or are separated by firewalls.

In general, Cloudera recommends setting up the Active Directory domain controller (Microsoft Server Domain Services) on the same subnet as the cluster and never over a WAN connection. Separating the cluster from the KDC running on the Active Directory domain controller results in considerable latency and affects cluster performance.

Troubleshooting cluster operations when Active Directory is being used for Kerberos authentication requires administrative access to the Microsoft Server Domain Services instance. Administrators may need to enable Kerberos event logging on the Microsoft Server KDC to resolve issues.

Deleting Cloudera Manager roles or nodes requires manually deleting the associate Active Directory accounts. Cloudera Manager cannot delete entries from Active Directory.

Identity Integration with Active Directory

A core requirement for enabling Kerberos security in the platform is that users have accounts on all cluster processing nodes. Commercial products such as Centrify or Quest Authentication Services (QAS) provide integration of all cluster hosts for user and group resolution to Active Directory. These tools support automated Kerberos authentication on login by users to a Linux host with AD. For sites not using Active Directory, or sites wanting to use an open source solution, the Site Security Services Daemon (SSSD) can be used with either AD or OpenLDAP compatible directory services and MIT Kerberos for the same needs.

For third-party providers, you may have to purchase licenses from the respective vendors. This procedure requires some planning as it takes time to procure these licenses and deploy these products on a cluster. Care should be taken to ensure that the identity management product does not associate the service principal names (SPNs) with the host principals when the computers are joined to the AD domain. For example, Centrify by default associates the HTTP SPN with the host principal. So the HTTP SPN should be specifically excluded when the hosts are joined to the domain.

You will also need to complete the following setup tasks in AD:

- **Active Directory Organizational Unit (OU) and OU user** - A separate OU in Active Directory should be created along with an account that has privileges to create additional accounts in that OU.
- **Enable SSL for AD** - Cloudera Manager should be able to connect to AD on the LDAPS (TCP 636) port.
- **Principals and Keytabs** - In a direct-to-AD deployment that is set up using the Kerberos wizard, by default, all required principals and keytabs will be created, deployed and managed by Cloudera Manager. However, if for some reason you cannot allow Cloudera Manager to manage your direct-to-AD deployment, then unique accounts should be manually created in AD for each service running on each host and keytab files must be provided for the same. These accounts should have the AD User Principal Name (UPN) set to `service/fqdn@REALM`, and the Service Principal Name (SPN) set to `service/fqdn`. The principal name in the keytab files should be the UPN of the account. The keytab files should follow the naming convention: `servicename_fqdn.keytab`. The following principals and keytab files must be created for each host they run on: Hadoop Users (user:group) and Kerberos Principals on page 103.
- **AD Bind Account** - Create an AD account that will be used for LDAP bindings in Hue, Cloudera Manager and Cloudera Navigator.
- **AD Groups for Privileged Users** - Create AD groups and add members for the authorized users, HDFS admins and HDFS superuser groups.
  - Authorized users – A group consisting of all users that need access to the cluster
  - HDFS admins – Groups of users that will run HDFS administrative commands

- HDFS super users – Group of users that require superuser privilege, that is, read/wwrite access to all data and directories in HDFS

  Putting regular users into the HDFS superuser group is *not* recommended. Instead, an account that administrators escalate issues to, should be part of the HDFS superuser group.

- **AD Groups for Role-Based Access to Cloudera Manager and Cloudera Navigator** - Create AD groups and add members to these groups so you can later configure role-based access to Cloudera Manager and Cloudera Navigator.
- **AD Test Users and Groups** - At least one existing AD user and the group that the user belongs to should be provided to test whether authorization rules work as expected.

Using TLS/SSL for Secure Keytab Distribution

The Kerberos keytab file is transmitted among the hosts in the Cloudera Manager cluster, between Cloudera Manager Server and Cloudera Manager Agent hosts. To keep this sensitive data secure, configure Cloudera Manager Server and the Cloudera Manager Agent hosts for encrypted communications using TLS/SSL. See Encrypting Data in Transit on page 191 for details.

Using the Wizard or Manual Process to Configure Kerberos Authentication

Cloudera does not provide a Kerberos implementation but uses an existing Kerberos deployment to authenticate services and users. The Kerberos server may be set up exclusively for use by the cluster (for example, Local MIT KDC on page 17) or may be a distributed Kerberos deployment used by other applications in the organization.

Regardless of the deployment model, the Kerberos instance must be operational before the cluster can be configured to use it. In addition, the cluster itself should also be operational and ideally, configured to use TLS/SSL for Cloudera Manager Server and Cloudera Manager Agent hosts, as mentioned above.

When you are ready to integrate the cluster with your organization's MIT KDC or Active Directory KDC, you can do so using the wizard provided in Cloudera Manager Server or by following a manual process, as follows:

- Enabling Kerberos Authentication Using the Wizard on page 51
- How to Configure Clusters for Kerberos Authentication

Authentication Mechanisms used by Cluster Components

Component or Product	Authentication Mechanism Supported
Accumulo	Kerberos (partial)
Backup and Disaster Recovery	Kerberos (used to authenticate Cloudera Manager to Kerberos-protected services), LDAP, SAML
Cloudera Manager	Kerberos (used to authenticate Cloudera Manager to Kerberos-protected services), LDAP, SAML
Cloudera Navigator	Active Directory, OpenLDAP, SAML
Flume	Kerberos (starting CDH 5.4)
HBase	Kerberos, user-based authentication required for HBase Thrift and REST clients
HDFS	Kerberos, SPNEGO (HttpFS)
HiveServer	None
HiveServer2	Kerberos, LDAP, Custom/pluggable authentication
Hive Metastore	Kerberos
Hue	Kerberos, LDAP, SAML, Custom/pluggable authentication
Impala	Kerberos, LDAP, SPNEGO (Impala Web Console)

Component or Product	Authentication Mechanism Supported
Kudu	Kerberos
MapReduce	Kerberos *(also see HDFS)*
Oozie	Kerberos, SPNEGO
Pig	Kerberos
Search	Kerberos, SPNEGO
Sentry	Kerberos
Spark	Kerberos
Sqoop	Kerberos
Sqoop2	Kerberos (as of CDH 5.4)
YARN	Kerberos *(also see HDFS)*
Zookeeper	Kerberos

Encryption Overview

Encryption is a process that uses digital keys to encode various components—text, files, databases, passwords, applications, or network packets, for example—so that only the appropriate entity (user, system process, and so on) can decode (decrypt) the item and view, modify, or add to the data. Cloudera provides encryption mechanisms to protect data persisted to disk or other storage media (***data at rest encryption*** or simply, data encryption) and as it moves over the network (***data in transit encryption***).

Data encryption is mandatory in government, health, finance, education, and many other environments. For example, the Federal Information Security Management Act (FISMA) governs patient privacy concerns and the Payment Card Industry Data Security Standard (PCI DSS) regulates information security for credit-card processors. These are just two examples.

The vast quantity of data contained in Cloudera clusters, deployed using many different components, must nonetheless support whatever degree of privacy, confidentiality, and data integrity is required by the use case. The encryption mechanisms supported by Cloudera and discussed in this overview aim to do just that.

Protecting Data At-Rest

Protecting data at rest typically means encrypting the data when it is stored on disk and letting authorized users and processes—and only authorized users and processes—to decrypt the data when needed for the application or task at hand. With data-at-rest encryption, encryption keys must be distributed and managed, keys should be rotated or changed on a regular basis (to reduce the risk of having keys compromised), and many other factors complicate the process.

However, encrypting data alone may not be sufficient. For example, administrators and others with sufficient privileges may have access to personally identifiable information (PII) in log files, audit data, or SQL queries. Depending on the specific use case—in hospital or financial environment, the PII may need to be redacted from all such files, to ensure that users with privileges on the logs and queries that might contain sensitive data are nonetheless unable to view that data when they should not.

Cloudera provides complementary approaches to encrypting data at rest, and provides mechanisms to mask PII in log files, audit data, and SQL queries.

Encryption Options Available

Cloudera provides several mechanisms to ensure that sensitive data is secure. CDH provides transparent HDFS encryption, ensuring that all sensitive data is encrypted before being stored on disk. HDFS encryption when combined with the enterprise-grade encryption key management of Navigator Key Trustee enables regulatory compliance for most

enterprises. For Cloudera Enterprise, HDFS encryption can be augmented by Navigator Encrypt to secure metadata, in addition to data. Cloudera clusters that use these solutions run as usual and have very low performance impact, given that data nodes are encrypted in parallel. As the cluster grows, encryption grows with it.

Additionally, this transparent encryption is optimized for the Intel chipset for high performance. Intel chipsets include AES-NI co-processors, which provide special capabilities that make encryption workloads run extremely fast. Cloudera leverages the latest Intel advances for even faster performance.

The Key Trustee KMS, used in conjunction with Key Trustee Server and Key HSM, provides HSM-based protection of stored key material. The Key Trustee KMS generates encryption zone key material locally on the KMS and then encrypts this key material using an HSM-generated key. Navigator HSM KMS services, in contrast, rely on the HSM for all encryption zone key generation and storage. When using the Navigator HSM KMS, encryption zone key material originates on the HSM and never leaves the HSM. This allows for the highest level of key isolation, but requires some overhead for network calls to the HSM for key generation, encryption and decryption operations. The Key Trustee KMS remains the recommended key management solution for HDFS encryption for most production scenarios.

The figure below shows an example deployment that uses:

- Cloudera Transparent HDFS Encryption to encrypt data stored on HDFS
- Navigator Encrypt for all other data (including metadata, logs, and spill data) associated with Cloudera Manager, Cloudera Navigator, Hive, and HBase
- Navigator Key Trustee for robust, fault-tolerant key management

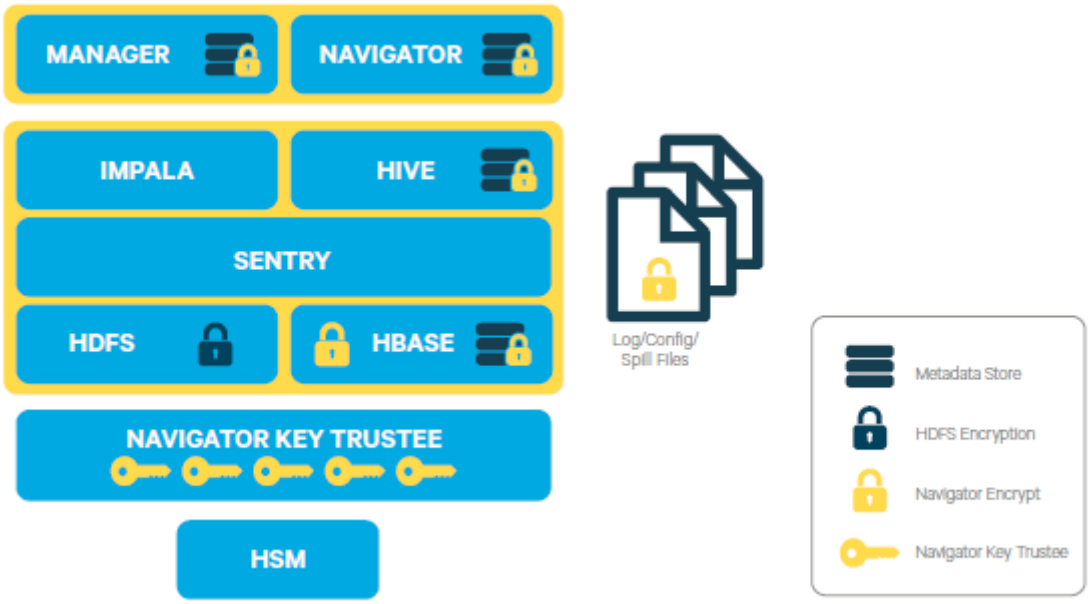

In addition to applying encryption to the data layer of a Cloudera cluster, encryption can also be applied at the network layer, to encrypt communications among nodes of the cluster. See Encryption Mechanisms Overview on page 28 for more information.

Encryption does not prevent administrators with full access to the cluster from viewing sensitive data. To obfuscate sensitive data, including PII, the cluster can be configured for data redaction.

Data Redaction for Cloudera Clusters

Redaction is a process that obscures data. It can help organizations comply with industry regulations and standards, such as PCI (Payment Card Industry) and HIPAA, by obfuscating personally identifiable information (PII) so that is not usable except by those whose jobs require such access. For example, HIPAA legislation requires that patient PII not be available to anyone other than appropriate physician (and the patient), and that any patient's PII cannot be used to determine or associate an individual's identity with health data. Data redaction is one process that can help ensure

this privacy, by transforming PII to meaningless patterns—for example, transforming U.S. social security numbers to `XXX-XX-XXXX` strings.

Data redaction works separately from Cloudera encryption techniques, which do not preclude administrators with full access to the cluster from viewing sensitive user data. It ensures that cluster administrators, data analysts, and others cannot see PII or other sensitive data that is not within their job domain and at the same time, it does not prevent users with appropriate permissions from accessing data to which they have privileges.

See How to Enable Sensitive Data Redaction for details.

Protecting Data In-Transit

For data-in-transit, implementing data protection and encryption is relatively easy. Wire encryption is built into the Hadoop stack, such as SSL, and typically does not require external systems. This data-in-transit encryption is built using session-level, one-time keys, by means of a session handshake with immediate and subsequent transmission. Thus, data-in-transit avoids much of the key management issues associated with data-at-rest due the temporal nature of the keys, but it does rely on proper authentication; a certificate compromise is an issue with authentication, but can compromise wire encryption. As the name implies, data-in-transit covers the secure transfer and intermediate storage of data. This applies to all process-to-process communication, within the same node or between nodes. There are three primary communication channels:

- **HDFS Transparent Encryption:** Data encrypted using HDFS Transparent Encryption is protected end-to-end. Any data written to and from HDFS can only be encrypted or decrypted by the client. HDFS does not have access to the unencrypted data or the encryption keys. This supports both, at-rest encryption as well as in-transit encryption.
- **Data Transfer:** The first channel is data transfer, including the reading and writing of data blocks to HDFS. Hadoop uses a SASL-enabled wrapper around its native direct TCP/IP-based transport, called `DataTransportProtocol`, to secure the I/O streams within an DIGEST-MD5 envelope (For steps, see Configuring Encrypted Transport for HDFS on page 235). This procedure also employs secured HadoopRPC (see Remote Procedure Calls) for the key exchange. The HttpFS REST interface, however, does not provide secure communication between the client and HDFS, only secured authentication using SPNEGO.

  For the transfer of data between DataNodes during the shuffle phase of a MapReduce job (that is, moving intermediate results between the Map and Reduce portions of the job), Hadoop secures the communication channel with HTTP Secure (HTTPS) using Transport Layer Security (TLS). See Encrypted Shuffle and Encrypted Web UIs on page 227.

- **Remote Procedure Calls:** The second channel is system calls to remote procedures (RPC) to the various systems and frameworks within a Hadoop cluster. Like data transfer activities, Hadoop has its own native protocol for RPC, called HadoopRPC, which is used for Hadoop API client communication, intra-Hadoop services communication, as well as monitoring, heartbeats, and other non-data, non-user activity. HadoopRPC is SASL-enabled for secured transport and defaults to Kerberos and DIGEST-MD5 depending on the type of communication and security settings. For steps, see Configuring Encrypted Transport for HDFS on page 235.
- **User Interfaces:** The third channel includes the various web-based user interfaces within a Hadoop cluster. For secured transport, the solution is straightforward; these interfaces employ HTTPS.

TLS/SSL Certificates Overview

Certificates can be signed in one three different ways:

Type	Usage Note
Public CA-signed certificates	**Recommended**. Using certificates signed by a trusted public CA simplifies deployment because the default Java client already trusts most public CAs. Obtain certificates from one of the trusted well-known (public) CAs, such as Symantec and Comodo, as detailed in Generate TLS Certificates on page 195
Internal CA-signed certificates	Obtain certificates from your organization's internal CA if your organization has its own. Using an internal CA can reduce costs (although cluster configuration may require establishing the trust chain for certificates signed by an internal CA, depending on your IT infrastructure).

Type	Usage Note
	See Configuring TLS Encryption for Cloudera Manager on page 195 for information about establishing trust as part of configuring a Cloudera Manager cluster.
Self-signed certificates	Not recommended for production deployments. Using self-signed certificates requires configuring each client to trust the specific certificate (in addition to generating and distributing the certificates). However, self-signed certificates are fine for non-production (testing or proof-of-concept) deployments. See How to Use Self-Signed Certificates for TLS on page 410 for details.

For more information on setting up SSL/TLS certificates, see Encrypting Data in Transit on page 191.

TLS/SSL Encryption for CDH Components

Cloudera recommends securing a cluster using Kerberos authentication before enabling encryption such as SSL on a cluster. If you enable SSL for a cluster that does not already have Kerberos authentication configured, a warning will be displayed.

Hadoop services differ in their use of SSL as follows:

- HDFS, MapReduce, and YARN daemons act as both SSL servers and clients.
- HBase daemons act as SSL servers only.
- Oozie daemons act as SSL servers only.
- Hue acts as an SSL client to all of the above.

Daemons that act as SSL servers load the keystores when starting up. When a client connects to an SSL server daemon, the server transmits the certificate loaded at startup time to the client, which then uses its truststore to validate the server's certificate.

For information on setting up SSL/TLS for CDH services, see Configuring TLS/SSL Encryption for CDH Services on page 204.

Data Protection within Hadoop Projects

The table below lists the various encryption capabilities that can be leveraged by CDH components and Cloudera Manager.

Project	Encryption for Data-in-Transit	Encryption for Data-at-Rest (HDFS Encryption + Navigator Encrypt + Navigator Key Trustee)
HDFS	SASL (RPC), SASL (DataTransferProtocol)	Yes
MapReduce	SASL (RPC), HTTPS (encrypted shuffle)	Yes
YARN	SASL (RPC)	Yes
Accumulo	Partial - Only for RPCs and Web UI (Not directly configurable in Cloudera Manager)	Yes
Flume	TLS (Avro RPC)	Yes
HBase	SASL - For web interfaces, inter-component replication, the HBase shell and the REST, Thrift 1 and Thrift 2 interfaces	Yes
HiveServer2	SASL (Thrift), SASL (JDBC), TLS (JDBC, ODBC)	Yes
Hue	TLS	Yes
Impala	TLS or SASL between impalad and clients, but not between daemons	

Project	Encryption for Data-in-Transit	Encryption for Data-at-Rest (HDFS Encryption + Navigator Encrypt + Navigator Key Trustee)
Oozie	TLS	Yes
Pig	N/A	Yes
Search	TLS	Yes
Sentry	SASL (RPC)	Yes
Spark	None	Yes
Sqoop	Partial - Depends on the RDBMS database driver in use	Yes
Sqoop2	Partial - You can encrypt the JDBC connection depending on the RDBMS database driver	Yes
ZooKeeper	SASL (RPC)	No
Cloudera Manager	TLS - Does not include monitoring	Yes
Cloudera Navigator	TLS - *Also see Cloudera Manager*	Yes
Backup and Disaster Recovery	TLS - *Also see Cloudera Manager*	Yes

Encryption Mechanisms Overview

Data at rest and data in transit encryption function at different technology layers of the cluster:

Layer	Description
Application	Applied by the HDFS client software, HDFS Transparent Encryption lets you encrypt specific folders contained in HDFS. To securely store the required encryption keys, Cloudera recommends using Cloudera Navigator Key Trustee Server in conjunction with HDFS encryption. See Enabling HDFS Encryption Using Navigator Key Trustee Server on page 247 for details. Data stored temporarily on the local filesystem outside HDFS by CDH components (including Impala, MapReduce, YARN, or HBase) can also be encrypted. See Configuring Encryption for Data Spills on page 351 for details.
Operating System	At the Linux OS filesystem layer, encryption can be applied to an entire volume. For example, Cloudera Navigator Encrypt can encrypt data inside and outside HDFS, such as temp/spill files, configuration files, and databases that store metadata associated with a CDH cluster. Cloudera Navigator Encrypt operates as a Linux kernel module, part of the operating system. Navigator Encrypt requires a license for Cloudera Navigator and must be configured to use Navigator Key Trustee Server.
Network	Network communications between client processes and server processes (HTTP, RPC, or TCP/IP services) can be encrypted using industry-standard TLS/SSL as detailed in Encrypting Data in Transit on page 191.

Here are some good starting places for more information about encryption for Cloudera clusters:

- Data at rest encryption
  - Cloudera Navigator Encrypt
  - HDFS Transparent Encryption
  - How to Configure Encryption for Amazon S3

- Data in transit encryption
  - How to Configure TLS Encryption for Cloudera Manager (a complete step-by-step guide)

Authorization Overview

Authorization is one of the fundamental security requirements of any computing environment. Its goal is to ensure that only the appropriate people or processes can access, view, use, control, or change specific resources, services, or data. In any cluster deployed to meet specific workloads using various CDH components (Hive, HDFS, Impala, and so on), different authorization mechanisms can ensure that only authorized users or processes can access data, systems, and other resources as needed. Ideally, authorization mechanisms can leverage the authentication mechanisms, so that when users login to a system—a cluster, for example—they are transparently authorized based on their identity across the system for the applications, data, and other resources they are authorized to use.

For example, Cloudera CDH clusters can be configured to leverage the user and group accounts that exist in the organization's Active Directory (or other LDAP-accessible directory) instance.

The various possible configurations and integrations are discussed later in this guide.

Authorization Mechanisms in Hadoop

Hadoop supports several authorization mechanisms, including:

- Traditional POSIX-style permissions on files and directories. Each directory and file has a single owner and group with basic permissions that can be set to read, write, execute (at the file level). Directories have an additional permission that enables access to child directories.
- Access Control Lists (ACL) for management of services and resources. For example, Apache HBase uses ACLs to authorize various operations (`READ, WRITE, CREATE, ADMIN`) by column, column family, and column family qualifier. HBase ACLs are granted and revoked to users and groups. Fine-grained permissions can be applied to HDFS files using HDFS Extended ACLs on page 175 which enable setting different permissions for specific named users and named groups.
- Role-Based Access Control (RBAC) for certain services with advanced access controls to data. Apache Sentry provides role-based access control for the cluster, including access to Hive, Impala, and Solr services. Cloudera recommends using the database-backed Sentry Service to configure permissions rather than using policy files. However, both approaches are supported. SeeAuthorization With Apache Sentry on page 182 for more information.

POSIX Permissions

Most services running on Hadoop clusters, such as the command-line interface (CLI) or client applications that use Hadoop API, directly access data stored within HDFS. HDFS uses POSIX-style permissions for directories and files; each directory and file is assigned a single owner and group. Each assignment has a basic set of permissions available; file permissions are read, write, and execute, and directories have an additional permission to determine access to child directories.

Ownership and group membership for a given HDFS asset determines a user’s privileges. If a given user fails either of these criteria, they are denied access. For services that may attempt to access more than one file, such as MapReduce, Cloudera Search, and others, data access is determined separately for each file access attempt. File permissions in HDFS are managed by the NameNode.

Access Control Lists

Hadoop also maintains general access controls for the services themselves in addition to the data within each service and in HDFS. Service access control lists (ACL) are typically defined within the global hadoop-policy.xml file and range from NameNode access to client-to-DataNode communication. In the context of MapReduce and YARN, user and group identifiers form the basis for determining permission for job submission or modification.

In addition, with MapReduce and YARN, jobs can be submitted using queues controlled by a scheduler, which is one of the components comprising the resource management capabilities within the cluster. Administrators define permissions to individual queues using ACLs. ACLs can also be defined on a job-by-job basis. Like HDFS permissions,

local user accounts and groups must exist on each executing server, otherwise the queues will be unusable except by superuser accounts.

Apache HBase also uses ACLs for data-level authorization. HBase ACLs authorize various operations (READ, WRITE, CREATE, ADMIN) by column, column family, and column family qualifier. HBase ACLs are granted and revoked to both users and groups. Local user accounts are required for proper authorization, similar to HDFS permissions.

Apache ZooKeeper also maintains ACLs to the information stored within the DataNodes of a ZooKeeper data tree.

Role-Based Access Control with Apache Sentry

For fine-grained access to data accessible using schema—that is, data structures described by the Apache Hive Metastore and used by computing engines like Hive and Impala, as well as collections and indices within Cloudera Search—CDH supports Apache Sentry, which offers a role-based privilege model for this data and its given schema. Apache Sentry is a role-based authorization module for Hadoop. Sentry lets you configure explicit privileges for authenticated users and applications on a Hadoop cluster. Sentry is integrated natively with Apache Hive, Apache Solr, HDFS (for Hive table data), Hive Metastore/HCatalog, and Impala.

Sentry comprises a pluggable authorization engine for Hadoop components. It lets you define authorization rules to validate requests for access to resources from users or applications. The authorization holds the rules and privileges, but it's the engines that apply the rules at runtime.

Architecture Overview

Sentry Components

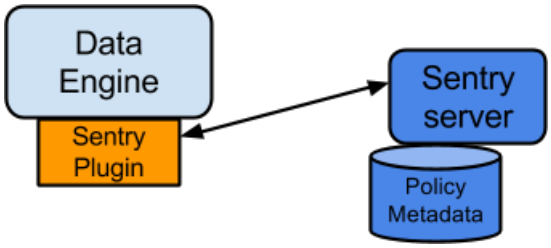

The subsystems nvolved in the authorization process:

- **Sentry Server**—An RPC server that manages authorization metadata stored in a database. It supports interfaces to securely retrieve and manipulate the metadata.
- **Data Engine**

  This is a data processing application such as Hive or Impala that needs to authorize access to data or metadata resources. The data engine loads the Sentry plugin and all client requests for accessing resources are intercepted and routed to the Sentry plugin for validation.

- **Sentry Plugin**

  The Sentry plugin runs in the data engine. It offers interfaces to manipulate authorization metadata stored in the Sentry Server, and includes the authorization policy engine that evaluates access requests using the authorization metadata retrieved from the server.

Key Concepts

- Authentication - Verifying credentials to reliably identify a user
- Authorization - Limiting the user's access to a given resource
- User - Individual identified by underlying authentication system
- Group - A set of users, maintained by the authentication system
- Privilege - An instruction or rule that allows access to an object
- Role - A set of privileges; a template to combine multiple access rules
- Authorization models - Defines the objects to be subject to authorization rules and the granularity of actions allowed. For example, in the SQL model, the objects can be databases or tables, and the actions are `SELECT`,

`INSERT`, and `CREATE`. For the Search model, the objects are indexes, collections and documents; the access modes are query and update.

User Identity and Group Mapping

Sentry relies on underlying authentication systems such as Kerberos or LDAP to identify the user. It also uses the group mapping mechanism configured in Hadoop to ensure that Sentry sees the same group mapping as other components of the Hadoop ecosystem.

Consider users Alice and Bob who belong to an Active Directory (AD) group called `finance-department`. Bob also belongs to a group called `finance-managers`. In Sentry, you first create roles and then grant privileges to these roles. For example, you can create a role called Analyst and grant `SELECT` on tables Customer and Sales to this role.

The next step is to join these authentication entities (users and groups) to authorization entities (roles). This can be done by granting the Analyst role to the `finance-department` group. Now Bob and Alice who are members of the `finance-department` group get `SELECT` privilege to the Customer and Sales tables.

Role-based access control (RBAC) is a powerful mechanism to manage authorization for a large set of users and data objects in a typical enterprise. New data objects get added or removed, users join, move, or leave organisations all the time. RBAC makes managing this a lot easier. Hence, as an extension of the discussed previously, if Carol joins the Finance Department, all you need to do is add her to the `finance-department` group in AD. This will give Carol access to data from the Sales and Customer tables.

Unified Authorization

Another important aspect of Sentry is the unified authorization. The access control rules once defined, work across multiple data access tools. For example, being granted the Analyst role in the previous example will allow Bob, Alice, and others in the `finance-department` group to access table data from SQL engines such as Hive and Impala, as well as using MapReduce, Pig applications or metadata access using HCatalog.

Sentry Integration with the Hadoop Ecosystem

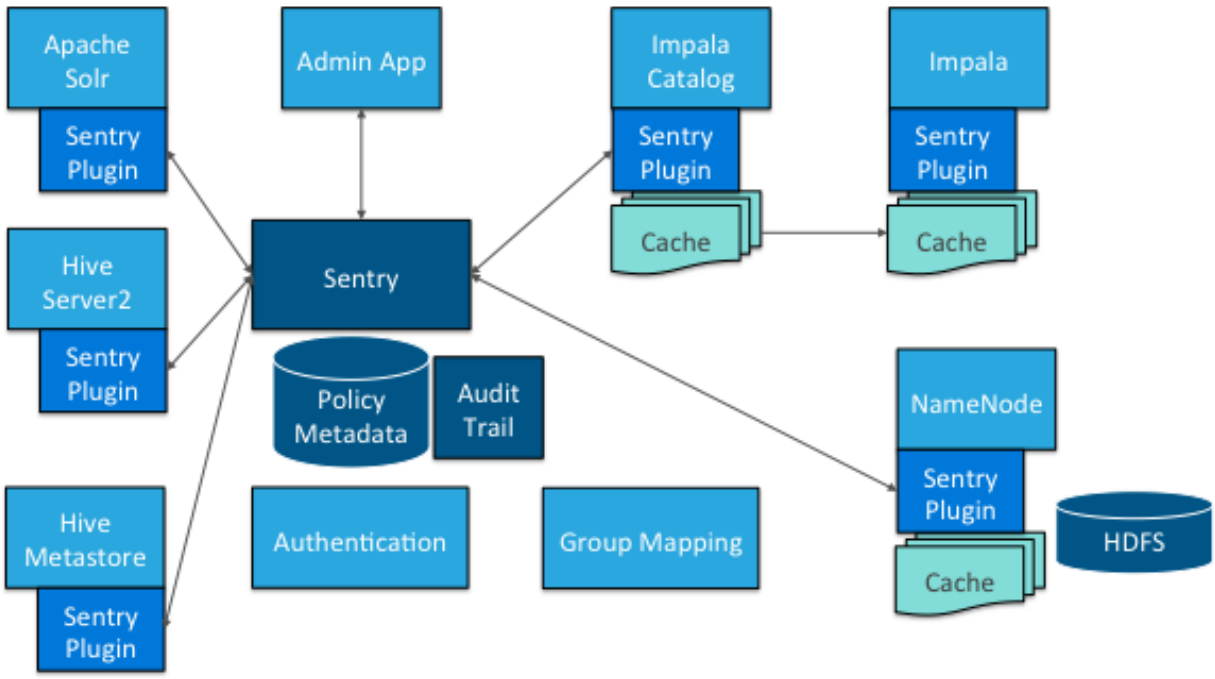

As illustrated above, Apache Sentry works with multiple Hadoop components. At the heart you have the Sentry Server which stores authorization metadata and provides APIs for tools to retrieve and modify this metadata securely.

Note that the Sentry Server only facilitates the metadata. The actual authorization decision is made by a policy engine which runs in data processing applications such as Hive or Impala. Each component loads the Sentry plugin which includes the service client for dealing with the Sentry service and the policy engine to validate the authorization request.

Hive and Sentry

Consider an example where Hive gets a request to access an object in a certain mode by a client. If Bob submits the following Hive query:

```
select * from production.sales
```

Hive will identify that user Bob is requesting SELECT access to the Sales table. At this point Hive will ask the Sentry plugin to validate Bob's access request. The plugin will retrieve Bob's privileges related to the Sales table and the policy engine will determine if the request is valid.

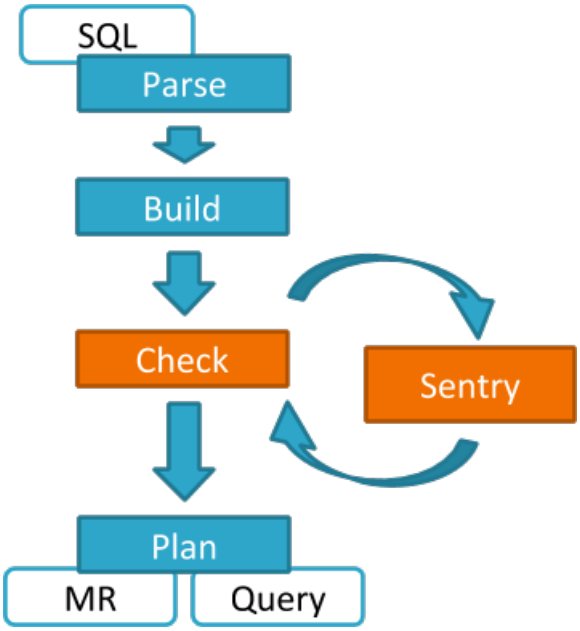

Hive works with both, the Sentry service and policy files. Cloudera recommends you use the Sentry service which makes it easier to manage user privileges. For more details and instructions, see Managing the Sentry Service *or* Sentry Policy File Authorization.

Impala and Sentry

Authorization processing in Impala is similar to that in Hive. The main difference is caching of privileges. Impala's Catalog server manages caching schema metadata and propagating it to all Impala server nodes. This Catalog server caches Sentry metadata as well. As a result, authorization validation in Impala happens locally and much faster.

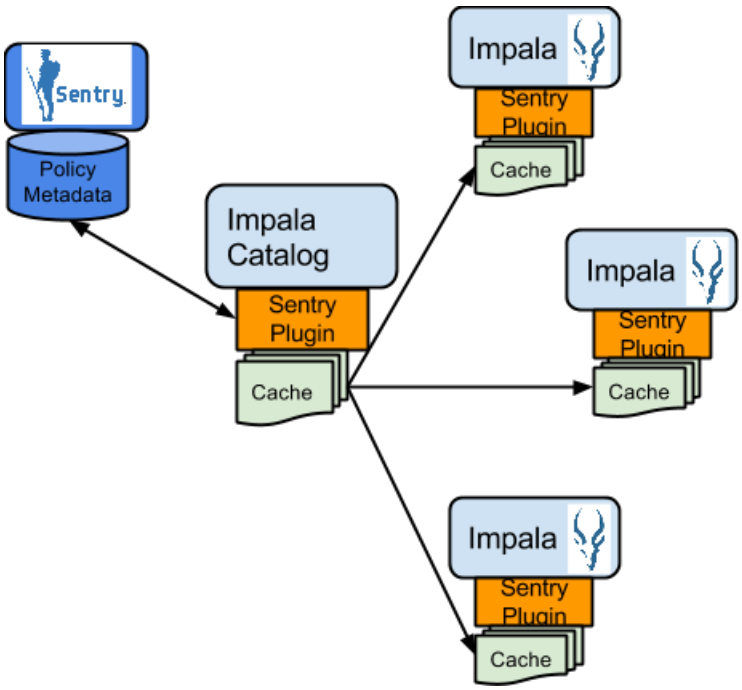

For detailed documentation, see Enabling Sentry Authorization for Impala.

Sentry-HDFS Synchronization

Sentry-HDFS authorization is focused on Hive warehouse data - that is, any data that is part of a table in Hive or Impala. The real objective of this integration is to expand the same authorization checks to Hive warehouse data being accessed from any other components such as Pig, MapReduce or Spark. At this point, this feature does not replace HDFS ACLs. Tables that are not associated with Sentry will retain their old ACLs.

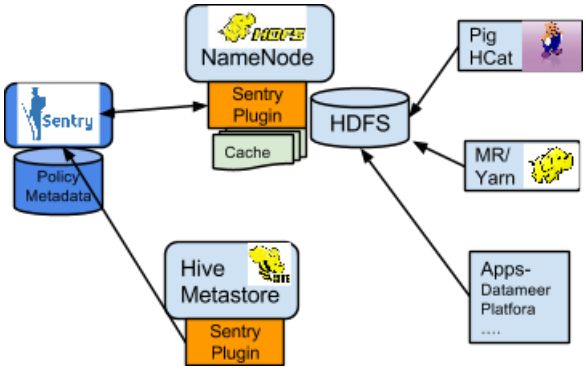

The mapping of Sentry privileges to HDFS ACL permissions is as follows:

- SELECT privilege -> Read access on the file.
- INSERT privilege -> Write access on the file.
- ALL privilege -> Read and Write access on the file.

The NameNode loads a Sentry plugin that caches Sentry privileges as well Hive metadata. This helps HDFS to keep file permissions and Hive tables privileges in sync. The Sentry plugin periodically polls the Sentry and Metastore to keep the metadata changes in sync.

For example, if Bob runs a Pig job that is reading from the Sales table data files, Pig will try to get the file handle from HDFS. At that point the Sentry plugin on the NameNode will figure out that the file is part of Hive data and overlay Sentry privileges on top of the file ACLs. As a result, HDFS will enforce the same privileges for this Pig client that Hive would apply for a SQL query.

For HDFS-Sentry synchronization to work, you *must* use the Sentry service, not policy file authorization. See Synchronizing HDFS ACLs and Sentry Permissions, for more details.

Search and Sentry

Sentry can apply a range of restrictions to various Search tasks, such accessing data or creating collections. These restrictions are consistently applied, regardless of the way users attempt to complete actions. For example, restricting access to data in a collection restricts that access whether queries come from the command line, from a browser, or through the admin console.

With Search, Sentry stores its privilege policies in a policy file (for example, sentry-provider.ini) which is stored in an HDFS location such as `hdfs://ha-nn-uri/user/solr/sentry/sentry-provider.ini`.

Sentry with Search does not support multiple policy files for multiple databases. However, you must use a separate policy file for each Sentry-enabled service. For example, Hive and Search were using policy file authorization, using a combined Hive and Search policy file would result in an invalid configuration and failed authorization on both services.

> **Note:** While Hive and Impala are compatible with the database-backed Sentry service, Search still uses Sentry's policy file authorization. Note that it is possible for a single cluster to use both, the Sentry service (for Hive and Impala as described above) and Sentry policy files (for Solr).

For detailed documentation, see Configuring Sentry Authorization for Cloudera Search.

Authorization Administration

The Sentry Server supports APIs to securely manipulate roles and privileges. Both Hive and Impala support SQL statements to manage privileges natively. Sentry assumes that HiveServer2 and Impala run as superusers, usually called `hive` and `impala`. To initiate top-level permissions for Sentry, an admin must login as a superuser.You can use either Beeline or the Impala shell to execute the following sample statement:

```
GRANT ROLE Analyst TO GROUP finance-managers
```

Using Hue to Manage Sentry Permissions

Hue supports a Security app to manage Sentry authorization. This allows users to explore and change table permissions. Here is a video blog that demonstrates its functionality.

Integration with Authentication Mechanisms for Identity Management

Like many distributed systems, Hadoop projects and workloads often consist of a collection of processes working in concert. In some instances, the initial user process conducts authorization throughout the entirety of the workload or job's lifecycle. But for processes that spawn additional processes, authorization can pose challenges. In this case, the spawned processes are set to execute as if they were the authenticated user, that is, setuid, and thus only have the privileges of that user. The overarching system requires a mapping to the authenticated principal and the user account must exist on the local host system for the setuid to succeed.

> **Important:**
>
> - Cloudera strongly recommends *against* using Hadoop's `LdapGroupsMapping` provider. `LdapGroupsMapping` should only be used in cases where OS-level integration is not possible. Production clusters require an identity provider that works well with all applications, not just Hadoop. Hence, often the preferred mechanism is to use tools such as SSSD, VAS or Centrify to replicate LDAP groups.
> - Cloudera does not support the use of Winbind in production environments. Winbind uses an inefficient approach to user/group mapping, which may lead to low performance or cluster failures as the size of the cluster, and the number of users and groups increases.
>
> Irrespective of the mechanism used, user/group mappings must be applied consistently across all cluster hosts for ease with maintenance.

System and Service Authorization - Several Hadoop services are limited to inter-service interactions and are not intended for end-user access. These services do support authentication to protect against unauthorized or malicious users. However, any user or, more typically, another service that has login credentials and can authenticate to the service is authorized to perform all actions allowed by the target service. Examples include ZooKeeper, which is used by internal systems such as YARN, Cloudera Search, and HBase, and Flume, which is configured directly by Hadoop administrators and thus offers no user controls.

The authenticated Kerberos principals for these "system" services are checked each time they access other services such as HDFS, HBase, and MapReduce, and therefore must be authorized to use those resources. Thus, the fact that Flume does not have an explicit authorization model does not imply that Flume has unrestricted access to HDFS and other services; the Flume service principals still must be authorized for specific locations of the HDFS file system. Hadoop administrators can establish separate system users for a services such as Flume to segment and impose access rights to only the parts of the file system for a specific Flume application.

Authorization within Hadoop Projects

Project	Authorization Capabilities
HDFS	File Permissions, Sentry*
MapReduce	File Permissions, Sentry*
YARN	File Permissions, Sentry*
Accumulo	
Flume	None
HBase	HBase ACLs
HiveServer2	File Permissions, Sentry
Hue	Hue authorization mechanisms (assigning permissions to Hue apps)
Impala	Sentry
Oozie	ACLs
Pig	File Permissions, Sentry*
Search	File Permissions, Sentry
Sentry	N/A
Spark	File Permissions, Sentry*
Sqoop	N/A
Sqoop2	None
ZooKeeper	ACLs
Cloudera Manager	Cloudera Manager roles
Cloudera Navigator	Cloudera Navigator roles
Backup and Disaster Recovery	N/A

* Sentry HDFS plug-in; when enabled, Sentry enforces its own access permissions over files that are part of tables defined in the Hive Metastore.

Auditing and Data Governance Overview

Organizations of all kinds want to understand where data in their clusters is coming from and how it is used. Cloudera Navigator Data Management component is a fully integrated data management and security tool for the Hadoop

platform that has been designed to meet compliance, data governance, and auditing needs of global enterprises. Without Cloudera Navigator, Hadoop clusters rely primarily on log files for auditing. However, log files are not an enterprise-class real-time auditing or monitoring solution. For example, log files can be corrupted by a system crash during a write commit.

Cloudera Navigator Data Management

Cloudera Navigator captures a complete and immutable record of all system activity. An audit trail can be used to determine the particulars—the who, what, where, and when—of a data breach or attempted breach. Auditing can be used to not only identify a rogue administrator who deleted user data, for example, but can also be used to recover data from a backup. Enterprises that must prove they are in compliance with HIPAA (Health Insurance Portability and Accountability Act), PCI (Payment Card Industry Data Security Standard), or other regulations associated with sensitive or personally identifiable data (PII) are required to produce auditing records when asked by government or other officials, such as banking regulators.

Auditing also serves to provide a historical record of data and context for data forensics. Data stewards and curators can use auditing reports to determine consumption and use patterns across various data sets by different user communities, for optimizing data access.

This section provides a brief overview of functionality of Cloudera Navigator. For complete details, see Cloudera Navigator Data Management.

Auditing

While Hadoop has historically lacked centralized cross-component audit capabilities, products such as Cloudera Navigator add secured, real-time audit components to key data and access frameworks. Using Cloudera Navigator, administrators can configure, collect, and view audit events, to understand who accessed what data and how.

Cloudera Navigator also lets administrators generate reports that list the HDFS access permissions granted to groups. Cloudera Navigator tracks access permissions and actual accesses to all entities in HDFS, Hive, HBase, Impala, Sentry, and Solr, and the Cloudera Navigator Metadata Server itself to help answer questions such as - who has access to which entities, which entities were accessed by a user, when was an entity accessed and by whom, what entities were accessed using a service, and which device was used to access. Cloudera Navigator auditing supports tracking access to:

- HDFS entities accessed by HDFS, Hive, HBase, Impala, and Solr services
- HBase and Impala
- Hive metadata
- Sentry
- Solr
- Cloudera Navigator Metadata Server

Data collected from these services also provides visibility into usage patterns for users, ability to see point-in-time permissions and how they have changed (leveraging Sentry), and review and verify HDFS permissions. Cloudera Navigator also provides out-of-the-box integration with leading enterprise metadata, lineage, and SIEM applications.

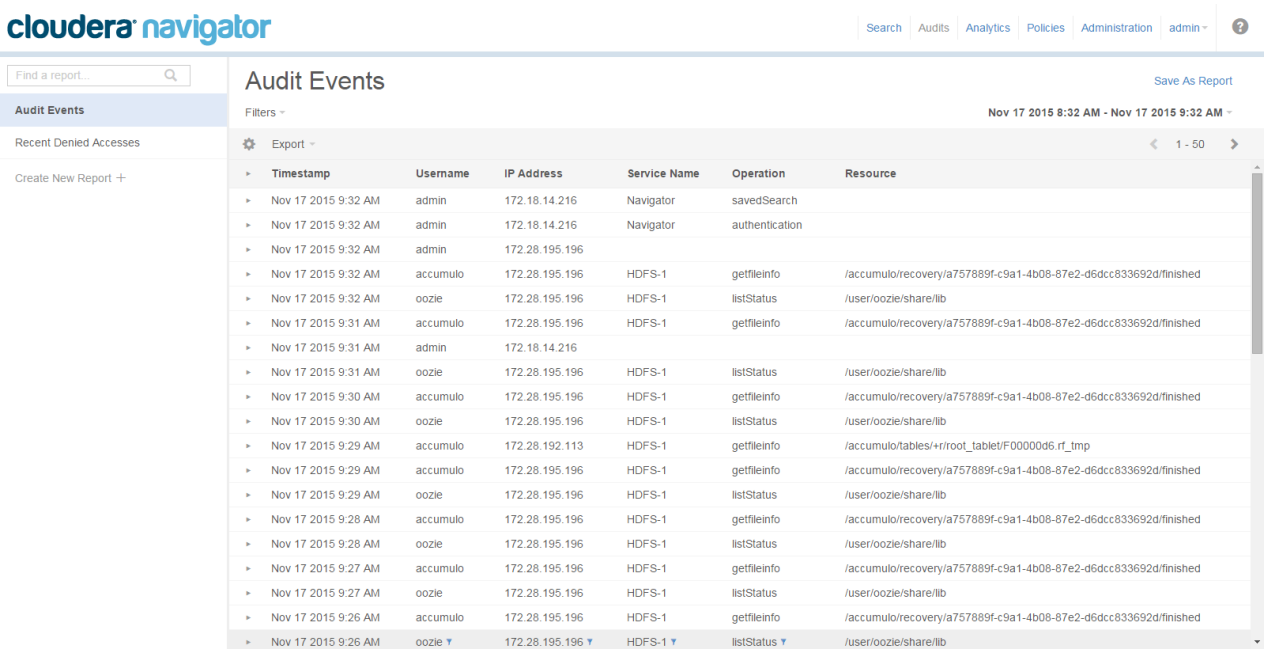

Metadata Management

Cloudera Navigator features complete metadata storage and supports data discovery. It consolidates technical metadata for all cluster data and enables automatic tagging of data based on the external sources entering the cluster. The consolidated metadata store is searchable through the Cloudera Navigator console, a web-based unified interface.

In addition, Cloudera Navigator supports user-defined metadata that can be applied to files, tables, and individual columns, to identify data assets for business context. The result is that data stewards can devise appropriate classification schemes for specific business purposes and data is more easily discovered and located by users.

Furthermore, policies can be used to automatically classify and applying metadata to cluster data based on arrival, scheduled interval, or other trigger.

Lineage

Cloudera Navigator lineage is a visualization tool for tracing data and its transformations from upstream to downstream through the cluster. Lineage can show the transforms that produced upstream data sources and the effect the data has on downstream artifacts, to the column level. Cloudera Navigator tracks lineage of HDFS files, datasets, and directories, Hive tables and columns, MapReduce and YARN jobs, Hive queries, Impala queries, Pig scripts, Oozie workflows, Spark jobs, and Sqoop jobs.

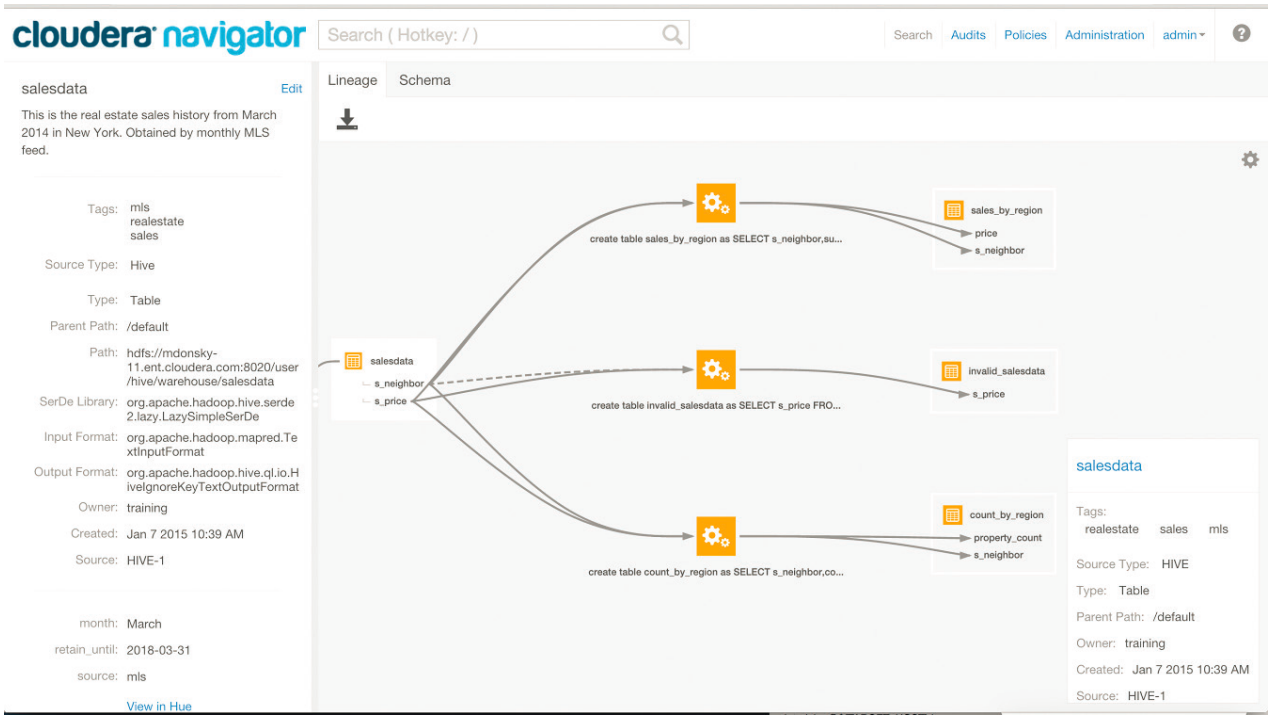

Integration within the Enterprise

Monitoring and reporting are only part a subset of enterprise auditing infrastructure capabilities. Enterprise policies may mandate that audit data route be able to integrate with existing enterprise SIEM (security information and event management) applications and other tools. Toward that end, Cloudera Navigator can deliver or export audit data through several mechanisms:

- Using `syslog` as a mediator between raw-event stream generated by Hadoop cluster and the SIEM tools.
- Using a REST API for custom enterprise tools.
- Exporting data to CSV or other text file.

Auditing and Components

The table below details auditing capabilities of Cloudera Manager and CDH components.

Component	Auditing Capabilities
HDFS	Events captured by Cloudera Navigator (including security events)
MapReduce	Inferred through HDFS
YARN	Inferred through HDFS
Accumulo	Log Files - Partial inclusion of security events; does not include non-bulk writes
Flume	Log Files
HBase	Audit events captured by Cloudera Navigator (including security events)
HiveServer2	Audit events captured by Cloudera Navigator
Hue	Inferred through underlying components
Impala	Audit events captured by Cloudera Navigator
Kafka	Log Files
Oozie	Log Files

Component	Auditing Capabilities
Pig	Inferred through HDFS
Search	Log Files
Sentry	Audit events captured by Cloudera Navigator
Spark	Inferred through HDFS
Sqoop	Log Files
Sqoop2	Log Files (including security events)
ZooKeeper	Log Files
Cloudera Manager	Audit events captured by Cloudera Navigator (partial capture of security events)
Cloudera Navigator	Audit events captured by Cloudera Navigator itself
Backup and Disaster Recovery	None

Security Events

Security events include a machine readable log of the following activities:

- User data read
- User data written
- Permission changes
- Configuration changes
- Login attempts
- Escalation of privileges
- Session Tracking
- Key Operations (Key Trustee)

Authentication

Authentication is a process that requires users and services to prove their identity when trying to access a system resource. Organizations typically manage user identity and authentication through various time-tested technologies, including Lightweight Directory Access Protocol (LDAP) for identity, directory, and other services, such as group management, and Kerberos for authentication.

Cloudera clusters support integration with both of these technologies. For example, organizations with existing LDAP directory services such as Active Directory (included in Microsoft Windows Server as part of its suite of Active Directory Services) can leverage the organization's existing user accounts and group listings instead of creating new accounts throughout the cluster. Using an external system such as Active Directory or OpenLDAP is required to support the user role authorization mechanism implemented in Cloudera Navigator.

For authentication, Cloudera supports integration with MIT Kerberos and with Active Directory. Microsoft Active Directory supports Kerberos for authentication in addition to its identity management and directory functionality, that is, LDAP.

Kerberos provides ***strong authentication***, *strong* meaning that cryptographic mechanisms—rather than passwords alone—are used in the exchange between requesting process and service during the authentication process.

These systems are not mutually exclusive. For example, Microsoft Active Directory is an LDAP directory service that also provides Kerberos authentication services, and Kerberos credentials can be stored and managed in an LDAP directory service. Cloudera Manager Server, CDH nodes, and Cloudera Enterprise components, such as Cloudera Navigator, Apache Hive, Hue, and Impala, can all make use of Kerberos authentication.

> **Note:** Cloudera does not provide a Kerberos implementation but rather can use MIT Kerberos or Microsoft Server Active Directory service and its KDC for authentication.

Configuring the cluster to use Kerberos requires having administrator privileges—and access to—the Kerberos server Key Distribution Center (KDC). The process may require debugging issues between the Cloudera Manager cluster and the KDC.

> **Note:** Integrating clusters to use Microsoft Active Directory as a KDC requires the Windows registry setting for `AllowTgtSessionKey` to be disabled (set to 0). If this registry key has already been enabled, users and credentials are not created—despite the "Successful" message at the end of the configuration/integration process. Before configuring Active Directory for use as KDC, check the value of `AllowTgtSessionKey` on the Active Directory instance and reset to 0 if necessary. See Registry Key to Allow Session Keys to Be Sent in Kerberos Ticket-Granting-Ticket at Microsoft for details.

On each host operating system underlying each node in a cluster, local Linux `user:group` accounts are created during installation of Cloudera Server and CDH services. To apply per-node authentication and authorization mechanism consistently across all the nodes of a cluster, local `user:group` accounts are mapped to user accounts and groups in an LDAP-compliant directory service, such as Active Directory or OpenLDAP. See Configuring LDAP Group Mappings on page 179 for details.

To facilitate the authentication process from each host system (node in the cluster) to the LDAP directory, Cloudera recommends using additional software mechanisms such as SSSD (Systems Security Services Daemon) or Centrify Server Suite. See the Centrify guide Identity and Access management for Cloudera for details.

Kerberos Security Artifacts Overview

Cloudera recommends using Kerberos for authentication because native Hadoop authentication alone checks only for valid `user:group` membership in the context of HDFS, but does not authenticate users or services across all network resources, as does Kerberos. Unlike other mechanisms that may be far easier to deploy (see Kerberos Deployment

Models on page 17 for details), the Kerberos protocol authenticates a requesting user or service for a specific period of time only, and each service that the user may want to use requires the appropriate Kerberos artifact in the context of the protocol. This section describes how Cloudera clusters use some of these artifacts, such as Kerberos principals and keytabs for user authentication, and how delegation tokens are used by the system to authenticate jobs on behalf of authenticated users at runtime.

Kerberos Principals

Each user and service that needs to authenticate to Kerberos needs a ***principal***, an entity that uniquely identifies the user or service in the context of possibly multiple Kerberos servers and related subsystems. A principal includes up to three pieces of identifying information, starting with the user or service name (called a ***primary***). Typically, the primary portion of the principal consists of the user account name from the operating system, such as `jcarlos` for the user's Unix account or `hdfs` for the Linux account associated with the service daemon on the host underlying cluster node.

Principals for users typically consist solely of the primary and the Kerberos ***realm*** name. The realm is a logical grouping of principals tied to the same Key Distribution Center (KDC) which is configured with many of the same properties, such as supported encryption algorithms. Large organizations may use realms as means of delegating administration to various groups or teams for specific sets of users or functions and distributing the authentication-processing tasks across multiple servers.

Standard practice is to use your organization's domain name as the Kerberos realm name (in all uppercase characters) to easily distinguish it as part of a Kerberos principal, as shown in this user principal pattern:

```
username@REALM.EXAMPLE.COM
```

The combination of the primary and the realm name can distinguish one user from another. For example, `jcarlos@SOME-REALM.EXAMPLE.COM` and `jcarlos@ANOTHER-REALM.EXAMPLE.COM` may be unique individuals within the same organization.

For **service role instance identities**, the primary is the Unix account name used by Hadoop daemons (`hdfs`, `mapred`, and so on) followed by an ***instance*** name that identifies the specific host on which the service runs. For example, `hdfs/hostname.fqdn.example.com@SOME-REALM.EXAMPLE.COM` is an example of the principal for an HDFS service instance. The forward slash (/) separates the primary and the instance names using this basic pattern:

```
service-name/hostname.fqdn.example.com@REALM.EXAMPLE.COM
```

The HTTP principal needed for Hadoop web service interfaces does not have a Unix local account for its primary but rather is `HTTP`.

An instance name can also identify users with special roles, such as administrators. For example, the principal `jcarlos@SOME-REALM.COM` and the principal `jcarlos/admin@SOME-REALM.COM` each have their own passwords and privileges, and they may or may not be the same individual.

For example, the principal for the HDFS service role instance running on a cluster in an organization with realms for each geographical location might be as follows:

```
hdfs/hostname.fqdn.example.com@OAKLAND.EXAMPLE.COM
```

Generally, the service name is the Unix account name used by the given service role instance, such as `hdfs` or `mapred`, as shown above. The HTTP principal for securing web authentication to Hadoop service web interfaces has no Unix account, so the primary for the principal is `HTTP`.

Kerberos Keytabs

A ***keytab*** is a file that contains the principal and the encrypted key for the principal. A keytab file for a Hadoop daemon is unique to each host since the principal names include the hostname. This file is used to authenticate a principal on a host to Kerberos without human interaction or storing a password in a plain text file. Because having access to the keytab file for a principal allows one to act as that principal, access to the keytab files should be tightly secured.

They should be readable by a minimal set of users, should be stored on local disk, and should not be included in host backups, unless access to those backups is as secure as access to the local host.

Delegation Tokens

Users in a Hadoop cluster authenticate themselves to the NameNode using their Kerberos credentials. However, once the user is authenticated, each job subsequently submitted must also be checked to ensure it comes from an authenticated user. Since there could be a time gap between a job being submitted and the job being executed, during which the user could have logged off, user credentials are passed to the NameNode using delegation tokens that can be used for authentication in the future.

Delegation tokens are a secret key shared with the NameNode, that can be used to impersonate a user to get a job executed. While these tokens can be renewed, new tokens can only be obtained by clients authenticating to the NameNode using Kerberos credentials. By default, delegation tokens are only valid for a day. However, since jobs can last longer than a day, each token specifies a NodeManager as a *renewer* which is allowed to renew the delegation token once a day, until the job completes, or for a maximum period of 7 days. When the job is complete, the NodeManager requests the NameNode to cancel the delegation token.

Token Format

The NameNode uses a random `masterKey` to generate delegation tokens. All active tokens are stored in memory with their expiry date (`maxDate`). Delegation tokens can either expire when the current time exceeds the expiry date, or, they can be canceled by the owner of the token. Expired or canceled tokens are then deleted from memory. The `sequenceNumber` serves as a unique ID for the tokens. The following section describes how the Delegation Token is used for authentication.

```
TokenID = {ownerID, renewerID, issueDate, maxDate, sequenceNumber}
TokenAuthenticator = HMAC-SHA1(masterKey, TokenID)
Delegation Token = {TokenID, TokenAuthenticator}
```

Authentication Process

To begin the authentication process, the client first sends the TokenID to the NameNode. The NameNode uses this TokenID and the `masterKey` to once again generate the corresponding TokenAuthenticator, and consequently, the Delegation Token. If the NameNode finds that the token already exists in memory, and that the current time is less than the expiry date (`maxDate`) of the token, then the token is considered valid. If valid, the client and the NameNode will then authenticate each other by using the TokenAuthenticator that they possess as the secret key, and MD5 as the protocol. Since the client and NameNode do not actually exchange TokenAuthenticators during the process, even if authentication fails, the tokens are not compromised.

Token Renewal

Delegation tokens must be renewed periodically by the designated renewer (`renewerID`). For example, if a NodeManager is the designated renewer, the NodeManager will first authenticate itself to the NameNode. It will then send the token to be authenticated to the NameNode. The NameNode verifies the following information before renewing the token:

- The NodeManager requesting renewal is the same as the one identified in the token by `renewerID`.
- The TokenAuthenticator generated by the NameNode using the `TokenID` and the `masterKey` matches the one previously stored by the NameNode.
- The current time must be less than the time specified by `maxDate`.

If the token renewal request is successful, the NameNode sets the new expiry date to `min(current time+renew period, maxDate)`. The tokens are persisted into NameNode metadata (`FSImage` and edit logs). If the NameNode was restarted at any time, it will have lost all previous tokens from memory. In this case, the token will be loaded into memory again from the NameNode metadata. Token renewals and cancellation persist, so NameNode restart does not impact the lifetime of the tokens.

A designated renewer can renew tokens as long as the token is not expired. The designated renewer can also cancel tokens. Shortly (by default, one hour) after the token is expired, or immediately after it is canceled, it is removed from the NameNode. Afterward, the NameNode cannot distinguish between a token that was canceled, or has expired.

Configuring Authentication in Cloudera Manager

Cloudera clusters can be configured to use Kerberos for authentication using a manual configuration process or by using the configuration wizard available from the Cloudera Manager Admin Console. Cloudera recommends using the wizard because it automates many of the configuration and deployment tasks. In addition, enabling Kerberos the cluster using the wizard also enables Kerberos authentication for all CDH components set up on the cluster, so you do not need to enable authentication for CDH as detailed in the Configuring Authentication in CDH Using the Command Line section.

> **Important:** Cloudera recommends configuring clusters to use Kerberos authentication after the Cloudera Manager Server has been configured for TLS/SSL. See How to Configure TLS Encryption for Cloudera Manager for details.

Cloudera Manager Kerberos Wizard Overview

The Cloudera Manager Kerberos wizard starts by verifying various details of the Kerberos instance that will be used for the cluster. Before using the wizard, be sure to gather all the details about the Kerberos service or engage the Kerberos administrator's help during this process. The details of the Kerberos instance are many and you will need to enter them in the wizard's pages.

The wizard requires a working KDC, either an MIT KDC or an Active Directory KDC. For configuration ease, the KDC should be set up and working prior to starting the wizard. Administrator-level privileges to the Kerberos instance are required to complete the prompts of the wizard, so obtain help from the Kerberos administrator if you do not have privileges.

Given the information provided to the wizard entry screens, the configuration wizard does the following:

- Configures the necessary properties in all configuration files—`core-site.xml`, `hdfs-site.xml`, `mapred-site.xml`, and `taskcontroller.cfg`—to identify Kerberos as the authentication mechanism for the cluster
- Configures the necessary properties in the `oozie-site.xml` and `hue.ini` files for Oozie and Hue for Kerberos authentication
- Creates principal and keytab files for core system users, such as `hdfs` and `mapred`, and for CDH services
- Distributes the keytab files to each host in the cluster
- Creates keytab files for `oozie` and hue users and deploys to the appropriate hosts that support these client-focused services
- Distributes a configured `krb5.conf` to all nodes in the cluster
- Stops all services
- Deploys client configurations
- Restarts all services throughout the cluster

Keytab file for...	Principals
`hdfs`	`hdfs`, `host`
`mapred`	`mapred`, `host`
`oozie`	`oozie`, `HTTP`
hue	hue

The `host` principal is the same in both `hdfs` and `mapred` keytab files.

After making the configuration changes and deploying the keytabs, and configuration files to the appropriate nodes in the cluster, Cloudera Manager starts all services to stand up the cluster.

- To use the Kerberos configuration wizard, see Enabling Kerberos Authentication Using the Wizard on page 51.
- To configure Kerberos authentication manually, see the Security How-To Guides on page 365.

Cloudera Manager User Accounts

Minimum Required Role: User Administrator (also provided by **Full Administrator**)

Access to Cloudera Manager features is controlled by user accounts. A user account identifies how a user is authenticated and determines what privileges are granted to the user.

When you are logged in to the Cloudera Manager Admin Console, the username you are logged in as is located at the far right of the top navigation bar—for example, if you are logged in as *admin* you will see 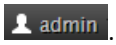

.

A user with the User Administrator or Full Administrator role manages user accounts through the **Administration** > **Users** page. View active user sessions on the **User Sessions** tab.

User Authentication

Cloudera Manager provides several mechanisms for authenticating users. You can configure Cloudera Manager to authenticate users against the Cloudera Manager database or against an external authentication service. The external authentication service can be an LDAP server (Active Directory or an OpenLDAP compatible directory), or you can specify another external service. Cloudera Manager also supports using the Security Assertion Markup Language (SAML) to enable single sign-on.

If you are using LDAP or another external service, you can configure Cloudera Manager so that it can use both methods of authentication (internal database and external service), and you can determine the order in which it performs these searches. If you select an external authentication mechanism, Full Administrator users can always authenticate against the Cloudera Manager database. This prevents locking everyone out if the authentication settings are misconfigured, such as with a bad LDAP URL.

With external authentication, you can restrict login access to members of specific groups, and can specify groups whose members are automatically given Full Administrator access to Cloudera Manager.

Users accounts in the Cloudera Manager database page show **Cloudera Manager** in the User Type column. User accounts in an LDAP directory or other external authentication mechanism show **External** in the User Type column.

User Roles

User accounts include the user's role, which determines the Cloudera Manager features visible to the user and the actions the user can perform. All tasks in the Cloudera Manager documentation indicate which role is required to perform the task. For more information about user roles, see Cloudera Manager User Roles on page 172.

Determining the Role of the Currently Logged in User

1. Click the logged-in username at the far right of the top navigation bar.
2. Select **My Profile**. The role displays. For example:

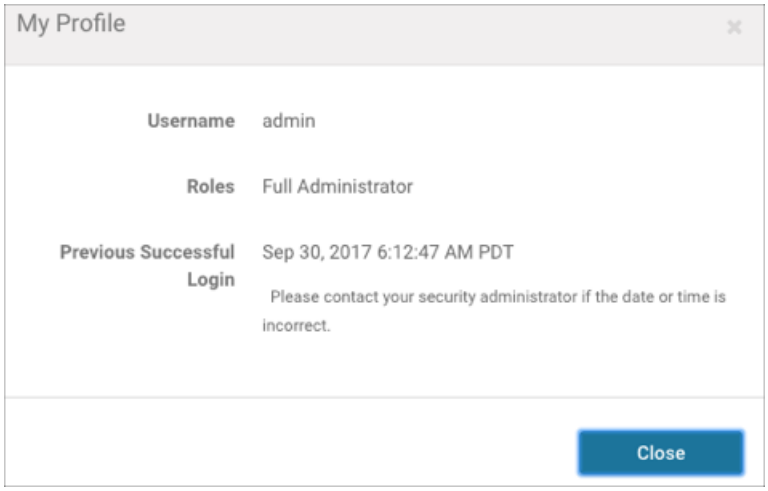

Changing the Logged-In Internal User Password

1. Click the logged-in username at the far right of the top navigation bar and select **Change Password**.
2. Enter the current password and a new password twice, and then click **OK**.

Adding an Internal User Account

1. Select **Administration** > **Users**.
2. Click the **Add User** button.
3. Enter a username and password.
4. In the Role drop-down menu, select a role for the new user.
5. Click **Add**.

Assigning User Roles

1. Select **Administration** > **Users**.
2. Check the checkbox next to one or more usernames.
3. Select **Actions for Selected** > **Assign User Roles**.
4. In the drop-down menu, select the role.
5. Click the **Assign Role** button.

Changing an Internal User Account Password

1. Select **Administration** > **Users**.
2. Click the **Change Password** button next to a username with User Type **Cloudera Manager**.
3. Type the new password and repeat it to confirm.
4. Click the **Update** button to make the change.

Deleting Internal User Accounts

1. Select **Administration** > **Users**.
2. Check the checkbox next to one or more usernames with User Type **Cloudera Manager**.
3. Select **Actions for Selected** > **Delete**.
4. Click the **OK** button. (There is no confirmation of the action.)

Viewing User Sessions

1. Select **Administration** > **Users**.
2. Click the tab **User Sessions**.

Configuring External Authentication for Cloudera Manager

Required Role: Full Administrator

> **Important:** This feature requires a Cloudera Enterprise license. It is not available in Cloudera Express. See Managing Licenses for more information.

Cloudera Manager supports user authentication against an internal database and against an external service. The following sections describe how to configure the supported external services.

Configuring Authentication Using Active Directory

1. Log in to Cloudera Manager Admin Console.
2. Select **Administration** > **Settings**.
3. Select **External Authentication** for the **Category** filter to display the settings.
4. For **Authentication Backend Order**, select the order in which Cloudera Manager should look up authentication credentials for login attempts.
5. For **External Authentication Type**, select **Active Directory**.
6. In the **LDAP URL** property, provide the URL of the Active Directory server.
7. In the **Active Directory Domain** property, provide the domain to authenticate against.

   LDAP URL and **Active Directory** are the only settings required to allow anyone in Active Directory to log in to Cloudera Manager.

   For example, if you set LDAP URL to `ldap://adserver.example.com` and the Active Directory Domain to `ADREALM.EXAMPLE.COM`, users can log into Cloudera Manager using just their username, such as `sampleuser`. They no longer require the complete string: `sampleuser@ADREALM.EXAMPLE.COM`.

8. In the **LDAP User Groups** property, optionally provide a comma-separated list of case-sensitive LDAP group names. If this list is provided, only users who are members of one or more of the groups in the list will be allowed to log into Cloudera Manager. If this property is left empty, all authenticated LDAP users will be able to log into Cloudera Manager. For example, if there is a group called `CN=ClouderaManagerUsers,OU=Groups,DC=corp,DC=com`, add the group name `ClouderaManagerUsers` to the **LDAP User Groups** list to allow members of that group to log in to Cloudera Manager.
9. To automatically assign a role to users when they log in, provide a comma-separated list of LDAP group names in the following properties:

   - LDAP Full Administrator Groups
   - LDAP User Administrator Groups
   - LDAP Cluster Administrator Groups
   - LDAP BDR Administrator Groups
   - LDAP Configurator Groups
   - LDAP Key Administrator Groups
   - LDAP Navigator Administrator Groups
   - LDAP Operator Groups
   - LDAP Limited Operator Groups
   - LDAP Auditor Groups
   - LDAP Dashboard Groups

   If you specify groups in these properties, users must also be a member of at least one of the groups specified in the **LDAP User Groups** property or they will not be allowed to log in. If these properties are left empty, users will be assigned to the Read-Only role and any other role assignment must be performed manually by an Administrator.

   > **Note:** Users added to LDAP groups are not automatically assigned user roles from the internal Cloudera Manager database so they are granted Read-Only access.

10. Restart the Cloudera Manager Server.

Configuring Authentication Using an LDAP-compliant Identity Service

An LDAP-compliant identity/directory service, such as OpenLDAP, provides different options for enabling Cloudera Manager to look-up user accounts and groups in the directory:

- Use a single Distinguished Name (DN) as a base and provide a pattern (**Distinguished Name Pattern**) for matching user names in the directory, or
- Search filter options let you search for a particular user based on somewhat broader search criteria – for example Cloudera Manager users could be members of different groups or organizational units (OUs), so a single pattern does not find all those users. Search filter options also let you find all the groups to which a user belongs, to help determine if that user should have login or admin access.

1. Log in to Cloudera Manager Admin Console.
2. Select **Administration** > **Settings**.
3. Select **External Authentication** for the **Category** filter to display the settings.
4. For **Authentication Backend Order**, select the order in which Cloudera Manager should look up authentication credentials for login attempts.
5. For **External Authentication Type**, select **LDAP**.
6. In the **LDAP URL** property, provide the URL of the LDAP server and (optionally) the base Distinguished Name (DN) (the search base) as part of the URL — for example `ldap://ldap-server.corp.com/dc=corp,dc=com`.
7. If your server does not allow anonymous binding, provide the user DN and password to be used to bind to the directory. These are the **LDAP Bind User Distinguished Name** and **LDAP Bind Password** properties. By default, Cloudera Manager assumes anonymous binding.
8. Use one of the following methods to search for users and groups:

   - You can search using User or Group search filters, using the **LDAP User Search Base**, **LDAP User Search Filter**, **LDAP Group Search Base** and **LDAP Group Search Filter** settings. These allow you to combine a base DN with a search filter to allow a greater range of search targets.

     For example, if you want to authenticate users who may be in one of multiple OUs, the search filter mechanism will allow this. You can specify the User Search Base DN as `dc=corp,dc=com` and the user search filter as `uid={0}`. Then Cloudera Manager will search for the user anywhere in the tree starting from the Base DN. Suppose you have two OUs—`ou=Engineering` and `ou=Operations`—Cloudera Manager will find User "foo" if it exists in either of these OUs, that is, `uid=foo,ou=Engineering,dc=corp,dc=com` or `uid=foo,ou=Operations,dc=corp,dc=com`.

     You can use a user search filter along with a DN pattern, so that the search filter provides a fallback if the DN pattern search fails.

     The Groups filters let you search to determine if a DN or username is a member of a target group. In this case, the filter you provide can be something like `member={0}` where `{0}` will be replaced with the **DN** of the user you are authenticating. For a filter requiring the username, `{1}` may be used, as `memberUid={1}`. This will return a list of groups the user belongs to, which will be compared to the list in the group properties discussed in step 8 of Configuring Authentication Using Active Directory on page 46.

     OR

   - Alternatively, specify a single base Distinguished Name (DN) and then provide a "Distinguished Name Pattern" in the **LDAP Distinguished Name Pattern** property.

     Use `{0}` in the pattern to indicate where the username should go. For example, to search for a distinguished name where the `uid` attribute is the username, you might provide a pattern similar to `uid={0},ou=People,dc=corp,dc=com`. Cloudera Manager substitutes the name provided at login into this pattern and performs a search for that specific user. So if a user provides the username "foo" at the Cloudera Manager login page, Cloudera Manager will search for the DN `uid=foo,ou=People,dc=corp,dc=com`.

If you provided a base DN along with the URL, the pattern only needs to specify the rest of the DN pattern. For example, if the URL you provide is `ldap://ldap-server.corp.com/dc=corp,dc=com`, and the pattern is `uid={0},ou=People`, then the search DN will be `uid=foo,ou=People,dc=corp,dc=com`.

9. Restart the Cloudera Manager Server.

Configuring Cloudera Manager to Use LDAPS

If the LDAP server certificate has been signed by a trusted Certificate Authority, steps 1 and 2 below may not be necessary.

1. Copy the CA certificate file to the Cloudera Manager Server host.
2. Import the CA certificate(s) from the CA certificate file to the local truststore. The default truststore is located in the `$JAVA_HOME/jre/lib/security/cacerts` file. This contains the default CA information shipped with the JDK. Create an alternate default file called `jssecacerts` in the same location as the `cacerts` file. You can now safely append CA certificates for any private or public CAs not present in the default `cacerts` file, while keeping the original file intact.

   For our example, we will follow this recommendation by copying the default `cacerts` file into the new `jssecacerts` file, and then importing the CA certificate to this alternate truststore.

```
cp $JAVA_HOME/jre/lib/security/cacerts $JAVA_HOME/jre/lib/security/jssecacerts
```

```
$ /usr/java/latest/bin/keytool -import -alias nt_domain_name
-keystore /usr/java/latest/jre/lib/security/jssecacerts -file path_to_CA_cert
```

> **Note:** The default password for the `cacerts` store is `changeit`. The `-alias` does not always need to be the domain name.

   Alternatively, you can use the Java options: `javax.net.ssl.trustStore` and `javax.net.ssl.trustStorePassword`. Open the `/etc/default/cloudera-scm-server` file and add the following options:

```
export CMF_JAVA_OPTS="-Xmx2G -XX:MaxPermSize=256m -XX:+HeapDumpOnOutOfMemoryError
-XX:HeapDumpPath=/tmp
-Djavax.net.ssl.trustStore=/usr/java/default/jre/lib/security/jssecacerts
-Djavax.net.ssl.trustStorePassword=changeit"
```

3. Configure the **LDAP URL** property to use `ldaps://ldap_server` instead of `ldap://ldap_server`.
4. Restart the Cloudera Manager Server.

Configuring Authentication Using an External Program

Cloudera Manager can use a custom external authentication program,. Typically, this may be a custom script that interacts with a custom authentication service. Cloudera Manager will call the external program with the username as the first command line argument. The password is passed over `stdin`. Cloudera Manager assumes the program will return the following exit codes identifying the user role for a successful authentication:

- 0 - Read-Only
- 1 - Full Administrator
- 2 - Limited Operator
- 3 - Operator
- 4 - Configurator
- 5 - Cluster Administrator
- 6 - BDR Administrator
- 7 - Navigator Administrator
- 8 - User Administrator

- 9 - Auditor
- 10 - Key Administrator
- 11 - Dashboard

and a negative value is returned for a failure to authenticate.

To configure authentication using an external program:

1. Log in to Cloudera Manager Admin Console.
2. Select **Administration** > **Settings**.
3. Select **External Authentication** for the **Category** filter to display the settings.
4. For **External Authentication Type**, select **External Program**.
5. Provide a path to the external program in the **External Authentication Program Path** property.

Configuring Authentication Using SAML

Cloudera Manager supports the Security Assertion Markup Language (SAML), an XML-based open standard data format for exchanging authentication and authorization data between parties, in particular, between an identity provider (IDP) and a service provider (SP). The SAML specification defines three roles: the principal (typically a user), the IDP, and the SP. In the use case addressed by SAML, the principal (user agent) requests a service from the service provider. The service provider requests and obtains an identity assertion from the IDP. On the basis of this assertion, the SP can make an access control decision—in other words it can decide whether to perform some service for the connected principal.

The primary SAML use case is called web browser single sign-on (SSO). A user wielding a user agent (usually a web browser) requests a web resource protected by a SAML SP. The SP, wanting to know the identity of the requesting user, issues an authentication request to a SAML IDP through the user agent. In the context of this terminology, Cloudera Manager operates as a SP. This topic discusses the Cloudera Manager part of the configuration process; it assumes that you are familiar with SAML and SAML configuration in a general sense, and that you have a functioning IDP already deployed.

Note:

- Cloudera Manager supports both SP- and IDP-initiated SSO.
- The logout action in Cloudera Manager will send a single-logout request to the IDP.
- SAML authentication has been tested with specific configurations of SiteMinder and Shibboleth. While SAML is a standard, there is a great deal of variability in configuration between different IDP products, so it is possible that other IDP implementations, or other configurations of SiteMinder and Shibboleth, may not interoperate with Cloudera Manager.
- To bypass SSO if SAML configuration is incorrect or not working, you can login using a Cloudera Manager local account using the URL: `http://cm_host:7180/cmf/localLogin`

Setting up Cloudera Manager to use SAML requires the following steps.

Preparing Files

You will need to prepare the following files and information, and provide these to Cloudera Manager:

- A Java keystore containing a private key for Cloudera Manager to use to sign/encrypt SAML messages. For guidance on creating Java keystores, see Understanding Keystores and Truststores on page 193.
- The SAML metadata XML file from your IDP. This file must contain the public certificates needed to verify the sign/encrypt key used by your IDP per the SAML Metadata Interoperability Profile. For example, if you are using the Shibboleth IdP, the metadata file is available at: `https://<IdPHOST>:8080/idp/shibboleth`.

  Note: For guidance on how to obtain the metadata XML file from your IDP, either contact your IDP administrator or consult the documentation for the version of the IDP you are using.

- The entity ID that should be used to identify the Cloudera Manager instance

- How the user ID is passed in the SAML authentication response:
  - As an attribute. If so, what identifier is used.
  - As the NameID.
- The method by which the Cloudera Manager role will be established:
  - From an attribute in the authentication response:
    - What identifier will be used for the attribute
    - What values will be passed to indicate each role
  - From an external script that will be called for each use:
    - The script takes user ID as $1
    - The script sets an exit code to reflect successful authentication of the assigned role:
      - 0 - Full Administrator
      - 1 - Read-Only
      - 2 - Limited Operator
      - 3 - Operator
      - 4 - Configurator
      - 5 - Cluster Administrator
      - 6 - BDR Administrator
      - 7 - Navigator Administrator
      - 8 - User Administrator
      - 9 - Auditor
      - 10 - Key Administrator
      - 11 - Dashboard

      and a negative value is returned for a failure to authenticate.

Configuring Cloudera Manager

1. Log in to Cloudera Manager Admin Console.
2. Select **Administration** > **Settings**.
3. Select **External Authentication** for the **Category** filter to display the settings.
4. Set the **External Authentication Type** property to **SAML** (the Authentication Backend Order property is ignored for SAML).
5. Set the **Path to SAML IDP Metadata File** property to point to the IDP metadata file.
6. Set the **Path to SAML Keystore File** property to point to the Java keystore prepared earlier.
7. In the **SAML Keystore Password** property, set the keystore password.
8. In the **Alias of SAML Sign/Encrypt Private Key** property, set the alias used to identify the private key for Cloudera Manager to use.
9. In the **SAML Sign/Encrypt Private Key Password** property, set the private key password.
10. Set the **SAML Entity ID** property if:
    - There is more than one Cloudera Manager instance being used with the same IDP (each instance needs a different entity ID).
    - Entity IDs are assigned by organizational policy.
11. In the **Source of User ID in SAML Response** property, set whether the user ID will be obtained from an attribute or the NameID.

    If an attribute will be used, set the attribute name in the **SAML attribute identifier for user ID** property. The default value is the normal OID used for user IDs and so may not need to be changed.
12. In the **SAML Role assignment mechanism** property, set whether the role assignment will be done from an attribute or an external script.

- If an attribute will be used:
  - In the **SAML attribute identifier for user role** property, set the attribute name if necessary. The default value is the normal OID used for OrganizationalUnits and so may not need to be changed.
  - In the **SAML Attribute Values for Roles** property, set which attribute values will be used to indicate the user role.
- If an external script will be used, set the path to that script in the **Path to SAML Role Assignment Script** property. Make sure that the script is executable (an executable binary is fine - it doesn't need to be a shell script).

13. Save the changes. Cloudera Manager will run a set of validations that ensure it can find the metadata XML and the keystore, and that the passwords are correct. If you see a validation error, correct the problem before proceeding.
14. Restart the Cloudera Manager Server.

Configuring the IDP

After the Cloudera Manager Server is restarted, it will attempt to redirect to the IDP login page instead of showing the normal CM page. This may or may not succeed, depending on how the IDP is configured. In either case, the IDP will need to be configured to recognize CM before authentication will actually succeed. The details of this process are specific to each IDP implementation - refer to your IDP documentation for details. If you are using the Shibboleth IdP, information on configuring the IdP to communicate with a Service Provider is available here.

1. Download the Cloudera Manager's SAML metadata XML file from `http://hostname:7180/saml/metadata`.
2. Inspect the metadata file and ensure that any URLs contained in the file can be resolved by users' web browsers. The IDP will redirect web browsers to these URLs at various points in the process. If the browser cannot resolve them, authentication will fail. If the URLs are incorrect, you can manually fix the XML file or set the Entity Base URL in the CM configuration to the right value, and then re-download the file.
3. Provide this metadata file to your IDP using whatever mechanism your IDP provides.
4. Ensure that the IDP has access to whatever public certificates are necessary to validate the private key that was provided to Cloudera Manager earlier.
5. Ensure that the IDP is configured to provide the User ID and Role using the attribute names that Cloudera Manager was configured to expect, if relevant.
6. Ensure the changes to the IDP configuration have taken effect (a restart may be necessary).

Verifying Authentication and Authorization

1. Return to the Cloudera Manager Admin Console and refresh the login page.
2. Attempt to log in with credentials for a user that is entitled. The authentication should complete and you should see the **Home** > **Status** tab.
3. If authentication fails, you will see an IDP provided error message. Cloudera Manager is not involved in this part of the process, and you must ensure the IDP is working correctly to complete the authentication.
4. If authentication succeeds but the user is not authorized to use Cloudera Manager, they will be taken to an error page by Cloudera Manager that explains the situation. If an user who should be authorized sees this error, then you will need to verify their role configuration, and ensure that it is being properly communicated to Cloudera Manager, whether by attribute or external script. The Cloudera Manager log will provide details on failures to establish a user's role. If any errors occur during role mapping, Cloudera Manager will assume the user is unauthorized.

Enabling Kerberos Authentication Using the Wizard

Required Role: Cluster Administrator or Full Administrator

Cloudera Manager provides a wizard for integrating your organization's Kerberos instance with your cluster to provide authentication services.

Kerberos must already be deployed in your organization and the Kerberos key distribution center (KDC) must be ready to use, with a realm established. For Hue and Oozie, the Kerberos realm must support renewable tickets.

> **Important:** Before integrating Kerberos with your cluster, configure TLS/SSL encryption between Cloudera Manager Server and all Cloudera Manager Agent host systems in the cluster. During the Kerberos integration process, Cloudera Manager Server sends keytab files to the Cloudera Manager Agent hosts, and TLS/SSL encrypts the network communication so these files are protected. See Configuring TLS Encryption for Cloudera Manager on page 195 for details.

Cloudera Manager clusters can be integrated with MIT Kerberos or with Microsoft Active Directory:

- See MIT Kerberos home and MIT Kerberos 5 Release 1.8.6 documentation for more information about MIT Kerberos.
- See Using a Centralized Active Directory Service on page 20 and Microsoft Active Directory documentation for more information about using Active Directory as a KDC.

For Active Directory, you must have administrative privileges to the Active Directory instance for initial setup and for on-going management, or you will need to have the help of your AD administrator prior to and during the integration process. For example, administrative access is needed to access the Active Directory KDC, create principals, and troubleshoot Kerberos TGT/TGS-ticket-renewal and take care of any other issues that may arise.

Kerberos client OS-specific packages must be installed on all cluster hosts and client hosts that will authenticate using Kerberos.

OS	Packages Required
RHEL 7 Compatible, **RHEL 6 Compatible**, **RHEL 5 Compatible**	• `openldap-clients` on the Cloudera Manager Server host • `krb5-workstation`, `krb5-libs` on ALL hosts
SLES	• `openldap2-client` on the Cloudera Manager Server host • `krb5-client` on ALL hosts
Ubuntu or Debian	• `ldap-utils` on the Cloudera Manager Server host • `krb5-user` on ALL hosts
Windows	• `krb5-workstation`, `krb5-libs` on ALL hosts

See Before you Begin Using the Wizard on page 56 for more information.

Cloudera supports the Kerberos version that ships with each supported operating system listed in CDH and Cloudera Manager Supported Operating Systems.

Step 1: Install Cloudera Manager and CDH

Cloudera strongly recommends that you install and configure the Cloudera Manager Server and Cloudera Manager Agents and CDH to set up a fully-functional CDH cluster *before* trying to configure Kerberos authentication for the cluster.

Required user:group Settings for Security

When you install the CDH packages and the Cloudera Manager Agents on your cluster hosts, Cloudera Manager takes some steps to provide system security such as creating the OS `user:group` accounts and setting directory permissions as shown in the table below. These user accounts and directory permissions work with the Hadoop Kerberos security requirements.

This User	Runs These Roles
`hdfs`	NameNode, DataNodes, and Secondary NameNode
`mapred`	JobTracker and TaskTrackers (MR1) and Job History Server (YARN)
`yarn`	ResourceManager and NodeManagers (YARN)
`oozie`	Oozie Server

This User	Runs These Roles
hue	Hue Server, Beeswax Server, Authorization Manager, and Job Designer

The `hdfs` user has HDFS superuser privileges.

When you install the Cloudera Manager Server on the server host, a new Unix user account called `cloudera-scm` is created automatically to support security. The Cloudera Manager Server uses this account to create host principals and deploy the keytabs on your cluster.

Depending on whether you installed CDH and Cloudera Manager at the same time or not, use one of the following sections for information on configuring directory ownerships on cluster hosts:

New Installation, Cloudera Manager and CDH Together

Installing a new Cloudera Manager cluster with CDH components at the same time can save you some of the user:group configuration required if you install them separately. The installation process creates the necessary user accounts on the Linux host system for the service daemons. At the end of the installation process when each cluster node starts up, the Cloudera Manager Agent process on the host automatically configures the directory ownership as shown in the table below, and the Hadoop daemons can then automatically set permissions for their respective directories. Do not change the directory owners on the cluster. They must be configured exactly as shown below.

Directory Specified in this Property	Owner
`dfs.name.dir`	`hdfs:hadoop`
`dfs.data.dir`	`hdfs:hadoop`
`mapred.local.dir`	`mapred:hadoop`
`mapred.system.dir` in HDFS	`mapred:hadoop`
`yarn.nodemanager.local-dirs`	`yarn:yarn`
`yarn.nodemanager.log-dirs`	`yarn:yarn`
`oozie.service.StoreService.jdbc.url` (if using Derby)	`oozie:oozie`
`[[database]] name`	`hue:hue`
`javax.jdo.option.ConnectionURL`	`hue:hue`

Existing CDH Cluster Installed Before Cloudera Manager Installation

If you have been using HDFS and running MapReduce jobs in an existing installation of CDH before Cloudera Manager was installed, you must manually configure the directory ownership as shown in the table above to enable the Hadoop daemons to set appropriate permissions on each directory. Configure directory user:group ownership exactly as shown in the table.

Step 2: Installing JCE Policy File for AES-256 Encryption

> **Note:** This step is not required when using JDK 1.8.0_161 or greater. JDK 1.8.0_161 enables unlimited strength encryption by default.

By default, CentOS and Red Hat Enterprise Linux 5.5 (and higher) use AES-256 encryption for Kerberos tickets, so the Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction Policy File must be installed on all cluster hosts as detailed below. Alternatively, the Kerberos instance can be modified to not use AES-256.

To install the JCE Policy file on the host system at the OS layer:

1. Download the `jce_policy-x.zip`.

2. Unzip the file.
3. Follow the steps in the `README.txt` to install it.

To use Cloudera Manager to install the JCE policy file:

1. Log in to the Cloudera Manager Admin Console.
2. Select **Hosts** > **All Hosts**.
3. Click the **Re-run Upgrade Wizard** and select the option to have Cloudera Manager install the JCE Policy file.

Alternative: Disable AES-256 encryption from the Kerberos instance:

1. Remove `aes256-cts:normal` from the `supported_enctypes` field of the `kdc.conf` or `krb5.conf` file.
2. Restart the Kerberos KDC and the kadmin server so the changes take effect.

The keys of relevant principals, such as Ticket Granting Ticket principal (krbtgt/REALM@REALM), might need to change.

> **Note:** If AES-256 remains in use despite disabling it, it may be because the `aes256-cts:normal` setting existed when the Kerberos database was created. To resolve this issue, create a new Kerberos database and then restart both the KDC and the kadmin server.

To verify the type of encryption used in your cluster:

1. **For MIT KDC:** On the local KDC host, type this command in the kadmin.local or kadmin shell to create a test principal:

```
kadmin:  addprinc test
```

   For Active Directory: Create a new AD account with the name, `test`.

2. On a cluster host, type this command to start a Kerberos session as test:

```
$ kinit test
```

3. On a cluster host, type this command to view the encryption type in use:

```
$ klist -e
```

   If AES is being used, output like the following is displayed after you type the `klist` command (note that `AES-256` is included in the output):

```
Ticket cache: FILE:/tmp/krb5cc_0
Default principal: test@Cloudera Manager
Valid starting     Expires            Service principal
05/19/11 13:25:04  05/20/11 13:25:04  krbtgt/Cloudera Manager@Cloudera Manager
    Etype (skey, tkt): AES-256 CTS mode with 96-bit SHA-1 HMAC, AES-256 CTS mode with
96-bit SHA-1 HMAC
```

Step 3: Create the Kerberos Principal for Cloudera Manager Server

At the end of the integration process using the configuration wizard, Cloudera Manager Server will create host principals and deploy keytabs for all services configured on the cluster, which means that Cloudera Manager Server needs its own principal and have privileges to create these other accounts.

> **Note:** This task requires administrator privileges on the Kerberos instance. If you do not have administrator privileges, ask your Kerberos administrator to setup the principal and password for you. You will need to enter the principal name and password into the wizard later in the process (see Import KDC Account Manager Credentials on page 57).

If an administrator principal to act on behalf of Cloudera Manager cannot be created on the Kerberos KDC for whatever reason, Cloudera Manager will not be able to create or manage principals and keytabs for CDH services. That means

these principals must be created manually on the Kerberos KDC and then imported (retrieved) by Cloudera Manager. See Using a Custom Kerberos Keytab Retrieval Script on page 67 for details about this process.

Creating the Cloudera Manager Principal

The steps below summarize the process of adding a principal specifically for Cloudera Manager Server to an MIT KDC and an Active Directory KDC. See documentation from MIT, Microsoft, or the appropriate vendor for more detailed information.

Creating a Principal in Active Directory

Check your Microsoft documentation for specific details for your Active Directory KDC. The general process is as follows:

1. Create an Organizational Unit (OU) in your Active Directory KDC service that will contain the principals for use by the CDH cluster.
2. Add a new user account to Active Directory, for example, *`username`*`@EXAMPLE.COM`. Set the password for the user to never expire.
3. Use the **Delegate Control** wizard in Active Directory to grant this new user permission to **Create, Delete, and Manage User Accounts** in the OU created in step 1. Make sure that these permissions are only applied to that specific OU, and nowhere else.

Creating a Principal in an MIT KDC

For MIT Kerberos, administrator principals are defined in the `/var/kerberos/krb5kdc/kadm5.acl` file on the KDC host. The default entry is:

```
*/admin@EXAMPLE.COM      *
```

In this example, principals that include the instance name `admin` designate a user account as an administrator, such as *`jdoe`*`/admin@`*`EXAMPLE.COM`*.

If you modify the `kadm5.acl` file, such as replacing `EXAMPLE.COM` with your realm name, make sure to restart the `kadmin` service:

- **RHEL 7 Compatible:**

```
systemctl restart kadmin
```

- **All Others:**

```
service kadmin restart
```

Create the Cloudera Manager Server administrator principal as shown below, using the `admin` instance name and your realm name. If your `kadm5.acl` file defines a different pattern for administrators, make sure that the principal you create matches that pattern.

For MIT Kerberos KDC on a remote host:

> **Note:** To connect to a remote KDC using the `kadmin` command, your currently authenticated Kerberos principal must be an existing Kerberos administrator. If you do not have an existing admin principal, you must run `kadmin.local` as described below.

```
kadmin
kadmin: addprinc -pw password cloudera-scm/admin@EXAMPLE.COM
```

For MIT Kerberos KDC on the local host:

> **Note:** The `kadmin.local` command must be run as `root` on the same host as the KDC.

```
kadmin.local
kadmin.local: addprinc -pw password cloudera-scm/admin@EXAMPLE.COM
```

Step 4: Enabling Kerberos Using the Wizard

Minimum Required Role: Full Administrator

To start the Kerberos wizard:

1. Go to the Cloudera Manager Admin Console and click ▼ to the right of the cluster for which you want to enable Kerberos authentication.
2. Select **Enable Kerberos**.

Before you Begin Using the Wizard

The Welcome page lists steps you should have completed before starting the wizard.

- Set up a working KDC. Cloudera Manager supports authentication with MIT KDC and Active Directory.
- Configure the KDC to allow renewable tickets with non-zero ticket lifetimes.

  Active Directory allows renewable tickets with non-zero lifetimes by default. You can verify this by checking **Domain Security Settings** > **Account Policies** > **Kerberos Policy** in Active Directory.

  For MIT KDC, make sure you have the following lines in the `kdc.conf`.

```
max_life = 1d
max_renewable_life = 7d
```

- If you are using Active Directory, make sure LDAP over TLS/SSL (LDAPS) is enabled for the Domain Controllers.
- Hostnames *must* be in lowercase. If you use uppercase letters in any hostname, the cluster services will not start after enabling Kerberos.
- Install the OS-specific packages for your cluster listed in the table:

OS	Packages Required
RHEL 7 Compatible, RHEL 6 Compatible, RHEL 5 Compatible	• `openldap-clients` on the Cloudera Manager Server host • `krb5-workstation`, `krb5-libs` on ALL hosts
SLES	• `openldap2-client` on the Cloudera Manager Server host • `krb5-client` on ALL hosts
Ubuntu or Debian	• `ldap-utils` on the Cloudera Manager Server host • `krb5-user` on ALL hosts
Windows	• `krb5-workstation`, `krb5-libs` on ALL hosts

- Create an account for Cloudera Manager that has the permissions to create other accounts in the KDC. This should have been completed as part of Step 3: Create the Kerberos Principal for Cloudera Manager Server on page 54.

Important:

If YARN Resource Manager HA has been enabled in a non-secure cluster, before enabling Kerberos you must clear the StateStore znode in ZooKeeper, as follows:

1. Go to the Cloudera Manager Admin Console home page, click to the right of the YARN service and select **Stop**.
2. When you see a **Finished** status, the service has stopped.
3. Go to the YARN service and select **Actions** > **Format State Store**.
4. When the command completes, click **Close**.

Once you are able to check all the items on this list, click **Continue**.

KDC Information

On this page, select the KDC type you are using, MIT KDC or Active Directory, and complete the fields as applicable to enable Cloudera Manager to generate principals/accounts for the CDH services running on the cluster.

Note:

- If you are using AD and have multiple Domain Controllers behind a Load Balancer, enter the name of the Load Balancer in the **KDC Server Host** field and any *one* of the Domain Controllers in **Active Directory Domain Controller Override**. Hadoop daemons will use the Load Balancer for authentication, but Cloudera Manager will use the override for creating accounts.
- If you have multiple Domain Controllers (in case of AD) or MIT KDC servers, only enter the name of any *one* of them in the **KDC Server Host** field. Cloudera Manager will use that server only for creating accounts. If you choose to use Cloudera Manager to manage `krb5.conf`, you can specify the rest of the Domain Controllers using Safety Valve as explained below.
- Make sure the entries for the **Kerberos Encryption Types** field matches what your KDC supports.
- If you are using an Active Directory KDC, you can configure Active Directory account properties such as `objectClass` and `accountExpires` directly from the Cloudera Manager UI. You can also enable Cloudera Manager to delete existing AD accounts so that new ones can be created when Kerberos credentials are being regenerated. See Viewing or Regenerating Kerberos Credentials Using Cloudera Manager on page 67.

Click **Continue** to proceed.

KRB5 Configuration

Manage krb5.conf through Cloudera Manager allows you to choose whether Cloudera Manager should deploy the `krb5.conf` on your cluster or not. If left unchecked, you must ensure that the `krb5.conf` is deployed on all hosts in the cluster, including the Cloudera Manager Server's host.

If you check **Manage krb5.conf through Cloudera Manager**, this page will let you configure the properties that will be emitted in it. In particular, the safety valves on this page can be used to configure cross-realm authentication. More information can be found at Configuring a Dedicated MIT KDC for Cross-Realm Trust on page 163.

Note: Cloudera Manager is unable to use a non-default realm. You must specify the default realm.

Click **Continue** to proceed.

Import KDC Account Manager Credentials

Enter the username and password for the user that can create principals for CDH cluster in the KDC. This is the user/principal you created in Step 3: Create the Kerberos Principal for Cloudera Manager Server on page 54. Cloudera Manager encrypts the username and password into a keytab and uses it as needed to create new principals.

> **Note:** Enter the REALM portion of the principal in upper-case only to conform to Kerberos convention.

> **Note:**
>
> Many enterprises employ policies that require all passwords to be changed after a particular number of days. If you must change the password in Cloudera Manager for the Account Manager, then:
>
> 1. In the Cloudera Manager Admin Console, select **Administration** > **Security**.
> 2. Go to the **Kerberos Credentials** tab and click **Import Kerberos Account Manager Credentials**.
> 3. In the **Import Kerberos Account Manager Credentials** dialog box, enter the username and password for the user that can create principals for CDH cluster in the KDC.

Click **Continue** to proceed.

(Optional) Configuring Custom Kerberos Principals

Starting with Cloudera Manager 5.4, you can configure custom service principals for CDH services. Before you begin making configuration changes, see Customizing Kerberos Principals on page 63 for some additional configuration changes required and limitations.

Configure HDFS DataNode Ports

On this page, specify the privileged ports needed by the DataNode's Transceiver Protocol and the HTTP Web UI in a secure cluster.

Use the checkbox to confirm you are ready to restart the cluster. Click **Continue**.

Enabling Kerberos

This page lets you track the progress made by the wizard as it first stops all services on your cluster, deploys the `krb5.conf`, generates keytabs for other CDH services, deploys client configuration and finally restarts all services. Click **Continue**.

Congratulations

The final page lists the cluster(s) for which Kerberos has been successfully enabled. Click **Finish** to return to the Cloudera Manager Admin Console home page.

Step 5: Create the HDFS Superuser

To be able to create home directories for users, you will need access to the HDFS superuser account. (CDH automatically created the HDFS superuser account on each cluster host during CDH installation.) When you enabled Kerberos for the HDFS service, you lost access to the default HDFS superuser account using `sudo -u hdfs` commands. Cloudera recommends you use a different user account as the superuser, not the default `hdfs` account.

Designating a Non-Default Superuser Group

To designate a different group of superusers instead of using the default `hdfs` account, follow these steps:

1. Go to the Cloudera Manager Admin Console and navigate to the HDFS service.
2. Click the **Configuration** tab.
3. Select **Scope** > **HDFS (Service-Wide)**.
4. Select **Category** > **Security**.
5. Locate the **Superuser Group** property and change the value to the appropriate group name for your environment. For example, *<superuser>*.
6. Click **Save Changes** to commit the changes.
7. Restart the HDFS service.

To enable your access to the superuser account now that Kerberos is enabled, you must now create a Kerberos principal or an Active Directory user whose first component is *<superuser>*:

If you are using Active Directory

Add a new user account to Active Directory, `<superuser>@YOUR-REALM.COM`. The password for this account should be set to never expire.

If you are using MIT KDC

1. In the `kadmin.local` or `kadmin` shell, type the following command to create a Kerberos principal called *<superuser>*:

```
kadmin:  addprinc <superuser>@YOUR-LOCAL-REALM.COM
```

   This command prompts you to create a password for the *<superuser>* principal. You should use a strong password because having access to this principal provides superuser access to all of the files in HDFS.
2. To run commands as the HDFS superuser, you must obtain Kerberos credentials for the *<superuser>* principal. To do so, run the following command and provide the appropriate password when prompted.

```
$ kinit <superuser>@YOUR-LOCAL-REALM.COM
```

Step 6: Get or Create a Kerberos Principal for Each User Account

Now that Kerberos is configured and enabled on your cluster, you and every other Hadoop user must have a Kerberos principal or keytab to obtain Kerberos credentials to be allowed to access the cluster and use the Hadoop services. In the next step of this procedure, you will need to create your own Kerberos principals to verify that Kerberos security is working on your cluster. If you and the other Hadoop users already have a Kerberos principal or keytab, or if your Kerberos administrator can provide them, you can skip ahead to the next step.

The following instructions explain how to create a Kerberos principal for a user account.

If you are using Active Directory

Add a new AD user account for each new user that should have access to the cluster. You do not need to make any changes to existing user accounts.

If you are using MIT KDC

1. In the `kadmin.local` or `kadmin` shell, use the following command to create user principals by replacing `YOUR-LOCAL-REALM.COM` with the name of your realm, and replacing `USERNAME` with a username:

```
kadmin:  addprinc USERNAME@YOUR-LOCAL-REALM.COM
```

2. Enter and re-enter a password when prompted.

Step 7: Prepare the Cluster for Each User

Before you and other users can access the cluster, there are a few tasks you must do to prepare the hosts for each user.

1. Make sure all hosts in the cluster have a Linux user account with the same name as the first component of that user's principal name. For example, the Linux account `joe` should exist on every box if the user's principal name is `joe@YOUR-REALM.COM`. You can use LDAP for this step if it is available in your organization.

> **Note:** Each account must have a user ID that is greater than or equal to 1000. In the `/etc/hadoop/conf/taskcontroller.cfg` file, the default setting for the `banned.users` property is `mapred`, `hdfs`, and `bin` to prevent jobs from being submitted using those user accounts. The default setting for the `min.user.id` property is `1000` to prevent jobs from being submitted with a user ID less than 1000, which are conventionally Unix super users.

2. Create a subdirectory under `/user` on HDFS for each user account (for example, `/user/joe`). Change the owner and group of that directory to be the user.

```
$ hadoop fs -mkdir /user/joe
$ hadoop fs -chown joe /user/joe
```

> **Note:** `sudo -u hdfs` is not included in the commands above. This is because it is not required if Kerberos is enabled on your cluster. You will, however, need to have Kerberos credentials for the HDFS super user to successfully run these commands. For information on gaining access to the HDFS super user account, see Step 5: Create the HDFS Superuser on page 58.

Step 8: Verify that Kerberos Security is Working

After you have Kerberos credentials, you can verify that Kerberos security is working on your cluster by trying to run MapReduce jobs. To confirm, try launching a sleep or a pi job from the provided Hadoop examples (`/usr/lib/hadoop/hadoop-examples.jar`).

> **Note:**
>
> This section assumes you have a fully-functional CDH cluster and you have been able to access HDFS and run MapReduce jobs before you followed these instructions to configure and enable Kerberos on your cluster. If you have not already done so, you should at a minimum use the Cloudera Manager Admin Console to generate a client configuration file to enable you to access the cluster. For instructions, see Deploying Client Configuration Files.

To verify that Kerberos security is working:

1. Acquire Kerberos credentials for your user account.

```
$ kinit USERNAME@YOUR-LOCAL-REALM.COM
```

2. Enter a password when prompted.
3. Submit a sample pi calculation as a test MapReduce job. Use the following command if you use a package-based setup for Cloudera Manager:

```
$ hadoop jar /usr/lib/hadoop-0.20-mapreduce/hadoop-examples.jar pi 10 10000
Number of Maps = 10
Samples per Map = 10000
...
Job Finished in 38.572 seconds
Estimated value of Pi is 3.14120000000000000000
```

If you have a parcel-based setup, use the following command instead:

```
$ hadoop jar /opt/cloudera/parcels/CDH/lib/hadoop-0.20-mapreduce/hadoop-examples.jar pi
 10 10000
Number of Maps = 10
Samples per Map = 10000
...
Job Finished in 30.958 seconds
Estimated value of Pi is 3.14120000000000000000
```

You have now verified that Kerberos security is working on your cluster.

Important:

Running a MapReduce job will fail if you do not have a valid Kerberos ticket in your credentials cache. You can examine the Kerberos tickets currently in your credentials cache by running the `klist` command. You can obtain a ticket by running the `kinit` command and either specifying a keytab file containing credentials, or entering the password for your principal. If you do not have a valid ticket, you will receive an error such as:

```
11/01/04 12:08:12 WARN ipc.Client:
Exception encountered while connecting to the server :
javax.security.sasl.SaslException:GSS initiate failed
[Caused by GSSException: No valid credentials provided (Mechanism level:
 Failed to find any
Kerberos tgt)]
Bad connection to FS. command aborted. exception: Call to
nn-host/10.0.0.2:8020 failed on local exception:
java.io.IOException:javax.security.sasl.SaslException: GSS initiate
failed
[Caused by GSSException: No valid credentials provided
(Mechanism level: Failed to find any Kerberos tgt)]
```

Step 9: (Optional) Enable Authentication for HTTP Web Consoles for Hadoop Roles

Minimum Required Role: Configurator (also provided by **Cluster Administrator, Full Administrator**)

Authentication for access to the HDFS, MapReduce, and YARN roles' web consoles can be enabled using a configuration option for the appropriate service. To enable this authentication:

1. From the **Clusters** tab, select the service (HDFS, MapReduce, or YARN) for which you want to enable authentication.
2. Click the **Configuration** tab.
3. Select **Scope** > ***service name*** **Service-Wide**.
4. Select **Category** > **Security**.
5. Type `Enable Kerberos` in the Search box.
6. Select **Enable Kerberos Authentication for HTTP Web-Consoles**.
7. Click **Save Changes** to commit the changes.
8. When the command finishes, restart all roles of that service.

Enabling SPNEGO as an Authentication Backend for Hue

1. In Cloudera Manager, set the authentication backend to `SpnegoDjangoBackend`.
   a. Go to the Cloudera Manager Admin Console. From the **Clusters** tab, select the Hue service.
   b. Click the **Configuration** tab.
   c. Select **Scope** > **Service-Wide**.
   d. Select **Category** > **Security**.
   e. Locate the **Authentication Backend** property and select `desktop.auth.backend.SpnegoDjangoBackend`.
   f. Click **Save Changes**.
2. Restart the Hue service.
3. If you are using an external load balancer, perform the following steps, otherwise skip the remaining steps. Cloudera Manager creates these configuration automatically:
   a. On the host running the Hue Kerberos Ticket Renewer, switch to the `KT_RENEWER` process directory. For example:

```
cd /var/run/cloudera-scm-agent/process/`ls -lrt /var/run/cloudera-scm-agent/process/
  \
| awk '{print $9}' |grep KT_RENEWER| tail -1`/
```

b. Verify that the Hue keytab includes the HTTP principal.

```
klist -kte ./hue.keytab

Keytab name: FILE:./hue.keytab
KVNO Timestamp Principal
---- ----------------- --------------------------------------------------------
1 03/09/15 20:20:35 hue/host-10-16-8-168.openstacklocal@EXAMPLE.CLOUDERA.COM
(aes128-cts-hmac-sha1-96)
1 03/09/15 20:20:36 HTTP/host-10-16-8-168.openstacklocal@EXAMPLE.CLOUDERA.COM
(aes128-cts-hmac-sha1-96)
```

c. Copy the `hue.keytab` file to `/var/lib/hue` and change ownership to the hue user and group.

```
$ cp ./hue.keytab /var/lib/hue/
$ chown hue:hue /var/lib/hue/hue.keytab
```

d. Go to the Cloudera Manager Admin Console. From the **Clusters** tab, select the Hue service.
e. Click the **Configuration** tab.
f. Select **Scope** > **Service-Wide**.
g. Select **Category** > **Advanced**.
h. Locate the **Hue Service Environment Advanced Configuration Snippet (Safety Valve)** property and add the following line:

```
KRB5_KTNAME=/var/lib/hue/hue.keytab
```

i. Click **Save Changes** to commit the changes.
j. Restart the Hue service.

Kerberos Authentication for Single User Mode and Non-Default Users

The steps described in this topic are only applicable in the following cases:

- You are running the Cloudera Manager in the single user mode. In this case, configure all the services described in the table below.

  OR

- You are running one or more CDH services with non-default users. This means if you have modified the default value for the **System User** property for any service in Cloudera Manager, you must only perform the command (as described below) corresponding to that service, to be able to successfully run jobs with the non-default user.

MapReduce	Configure the `mapred.system.dir` directory to be owned by the `mapred` user. `sudo -u hdfs hadoop fs -chown mapred:hadoop ${mapred.system.dir}` By default, `mapred.system.dir` is `/tmp/mapred/system`.
HBase	Give the `hbase` user ownership of the HBase root directory: `sudo -u hdfs hadoop fs -chown -R hbase ${hbase.rootdir}` By default, `hbase.rootdir` is `/hbase`.
Hive	Give the `hive` user ownership of the `/user/hive` directory. `sudo -u hdfs hadoop fs -chown hive /user/hive`

YARN	For every NodeManager host, for each path in `yarn.nodemanager.local-dirs`, run: `rm -rf ${yarn.nodemanager.local-dirs}/usercache/*` This removes the `/usercache` directory that contains intermediate data stored for previous jobs.

Customizing Kerberos Principals

By default, the Cloudera Manager Kerberos wizard configures CDH services to use the same Kerberos principals as the default process users. For example, the `hdfs` principal for the HDFS service, and the `hive` principal for the Hive service. The advantage to this is that when Kerberos is enabled, no HDFS directory permissions need to be changed for the new principals. However, starting with Cloudera Manager 5.4, you can configure custom service principals for CDH services.

Important Considerations

- Using different Kerberos principals for different services will make it easier to track the HDFS directories being accessed by each service.
- If you are using `ShellBasedUnixGroupsMapping` to obtain user-group mappings, ensure you have the UNIX accounts for the principals present on all hosts of the cluster.

Configuring Directory Permissions

Configure the following HDFS directories to give their corresponding custom service principals `read, write` and `execute` permissions.

Service	HDFS Directory
Accumulo	• HDFS Directory • `/user/`*`principal`
HBase	HBase Root Directory
Hive	• Hive Warehouse Directory • `/user/`*`principal`
Impala	`/user/`*`principal`
MapReduce v1	`/tmp/mapred`
Oozie	Oozie ShareLib Root Directory
Solr	HDFS Data Directory
Spark on YARN	• `/user/`*`principal` • Spark History Location • Spark Jar Location
Sqoop2	`/user/`*`principal`

Configuring CDH Services

The following services will require additional settings if you are using custom principals:

- **HDFS** - If you have enabled synchronization of HDFS and Sentry permissions, add the Hive and Impala principals to the **Sentry Authorization Provider Group** property.
  1. Go to the HDFS service.
  2. Click Configuration.
  3. Select **Scope** > **HDFS Service-Wide**.

4. Select **Category** > **Security**.
5. Locate the **Sentry Authorization Provider Group** property and add the custom Hive and Impala principals.
6. Click **Save Changes**.

- **YARN** - The principals used by YARN daemons should be part of hadoop group so that they are allowed to read JobHistory Server data.
- **Impala** - If you are running the Hue service with a custom principal, configure Impala to allow the Hue principal to impersonate other users.

  1. Go to the Impala service.
  2. Click Configuration.
  3. Select **Scope** > **Impala (Service-Wide)**.
  4. Select **Category** > **Policy File-Based Sentry**.
  5. Locate the **Proxy User Configuration** property and add the custom Hue principal.
  6. Click **Save Changes**.

- **Hive** - If the Sentry service is enabled, allow the Kerberos principals used by Hive, Impala, Hue, HDFS and the Service Monitor to bypass Sentry authorization in the Hive metastore.

  1. Go to the Hive service.
  2. Click Configuration.
  3. Select **Scope** > **Impala (Service-Wide)**.
  4. Select **Category** > **Policy File-Based Sentry**.
  5. Locate the **Bypass Sentry Authorization Users** property and add the custom Hive, Impala, Hue and HDFS principals to the list.
  6. Click **Save Changes**.

- **Spark on YARN** - The principal used by the Spark service should be part of the `spark` group.
- **Sentry** - Allow the Hive, Impala, Hue and HDFS principals to connect to the Sentry service.

  1. Go to the Sentry service.
  2. Click Configuration.
  3. Search for the **Allowed Connecting Users** property and add the custom Hive, Impala, Hue and HDFS principals to the list.
  4. Search for the **Admin Groups** property and include the groups to which the Hive, Impala, and Hue principals belong.
  5. Click **Save Changes**.

- **Cloudera Management Service** - Configure the Reports Manager principal and the Navigator principal for HDFS as HDFS superusers.

  1. Go to the Cloudera Management Service.
  2. Click Configuration.
  3. Search for *kerberos*.
  4. Locate the **Reports Manager Kerberos Principal** property and set it to a principal with administrative and superuser privileges on all HDFS services.
  5. Locate the **Navigator Kerberos Principal for HDFS** property and set it to a principal with administrative and superuser privileges on all HDFS services.
  6. Click **Save Changes**.

Incompatibilities

The following features do not work with custom principals:

- Llama must always use the default Kerberos principal `llama`.
- If you are using MapReduce v1, the Activity Monitor and Cloudera Navigator should use the same principal as the Hue service.

- If you are using the Java KeyStore KMS or KeyTrustee KMS with a custom principal, you will need to add the proxy user for the custom principal to the `kms-site.xml` safety valve.

  For example, if you've replaced the default `oozie` principal with `oozieprinc`, add the `hadoop.kms.proxyuser.oozieprinc.groups` and `hadoop.kms.proxyuser.oozieprinc.hosts` properties to the `kms-site.xml` safety valve.

Managing Kerberos Credentials Using Cloudera Manager

Minimum Required Role: Full Administrator

When Kerberos authentication is enabled for HDFS and MapReduce service instances, Cloudera Manager starts creating Kerberos principals for each role instance on the cluster at the end of the configuration process. Depending on the number of hosts and the number of HDFS and MapReduce role instances in the cluster, the process may take anywhere from a few seconds to several minutes.

After the process completes, view the list of Kerberos principals created for the cluster by using the Cloudera Manager Admin Console. Every host with HDFS and MapReduce role instances should have Kerberos principals.

If no principals have been created after 10 minutes, there may be an issue with the process. See Kerberos Credential-Generation Issues on page 421 to troubleshoot.

Important:

- Do not regenerate the principals for your cluster unless you have made a global configuration change, such as changing the encryption type.
- Regenerate principals using the Cloudera Manager Admin Console only. Do not use `kadmin` shell.
- For an MIT KDC, Configuring a Dedicated MIT KDC for Cross-Realm Trust on page 163 to avoid invalidating existing host keytabs.

Updating Kerberos Credentials in Cloudera Manager

If you change the user name or the password (or both) in the Active Directory KDC for the account used by Cloudera Manager for Kerberos authentication, you must also change it in Cloudera Manager. These credentials were stored during the Kerberos integration process (see Step 3: Add the Credentials for the Principal to the Cluster on page 379).

1. Log in to Cloudera Manager Admin Console.
2. Select **Administration** > **Security**.
3. Click the **Kerberos Credentials** tab.
4. Click the **Import Kerberos Account Manager Credentials** button.
5. Enter the new user and password for the principal added to the Kerberos KDC in Step 2. Create Principal for Cloudera Manager Server in the Kerberos KDC on page 378.

Managing Active Directory Account Properties

If you are using an Active Directory KDC, Cloudera Manager 5.8 (and higher) lets you configure Active Directory accounts and customize the credential regeneration process using the Cloudera Manager Admin Console. You can also use Cloudera Manager to configure the encryption types to be used by your Active Directory account. Once you modify any Active Directory account properties, you must regenerate Kerberos credentials to reflect those changes. The credential regeneration process requires you to delete existing accounts before new ones are created.

By default, Cloudera Manager does *not* delete accounts in Active Directory, which means that to regenerate Kerberos principals contained in Active Directory, you need to manually delete the existing Active Directory accounts. You can either delete and regenerate *all* existing Active Directory accounts, or only delete those with the userPrincipalName (or login name) that you will later manually select for regeneration. If the accounts haven't already been deleted manually, the regeneration process will throw an error message saying that deletion of accounts is required before you proceed.

Modifying Active Directory Account Properties Using Cloudera Manager

If you are using an Active Directory KDC, you can configure Active Directory account properties such as `objectClass` and `accountExpires` directly from the Cloudera Manager Admin Console. Any changes to these properties will be reflected in the regenerated Kerberos credentials.

To configure Active Diretory account properties:

1. Go to the Cloudera Manager Admin Console and click the **Administration** tab.
2. Select **Administration** > **Settings**.
3. Click the **Kerberos** category.
4. Locate the **Active Directory Account Properties** and edit as required. By default, the property will be set to:

```
accountExpires=0,objectClass=top,objectClass=person,objectClass=organizationalPerson,objectClass=user
```

5. Locate the **Active Directory Password Properties** and edit the field as needed. By default, the property will be set to:

```
length=12,minLowerCaseLetters=2,minUpperCaseLetters=2,minDigits=2,minSpaces=0,minSpecialChars=0,specialChars=?.!$%^*()-_+=~
```

6. Click **Save Changes**.
7. Regenerate new credentials with the new properties.

Enabling Credential Regeneration for Active Directory Accounts Using Cloudera Manager

To avoid having to delete accounts manually, enable the **Active Directory Delete Accounts on Credential Regeneration** property. By default, this property is disabled. After enabling this feature, Cloudera Manager will delete existing Active Directory accounts automatically when new ones are created during regeneration.

1. Go to the Cloudera Manager Admin Console and click the **Administration** tab.
2. Select **Administration** > **Settings**.
3. Click the **Kerberos** category.
4. Locate the **Active Directory Delete Accounts on Credential Regeneration** and check the property to activate this capability.
5. Click **Save Changes**.

Configuring Encryption Types for Active Directory KDC Using Cloudera Manager

Cloudera Manager allows you to configure the encryption types (or `enctype`) used by an Active Directory KDC to protect its data. Cloudera supports the following encryption types:

- `rc4-hmac`
- `aes128-cts`
- `aes256-cts`
- `des-cbc-crc`
- `des-cbc-md5`

To configure encryption types for an Active Directory KDC:

1. Go to the Cloudera Manager Admin Console and click the **Administration** tab.
2. Select **Administration** > **Settings**.
3. Click the **Kerberos** category.
4. Locate the **Kerberos Encryption Types** and click ✚ to add the encryption types you want Active Directory to use (see the list above for supported encryption types `enctypes`).
5. Check the checkbox for the **Active Directory Set Encryption Types** property. This will automatically set the Cloudera Manager AD account to use the encryption types configured in the previous step.
6. Click **Save Changes**.

Moving Kerberos Principals to Another OU Within Active Directory

If you have a Kerberized cluster configured with an Active Directory KDC, you can use the following steps to move the Kerberos principals from one AD Organizational Unit (OU) to another.

1. Create the new OU on the Active Directory Server.
2. Use AD's Delegate Control wizard to set the permissions on the new OU such that the configured Cloudera Manager admin account has the ability to **Create, Delete and Manage User Accounts** within this OU.
3. Stop the cluster.
4. Stop the Cloudera Management Service.
5. In Active Directory, move all the Cloudera Manager and CDH components' user accounts to the new OU.
6. Go to Cloudera Manager and go to **Administration** > **Security**.
7. Go to the **Kerberos Credentials** tab and click **Configuration**.
8. Select **Scope** > **Settings**.
9. Select **Category** > **Kerberos**.
10. Locate the **Active Directory Suffix** property and edit the value to reflect the new OU name.
11. Click **Save Changes**.

Viewing or Regenerating Kerberos Credentials Using Cloudera Manager

To view Kerberos principals for the cluster from Cloudera Manager for either MIT Kerberos or Active Diretory:

1. Log in to Cloudera Manager Admin Console.
2. Select **Administration** > **Security**.
3. Click the **Kerberos Credentials** tab. The currently configured Kerberos principals for services running on the cluster display, such as:
   - For HDFS, principals `hdfs/hostname` and `host/hostname`
   - For MapReduce, principals `mapred/hostname` and `host/hostname`

To regenerate any principals (if necessary):

- Select the principal from the list.
- Click **Regenerate**.

Running the Security Inspector

The Security Inspector uses the Host Inspector to run a security-related set of commands on the hosts in your cluster. It reports on matters such as how Java is configured for encryption and on the default realms configured on each host:

1. Select **Administration** > **Security**.
2. Click **Security Inspector**. Cloudera Manager begins several tasks to inspect the managed hosts.
3. After the inspection completes, click **Download Result Data** or **Show Inspector Results** to review the results.

Using a Custom Kerberos Keytab Retrieval Script

The Cloudera Manager Kerberos setup procedure requires you to create an administrator account for the Cloudera Manager user. Cloudera Manager will then connect to your KDC and use this admin account to generate principals and keytabs for the remaining CDH services. If for some reason, you cannot create a Cloudera Manager administrator account on your KDC with the privileges to create other principals and keytabs for CDH services, then these will need to be created manually.

Cloudera Manager gives you the option to use a custom script to retrieve keytabs from the local filesystem. To use a custom Kerberos keytab retrieval script:

1. The KDC administrators should create the required principals and keytabs, and store them securely on the Cloudera Manager Server host.
2. Create the keytab retrieval script. Your script should take two arguments: a full principal name for which it should retrieve a keytab, and a destination to which it can write the keytab. The script must be executable by the Cloudera Manager admin user, `cloudera-scm`. Depending on the principal name input by Cloudera Manager, the script

should locate the corresponding keytab on the Cloudera Manager Server host (stored in step 1), and copy it into a location accessible to the `cloudera-scm` user. Here is a simple example:

```
#!/bin/bash

# Cloudera Manager will input a destination path
DEST="$1"

# Cloudera Manager will input the principal name in the format: <service>/<fqdn>@REALM
PRINC="$2"

# Assuming the '<service>_<fqdn>@REALM.keytab' naming convention for keytab files
IN=$(echo $PRINC | sed -e 's/\//_/')
SRC="/keytabs/${IN}.keytab"

# Copy the keytab to the destination input by Cloudera Manager
cp -v $SRC $DEST
```

Note that the script will change according to the keytab naming convention followed by your organization.

3. Configure the location for the script in Cloudera Manager:
   a. Go to the Cloudera Manager Admin console.
   b. Select **Administration** > **Settings**.
   c. Select **Category** > **Kerberos**.
   d. Locate the **Custom Kerberos Keytab Retrieval Script** and set it to point to the script created in step 2.
   e. Click **Save Changes**.
4. Once the **Custom Kerberos Keytab Retrieval Script** property is set, whenever Cloudera Manager needs a keytab, it will ignore all other Kerberos configuration and run the keytab retrieval script to copy the required keytab to the desired destination.
5. Cloudera Manager can now distribute the keytab to the services that need access to it.

> **Note:** The Cloudera Navigator web server accesses HDFS and Hue using the keytabs corresponding to those principals; however the custom script does not move these additional keytabs to the Navigator Metadata Server. To complete the setup for Navigator, move keytabs for HDFS and Hue principals to the Navigator Metadata Server host manually. See Configuring Navigator Metadata Server for Kerberos and Custom Keytab Retrieval.

Adding Trusted Realms to the Cluster

The Kerberos instance associated with a given cluster has its *REALM-NAME* specified as the `default_realm` in the Kerberos configuration file (`krb5.conf`) on the cluster's NameNode. Rules defined in the `hadoop.security.auth_to_local` property translate the Kerberos principals to local account names at the host level (see Hadoop Users (user:group) and Kerberos Principals on page 103 and Mapping Kerberos Principals to Short Names on page 108). The default rules simply remove the `@REALM` portion of the Kerberos principal and leave the short name.

To allow principals from other realms to use the cluster, the trusted realms must be specified in Cloudera Manager. For example, the Kerberos realm used by your cluster may have a trust relationship to a central Active Directory or MIT Kerberos realm. Add the central realm to the cluster as detailed in the following steps so that Cloudera Manager can apply the appropriate mapping rules to the principals from the trusted realm to access the cluster's services.

To specify trusted realms using Cloudera Manager:

1. Log in to the Cloudera Manager Admin Console.
2. Select **Clusters** > **HDFS Service**.
3. Click the **Configuration** tab.
4. Select **HDFS (Service-Wide)** for the **Scope** filter.
5. Select **Security** for the **Category** filter.

6. In the Search field, type `Kerberos Realms` to find the **Trusted Kerberos Realms** and **Additional Rules to Map Kerberos Principals to Short Names** settings.
7. Add other Kerberos realms that the cluster's realm can trust. Use all upper-case letters to specify the REALM name for MIT Kerberos or Active Directory realms:

```
ANOTHER-REALM.EXAMPLE.COM
```

   To add multiple realms, click the plus (+) button.
8. Click **Save Changes**.

For each trusted realm identified in **Trusted Kerberos Realms**, default mapping rules automatically strip the REALM name. To customize the mapping rules, specify additional rules in the **Additional Rules to Map Kerberos Principals to Short Names** setting, one rule per line. Cloudera Manager will wrap each rule in the appropriate XML tags and add to the generated `core-site.xml` file. To create custom rules and translate translate mixed-case Kerberos principals to lower-case Hadoop usernames, see Mapping Rule Syntax on page 108.

If you specify custom mapping rules for a Kerberos realm using the **Additional Rules to Map Kerberos Principals to Short Names** setting, ensure that the same realm is not specified in the **Trusted Kerberos Realms** setting. If it is, the auto-generated rule (which only strips the realm from the principal and does no additional transformations) takes precedent, and the custom rule is ignored.

For these changes to take effect, you must restart the cluster and redeploy the client configuration, as follows:

1. On the Cloudera Manager Admin Console, **Clusters** > **Cluster-*n*** to choose cluster-wide actions.
2. From the **Actions** drop-down button, select **Deploy Client Configuration**.

Using Auth-to-Local Rules to Isolate Cluster Users

By default, the Hadoop auth-to-local rules map a principal of the form `<username>/<hostname>@<REALM>` to `<username>`. This means if there are multiple clusters in the same realm, then principals associated with hosts of one cluster would map to the same user in all other clusters.

For example, if you have two clusters, `cluster1-host-[1..4].example.com` and `cluster2-host-[1..4].example.com`, that are part of the same Kerberos realm, `EXAMPLE.COM`, then the `cluster2` principal, `hdfs/cluster2-host1.example.com@EXAMPLE.COM`, will map to the `hdfs` user even on `cluster1` hosts.

To prevent this, use auth-to-local rules as follows to ensure only principals containing hostnames of `cluster1` are mapped to legitimate users.

1. Go to the **HDFS Service** > **Configuration** tab.
2. Select **Scope** > **HDFS (Service-Wide)**.
3. Select **Category** > **Security**.
4. In the Search field, type `Additional Rules` to find the **Additional Rules to Map Kerberos Principals to Short Names** settings.
5. Additional mapping rules can be added to the **Additional Rules to Map Kerberos Principals to Short Names** property. These rules will be inserted before the rules generated from the list of trusted realms (configured above) and before the default rule.

```
RULE:[2:$1/$2@$0](hdfs/cluster1-host1.example.com@EXAMPLE.COM)s/(.*)@EXAMPLE.COM/hdfs/
RULE:[2:$1/$2@$0](hdfs/cluster1-host2.example.com@EXAMPLE.COM)s/(.*)@EXAMPLE.COM/hdfs/
RULE:[2:$1/$2@$0](hdfs/cluster1-host3.example.com@EXAMPLE.COM)s/(.*)@EXAMPLE.COM/hdfs/
RULE:[2:$1/$2@$0](hdfs/cluster1-host4.example.com@EXAMPLE.COM)s/(.*)@EXAMPLE.COM/hdfs/
RULE:[2:$1/$2@$0](hdfs.*@EXAMPLE.COM)s/(.*)@EXAMPLE.COM/nobody/
```

   In the example, the principal `hdfs/<hostname>@REALM` is mapped to the `hdfs` user if `<hostname>` is one of the cluster hosts. Otherwise it gets mapped to nobody, thus ensuring that principals from other clusters do not have access to `cluster1`.

If the cluster hosts can be represented with a regular expression, that expression can be used to make the configuration easier and more conducive to scaling. For example:

```
RULE:[2:$1/$2@$0](hdfs/cluster1-host[1-4].example.com@EXAMPLE.COM)s/(.*)@EXAMPLE.COM/hdfs/
RULE:[2:$1/$2@$0](hdfs.*@EXAMPLE.COM)s/(.*)@EXAMPLE.COM/nobody/
```

Note:

It is not possible to use alternatives in capturing or non-capturing groups within the matching portion of the rule, because the use of round brackets in the expression is not currently supported. For example, the following rule would result in an error:

```
RULE:[2:$1/$2@$0](hdfs/cluster1-(hosta|hostb|hostc).example.com@EXAMPLE.COM)s/(.*)@EXAMPLE.COM/hdfs/
```

6. Click **Save Changes**.
7. Restart the HDFS service and any dependent services.

Configuring Authentication for Cloudera Navigator

Cloudera Manager Server has an ***internal*** authentication mechanism, a database repository of user accounts that can be used to create user accounts. As an alternative to using the internal database, Cloudera Manager and Cloudera Navigator can be configured to use ***external*** authentication mechanisms.

Cloudera Manager Server and Cloudera Navigator each have their own user role schemes for granting privileges to system features and functions. Cloudera Manager user roles can be applied to user accounts as they are created in the internal repository. Cloudera Navigator user roles are applied to groups defined in the external system for use by Cloudera Navigator. The only user role that can be effectively applied to an account created in the Cloudera Manager internal repository is that of Navigator Administrator, which grants the user account privileges as a Full Administrator on the Cloudera Navigator services (Navigator Metadata Server, Navigator Audit Server).

In other words, assigning Cloudera Navigator user roles to user accounts requires using an external authentication mechanism.

Cloudera Navigator and External Authentication

To support its user role-based authorization scheme, Cloudera Navigator integrates with **external authentication mechanisms**. External authentication mechanisms include:

- LDAP-compliant identity/authentication services, such as Active Directory and OpenLDAP
- SAML-based SSO solutions, such as Shibboleth and SiteMinder

Cloudera Manager Server has its own internal authentication mechanism, a database repository of user accounts. However, the user accounts defined in the internal Cloudera Manager account repository cannot be assigned Cloudera Navigator user roles. The only user role that can be effectively applied to an account created in the Cloudera Manager internal repository is that of Navigator Administrator. In other words, assigning Cloudera Navigator user roles to user accounts requires using one of the external authentication mechanisms detailed in this section.

Cloudera Manager and Cloudera Navigator have their own distinct sets of user roles. Cloudera Manager and Cloudera Navigator can be configured to use external authentication mechanisms. The organization may have a central Active Directory or other LDAP-identity service used by Cloudera Manager and Cloudera Navigator for external authentication, but the relationship between each of these to the external system functions independently. That means a Cloudera Manager user that successfully authenticates to the external LDAP system cannot log in to Cloudera Navigator using that same authentication token.

How it Works: Cloudera Navigator and External Authentication

At runtime, the Navigator Metadata Server role instance (the daemon) forwards login requests from Cloudera Navigator users to the external authentication mechanism which has a repository containing user accounts and groups that have

been setup for Cloudera Navigator users. The groups have had specific Cloudera Navigator user role assigned to them, so once users authenticate to the external system, they can use the features of Cloudera Navigator console as specified for their group.

All this occurs transparently to Cloudera Navigator users, assuming Cloudera Navigator has been correctly configured as detailed in the appropriate section for the external mechanism—Active Directory, OpenLDAP, or SAML—as detailed in this section.

Configuring Cloudera Navigator for Active Directory

To configure Cloudera Navigator for external authentication:

1. Log in to Cloudera Manager Admin Console.
2. Select **Clusters** > **Cloudera Management Service**.
3. Click the **Configuration** tab.
4. Select **Navigator Metadata Server** for the **Scope** filter.
5. Select **External Authentication** for the **Category** filter.
6. Leave the **Authentication Backend Order** set to the default value—**Cloudera Manager Only** until *after* the external system has been successfully configured for Cloudera Navigator (as detailed in these steps) and user accounts in Active Directory instance are members of groups that have been granted Cloudera Navigator user role privileges. When Cloudera Navigator receives a login request, it checks user repositories in the order specified. Checking only the external system before having user accounts and roles configured can result in authentication failures.
   - If user accounts and groups for Cloudera Navigator already exist in the Active Directory and a group with privileges for Cloudera Manager Full Administrator or Navigator Administrator user roles contains user accounts—so that the system can be managed—the order can be set to **External then Cloudera Manager** or External Only.
7. Configure the remaining settings for the Active Directory instance as detailed in the table.

Property	Description and usage note
External Authentication Type	Active Directory
LDAP URL	Full path to the Active Directory instance, including the protocol specifier, `ldap` or `ldaps` (for TLS/SSL). Not necessary to specify port number if the Active Directory service is hosted using the default ports—389 (LDAP), 636 (LDAPS). For example: `ldap://ad-srv.ldap-srvs.subnet.example.com`
LDAP Bind User Distinguished Name	The user name that connects to the Active Directory service to look up login requests on behalf of Cloudera Navigator. Enter either the complete user principal name or just the short name. For example, `cn-admin@EXAMPLE.COM` or `cn-admin`. For Active Directory, this distinguished name (DN) corresponds to the `sAMAcountName`.
LDAP Bind Password	Enter the password used to log in to the Active Directory instance using the DN specified for the bind user.
Active Directory Domain	The fully-qualified domain name of the Active Directory domain controller host system. This is the service to which the bind operation For example: *`ldap-servs.subnet.example.com`*
LDAP Distinguished Name Pattern	Leave blank if **LDAP User Search Base** is set.

Property	Description and usage note
LDAP User Search Base	Specify the organizational unit (OU) and domain component (DC) properties for the LDAP search tree. For example: `ou=nav_people,dc=ldap-srvs,dc=subnet,dc=example,dc=com`
LDAP User Search Filter	Optional.
LDAP Group Search Base	`ou=nav_groups,dc=ldap-srvs,dc=subnet,dc=example,dc=com`
LDAP Group Search Filter For Logged In User	Optional.
LDAP Groups Search Filter	`(&(objectClass=groupOfNames)(cn=*{0}*))`

8. Click **Save Changes**.
9. Restart the Navigator Metadata Service:
   - From **Cloudera Management Service**, click the **Instances** tab.
   - Select **Navigator Metadata Service** from among the instances listed.
   - Click the **Actions for Selected** button and select **Restart**.

Configuring Cloudera Navigator for LDAP

There is more than one way to configure an LDAP-compatible identity/authentication service to authenticate Navigator users and to allow Navigator to access group membership. The most commonly used LDAP authentication mode is **search-bind**.

Alternatively, you can configure Navigator to use **bind-direct** LDAP authentication. Note that this method assumes all users are in one directory tree; it doesn't have the flexibility to identify multiple LDAP directory trees.

Considerations for Configuring Navigator with LDAPS (TLS/SSL)

It is important that you configure TLS/SSL for Cloudera Manager as part of using LDAP for authenticating users. If you don't have secure network communication, you risk exposing your LDAP Distinguished Name credentials on the network.

For information on how to generate and distribute certificate files, see Generate TLS Certificates on page 195.

Configuring LDAP as Navigator's External Authentication

1. Make sure that you have the TLS truststore for communicating with the LDAP server configured on the host where the Navigator Metadata Server role is running.
2. Decide which LDAP authentication method you want to use to locate user entries in the directory.

   In enterprise environments, the more common method is Search-Bind because it allows a flexible configuration for determining the Distinguished Name based on a user's login. In some cases, Direct Bind may be used when all users must be under a single branch in the LDAP directory. There's more information about the requirements for configuration properties for these modes in the table below.

3. Select **Clusters** > **Cloudera Management Service**.
4. Click the **Configuration** tab.
5. Select **Scope** > **Navigator Metadata Server**.
6. Select **Category** > **External Authentication**.
7. In the **External Authentication Type** property, select **LDAP**.
8. Configure the remaining settings for the LDAP server instance as detailed in the table.

Property	Description and usage note
External Authentication Type	LDAP
LDAP URL	Full path to the LDAP server instance, including the protocol specifier, `ldap` or `ldaps` (for TLS/SSL). It is not necessary to specify port number if the Active Directory service is hosted using the default ports—389 (LDAP), 636 (LDAPS). For example: `ldaps://ldap-server.corp.com` If you have not configured TLS, the URL would start with `ldap://`.
Authentication Backend Order	To instruct Navigator to refer to the LDAP server to authenticate users, set this property to **External then Cloudera Manager** or **External Only**. The "External then Cloudera Manager" option allows administrators to log into Navigator in the case where LDAP logins are failing. If you need to debug connectivity issues without LDAP involved, you can log in with a Cloudera Manager-only user account.
LDAP Bind User Distinguished Name	The LDAP account that has permission to query the LDAP database of user accounts on behalf of Cloudera Navigator. The bind user can be specified as the full distinguished name (DN) (`cn=account,ou=people,dc=corp,dc=region`) or as only the common name (*`user@domain`*). Use the same format as the string used for Cloudera Manager LDAP configuration. This property is not required if your LDAP server accepts anonymous binding.
LDAP Bind Password	The password for the bind user.
LDAP Group Search Base	Specify the organizational unit (OU) and domain component (DC) properties for the LDAP search tree where Navigator searches for groups. For example, this search base starts the group membership search to an organizational unit configured specifically for Navigator users and Navigator searches this subtree: `ou=nav_people,dc=ldap-srvs,dc=subnet,dc=example,dc=com` The Group Search Base is used for both the Groups Search Filter and the Group Search Filter for Logged in User.
LDAP Group Search Filter For Logged In User	Optional. Specifies how a user is determined to be a member of a group. Defaults to: `member={0}` where `{0}` is replaced with the **DN** of the user you are authenticating. For a filter requiring the username, use `{1}`, as `memberUid={1}`.
LDAP Groups Search Filter	Specify this filter to refine the scope of LDAP groups to associate with Navigator roles. LDAP groups and users belonging to these groups inherit the Navigator roles. The roles are then used for authorization. For example, to limit the groups to groups with the current user as a member, specify `(&(objectClass=group)(cn=*{0}*))` The Groups Search Filter is combined with the Group Search Base to define the group lookup.
Search-Bind Authentication Mode	

Property	Description and usage note
LDAP User Search Base	Specify the organizational unit (OU) and domain component (DC) properties for the LDAP search tree where Navigator searches to authenticate users. You can also specify a User Search Filter to further reduce the scope of the search. For example, the following User Search Base string limits Navigator authentication to users inside the specified OU and DCs: `ou=nav_people,dc=ldap-srvs,dc=subnet,dc=example,dc=com` If you leave out OUs from the search base, Navigator authenticates users from any OU. For example, with the User Search Base `dc=corp,dc=com` and the user search filter as `uid={0}`, Navigator searches for the user anywhere in the tree starting from the User Search Base. If LDAP includes two OUs—`ou=Engineering` and `ou=Operations`—Navigator finds user "foo" if it exists in either of these OUs, that is, `uid=foo,ou=Engineering,dc=corp,dc=com` or `uid=foo,ou=Operations,dc=corp,dc=com`.
LDAP User Search Filter	Used with the User Search Base to further limit the scope of the search for a directory entry that matches the credentials of the user logging into Navigator. Use a user search filter along with a DN pattern so that the search filter provides a fallback if the DN pattern search fails. For example, set the filter to `uid={0}` to use the username provided at login. For Active Directory's LDAP server, use `sAMAccountName={0}`.
Direct-Bind Authentication Mode	
LDAP Distinguished Name Pattern	Direct-bind authentication can be used if search is not required to determine the DN needed to bind to the LDAP server. Leave this property blank if **LDAP User Search Base** is set. To use this authentication mode, all users must be under a single branch in the LDAP directory. For example, to search for a distinguished name where the `uid` attribute is the username at login, you might provide a pattern such as `uid={0},ou=People,dc=corp,dc=com` where `{0}` indicates the username of the authenticating user. If a user provides the username "foo" at the Navigator login page, Navigator searches for the DN `uid=foo,ou=People,dc=corp,dc=com`.

9. Click **Save Changes**.
10. Restart the Navigator Metadata Server.

Configuring Cloudera Navigator for SAML

Cloudera Navigator supports SAML (Security Assertion Markup Language), an XML-based open standard data format for exchanging authentication and authorization details between identity providers and service providers. One of the main benefits of SAML is that it enables Single Sign-on (SSO) for browser-based clients, such as the Cloudera Navigator console. That means that you can integrate Cloudera Navigator with SSO solutions such as Shibboleth or CA Single Sign-On (formerly, SiteMinder).

> **Note:** Not all SAML identity providers are interoperable. Cloudera Navigator has been tested with Shibboleth and CA Single Sign-On.

Overview of SAML and SSO

The steps below assume you have a functioning SAML IDP already deployed. Here is a brief summary of some of the high level details as background to the configuration tasks:

- An ***identity provider*** or IDP is one of the main functions provided by an organization's SAML/SSO solution. The IDP provides identity assertions (tokens) to service providers that want to identify users when those users request access. service.
- A ***service provider*** or SP, such as Cloudera Navigator, protects itself from unauthorized access by checking the identity of users requesting the service against the IDP. When the SP gets back the assertion from the IDP, the service gives the requesting user the level of access for that user.
- The service's users or ***principals*** obtain access to the SP when they open their browsers to the URL of the service. Transparently to users, the SP—Cloudera Navigator, but specifically the web service hosted on the Navigator Metadata Server—sends an authentication request to the IDP through the user agent (browser) and obtains an identity assertion back from the IDP.

There are two properties that control whether the SSO process is initiated from the IDP or the SP (Navigator):

- SAML Login URL `nav.saml.login.url`
- Skip Authorization Check `nav.auth.skip_saml_auth_check` (set in **Navigator Metadata Server Advanced Configuration Snippet (Safety Valve) for cloudera-navigator.properties**)

As shown in the following table, when SAML authentication is enabled, the default Navigator login URL always produces a service-provider initiated SSO process (SP initiated SSO); in this process, users log into the Navigator login page and Navigator sends the authentication request to the identity provider. Specifying a SAML Login URL produces an identity provider-initiated SSO process (IdP initiated SSO); in this process, users log into the identity provider's login page and are redirected to the Navigator console and logged in.

SAML Login URL Property	Skip Authorization Check Property	No addition to the login URL (/)	/login.html	/locallogin.html
Set	True	SP initiated SSO	IdP initiated SSO	Login page
Set	False	SP initiated SSO	IdP initiated SSO	IdP initiated SSO
Not set	True	SP initiated SSO	SP initiated SSO	Login page
Not set	False	SP initiated SSO	SP initiated SSO	SP initiated SSO

See the OASIS SAML Wiki for more information about SAML.

Preparing Files

You must obtain the following files and information and provide to Cloudera Navigator:

- A Java keystore containing a private key for Cloudera Navigator to use to sign/encrypt SAML messages.
- The SAML metadata XML file from your IDP. This file must contain the public certificates needed to verify the sign/encrypt key used by your IDP per the SAML Metadata Interoperability Profile.
- The entity ID that should be used to identify the Navigator Metadata Server instance.
- How the user ID is passed in the SAML authentication response:
  - As an attribute. If so, what identifier is used.
  - As the NameID.
- The method by which the Cloudera Navigator role will be established:
  - From an attribute in the authentication response:
    - What identifier will be used for the attribute
    - What values will be passed to indicate each role
  - From an external script that will be called for each use:
    - The script takes user ID as $1

- The script must assign an exit code to reflect successful authentication of the assigned role:
  - 0 - Full Administrator
  - 1 - User Administrator
  - 2 - Auditing Viewer
  - 4 - Lineage Viewer
  - 8 - Metadata Administrator
  - 16 - Policy Viewer
  - 32 - Policy Administrator
  - 64 - Custom Metadata Administrator
  - A negative value is returned for a failure to authenticate

  To assign more than one role, add the numbers for the roles. For example, to assign the Policy Viewer and User Administrator roles, the exit code should be 17.

Configuring the Navigator Metadata Server

1. Select **Clusters** > **Cloudera Management Service**.
2. Click the **Configuration** tab.
3. Select **Navigator Metadata Server** from the **Scope** filter.
4. Select **External Authentication** from the **Category** filter.
5. Type `SAML` in the Search box to display only the SAML relevant settings.
6. Enter the values for the properties in the table based on your SAML implementation.

Property	Description and usage note
Authentication Backend Order	Set to **External then Cloudera Manager**.
External Authentication Type	SAML
Path to SAML IDP Metadata File	Set to the location (complete path) of the metadata file obtained from the IDP.
Path to SAML Keystore File	Path to the Java keystore file containing the Cloudera Navigator private key (prepared above).
SAML Keystore Password	Enter the SAML keystore password.
Alias of SAML Sign/Encrypt Private Key	Enter the alias used to identify the private key for Cloudera Navigator to use.
SAML Sign/Encrypt Private Key Password	Enter the password for the sign/encrypt private key.
SAML Entity ID	Default setting is `clouderaNavigator`. Leave set to the default unless more than one Cloudera Navigator instance is using the same IDP. Each Cloudera Navigator instance needs a unique entity ID as assigned by organizational policy.
SAML Entity Base URL	
SAML Response Binding	HTTP-Artifact (selected by default), or HTTP-Post
SAML Login URL	(IDP-initiated SSO only) If your IDP does not support SP-initiated SSO (very uncommon), specify the user login URL for the IDP.
Source of User ID in SAML Response	Attribute (selected by default), or NameID—Specifies the source of the user ID, an attribute or NameID. For attribute, also set the attribute name in the **SAML Attribute Identifier for User ID** property.

Property	Description and usage note
SAML Attribute Identifier for User ID	`urn:oid:0.9.2342.19200300.100.1.1` (Default) The standard object identifier (OID) for user IDs. This setting is used only when **Source of User ID in SAML Response** specifies Attribute.
SAML Role Assignment Mechanism	Attribute (selected by default), or Script—Specifies how user roles are assigned to authenticated user: • Attribute—Set the **SAML Attribute Identifier for User Role** and the **SAML Attribute Values for Roles** properties. • Script—A binary or shell script executable that assigns user roles. Set the path to the executable in Path to SAML Role Assignment Script.
SAML Attribute Identifier for User Role	`urn:oid:2.5.4.11` (Default) The standard OID typically used for OrganizationalUnits. Can be left to this setting.
SAML Attribute Values for Roles	Set the attribute values that will be used to indicate the user roles. For more than one role, the attribute can return comma-separated values, such as "role1, role2".
Path to SAML Role Assignment Script	Path to executable (binary or shell script) that assigns Cloudera Navigator user roles upon authentication. Required when **SAML Role Assignment Mechanism** specifies **Script**.

7. Click **Save Changes**.
8. Restart the Navigator Metadata Server role.

Configuring the Identity Provider

After Cloudera Navigator restarts, attempted logins to Cloudera Navigator are redirected to the identity provider's login page. Authentication cannot succeed until the IDP is configured for Cloudera Navigator. The configuration details are specific to the IDP, but in general you must download the SAML file from your Cloudera Navigator instance and perform the other steps below.

1. Download Cloudera Navigator's SAML metadata XML file from your Cloudera Navigator instance:

```
http://hostname:7187/saml/metadata
```

2. Inspect the metadata file and ensure that any URLs contained in the file can be resolved by users' web browsers. The IDP will redirect web browsers to these URLs at various points in the process. If the browser cannot resolve them, authentication will fail. If the URLs are incorrect, you can manually fix the XML file or set the **SAML Entity Base URL** property in the Navigator Metadata Server configuration to the correct value, and then re-download the file.
3. Provide this metadata file to your IDP using whatever mechanism your IDP provides.
4. Ensure that the IDP has access to whatever public certificates are necessary to validate the private key that was provided by Cloudera Navigator earlier.
5. Ensure that the IDP is configured to provide the User ID and Role using the attribute names that Cloudera Navigator was configured to expect, if relevant.
6. Ensure the changes to the IDP configuration have taken effect (a restart may be necessary).

Testing Cloudera Navigator and the SSO Setup

1. Return to the Cloudera Navigator home page at: `http://hostname:7187/`.
2. Attempt to log in with credentials for a user that is entitled. The authentication should complete and you should see the Home page.
3. If authentication fails, you will see an IDP provided error message. Cloudera Navigator is not involved in this part of the process, and you must ensure the IDP is working correctly to complete the authentication.
4. If authentication succeeds but the user is not authorized to use Cloudera Navigator, they will be taken to an error page that explains the situation. If a user who should be authorized sees this error, then you will need to verify

their role configuration, and ensure that it is being properly communicated to the Navigator Metadata Server, whether by attribute or external script. The Cloudera Navigator log will provide details on failures to establish a user's role. If any errors occur during role mapping, Cloudera Navigator will assume the user is unauthorized.

Bypassing SAML SSO

The SAML-based SSO can be bypassed by accessing the Cloudera Navigator login page directly:

```
http://fqdn-1.example.com:7187/locallogin.html
```

You can turn off this bypass by setting the Skip Authorization Check property (`nav.auth.skip_saml_auth_check`) in the **Navigator Metadata Server Advanced Configuration Snippet (Safety Valve) for cloudera-navigator.properties**.

Configuring Groups for Cloudera Navigator

Cloudera Navigator user role privileges are applied to groups contained in the external authentication mechanism. In the external system using the appropriate management tool for the platform, create groups comprising all the user accounts to which a specific set of privileges should apply. For example, create one group for auditors and another group for data stewards. Create a different group for administrators. To each group, add the user accounts of people who should have the same set of Cloudera Navigator privileges. The user accounts and groups must exist in the external authentication mechanism before the user roles can be applied to them.

Once a group exists in the external authentication system, one or more Cloudera Navigator user roles can be assigned to the group using the Cloudera Navigator console. See Configuring Navigator User Roles in the Cloudera Navigator Data Management guide

Configuring Authentication in CDH Using the Command Line

The security features in CDH 5 enable Hadoop to prevent malicious user impersonation. The Hadoop daemons leverage Kerberos to perform user authentication on all remote procedure calls (RPCs). Group resolution is performed on the Hadoop master nodes, NameNode, JobTracker and ResourceManager to guarantee that group membership cannot be manipulated by users. Map tasks are run under the user account of the user who submitted the job, ensuring isolation there. In addition to these features, new authorization mechanisms have been introduced to HDFS and MapReduce to enable more control over user access to data.

The security features in CDH 5 meet the needs of most Hadoop customers because typically the cluster is accessible only to trusted personnel. In particular, Hadoop's current threat model assumes that users cannot:

1. Have `root` access to cluster machines.
2. Have `root` access to shared client machines.
3. Read or modify packets on the network of the cluster.

> **Note:**
>
> CDH 5 supports encryption of all user data sent over the network. For configuration instructions, see Configuring Encrypted Shuffle, Encrypted Web UIs, and Encrypted HDFS Transport.
>
> Note also that there is no built-in support for on-disk encryption.

Enabling Kerberos Authentication for Hadoop Using the Command Line

Important:

These instructions assume you know how to install and configure Kerberos, you already have a working Kerberos Key Distribution Center (KDC) and realm setup, and that you've installed the Kerberos user packages on all cluster machines and machines which will be used to access the cluster. Furthermore, Oozie and Hue require that the realm support renewable tickets. For more information about installing and configuring Kerberos, see:

- MIT Kerberos Home
- MIT Kerberos Documentation
- Kerberos Explained
- Microsoft Kerberos Overview
- Microsoft Kerberos in Windows Server 2008
- Microsoft Kerberos in Windows Server 2003

Kerberos security in CDH 5 has been tested with the following version of MIT Kerberos 5:

- krb5-1.6.1 on Red Hat Enterprise Linux 5 and CentOS 5

Kerberos security in CDH 5 is supported with the following versions of MIT Kerberos 5:

- krb5-1.6.3 on SUSE Linux Enterprise Server (SLES) 11 Service Pack 1
- krb5-1.8.1 on Ubuntu
- krb5-1.8.2 on Red Hat Enterprise Linux 6 and CentOS 6
- krb5-1.9 on Red Hat Enterprise Linux 6.1

Note: The `krb5-server` package includes a `logrotate` policy file to rotate log files monthly. To take advantage of this, install the `logrotate` package. No additional configuration is necessary.

If you want to enable Kerberos SPNEGO-based authentication for the Hadoop web interfaces, see the Hadoop Auth, Java HTTP SPNEGO Documentation.

Here are the general steps to configuring secure Hadoop, each of which is described in more detail in the following sections:

Step 1: Install CDH 5

Cloudera strongly recommends that you set up a fully-functional CDH 5 cluster before you begin configuring it to use Hadoop's security features. When a secure Hadoop cluster is not configured correctly, the resulting error messages are in a preliminary state, so it's best to start implementing security after you are sure your Hadoop cluster is working properly without security.

For information about installing and configuring Hadoop and CDH 5 components, and deploying them on a cluster, see Cloudera Installation.

Step 2: Verify User Accounts and Groups in CDH 5 Due to Security

Note: CDH 5 introduces a new version of MapReduce: MapReduce 2.0 (MRv2) built on the YARN framework. In this document, we refer to this new version as YARN. CDH 5 also provides an implementation of the previous version of MapReduce, referred to as MRv1 in this document.

- If you are using MRv1, see Step 2a (MRv1 only): Verify User Accounts and Groups in MRv1 on page 80 for configuration information.
- If you are using YARN, see Step 2b (YARN only): Verify User Accounts and Groups in YARN on page 81 for configuration information.

Step 2a (MRv1 only): Verify User Accounts and Groups in MRv1

> **Note:** If you are using YARN, skip this step and proceed to Step 2b (YARN only): Verify User Accounts and Groups in YARN.

During CDH 5 package installation of MRv1, the following Unix user accounts are automatically created to support security:

This User	Runs These Hadoop Programs
`hdfs`	HDFS: NameNode, DataNodes, Secondary NameNode (or Standby NameNode if you are using HA)
`mapred`	MRv1: JobTracker and TaskTrackers

The `hdfs` user also acts as the HDFS superuser.

The `hadoop` user no longer exists in CDH 5. If you currently use the `hadoop` user to run applications as an HDFS super-user, you should instead use the new `hdfs` user, or create a separate Unix account for your application such as `myhadoopapp`.

MRv1: Directory Ownership in the Local File System

Because the HDFS and MapReduce services run as different users, you must be sure to configure the correct directory ownership of the following files on the local filesystem of each host:

File System	Directory	Owner	Permissions
Local	`dfs.namenode.name.dir` (`dfs.name.dir` is deprecated but will also work)	`hdfs:hdfs`	drwx------
Local	`dfs.datanode.data.dir` (`dfs.data.dir` is deprecated but will also work)	`hdfs:hdfs`	drwx------
Local	`mapred.local.dir`	`mapred:mapred`	drwxr-xr-x

See also Deploying MapReduce v1 (MRv1) on a Cluster.

You must also configure the following permissions for the HDFS and MapReduce log directories (the default locations in `/var/log/hadoop-hdfs` and `/var/log/hadoop-0.20-mapreduce`), and the `$MAPRED_LOG_DIR/userlogs/` directory:

File System	Directory	Owner	Permissions
Local	`HDFS_LOG_DIR`	`hdfs:hdfs`	drwxrwxr-x
Local	`MAPRED_LOG_DIR`	`mapred:mapred`	drwxrwxr-x
Local	`userlogs` directory in `MAPRED_LOG_DIR`	`mapred:anygroup`	*permissions will be set automatically at daemon start time*

[1] **In CDH 5, package installation and the Hadoop daemons will automatically configure the correct permissions for you if you configure the directory ownership correctly as shown in the table above.**

MRv1: Directory Ownership on HDFS

The following directories on HDFS must also be configured as follows:

File System	Directory	Owner	Permissions
HDFS	`mapreduce.jobtracker.system.dir` (`mapred.system.dir` is deprecated but will also work)	`mapred:hadoop`	drwx------
HDFS	/ (root directory)	`hdfs:hadoop`	drwxr-xr-x

MRv1: Changing the Directory Ownership on HDFS

- If Hadoop security is enabled, use `kinit hdfs` to obtain Kerberos credentials for the `hdfs` user by running the following commands before changing the directory ownership on HDFS:

```
$ sudo -u hdfs kinit -k -t hdfs.keytab hdfs/fully.qualified.domain.name@YOUR-REALM.COM
```

If `kinit hdfs` does not work initially, run `kinit -R` after running `kinit` to obtain credentials. (For more information, see Error Messages and Various Failures on page 412). To change the directory ownership on HDFS, run the following commands. Replace the example `/mapred/system` directory in the commands below with the HDFS directory specified by the `mapreduce.jobtracker.system.dir` (or `mapred.system.dir`) property in the `conf/mapred-site.xml` file:

```
$ sudo -u hdfs hadoop fs -chown mapred:hadoop /mapred/system
$ sudo -u hdfs hadoop fs -chown hdfs:hadoop /
$ sudo -u hdfs hadoop fs -chmod -R 700 /mapred/system
$ sudo -u hdfs hadoop fs -chmod 755 /
```

- In addition (whether or not Hadoop security is enabled) create the `/tmp` directory. For instructions on creating `/tmp` and setting its permissions, see these instructions.

Step 2b (YARN only): Verify User Accounts and Groups in YARN

> **Note:** If you are using MRv1, skip this step and proceed to Step 3: If you are Using AES-256 Encryption, Install the JCE Policy File on page 83.

During CDH 5 package installation of MapReduce 2.0 (YARN), the following Unix user accounts are automatically created to support security:

This User	Runs These Hadoop Programs
`hdfs`	HDFS: NameNode, DataNodes, Standby NameNode (if you are using HA)
`yarn`	YARN: ResourceManager, NodeManager
`mapred`	YARN: MapReduce JobHistory Server

> **Important:** The HDFS and YARN daemons must run as different Unix users; for example, `hdfs` and `yarn`. The MapReduce JobHistory server must run as user `mapred`. Having all of these users share a common Unix group is recommended; for example, hadoop.

[2] **When starting up, MapReduce sets the permissions for the `mapreduce.jobtracker.system.dir` (or `mapred.system.dir`) directory in HDFS, assuming the user `mapred` owns that directory.**

YARN: Directory Ownership in the Local Filesystem

Because the HDFS and MapReduce services run as different users, you must be sure to configure the correct directory ownership of the following files on the local filesystem of each host:

File System	Directory	Owner	Permissions (see Footnote 1)
Local	`dfs.namenode.name.dir` (`dfs.name.dir` is deprecated but will also work)	`hdfs:hdfs`	drwx------
Local	`dfs.datanode.data.dir` (`dfs.data.dir` is deprecated but will also work)	`hdfs:hdfs`	drwx------
Local	`yarn.nodemanager.local-dirs`	`yarn:yarn`	drwxr-xr-x
Local	`yarn.nodemanager.log-dirs`	`yarn:yarn`	drwxr-xr-x
Local	`container-executor`	`root:yarn`	--Sr-s---
Local	`conf/container-executor.cfg`	`root:yarn`	r--------

Important: Configuration changes to the Linux container executor could result in local NodeManager directories (such as `usercache`) being left with incorrect permissions. To avoid this, when making changes using either Cloudera Manager or the command line, first manually remove the existing NodeManager local directories from all configured local directories (`yarn.nodemanager.local-dirs`), and let the NodeManager recreate the directory structure.

You must also configure the following permissions for the HDFS, YARN and MapReduce log directories (the default locations in `/var/log/hadoop-hdfs`, `/var/log/hadoop-yarn` and `/var/log/hadoop-mapreduce`):

File System	Directory	Owner	Permissions
Local	`HDFS_LOG_DIR`	`hdfs:hdfs`	drwxrwxr-x
Local	`$YARN_LOG_DIR`	`yarn:yarn`	drwxrwxr-x
Local	`MAPRED_LOG_DIR`	`mapred:mapred`	drwxrwxr-x

YARN: Directory Ownership on HDFS

The following directories on HDFS must also be configured as follows:

File System	Directory	Owner	Permissions
HDFS	/ (root directory)	`hdfs:hadoop`	drwxr-xr-x
HDFS	`yarn.nodemanager.remote-app-log-dir`	`yarn:hadoop`	drwxrwxrwxt
HDFS	`mapreduce.jobhistory.intermediate-done-dir`	`mapred:hadoop`	drwxrwxrwxt

[3] **In CDH 5, package installation and the Hadoop daemons will automatically configure the correct permissions for you if you configure the directory ownership correctly as shown in the two tables above. See also Deploying MapReduce v2 (YARN) on a Cluster.**

File System	Directory	Owner	Permissions
HDFS	`mapreduce.jobhistory.done-dir`	`mapred:hadoop`	drwxr-x---

YARN: Changing the Directory Ownership on HDFS

If Hadoop security is enabled, use `kinit hdfs` to obtain Kerberos credentials for the `hdfs` user by running the following commands:

```
$ sudo -u hdfs kinit -k -t hdfs.keytab hdfs/fully.qualified.domain.name@YOUR-REALM.COM
$ hadoop fs -chown hdfs:hadoop /
$ hadoop fs -chmod 755 /
```

If `kinit hdfs` does not work initially, run `kinit -R` after running `kinit` to obtain credentials. See Error Messages and Various Failures on page 412. To change the directory ownership on HDFS, run the following commands:

```
$ sudo -u hdfs hadoop fs -chown hdfs:hadoop /
$ sudo -u hdfs hadoop fs -chmod 755 /
$ sudo -u hdfs hadoop fs -chown yarn:hadoop [yarn.nodemanager.remote-app-log-dir]
$ sudo -u hdfs hadoop fs -chmod 1777 [yarn.nodemanager.remote-app-log-dir]
$ sudo -u hdfs hadoop fs -chown mapred:hadoop [mapreduce.jobhistory.intermediate-done-dir]
$ sudo -u hdfs hadoop fs -chmod 1777 [mapreduce.jobhistory.intermediate-done-dir]
$ sudo -u hdfs hadoop fs -chown mapred:hadoop [mapreduce.jobhistory.done-dir]
$ sudo -u hdfs hadoop fs -chmod 750 [mapreduce.jobhistory.done-dir]
```

- In addition (whether or not Hadoop security is enabled) create the `/tmp` directory. For instructions on creating `/tmp` and setting its permissions, see Step 7: If Necessary, Create the HDFS /tmp Directory.
- In addition (whether or not Hadoop security is enabled), change permissions on the `/user/history` Directory. See Step 8: Create the history Directory and Set Permissions.

Step 3: If you are Using AES-256 Encryption, Install the JCE Policy File

> **Note:** This step is not required when using JDK 1.8.0_161 or greater. JDK 1.8.0_161 enables unlimited strength encryption by default.

If you are using CentOS/Red Hat Enterprise Linux 5.6 or higher, or Ubuntu, which use AES-256 encryption by default for tickets, you must install the Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction Policy File on all cluster and Hadoop user machines. For JCE Policy File installation instructions, see the `README.txt` file included in the `jce_policy-x.zip` file.

Alternatively, you can configure Kerberos to not use AES-256 by removing `aes256-cts:normal` from the `supported_enctypes` field of the `kdc.conf` or `krb5.conf` file. After changing the `kdc.conf` file, you must restart both the KDC and the kadmin server for those changes to take affect. You may also need to re-create or change the password of the relevant principals, including potentially the Ticket Granting Ticket principal (krbtgt/REALM@REALM). If AES-256 is still used after completing steps, the `aes256-cts:normal` setting existed when the Kerberos database was created. To fix this, create a new Kerberos database and then restart both the KDC and the kadmin server.

To verify the type of encryption used in your cluster:

1. On the local KDC host, type this command to create a test principal:

```
$ kadmin -q "addprinc test"
```

2. On a cluster host, type this command to start a Kerberos session as test:

```
$ kinit test
```

3. On a cluster host, type this command to view the encryption type in use:

```
$ klist -e
```

If AES is being used, output like the following is displayed after you type the `klist` command; note that `AES-256` is included in the output:

```
Ticket cache: FILE:/tmp/krb5cc_0
Default principal: test@SCM
Valid starting     Expires            Service principal
05/19/11 13:25:04  05/20/11 13:25:04  krbtgt/SCM@SCM
    Etype (skey, tkt): AES-256 CTS mode with 96-bit SHA-1 HMAC, AES-256 CTS mode with
96-bit SHA-1 HMAC
```

Step 4: Create and Deploy the Kerberos Principals and Keytab Files

A Kerberos principal is used in a Kerberos-secured system to represent a unique identity. Kerberos assigns tickets to Kerberos principals to enable them to access Kerberos-secured Hadoop services. For Hadoop, the principals should be of the format `username/fully.qualified.domain.name@YOUR-REALM.COM`. In this guide, the term `username` in the `username/fully.qualified.domain.name@YOUR-REALM.COM` principal refers to the username of an existing Unix account, such as `hdfs` or `mapred`.

A keytab is a file containing pairs of Kerberos principals and an encrypted copy of that principal's key. The keytab files are unique to each host since their keys include the hostname. This file is used to authenticate a principal on a host to Kerberos without human interaction or storing a password in a plain text file. Because having access to the keytab file for a principal allows one to act as that principal, access to the keytab files should be tightly secured. They should be readable by a minimal set of users, should be stored on local disk, and should not be included in machine backups, unless access to those backups is as secure as access to the local machine.

Important:

For both MRv1 and YARN deployments: *On every machine in your cluster*, there must be a keytab file for the `hdfs` user and a keytab file for the `mapred` user. The `hdfs` keytab file must contain entries for the `hdfs` principal and a `HTTP` principal, and the `mapred` keytab file must contain entries for the `mapred` principal and a `HTTP` principal. On each respective machine, the `HTTP` principal will be the same in both keytab files.

In addition, for YARN deployments only: *On every machine in your cluster*, there must be a keytab file for the `yarn` user. The `yarn` keytab file must contain entries for the `yarn` principal and a `HTTP` principal. On each respective machine, the `HTTP` principal in the `yarn` keytab file will be the same as the `HTTP` principal in the `hdfs` and `mapred` keytab files.

Note:

The following instructions illustrate an example of creating keytab files for MIT Kerberos. If you are using another version of Kerberos, refer to your Kerberos documentation for instructions. You may use either `kadmin` or `kadmin.local` to run these commands.

When to Use kadmin.local and kadmin

When creating the Kerberos principals and keytabs, you can use kadmin.local or kadmin depending on your access and account:

- If you have root access to the KDC machine, but you do not have a Kerberos admin account, use kadmin.local.
- If you do not have root access to the KDC machine, but you do have a Kerberos admin account, use kadmin.
- If you have both root access to the KDC machine and a Kerberos admin account, you can use either one.

To start kadmin.local (on the KDC machine) or kadmin from any machine, run this command:

```
$ sudo kadmin.local
```

OR:

```
$ kadmin
```

> **Note:**
>
> In this guide, `kadmin` is shown as the prompt for commands in the kadmin shell, but you can type the same commands at the `kadmin.local` prompt in the kadmin.local shell.

> **Note:**
>
> Running kadmin.local may prompt you for a password because it is being run via sudo. You should provide your Unix password. Running kadmin may prompt you for a password because you need Kerberos admin privileges. You should provide your Kerberos admin password.

To create the Kerberos principals

> **Important:**
>
> If you plan to use Oozie, Impala, or the Hue Kerberos ticket renewer in your cluster, you must configure your KDC to allow tickets to be renewed, and you must configure `krb5.conf` to request renewable tickets. Typically, you can do this by adding the `max_renewable_life` setting to your realm in `kdc.conf`, and by adding the `renew_lifetime` parameter to the `libdefaults` section of `krb5.conf`. For more information about renewable tickets, see the Kerberos documentation.

Do the following steps for every host in your cluster. Run the commands in the kadmin.local or kadmin shell, replacing the `fully.qualified.domain.name` in the commands with the fully qualified domain name of each host. Replace `YOUR-REALM.COM` with the name of the Kerberos realm your Hadoop cluster is in.

1. In the kadmin.local or kadmin shell, create the `hdfs` principal. This principal is used for the NameNode, Secondary NameNode, and DataNodes.

```
kadmin:  addprinc -randkey hdfs/fully.qualified.domain.name@YOUR-REALM.COM
```

> **Note:**
>
> If your Kerberos administrator or company has a policy about principal names that does not allow you to use the format shown above, you can work around that issue by configuring the `<kerberos principal>` to `<short name>` mapping that is built into Hadoop. For more information, see Configuring the Mapping from Kerberos Principals to Short Names.

2. Create the `mapred` principal. If you are using MRv1, the `mapred` principal is used for the JobTracker and TaskTrackers. If you are using YARN, the `mapred` principal is used for the MapReduce Job History Server.

```
kadmin:  addprinc -randkey mapred/fully.qualified.domain.name@YOUR-REALM.COM
```

3. **YARN only:** Create the `yarn` principal. This principal is used for the ResourceManager and NodeManager.

```
kadmin:  addprinc -randkey yarn/fully.qualified.domain.name@YOUR-REALM.COM
```

4. Create the `HTTP` principal.

```
kadmin:  addprinc -randkey HTTP/fully.qualified.domain.name@YOUR-REALM.COM
```

> **Important:**
>
> The HTTP principal must be in the format `HTTP/fully.qualified.domain.name@YOUR-REALM.COM`. The first component of the principal must be the literal string "HTTP". This format is standard for HTTP principals in SPNEGO and is hard-coded in Hadoop. It cannot be deviated from.

To create the Kerberos keytab files

> **Important:**
>
> The instructions in this section for creating keytab files require using the Kerberos `norandkey` option in the `xst` command. If your version of Kerberos does not support the `norandkey` option, or if you cannot use `kadmin.local`, then use these alternate instructions to create appropriate Kerberos keytab files. After using those alternate instructions to create the keytab files, continue with the next section To deploy the Kerberos keytab files.

Do the following steps for every host in your cluster. Run the commands in the kadmin.local or kadmin shell, replacing the `fully.qualified.domain.name` in the commands with the fully qualified domain name of each host:

1. Create the `hdfs` keytab file that will contain the `hdfs` principal and `HTTP` principal. This keytab file is used for the NameNode, Secondary NameNode, and DataNodes.

```
kadmin:  xst -norandkey -k hdfs.keytab hdfs/fully.qualified.domain.name
HTTP/fully.qualified.domain.name
```

2. Create the `mapred` keytab file that will contain the `mapred` principal and `HTTP` principal. If you are using MRv1, the `mapred` keytab file is used for the JobTracker and TaskTrackers. If you are using YARN, the `mapred` keytab file is used for the MapReduce Job History Server.

```
kadmin:  xst -norandkey -k mapred.keytab mapred/fully.qualified.domain.name
HTTP/fully.qualified.domain.name
```

3. **YARN only:** Create the `yarn` keytab file that will contain the `yarn` principal and `HTTP` principal. This keytab file is used for the ResourceManager and NodeManager.

```
kadmin:  xst -norandkey -k yarn.keytab yarn/fully.qualified.domain.name
HTTP/fully.qualified.domain.name
```

4. Use `klist` to display the keytab file entries; a correctly-created hdfs keytab file should look something like this:

```
$ klist -e -k -t hdfs.keytab
Keytab name: WRFILE:hdfs.keytab
slot KVNO Principal
---- ---- ---------------------------------------------------------------------
   1    7    HTTP/fully.qualified.domain.name@YOUR-REALM.COM (DES cbc mode with CRC-32)

   2    7    HTTP/fully.qualified.domain.name@YOUR-REALM.COM (Triple DES cbc mode with
 HMAC/sha1)
   3    7    hdfs/fully.qualified.domain.name@YOUR-REALM.COM (DES cbc mode with CRC-32)

   4    7    hdfs/fully.qualified.domain.name@YOUR-REALM.COM (Triple DES cbc mode with
 HMAC/sha1)
```

5. Continue with the next section To deploy the Kerberos keytab files.

To deploy the Kerberos keytab files

On every node in the cluster, repeat the following steps to deploy the `hdfs.keytab` and `mapred.keytab` files. If you are using YARN, you will also deploy the `yarn.keytab` file.

1. On the host machine, copy or move the keytab files to a directory that Hadoop can access, such as `/etc/hadoop/conf`.

   a. **If you are using MRv1:**

```
$ sudo mv hdfs.keytab mapred.keytab /etc/hadoop/conf/
```

   If you are using YARN:

```
$ sudo mv hdfs.keytab mapred.keytab yarn.keytab /etc/hadoop/conf/
```

   b. Make sure that the `hdfs.keytab` file is only readable by the `hdfs` user, and that the `mapred.keytab` file is only readable by the `mapred` user.

```
$ sudo chown hdfs:hadoop /etc/hadoop/conf/hdfs.keytab
$ sudo chown mapred:hadoop /etc/hadoop/conf/mapred.keytab
$ sudo chmod 400 /etc/hadoop/conf/*.keytab
```

> **Note:**
>
> To enable you to use the same configuration files on every host, Cloudera recommends that you use the same name for the keytab files on every host.

   c. **YARN only:** Make sure that the `yarn.keytab` file is only readable by the `yarn` user.

```
$ sudo chown yarn:hadoop /etc/hadoop/conf/yarn.keytab
$ sudo chmod 400 /etc/hadoop/conf/yarn.keytab
```

> **Important:**
>
> If the NameNode, Secondary NameNode, DataNode, JobTracker, TaskTrackers, HttpFS, or Oozie services are configured to use Kerberos HTTP SPNEGO authentication, and two or more of these services are running on the same host, then all of the running services must use the same HTTP principal and keytab file used for their HTTP endpoints.

Step 5: Shut Down the Cluster

To enable security in CDH, you must stop all Hadoop daemons in your cluster and then change some configuration properties. You must stop all daemons in the cluster because after one Hadoop daemon has been restarted with the configuration properties set to enable security, daemons running without security enabled will be unable to communicate with that daemon. This requirement to shut down all daemons makes it impossible to do a rolling upgrade to enable security on a Hadoop cluster.

To shut down the cluster, run the following command on every node in your cluster (as root):

```
$ for x in `cd /etc/init.d ; ls hadoop-*` ; do sudo service $x stop ; done
```

Step 6: Enable Hadoop Security

Cloudera recommends that all of the Hadoop configuration files throughout the cluster have the same contents.

To enable Hadoop security, add the following properties to the `core-site.xml` file *on every machine* in the cluster:

```
<property>
  <name>hadoop.security.authentication</name>
  <value>kerberos</value> <!-- A value of "simple" would disable security. -->
</property>
```

```
<property>
  <name>hadoop.security.authorization</name>
  <value>true</value>
</property>
```

Enabling Service-Level Authorization for Hadoop Services

Service-level authorizations prevent users from accessing a cluster at the course-grained level. For example, when Authorized Users and Authorized Groups are setup properly, an unauthorized user cannot use the hdfs shell to list the contents of HDFS. This also limits the exposure of world-readable files to an explicit set of users instead of all authenticated users, which could be, for example, every user in Active Directory.

The `hadoop-policy.xml` file maintains access control lists (ACL) for Hadoop services. Each ACL consists of comma-separated lists of users and groups separated by a space. For example:

```
user_a,user_b group_a,group_b
```

If you only want to specify a set of users, add a comma-separated list of users followed by a blank space. Similarly, to specify only authorized groups, use a blank space at the beginning. A `*` can be used to give access to all users.

For example, to give users, ann, bob, and groups, `group_a`, `group_b` access to Hadoop's DataNodeProtocol service, modify the `security.datanode.protocol.acl` property in `hadoop-policy.xml`. Similarly, to give all users access to the InterTrackerProtocol service, modify `security.inter.tracker.protocol.acl` as follows:

```
<property>
    <name>security.datanode.protocol.acl</name>
    <value>ann,bob group_a,group_b</value>
    <description>ACL for DatanodeProtocol, which is used by datanodes to
    communicate with the namenode.</description>
</property>

<property>
    <name>security.inter.tracker.protocol.acl</name>
    <value>*</value>
    <description>ACL for InterTrackerProtocol, which is used by tasktrackers to
    communicate with the jobtracker.</description>
</property>
```

For more details, see Service-Level Authorization in Hadoop.

Step 7: Configure Secure HDFS

When following the instructions in this section to configure the properties in the `hdfs-site.xml` file, keep the following important guidelines in mind:

- The properties for each daemon (NameNode, Secondary NameNode, and DataNode) must specify both the HDFS and HTTP principals, as well as the path to the HDFS keytab file.
- The Kerberos principals for the NameNode, Secondary NameNode, and DataNode are configured in the `hdfs-site.xml` file. The same `hdfs-site.xml` file with *all three* of these principals must be installed on every host machine in the cluster. That is, it is not sufficient to have the NameNode principal configured on the NameNode host machine only. This is because, for example, the DataNode must know the principal name of the NameNode in order to send heartbeats to it. Kerberos authentication is bi-directional.
- The special string `_HOST` in the properties is replaced at run-time by the fully qualified domain name of the host machine where the daemon is running. This requires that reverse DNS is properly working on all the hosts configured this way. You may use `_HOST` only as the entirety of the second component of a principal name. For example, hdfs/_HOST@YOUR-REALM.COM is valid, but hdfs._HOST@YOUR-REALM.COM and hdfs/_HOST.example.com@YOUR-REALM.COM are not.
- When performing the `_HOST` substitution for the Kerberos principal names, the NameNode determines its own hostname based on the configured value of `fs.default.name`, whereas the DataNodes determine their hostnames based on the result of reverse DNS resolution on the DataNode hosts. Likewise, the JobTracker uses the configured

value of `mapred.job.tracker` to determine its hostname whereas the TaskTrackers, like the DataNodes, use reverse DNS.

- The `dfs.datanode.address` and `dfs.datanode.http.address` port numbers for the DataNode *must* be below 1024, because this provides part of the security mechanism to make it impossible for a user to run a map task which impersonates a DataNode. The port numbers for the NameNode and Secondary NameNode can be anything you want, but the default port numbers are good ones to use.

To configure secure HDFS

Add the following properties to the `hdfs-site.xml` file *on every machine* in the cluster. Replace these example values shown below with the correct settings for your site: *path to the HDFS keytab*, *YOUR-REALM.COM*, *fully qualified domain name of NN*, and *fully qualified domain name of 2NN*

```
<!-- General HDFS security config -->
<property>
  <name>dfs.block.access.token.enable</name>
  <value>true</value>
</property>

<!-- NameNode security config -->
<property>
  <name>dfs.namenode.keytab.file</name>
  <value>/etc/hadoop/conf/hdfs.keytab</value> <!-- path to the HDFS keytab -->
</property>
<property>
  <name>dfs.namenode.kerberos.principal</name>
  <value>hdfs/_HOST@YOUR-REALM.COM</value>
</property>
<property>
  <name>dfs.namenode.kerberos.internal.spnego.principal</name>
  <value>HTTP/_HOST@YOUR-REALM.COM</value>
</property>

<!-- Secondary NameNode security config -->
<property>
  <name>dfs.secondary.namenode.keytab.file</name>
  <value>/etc/hadoop/conf/hdfs.keytab</value> <!-- path to the HDFS keytab -->
</property>
<property>
  <name>dfs.secondary.namenode.kerberos.principal</name>
  <value>hdfs/_HOST@YOUR-REALM.COM</value>
</property>
<property>
  <name>dfs.secondary.namenode.kerberos.internal.spnego.principal</name>
  <value>HTTP/_HOST@YOUR-REALM.COM</value>
</property>

<!-- DataNode security config -->
<property>
  <name>dfs.datanode.data.dir.perm</name>
  <value>700</value>
</property>
<property>
  <name>dfs.datanode.address</name>
  <value>0.0.0.0:1004</value>
</property>
<property>
  <name>dfs.datanode.http.address</name>
  <value>0.0.0.0:1006</value>
</property>
<property>
  <name>dfs.datanode.keytab.file</name>
  <value>/etc/hadoop/conf/hdfs.keytab</value> <!-- path to the HDFS keytab -->
</property>
<property>
  <name>dfs.datanode.kerberos.principal</name>
  <value>hdfs/_HOST@YOUR-REALM.COM</value>
</property>

<!-- Web Authentication config -->
<property>
```

```
  <name>dfs.web.authentication.kerberos.principal</name>
  <value>HTTP/_HOST@YOUR_REALM</value>
 </property>
```

To enable TLS/SSL for HDFS

Add the following property to `hdfs-site.xml` on *every machine* in your cluster.

```
<property>
<name>dfs.http.policy</name>
<value>HTTPS_ONLY</value>
</property>
```

Optional Step 8: Configuring Security for HDFS High Availability

CDH 5 supports the HDFS High Availability (HA) feature with Kerberos security enabled. There are two use cases that affect security for HA:

- If you are not using Quorum-based Storage (see Software Configuration for Quorum-based Storage), then no extra configuration for HA is necessary if automatic failover is not enabled. If automatic failover is enabled then access to ZooKeeper should be secured. See the Software Configuration for Shared Storage Using NFS documentation for details.
- If you are using Quorum-based Storage, then you must configure security for Quorum-based Storage by following the instructions in this section.

To configure security for Quorum-based Storage:

Add the following Quorum-based Storage configuration properties to the `hdfs-site.xml` file on all of the machines in the cluster:

```
<property>
  <name>dfs.journalnode.keytab.file</name>
  <value>/etc/hadoop/conf/hdfs.keytab</value> <!-- path to the HDFS keytab -->
</property>
<property>
  <name>dfs.journalnode.kerberos.principal</name>
  <value>hdfs/_HOST@YOUR-REALM.COM</value>
</property>
<property>
  <name>dfs.journalnode.kerberos.internal.spnego.principal</name>
  <value>HTTP/_HOST@YOUR-REALM.COM</value>
</property>
```

Note:

If you already have principals and keytabs created for the machines where the JournalNodes are running, then you should reuse those principals and keytabs in the configuration properties above. You will likely have these principals and keytabs already created if you are collocating a JournalNode on a machine with another HDFS daemon.

Optional Step 9: Configure secure WebHDFS

Note:

If you are not using WebHDFS, you can skip this step.

Security for WebHDFS is disabled by default. If you want use WebHDFS with a secure cluster, this is the time to enable and configure it.

To configure secure WebHDFS:

1. If you have not already done so, enable WebHDFS by adding the following property to the `hdfs-site.xml` file *on every machine* in the cluster.

```
<property>
  <name>dfs.webhdfs.enabled</name>
  <value>true</value>
</property>
```

2. Add the following properties to the `hdfs-site.xml` file *on every machine* in the cluster. Replace the example values shown below with the correct settings for your site.

```
<property>
  <name>dfs.web.authentication.kerberos.principal</name>
  <value>HTTP/_HOST@YOUR-REALM.COM</value>
</property>

<property>
  <name>dfs.web.authentication.kerberos.keytab</name>
  <value>/etc/hadoop/conf/HTTP.keytab</value> <!-- path to the HTTP keytab -->
</property>
```

Optional Step 10: Configuring a secure HDFS NFS Gateway

To deploy a Kerberized HDFS NFS gateway, add the following configuration properties to `hdfs-site.xml` on the NFS server.

```
<property>
<name>dfs.nfs.keytab.file</name>
<value>/etc/hadoop/conf/hdfs.keytab</value> <!-- path to the HDFS or NFS gateway keytab
 -->
</property>

<property>
<name>dfs.nfs.kerberos.principal</name>
<value>hdfs/_HOST@YOUR-REALM.COM</value>
</property>
```

Potential Insecurities with a Kerberized NFS Gateway

When configuring an NFS gateway in a secure cluster, the gateway accesses the contents of HDFS using the HDFS service principals. However, authorization for end users is handled by comparing the end user's UID/GID against the UID/GID of the files on the NFS mount. No Kerberos is involved in authenticating the user first.

Because HDFS metadata doesn't have any UIDs/GIDs, only names and groups, the NFS gateway maps user names and group names to UIDs and GIDs. The user names and group names used for this mapping are derived from the local users of the host where the NFS gateway is running. The mapped IDs are then presented to the NFS client for authorization. The NFS client performs the authorization locally, comparing the UID/GID presented by the NFS Gateway to the IDs of the users on the remote host.

The main risk with this procedure is that it's quite possible to create local users with UIDs that were previously associated with any superusers. For example, users with access to HDFS can view the directories that belong to the `hdfs` user, and they can also access the underlying metadata to obtain the associated UID. Assuming the directories owned by `hdfs` have their UID set to `xyz`, a malicious user could create a new local user on the NFS gateway host with the UID set to `xyz`. This local user will now be able to freely access the `hdfs` user's files.

Solutions:

- Set the NFS Gateway property, **Allowed Hosts and Privileges**, to allow only those NFS clients that are trusted and managed by the Hadoop administrators.
  1. Go to the Cloudera Manager Admin Console and navigate to the HDFS service.
  2. Click the **Configuration** tab.
  3. Select **Scope** > **NFS Gateway**.

4. Select **Category** > **Main**.
5. Locate the **Allowed Hosts and Privileges** property and set it to a list of trusted host names and access privileges (`ro` - read-only, `rw` - read/write). For example:

```
192.168.0.0/22 rw
host1.example.org ro
```

   The current default setting of this property is `* rw`, which is a security risk because it lets everybody map the NFS export in read-write mode.
6. Click **Save Changes** to commit the changes.

- Specify a user with restricted privileges for the `dfs.nfs.kerberos.principal` property, so that the NFS gateway has limited access to the NFS contents. The current default setting for this property is `hdfs/_HOST@YOUR-REALM.COM</value>`, which gives the NFS gateway unrestricted access to HDFS.

Step 11: Set Variables for Secure DataNodes

In order to allow DataNodes to start on a secure Hadoop cluster, you must set the following variables on all DataNodes in `/etc/default/hadoop-hdfs-datanode`.

```
export HADOOP_SECURE_DN_USER=hdfs
export HADOOP_SECURE_DN_PID_DIR=/var/lib/hadoop-hdfs
export HADOOP_SECURE_DN_LOG_DIR=/var/log/hadoop-hdfs
export JSVC_HOME=/usr/lib/bigtop-utils/
```

Note:

Depending on the version of Linux you are using, you may not have the `/usr/lib/bigtop-utils` directory on your system. If that is the case, set the `JSVC_HOME` variable to the `/usr/libexec/bigtop-utils` directory by using this command:

```
export JSVC_HOME=/usr/libexec/bigtop-utils
```

Step 12: Start up the NameNode

The NameNode can now be started. Use the `service` command to run the `/etc/init.d` script.

```
$ sudo service hadoop-hdfs-namenode start
```

As start-up proceeds, various messages get written to the logs, such as the following:

```
10/10/25 17:01:46 INFO security.UserGroupInformation:
Login successful for user hdfs/fully.qualified.domain.name@YOUR-REALM.COM using keytab
 file /etc/hadoop/conf/hdfs.keytab
```

and also:

```
12/05/23 18:18:31 INFO http.HttpServer: Adding Kerberos (SPNEGO) filter to
getDelegationToken
12/05/23 18:18:31 INFO http.HttpServer: Adding Kerberos (SPNEGO) filter to
renewDelegationToken
12/05/23 18:18:31 INFO http.HttpServer: Adding Kerberos (SPNEGO) filter to
cancelDelegationToken
12/05/23 18:18:31 INFO http.HttpServer: Adding Kerberos (SPNEGO) filter to fsck
12/05/23 18:18:31 INFO http.HttpServer: Adding Kerberos (SPNEGO) filter to getimage
12/05/23 18:18:31 INFO http.HttpServer: Jetty bound to port 50070
12/05/23 18:18:31 INFO mortbay.log: jetty-6.1.26
12/05/23 18:18:31 INFO server.KerberosAuthenticationHandler: Login using keytab
/etc/hadoop/conf/hdfs.keytab, for principal
HTTP/fully.qualified.domain.name@YOUR-REALM.COM
12/05/23 18:18:31 INFO server.KerberosAuthenticationHandler: Initialized, principal
```

```
[HTTP/fully.qualified.domain.name@YOUR-REALM.COM] from keytab
[/etc/hadoop/conf/hdfs.keytab]
```

Verify that the NameNode has started by opening a web browser to the host and port, as follows:

```
http://fqdn.hostname.example.com:50070/
```

Cloudera recommends testing that the NameNode is working properly by performing a metadata-only HDFS operation, which will now require correct Kerberos credentials. For example:

```
$ hadoop fs -ls
```

Information about the kinit Command

Important:

Running the `hadoop fs -ls` command will fail if you do not have a valid Kerberos ticket in your credentials cache. You can examine the Kerberos tickets currently in your credentials cache by running the `klist` command. You can obtain a ticket by running the `kinit` command and either specifying a keytab file containing credentials, or entering the password for your principal. If you do not have a valid ticket, you will receive an error such as:

```
11/01/04 12:08:12 WARN ipc.Client: Exception encountered while connecting
 to the server : javax.security.sasl.SaslException:
 GSS initiate failed [Caused by GSSException: No valid credentials
provided (Mechanism level: Failed to find any Kerberos tgt)]
Bad connection to FS. command aborted. exception: Call to
nn-host/10.0.0.2:8020 failed on local exception: java.io.IOException:
javax.security.sasl.SaslException: GSS initiate failed [Caused by
GSSException: No valid credentials provided (Mechanism level: Failed to
 find any Kerberos tgt)]
```

Note:

The `kinit` command must either be on the path for user accounts running the Hadoop client, or else the `hadoop.kerberos.kinit.command` parameter in `core-site.xml` must be manually configured to the absolute path to the `kinit` command.

Note:

For MIT Kerberos 1.8.1 or higher, a bug in versions of the Oracle JDK 6 Update 26 and higher causes Java to be unable to read the Kerberos credentials cache even after you have successfully obtained a Kerberos ticket using `kinit`. To workaround this bug, run `kinit -R` after running `kinit` initially to obtain credentials. Doing so will cause the ticket to be renewed, and the credentials cache rewritten in a format which Java can read. See Troubleshooting Java and Kerberos issues for more information.

Step 12: Start up a DataNode

Begin by starting one DataNode only to make sure it can properly connect to the NameNode. Use the `service` command to run the `/etc/init.d` script.

```
$ sudo service hadoop-hdfs-datanode start
```

You'll see some extra information in the logs such as:

```
10/10/25 17:21:41 INFO security.UserGroupInformation:
Login successful for user hdfs/fully.qualified.domain.name@YOUR-REALM.COM using keytab
 file /etc/hadoop/conf/hdfs.keytab
```

If you can get a single DataNode running and you can see it registering with the NameNode in the logs, then start up all the DataNodes. You should now be able to do all HDFS operations.

Step 14: Set the Sticky Bit on HDFS Directories

This step is optional but strongly recommended for security. In CDH 5, HDFS file permissions have support for the sticky bit. The sticky bit can be set on directories, preventing anyone except the superuser, directory owner, or file owner from deleting or moving the files within the directory. Setting the sticky bit for a file has no effect. This is useful for directories such as `/tmp` which previously had to be set to be world-writable. To set the sticky bit on the `/tmp` directory, run the following command:

```
$ sudo -u hdfs kinit -k -t hdfs.keytab hdfs/fully.qualified.domain.name@YOUR-REALM.COM
$ sudo -u hdfs hadoop fs -chmod 1777 /tmp
```

After running this command, the permissions on `/tmp` will appear as shown below. (Note the "`t`" instead of the final "`x`".)

```
$ hadoop fs -ls /
Found 2 items
drwxrwxrwt - hdfs supergroup 0 2011-02-14 15:55 /tmp
drwxr-xr-x - hdfs supergroup 0 2011-02-14 14:01 /user
```

Step 15: Start up the Secondary NameNode (if used)

At this point, you should be able to start the Secondary NameNode if you are using one:

```
$ sudo service hadoop-hdfs-secondarynamenode start
```

> **Note:**
>
> If you are using HDFS HA, do not use the Secondary NameNode. See Configuring HDFS High Availability for instructions on configuring and deploying the Standby NameNode.

You'll see some extra information in the logs such as:

```
10/10/26 12:03:18 INFO security.UserGroupInformation:
Login successful for user hdfs/fully.qualified.domain.name@YOUR-REALM using keytab file
 /etc/hadoop/conf/hdfs.keytab
```

and:

```
12/05/23 18:33:06 INFO http.HttpServer: Adding Kerberos (SPNEGO) filter to getimage
12/05/23 18:33:06 INFO http.HttpServer: Jetty bound to port 50090
12/05/23 18:33:06 INFO mortbay.log: jetty-6.1.26
12/05/23 18:33:06 INFO server.KerberosAuthenticationHandler: Login using keytab
/etc/hadoop/conf/hdfs.keytab, for principal
HTTP/fully.qualified.domain.name@YOUR-REALM.COM
12/05/23 18:33:06 INFO server.KerberosAuthenticationHandler: Initialized, principal
[HTTP/fully.qualified.domain.name@YOUR-REALM.COM] from keytab
[/etc/hadoop/conf/hdfs.keytab]
```

You should make sure that the Secondary NameNode not only starts, but that it is successfully checkpointing.

If you're using the `service` command to start the Secondary NameNode from the `/etc/init.d scripts`, Cloudera recommends setting the property `fs.checkpoint.period` in the `hdfs-site.xml` file to a very low value (such as 5), and then monitoring the Secondary NameNode logs for a successful startup and checkpoint. Once you are satisfied that the Secondary NameNode is checkpointing properly, you should reset the `fs.checkpoint.period` to a reasonable value, or return it to the default, and then restart the Secondary NameNode.

You can make the Secondary NameNode perform a checkpoint by doing the following:

```
$ sudo -u hdfs hdfs secondarynamenode -checkpoint force
```

Note that this will not cause a running Secondary NameNode to checkpoint, but rather will start up a Secondary NameNode that will immediately perform a checkpoint and then shut down. This can be useful for debugging.

> **Note:**
>
> If you encounter errors during Secondary NameNode checkpointing, it may be helpful to enable Kerberos debugging output. See Enabling Debugging for details.

Step 16: Configure Either MRv1 Security or YARN Security

At this point, you are ready to configure either MRv1 Security or YARN Security.

- If you are using MRv1, do the steps in Configuring MRv1 Security to configure, start, and test secure MRv1.
- If you are using YARN, do the steps in Configuring YARN Security to configure, start, and test secure YARN.

Configuring MRv1 Security

If you are using YARN, skip this section and see Configuring YARN Security.

If you are using MRv1, do the following steps to configure, start, and test secure MRv1.

1. Step 1: Configure Secure MRv1 on page 95
2. Step 2: Start up the JobTracker on page 97
3. Step 3: Start up a TaskTracker on page 97
4. Step 4: Try Running a Map/Reduce Job on page 97

Step 1: Configure Secure MRv1

Keep the following important information in mind when configuring secure MapReduce:

- The properties for JobTracker and TaskTracker must specify the mapred principal, as well as the path to the `mapred` keytab file.
- The Kerberos principals for the JobTracker and TaskTracker are configured in the `mapred-site.xml` file. The same `mapred-site.xml` file with *both* of these principals must be installed on every host machine in the cluster. That is, it is not sufficient to have the JobTracker principal configured on the JobTracker host machine only. This is because, for example, the TaskTracker must know the principal name of the JobTracker to securely register with the JobTracker. Kerberos authentication is bi-directional.
- Do not use `${user.name}` in the value of the `mapred.local.dir` or `hadoop.log.dir` properties in `mapred-site.xml`. Doing so can prevent tasks from launching on a secure cluster.
- Make sure that each user who will be running MRv1 jobs exists on all cluster hosts (that is, on every host that hosts any MRv1 daemon).
- Make sure the value specified for `mapred.local.dir` is identical in `mapred-site.xml` and `taskcontroller.cfg`. If the values are different, this error message is returned.
- Make sure the value specified in `taskcontroller.cfg` for `hadoop.log.dir` is the same as what the Hadoop daemons are using, which is `/var/log/hadoop-0.20-mapreduce` by default and can be configured in `mapred-site.xml`. If the values are different, this error message is returned.

To configure secure MapReduce:

1. Add the following properties to the `mapred-site.xml` file *on every machine* in the cluster:

```
<!-- JobTracker security configs -->
<property>
  <name>mapreduce.jobtracker.kerberos.principal</name>
  <value>mapred/_HOST@YOUR-REALM.COM</value>
</property>
<property>
  <name>mapreduce.jobtracker.keytab.file</name>
  <value>/etc/hadoop/conf/mapred.keytab</value> <!-- path to the MapReduce keytab -->
</property>

<!-- TaskTracker security configs -->
<property>
  <name>mapreduce.tasktracker.kerberos.principal</name>
  <value>mapred/_HOST@YOUR-REALM.COM</value>
</property>
<property>
  <name>mapreduce.tasktracker.keytab.file</name>
  <value>/etc/hadoop/conf/mapred.keytab</value> <!-- path to the MapReduce keytab -->
</property>

<!-- TaskController settings -->
<property>
  <name>mapred.task.tracker.task-controller</name>
  <value>org.apache.hadoop.mapred.LinuxTaskController</value>
</property>
<property>
  <name>mapreduce.tasktracker.group</name>
  <value>mapred</value>
</property>
```

2. Create a file called `taskcontroller.cfg` that contains the following information:

```
hadoop.log.dir=<Path to Hadoop log directory. Should be same value used to start the
TaskTracker. This is required to set proper permissions on the log files so that they
can be written to by the user's tasks and read by the TaskTracker for serving on the
web UI.>
mapreduce.tasktracker.group=mapred
banned.users=mapred,hdfs,bin
min.user.id=1000
```

> **Note:**
>
> The default setting for the `banned.users` property in the `taskcontroller.cfg` file is `mapred`, `hdfs`, and `bin` to prevent jobs from being submitted using those user accounts. The default setting for the `min.user.id` property is `1000` to prevent jobs from being submitted with a user ID less than 1000, which are conventionally Unix super users. Some operating systems such as CentOS 5 use a default value of 500 and above for user IDs, not 1000. If this is the case on your system, change the default setting for the `min.user.id` property to `500`. If there are user accounts on your cluster that have a user ID less than the value specified for the `min.user.id` property, the TaskTracker returns an error code of 255.

3. The path to the `taskcontroller.cfg` file is determined relative to the location of the `task-controller` binary. Specifically, the path is `<path of task-controller binary>/../../conf/taskcontroller.cfg`. If you installed the CDH 5 package, this path will always correspond to `/etc/hadoop/conf/taskcontroller.cfg`.

> **Note:**
>
> For more information about the `task-controller` program, see Information about Other Hadoop Security Programs.

> **Important:**
>
> The same `mapred-site.xml` file and the same `hdfs-site.xml` file must both be installed on every host machine in the cluster so that the NameNode, Secondary NameNode, DataNode, JobTracker and TaskTracker can all connect securely with each other.

Step 2: Start up the JobTracker

You are now ready to start the JobTracker.

If you're using the `/etc/init.d/hadoop-0.20-mapreduce-jobtracker` script, then you can use the `service` command to run it now:

```
$ sudo service hadoop-0.20-mapreduce-jobtracker start
```

You can verify that the JobTracker is working properly by opening a web browser to `http://machine:50030/` where *machine* is the name of the machine where the JobTracker is running.

Step 3: Start up a TaskTracker

You are now ready to start a TaskTracker.

If you're using the `/etc/init.d/hadoop-0.20-mapreduce-tasktracker` script, then you can use the `service` command to run it now:

```
$ sudo service hadoop-0.20-mapreduce-tasktracker start
```

Step 4: Try Running a Map/Reduce Job

You should now be able to run Map/Reduce jobs. To confirm, try launching a sleep or a pi job from the provided Hadoop examples (`/usr/lib/hadoop-0.20-mapreduce/hadoop-examples.jar`). You need Kerberos credentials to do so.

> **Important:**
>
> Remember that the user who launches the job must exist on every host.

Configuring YARN Security

This page explains how to configure, start, and test secure YARN. For instructions on MapReduce1, see Configuring MRv1 Security.

1. Configure Secure YARN.
2. Start up the ResourceManager.
3. Start up the NodeManager.
4. Start up the MapReduce Job History Server.
5. Try Running a Map/Reduce YARN Job.
6. (Optional) Configure YARN for Long-running Applications

Step 1: Configure Secure YARN

Before you start:

- The Kerberos principals for the ResourceManager and NodeManager are configured in the `yarn-site.xml` file. The same `yarn-site.xml` file must be installed on every host machine in the cluster.
- Make sure that each user who runs YARN jobs exists on all cluster nodes (that is, on every node that hosts any YARN daemon).

To configure secure YARN:

1. Add the following properties to the `yarn-site.xml` file *on every machine* in the cluster:

```
<!-- ResourceManager security configs -->
<property>
  <name>yarn.resourcemanager.keytab</name>
  <value>/etc/hadoop/conf/yarn.keytab</value> <!-- path to the YARN keytab -->
</property>
<property>
  <name>yarn.resourcemanager.principal</name>
  <value>yarn/_HOST@YOUR-REALM.COM</value>
</property>

<!-- NodeManager security configs -->
<property>
  <name>yarn.nodemanager.keytab</name>
  <value>/etc/hadoop/conf/yarn.keytab</value> <!-- path to the YARN keytab -->
</property>
<property>
  <name>yarn.nodemanager.principal</name>
  <value>yarn/_HOST@YOUR-REALM.COM</value>
</property>
<property>
  <name>yarn.nodemanager.container-executor.class</name>
  <value>org.apache.hadoop.yarn.server.nodemanager.LinuxContainerExecutor</value>
</property>
<property>
  <name>yarn.nodemanager.linux-container-executor.group</name>
  <value>yarn</value>
</property>

<!-- To enable TLS/SSL -->
<property>
  <name>yarn.http.policy</name>
  <value>HTTPS_ONLY</value>
</property>
```

2. Add the following properties to the `mapred-site.xml` file *on every machine* in the cluster:

```
<!-- MapReduce JobHistory Server security configs -->
<property>
  <name>mapreduce.jobhistory.address</name>
  <value>host:port</value> <!-- Host and port of the MapReduce JobHistory Server; default
 port is 10020  -->
</property>
<property>
  <name>mapreduce.jobhistory.keytab</name>
  <value>/etc/hadoop/conf/mapred.keytab</value> <!-- path to the MAPRED keytab for the
 JobHistory Server -->
</property>
<property>
  <name>mapreduce.jobhistory.principal</name>
  <value>mapred/_HOST@YOUR-REALM.COM</value>
</property>

<!-- To enable TLS/SSL -->
<property>
  <name>mapreduce.jobhistory.http.policy</name>
  <value>HTTPS_ONLY</value>
</property>
```

3. Create a file called `container-executor.cfg` for the Linux Container Executor program that contains the following information:

```
yarn.nodemanager.local-dirs=<comma-separated list of paths to local NodeManager
directories. Should be same values specified in yarn-site.xml. Required to validate
paths passed to container-executor in order.>
yarn.nodemanager.linux-container-executor.group=yarn
yarn.nodemanager.log-dirs=<comma-separated list of paths to local NodeManager log
directories. Should be same values specified in yarn-site.xml. Required to set proper
permissions on the log files so that they can be written to by the user's containers
and read by the NodeManager for log aggregation.
```

```
banned.users=hdfs,yarn,mapred,bin
min.user.id=1000
```

> **Note:**
>
> In the `container-executor.cfg` file, the default setting for the `banned.users` property is `hdfs`, `yarn`, `mapred`, and `bin` to prevent jobs from being submitted using those user accounts. The default setting for the `min.user.id` property is `1000` to prevent jobs from being submitted with a user ID less than 1000, which are conventionally Unix super users. Some operating systems such as CentOS 5 use a default value of 500 and above for user IDs, not 1000. If this is the case on your system, change the default setting for the `min.user.id` property to `500`. If there are user accounts on your cluster that have a user ID less than the value specified for the `min.user.id` property, the NodeManager returns an error code of 255.

4. The path to the `container-executor.cfg` file is determined relative to the location of the container-executor binary. Specifically, the path is `<dirname of container-executor binary>/../etc/hadoop/container-executor.cfg`. If you installed the CDH 5 package, this path will always correspond to `/etc/hadoop/conf/container-executor.cfg`.

> **Note:**
>
> The `container-executor` program requires that the paths including and leading up to the directories specified in `yarn.nodemanager.local-dirs` and `yarn.nodemanager.log-dirs` to be set to 755 permissions as shown in this table on permissions on directories.

5. Verify that the ownership and permissions of the `container-executor` program corresponds to:

```
---Sr-s--- 1 root yarn 36264 May 20 15:30 container-executor
```

> **Note:** For more information about the Linux Container Executor program, see Information about Other Hadoop Security Programs.

Step 2: Start the ResourceManager

You are now ready to start the ResourceManager.

> **Note:** Always start ResourceManager before starting NodeManager.

If you're using the `/etc/init.d/hadoop-yarn-resourcemanager` script, then you can use the `service` command to run it now:

```
$ sudo service hadoop-yarn-resourcemanager start
```

You can verify that the ResourceManager is working properly by opening a web browser to http://`host`:8088/ where `host` is the name of the machine where the ResourceManager is running.

Step 3: Start the NodeManager

You are now ready to start the NodeManager.

If you're using the `/etc/init.d/hadoop-yarn-nodemanager` script, then you can use the `service` command to run it now:

```
$ sudo service hadoop-yarn-nodemanager start
```

You can verify that the NodeManager is working properly by opening a web browser to http://`host`:8042/ where `host` is the name of the machine where the NodeManager is running.

Step 4: Start the MapReduce Job History Server

You are now ready to start the MapReduce JobHistory Server.

If you're using the `/etc/init.d/hadoop-mapreduce-historyserver` script, then you can use the `service` command to run it now:

```
$ sudo service hadoop-mapreduce-historyserver start
```

You can verify that the MapReduce JobHistory Server is working properly by opening a web browser to http://`host`:19888/ where `host` is the name of the machine where the MapReduce JobHistory Server is running.

Step 5: Try Running a Map/Reduce YARN Job

You should now be able to run Map/Reduce jobs. To confirm, try launching a sleep or a pi job from the provided Hadoop examples (`/usr/lib/hadoop-mapreduce/hadoop-mapreduce-examples.jar`). You need Kerberos credentials to do so.

> **Important:** The user who launches the job must exist on every node.

To try running a MapReduce job using YARN, set the `HADOOP_MAPRED_HOME` environment variable and then submit the job. For example:

```
$ export HADOOP_MAPRED_HOME=/usr/lib/hadoop-mapreduce
$ /usr/bin/hadoop jar /usr/lib/hadoop-mapreduce/hadoop-mapreduce-examples.jar pi 10
10000
```

Step 6: (Optional) Configure YARN for Long-running Applications

Long-running applications such as Spark Streaming jobs will need additional configuration since the default settings only allow the `hdfs` user's delegation tokens a maximum lifetime of 7 days which is not always sufficient.

You can work around this by configuring the ResourceManager as a proxy user for the corresponding HDFS NameNode so that the ResourceManager can request new tokens when the existing ones are past their maximum lifetime. YARN will then be able to continue performing localization and log-aggregation on behalf of the `hdfs` user.

Set the following property in `yarn-site.xml` to `true`:

```
<property>
<name>yarn.resourcemanager.proxy-user-privileges.enabled</name>
<value>true</value>
</property>
```

Configure the following properties in `core-site.xml` on the HDFS NameNode. You can use a more restrictive configuration by specifying hosts/groups instead of `*` as in the example below.

```
<property>
<name>hadoop.proxyuser.yarn.hosts</name>
<value>*</value>
</property>

<property>
<name>hadoop.proxyuser.yarn.groups</name>
```

```
<value>*</value>
</property>
```

FUSE Kerberos Configuration

This section describes how to use FUSE (Filesystem in Userspace) and CDH with Kerberos security on your Hadoop cluster. FUSE enables you to mount HDFS, which makes HDFS files accessible just as if they were UNIX files.

To use FUSE and CDH with Kerberos security, follow these guidelines:

- For each HDFS user, make sure that there is a UNIX user with the same name. If there isn't, some files in the FUSE mount point will appear to be owned by a non-existent user. Although this is harmless, it can cause confusion.
- When using Kerberos authentication, users must run `kinit` before accessing the FUSE mount point. Failure to do this will result in I/O errors when the user attempts to access the mount point. For security reasons, it is not possible to list the files in the mount point without first running `kinit`.
- When a user runs `kinit`, all processes that run as that user can use the Kerberos credentials. It is not necessary to run `kinit` in the same shell as the process accessing the FUSE mount point.

Using kadmin to Create Kerberos Keytab Files

If your version of Kerberos does not support the Kerberos `-norandkey` option in the `xst` command, or if you must use `kadmin` because you cannot use `kadmin.local`, then you can use the following procedure to create Kerberos keytab files. Using the `-norandkey` option when creating keytabs is optional and a convenience, but it is not required.

Important:

For both MRv1 and YARN deployments: *On every machine in your cluster*, there must be a keytab file for the `hdfs` user and a keytab file for the `mapred` user. The `hdfs` keytab file must contain entries for the `hdfs` principal and an `HTTP` principal, and the `mapred` keytab file must contain entries for the `mapred` principal and an `HTTP` principal. On each respective machine, the `HTTP` principal will be the same in both keytab files.

In addition, for YARN deployments only: *On every machine in your cluster*, there must be a keytab file for the `yarn` user. The `yarn` keytab file must contain entries for the `yarn` principal and an `HTTP` principal. On each respective machine, the `HTTP` principal in the `yarn` keytab file will be the same as the `HTTP` principal in the `hdfs` and `mapred` keytab files.

For instructions, see To create the Kerberos keytab files on page 101.

Note:

These instructions illustrate an example of creating keytab files for MIT Kerberos. If you are using another version of Kerberos, refer to your Kerberos documentation for instructions. You can use either `kadmin` or `kadmin.local` to run these commands.

To create the Kerberos keytab files

Do the following steps for every host in your cluster, replacing the `fully.qualified.domain.name` in the commands with the fully qualified domain name of each host:

1. Create the `hdfs` keytab file, which contains an entry for the `hdfs` principal. This keytab file is used for the NameNode, Secondary NameNode, and DataNodes.

```
$ kadmin
kadmin:  xst -k hdfs-unmerged.keytab hdfs/fully.qualified.domain.name
```

2. Create the `mapred` keytab file, which contains an entry for the `mapred` principal. If you are using MRv1, the `mapred` keytab file is used for the JobTracker and TaskTrackers. If you are using YARN, the `mapred` keytab file is used for the MapReduce Job History Server.

```
kadmin:  xst -k mapred-unmerged.keytab mapred/fully.qualified.domain.name
```

3. **YARN only:** Create the `yarn` keytab file, which contains an entry for the `yarn` principal. This keytab file is used for the ResourceManager and NodeManager.

```
kadmin:  xst -k yarn-unmerged.keytab yarn/fully.qualified.domain.name
```

4. Create the `http` keytab file, which contains an entry for the `HTTP` principal.

```
kadmin:  xst -k http.keytab HTTP/fully.qualified.domain.name
```

5. Use the `ktutil` command to merge the previously-created keytabs:

```
$ ktutil
ktutil:  rkt hdfs-unmerged.keytab
ktutil:  rkt http.keytab
ktutil:  wkt hdfs.keytab
ktutil:  clear
ktutil:  rkt mapred-unmerged.keytab
ktutil:  rkt http.keytab
ktutil:  wkt mapred.keytab
ktutil:  clear
ktutil:  rkt yarn-unmerged.keytab
ktutil:  rkt http.keytab
ktutil:  wkt yarn.keytab
```

This procedure creates three new files: `hdfs.keytab`, `mapred.keytab` and `yarn.keytab`. These files contain entries for the `hdfs` and `HTTP` principals, the `mapred` and `HTTP` principals, and the `yarn` and `HTTP` principals respectively.

6. Use `klist` to display the keytab file entries. For example, a correctly-created `hdfs` keytab file should look something like this:

```
$ klist -e -k -t hdfs.keytab
Keytab name: WRFILE:hdfs.keytab
slot KVNO Principal
---- ---- ---------------------------------------------------------------------
   1    7    HTTP/fully.qualified.domain.name@YOUR-REALM.COM (DES cbc mode with CRC-32)

   2    7    HTTP/fully.qualified.domain.name@YOUR-REALM.COM (Triple DES cbc mode with
 HMAC/sha1)
   3    7    hdfs/fully.qualified.domain.name@YOUR-REALM.COM (DES cbc mode with CRC-32)

   4    7    hdfs/fully.qualified.domain.name@YOUR-REALM.COM (Triple DES cbc mode with
 HMAC/sha1)
```

7. To verify that you have performed the merge procedure correctly, make sure you can obtain credentials as both the `hdfs` and `HTTP` principals using the single merged keytab:

```
$ kinit -k -t hdfs.keytab hdfs/fully.qualified.domain.name@YOUR-REALM.COM
$ kinit -k -t hdfs.keytab HTTP/fully.qualified.domain.name@YOUR-REALM.COM
```

If either of these commands fails with an error message such as "`kinit: Key table entry not found while getting initial credentials`", then something has gone wrong during the merge procedure. Go back to step 1 of this document and verify that you performed all the steps correctly.

8. To continue the procedure of configuring Hadoop security in CDH 5, follow the instructions in the section To deploy the Kerberos keytab files.

Hadoop Users (user:group) and Kerberos Principals

During the Cloudera Manager/CDH installation process, several Linux user accounts and groups are created by default. These are listed in the table below. Integrating the cluster to use Kerberos for authentication requires creating Kerberos principals and keytabs for these user accounts.

> **Note:** Cloudera Manager 5.3 (and later releases) can be deployed in ***single user mode***. In single user mode, Hadoop users and groups are subsumed by `cloudera-scm:cloudera-scm`. Cloudera Manager starts all Cloudera Manager Agent processes and services running on the nodes in the cluster as a unit owned by this `cloudera-scm:cloudera-scm`. Single user mode is not recommended for production clusters.

Table 1: Users and Groups

Component (Version)	Unix User ID	Groups	Functionality
Cloudera Manager (all versions)	`cloudera-scm`	`cloudera-scm`	Clusters managed by Cloudera Manager run Cloudera Manager Server, monitoring roles, and other Cloudera Server processes as cloudera-scm. Requires keytab file named `cmf.keytab` because name is hard-coded in Cloudera Manager.
Apache Accumulo (Accumulo 1.4.3 and higher)	`accumulo`	`accumulo`	Accumulo processes run as this user.
Apache Avro	~	~	No special user:group.
Apache Flume	`flume`	`flume`	The sink that writes to HDFS as user must have write privileges.
Apache HBase	`hbase`	`hbase`	The Master and the RegionServer processes run as this user.
HDFS	`hdfs`	`hdfs, hadoop`	The NameNode and DataNodes run as this user, and the HDFS root directory as well as the directories used for edit logs should be owned by it.
Apache Hive	`hive`	`hive`	The HiveServer2 process and the Hive Metastore processes run as this user. A user must be defined for Hive access to its Metastore DB (for example, MySQL or Postgres) but it can be any identifier and does not correspond to a Unix uid. This is `javax.jdo.option.ConnectionUserName` in hive-site.xml.
Apache HCatalog	`hive`	`hive`	The WebHCat service (for REST access to Hive functionality) runs as the `hive` user.
HttpFS	`httpfs`	`httpfs`	The HttpFS service runs as this user. See HttpFS Security Configuration for instructions on how to generate the merged `httpfs-http.keytab` file.
Hue	`hue`	`hue`	Hue services run as this user.

Component (Version)	Unix User ID	Groups	Functionality
Hue Load Balancer (Cloudera Manager 5.5 and higher)	apache	apache	The Hue Load balancer has a dependency on the apache2 package that uses the apache user name. Cloudera Manager does not run processes using this user ID.
Impala	`impala`	`impala, hive`	Impala services run as this user.
Apache Kafka (Cloudera Distribution of Kafka 1.2.0)	`kafka`	`kafka`	Kafka brokers and mirror makers run as this user.
Java KeyStore KMS (CDH 5.2.1 and higher)	`kms`	`kms`	The Java KeyStore KMS service runs as this user.
Key Trustee KMS (CDH 5.3 and higher)	`kms`	`kms`	The Key Trustee KMS service runs as this user.
Key Trustee Server (CDH 5.4 and higher)	`keytrustee`	`keytrustee`	The Key Trustee Server service runs as this user.
Kudu	`kudu`	`kudu`	Kudu services run as this user.
Llama	`llama`	`llama`	Llama runs as this user.
Apache Mahout	~	~	No special users.
MapReduce	`mapred`	`mapred, hadoop`	Without Kerberos, the JobTracker and tasks run as this user. The LinuxTaskController binary is owned by this user for Kerberos.
Apache Oozie	`oozie`	`oozie`	The Oozie service runs as this user.
Parquet	~	~	No special users.
Apache Pig	~	~	No special users.
Cloudera Search	`solr`	`solr`	The Solr processes run as this user.
Apache Spark	`spark`	`spark`	The Spark History Server process runs as this user.
Apache Sentry	`sentry`	`sentry`	The Sentry service runs as this user.
Apache Sqoop	`sqoop`	`sqoop`	This user is only for the Sqoop1 Metastore, a configuration option that is not recommended.
Apache Sqoop2	`sqoop2`	`sqoop, sqoop2`	The Sqoop2 service runs as this user.
Apache Whirr	~	~	No special users.
YARN	`yarn`	`yarn, hadoop`	Without Kerberos, all YARN services and applications run as this user. The LinuxContainerExecutor binary is owned by this user for Kerberos.
Apache ZooKeeper	`zookeeper`	`zookeeper`	The ZooKeeper processes run as this user. It is not configurable.

Keytabs and Keytab File Permissions

Linux `user` accounts, such as `hdfs`, `flume`, or `mapred` are mapped to the `username` portion of the Kerberos principal names, as follows:

```
username/fqdn.example.com@YOUR-REALM.COM
```

For example, the Kerberos principal for Apache Flume would be:

```
flume/fqdn.example.com@YOUR-REALM.COM
```

Keytabs that contain multiple principals are merged automatically from individual keytabs by Cloudera Manager. If you do not use Cloudera Manager, you must merge the keytabs manually.

The table below lists the usernames to use for Kerberos principal names.

Table 2: Clusters Managed by Cloudera Manager

Component (Unix User ID)	Service	Kerberos Principals	Filename (*.keytab)	Keytab File Owner	Keytab File Group	File Permission (octal)
Cloudera Manager (`cloudera-scm`)	NA	cloudera-scm	cmf	cloudera-scm	cloudera-scm	600
Cloudera Management Service (`cloudera-scm`)	cloudera-mgmt-REPORTSMANAGER	hdfs	headlamp	cloudera-scm	cloudera-scm	600
Cloudera Management Service (`cloudera-scm`)	cloudera-mgmt-SERVICEMONITOR, cloudera-mgmt-ACTIVITYMONITOR	hue	cmon	cloudera-scm	cloudera-scm	600
Cloudera Management Service (`cloudera-scm`)	cloudera-mgmt-HOSTMONITOR	N/A	N/A	N/A	N/A	N/A
Apache Accumulo (`accumulo`)	accumulo16-ACCUMULO16_MASTER	accumulo	accumulo16	cloudera-scm	cloudera-scm	600
	accumulo16-ACCUMULO16_TRACER					
	accumulo16-ACCUMULO16_MONITOR					
	accumulo16-ACCUMULO16_GC					
	accumulo16-ACCUMULO16_TSERVER					
Flume (`flume`)	flume-AGENT	flume	flume	cloudera-scm	cloudera-scm	600
HBase (`hbase`)	hbase-HBASETHRIFTSERVER	HTTP	HTTP	cloudera-scm	cloudera-scm	600
	hbase-REGIONSERVER	hbase	hbase			
	hbase-HBASERESTSERVER					
	hbase-MASTER					
HDFS (`hdfs`)	hdfs-NAMENODE	hdfs, HTTP	hdfs	cloudera-scm	cloudera-scm	600
	hdfs-DATANODE					

Component (Unix User ID)	Service	Kerberos Principals	Filename (*.keytab)	Keytab File Owner	Keytab File Group	File Permission (octal)
	hdfs-SECONDARYNAMENODE					
Hive (`hive`)	hive-HIVESERVER2	hive	hive	cloudera-scm	cloudera-scm	600
	hive-WEBHCAT	HTTP	HTTP			
	hive-HIVEMETASTORE	hive	hive			
HttpFS (`httpfs`)	hdfs-HTTPFS	httpfs	httpfs	cloudera-scm	cloudera-scm	600
Hue (`hue`)	hue-KT_RENEWER	hue	hue	cloudera-scm	cloudera-scm	600
Impala (`impala`)	impala-STATESTORE	impala	impala	cloudera-scm	cloudera-scm	600
	impala-CATALOGSERVER					
	impala-IMPALAD					
Java KeyStore KMS (`kms`)	kms-KMS	HTTP	kms	cloudera-scm	cloudera-scm	600
Apache Kafka (`kafka`)	kafka-KAFKA_BROKER	kafka	kafka	kafka	kafka	600
Apache Kafka (`kafka`)	kafka-KAFKA_MIRROR_MAKER	kafka_mirror_maker	kafka	kafka	kafka	600
Key Trustee KMS (`kms`)	keytrustee-KMS_KEYTRUSTEE	HTTP	keytrustee	cloudera-scm	cloudera-scm	600
Llama (`llama`)	impala-LLAMA	llama, HTTP	llama	cloudera-scm	cloudera-scm	600
MapReduce (`mapred`)	mapreduce-JOBTRACKER	mapred, HTTP	mapred	cloudera-scm	cloudera-scm	600
	mapreduce- TASKTRACKER					
Oozie (`oozie`)	oozie-OOZIE_SERVER	oozie, HTTP	oozie	cloudera-scm	cloudera-scm	600
Search (`solr`)	solr-SOLR_SERVER	solr, HTTP	solr	cloudera-scm	cloudera-scm	600
Sentry (`sentry`)	sentry-SENTRY_SERVER	sentry	sentry	cloudera-scm	cloudera-scm	600
Spark (`spark`)	spark_on_yarn-SPARK_YARN_HISTORY_SERVER	spark	spark	cloudera-scm	cloudera-scm	600
YARN (`yarn`)	yarn-NODEMANAGER	yarn, HTTP	yarn	cloudera-scm	cloudera-scm	644
	yarn- RESOURCEMANAGER					600
	yarn-JOBHISTORY					600
ZooKeeper (`zookeeper`)	zookeeper-server	zookeeper	zookeeper	cloudera-scm	cloudera-scm	600

Table 3: CDH Clusters Not Managed by Cloudera Manager

Component (Unix User ID)	Service	Kerberos Principals	Filename (*.keytab)	Keytab File Owner	Keytab File Group	File Permission (octal)
Apache Accumulo (`accumulo`)	accumulo16-ACCUMULO16_MASTER	accumulo	accumulo16	accumulo	accumulo	600

Component (Unix User ID)	Service	Kerberos Principals	Filename (*.keytab)	Keytab File Owner	Keytab File Group	File Permission (octal)
	accumulo16-ACCUMULO16_TRACER					
	accumulo16-ACCUMULO16_MONITOR					
	accumulo16-ACCUMULO16_GC					
	accumulo16-ACCUMULO16_TSERVER					
Flume (`flume`)	flume-AGENT	flume	flume	flume	flume	600
HBase (`hbase`)	hbase-HBASETHRIFTSERVER	HTTP	HTTP	hbase	hbase	600
	hbase-REGIONSERVER	hbase	hbase			
	hbase-HBASERESTSERVER					
	hbase-MASTER					
HDFS (`hdfs`)	hdfs-NAMENODE	hdfs, HTTP	hdfs	hdfs	hdfs	600
	hdfs-DATANODE					
	hdfs- SECONDARYNAMENODE					
Hive (`hive`)	hive-HIVESERVER2	hive	hive	hive	hive	600
	hive-WEBHCAT	HTTP	HTTP			
	hive-HIVEMETASTORE	hive	hive			
HttpFS (`httpfs`)	hdfs-HTTPFS	httpfs	httpfs	httpfs	httpfs	600
Hue (hue)	hue-KT_RENEWER	hue	hue	hue	hue	600
Impala (`impala`)	impala-STATESTORE	impala	impala	impala	impala	600
	impala-CATALOGSERVER					
	impala-IMPALAD					
Llama (`llama`)	impala-LLAMA	llama, HTTP	llama	llama	llama	600
Java KeyStore KMS (`kms`)	kms-KMS	HTTP	kms	kms	kms	600
Apache Kafka (`kafka`)	kafka-KAFKA_BROKER	kafka	kafka	kafka	kafka	600
Apache Kafka (`kafka`)	kafka-MIRROR_MAKER	kafka_mirror_maker	kafka	kafka	kafka	600
Key Trustee KMS (`kms`)	kms-KEYTRUSTEE	HTTP	kms	kms	kms	600
MapReduce (`mapred`)	mapreduce-JOBTRACKER	mapred, HTTP	mapred	mapred	hadoop	600
	mapreduce- TASKTRACKER					
Oozie (`oozie`)	oozie-OOZIE_SERVER	oozie, HTTP	oozie	oozie	oozie	600
Search (`solr`)	solr-SOLR_SERVER	solr, HTTP	solr	solr	solr	600
Sentry (`sentry`)	sentry-SENTRY_SERVER	sentry	sentry	sentry	sentry	600

Component (Unix User ID)	Service	Kerberos Principals	Filename (*.keytab)	Keytab File Owner	Keytab File Group	File Permission (octal)
Spark (`spark`)	spark_on_yarn-SPARK_YARN_HISTORY_SERVER	spark	spark	spark	spark	600
YARN (`yarn`)	yarn-NODEMANAGER	yarn, HTTP	yarn	yarn	hadoop	644
	yarn- RESOURCEMANAGER					600
	yarn-JOBHISTORY					600
ZooKeeper (`zookeeper`)	zookeeper-server	zookeeper	zookeeper	zookeeper	zookeeper	600

Mapping Kerberos Principals to Short Names

You configure the mapping from Kerberos principals to short names in the `hadoop.security.auth_to_local` property setting in the `core-site.xml` file. Kerberos has this support natively, and Hadoop's implementation uses Kerberos's configuration language to specify the mapping.

A mapping consists of a set of rules that are evaluated in the order listed in the `hadoop.security.auth_to_local` property. The first rule that matches a principal name is used to map that principal name to a short name. Any later rules in the list that match the same principal name are ignored.

You specify the mapping rules on separate lines in the `hadoop.security.auth_to_local` property as follows:

```
<property>
  <name>hadoop.security.auth_to_local</name>
  <value>
  RULE:[<principal translation>](<acceptance filter>)<short name substitution>
  RULE:[<principal translation>](<acceptance filter>)<short name substitution>
  DEFAULT
  </value>
</property>
```

Note: You can also set `auth_to_local` rules in `krb5.conf` using MIT Kerberos. If you do, be aware that the format is different, and the rules described below will fail if used in this context. For details about the correct `krb5.conf` format, see MIT Kerberos krb5.conf.

Mapping Rule Syntax

To specify a mapping rule, use the prefix string `RULE:` followed by three sections—principal translation, acceptance filter, and short name substitution—described in more detail below. The syntax of a mapping rule is:

```
RULE:[<principal translation>](<acceptance filter>)<short name substitution>
```

Principal Translation

The first section of a rule, `<principal translation>`, performs the matching of the principal name to the rule. If there is a match, the principal translation also does the initial translation of the principal name to a short name. In the `<principal translation>` section, you specify the number of components in the principal name and the pattern you want to use to translate those principal component(s) and realm into a short name. In Kerberos terminology, a principal name is a set of components separated by slash ("/") characters.

The principal translation is composed of two parts that are both specified within "[ ]" using the following syntax:

```
[<number of components in principal name>:<initial specification of short name>]
```

where:

<number of components in principal name> – This first part specifies the number of components in the principal name (not including the realm) and must be `1` or `2`. A value of `1` specifies principal names that have a single component (for example, `hdfs`), and `2` specifies principal names that have two components (for example, `hdfs/fully.qualified.domain.name`). A principal name that has only one component will only match single-component rules, and a principal name that has two components will only match two-component rules.

<initial specification of short name> – This second part specifies a pattern for translating the principal component(s) and the realm into a short name. The variable $0 translates the realm, $1 translates the first component, and $2 translates the second component.

Here are some examples of principal translation sections. These examples use `atm@YOUR-REALM.COM` and `atm/fully.qualified.domain.name@YOUR-REALM.COM` as principal name inputs:

This Principal Translation	Translates atm@YOUR-REALM.COM into this short name	Translates atm/fully.qualified.domain.name@YOUR-REALM.COM into this short name
`[1:$1@$0]`	atm@YOUR-REALM.COM	**Rule does not match**[1]
`[1:$1]`	atm	**Rule does not match**[1]
`[1:$1.foo]`	atm.foo	**Rule does not match**[1]
`[2:$1/$2@$0]`	**Rule does not match**[2]	atm/fully.qualified.domain.name@YOUR-REALM.COM
`[2:$1/$2]`	**Rule does not match**[2]	atm/fully.qualified.domain.name
`[2:$1@$0]`	**Rule does not match**[2]	atm@YOUR-REALM.COM
`[2:$1]`	**Rule does not match**[2]	atm

Footnotes:

[1]Rule does not match because there are two components in principal name `atm/fully.qualified.domain.name@YOUR-REALM.COM`

[2]Rule does not match because there is one component in principal name `atm@YOUR-REALM.COM`

Acceptance Filter

The second section of a rule, (`<acceptance filter>`), matches the translated short name from the principal translation (that is, the output from the first section). The acceptance filter is specified in "( )" characters and is a standard regular expression. A rule matches only if the specified regular expression matches the entire translated short name from the principal translation. That is, there's an implied ^ at the beginning of the pattern and an implied $ at the end.

Short Name Substitution

The third and final section of a rule is the (`<short name substitution>`). If there is a match in the second section, the acceptance filter, the (`<short name substitution>`) section does a final translation of the short name from the first section. This translation is a sed replacement expression (`s/.../.../g`) that translates the short name from the first section into the final short name string. The short name substitution section is optional. In many cases, it is sufficient to use the first two sections only.

Converting Principal Names to Lowercase

In some organizations, naming conventions result in mixed-case usernames (for example, `John.Doe`) or even uppercase usernames (for example, `JDOE`) in Active Directory or LDAP. This can cause a conflict when the Linux username and HDFS home directory are lowercase.

To convert principal names to lowercase, append `/L` to the rule.

Example Rules

Suppose all of your service principals are either of the form `App.service-name/fully.qualified.domain.name@YOUR-REALM.COM` or

App.service-name@YOUR-REALM.COM, and you want to map these to the short name string service-name. To do this, your rule set would be:

```
<property>
  <name>hadoop.security.auth_to_local</name>
  <value>
 RULE:[1:$1](App\..*)s/App\.(.*)/$1/g
  RULE:[2:$1](App\..*)s/App\.(.*)/$1/g
  DEFAULT
  </value>
</property>
```

The first $1 in each rule is a reference to the first component of the full principal name, and the second $1 is a regular expression back-reference to text that is matched by (.*).

In the following example, suppose your company's naming scheme for user accounts in Active Directory is FirstnameLastname (for example, JohnDoe), but user home directories in HDFS are /user/firstnamelastname. The following rule set converts user accounts in the CORP.EXAMPLE.COM domain to lowercase.

```
<property>
  <name>hadoop.security.auth_to_local</name>
  <value>
  RULE:[2:$1@$0](HTTP@\QCORP.EXAMPLE.COM\E$)s/@\QCORP.EXAMPLE.COM\E$//
  RULE:[1:$1@$0](.*@\QCORP.EXAMPLE.COM\E$)s/@\QCORP.EXAMPLE.COM\E$///L
  RULE:[2:$1@$0](.*@\QCORP.EXAMPLE.COM\E$)s/@\QCORP.EXAMPLE.COM\E$///L
  DEFAULT
  </value>
</property>
```

In this example, the JohnDoe@CORP.EXAMPLE.COM principal becomes the johndoe HDFS user.

Default Rule

You can specify an optional default rule called DEFAULT (see example above). The default rule reduces a principal name down to its first component only. For example, the default rule reduces the principal names atm@YOUR-REALM.COM or atm/fully.qualified.domain.name@YOUR-REALM.COM down to atm, assuming that the default domain is YOUR-REALM.COM.

The default rule applies only if the principal is in the default realm.

If a principal name does not match any of the specified rules, the mapping for that principal name will fail.

Testing Mapping Rules

You can test mapping rules for a long principal name by running:

```
$ hadoop org.apache.hadoop.security.HadoopKerberosName name1 name2 name3
```

Configuring Authentication for Other Components

The pages in this section provide configuration steps for many components that may or may not be configured by default during Kerberos configuration. Some components use native authentication rather than or in addition to Kerberos.

Flume Authentication

Flume agents can store data on an HDFS filesystem. For clusters configured to use Kerberos authentication, Flume requires a Kerberos principal and keytab to authenticate to the cluster, which then interacts with HDFS, MapReduce, and other cluster services on behalf of the Flume agent.

Enabling Flume to use Kerberos authentication on a cluster assumes that cluster has been configured to integrate with Kerberos. See Configuring Hadoop Security in CDH 5 for details.

The steps below have been tested with CDH 5 and MIT Kerberos 5 only. The discussion includes an example of configuring user `flume` for Kerberos authentication as an HDFS client. Configuring authentication between Flume agents is not covered.

Configuring Kerberos for Flume Sinks

Writing as a single user for all HDFS sinks in a given Flume agent

The Hadoop services require a three-part principal that has the form of `username/fully.qualified.domain.name@YOUR-REALM.COM`. Cloudera recommends using `flume` as the first component and the fully qualified domain name of the host machine as the second. Assuming that Kerberos and security-enabled Hadoop have been properly configured on the Hadoop cluster itself, you must add the following parameters to the Flume agent's `flume.conf` configuration file, which is typically located at `/etc/flume-ng/conf/flume.conf`:

```
agentName.sinks.sinkName.hdfs.kerberosPrincipal =
flume/fully.qualified.domain.name@YOUR-REALM.COM
agentName.sinks.sinkName.hdfs.kerberosKeytab = /etc/flume-ng/conf/flume.keytab
```

where:

`agentName` is the name of the Flume agent being configured, which in this release defaults to the value "agent". `sinkName` is the name of the HDFS sink that is being configured. The respective sink's `type` must be `HDFS`. These properties can also be set using the substitution strings *$KERBEROS_PRINCIPAL* and *$KERBEROS_KEYTAB*, respectively.

In the previous example, `flume` is the first component of the principal name, `fully.qualified.domain.name` is the second, and `YOUR-REALM.COM` is the name of the Kerberos realm your Hadoop cluster is in. The `/etc/flume-ng/conf/flume.keytab` file contains the keys necessary for `flume/fully.qualified.domain.name@YOUR-REALM.COM` to authenticate with other services.

Flume and Hadoop also provide a simple keyword, `_HOST`, that expands to the fully qualified domain name of the host machine where the service is running, so you can use one `flume.conf` file with the same `hdfs.kerberosPrincipal` value on all of your agent host machines.

```
agentName.sinks.sinkName.hdfs.kerberosPrincipal = flume/_HOST@YOUR-REALM.COM
```

Writing as different users across multiple HDFS sinks in a single Flume agent

Hadoop users, such as secure impersonation of Hadoop users (similar to "sudo" in UNIX). This is implemented in a way similar to how Oozie implements secure user impersonation.

The following steps to set up secure impersonation from Flume to HDFS assume your cluster is configured using Kerberos. (However, impersonation also works on non-Kerberos secured clusters, and Kerberos-specific aspects should be omitted in that case.)

1. Configure Hadoop to allow impersonation. Add the following configuration properties to your `core-site.xml`.

```
<property>
     <name>hadoop.proxyuser.flume.groups</name>
     <value>group1,group2</value>
     <description>Allow the flume user to impersonate any members of group1 and
group2</description>
</property>
<property>
     <name>hadoop.proxyuser.flume.hosts</name>
     <value>host1,host2</value>
    <description>Allow the flume user to connect only from host1 and host2 to impersonate
 a user</description>
</property>
```

   You can use the wildcard character `*` to enable impersonation of any user from any host. For more information, see Secure Impersonation.

2. Set up a Kerberos keytab for the Kerberos principal and host Flume is connecting to HDFS from. This user must match the Hadoop configuration in the preceding step. For instructions, see Configuring Hadoop Security in CDH 5.
3. Configure the HDFS sink with the following configuration options:
4. `hdfs.kerberosPrincipal` - fully qualified principal. Note: `_HOST` will be replaced by the hostname of the local machine (only in-between the `/` and `@` characters)
5. `hdfs.kerberosKeytab` - location on the local machine of the keytab containing the user and host keys for the above principal
6. `hdfs.proxyUser` - the proxy user to impersonate

Example snippet (the majority of the HDFS sink configuration options have been omitted):

```
agent.sinks.sink-1.type = HDFS
agent.sinks.sink-1.hdfs.kerberosPrincipal = flume/_HOST@YOUR-REALM.COM
agent.sinks.sink-1.hdfs.kerberosKeytab = /etc/flume-ng/conf/flume.keytab
agent.sinks.sink-1.hdfs.proxyUser = weblogs

agent.sinks.sink-2.type = HDFS
agent.sinks.sink-2.hdfs.kerberosPrincipal = flume/_HOST@YOUR-REALM.COM
agent.sinks.sink-2.hdfs.kerberosKeytab = /etc/flume-ng/conf/flume.keytab
agent.sinks.sink-2.hdfs.proxyUser = applogs
```

In the above example, the flume Kerberos principal impersonates the user `weblogs` in `sink-1` and the user `applogs` in `sink-2`. This will only be allowed if the Kerberos KDC authenticates the specified principal (`flume` in this case), and the if NameNode authorizes impersonation of the specified proxy user by the specified principal.

Limitations

Flume does not support using multiple Kerberos principals or keytabs in the same agent. Creating files on HDFS as different users requires impersonation. To impersonate various other users, configure a single principal in Hadoop to impersonate all other user accounts.

In addition, the same keytab path must be used across all HDFS sinks in the same agent. Attempting to configure multiple principals or keytabs in the same agent raises a Flume error message:

```
Cannot use multiple kerberos principals in the same agent. Must restart agent to use
new principal or keytab.
```

Configuring Kerberos for Flume Thrift Source and Sink Using Cloudera Manager

The Thrift source can be configured to start in secure mode by enabling Kerberos authentication. To communicate with a secure Thrift source, the Thrift sink should also be operating in secure mode.

1. Open the Cloudera Manager Admin Console and go to the **Flume** service.
2. Click the **Configuration** tab.
3. Select **Scope** > **Agent**.
4. Select **Category** > **Main**.
5. Edit the **Configuration File** property and add the Thrift source and sink properties listed in the tables below to the configuration file.

Table 4: Thrift Source Properties

Property	Description
`kerberos`	Set to `true` to enable Kerberos authentication. The `agent-principal` and `agent-keytab` properties are required for successful authentication. The Thrift source in secure mode, will accept connections only from Thrift sinks that have Kerberos-enabled and are successfully authenticated to the KDC.

Property	Description
`agent-principal`	The Kerberos principal used by the Thrift Source to authenticate to the KDC.
`agent-keytab`	The path to the keytab file used by the Thrift Source in combination with the `agent-principal` to authenticate to the KDC.

Table 5: Thrift Sink Properties

Property	Description
`kerberos`	Set to `true` to enable Kerberos authentication. In Kerberos mode, client-principal, client-keytab and server-principal are required for successful authentication and communication to a Kerberos enabled Thrift Source.
`client-principal`	The principal used by the Thrift Sink to authenticate to the Kerberos KDC.
`client-keytab`	The path to the keytab file used by the Thrift Sink in combination with the `client-principal` to authenticate to the KDC.
`server-principal`	The principal of the Thrift Source to which this Thrift Sink connects.

Note: Since Cloudera Manager generates the Flume keytab files for you, and the locations of the keytab files cannot be known beforehand, substitution variables are required for Flume. Cloudera Manager provides two Flume substitution variables called `$KERBEROS_PRINCIPAL` and `$KERBEROS_KEYTAB` to configure the principal name and the keytab file path respectively on each host.

Make sure you are configuring these properties for ***each*** Thrift source and sink instance managed by Cloudera Manager. For example, for agent `a1`, source `r1`, and sink `k1`, you would add the following properties:

```
# Kerberos properties for Thrift source s1
a1.sources.r1.kerberos=true
a1.sources.r1.agent-principal=<source_principal>
a1.sources.r1.agent-keytab=<path/to/source/keytab>

# Kerberos properties for Thrift sink k1
a1.sinks.k1.kerberos=true
a1.sinks.k1.client-principal=<sink_principal>
a1.sinks.k1.client-keytab=<path/to/sink/keytab>
a1.sinks.k1.server-principal=<path/to/source/keytab>
```

6. Click **Save Changes** to commit the changes.
7. Restart the Flume service.

Configuring Kerberos for Flume Thrift Source and Sink Using the Command Line

The Thrift source can be configured to start in secure mode by enabling Kerberos authentication. To communicate with a secure Thrift source, the Thrift sink should also be operating in secure mode.

The following tables list the properties that must be configured in the `/etc/flume-ng/conf/flume.conf` file to enable Kerberos for Flume's Thrift source and sink instances.

Table 6: Thrift Source Properties

Property	Description
`kerberos`	Set to `true` to enable Kerberos authentication. The `agent-principal` and `agent-keytab` properties are required for successful authentication. The Thrift source in secure mode, will accept connections only from Thrift sinks that have Kerberos-enabled and are successfully authenticated to the KDC.
`agent-principal`	The Kerberos principal used by the Thrift Source to authenticate to the KDC.
`agent-keytab`	The path to the keytab file used by the Thrift Source in combination with the `agent-principal` to authenticate to the KDC.

Table 7: Thrift Sink Properties

Property	Description
`kerberos`	Set to `true` to enable Kerberos authentication. In Kerberos mode, client-principal, client-keytab and server-principal are required for successful authentication and communication to a Kerberos enabled Thrift Source.
`client-principal`	The principal used by the Thrift Sink to authenticate to the Kerberos KDC.
`client-keytab`	The path to the keytab file used by the Thrift Sink in combination with the `client-principal` to authenticate to the KDC.
`server-principal`	The principal of the Thrift Source to which this Thrift Sink connects.

Make sure you are configuring these properties for ***each*** Thrift source and sink instance. For example, for agent `a1`, source `r1`, and sink `k1`, you would add the following properties:

```
# Kerberos properties for Thrift source s1
a1.sources.r1.kerberos=true
a1.sources.r1.agent-principal=<source_principal>
a1.sources.r1.agent-keytab=<path/to/source/keytab>

# Kerberos properties for Thrift sink k1
a1.sinks.k1.kerberos=true
a1.sinks.k1.client-principal=<sink_principal>
a1.sinks.k1.client-keytab=<path/to/sink/keytab>
a1.sinks.k1.server-principal=<path/to/source/keytab>
```

Configure these sets of properties for as many instances of the Thrift source and sink as needed to enable Kerberos.

Flume Account Requirements

This section provides an overview of the account and credential requirements for Flume to write to a Kerberized HDFS. Note the distinctions between the Flume agent machine, DataNode machine, and NameNode machine, as well as the `flume` Unix user account versus the `flume` Hadoop/Kerberos user account.

- Each Flume agent machine that writes to HDFS (using a configured HDFS sink) needs a Kerberos principal of the form:

```
flume/fully.qualified.domain.name@YOUR-REALM.COM
```

  where `fully.qualified.domain.name` is the fully qualified domain name of the given Flume agent host machine, and `YOUR-REALM.COM` is the Kerberos realm.

- Each Flume agent machine that writes to HDFS does *not* need to have a `flume` Unix user account to write files owned by the `flume` Hadoop/Kerberos user. Only the keytab for the `flume` Hadoop/Kerberos user is required on the Flume agent machine.
- DataNode machines do *not* need Flume Kerberos keytabs and also do *not* need the `flume` Unix user account.
- TaskTracker (MRv1) or NodeManager (YARN) machines need a `flume` Unix user account *if and only if* MapReduce jobs are being run as the `flume` Hadoop/Kerberos user.
- The NameNode machine needs to be able to resolve the groups of the `flume` user. The groups of the `flume` user on the NameNode machine are mapped to the Hadoop groups used for authorizing access.
- The NameNode machine does *not* need a Flume Kerberos keytab.

Testing the Flume HDFS Sink Configuration

To test whether your Flume HDFS sink is properly configured to connect to your secure HDFS cluster, you must run data through Flume. An easy way to do this is to configure a Netcat source, a Memory channel, and an HDFS sink. Start Flume with that configuration, and use the nc command (available freely online and with many UNIX distributions) to send events to the Netcat source port. The resulting events should appear on HDFS in the configured location. If the events do not appear, check the Flume log at `/var/log/flume-ng/flume.log` for any error messages related to Kerberos.

Writing to a Secure HBase Cluster

Before you write to a secure HBase cluster, be aware of the following:

- Flume must be configured to use Kerberos security as documented above, and HBase must be configured to use Kerberos security as documented in HBase Security Configuration.
- The `hbase-site.xml` file, which must be configured to use Kerberos security, must be in Flume's classpath or `HBASE_HOME/conf`.
- HBaseSink `org.apache.flume.sink.hbase.HBaseSink` supports secure HBase, but AsyncHBaseSink `org.apache.flume.sink.hbase.AsyncHBaseSink` does not.
- The Flume HBase sink takes the `kerberosPrincipal` and `kerberosKeytab` parameters:
  - `kerberosPrincipal` – specifies the Kerberos principal to be used
  - `kerberosKeytab` – specifies the path to the Kerberos keytab
  - These are defined as:

```
agent.sinks.hbaseSink.kerberosPrincipal = flume/fully.qualified.domain.name@YOUR-REALM.COM
agent.sinks.hbaseSink.kerberosKeytab = /etc/flume-ng/conf/flume.keytab
```

  - You can use the `$KERBEROS_PRINCIPAL` and `$KERBEROS_KEYTAB` substitution variables to configure the principal name and the keytab file path. See the following documentation for steps on how to configure the substitution variables: Use Substitution Variables for the Kerberos Principal and Keytab.

- If HBase is running with the AccessController coprocessor, the `flume` user (or whichever user the agent is running as) must have permissions to write to the same table and the column family that the sink is configured to write to. You can grant permissions using the `grant` command from HBase shell as explained in HBase Security Configuration.
- The Flume HBase Sink does not currently support impersonation; it will write to HBase as the user the agent is being run as.
- If you want to use HDFS Sink and HBase Sink to write to HDFS and HBase from the same agent respectively, both sinks have to use the same principal and keytab. If you want to use different credentials, the sinks have to be on different agents.
- Each Flume agent machine that writes to HBase (using a configured HBase sink) needs a Kerberos principal of the form:

```
flume/fully.qualified.domain.name@YOUR-REALM.COM
```

where `fully.qualified.domain.name` is the fully qualified domain name of the given Flume agent host machine, and YOUR-REALM.COM is the Kerberos realm.

Using Substitution Variables for Flume Kerberos Principal and Keytab

Required Role: Configurator (also provided by **Cluster Administrator, Full Administrator**)

Flume instances running on clusters that use Kerberos for authentication require the Kerberos principal and keytab for the HDFS sink or HBase sink to be configured in the `flume.conf` file. Cloudera Manager generates Flume keytab files automatically. However, the location of the files is not known in advance, so Cloudera Manager provides two substitution variables to use for these Kerberos artifacts as shown in the table:

Kerberos	Description	Substitution variable
kerberosPrincipal	Fully qualified principal name, such as *service/user@DOMAIN*	$KERBEROS_PRINCIPAL
kerberosKeytab	Location on the host to keytabs for kerberos principal	$KERBEROS_KEYTAB

Use the Flume substitution variables for the principal name and the keytab file path on each host. Here are some usage examples:

- HDFS Sink Example on page 116
- HBase Sink Example on page 116
- Configuring Flume Substitution Variables Using Cloudera Manager Admin Console on page 116

HDFS Sink Example

The following example shows an HDFS sink configuration in the `flume.conf` file (the majority of the HDFS sink configuration options have been omitted):

```
agent.sinks.sink-1.type = HDFS
agent.sinks.sink-1.hdfs.kerberosPrincipal = flume/_HOST@YOUR-REALM.COM
agent.sinks.sink-1.hdfs.kerberosKeytab = /etc/flume-ng/conf/flume.keytab
agent.sinks.sink-1.hdfs.proxyUser = weblogs
```

The text below shows the same configuration options with the substitution variables:

```
agent.sinks.sink-1.type = hdfs
agent.sinks.sink-1.hdfs.kerberosPrincipal = $KERBEROS_PRINCIPAL
agent.sinks.sink-1.hdfs.kerberosKeytab = $KERBEROS_KEYTAB
agent.sinks.sink-1.hdfs.proxyUser = weblogs
```

HBase Sink Example

The following example shows an HBase sink configuration in the `flume.conf` file (the majority of the HBase sink configuration options have been omitted):

```
agent.sinks.sink-1.type = hbase
agent.sinks.sink-1.kerberosPrincipal = flume/_HOST@YOUR-REALM.COM
agent.sinks.sink-1.kerberosKeytab = /etc/flume-ng/conf/flume.keytab
```

The text below shows the same configuration options with the substitution variables:

```
agent.sinks.sink-1.type = hbase
agent.sinks.sink-1.kerberosPrincipal = $KERBEROS_PRINCIPAL
agent.sinks.sink-1.kerberosKeytab = $KERBEROS_KEYTAB
```

Configuring Flume Substitution Variables Using Cloudera Manager Admin Console

Complete the following steps to have Cloudera Manager add these variables to the `flume.conf` file on every host that Cloudera Manager manages.

1. Log in to the Cloudera Manager Admin Console.
2. Select **Clusters** > **Flume** service.
3. Click the **Configuration** tab.
4. Select **Agent** for the **Scope** filter.
5. In the **Configuration File** property, set the agent sinks used by your cluster equal to the substitution variables. For example, for an HDFS sink, the values would be specified as follows:

```
agent.sinks.sink-1.type = hdfs
agent.sinks.sink-1.hdfs.kerberosPrincipal = $KERBEROS_PRINCIPAL
agent.sinks.sink-1.hdfs.kerberosKeytab = $KERBEROS_KEYTAB
agent.sinks.sink-1.hdfs.proxyUser = weblogs
```

6. Click **Save Changes**.

HBase Authentication

To configure HBase security, complete the following tasks:

1. **Configure HBase Authentication**: You must establish a mechanism for HBase servers and clients to securely identify themselves with HDFS, ZooKeeper, and each other. This ensures that hosts are who they claim to be.

   Note:

   - To enable HBase to work with Kerberos security, you must perform the installation and configuration steps in Configuring Hadoop Security in CDH 5 and ZooKeeper Security Configuration.
   - Although an HBase Thrift server can connect to a secured Hadoop cluster, access is not secured from clients to the HBase Thrift server. To encrypt communication between clients and the HBase Thrift Server, see Configuring TLS/SSL for HBase Thrift Server on page 208.

   The following sections describe how to use Apache HBase and CDH 5 with Kerberos security:

   - Configuring Kerberos Authentication for HBase on page 117
   - Configuring Secure HBase Replication on page 123
   - Configuring the HBase Client TGT Renewal Period on page 124

2. **Configure HBase Authorization**: You must establish rules for the resources that clients are allowed to access. For more information, see Configuring HBase Authorization on page 186.

Using the Hue HBase App

Hue includes an HBase App that allows you to interact with HBase through a Thrift proxy server. Because Hue sits between the Thrift server and the client, the Thrift server assumes that all HBase operations come from the hue user and not the client. To ensure that users in Hue are only allowed to perform HBase operations assigned to their own credentials, and not those of the hue user, you must enable HBase impersonation. For more information about the how to enable doAs Impersonation for the HBase Browser Application, see Enabling the HBase Browser Application with doAs Impersonation.

Configuring Kerberos Authentication for HBase

Using Kerberos for authentication for the HBase component requires that you also use Kerberos authentication for ZooKeeper. This means that HBase Master, RegionServer, and client hosts must each have a Kerberos principal for authenticating to the ZooKeeper ensemble. The steps below provide the details. Before you start, be sure that:

- Kerberos is enabled for the cluster, as detailed in Enabling Kerberos Authentication Using the Wizard.
- Kerberos principals for Cloudera Manager Server, HBase, and ZooKeeper hosts exist and are available for use. See Managing Kerberos Credentials Using Cloudera Manager on page 65 for details.

Cloudera Manager automatically configures authentication between HBase to ZooKeeper and sets up the HBase Thrift gateway to support impersonation (doAs). However, you must manually configure the HBase REST service for Kerberos (it currently uses Simple authentication by default, instead of Kerberos). See Configure HBase REST Server for Kerberos Authentication on page 118 for details.

> **Note:** Impersonation (doAs) cannot be used with Thrift framed transport (TFramedTransport) because SASL does not work with Thrift framed transport.

You can use either Cloudera Manager or the command line to configure Kerberos authentication for HBase. Using Cloudera Manager simplifies the process, but both approaches are detailed below. This page includes these topics:

Configuring Kerberos Authentication for HBase Using Cloudera Manager

Cloudera Manager simplifies the task of configuring Kerberos authentication for HBase.

Configure HBase Servers to Authenticate with a Secure HDFS Cluster Using Cloudera Manager

Required Role: **Cluster Administrator** or **Full Administrator**

1. Log on to Cloudera Manager Admin Console.
2. Go to the HBase service (select **Clusters** > **HBASE**).
3. Click the **Configuration** tab.
4. Under the Scope filter, click **HBase (Service-Wide)**.
5. Under the Category filter, click **Security**.
6. Ensure the Kerberos principal for the HBase service was generated.
7. Find the **HBase Secure Authentication** property (type "HBase Secure" in the Search box, if necessary), and confirm (or enter) the principal to use for HBase.
8. Select **kerberos** as the authentication type.
9. Click **Save Changes**.
10. Restart the role.
11. Restart the service. Select **Restart** from the **Actions** drop-down menu adjacent to HBASE-n (Cluster).

Configure HBase Servers and Clients to Authenticate with a Secure ZooKeeper

As mentioned above, secure HBase also requires secure ZooKeeper. The various HBase host systems—Master, RegionServer, and client—must have a principal to use to authenticate to the secure ZooKeeper ensemble. This is handled transparently by Cloudera Manager when you enable Kerberos as detailed above.

Configure HBase REST Server for Kerberos Authentication

Currently, the HBase REST Server uses Simple (rather than Kerberos) authentication by default. You must manually modify the setting using the Cloudera Manager Admin Console, as follows:

1. Log on to Cloudera Manager Admin Console.
2. Select **Clusters** > **HBASE**.
3. Click the **Configuration** tab.
4. Under the Scope filter, click **HBase (Service-Wide)**.
5. Under the Category filter, click **Security**.
6. Find the **HBase REST Authentication** property:

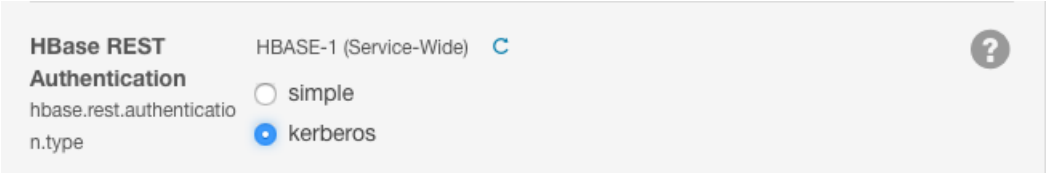

7. Click **kerberos** to select Kerberos instead of simple authentication.
8. Click **Save Changes**.
9. Restart the role.

10. Restart the service. Select **Restart** from the **Actions** drop-down menu adjacent to HBASE-n (Cluster).

Configuring Kerberos Authentication for HBase Using the Command Line

> **Important:**
> - Follow these command-line instructions on systems that do not use Cloudera Manager.
> - This information applies specifically to CDH 5.14.x. See Cloudera Documentation for information specific to other releases.

Configure HBase Servers to Authenticate with a Secure HDFS Cluster Using the Command Line

To configure HBase servers to authenticate with a secure HDFS cluster, do the following:

Enable HBase Authentication

Set the `hbase.security.authentication` property to `kerberos` in `hbase-site.xml` on every host acting as an HBase master, RegionServer, or client. In CDH 5, `hbase.rpc.engine` is automatically detected and does not need to be set.

```
<property>
    <name>hbase.security.authentication</name>
    <value>kerberos</value>
</property>
```

Configure HBase Kerberos Principals

To run on a secure HDFS cluster, HBase must authenticate itself to the HDFS services. HBase acts as a Kerberos principal and needs Kerberos credentials to interact with the Kerberos-enabled HDFS daemons. You can authenticate a service by using a keytab file, which contains a key that allows the service to authenticate to the Kerberos Key Distribution Center (KDC).

1. Create a service principal for the HBase server using the following syntax. This principal is used to authenticate the HBase server with the HDFS services. Cloudera recommends using `hbase` as the username.

```
$ kadmin
kadmin: addprinc -randkey hbase/fully.qualified.domain.name@YOUR-REALM.COM
```

   `fully.qualified.domain.name` is the host where the HBase server is running, and `YOUR-REALM` is the name of your Kerberos realm.

2. Create a keytab file for the HBase server.

```
$ kadmin
kadmin: xst -k hbase.keytab hbase/fully.qualified.domain.name
```

3. Copy the `hbase.keytab` file to the `/etc/hbase/conf` directory on the HBase server host. The owner of the `hbase.keytab` file should be the `hbase` user, and the file should have owner-only read permissions—that is, assign the file `0400` permissions and make it owned by `hbase:hbase`.

```
-r--------  1 hbase   hbase     1343 2012-01-09 10:39  hbase.keytab
```

4. To test that the keytab file was created properly, try to obtain Kerberos credentials as the HBase principal using only the keytab file. Substitute your `fully.qualified.domain.name` and realm in the following command:

```
$ kinit -k -t /etc/hbase/conf/hbase.keytab
hbase/fully.qualified.domain.name@YOUR-REALM.COM
```

5. In the `/etc/hbase/conf/hbase-site.xml` configuration file on *all* cluster hosts running the HBase daemon, add the following lines:

```
<property>
  <name>hbase.regionserver.kerberos.principal</name>
  <value>hbase/_HOST@YOUR-REALM.COM</value>
</property>

<property>
  <name>hbase.regionserver.keytab.file</name>
  <value>/etc/hbase/conf/hbase.keytab</value>
</property>

<property>
  <name>hbase.master.kerberos.principal</name>
  <value>hbase/_HOST@YOUR-REALM.COM</value>
</property>

<property>
<name>hbase.master.keytab.file</name>
<value>/etc/hbase/conf/hbase.keytab</value>
</property>
```

Configure HBase Servers and Clients to Authenticate with a Secure ZooKeeper

To run a secure HBase, you must also use a secure ZooKeeper. To use a secure ZooKeeper, each HBase host machine (Master, RegionServer, and client) must have a principal that allows it to authenticate with your secure ZooKeeper ensemble. The steps below assume that:

- ZooKeeper has been secured per the steps in ZooKeeper Security Configuration.
- ZooKeeper is *not* managed by HBase.
- You have successfully completed steps above (Enable HBase Authentication, Configure HBase Kerberos Principals) and have principal and keytab files in place for every HBase server and client.

To configure HBase Servers and clients to authenticate to ZooKeeper, you must:

Configure HBase JVMs (all Masters, RegionServers, and clients) to Use JAAS

1. On each host, set up a Java Authentication and Authorization Service (JAAS) by creating a `/etc/hbase/conf/zk-jaas.conf` file that contains the following:

```
   Client {
      com.sun.security.auth.module.Krb5LoginModule required
      useKeyTab=true
      useTicketCache=false
      keyTab="/etc/hbase/conf/hbase.keytab"
      principal="hbase/fully.qualified.domain.name@<YOUR-REALM>";
   };
```

2. Modify the `hbase-env.sh` file on HBase server and client hosts to include the following:

```
   export HBASE_OPTS="$HBASE_OPTS
-Djava.security.auth.login.config=/etc/hbase/conf/zk-jaas.conf"
   export HBASE_MANAGES_ZK=false
```

3. Restart the HBase cluster.

Configure the HBase Servers (Masters and RegionServers) to Use Authentication to Connect to ZooKeeper

Note: These steps are required for command-line configuration only. Cloudera Manager does this automatically.

1. Update your `hbase-site.xml` on each HBase server host with the following properties:

```
<configuration>
   <property>
      <name>hbase.zookeeper.quorum</name>
      <value>$ZK_NODES</value>
   </property>
   <property>
      <name>hbase.cluster.distributed</name>
      <value>true</value>
   </property>
</configuration>
```

   $ZK_NODES is the comma-separated list of hostnames of the ZooKeeper Quorum hosts that you configured according to the instructions in ZooKeeper Security Configuration.

2. Add the following lines to the ZooKeeper configuration file `zoo.cfg`:

```
kerberos.removeHostFromPrincipal=true
kerberos.removeRealmFromPrincipal=true
```

3. Restart ZooKeeper.

Configure Authentication for the HBase REST and Thrift Gateways

By default, the REST gateway does not support impersonation, but accesses HBase as a statically configured user. The actual user who initiated the request is not tracked. With impersonation, the REST gateway user is a proxy user. The HBase server records the actual user who initiates each request and uses this information to apply authorization.

1. Enable support for proxy users by adding the following properties to `hbase-site.xml`. Substitute the REST gateway proxy user for *$USER*, and the allowed group list for *$GROUPS*.

```
<property>
  <name>hbase.security.authorization</name>
  <value>true</value>
</property>
<property>
  <name>hadoop.proxyuser.$USER.groups</name>
  <value>$GROUPS</value>
</property>
<property>
  <name>hadoop.proxyuser.$USER.hosts</name>
  <value>$GROUPS</value>
</property>
```

2. Enable REST gateway impersonation by adding the following to the `hbase-site.xml` file for every REST gateway:

```
<property>
  <name>hbase.rest.authentication.type</name>
  <value>kerberos</value>
</property>
<property>
  <name>hbase.rest.authentication.kerberos.principal</name>
  <value>HTTP/fully.qualified.domain.name@<YOUR-REALM/value>
</property>
<property>
  <name>hbase.rest.authentication.kerberos.keytab</name>
  <value>/etc/hbase/conf/hbase.keytab</value>
</property>
```

3. Add the following properties to `hbase-site.xml` for each Thrift gateway, replacing the Kerberos principal with a valid value:

```
<property>
  <name>hbase.thrift.keytab.file</name>
  <value>/etc/hbase/conf/hbase.keytab</value>
</property>
```

```
<property>
  <name>hbase.thrift.kerberos.principal</name>
  <value>hbase/fully.qualified.domain.name@<YOUR-REALM</value>
</property>
<property>
  <name>hbase.thrift.security.qop</name>
  <value>auth</value>
</property>
```

The value for the property `hbase.thrift.security.qop` can be one of the following:

- `auth-conf`—Authentication, integrity, and confidentiality checking
- `auth-int`—Authentication and integrity checking
- `auth`—Authentication checking only

4. To use the Thrift API principal to interact with HBase, add the `hbase.thrift.kerberos.principal` to the *acl* table. For example, to provide administrative access to the Thrift API principal `thrift_server`, run an HBase Shell command like the following:

```
hbase> grant'thrift_server', 'RWCA'
```

5. Optional: Configure HTTPS transport for Thrift by configuring the following parameters, substituting the placeholders with actual values:

```
<property>
  <name>hbase.thrift.ssl.enabled</name>
  <value>true</value>
</property>
<property>
  <name>hbase.thrift.ssl.keystore.store</name>
  <value>LOCATION_OF_KEYSTORE</value>
</property>
<property>
  <name>hbase.thrift.ssl.keystore.password</name>
  <value>KEYSTORE_PASSWORD</value>
</property>
<property>
  <name>hbase.thrift.ssl.keystore.keypassword</name>
  <value>LOCATION_OF_KEYSTORE_KEY_PASSWORD</value>
</property>
```

The Thrift gateway authenticates with HBase using the supplied credential. No authentication is performed by the Thrift gateway itself. All client access through the Thrift gateway uses the gateway's credential, and all clients have its privileges.

Configure doAs Impersonation for the HBase Thrift Gateway

> **Note:** If you use framed transport, you cannot use doAs impersonation, because SASL does not work with Thrift framed transport.

doAs Impersonation provides a flexible way to use the same client to impersonate multiple principals. doAs is supported only in Thrift 1, not Thrift 2.

Enable doAs support by adding the following properties to `hbase-site.xml` on each Thrift gateway:

```
<property>
  <name>hbase.regionserver.thrift.http</name>
  <value>true</value>
</property>
<property>
```

```
  <name>hbase.thrift.support.proxyuser</name>
  <value>true/value>
</property>
```

See the demo client for information on using doAs impersonation in your client applications.

Start HBase

If the configuration worked, you see something similar to the following in the HBase Master and RegionServer logs when you start the cluster:

```
INFO zookeeper.ZooKeeper: Initiating client connection,
connectString=ZK_QUORUM_SERVER:2181 sessionTimeout=180000 watcher=master:60000
INFO zookeeper.ClientCnxn: Opening socket connection to server /ZK_QUORUM_SERVER:2181
INFO zookeeper.RecoverableZooKeeper: The identifier of this process is
PID@ZK_QUORUM_SERVER
INFO zookeeper.Login: successfully logged in.
INFO client.ZooKeeperSaslClient: Client will use GSSAPI as SASL mechanism.
INFO zookeeper.Login: TGT refresh thread started.
INFO zookeeper.ClientCnxn: Socket connection established to ZK_QUORUM_SERVER:2181,
initiating session
INFO zookeeper.Login: TGT valid starting at:        Sun Apr 08 22:43:59 UTC 2012
INFO zookeeper.Login: TGT expires:                  Mon Apr 09 22:43:59 UTC 2012
INFO zookeeper.Login: TGT refresh sleeping until: Mon Apr 09 18:30:37 UTC 2012
INFO zookeeper.ClientCnxn: Session establishment complete on server ZK_QUORUM_SERVER:2181,
 sessionid = 0x134106594320000, negotiated timeout = 180000
```

Configuring Secure HBase Replication

If you are using HBase Replication and you want to make it secure, read this section for instructions. Before proceeding, you should already have configured HBase Replication by following the instructions in the HBase Replication section of the *CDH 5 Installation Guide*.

To configure secure HBase replication, you must configure cross realm support for Kerberos, ZooKeeper, and Hadoop.

> **Note:** HBase peer-to-peer replication from a non-Kerberized cluster to a Kerberized cluster is not supported.

To configure secure HBase replication:

1. Create krbtgt principals for the two realms. For example, if you have two realms called `ONE.COM` and `TWO.COM`, you need to add the following principals: `krbtgt/ONE.COM@TWO.COM` and `krbtgt/TWO.COM@ONE.COM`. Add these two principals at both realms. There must be at least one common encryption mode between these two realms.

```
kadmin: addprinc -e "<enc_type_list>" krbtgt/ONE.COM@TWO.COM
kadmin: addprinc -e "<enc_type_list>" krbtgt/TWO.COM@ONE.COM
```

2. Add rules for creating short names in Zookeeper. To do this, add a system level property in `java.env`, defined in the `conf` directory. Here is an example rule that illustrates how to add support for the realm called `ONE.COM`, and have two members in the principal (such as `service/instance@ONE.COM`):

```
-Dzookeeper.security.auth_to_local=RULE:[2:\$1@\$0](.*@\\QONE.COM\\E$)s/@\\QONE.COM\\E$//DEFAULT
```

   The above code example adds support for the `ONE.COM` realm in a different realm. So, in the case of replication, you must add a rule for the primary cluster realm in the replica cluster realm. `DEFAULT` is for defining the default rule.

3. Add rules for creating short names in the Hadoop processes. To do this, add the `hadoop.security.auth_to_local` property in the `core-site.xml` file in the replica cluster. For example, to add support for the `ONE.COM` realm:

```
<property>
  <name>hadoop.security.auth_to_local</name>
  <value>
    RULE:[2:$1@$0](.*@\QONE.COM\E$)s/@\QONE.COM\E$//
  DEFAULT
  </value>
</property>
```

For more information about adding rules, see Configuring the Mapping from Kerberos Principals to Short Names.

Configuring the HBase Client TGT Renewal Period

An HBase client user must also have a Kerberos principal which typically has a password that only the user knows. You should configure the `maxrenewlife` setting for the client's principal to a value that allows the user enough time to finish HBase client processes before the ticket granting ticket (TGT) expires. For example, if the HBase client processes require up to four days to complete, you should create the user's principal and configure the `maxrenewlife` setting by using this command:

```
kadmin: addprinc -maxrenewlife 4days
```

HCatalog Authentication

This section describes how to configure HCatalog in CDH 5 with Kerberos security in a Hadoop cluster:

For more information about HCatalog see Installing and Using HCatalog.

Before You Start

Secure Web HCatalog requires a running remote Hive metastore service configured in secure mode. See Hive MetaStoreServer Security Configuration for instructions. Running secure WebHCat with an embedded repository is not supported.

Step 1: Create the HTTP keytab file

You need to create a `keytab` file for WebHCat. Follow these steps:

1. Create the file:

```
kadmin: addprinc -randkey HTTP/fully.qualified.domain.name@YOUR-REALM.COM
kadmin: xst -k HTTP.keytab HTTP/fully.qualified.domain.name
```

2. Move the file into the WebHCat configuration directory and restrict its access exclusively to the `hcatalog` user:

```
$ mv HTTP.keytab /etc/webhcat/conf/
$ chown hcatalog /etc/webhcat/conf/HTTP.keytab
$ chmod 400 /etc/webhcat/conf/HTTP.keytab
```

Step 2: Configure WebHCat to Use Security

Create or edit the WebHCat configuration file `webhcat-site.xml` in the configuration directory and set following properties:

Property	Value
templeton.kerberos.secret	Any random value
templeton.kerberos.keytab	`/etc/webhcat/conf/HTTP.keytab`
templeton.kerberos.principal	HTTP/fully.qualified.domain.name@YOUR-REALM.COM

Example configuration:

```
<property>
    <name>templeton.kerberos.secret</name>
    <value>SuPerS3c3tV@lue!</value>
  </property>

  <property>
    <name>templeton.kerberos.keytab</name>
    <value>/etc/webhcat/conf/HTTP.keytab</value>
  </property>

  <property>
    <name>templeton.kerberos.principal</name>
    <value>HTTP/fully.qualified.domain.name@YOUR-REALM.COM</value>
  </property>
```

Step 3: Create Proxy Users

WebHCat needs access to your NameNode to work properly, and so you must configure Hadoop to allow impersonation from the `hcatalog` user. To do this, edit your `core-site.xml` configuration file and set the `hadoop.proxyuser.HTTP.hosts` and `hadoop.proxyuser.HTTP.groups` properties to specify the hosts from which HCatalog can do the impersonation and what users can be impersonated. You can use the value `*` for "any".

Example configuration:

```
  <property>
    <name>hadoop.proxyuser.HTTP.hosts</name>
    <value>*</value>
  </property>
  <property>
    <name>hadoop.proxyuser.HTTP.groups</name>
    <value>*</value>
  </property>
```

Step 4: Verify the Configuration

After restarting WebHcat you can verify that it is working by using `curl` (you may need to run `kinit` first):

```
$ curl --negotiate -i -u :
'http://fully.qualified.domain.name:50111/templeton/v1/ddl/database'
```

Hive Authentication

Hive authentication involves configuring Hive metastore, HiveServer2, and all Hive clients to use your deployment of LDAP/Active Directory Kerberos on your cluster.

Here is a summary of the status of Hive authentication in CDH 5:

- HiveServer2 supports authentication of the Thrift client using Kerberos or user/password validation backed by LDAP. For configuration instructions, see HiveServer2 Security Configuration.
- Earlier versions of HiveServer do not support Kerberos authentication for clients. However, the Hive MetaStoreServer does support Kerberos authentication for Thrift clients. For configuration instructions, see Hive MetaStoreServer Security Configuration.

  See also: Using Hive to Run Queries on a Secure HBase Server on page 133

For authorization, Hive uses Apache Sentry to enable role-based, fine-grained authorization for HiveServer2. See Apache Sentry Overview.

> **Important:** Cloudera does not support Apache Ranger or Hive's native authorization frameworks for configuring access control in Hive. Use Cloudera-supported Apache Sentry instead.

HiveServer2 Security Configuration

HiveServer2 supports authentication of the Thrift client using the following methods:

- Kerberos authentication
- LDAP authentication

  Starting with CDH 5.7, clusters running LDAP-enabled HiveServer2 deployments also accept Kerberos authentication. This ensures that users are not forced to enter usernames/passwords manually, and are able to take advantage of the multiple authentication schemes SASL offers. In CDH 5.6 and lower, HiveServer2 stops accepting delegation tokens when any alternate authentication is enabled.

Kerberos authentication is supported between the Thrift client and HiveServer2, and between HiveServer2 and secure HDFS. LDAP authentication is supported only between the Thrift client and HiveServer2.

To configure HiveServer2 to use one of these authentication modes, configure the `hive.server2.authentication` configuration property.

> **Note:** To configure dual authentication (both Kerberos and LDAP), configure Kerberos for HiveServer2 and set the `hive.server2.authentication` property to `LDAP`.

Enabling Kerberos Authentication for HiveServer2

If you configure HiveServer2 to use Kerberos authentication, HiveServer2 acquires a Kerberos ticket during startup. HiveServer2 requires a principal and keytab file specified in the configuration. Client applications (for example, JDBC or Beeline) must have a valid Kerberos ticket before initiating a connection to HiveServer2.

Configuring HiveServer2 for Kerberos-Secured Clusters

To enable Kerberos Authentication for HiveServer2, add the following properties in the `/etc/hive/conf/hive-site.xml` file:

```
<property>
  <name>hive.server2.authentication</name>
  <value>KERBEROS</value>
</property>
<property>
  <name>hive.server2.authentication.kerberos.principal</name>
  <value>hive/_HOST@YOUR-REALM.COM</value>
</property>
<property>
  <name>hive.server2.authentication.kerberos.keytab</name>
  <value>/etc/hive/conf/hive.keytab</value>
</property>
```

where:

- `hive.server2.authentication` is a client-facing property that controls the type of authentication HiveServer2 uses for connections to clients. In this case, HiveServer2 uses Kerberos to authenticate incoming clients.
- The `_HOST@YOUR-REALM.COM` value in the example above is the Kerberos principal for the host where HiveServer2 is running. The string `_HOST` in the properties is replaced at run time by the fully qualified domain name (FQDN) of the host machine where the daemon is running. Reverse DNS must be working on all the hosts configured this way. Replace `YOUR-REALM.COM` with the name of the Kerberos realm your Hadoop cluster is in.
- The `/etc/hive/conf/hive.keytab` value in the example above is a keytab file for that principal.

If you configure HiveServer2 to use both Kerberos authentication and secure impersonation, JDBC clients and Beeline can specify an alternate session user. If these clients have proxy user privileges, HiveServer2 impersonates the alternate user instead of the one connecting. The alternate user can be specified by the JDBC connection string `proxyUser=userName`

Configuring JDBC Clients for Kerberos Authentication with HiveServer2 (Using the Apache Hive Driver in Beeline)

JDBC-based clients must include `principal=<hive.server2.authentication.principal>` in the JDBC connection string. For example:

```
String url =
"jdbc:hive2://node1:10000/default;principal=hive/HiveServer2Host@YOUR-REALM.COM"
Connection con = DriverManager.getConnection(url);
```

where `hive` is the principal configured in `hive-site.xml` and `HiveServer2Host` is the host where HiveServer2 is running.

See the database drivers section on the Cloudera downloads web page to download and install the driver.

Using Beeline to Connect to a Secure HiveServer2

Use the following command to start `beeline` and connect to a secure HiveServer2 process. In this example, the HiveServer2 process is running on `localhost` at port `10000`:

```
$ /usr/lib/hive/bin/beeline
beeline> !connect
jdbc:hive2://localhost:10000/default;principal=hive/HiveServer2Host@YOUR-REALM.COM
0: jdbc:hive2://localhost:10000/default>
```

For more information about the Beeline CLI, see Using the Beeline CLI.

For instructions on encrypting communication with the ODBC/JDBC drivers, see Configuring Encrypted Communication Between HiveServer2 and Client Drivers on page 210.

Using LDAP Username/Password Authentication with HiveServer2

As an alternative to Kerberos authentication, you can configure HiveServer2 to use user and password validation backed by LDAP. The client sends a username and password during connection initiation. HiveServer2 validates these credentials using an external LDAP service.

You can enable LDAP Authentication with HiveServer2 using Active Directory or OpenLDAP.

> **Important:** When using LDAP username/password authentication with HiveServer2, you must enable encrypted communication between HiveServer2 and its client drivers to avoid sending cleartext passwords. For instructions, see Configuring Encrypted Communication Between HiveServer2 and Client Drivers on page 210. To avoid sending LDAP credentials over a network in cleartext, see Configuring LDAPS Authentication with HiveServer2 on page 129.

Enabling LDAP Authentication with HiveServer2 using Active Directory

- **For managed clusters, use Cloudera Manager:**
  1. In the Cloudera Manager Admin Console, click **Hive** in the list of components, and then select the **Configuration** tab.
  2. Type "ldap" in the Search text box to locate the LDAP configuration fields.
  3. Check **Enable LDAP Authentication**.
  4. Enter the **LDAP URL** in the format `ldap[s]://<host>:<port>`
  5. Enter the **Active Directory Domain** for your environment.
  6. Click **Save Changes**.

- **For unmanaged clusters, use the command line:**

  Add the following properties to the `hive-site.xml`:

```
<property>
  <name>hive.server2.authentication</name>
  <value>LDAP</value>
</property>
<property>
  <name>hive.server2.authentication.ldap.url</name>
  <value>LDAP_URLL</value>
</property>
<property>
  <name>hive.server2.authentication.ldap.Domain</name>
  <value>AD_DOMAIN_ADDRESS</value>
</property>
```

  Where:

  The *LDAP_URL* value is the access URL for your LDAP server. For example, `ldap[s]://<host>:<port>`

Enabling LDAP Authentication with HiveServer2 using OpenLDAP

To enable LDAP authentication using OpenLDAP, include the following properties in `hive-site.xml`:

```
<property>
  <name>hive.server2.authentication</name>
  <value>LDAP</value>
</property>
<property>
  <name>hive.server2.authentication.ldap.url</name>
  <value>LDAP_URL</value>
</property>
<property>
  <name>hive.server2.authentication.ldap.baseDN</name>
  <value>LDAP_BaseDN</value>
</property>
```

where:

- The `LDAP_URL` value is the access URL for your LDAP server.
- The `LDAP_BaseDN` value is the base LDAP DN for your LDAP server; for example, `ou=People,dc=example,dc=com`.

Configuring JDBC Clients for LDAP Authentication with HiveServer2

The JDBC client requires a connection URL that includes the user name and password in the JDBC connection string. For example:

```
String url ="jdbc:hive2://#<host>:#<port>/#<dbName>; \
          user=<user name>; \
          password=<password>"
          Connection con = DriverManager.getConnection(url);
```

Enabling LDAP Authentication for HiveServer2 in Hue

Enable LDAP authentication with HiveServer2 by setting the following properties under the `[beeswax]` section in `hue.ini`.

`auth_username`	LDAP username of Hue user to be authenticated.
`auth_password`	LDAP password of Hue user to be authenticated.

Hive uses these login details to authenticate to LDAP. The Hive service trusts that Hue has validated the user being impersonated.

Configuring LDAPS Authentication with HiveServer2

HiveServer2 supports LDAP username/password authentication for clients. Clients send LDAP credentials to HiveServer2 which in turn verifies them against the configured LDAP provider, such as OpenLDAP or Microsoft Active Directory. Most implementations now support LDAPS (LDAP over TLS/SSL), an authentication protocol that uses TLS/SSL to encrypt communication between the LDAP service and its client (in this case, HiveServer2) to avoid sending LDAP credentials in cleartext.

To configure the LDAPS service with HiveServer2:

1. Import the LDAP server CA certificate or the server certificate into a truststore on the HiveServer2 host. If you import the CA certificate, HiveServer2 will trust any server with a certificate issued by the LDAP server's CA. If you only import the server certificate, HiveServer2 trusts only that server. See Understanding Keystores and Truststores on page 193 for more details.
2. Make sure the truststore file is readable by the `hive` user.
3. Set the `hive.server2.authentication.ldap.url` configuration property in `hive-site.xml` to the LDAPS URL. For example, `ldaps://sample.myhost.com`.

   > **Note:** The URL scheme should be `ldaps` and *not* `ldap`.

4. If this is a managed cluster, in Cloudera Manager, go to the Hive service and select **Configuration**. Under Category, select **Security**. In the right panel, search for **HiveServer2 TLS/SSL Certificate Trust Store File**, and add the path to the truststore file that you created in step 1.

   If you are using an unmanaged cluster, set the environment variable `HADOOP_OPTS` as follows:

```
HADOOP_OPTS="-Djavax.net.ssl.trustStore=<trustStore-file-path>
-Djavax.net.ssl.trustStorePassword=<trustStore-password>"
```

5. Restart HiveServer2.

Pluggable Authentication

Pluggable authentication allows you to provide a custom authentication provider for HiveServer2.

To enable pluggable authentication:

1. Set the following properties in `/etc/hive/conf/hive-site.xml`:

```
<property>
  <name>hive.server2.authentication</name>
  <value>CUSTOM</value>
  <description>Client authentication types.
  NONE: no authentication check
  LDAP: LDAP/AD based authentication
  KERBEROS: Kerberos/GSSAPI authentication
  CUSTOM: Custom authentication provider
  (Use with property hive.server2.custom.authentication.class)
  </description>
</property>

<property>
  <name>hive.server2.custom.authentication.class</name>
  <value>pluggable-auth-class-name</value>
  <description>
  Custom authentication class. Used when property
  'hive.server2.authentication' is set to 'CUSTOM'. Provided class
  must be a proper implementation of the interface
  org.apache.hive.service.auth.PasswdAuthenticationProvider. HiveServer2
  will call its Authenticate(user, passed) method to authenticate requests.
  The implementation may optionally extend the Hadoop's
  org.apache.hadoop.conf.Configured class to grab Hive's Configuration object.
```

```
  </description>
</property>
```

2. Make the class available in the CLASSPATH of HiveServer2.

Trusted Delegation with HiveServer2

HiveServer2 determines the identity of the connecting user from the authentication subsystem (Kerberos or LDAP). Any new session started for this connection runs on behalf of this connecting user. If the server is configured to proxy the user at the Hadoop level, then all MapReduce jobs and HDFS accesses will be performed with the identity of the connecting user. If Apache Sentry is configured, then this connecting userid can also be used to verify access rights to underlying tables and views.

Users with Hadoop superuser privileges can request an alternate user for the given session. HiveServer2 checks that the connecting user can proxy the requested userid, and if so, runs the new session as the alternate user. For example, the Hadoop superuser hue can request that a connection's session be run as user bob.

Alternate users for new JDBC client connections are specified by adding the `hive.server2.proxy.user=`*`alternate_user_id`* property to the JDBC connection URL. For example, a JDBC connection string that lets user hue run a session as user bob would be as follows:

```
# Login as super user Hue
kinit hue -k -t hue.keytab hue@MY-REALM.COM

# Connect using following JDBC connection string
#
jdbc:hive2://myHost.myOrg.com:10000/default;principal=hive/_HOST@MY-REALM.COM;hive.server2.proxy.user=bob
```

The connecting user must have Hadoop-level proxy privileges over the alternate user.

HiveServer2 Impersonation

> **Important:** This is not the recommended method to implement HiveServer2 authorization. Cloudera recommends you use Sentry to implement this instead.

HiveServer2 impersonation lets users execute queries and access HDFS files as the connected user rather than as the super user. Access policies are applied at the file level using the HDFS permissions specified in ACLs (access control lists). Enabling HiveServer2 impersonation bypasses Sentry from the end-to-end authorization process. Specifically, although Sentry enforces access control policies on tables and views within the Hive warehouse, it does not control access to the HDFS files that underlie the tables. This means that users without Sentry permissions to tables in the warehouse may nonetheless be able to bypass Sentry authorization checks and execute jobs and queries against tables in the warehouse as long as they have permissions on the HDFS files supporting the table.

To configure Sentry correctly, restrict ownership of the Hive warehouse to `hive:hive` and disable Hive impersonation.

To enable impersonation in HiveServer2:

1. Add the following property to the `/etc/hive/conf/hive-site.xml` file and set the value to `true`. (The default value is `false`.)

```
<property>
  <name>hive.server2.enable.impersonation</name>
  <description>Enable user impersonation for HiveServer2</description>
  <value>true</value>
</property>
```

2. In HDFS or MapReduce configurations, add the following property to the `core-site.xml` file:

```
<property>
  <name>hadoop.proxyuser.hive.hosts</name>
  <value>*</value>
</property>
<property>
```

```
  <name>hadoop.proxyuser.hive.groups</name>
  <value>*</value>
</property>
```

See also File System Permissions.

Securing the Hive Metastore

> **Note:** This is not the recommended method to protect the Hive Metastore. Cloudera recommends you use Sentry to implement this instead.

To prevent users from accessing the Hive metastore and the Hive metastore database using any method other than through HiveServer2, the following actions are recommended:

- Add a firewall rule on the metastore service host to allow access to the metastore port only from the HiveServer2 host. You can do this using iptables.
- Grant access to the metastore database only from the metastore service host. This is specified for MySQL as:

```
GRANT ALL PRIVILEGES ON metastore.* TO 'hive'@'metastorehost';
```

  where `metastorehost` is the host where the metastore service is running.

- Make sure users who are not admins cannot log on to the host on which HiveServer2 runs.

Disabling the Hive Security Configuration

Hive's security related metadata is stored in the configuration file `hive-site.xml`. The following sections describe how to disable security for the Hive service.

Disable Client/Server Authentication

To disable client/server authentication, set `hive.server2.authentication` to `NONE`. For example,

```
<property>
  <name>hive.server2.authentication</name>
  <value>NONE</value>
  <description>
    Client authentication types.
       NONE: no authentication check
       LDAP: LDAP/AD based authentication
       KERBEROS: Kerberos/GSSAPI authentication
       CUSTOM: Custom authentication provider
               (Use with property hive.server2.custom.authentication.class)
  </description>
</property>
```

Disable Hive Metastore security

To disable Hive Metastore security, perform the following steps:

- Set the `hive.metastore.sasl.enabled` property to `false` in all configurations, the metastore service side as well as for all clients of the metastore. For example, these might include HiveServer2, Impala, Pig and so on.
- Remove or comment the following parameters in `hive-site.xml` for the metastore service. Note that this is a server-only change.
  - `hive.metastore.kerberos.keytab.file`
  - `hive.metastore.kerberos.principal`

Disable Underlying Hadoop Security

If you also want to disable the underlying Hadoop security, remove or comment out the following parameters in `hive-site.xml`.

- `hive.server2.authentication.kerberos.keytab`
- `hive.server2.authentication.kerberos.principal`

Hive Metastore Server Security Configuration

> **Important:**
>
> This section describes how to configure security for the Hive metastore server. If you are using HiveServer2, see HiveServer2 Security Configuration.

Here is a summary of Hive metastore server security in CDH 5:

- No additional configuration is required to run Hive on top of a security-enabled Hadoop cluster in standalone mode using a local or embedded metastore.
- HiveServer does not support Kerberos authentication for clients. While it is possible to run HiveServer with a secured Hadoop cluster, doing so creates a security hole since HiveServer does not authenticate the Thrift clients that connect to it. Instead, you can use HiveServer2 HiveServer2 Security Configuration.
- The Hive metastore server supports Kerberos authentication for Thrift clients. For example, you can configure a standalone Hive metastore server instance to force clients to authenticate with Kerberos by setting the following properties in the `hive-site.xml` configuration file used by the metastore server:

```
<property>
  <name>hive.metastore.sasl.enabled</name>
  <value>true</value>
  <description>If true, the metastore thrift interface will be secured with SASL. Clients
 must authenticate with Kerberos.</description>
</property>

<property>
  <name>hive.metastore.kerberos.keytab.file</name>
  <value>/etc/hive/conf/hive.keytab</value>
  <description>The path to the Kerberos Keytab file containing the metastore thrift
server's service principal.</description>
</property>

<property>
  <name>hive.metastore.kerberos.principal</name>
  <value>hive/_HOST@YOUR-REALM.COM</value>
  <description>The service principal for the metastore thrift server. The special string
 _HOST will be replaced automatically with the correct host name.</description>
</property>
```

> **Note:**
>
> The values shown above for the `hive.metastore.kerberos.keytab.file` and `hive.metastore.kerberos.principal` properties are examples which you will need to replace with the appropriate values for your cluster. Also note that the Hive keytab file should have its access permissions set to `600` and be owned by the same account that is used to run the Metastore server, which is the `hive` user by default.

- Requests to access the metadata are fulfilled by the Hive metastore impersonating the requesting user. This includes read access to the list of databases, tables, properties of each table such as their HDFS location and file type. You can restrict access to the Hive metastore service by allowing it to impersonate only a subset of Kerberos users. This can be done by setting the `hadoop.proxyuser.hive.groups` property in `core-site.xml` on the Hive metastore host.

  For example, if you want to give the `hive` user permission to impersonate members of groups `hive` and `user1`:

```
<property>
<name>hadoop.proxyuser.hive.groups</name>
```

```
<value>hive,user1</value>
</property>
```

In this example, the Hive metastore can impersonate users belonging to *only* the `hive` and `user1` groups. Connection requests from users not belonging to these groups will be rejected.

Using Hive to Run Queries on a Secure HBase Server

To use Hive to run queries on a secure HBase Server, you must set the following `HIVE_OPTS` environment variable:

```
env HIVE_OPTS="-hiveconf hbase.security.authentication=kerberos -hiveconf
hbase.master.kerberos.principal=hbase/_HOST@YOUR-REALM.COM -hiveconf
hbase.regionserver.kerberos.principal=hbase/_HOST@YOUR-REALM.COM -hiveconf
hbase.zookeeper.quorum=zookeeper1,zookeeper2,zookeeper3" hive
```

where:

- You replace `YOUR-REALM` with the name of your Kerberos realm
- You replace `zookeeper1`, `zookeeper2`, `zookeeper3` with the names of your ZooKeeper servers. The `hbase.zookeeper.quorum` property is configured in the `hbase-site.xml` file.
- The special string `_HOST` is replaced at run-time by the fully qualified domain name of the host machine where the HBase Master or RegionServer is running. This requires that reverse DNS is properly working on all the hosts configured this way.

In the following, `_HOST` is the name of the host where the HBase Master is running:

```
-hiveconf hbase.master.kerberos.principal=hbase/_HOST@YOUR-REALM.COM
```

In the following, `_HOST` is the hostname of the HBase RegionServer that the application is connecting to:

```
-hiveconf hbase.regionserver.kerberos.principal=hbase/_HOST@YOUR-REALM.COM
```

Note:

You can also set the `HIVE_OPTS` environment variable in your shell profile.

HttpFS Authentication

This section describes how to configure HttpFS CDH 5 with Kerberos security on a Hadoop cluster:

- Configuring the HttpFS Server to Support Kerberos Security on page 134
- Using curl to access an URL Protected by Kerberos HTTP SPNEGO on page 135

For more information about HttpFS, see https://archive.cloudera.com/cdh5/cdh/5/hadoop/hadoop-hdfs-httpfs/index.html.

Important:

To enable HttpFS to work with Kerberos security on your Hadoop cluster, make sure you perform the installation and configuration steps in Configuring Hadoop Security in CDH 5.

Important:

If the NameNode, Secondary NameNode, DataNode, JobTracker, TaskTrackers, ResourceManager, NodeManagers, HttpFS, or Oozie services are configured to use Kerberos HTTP SPNEGO authentication, and two or more of these services are running on the same host, then all of the running services must use the same HTTP principal and keytab file used for their HTTP endpoints.

Configuring the HttpFS Server to Support Kerberos Security

1. Create an HttpFS service user principal that is used to authenticate with the Hadoop cluster. The syntax of the principal is: `httpfs/<fully.qualified.domain.name>@<YOUR-REALM>` where: `fully.qualified.domain.name` is the host where the HttpFS server is running `YOUR-REALM` is the name of your Kerberos realm

```
kadmin: addprinc -randkey httpfs/fully.qualified.domain.name@YOUR-REALM.COM
```

2. Create a HTTP service user principal that is used to authenticate user requests coming to the HttpFS HTTP web-services. The syntax of the principal is: `HTTP/<fully.qualified.domain.name>@<YOUR-REALM>` where: `'fully.qualified.domain.name'` is the host where the HttpFS server is running `YOUR-REALM` is the name of your Kerberos realm

```
kadmin: addprinc -randkey HTTP/fully.qualified.domain.name@YOUR-REALM.COM
```

> **Important:**
>
> The HTTP/ component of the HTTP service user principal must be upper case as shown in the syntax and example above.

3. Create keytab files with both principals.

```
$ kadmin
kadmin: xst -k httpfs.keytab httpfs/fully.qualified.domain.name
kadmin: xst -k http.keytab HTTP/fully.qualified.domain.name
```

4. Merge the two keytab files into a single keytab file:

```
$ ktutil
ktutil: rkt httpfs.keytab
ktutil: rkt http.keytab
ktutil: wkt httpfs-http.keytab
```

5. Test that credentials in the merged keytab file work. For example:

```
$ klist -e -k -t httpfs-http.keytab
```

6. Copy the `httpfs-http.keytab` file to the HttpFS configuration directory. The owner of the `httpfs-http.keytab` file should be the `httpfs` user and the file should have owner-only read permissions.
7. Edit the HttpFS server `httpfs-site.xml` configuration file in the HttpFS configuration directory by setting the following properties:

Property	Value
httpfs.authentication.type	kerberos
httpfs.hadoop.authentication.type	kerberos
httpfs.authentication.kerberos.principal	HTTP/<HTTPFS-HOSTNAME>@<YOUR-REALM.COM>
httpfs.authentication.kerberos.keytab	/etc/hadoop-httpfs/conf/httpfs-http.keytab
httpfs.hadoop.authentication.kerberos.principal	httpfs/<HTTPFS-HOSTNAME>@<YOUR-REALM.COM>
httpfs.hadoop.authentication.kerberos.keytab	/etc/hadoop-httpfs/conf/httpfs-http.keytab
httpfs.authentication.kerberos.name.rules	Use the value configured for 'hadoop.security.auth_to_local' in 'core-site.xml'

> **Important:**
> You must restart the HttpFS server to have the configuration changes take effect.

Using curl to access an URL Protected by Kerberos HTTP SPNEGO

> **Important:**
> Your version of `curl` must support GSS and be capable of running `curl -V`.

To configure curl to access an URL protected by Kerberos HTTP SPNEGO:

1. Run `curl -V`:

```
$ curl -V
curl 7.19.7 (universal-apple-darwin10.0) libcurl/7.19.7 OpenSSL/0.9.8l
zlib/1.2.3
Protocols: tftp ftp telnet dict ldap http file https ftps
Features: GSS-Negotiate IPv6 Largefile NTLM SSL libz
```

2. Login to the KDC using `kinit`.

```
$ kinit
Please enter the password for tucu@LOCALHOST:
```

3. Use `curl` to fetch the protected URL:

```
$ curl --cacert
             /path/to/truststore.pem --negotiate -u : -b ~/cookiejar.txt -c
~/cookiejar.txt https://localhost:14000/webhdfs/v1/?op=liststatus
```

where:

- The `--cacert` option is required if you are using TLS/SSL certificates that curl does not recognize by default.
- The `--negotiate` option enables SPNEGO in `curl`.
- The `-u :` option is required but the username is ignored (the principal that has been specified for `kinit` is used).
- The `-b` and `-c` options are used to store and send HTTP cookies.
- Cloudera does not recommend using the `-k` or `--insecure` option as it turns off curl's ability to verify the certificate.

Hue Authentication

This page describes properties in the Hue configuration file, `hue.ini`, that support authentication and Hue security in general.

For information on configuring Hue with Kerberos encrypting session communication, and enabling single sign-on with SAML, see:

- Configuring Kerberos Authentication for Hue on page 137
- Authenticate Hue Users with LDAP and Synchronize Hue with LDAP Server
- Authenticate Hue Users with SAML
- Authorize Hue User Groups with Sentry

Enabling LDAP Authentication with HiveServer2 and Impala

LDAP authentication with HiveServer2 and Impala can be enabled by setting the following properties under their respective sections in `hue.ini`, `[beeswax]` and `[impala]`.

`auth_username`	LDAP username of Hue user to be authenticated.
`auth_password`	LDAP password of Hue user to be authenticated.

These login details are only used by Impala and Hive to authenticate to LDAP. The Impala and Hive services trust Hue to have already validated the user being impersonated, rather than simply passing on the credentials.

Securing Sessions

When a session expires, the screen blurs and the user is automatically logged out of the Hue Web UI. Logging on returns the user to same location.

Session Timeout

User sessions are controlled with the `ttl` (time-to-live) property under `[desktop]`> `[[session]]` in `hue.ini`. After n seconds, the session expires whether active or not.

`ttl`	The cookie with the users session ID expires after n seconds. Default: `ttl=1209600` which is 60*60*24*14 seconds or 2 weeks

Idle Session Timeout

Idle sessions are controlled with the idle_session_timeout property under `[desktop]` > `[[auth]]` in `hue.ini`. Sessions that are idle for n seconds, expire. You can disable this property by setting it to a negative value.

`idle_session_timeout`	User session IDs expire after idle for n seconds. A negative value means idle sessions do not expire. `idle_session_timeout=900` means that sessions expire after being idle for 15 minutes `idle_session_timeout=-1` means that idle sessions do not expire (until `ttl`)

Secure Login

Login properties are set in `hue.ini` under `[desktop]` > `[[auth]]`. They are based on django-axes 1.5.0.

`change_default_password`	If true, users must change password on first login. Must enable `backend=desktop.auth.backend.AllowFirstUserDjangoBackend`
`expires_after`	User accounts are disabled n seconds after logout. If negative, user sessions never expire.
`expire_superusers`	Apply `expires_after` to superusers.
`login_cooloff_time`	Failed logins are forgotten after n seconds.
`login_failure_limit`	Number of login attempts allowed before a record is created for failed logins.
`login_lock_out_at_failure`	If true, lock out IP after exceeding `login_failure_limit`. If `login_lock_out_by_combination_user_and_ip=true`, lock out IP and user. If `login_lock_out_use_user_agent=true`, also lock out user agent.
`login_lock_out_by_combination_user_and_ip`	If true, lock out IP *and* user.
`login_lock_out_use_user_agent`	If true, lock out user agent (such as a browser).

Secure Cookies

Secure session cookies can be enabled by specifying the `secure` configuration property under the `[desktop]> [[session]]` section in `hue.ini`. Additionally, you can set the `http_only` flag for cookies containing users' session IDs.

`secure`	The cookie with the user session ID is secure. Should only be enabled with HTTPS. Default: `false`
`http_only`	The cookie with the user session ID uses the HTTP only flag. Default: `true` *If the HttpOnly flag is included in the HTTP response header, the cookie cannot be accessed through a client side script.*
`expire_at_browser_close`	Use session-length cookies. Logs out the user when the browser window is closed. Default: `false`

Allowed HTTP Methods

You can specify the HTTP request methods that the server should respond to using the `http_allowed_methods` property under the `[desktop]` section in `hue.ini`.

`http_allowed_methods`	Default: `options,get,head,post,put,delete,connect`

Restricting the Cipher List

Cipher list support with HTTPS can be restricted by specifying the `ssl_cipher_list` configuration property under the `[desktop]` section in `hue.ini`.

`ssl_cipher_list`	Default: `!aNULL:!eNULL:!LOW:!EXPORT:!SSLv2`

URL Redirect Whitelist

Restrict the domains or pages to which Hue can redirect users. The `redirect_whitelist` property can be found under the `[desktop]` section in `hue.ini`.

`redirect_whitelist`	For example, to restrict users to your local domain and FQDN, the following value can be used: `^\/.*$,^http:\/\/www.mydomain.com\/.*$`

Oozie Permissions

Access to the Oozie dashboard and editor can be individually controlled in the Hue Web UI under **User Admin** > **Groups**.

Groups Property in UI	Description
oozie.dashboard_jobs_access	Enable Oozie Dashboard read-only access for all jobs. Default: `true`
oozie.disable_editor_access	Disable Oozie Editor access. Default: `false`

Configuring Kerberos Authentication for Hue

To configure the Hue server to use Kerberos for authentication:

1. Create a Hue user principal in the same realm as the cluster of the form:

```
kadmin: addprinc -randkey hue/hue.server.fully.qualified.domain.name@YOUR-REALM.COM
```

   where: hue is the principal the Hue server is running as, `hue.server.fully.qualified.domain.name` is the fully qualified domain name (FQDN) of your Hue server, `YOUR-REALM.COM` is the name of the Kerberos realm your Hadoop cluster is in

2. Create a keytab file for the Hue principal using the same procedure that you used to create the keytab for the `hdfs` or `mapred` principal for a specific host. You should name this file `hue.keytab` and put this keytab file in the directory `/etc/hue` on the machine running the Hue server. Like all keytab files, this file should have the most limited set of permissions possible. It should be owned by the user running the hue server (usually hue) and should have the permission `400`.
3. To test that the keytab file was created properly, try to obtain Kerberos credentials as the Hue principal using only the keytab file. Substitute your FQDN and realm in the following command:

```
$ kinit -k -t /etc/hue/hue.keytab
hue/hue.server.fully.qualified.domain.name@YOUR-REALM.COM
```

4. In the `/etc/hue/hue.ini` configuration file, add the following lines in the sections shown. Replace the `kinit_path` value, `/usr/kerberos/bin/kinit`, shown below with the correct path on the user's system.

```
[desktop]

 [[kerberos]]
 # Path to Hue's Kerberos keytab file
 hue_keytab=/etc/hue/hue.keytab
 # Kerberos principal name for Hue
 hue_principal=hue/FQDN@REALM
 # add kinit path for non root users
 kinit_path=/usr/kerberos/bin/kinit

[beeswax]
 # If Kerberos security is enabled, use fully qualified domain name (FQDN)
   ## hive_server_host=<FQDN of Hive Server>
 # Hive configuration directory, where hive-site.xml is located
   ## hive_conf_dir=/etc/hive/conf

[impala]
   ## server_host=localhost
 # The following property is required when impalad and Hue
 # are not running on the same host
   ## impala_principal=impala/impalad.hostname.domainname.com

[search]
 # URL of the Solr Server
   ## solr_url=http://localhost:8983/solr/
 # Requires FQDN in solr_url if enabled
   ## security_enabled=false

[hadoop]

 [[hdfs_clusters]]

 [[[default]]]
 # Enter the host and port on which you are running the Hadoop NameNode
 namenode_host=FQDN
 hdfs_port=8020
 http_port=50070
 security_enabled=true

 # Thrift plugin port for the name node
   ## thrift_port=10090

 # Configuration for YARN (MR2)
 # ------------------------------------------------------------------------
 [[yarn_clusters]]

 [[[default]]]
```

```
 # Enter the host on which you are running the ResourceManager
   ## resourcemanager_host=localhost
 # Change this if your YARN cluster is Kerberos-secured
   ## security_enabled=false

 # Thrift plug-in port for the JobTracker
   ## thrift_port=9290

[liboozie]
 # The URL where the Oozie service runs on. This is required in order for users to submit
 jobs.
   ## oozie_url=http://localhost:11000/oozie
 # Requires FQDN in oozie_url if enabled
   ## security_enabled=false
```

Important:

In the `/etc/hue/hue.ini` file, verify that:

- The `jobtracker_host` property is set to the fully-qualified domain name (FQDN) of the host running the JobTracker.
- The `fs.defaultfs` property under each [[hdfs_clusters]] section contains the FQDN of the file system access point, which is typically the NameNode.
- The `hive_conf_dir` property under the `[beeswax]` section points to a directory containing a valid `hive-site.xml` (either the original or a synced copy).
- The FQDN specified for HiveServer2 is the same as the FQDN specified for the `hue_principal` configuration property.

Also note that HiveServer2 currently does not support TLS/SSL when using Kerberos.

5. In the `/etc/hadoop/conf/core-site.xml` configuration file on each node in the cluster, add the following lines:

```
<!-- Hue security configuration -->
<property>
  <name>hue.kerberos.principal.shortname</name>
  <value>hue</value>
</property>
<property>
  <name>hadoop.proxyuser.hue.groups</name>
  <value>*</value> <!-- A group which all users of Hue belong to, or the wildcard value
 "*" -->
</property>
<property>
  <name>hadoop.proxyuser.hue.hosts</name>
  <value>hue.server.fully.qualified.domain.name</value>
</property>
```

Important:

Change the `/etc/hadoop/conf/core-site.xml` configuration file on *all* nodes in the cluster.

6. For Hue setups that include HttpFS for communication to Hadoop, add the following properties to `httpfs-site.xml`:

```
<property>
  <name>httpfs.proxyuser.hue.hosts</name>
  <value>fully.qualified.domain.name</value>
</property>
<property>
  <name>httpfs.proxyuser.hue.groups</name>
  <value>*</value>
</property>
```

7. Add the following properties to the Oozie server `oozie-site.xml` configuration file in the Oozie configuration directory:

```
<property>
   <name>oozie.service.ProxyUserService.proxyuser.hue.hosts</name>
   <value>*</value>
</property>
<property>
   <name>oozie.service.ProxyUserService.proxyuser.hue.groups</name>
   <value>*</value>
</property>
```

8. Restart the JobTracker to load the changes from the `core-site.xml` file.

```
$ sudo service hadoop-0.20-mapreduce-jobtracker restart
```

9. Restart Oozie to load the changes from the `oozie-site.xml` file.

```
$ sudo service oozie restart
```

10. Restart the NameNode, JobTracker, and all DataNodes to load the changes from the `core-site.xml` file.

```
$ sudo service hadoop-0.20-(namenode|jobtracker|datanode) restart
```

Enable Hue to Use Kerberos for Authentication

Minimum Required Role: Cluster Administrator (also provided by **Full Administrator**)

For Hue to work properly with a Cloudera cluster that uses Kerberos for authentication, the **Kerberos Ticket Renewer** must be added to the Hue service. Use the Cloudera Manager Admin Console to add the Kerberos Ticket Renewer role to each host with a Hue Server role instance.

The Hue Kerberos Ticket Renewer renews only those tickets created for the Hue service principal, hue/*hostname@REALM-NAME*. The Hue principal impersonates other users for applications within Hue such as the Job Browser, File Browser and so on. Other services, such as HDFS and MapReduce, do not use the Hue Kerberos Ticket Renewer but rather handle ticket renewal as needed using their own mechanisms.

Adding a Kerberos Ticket Renewer role instance in Cloudera Manager:

1. Go to the **Hue** service.
2. Click the **Instances** tab.
3. Click the **Add Role Instances** button.
4. Assign the **Kerberos Ticket Renewer** role instance to the same host as the Hue server.

   When the wizard status is **Finished**, the Kerberos Ticket Renewer role instance is configured. The Hue service now works with the secure Hadoop cluster.

5. Repeat these steps for each Hue Server role.

Troubleshooting the Kerberos Ticket Renewer:

If the Hue Kerberos Ticket Renewer does not start, check the configuration of your Kerberos Key Distribution Center (KDC). Look at the ticket renewal property, `maxrenewlife`, to ensure that the principals, `hue/<hostname>` and `krbtgt`, are renewable. If these principals are not renewable, run the following commands on the KDC to enable them:

```
kadmin.local: modprinc -maxrenewlife 90day krbtgt/YOUR_REALM.COM
kadmin.local: modprinc -maxrenewlife 90day +allow_renewable hue/<hostname>@YOUR-REALM.COM
```

Impala Authentication

Authentication is the mechanism to ensure that only specified hosts and users can connect to Impala. It also verifies that when clients connect to Impala, they are connected to a legitimate server. This feature prevents spoofing such as

impersonation (setting up a phony client system with the same account and group names as a legitimate user) and ***man-in-the-middle attacks*** (intercepting application requests before they reach Impala and eavesdropping on sensitive information in the requests or the results).

Impala supports authentication using either Kerberos or LDAP.

> **Note:** Regardless of the authentication mechanism used, Impala always creates HDFS directories and data files owned by the same user (typically `impala`). To implement user-level access to different databases, tables, columns, partitions, and so on, use the Sentry authorization feature, as explained in Enabling Sentry Authorization for Impala.

Once you are finished setting up authentication, move on to authorization, which involves specifying what databases, tables, HDFS directories, and so on can be accessed by particular users when they connect through Impala. See Enabling Sentry Authorization for Impala for details.

Enabling Kerberos Authentication for Impala

Impala supports Kerberos authentication. For background information on enabling Kerberos authentication, see the topic on Configuring Hadoop Security in the CDH 5 Security Guide.

When using Impala in a managed environment, Cloudera Manager automatically completes Kerberos configuration. In an unmanaged environment, create a Kerberos principal for each host running `impalad` or `statestored`. Cloudera recommends using a consistent format, such as `impala/_HOST@Your-Realm`, but you can use any three-part Kerberos server principal.

In Impala 2.0 and later, `user()` returns the full Kerberos principal string, such as `user@example.com`, in a Kerberized environment.

> **Note:** Regardless of the authentication mechanism used, Impala always creates HDFS directories and data files owned by the same user (typically `impala`). To implement user-level access to different databases, tables, columns, partitions, and so on, use the Sentry authorization feature, as explained in Enabling Sentry Authorization for Impala.

An alternative form of authentication you can use is LDAP, described in Enabling LDAP Authentication for Impala on page 144.

Requirements for Using Impala with Kerberos

On version 5 of Red Hat Enterprise Linux and comparable distributions, some additional setup is needed for the `impala-shell` interpreter to connect to a Kerberos-enabled Impala cluster:

```
sudo yum install python-devel openssl-devel python-pip
sudo pip-python install ssl
```

> **Important:**
>
> - If you plan to use Impala in your cluster, you must configure your KDC to allow tickets to be renewed, and you must configure `krb5.conf` to request renewable tickets. Typically, you can do this by adding the `max_renewable_life` setting to your realm in `kdc.conf`, and by adding the `renew_lifetime` parameter to the `libdefaults` section of `krb5.conf`.
>
> For more information about renewable tickets, see the Kerberos documentation.
>
> - The Impala Web UI does not support Kerberos authentication.
>
> - You cannot use the Impala resource management feature on a cluster that has Kerberos authentication enabled.

Start all `impalad` and `statestored` daemons with the `--principal` and `--keytab-file` flags set to the principal and full path name of the `keytab` file containing the credentials for the principal.

Impala supports the Cloudera ODBC driver and the Kerberos interface provided. To use Kerberos through the ODBC driver, the host type must be set depending on the level of the ODBD driver:

- `SecImpala` for the ODBC 1.0 driver.
- `SecBeeswax` for the ODBC 1.2 driver.
- Blank for the ODBC 2.0 driver or higher, when connecting to a secure cluster.
- `HS2NoSasl` for the ODBC 2.0 driver or higher, when connecting to a non-secure cluster.

To enable Kerberos in the Impala shell, start the `impala-shell` command using the `-k` flag.

To enable Impala to work with Kerberos security on your Hadoop cluster, make sure you perform the installation and configuration steps in Authentication in the CDH 5 Security Guide.

Configuring Impala to Support Kerberos Security

Enabling Kerberos authentication for Impala involves steps that can be summarized as follows:

- Creating service principals for Impala and the HTTP service. Principal names take the form: *`serviceName/fully.qualified.domain.name@KERBEROS.REALM`*
- Creating, merging, and distributing key tab files for these principals.
- Editing `/etc/default/impala` (in cluster not managed by Cloudera Manager), or editing the **Security** settings in the Cloudera Manager interface, to accommodate Kerberos authentication.

Enabling Kerberos for Impala

1. Create an Impala service principal, specifying the name of the OS user that the Impala daemons run under, the fully qualified domain name of each node running `impalad`, and the realm name. For example:

```
$ kadmin
kadmin: addprinc -requires_preauth -randkey
impala/impala_host.example.com@TEST.EXAMPLE.COM
```

2. Create an HTTP service principal. For example:

```
kadmin: addprinc -randkey HTTP/impala_host.example.com@TEST.EXAMPLE.COM
```

> **Note:** The HTTP component of the service principal must be uppercase as shown in the preceding example.

3. Create `keytab` files with both principals. For example:

```
kadmin: xst -k impala.keytab impala/impala_host.example.com
kadmin: xst -k http.keytab HTTP/impala_host.example.com
kadmin: quit
```

4. Use `ktutil` to read the contents of the two keytab files and then write those contents to a new file. For example:

```
$ ktutil
ktutil: rkt impala.keytab
ktutil: rkt http.keytab
ktutil: wkt impala-http.keytab
ktutil: quit
```

5. (Optional) Test that credentials in the merged keytab file are valid, and that the “renew until” date is in the future. For example:

```
$ klist -e -k -t impala-http.keytab
```

6. Copy the `impala-http.keytab` file to the Impala configuration directory. Change the permissions to be only read for the file owner and change the file owner to the `impala` user. By default, the Impala user and group are both named `impala`. For example:

```
$ cp impala-http.keytab /etc/impala/conf
$ cd /etc/impala/conf
$ chmod 400 impala-http.keytab
$ chown impala:impala impala-http.keytab
```

7. Add Kerberos options to the Impala defaults file, `/etc/default/impala`. Add the options for both the `impalad` and `statestored` daemons, using the `IMPALA_SERVER_ARGS` and `IMPALA_STATE_STORE_ARGS` variables. For example, you might add:

```
-kerberos_reinit_interval=60
-principal=impala_1/impala_host.example.com@TEST.EXAMPLE.COM
-keytab_file=/var/run/cloudera-scm-agent/process/3212-impala-IMPALAD/impala.keytab
```

For more information on changing the Impala defaults specified in `/etc/default/impala`, see Modifying Impala Startup Options.

Note: Restart `impalad` and `statestored` for these configuration changes to take effect.

Enabling Kerberos for Impala with a Proxy Server

A common configuration for Impala with High Availability is to use a proxy server to submit requests to the actual `impalad` daemons on different hosts in the cluster. This configuration avoids connection problems in case of machine failure, because the proxy server can route new requests through one of the remaining hosts in the cluster. This configuration also helps with load balancing, because the additional overhead of being the “coordinator node” for each query is spread across multiple hosts.

Although you can set up a proxy server with or without Kerberos authentication, typically users set up a secure Kerberized configuration. For information about setting up a proxy server for Impala, including Kerberos-specific steps, see Using Impala through a Proxy for High Availability.

Enabling Impala Delegation for Kerberos Users

See Configuring Impala Delegation for Hue and BI Tools on page 147 for details about the delegation feature that lets certain users submit queries using the credentials of other users.

Using TLS/SSL with Business Intelligence Tools

You can use Kerberos authentication, TLS/SSL encryption, or both to secure connections from JDBC and ODBC applications to Impala. See Configuring Impala to Work with JDBC and Configuring Impala to Work with ODBC for details.

Prior to CDH 5.7 / Impala 2.5, the Hive JDBC driver did not support connections that use both Kerberos authentication and SSL encryption. If your cluster is running an older release that has this restriction, to use both of these security features with Impala through a JDBC application, use the Cloudera JDBC Connector as the JDBC driver.

Enabling Access to Internal Impala APIs for Kerberos Users

For applications that need direct access to Impala APIs, without going through the HiveServer2 or Beeswax interfaces, you can specify a list of Kerberos users who are allowed to call those APIs. By default, the `impala` and `hdfs` users are the only ones authorized for this kind of access. Any users not explicitly authorized through the `internal_principals_whitelist` configuration setting are blocked from accessing the APIs. This setting applies to all the Impala-related daemons, although currently it is primarily used for HDFS to control the behavior of the catalog server.

Mapping Kerberos Principals to Short Names for Impala

In and higher, Impala recognizes the `auth_to_local` setting, specified through the HDFS configuration setting `hadoop.security.auth_to_local` or the Cloudera Manager setting **Additional Rules to Map Kerberos Principals**

to Short Names. This feature is disabled by default, to avoid an unexpected change in security-related behavior. To enable it:

- For clusters not managed by Cloudera Manager, specify `--load_auth_to_local_rules=true` in the `impalad` and `catalogd`configuration settings.
- For clusters managed by Cloudera Manager, select the **Use HDFS Rules to Map Kerberos Principals to Short Names** checkbox to enable the service-wide `load_auth_to_local_rules` configuration setting. Then restart the Impala service.

See Using Auth-to-Local Rules to Isolate Cluster Users for general information about this feature.

Kerberos-Related Memory Overhead for Large Clusters

On a kerberized cluster with high memory utilization, `kinit` commands executed after every `'kerberos_reinit_interval'` may cause out-of-memory errors, because executing the command involves a fork of the Impala process. The error looks similar to the following:

```
Failed to obtain Kerberos ticket for principal: <varname>principal_details</varname>
Failed to execute shell cmd: 'kinit -k -t <varname>keytab_details</varname>',
error was: Error(12): Cannot allocate memory
```

The following command changes the `vm.overcommit_memory` setting immediately on a running host. However, this setting is reset when the host is restarted.

```
echo 1 > /proc/sys/vm/overcommit_memory
```

To change the setting in a persistent way, add the following line to the `/etc/sysctl.conf` file:

```
vm.overcommit_memory=1
```

Then run `sysctl -p`. No reboot is needed.

Enabling LDAP Authentication for Impala

Authentication is the process of allowing only specified named users to access the server (in this case, the Impala server). This feature is crucial for any production deployment, to prevent misuse, tampering, or excessive load on the server. Impala uses LDAP for authentication, verifying the credentials of each user who connects through `impala-shell`, Hue, a Business Intelligence tool, JDBC or ODBC application, etc.

> **Note:** Regardless of the authentication mechanism used, Impala always creates HDFS directories and data files owned by the same user (typically `impala`). To implement user-level access to different databases, tables, columns, partitions, and so on, use the Sentry authorization feature, as explained in Enabling Sentry Authorization for Impala.

An alternative form of authentication you can use is Kerberos, described in Enabling Kerberos Authentication for Impala on page 141.

Requirements for Using Impala with LDAP

Authentication against LDAP servers is available in Impala 1.2.2 and higher. Impala 1.4.0 added the support for secure LDAP authentication through SSL and TLS.

The Impala LDAP support lets you use Impala with systems such as Active Directory that use LDAP behind the scenes.

Consideration for Connections Between Impala Components

Only the connections between clients and Impala can be authenticated by LDAP.

You must use the Kerberos authentication mechanism for connections between internal Impala components, such as between the `impalad`, `statestored`, and `catalogd` daemons. See Enabling Kerberos Authentication for Impala on page 141 on how to set up Kerberos for Impala.

Enabling LDAP in Command Line Interface

To enable LDAP authentication via a command line interface, start the `impalad` with the following startup options for:

`--enable_ldap_auth`

Enables LDAP-based authentication between the client and Impala.

`--ldap_uri`

Sets the URI of the LDAP server to use. Typically, the URI is prefixed with `ldap://`. You can specify secure SSL-based LDAP transport by using the prefix `ldaps://`. The URI can optionally specify the port, for example: `ldap://ldap_server.example.com:389` or `ldaps://ldap_server.example.com:636`. (389 and 636 are the default ports for non-SSL and SSL LDAP connections, respectively.)

Support for Custom Bind Strings

When Impala connects to LDAP it issues a bind call to the LDAP server to authenticate as the connected user. Impala clients, including the Impala shell, provide the short name of the user to Impala. This is necessary so that Impala can use Sentry for role-based access, which uses short names.

However, LDAP servers often require more complex, structured usernames for authentication. Impala supports three ways of transforming the short name (for example, `'henry'`) to a more complicated string. If necessary, specify one of the following configuration options when starting the `impalad` daemon.

`--ldap_domain`

Replaces the username with a string *`username@ldap_domain`*.

`--ldap_baseDN`

Replaces the username with a “distinguished name” (DN) of the form: `uid=`*`userid`*`,ldap_baseDN`. (This is equivalent to a Hive option).

`--ldap_bind_pattern`

This is the most general option, and replaces the username with the string *ldap_bind_pattern* where all instances of the string `#UID` are replaced with *userid*. For example, an `ldap_bind_pattern` of `"user=#UID,OU=foo,CN=bar"` with a username of `henry` will construct a bind name of `"user=henry,OU=foo,CN=bar"`.

The above options are mutually exclusive, and Impala does not start if more than one of these options are specified.

Secure LDAP Connections

To avoid sending credentials over the wire in cleartext, you must configure a secure connection between both the client and Impala, and between Impala and the LDAP server. The secure connection could use SSL or TLS.

Secure LDAP connections through SSL:

For SSL-enabled LDAP connections, specify a prefix of `ldaps://` instead of `ldap://`. Also, the default port for SSL-enabled LDAP connections is 636 instead of 389.

Secure LDAP connections through TLS:

TLS, the successor to the SSL protocol, is supported by most modern LDAP servers. Unlike SSL connections, TLS connections can be made on the same server port as non-TLS connections. To secure all connections using TLS, specify the following flags as startup options to the `impalad` daemon:

`--ldap_tls`

Tells Impala to start a TLS connection to the LDAP server, and to fail authentication if it cannot be done.

`--ldap_ca_certificate="/path/to/certificate/pem"`

Specifies the location of the certificate in standard `.PEM` format. Store this certificate on the local filesystem, in a location that only the `impala` user and other trusted users can read.

Enabling LDAP in Cloudera Manager

To enable LDAP authentication in Cloudera Manager:

1. In the Impala service, click the **Configuration** tab.
2. In the search box, type **ldap**.
3. Specify the values for the option fields. Each field lists the corresponding Impala startup flag. See the sections above for the corresponding flag if you need more information on a particular field.
4. Click **Save Changes**.
5. Restart the Impala service.

LDAP Authentication for impala-shell

To connect to Impala using LDAP authentication, you specify command-line options to the `impala-shell` command interpreter and enter the password when prompted.

`-l`

Enables LDAP authentication.

`-u`

Sets the user. Per Active Directory, the user is the short username, not the full LDAP distinguished name. If your LDAP settings include a search base, use the `--ldap_bind_pattern` on the `impalad` daemon to translate the short user name from `impala-shell` automatically to the fully qualified name.

`impala-shell` automatically prompts for the password.

See Configuring Impala to Work with JDBC for the format to use with the JDBC connection string for servers using LDAP authentication.

Enabling LDAP for Impala in Hue

1. Go to the Hue service.
2. Click the **Configuration** tab.
3. Select **Scope** > **Hue Server**.
4. Select **Category** > **Advanced**.
5. Add the following properties to the **Hue Server Advanced Configuration Snippet (Safety Valve) for hue_safety_valve_server.ini** property.

```
[impala]
auth_username=<LDAP username of Hue user to be authenticated>
auth_password=<LDAP password of Hue user to be authenticated>
```

6. Click **Save Changes**.

Enabling Impala Delegation for LDAP Users

See Configuring Impala Delegation for Hue and BI Tools on page 147 for details about the delegation feature that lets certain users submit queries using the credentials of other users.

LDAP Restrictions for Impala

The LDAP support is preliminary. It currently has only been tested against Active Directory.

Using Multiple Authentication Methods with Impala

Impala 2.0 and later automatically handles both Kerberos and LDAP authentication. Each `impalad` daemon can accept both Kerberos and LDAP requests through the same port. No special actions need to be taken if some users authenticate through Kerberos and some through LDAP.

Prior to Impala 2.0, you had to configure each `impalad` to listen on a specific port depending on the kind of authentication, then configure your network load balancer to forward each kind of request to a DataNode that was set up with the appropriate authentication type. Once the initial request was made using either Kerberos or LDAP authentication, Impala automatically handled the process of coordinating the work across multiple nodes and transmitting intermediate results back to the coordinator node.

Configuring Impala Delegation for Hue and BI Tools

When users submit Impala queries through a separate application, such as Hue or a business intelligence tool, typically all requests are treated as coming from the same user. In Impala 1.2 and higher, Impala supports "delegation" where users whose names you specify can delegate the execution of a query to another user. The query runs with the privileges of the delegated user, not the original authenticated user.

The name of the delegated user is passed using the HiveServer2 protocol configuration property `impala.doas.user` when the client connects to Impala.

Currently, the delegation feature is available only for Impala queries submitted through application interfaces such as Hue and BI tools. For example, Impala cannot issue queries using the privileges of the HDFS user.

The delegation feature is enabled by the startup option for `impalad`: `--authorized_proxy_user_config`.

The syntax for the option is:

```
--authorized_proxy_user_config=authenticated_user1=delegated_user1,delegated_user2,...;authenticated_user2=...
```

- The list of authorized users are delimited with `;`
- The list of delegated users are delimited with `,` by default.
- Use the `--authorized_proxy_user_config_delimiter` startup option to override the default user delimiter (the comma character) to another character.
- Wildcard (`*`) is supported to delegated to any users, e.g. `--authorized_proxy_user_config=hue=*`. Make sure to use single quotes or escape characters to ensure that any `*` characters do not undergo wildcard expansion when specified in command-line arguments.

When you start Impala with the `--authorized_proxy_user_config=authenticated_user=delegated_user` option:

- Authentication is based on the user on the left hand side (*authenticated_user*).
- Authorization is based on the right hand side user(s) (*delegated_user*).
- When opening a client connection, the client must provide a delegated username via the HiveServer2 protocol property,`impala.doas.user` or `DelegationUID`.
- It is not necessary for *authenticated_user* to have the permission to access/edit files.
- It is not necessary for the delegated users to have access to the service via Kerberos.
- *delegated_user* must exist in the OS.
- In Impala, `user()` returns *authenticated_user* and `effective_user()` returns the delegated user that the client specified.

The user delegation process works as follows:

1. The Impalad daemon starts with the following option:
   - `--authorized_proxy_user_config=authenticated_user=delegated_user`
2. A client connects to Impala via the HiveServer2 protocol with the `impala.doas.user` configuration property, e.g. connected user is *authenticated_user* with `impala.doas.user=delegated_user`.
3. The client user *authenticated_user* sends a request to Impala as the delegated user *delegated_user*.

4. Impala checks if *delegated_user* is in the list of authorized delegate users for the user *authenticated_user*.
5. If the user is an authorized delegated user for *authenticated_user*, the request is executed as the delegate user *delegated_user*.

See this Cloudera blog post for background information about the delegation capability in HiveServer2.

To set up authentication for the delegated users:

- On the server side, configure either user/password authentication through LDAP, or Kerberos authentication, for all the delegated users. See Enabling LDAP Authentication for Impala on page 144 or Enabling Kerberos Authentication for Impala on page 141 for details.
- On the client side, to learn how to enable delegation, consult the documentation for the ODBC driver you are using.

Enabling Delegation in Cloudera Manager

To enable delegation in Cloudera Manager:

1. Navigate to **Clusters** > **Impala** > **Configuration** > **Policy File-Based Sentry**.
2. In the **Proxy User Configuration** field, type the a semicolon-separated list of key=value pairs of authorized proxy users to the user(s) they can impersonate. The list of delegated users are delimited with a comma, e.g. **hue=user1, user2**.

   The user names should be given in the short form. The names are case-sensitive
3. Click **Save Changes** and then restart Impala service.

Llama Authentication

> **Note:**
>
> The use of the Llama component for integrated resource management within YARN is no longer supported with and higher. The Llama support code is removed entirely in and higher.
>
> For clusters running Impala alongside other data management components, you define static service pools to define the resources available to Impala and other components. Then within the area allocated for Impala, you can create dynamic service pools, each with its own settings for the Impala admission control feature.

This section describes how to configure Llama in CDH 5 with Kerberos security in a Hadoop cluster.

> **Note:** Llama has been tested only in a Cloudera Manager deployment. For information on using Cloudera Manager to configure Llama and Impala, see Installing Impala.

Configuring Llama to Support Kerberos Security

1. Create a Llama service user principal using the syntax: `llama/fully.qualified.domain.name@YOUR-REALM`. This principal is used to authenticate with the Hadoop cluster, where *fully.qualified.domain.name* is the host where Llama is running and *YOUR-REALM* is the name of your Kerberos realm:

```
$ kadmin
kadmin: addprinc -randkey
llama/fully.qualified.domain.name@YOUR-REALM
```

2. Create a keytab file with the Llama principal:

```
$ kadmin
kadmin: xst -k llama.keytab llama/fully.qualified.domain.name
```

3. Test that the credentials in the keytab file work. For example:

```
$ klist -e -k -t llama.keytab
```

4. Copy the `llama.keytab` file to the Llama configuration directory. The owner of the `llama.keytab` file should be the `llama` user and the file should have owner-only read permissions.
5. Edit the Llama `llama-site.xml` configuration file in the Llama configuration directory by setting the following properties:

Property	Value
`llama.am.server.thrift.security`	`true`
`llama.am.server.thrift.kerberos.keytab.file`	`llama/conf.keytab`
`llama.am.server.thrift.kerberos.server.principal.name`	`llama/`*`fully.qualified.domain.name`*
`llama.am.server.thrift.kerberos.notification.principal.name`	`impala`

6. Restart Llama to make the configuration changes take effect.

Oozie Authentication

This section describes how to configure Oozie CDH 5 with Kerberos security on a Hadoop cluster:

- Configuring Kerberos Authentication for the Oozie Server on page 149
- Configuring Oozie HA with Kerberos on page 151

> **Important:**
>
> To enable Oozie to work with Kerberos security on your Hadoop cluster, make sure you perform the installation and configuration steps in Configuring Hadoop Security in CDH 5. Also note that when Kerberos security is enabled in Oozie, a web browser that supports Kerberos HTTP SPNEGO is required to access the Oozie web-console (for example, Firefox, Internet Explorer or Chrome).

> **Important:**
>
> If the NameNode, Secondary NameNode, DataNode, JobTracker, TaskTrackers, ResourceManager, NodeManagers, HttpFS, or Oozie services are configured to use Kerberos HTTP SPNEGO authentication, and two or more of these services are running on the same host, then all of the running services must use the same HTTP principal and keytab file used for their HTTP endpoints.

Configuring Kerberos Authentication for the Oozie Server

1. Create a Oozie service user principal using the syntax: `oozie/<fully.qualified.domain.name>@<YOUR-REALM>`. This principal is used to authenticate with the Hadoop cluster. where: `fully.qualified.domain.name` is the host where the Oozie server is running `YOUR-REALM` is the name of your Kerberos realm.

```
kadmin: addprinc -randkey oozie/fully.qualified.domain.name@YOUR-REALM.COM
```

2. Create a HTTP service user principal using the syntax: `HTTP/<fully.qualified.domain.name>@<YOUR-REALM>`. This principal is used to authenticate user requests coming to the Oozie web-services. where: `fully.qualified.domain.name` is the host where the Oozie server is running `YOUR-REALM` is the name of your Kerberos realm.

```
kadmin: addprinc -randkey HTTP/fully.qualified.domain.name@YOUR-REALM.COM
```

> **Important:**
>
> The `HTTP/` component of the HTTP service user principal must be upper case as shown in the syntax and example above.

3. Create keytab files with both principals.

```
$ kadmin
kadmin: xst -k oozie.keytab oozie/fully.qualified.domain.name
kadmin: xst -k http.keytab HTTP/fully.qualified.domain.name
```

4. Merge the two keytab files into a single keytab file:

```
$ ktutil
ktutil: rkt oozie.keytab
ktutil: rkt http.keytab
ktutil: wkt oozie-http.keytab
```

5. Test that credentials in the merged keytab file work. For example:

```
$ klist -e -k -t oozie-http.keytab
```

6. Copy the `oozie-http.keytab` file to the Oozie configuration directory. The owner of the `oozie-http.keytab` file should be the `oozie` user and the file should have owner-only read permissions.
7. Edit the Oozie server `oozie-site.xml` configuration file in the Oozie configuration directory by setting the following properties:

> **Important:** You must restart the Oozie server to have the configuration changes take effect.

Property	Value
`oozie.service.HadoopAccessorService.kerberos.enabled`	`true`
`local.realm`	`<REALM>`
`oozie.service.HadoopAccessorService.keytab.file`	`/etc/oozie/conf/oozie-http.keytab` for a package installation, or `<EXPANDED_DIR>/conf/oozie-http.keytab` for a tarball installation
`oozie.service.HadoopAccessorService.kerberos.principal`	`oozie/<fully.qualified.domain.name>@<YOUR-REALM.COM>`
`oozie.authentication.type`	`kerberos`
`oozie.authentication.kerberos.principal`	`HTTP/<fully.qualified.domain.name>@<YOUR-REALM.COM>`
`oozie.authentication.kerberos.name.rules`	Use the value configured for `hadoop.security.auth_to_local` in `core-site.xml`

Configuring Oozie HA with Kerberos

> **Important:**
> - If you use Cloudera Manager, do not use these command-line instructions. Use the Cloudera Manager Kerberos wizard instead, which automates the steps described in this section. If you have already enabled Kerberos, Cloudera Manager will automatically generate Kerberos credentials for the new Oozie server. It will also regenerate credentials for any existing servers.
> - This information applies specifically to CDH 5.14.0. If you use a lower version of CDH, see the documentation for that version located at Cloudera Documentation.

In CDH 5, you can configure multiple active Oozie servers against the same database, providing high availability for Oozie. For instructions on setting up Oozie HA, see Oozie High Availability

Let's assume a setup with three hosts running Oozie servers: `host1.example.com`, `host2.example.com`, and `host3.example.com`. The Load Balancer which directs traffic to the Oozie servers is running on `oozie.example.com`. Perform the following steps to configure Kerberos authentication on this Oozie HA-enabled deployment:

1. Assuming your Kerberos realm is `EXAMPLE.COM`, create the following Kerberos principals:
   - `oozie/host1.example.com@EXAMPLE.COM`
   - `oozie/host2.example.com@EXAMPLE.COM`
   - `oozie/host3.example.com@EXAMPLE.COM`
   - `HTTP/host1.example.com@EXAMPLE.COM`
   - `HTTP/host2.example.com@EXAMPLE.COM`
   - `HTTP/host3.example.com@EXAMPLE.COM`
   - For the Load Balancer: `HTTP/oozie.example.com@EXAMPLE.COM`
2. On each host, create a keytab file with the corresponding `oozie` and `HTTP` principals from the list above. Each keytab file should also have the Load Balancer's HTTP principal. For example, the keytab file on `host1` would comprise:
   - `oozie/host1.example.com@EXAMPLE.COM`
   - `HTTP/host1.example.com@EXAMPLE.COM`
   - `HTTP/oozie.example.com@EXAMPLE.COM`
3. On each host, configure the following properties in `oozie-site.xml`:

```
<property>
    <name>oozie.authentication.kerberos.principal</name>
    <value>HTTP/<hostname>@$EXAMPLE.COM</value>
    <description>
        Indicates the Kerberos principal to be used for HTTP endpoint.
       The principal MUST start with 'HTTP/' as per Kerberos HTTP SPNEGO specification.

    </description>
</property>

<property>
    <name>oozie.authentication.kerberos.keytab</name>
    <value>${oozie.service.HadoopAccessorService.keytab.file}</value>
    <description>
        Location of the keytab file with the credentials for the principal.
        Referring to the same keytab file Oozie uses for its Kerberos credentials for
Hadoop.
    </description>
</property>
```

Solr Authentication

This section describes how to configure Solr to enable authentication.

When authentication is enabled, only specified hosts and users can connect to Solr. Authentication also verifies that clients connect to legitimate servers. This feature prevents spoofing such as impersonation and man-in-the-middle attacks. Search supports Kerberos and LDAP authentication.

Cloudera Search supports a variety of combinations of authentication protocols:

Table 8: Authentication Protocol Combinations

Solr Authentication	Use Case
No authentication	Insecure cluster
Kerberos only	The Hadoop cluster has Kerberos turned on and every user (or client) connecting to Solr has a Kerberos principal.
Kerberos and LDAP	The Hadoop cluster has Kerberos turned on. External Solr users (or clients) do not have Kerberos principals but do have identities in the LDAP server. Client authentication using LDAP requires that Kerberos is enabled for the cluster. Using LDAP alone is not supported.

Once you are finished setting up authentication, configure Sentry authorization. Authorization involves specifying which resources can be accessed by particular users when they connect through Search. See Configuring Sentry Authorization for Cloudera Search for details.

Enabling Kerberos Authentication for Solr

Solr supports Kerberos authentication. All necessary packages are installed when you install Search. To enable Kerberos, create principals and keytabs and then modify default configurations.

> **Important:** The following instructions only apply to configuring Kerberos in an unmanaged environment. If you are using Cloudera Manager, Kerberos configuration is automatically handled. For more information on enabling Kerberos using Cloudera Manager, see Configuring Authentication in Cloudera Manager on page 43.

Create Principals and Keytabs

Repeat this process on all Solr server hosts:

1. Create a Solr service user principal using the syntax: `solr/solr01.example.com@EXAMPLE.COM`. This principal is used to authenticate with the Hadoop cluster. Replace `solr01.example.com` with the Solr server host, and `EXAMPLE.COM` with your Kerberos realm.

   ```
   $ kadmin
   kadmin: addprinc -randkey solr/solr01.example.com@EXAMPLE.COM
   ```

2. Create a HTTP service user principal using the syntax: `HTTP/solr01.example.com@EXAMPLE.COM`. This principal is used to authenticate user requests coming to the Solr web-services. Replace `solr01.example.com` with the Solr server host, and `EXAMPLE.COM` with your Kerberos realm.

   ```
   kadmin: addprinc -randkey HTTP/solr01.example.com@EXAMPLE.COM
   ```

   > **Note:**
   >
   > The `HTTP/` component of the HTTP service user principal must be upper case as shown in the syntax and example above.

3. Create keytab files with both principals.

   ```
   kadmin: xst -norandkey -k solr.keytab solr/solr01.example.com \
   HTTP/solr01.example.com
   ```

4. Test that credentials in the merged keytab file work. For example:

   ```
   $ klist -e -k -t solr.keytab
   ```

5. Copy the `solr.keytab` file to the Solr configuration directory. The owner of the `solr.keytab` file should be the `solr` user and the file should have owner-only read permissions.

Modify Default Configurations

Repeat this process on all Solr server hosts:

1. Ensure that the following properties appear in `/etc/default/solr` or `/opt/cloudera/parcels/CDH-*/etc/default/solr` and that they are uncommented. Modify these properties to match your environment. The relevant properties to be uncommented and modified are:

   ```
   SOLR_AUTHENTICATION_TYPE=kerberos
   SOLR_AUTHENTICATION_SIMPLE_ALLOW_ANON=true
   SOLR_AUTHENTICATION_KERBEROS_KEYTAB=/etc/solr/conf/solr.keytab
   SOLR_AUTHENTICATION_KERBEROS_PRINCIPAL=HTTP/localhost@LOCALHOST
   SOLR_AUTHENTICATION_KERBEROS_NAME_RULES=DEFAULT
   SOLR_AUTHENTICATION_JAAS_CONF=/etc/solr/conf/jaas.conf
   ```

   > **Note:** Modify the values for these properties to match your environment. For example, the `SOLR_AUTHENTICATION_KERBEROS_PRINCIPAL=HTTP/localhost@LOCALHOST` must include the principal instance and Kerberos realm for your environment. That is often different from `localhost@LOCALHOST`.

2. Set `hadoop.security.auth_to_local` to match the value specified by `SOLR_AUTHENTICATION_KERBEROS_NAME_RULES` in `/etc/default/solr` or `/opt/cloudera/parcels/CDH-*/etc/default/solr`.

   > **Note:** For information on how to configure the rules, see Configuring the Mapping from Kerberos Principals to Short Names. For additional information on using Solr with HDFS, see Configuring Solr for Use with HDFS.

3. If using applications that use the `solrj` library, set up the Java Authentication and Authorization Service (JAAS).

   a. Create a `jaas.conf` file in the Solr configuration directory containing the following settings. This file and its location must match the `SOLR_AUTHENTICATION_JAAS_CONF` value. Make sure that you substitute a value for `principal` that matches your particular environment.

   ```
   Client {
     com.sun.security.auth.module.Krb5LoginModule required
     useKeyTab=true
     useTicketCache=false
     keyTab="/etc/solr/conf/solr.keytab"
     principal="solr/fully.qualified.domain.name@<YOUR-REALM>";
   };
   ```

Enabling LDAP Authentication for Solr

Before continuing, make sure that you have completed the steps in Enabling Kerberos Authentication for Solr on page 152. Solr supports LDAP authentication for external Solr client including:

- Command-line tools
- curl
- Web browsers
- Solr Java clients

In some cases, Solr does not support LDAP authentication. Use Kerberos authentication instead in these cases. Solr does not support LDAP authentication with:

- Search indexing components including the MapReduce indexer, Lily HBase indexer, or Flume.

- Solr internal requests such as those for replication or querying.
- Hadoop delegation token management requests such as `GETDELEGATIONTOKEN` or `RENEWDELEGATIONTOKEN`.

Configuring LDAP Authentication for Solr using Cloudera Manager

You can configure LDAP-based authentication using Cloudera Manager at the Solr service level.

1. Go to the **Solr** service.
2. Click the **Configuration** tab.
3. Select **Scope** > **Solr**
4. Select **Category** > **Security**
5. Select **Enable LDAP**.
6. Enter the LDAP URI in the **LDAP URI** property.
7. Configure only one of following mutually exclusive parameters:
   - **LDAP BaseDN**: Replaces the username with a "distinguished name" (DN) of the form: `uid=userid,ldap_baseDN`. Typically used for OpenLDAP server installation.

   -OR-

   - **Active Directory Domain**: Replaces the username with a string `username@ldap_domain`. Typically used for Active Directory server installation.

Configuring LDAP Authentication for Solr Using the Command Line

To enable LDAP authentication using the command line, configure the following environment variables in `/etc/default/solr`:

```
SOLR_AUTHENTICATION_HTTP_SCHEMES=Negotiate,Basic
SOLR_AUTHENTICATION_HTTP_DELEGATION_MGMT_SCHEMES=Negotiate
SOLR_AUTHENTICATION_HTTP_BASIC_HANDLER=ldap
SOLR_AUTHENTICATION_HTTP_NEGOTIATE_HANDLER=kerberos
SOLR_AUTHENTICATION_LDAP_PROVIDER_URL=ldap://www.example.com

# Specify value for only one of SOLR_AUTHENTICATION_LDAP_BASE_DN or
SOLR_AUTHENTICATION_LDAP_BIND_DOMAIN property.
SOLR_AUTHENTICATION_LDAP_BASE_DN=ou=Users,dc=example,dc=com
# SOLR_AUTHENTICATION_LDAP_BIND_DOMAIN=
# Required when using ‘Start TLS’ extension
# SOLR_AUTHENTICATION_LDAP_ENABLE_START_TLS=false
```

Securing LDAP Connections

You can secure communications using LDAP-based encryption.

To avoid sending credentials over the wire in clear-text, you must configure a secure connection between both the client and Solr, and between Solr and the LDAP server. The secure connection could use SSL or TLS.

Secure LDAP connections through SSL:

For SSL-enabled LDAP connections, specify a prefix of `ldaps://` instead of `ldap://`. Also, the default port for SSL-enabled LDAP connections is 636 instead of 389.

Secure LDAP connections through TLS:

TLS, the successor to the SSL protocol, is supported by most modern LDAP servers. Unlike SSL connections, TLS connections can be made on the same server port as non-TLS connections. You can enable xxx using Cloudera Manager.

1. Go to the **Solr** service.
2. Click the **Configuration** tab.
3. Select **Scope** > **Solr**
4. Select **Category** > **Security**
5. Select **Enable LDAP TLS**.
6. Import the LDAP server security certificate in the Solr Trust Store file:

a. Enter the location for the Solr Trust Store File in **Solr TLS/SSL Certificate Trust Store File**.
b. Enter the password for the Solr Trust Store File in **Solr TLS/SSL Certificate Trust Store Password**.

LDAP Client Configuration

Some HTTP clients such as curl or the Apache Http Java client must be configured to use a particular scheme. For example:

- curl tool supports using Kerberos or username/password authentication. Kerberos is activated using the `--negotiate` flag and username/password based authentication is activated using the `--basic` and `-u` flags.
- Apache HttpClient library can be configured to use specific authentication scheme. For more information, see the HTTP authentication chapter of Apache's HttpClient Tutorial.

Typically, web browsers automatically choose a preferred authentication scheme. For more information, see the HTTP authentication topic in The Chromium Projects.

To use LDAP authentication with Solr Java clients, `HttpClientConfigurer` needs to configured for Solr. This can either be done programmatically or using Java system properties.

For example, programmatic initialization might appear as:

SampleSolrClient.java

```
import org.apache.solr.client.solrj.impl.HttpClientUtil;
import org.apache.solr.client.solrj.impl.PreemptiveBasicAuthConfigurer;
import org.apache.solr.common.params.ModifiableSolrParams;

/**
 * This method initializes the Solr client to use LDAP authentication
 * This configuration is applicable to all Solr clients.
 * @param ldapUserName LDAP user name
 * @param ldapPassword  LDAP user password
 */
public static void initialize(String ldapUserName, String ldapPassword) {
  HttpClientUtil.setConfigurer(new PreemptiveBasicAuthConfigurer());
  ModifiableSolrParams params = new ModifiableSolrParams();
  params.set(HttpClientUtil.PROP_BASIC_AUTH_USER, ldapUserName);
  params.set(HttpClientUtil.PROP_BASIC_AUTH_PASS, ldapPassword);
  // Configure the JVM default parameters.
  PreemptiveBasicAuthConfigurer.setDefaultSolrParams(params);
}
```

For configuration using system properties, configure the following system properties:

Table 9: System properties configuration for LDAP authentication

System property	Description
`solr.httpclient.configurer`	Fully qualified classname of `HttpClientConfigurer` implementation. For example, `org.apache.solr.client.solrj.impl.PreemptiveBasicAuthConfigurer`.
`solr.httpclient.config`	Http client configuration properties file path. For example, `ldap-credentials.properties`.

For example, the entry in `ldap-credentials.properties` might appear as:

ldap-credentials.properties

```
httpBasicAuthUser=user1
httpBasicAuthPassword=passwd
```

Using Kerberos with Solr

The process of enabling Solr clients to authenticate with a secure Solr is specific to the client. This section demonstrates:

- Using Kerberos and curl
- Using solrctl
- Configuring SolrJ Library Usage
- This enables technologies including:
  - Command line solutions
  - Java applications
  - The MapReduceIndexerTool
- Configuring Flume Morphline Solr Sink Usage

Secure Solr requires that the CDH components that it interacts with are also secure. Secure Solr interacts with HDFS, ZooKeeper and optionally HBase, MapReduce, and Flume.

Using Kerberos and curl

You can use Kerberos authentication with clients such as `curl`. To use `curl`, begin by acquiring valid Kerberos credentials and then run the desired command. For example, you might use commands similar to the following:

```
$ kinit -kt username.keytab username
$ curl --negotiate -u foo:bar http://solrserver:8983/solr/
```

> **Note:** Depending on the tool used to connect, additional arguments may be required. For example, with `curl`, `--negotiate` and `-u` are required. The username and password specified with `-u` is not actually checked because Kerberos is used. As a result, any value such as foo:bar or even just : is acceptable. While any value can be provided for `-u`, note that the option is required. Omitting `-u` results in a 401 Unauthorized error, even though the `-u` value is not actually used.

Using solrctl

If you are using `solrctl` to manage your deployment in an environment that requires Kerberos authentication, you must have valid Kerberos credentials, which you can get using `kinit`. For more information on `solrctl`, see solrctl Reference

Configuring SolrJ Library Usage

If using applications that use the solrj library, begin by establishing a Java Authentication and Authorization Service (JAAS) configuration file.

Create a JAAS file:

- If you have already used `kinit` to get credentials, you can have the client use those credentials. In such a case, modify your `jaas-client.conf` file to appear as follows:

```
Client {
 com.sun.security.auth.module.Krb5LoginModule required
 useKeyTab=false
 useTicketCache=true
 principal="user/fully.qualified.domain.name@<YOUR-REALM>";
 };
```

  where `user/fully.qualified.domain.name@<YOUR-REALM>` is replaced with your credentials.
- If you want the client application to authenticate using a keytab, modify `jaas-client.conf` as follows:

```
Client {
 com.sun.security.auth.module.Krb5LoginModule required
 useKeyTab=true
 keyTab="/path/to/user.keytab"
 storeKey=true
 useTicketCache=false
```

```
 principal="user/fully.qualified.domain.name@EXAMPLE.COM";
};
```

Replace `/path/to/user.keytab` with the keytab file you want to use and `user/fully.qualified.domain.name@EXAMPLE.COM` with the principal in the keytab. If the principal omits the hostname, omit it in the `jaas-client.conf` file as well (for example, `jdoe@EXAMPLE.COM`).

Use the JAAS file to enable solutions:

- **Command line solutions**

  Set the property when invoking the program. For example, if you were using a jar, you might use:

```
java -Djava.security.auth.login.config=/home/user/jaas-client.conf -jar app.jar
```

- **Java applications**

  Set the Java system property `java.security.auth.login.config`. For example, if the JAAS configuration file is located on the filesystem as `/home/user/jaas-client.conf`. The Java system property `java.security.auth.login.config` must be set to point to this file. Setting a Java system property can be done programmatically, for example using a call such as:

```
System.setProperty("java.security.auth.login.config", "/home/user/jaas-client.conf");
```

- **The MapReduceIndexerTool**

  The MapReduceIndexerTool uses SolrJ to pass the JAAS configuration file. Using the MapReduceIndexerTool in a secure environment requires the use of the `HADOOP_OPTS` variable to specify the JAAS configuration file. For example, you might issue a command such as the following:

```
HADOOP_OPTS="-Djava.security.auth.login.config=/home/user/jaas.conf" \
hadoop jar MapReduceIndexerTool
```

- **Configuring the hbase-indexer CLI**

  Certain hbase-indexer CLI commands such as `replication-status` attempt to read ZooKeeper hosts owned by HBase. To successfully use these commands in Solr in a secure environment, specify a JAAS configuration file with the HBase principal in the `HBASE_INDEXER_OPTS` environment variable. For example, you might issue a command such as the following:

```
HBASE_INDEXER_OPTS="-Djava.security.auth.login.config=/home/user/hbase-jaas.conf" \
hbase-indexer replication-status
```

Configuring Flume Morphline Solr Sink Usage

Repeat this process on all Flume hosts:

1. If you have not created a keytab file, do so now at `/etc/flume-ng/conf/flume.keytab`. This file should contain the service principal `flume/<fully.qualified.domain.name>@<YOUR-REALM>`. See Flume Authentication on page 110 for more information.
2. Create a JAAS configuration file for flume at `/etc/flume-ng/conf/jaas-client.conf`. The file should appear as follows:

```
Client {
 com.sun.security.auth.module.Krb5LoginModule required
 useKeyTab=true
 useTicketCache=false
 keyTab="/etc/flume-ng/conf/flume.keytab"
 principal="flume/<fully.qualified.domain.name>@<YOUR-REALM>";
};
```

3. Add the flume JAAS configuration to the `JAVA_OPTS` in `/etc/flume-ng/conf/flume-env.sh`. For example, you might change:

```
JAVA_OPTS="-Xmx500m"
```

to:

```
JAVA_OPTS="-Xmx500m -Djava.security.auth.login.config=/etc/flume-ng/conf/jaas-client.conf"
```

Spark Authentication

Spark has an internal mechanism that authenticates executors with the driver controlling a given application. This mechanism is enabled using the Cloudera Manager Admin Console, as detailed in Enabling Spark Authentication.

When Spark on YARN is running on a secure cluster, users must authenticate to Kerberos before submitting jobs, as detailed in Running Spark Applications on Secure Clusters.

Enabling Spark Authentication

Minimum Required Role: Security Administrator (also provided by **Full Administrator**)

Spark has an internal mechanism that authenticates executors with the driver controlling a given application. This mechanism is enabled using the Cloudera Manager Admin Console, as detailed below. Cluster administrators can enable the `spark.authenticate` mechanism to authenticate the various processes that support a Spark application.

To enable this feature on the cluster:

1. Log into the Cloudera Manager Admin Console.
2. Select **Clusters** > **Spark** (or **Clusters** > **Spark_on_YARN**).
3. Click the **Configuration** menu.
4. Scroll down to the **Spark Authentication** setting, or search for `spark.authenticate` to find it.
5. In the Spark Authentication setting, click the checkbox next to the **Spark (Service-Wide)** property to activate the setting.
6. Click **Save Changes**.

> **Note:** If your cluster supports Spark 2, you must make the same change to the Spark 2 setting. For example:

7. Restart YARN:
   - Select **Clusters** > **YARN**.
   - Select **Restart** from the **Actions** drop-down selector.
8. Re-deploy the client configurations:
   - Select **Clusters** > ***Cluster_name***
   - Select **Deploy Client Configurations** from the **Actions** drop-down selector.

Running Spark Applications on Secure Clusters

Secure clusters are clusters that use Kerberos for authentication. For secure clusters, Spark History Server automatically uses Kerberos, so there's nothing to configure.

Users running Spark applications must first authenticate to Kerberos, using `kinit`, as follows:

```
$ kinit
ldap@EXAMPLE-REALM.COM:'s password:
```

After authenticating to Kerberos, users can submit their applications using `spark-submit` as usual, as shown below. This command submits one of the default Spark sample jobs using an environment variable as part of the path, so modify as needed for your own use:

```
$ spark-submit --class org.apache.spark.examples.SparkPi --master yarn \
--deploy-mode cluster $SPARK_HOME/lib/spark-examples.jar 10
```

Configuring Spark on YARN for Long-Running Applications

Long-running applications such as Spark Streaming jobs must be able to write to HDFS, which means that the `hdfs` user may need to delegate tokens possibly beyond the default lifetime. This workload type requires passing Kerberos principal and keytab to the `spark-submit` script using the `--principal` and `--keytab` parameters. The keytab is copied to the host running the ApplicationMaster, and the Kerberos login is renewed periodically by using the principal and keytab to generate the required delegation tokens needed for HDFS.

> **Note:** For secure distribution of the keytab to the ApplicationMaster host, the cluster should be configured for TLS/SSL communication for YARN and HDFS encryption .

Create the Spark Principal and Keytab File

These are needed for long-running applications running on Spark on YARN `cluster` mode only.

1. Create the `spark` principal and `spark.keytab` file:

```
kadmin: addprinc -randkey spark/fully.qualified.domain.name@YOUR-REALM.COM
kadmin: xst -k spark.keytab spark/fully.qualified.domain.name
```

See Step 4: Create and Deploy the Kerberos Principals and Keytab Files on page 84 for more information about Kerberos and its use with Cloudera clusters.

Sqoop 2 Authentication

This section describes how to configure Sqoop 2 with Kerberos security in a Hadoop cluster.

> **Note:** Sqoop 2 is being deprecated. Cloudera recommends using Sqoop 1.

Create the Sqoop 2 Principal and Keytab File

You need to create a `sqoop2.keytab` file for Sqoop 2. Follow these steps:

1. Create the principal and keytab file:

```
kadmin: addprinc -randkey sqoop2/fully.qualified.domain.name@YOUR-REALM.COM
kadmin: xst -k sqoop2.keytab sqoop2/fully.qualified.domain.name
```

2. Move the file into the Sqoop 2 configuration directory and restrict its access exclusively to the `sqoop2` user:

```
$ mv sqoop2.keytab /etc/sqoop2/conf/
$ chown sqoop2 /etc/sqoop2/conf/sqoop2.keytab
$ chmod 400 /etc/sqoop2/conf/sqoop2.keytab
```

For more details on creating Kerberos principals and keytabs, see Step 4: Create and Deploy the Kerberos Principals and Keytab Files on page 84.

Configure Sqoop 2 to Use Kerberos

Edit the Sqoop 2 configuration file `sqoop.properties` file in the `/etc/sqoop2/conf` directory and add the following properties:

```
org.apache.sqoop.authentication.type=KERBEROS
org.apache.sqoop.authentication.handler=org.apache.sqoop.security.KerberosAuthenticationHandler
org.apache.sqoop.authentication.kerberos.principal=sqoop2/fully.qualified.domain.name@YOUR-REALM.COM
org.apache.sqoop.authentication.kerberos.keytab=/etc/sqoop2/conf/sqoop2.keytab
```

Sqoop 1, Pig, and Whirr Security

Here is a summary of some security features for other CDH 5 components.

Pig

Supports security with no configuration required.

Sqoop 1

Sqoop1 does not support:

- TLS/SSL connections to Oracle, MySQL, or other databases
- Kerberos authentication to external databases

Whirr

Whirr does not support security in CDH 5.

ZooKeeper Authentication

As of Cloudera Manager 5.11, ZooKeeper supports mutual server-to-server (quorum peer) authentication using SASL (Simple Authentication and Security Layer), which provides a layer around Kerberos authentication. Server to server authentication among ZooKeeper servers in an ensemble mitigates the risk of spoofing by a rogue server on an unsecured network. For more information about quorum peer authentication and how the feature leverages ZooKeeper's SASL support, see the Cloudera Engineering Blog post, Hardening Apache ZooKeeper Security.

Client-to-server SASL-based authentication has been supported since Cloudera Manager/CDH 5.2 (ZooKeeper 3.4.0+). Follow the steps in Configuring ZooKeeper Server for Kerberos Authentication on page 160 and Configuring ZooKeeper Client Shell for Kerberos Authentication on page 162 to configure ZooKeeper to use this mechanism.

To configure client-server or server-server authentication for ZooKeeper, follow the appropriate steps below:

Requirements

Configuring ZooKeeper to use Kerberos for client-server or server-server authentication requires that your organization's Kerberos instance (MIT Kerberos, Microsoft Active Directory) be up and running, and reachable by the ZooKeeper server or client during the configuration processes detailed below. See Configuring Hadoop Security in CDH 5 for details.

Before enabling mutual authentication, the ZooKeeper servers in the cluster must be configured to authenticate using Kerberos.

> **Note:** Cloudera recommends that you ensure your ZooKeeper ensemble is working properly, before you attempt to integrate Kerberos authentication. See ZooKeeper Installation for details.

Configuring ZooKeeper Server for Kerberos Authentication

You can configure the ZooKeeper server for Kerberos authentication in Cloudera Manager or through the command line.

Using Cloudera Manager to Configure ZooKeeper Server for Kerberos Authentication

To set up the ZooKeeper server for Kerberos authentication in Cloudera Manager, complete the following steps:

1. In Cloudera Manager, open the ZooKeeper service.
2. Click the **Configuration** tab.
3. Enter **Kerberos** in the in the **Search** bar.
4. Find the **Enable Kerberos Authentication** property and select the check-box next to the ZooKeeper services that you want to configure for Kerberos authentication.

Using the Command Line to Configure ZooKeeper Server for Kerberos Authentication

Follow the steps below for each ZooKeeper server in the ensemble. To maintain consistency across ZooKeeper servers in the ensemble, use the same name for the keytab file you deploy to each server, for example, `zookeeper.keytab`. Each keytab file will contain its respective host's fully-qualified domain name (FQDN).

1. Create a service principal for the ZooKeeper server using the fully-qualified domain name (FQDN) of the host on which ZooKeeper server is running and the name of your Kerberos realm using the pattern `zookeeper/fqdn.example.com@ YOUR-REALM`. This principal will be used to authenticate the ZooKeeper server with the Hadoop cluster. Create this service principal as follows:

```
kadmin: addprinc -randkey zookeeper/fqdn.example.com@YOUR-REALM
```

2. Create a keytab file for the ZooKeeper server:

```
$ kadmin
kadmin: xst -k zookeeper.keytab zookeeper/fqdn.example.com@YOUR-REALM
```

> **Note:** For consistency across ZooKeeper Servers, use the same name for the keytab file you create for each subsequent ZooKeeper Server host system you configure using these steps, for example, `zookeeper.keytab`.

3. Copy the `zookeeper.keytab` file to the ZooKeeper configuration directory on the ZooKeeper server host, using the appropriate ZooKeeper configuration directory: `/etc/zookeeper/conf/`. The `zookeeper.keytab` file should be owned by the `zookeeper` user, with owner-only read permissions.
4. Add the following lines to the ZooKeeper configuration file `zoo.cfg`:

```
authProvider.1=org.apache.zookeeper.server.auth.SASLAuthenticationProvider
jaasLoginRenew=3600000
```

5. Set up the Java Authentication and Authorization Service (JAAS) by creating a `jaas.conf` file in the ZooKeeper configuration directory with the settings shown below, replacing *`fqdn.example.com`* with the ZooKeeper server's hostname.

```
Server {
  com.sun.security.auth.module.Krb5LoginModule required
  useKeyTab=true
  keyTab="/etc/zookeeper/conf/zookeeper.keytab"
  storeKey=true
  useTicketCache=false
  principal="zookeeper/fqdn.example.com
  @YOUR-REALM";
};
```

6. Add the following setting to the `java.env` file located in the ZooKeeper configuration directory, creating the file if necessary:

```
export JVMFLAGS="-Djava.security.auth.login.config=/etc/zookeeper/conf/jaas.conf"
```

7. Repeat these steps for each ZooKeeper server in the ensemble.
8. Restart the ZooKeeper server to have the configuration changes take effect. See ZooKeeper Installation for details.

Configuring ZooKeeper Client Shell for Kerberos Authentication

In addition to configuring ZooKeeper Server hosts to use Kerberos for authentication, you should also configure the ZooKeeper client shell (the ZooKeeper CLI) to authenticate to the ZooKeeper service using Kerberos credentials. As with the ZooKeeper Server, you must create a Kerberos principal for the client, as detailed below:

1. Create a Kerberos principal for the `zookeeper-client`, `zkcli@`*`YOUR-REALM`*, replacing *YOUR-REALM* with the name of your organization's Kerberos realm:

```
kadmin: addprinc -randkey zkcli@YOUR-REALM
```

2. Create a keytab file for the ZooKeeper client shell using the `-norandkey` option.

> **Note:** Not all versions of `kadmin` support the `-norandkey` option, in which case, simply omit this option from the command. Using the `kadmin` command without the `-norandkey` option invalidates previously exported keytabs and generates a new password.

```
$ kadmin
kadmin: xst -norandkey -k zkcli.keytab zkcli@YOUR-REALM
```

3. On the host running the ZooKeeper client shell, set up JAAS (Java Authentication and Authorization Service) in the configuration directory appropriate for your installation type:
   - Package installation: `/etc/zookeeper/conf/`
   - Tarball installation: *`EXPANDED_DIR`*`/conf`
4. Create a `jaas.conf` file containing the following settings:

```
Client {
  com.sun.security.auth.module.Krb5LoginModule required
  useKeyTab=true
  keyTab="/path/to/zkcli.keytab"
  storeKey=true
  useTicketCache=false
  principal="zkcli@YOUR-REALM";
};
```

5. In this same configuration directory, add the following setting to the `java.env` file, creating the file if necessary.

```
export JVMFLAGS="-Djava.security.auth.login.config=/etc/zookeeper/conf/jaas.conf"
```

Verifying the Configuration

After enabling Kerberos authentication and restarting the ZooKeeper cluster, you can verify that the authentication is working correctly by following these steps:

1. Start the ZooKeeper client, passing to it the name of a ZooKeeper server:

```
zookeeper-client -server fqdn.example.com:port
```

2. From the ZooKeeper CLI, create a protected `znode` using your ZooKeeper client principal:

```
create /znode1 znode1data sasl:zkcli@{{YOUR-REALM}}:cdwra
```

3. Verify the znode is created and the ACL is set correctly:

```
getAcl /znode1
```

The `getAcl` command returns the znode's scheme and permissions values. Verify that these are as expected.

Enabling Server-Server Mutual Authentication

As of Cloudera Manager 5.11, support for mutual authentication between ZooKeeper Servers can be enabled through the Cloudera Manager Admin Console. For secured networks, server-to-server authentication is considered an optional security enhancement, so the capability is disabled by default:

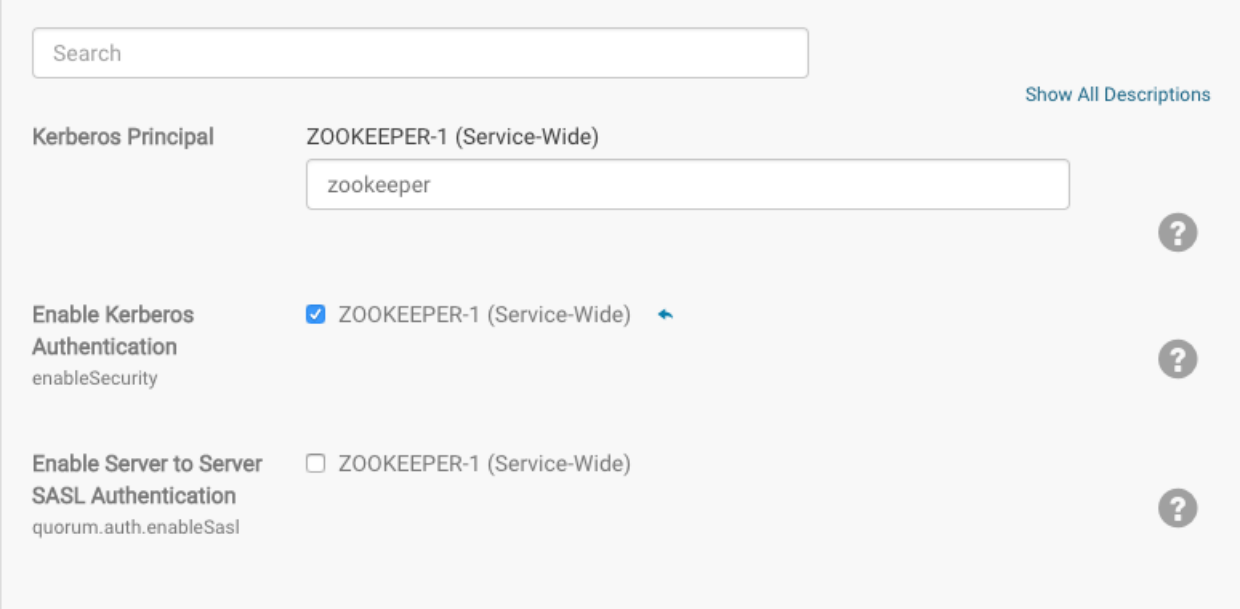

Server-to-server SASL authentication requires all servers in the ZooKeeper ensemble to authenticate using Kerberos, as detailed in Configuring ZooKeeper Server for Kerberos Authentication on page 160.

Assuming your cluster is already configured to authenticate using Kerberos, you can enable mutual authentication as follows:

1. Log into the Cloudera Manager Admin Console.
2. Select **Clusters** > **ZOOKEEPER-n**.
3. Click the **Configuration** tab.
4. Select **Category** > **Security** under the Filters menu to display the security properties.
5. Click the **Enable Server to Server SASL Authentication** box to select it.
6. Click **Save**.
7. Select **Restart** from the **Actions** drop-down to restart the cluster with this setting.

To disable server-to-server SASL authentication, simply return to the Cloudera Manager Admin Console page shown above, de-select **Enable Server to Server SASL Authentication** (by clicking the checked box), and restart the cluster.

Configuring a Dedicated MIT KDC for Cross-Realm Trust

Using Cloudera Manager to configure Kerberos authentication for the cluster creates several principals and keytabs automatically. Cloudera Manager also deploys the keytab files to every host in the cluster. See Hadoop Users (user:group) and Kerberos Principals on page 103 for complete listing.

> **Note:** The example below is specific for creating and deploying principals and keytab files for MIT Kerberos. See the appropriate documentation for other Kerberos implementations, such as Microsoft Active Directory, as needed.

Local and Remote Kerberos Admin Tools

Kerberos administrator commands can be run directly on the KDC server host or remotely, as shown in the table:

`kadmin.local`	Requires root access or Kerberos admin account. Use to log on directly to the KDC host.
`kadmin`	Use the logon to the KDC host system from another remote host over the network.

- To run Kerberos administration commands locally on the KRB host system:

```
$ sudo kadmin.local
```

  Enter your Linux system password (for the `sudo`).
- To run Kerberos administration commands from any host:

```
$ kadmin
```

  Enter your Kerberos administrator password.

> **Note:** Commands shown for the `kadmin` shell can also be run at the `kadmin.local` shell.

Setting up a Cluster-Dedicated KDC and Default Realm for the Hadoop Cluster

Cloudera has tested the following configuration approaches to Kerberos security for clusters managed by Cloudera Manager. For administration teams that are just getting started with Kerberos security, we recommend starting with these approaches to the configuration of KDC services for a number of reasons.

The number of Service Principal Names (SPNs) that are created and managed by the Cloudera Manager server for a CDH cluster can be significant, so it is important to realize the potential impact on cluster uptime and overall operations if you choose to manage keytabs manually instead. The Cloudera Manager server manages the creation of service keytabs on the proper hosts based on the current configuration of the database. Manual keytab management can be error prone and introduce delays when deploying or moving services within the cluster, especially under time-sensitive conditions.

Cloudera Manager creates SPNs within a KDC that it can access with the `kadmin` command based on configuration of the `/etc/krb5.conf` file on the Cloudera Manager host. SPNs are created with the format `service-name/host.fqdn.name@EXAMPLE.COM` where `service-name` is the relevant CDH service name such as hue or `hbase` or `hdfs`.

If your site already has a working KDC, and any existing principals share the same name as any of the principals that Cloudera Manager creates, the Cloudera Manager Server generates a new randomized key for those principals, and consequently causes existing keytabs to become invalid.

This is why Cloudera recommends using a dedicated local MIT Kerberos KDC and realm for the Hadoop cluster. You can set up a one-way cross-realm trust from the cluster-dedicated KDC and realm to your existing central MIT Kerberos KDC, or to an existing Active Directory realm. Using this method, there is no need to create Hadoop service principals in the central MIT Kerberos KDC or in Active Directory, but principals (users) in the central MIT KDC or in Active Directory can be authenticated to Hadoop. The steps to implement this approach are as follows:

1. Install and configure a cluster-dedicated MIT Kerberos KDC that will be managed by Cloudera Manager for creating and storing principals for the services supported by the cluster.

   > **Note:** The `krb5-server` package includes a `logrotate` policy file to rotate log files monthly. To take advantage of this, install the `logrotate` package. No additional configuration is necessary.

2. See the example `kdc.conf` and `krb5.conf` files in Sample Kerberos Configuration Files on page 425 for configuration considerations for the KDC and Kerberos clients.

3. Configure a default Kerberos realm for the cluster you want Cloudera Manager to manage and set up one-way cross-realm trust between the cluster-dedicated KDC and either your central KDC or Active Directory, using the appropriate steps:
   - Using a Cluster-Dedicated KDC with a Central MIT KDC on page 165
   - Using a Cluster-Dedicated MIT KDC with Active Directory on page 166

Cloudera strongly recommends the method above because:

- It requires minimal configuration in Active Directory.
- It is comparatively easy to script the creation of many principals and keytabs. A principal and keytab must be created for every daemon in the cluster, and in a large cluster this can be extremely onerous to do directly in Active Directory.
- There is no need to involve central Active Directory administrators to get service principals created.
- It allows for incremental configuration. The Hadoop administrator can completely configure and verify the functionality the cluster independently of integrating with Active Directory.

Using a Cluster-Dedicated KDC with a Central MIT KDC

> **Important:** If you plan to use Oozie or the Hue Kerberos Ticket Renewer in your cluster, you must configure your KDC to allow tickets to be renewed, and you must configure `krb5.conf` to request renewable tickets. Typically, you can do this by adding the `max_renewable_life` setting to your realm in `kdc.conf`, and by adding the `renew_lifetime` parameter to the `libdefaults` section of `krb5.conf`. about renewable tickets, see the Kerberos documentation. See the Sample Kerberos Configuration Files on page 425 for an example of configuration for ticket renewal.

1. In the `/var/kerberos/krb5kdc/kdc.conf` file on the local dedicated KDC server host, configure the default realm for the Hadoop cluster by substituting your Kerberos realm in the following `realms` property:

```
[realms]
 HADOOP.EXAMPLE.COM = {
```

2. In the `/etc/krb5.conf` file on all cluster hosts and all Hadoop client user hosts, configure the default realm for the Hadoop cluster by substituting your Kerberos realm in the following `realms` property. Also specify the local dedicated KDC server hostname in the `/etc/krb5.conf` file (for example, `kdc01.example.com`).

```
[libdefaults]
  default_realm = HADOOP.EXAMPLE.COM
[realms]
  HADOOP.EXAMPLE.COM = {
    kdc = kdc01.hadoop.example.com:88
    admin_server = kdc01.hadoop.example.com:749
    default_domain = hadoop.example.com
  }
  EXAMPLE.COM = {
    kdc = kdc01.example.com:88
    admin_server = kdc01.example.com:749
    default_domain = example.com
  }
[domain_realm]
  .hadoop.example.com = HADOOP.EXAMPLE.COM
  hadoop.example.com = HADOOP.EXAMPLE.COM
  .example.com = EXAMPLE.COM
  example.com = EXAMPLE.COM
```

3. To set up the cross-realm trust in the cluster-dedicated KDC, type the following command in the `kadmin.local` or `kadmin` shell on the cluster-dedicated KDC host to create a `krbtgt` principal. Substitute your cluster-dedicated KDC realm for `HADOOP.EXAMPLE.COM`, and substitute your central KDC realm for `EXAMPLE.COM`. Enter a trust

password when prompted. Note the password because you will need to enter the exact same password in the central KDC in the next step.

```
kadmin:  addprinc krbtgt/HADOOP.EXAMPLE.COM@EXAMPLE.COM
```

4. Each of your Hadoop client users must also place this information in their local `core-site.xml` file. The easiest way to do so is by using the Cloudera Manager Admin Console to generate a client configuration file.
5. To set up the cross-realm trust in the central KDC, type the same command in the `kadmin.local` or `kadmin` shell on the central KDC host to create the exact same `krbtgt` principal and password.

```
kadmin:  addprinc krbtgt/HADOOP.EXAMPLE.COM@EXAMPLE.COM
```

> **Important:** For a cross-realm trust to operate properly, both KDCs must have the same `krbtgt` principal and password, and both KDCs must be configured to use the same encryption type.

6. To properly translate principal names from the central KDC realm into the cluster-dedicated KDC realm for the Hadoop cluster, configure the **Trusted Kerberos Realms** property of the HDFS service.
   a. Open the Cloudera Manager Admin Console.
   b. Go to the HDFS service.
   c. Click the **Configuration** tab.
   d. Select **Scope** > **HDFS (Service Wide)**
   e. Select **Category** > **Security**.
   f. Type `Kerberos` in the **Search** box.
   g. Edit the **Trusted Kerberos Realms** property to add the name of your central KDC realm. If you need to use more advanced mappings which do more than just allow principals from another domain, you may enter them in the **Additional Rules to Map Kerberos Principals to Short Names** property. For more information about name mapping rules, see Mapping Kerberos Principals to Short Names on page 108.
7. Each of your Hadoop client users must also place this information in their local `core-site.xml` file. The easiest way to do so is by using the Cloudera Manager Admin Console to generate a client configuration file.
8. Later in this procedure, you will restart the services to have the configuration changes in `core-site.xml` take effect.

Using a Cluster-Dedicated MIT KDC with Active Directory

For Cloudera Manager clusters, the `openldap-clients` package must be installed on the Cloudera Manager Server host before configuring the cluster to use Kerberos for authentication.

On the Active Directory Server

1. On the Active Directory server host, type the following command to add the local realm trust to Active Directory:

```
netdom trust HADOOP.EXAMPLE.COM /Domain:EXAMPLE.COM /add /realm /passwordt:TrustPassword
```

2. On the Active Directory server host, type the following command to set the proper encryption type:

   Windows 2003 RC2

   Windows 2003 server installations do not support AES encryption for Kerberos. Therefore RC4 should be used. Please see the Microsoft reference documentation for more information.

```
ktpass /MITRealmName HADOOP.EXAMPLE.COM /TrustEncryp RC4
```

   Windows 2008

```
ksetup /SetEncTypeAttr HADOOP.EXAMPLE.COM <enc_type>
```

Where the <enc_type> parameter can be replaced with parameter strings for AES, DES, or RC4 encryption modes. For example, for AES encryption, replace <enc_type> with AES256-CTS-HMAC-SHA1-96 or AES128-CTS-HMAC-SHA1-96 and for RC4 encryption, replace with RC4-HMAC-MD5. See the Microsoft reference documentation for more information.

> **Important:** Make sure that the encryption type you specify is supported on both your version of Windows Active Directory and your version of MIT Kerberos.

On the MIT KDC Server

1. In the /var/kerberos/krb5kdc/kdc.conf file on the local dedicated KDC server host, configure the default realm for the Hadoop cluster by substituting your Kerberos realm in the following realms property:

```
[realms]
 HADOOP.EXAMPLE.COM = {
```

2. Each of your Hadoop client users must also place this information in their local core-site.xml file. The easiest way to do so is by using the Cloudera Manager Admin Console to generate a client configuration file.
3. On the local MIT KDC server host, type the following command in the kadmin.local or kadmin shell to add the cross-realm krbtgt principal:

```
kadmin:  addprinc -e "<keysalt_list>" krbtgt/HADOOP.EXAMPLE.COM@EXAMPLE.COM
```

where the <keysalt_list> parameter specifies the types of keys and their salt to be used for encryption of the password for this cross-realm krbtgt principal. It can be set to AES, or RC4 keytypes with a salt value of :normal. Note that DES is deprecated and should no longer be used. You can specify multiple keysalt types using the parameter in the command above. Make sure that at least one of the encryption types corresponds to the encryption types found in the tickets granted by the KDC in the remote realm. For an example of the values to use, see the examples based on the Active Directory functional domain level, below.

Examples by Active Directory Domain or Forest "Functional level"

Active Directory will, based on the Domain or Forest functional level, use encryption types supported by that release of the Windows Server operating system. It is not possible to use AES encryption types with an AD 2003 functional level. If you notice that DES encryption types are being used when authenticating or requesting service tickets to Active Directory then it might be necessary to enable weak encryption types in the /etc/krb5.conf. See Sample Kerberos Configuration Files on page 425 for an example.

- **Windows 2003**

```
kadmin: addprinc -e "rc4-hmac:normal" krbtgt/HADOOP.EXAMPLE.COM@EXAMPLE.COM
```

- **Windows 2008**

```
kadmin: addprinc -e "aes256-cts:normal aes128-cts:normal rc4-hmac:normal"
krbtgt/HADOOP.EXAMPLE.COM@EXAMPLE.COM
```

> **Note:** The cross-realm krbtgt principal that you add in this step must have *at least one entry* that uses the same encryption type as the tickets that are issued by the remote KDC. If there are no matching encryption types, principals in the local realm can successfully access the Hadoop cluster, but principals in the remote realm are unable to.

On All Cluster Hosts

1. In the `/etc/krb5.conf` file on all cluster hosts and all Hadoop client user hosts, configure both Kerberos realms. Note that `default_realm` should be configured as the local MIT Kerberos realm for the cluster. Your `krb5.conf` may contain more configuration properties than those demonstrated below. This example is provided to clarify configuration parameters. See Sample Kerberos Configuration Files on page 425 for more information.

```
[libdefaults]
  default_realm = HADOOP.EXAMPLE.COM
[realms]
  EXAMPLE.COM = {
    kdc = dc01.example.com:88
    admin_server = dc01.example.com:749
  }
  HADOOP.EXAMPLE.COM = {
    kdc = kdc01.hadoop.example.com:88
    admin_server = kdc01.hadoop.example.com:749
  }
[domain_realm]
  .hadoop.example.com = HADOOP.EXAMPLE.COM
  hadoop.example.com = HADOOP.EXAMPLE.COM
  .example.com = EXAMPLE.COM
  example.com = EXAMPLE.COM
```

2. Use one of the following methods to properly translate principal names from the Active Directory realm into the cluster-dedicated KDC realm for the Hadoop cluster.
   - **Using Cloudera Manager:** Configure the **Trusted Kerberos realms** property of the HDFS service:
     1. Open the Cloudera Manager Admin Console.
     2. Go to the HDFS service.
     3. Click the **Configuration** tab.
     4. Select **Scope** > **HDFS (Service Wide)**
     5. Select **Category** > **Security**.
     6. Type `Kerberos` in the **Search** box.
     7. Edit the **Trusted Kerberos Realms** property to add the name of your central KDC realm. If you need to use more advanced mappings which do more than just allow principals from another domain, you may enter them in the **Additional Rules to Map Kerberos Principals to Short Names** property. For more information about name mapping rules, see Mapping Kerberos Principals to Short Names on page 108.
   - **Using the Command Line:** Configure the `hadoop.security.auth_to_local` setting in the `core-site.xml` file on all of the cluster hosts. The following example translates all principal names with the realm `EXAMPLE.COM` into the first component of the principal name only. It also preserves the standard translation for the default realm (the cluster realm).

```
<property>
  <name>hadoop.security.auth_to_local</name>
  <value>
    RULE:[1:$1@$0](^.*@EXAMPLE\.COM$)s/^(.*)@EXAMPLE\.COM$/$1/g
    RULE:[2:$1@$0](^.*@EXAMPLE\.COM$)s/^(.*)@EXAMPLE\.COM$/$1/g
    DEFAULT
  </value>
</property>
```

Integrating MIT Kerberos and Active Directory

Several different subsystems are involved in servicing authentication requests, including the Key Distribution Center (KDC), Authentication Service (AS), and Ticket Granting Service (TGS). The more nodes in the cluster and the more services provided, the heavier the traffic between these services and the services running on the cluster.

As a general guideline, Cloudera recommends using a dedicated Active Directory instance (Microsoft Server Domain Services) for every 100-200 nodes in the cluster. However, this is just a loose guideline. Monitor utilization and deploy additional instances as needed to meet the demand.

By default, Kerberos uses UDP for client/server communication which is typically faster at delivering packets than TCP, but does not guarantee delivery. Additionally, using UDP packets that get too large are frequently dropped, as is the case when a user is a member of a large number of groups. To avoid this problem, force Kerberos to use TCP by modifying the Kerberos configuration file (`krb5.conf`) as follows:

```
[libdefaults]
udp_preference_limit = 1
...
```

This is especially useful if the domain controllers are not on the same subnet as the cluster or are separated by firewalls.

In general, Cloudera recommends setting up the Active Directory domain controller (Microsoft Server Domain Services) on the same subnet as the cluster and never over a WAN connection which results in considerable latency and affects cluster performance.

Troubleshooting cluster operations when Active Directory is being used for Kerberos authentication requires administrative access to the Microsoft Server Domain Services instance. Administrators may need to enable Kerberos event logging on the Microsoft Server KDC to resolve issues.

Deleting Cloudera Manager roles or nodes requires manually deleting the associate Active Directory accounts. Cloudera Manager cannot delete entries from Active Directory.

Integrating MIT Kerberos and Active Directory

Prior to release 5.1, clusters managed by Cloudera Manager could not integrate directly with a Microsoft Active Directory KDC. Rather, integrating the Cloudera Manager cluster with an Active Directory KDC required the additional setup of a local MIT KDC (local, meaning in the same subnet as the cluster).

> **Note:** The configuration process detailed below is not needed for Cloudera Manager clusters as of release 5.1 (and higher). To integrate the cluster with an Active Directory KDC, see Enabling Kerberos Authentication Using the Wizard on page 51.

The steps below can be used for Cloudera Manager clusters prior to release 5.1, or if the Active Directory KDC cannot be accessed directly for whatever reason. The setup process assumes that:

- An MIT Kerberos KDC is running in the same subnet as the cluster and that a Kerberos REALM is local to the cluster
- A Microsoft Server Active Directory instance (Microsoft Server Domain Services) is running elsewhere on the network, in its own Kerberos realm.

Given these two systems, you can then:

1. Create principals for all services running on the cluster in the MIT Kerberos realm local to the cluster.
2. Set up one-way cross-realm trust from the MIT Kerberos realm to the Active Directory realm, as detailed in Configuring a Local MIT Kerberos Realm to Trust Active Directory on page 169

The result of this setup is that Active Directory principals (users) can authenticate to the cluster without needing service principals.

Configuring a Local MIT Kerberos Realm to Trust Active Directory

On the Active Directory Server

1. Add the local realm trust to Active Directory with this command:

```
netdom trust YOUR-LOCAL-REALM.COMPANY.COM /Domain:AD-REALM.COMPANY.COM /add /realm
/passwordt:<TrustPassword>
```

2. Set the proper encryption type with this command:

On Windows 2003 RC2:

```
ktpass /MITRealmName YOUR-LOCAL-REALM.COMPANY.COM /TrustEncryp <enc_type>
```

On Windows 2008:

```
ksetup /SetEncTypeAttr YOUR-LOCAL-REALM.COMPANY.COM <enc_type>
```

The <enc_type> parameter specifies AES, DES, or RC4 encryption. Refer to the documentation for your version of Windows Active Directory to find the <enc_type> parameter string to use.

3. Get and verify the list of encryption types set with this command:

   On Windows 2008:

```
ksetup /GetEncTypeAttr YOUR-LOCAL-REALM.COMPANY.COM
```

> **Important:** Make sure the encryption type you specify is supported on both your version of Windows Active Directory and your version of MIT Kerberos.

On the MIT KDC Server

Type the following command in the kadmin.local or kadmin shell to add the cross-realm krbtgt principal. Use the same password you used in the `netdom` command on the Active Directory Server.

```
kadmin:  addprinc -e "enc_type_list"
krbtgt/YOUR-LOCAL-REALM.COMPANY.COM@AD-REALM.COMPANY.COM
```

where the *enc_type_list* parameter specifies the types of encryption this cross-realm krbtgt principal can support—AES, DES, or RC4. You can specify multiple encryption types using the parameter in the command above, what's important is that at least one of the encryption types corresponds to the encryption type found in the tickets granted by the KDC in the remote realm. For example:

```
kadmin:  addprinc -e "rc4-hmac:normal des3-hmac-sha1:normal"
krbtgt/YOUR-LOCAL-REALM.COMPANY.COM@AD-REALM.COMPANY.COM
```

> **Note:** The cross-realm krbtgt principal that you add in this step must have *at least one entry* that uses the same encryption type as the tickets that are issued by the remote KDC. If no entries have the same encryption type, then the problem you will see is that authenticating as a principal in the local realm will allow you to successfully run Hadoop commands, but authenticating as a principal in the remote realm will not allow you to run Hadoop commands.

On All of the Cluster Hosts

1. Verify that both Kerberos realms are configured on all of the cluster hosts. Note that the default realm and the domain realm should remain set as the MIT Kerberos realm which is local to the cluster.

```
[realms]
  AD-REALM.CORP.FOO.COM = {
    kdc = ad.corp.foo.com:88
    admin_server = ad.corp.foo.com:749
    default_domain = foo.com
  }
  CLUSTER-REALM.CORP.FOO.COM = {
    kdc = cluster01.corp.foo.com:88
    admin_server = cluster01.corp.foo.com:749
    default_domain = foo.com
  }
```

2. To properly translate principal names from the Active Directory realm into local names within Hadoop, you must configure the `hadoop.security.auth_to_local` setting in the `core-site.xml` file on all of the cluster machines. The following example translates all principal names with the realm `AD-REALM.CORP.FOO.COM` into the first component of the principal name only. It also preserves the standard translation for the default realm (the cluster realm).

```
<property>
  <name>hadoop.security.auth_to_local</name>
  <value>
    RULE:[1:$1@$0](^.*@AD-REALM\.CORP\.FOO\.COM$)s/^(.*)@AD-REALM\.CORP\.FOO\.COM$/$1/g

    RULE:[2:$1@$0](^.*@AD-REALM\.CORP\.FOO\.COM$)s/^(.*)@AD-REALM\.CORP\.FOO\.COM$/$1/g

    DEFAULT
  </value>
</property>
```

For more information about name mapping rules, see Mapping Kerberos Principals to Short Names on page 108.

Authorization

Authorization is concerned with who or what has access or control over a given resource or service. Since Hadoop merges together the capabilities of multiple varied, and previously separate IT systems as an enterprise data hub that stores and works on all data within an organization, it requires multiple authorization controls with varying granularities. In such cases, Hadoop management tools simplify setup and maintenance by:

- Tying all users to groups, which can be specified in existing LDAP or AD directories.
- Providing role-based access control for similar interaction methods, like batch and interactive SQL queries. For example, Apache Sentry permissions apply to Hive (HiveServer2) and Impala.

CDH currently provides the following forms of access control:

- Traditional POSIX-style permissions for directories and files, where each directory and file is assigned a single owner and group. Each assignment has a basic set of permissions available; file permissions are simply read, write, and execute, and directories have an additional permission to determine access to child directories.
- Extended Access Control Lists (ACLs) for HDFS that provide fine-grained control of permissions for HDFS files by allowing you to set different permissions for specific named users or named groups.
- Apache HBase uses ACLs to authorize various operations (`READ`, `WRITE`, `CREATE`, `ADMIN`) by column, column family, and column family qualifier. HBase ACLs are granted and revoked to both users and groups.
- Role-based access control with Apache Sentry. As of Cloudera Manager 5.1.x, Sentry permissions can be configured using either policy files or the database-backed Sentry service.
  - The Sentry service is the preferred way to set up Sentry permissions. See Managing the Sentry Service for more information.
  - For the policy file approach to configuring Sentry, see Sentry Policy File Authorization.

> **Important:** Cloudera does not support Apache Ranger or Hive's native authorization frameworks for configuring access control in Hive. Use Cloudera-supported Apache Sentry instead.

Cloudera Manager User Roles

Access to Cloudera Manager features is controlled by user accounts that specify an authentication mechanism and one or more user roles. User roles determine the tasks that an authenticated user can perform and the features visible to the user in the Cloudera Manager Admin Console. Documentation for Cloudera Manager administration and management tasks indicate user roles required to perform the task.

> **Note:** All possible user roles are available with Cloudera Enterprise. Cloudera Express provides Read-Only and Full Administrator user roles only. When a Cloudera Enterprise Data Hub Edition trial license expires, only users with Read-Only and Full Administrator roles can log in to Cloudera Manager. A Full Administrator must change user accounts with other roles to Read-Only or Full Administrator before such users can log in.

Displaying Roles for Current User Account Login

The user roles associated with a given login session are available at any time from the Cloudera Manager Admin Console menu. Assuming you are logged in to Cloudera Manager Admin Console, you can always verify the user roles associated with your current login as follows:

1. Select **My Profile** from the ***username*** drop-down menu, where *username* is the name of the logged in account (such as `admin`). The **My Profile** pop-up window displays the Username, Roles, and the date and time of the Last Successful Login.

2. Click **Close** to dismiss the message page.

User Roles

A Cloudera Manager user account can be assigned one of the following roles with associated permissions:

- **Auditor**
  - View configuration and monitoring information in Cloudera Manager.
  - View audit events.
- **Read-Only**
  - View configuration and monitoring information in Cloudera Manager.
  - View service and monitoring information.
  - View events and logs.
  - View replication jobs and snapshot policies.
  - View YARN applications and Impala queries.

  The Read-Only role does not allow the user to:

  - Add services or take any actions that affect the state of the cluster.
  - Use the HDFS file browser.
  - Use the HBase table browser.
  - Use the Solr Collection Statistics browser.
- **Dashboard**
  - Create, edit, or remove dashboards that belong to the user.
  - Add an existing chart or create a new chart to add to a dashboard that belongs to the user.
  - Perform the same tasks as the Read-Only role.
- **Limited Operator**
  - View configuration and monitoring information in Cloudera Manager.
  - View service and monitoring information.
  - Decommission hosts (except hosts running Cloudera Management Service roles).
  - Perform the same tasks as the Read-Only role.

  The Limited Operator role does not allow the user to add services or take any other actions that affect the state of the cluster.
- **Operator**
  - View configuration and monitoring information in Cloudera Manager.
  - View service and monitoring information.
  - Stop, start, and restart clusters, services (except the Cloudera Management Service), and roles.
  - Decommission and recommission hosts (except hosts running Cloudera Management Service roles).
  - Decommission and recommission roles (except Cloudera Management Service roles).
  - Start, stop, and restart KMS.
  - Perform the same tasks as the Read-Only role.

  The Operator role does not allow the user to add services, roles, or hosts, or take any other actions that affect the state of the cluster.
- **Configurator**
  - View configuration and monitoring information in Cloudera Manager.
  - Perform all Operator operations.
  - Configure services (except the Cloudera Management Service).
  - Enter and exit maintenance mode.

  - Manage dashboards (including Cloudera Management Service dashboards).
  - Start, stop, and restart KMS
  - Perform the same tasks as the Read-Only role.

- **Cluster Administrator** - Use all of the functionality available in Cloudera Manager and perform all actions *except* the following:
  - Administer Cloudera Navigator.
  - View replication schedules and snapshot policies.
  - View audit events.
  - Manage user accounts and configuration of external authentication.
  - Manage Full Administrator accounts.
  - Configure HDFS encryption, administer Key Trustee Server, and manage encryption keys.
  - Use the HDFS file browser, the HBase table browser, and the Solr Collection browser.
  - View the Directory Usage Report
  - View the HBase Statistics Page

  Unless otherwise noted above, the Cluster Administrator can view the data related to Cloudera Manager, such as file metadata. The Cluster Administrator cannot see things like the content of files stored by HDFS and other components.
- **BDR Administrator**
  - View configuration and monitoring information in Cloudera Manager.
  - View service and monitoring information.
  - Perform replication and define snapshot operations.
  - Use the HDFS file browser, the HBase table browser, and the Solr Collection browser.
  - View the Directory Usage Report
  - View the HBase Table Statistics Page
  - Perform the same tasks as the Read-Only role.
- **Navigator Administrator**
  - View configuration and monitoring information in Cloudera Manager.
  - View service and monitoring information.
  - Administer Cloudera Navigator.
  - View audit events.
  - Use the HDFS file browser, the HBase table browser, and the Solr Collection browser.
  - Perform the same tasks as the Read-Only role.
- **User Administrator**
  - View configuration and monitoring information in Cloudera Manager.
  - View service and monitoring information.
  - Manage user accounts and configuration of external authentication.
  - Use the HDFS file browser, the HBase table browser, and the Solr Collection browser.
  - Perform the same tasks as the Read-Only role.
- **Key Administrator**
  - View configuration and monitoring information in Cloudera Manager.
  - Configure HDFS encryption, administer Key Trustee Server, and manage encryption keys.
  - Start, stop, and restart KMS
  - Configure KMS ACLs
  - Use the HDFS file browser, the HBase table browser, and the Solr Collection browser.
  - Perform the same tasks as the Read-Only role.

- **Full Administrator** - Full Administrators have permissions to use all of the functionality available in Cloudera Manager and perform all actions on all clusters. Additionally, the Full Administrator can view the data related to Cloudera Manager, such as file metadata, snapshots, quotas, and file size. The Full Administrator cannot see things like the content of files stored by HDFS or other components.

Removing the Full Administrator User Role

Minimum Required Role: User Administrator (also provided by **Full Administrator**)

In some organizations, security policies may prohibit the use of the Full Administrator role. The Full Administrator role is created during Cloudera Manager installation, but you can remove it as long as you have at least one remaining user account with User Administrator privileges.

To remove the Full Administrator user role, perform the following steps.

1. Add at least one user account with User Administrator privileges, or ensure that at least one such user account already exists.
2. Ensure that there is only a single user account with Full Administrator privileges.
3. While logged in as the single remaining Full Administrator user, select your own user account and either delete it or assign it a new user role.

> **Warning:** After you delete the last Full Administrator account, you will be logged out immediately and will not be able to log in unless you have access to another user account. Also, it will no longer be possible to create or assign Full Administrators.

A consequence of removing the Full Administrator role is that some tasks may require collaboration between two or more users with different user roles. For example:

- If the machine that the Cloudera Navigator roles are running on needs to be replaced, the Cluster Administrator will want to move all the roles running on that machine to a different machine. The Cluster Administrator can move any non-Navigator roles by deleting and re-adding them, but would need a Navigator Administrator to perform the stop, delete, add, and start actions for the Cloudera Navigator roles.
- In order to take HDFS snapshots, snapshots must be enabled on the cluster by a Cluster Administrator, but the snapshots themselves must be taken by a BDR Administrator.

HDFS Extended ACLs

HDFS supports POSIX Access Control Lists (ACLs), as well as the traditional POSIX permissions model already supported. ACLs control access of HDFS files by providing a way to set different permissions for specific named users or named groups. They enhance the traditional permissions model by allowing users to define access control for arbitrary combinations of users and groups instead of a single owner/user or a single group.

Enabling HDFS Access Control Lists

By default, HDFS access control lists (ACLs) are disabled on a cluster. You can enable them using either Cloudera Manager or the command line.

The default ACL applies only to a directory, and all subdirectories and files created in that directory inherit the default ACL of the parent directory.

For example, if a directory has a default ACL entry of `default:group:analysts:rwx`, then all files created in the directory get a `group:analysts:rwx` entry, and subdirectories get both the default ACL and the access ACL copied over. To set a default ACL, simply prepend `default:` to the user or group entry in the `setfacl` command (see HDFS Extended ACL Example on page 177).

There is a maximum number/limit of 32 entries in a single HDFS extended ACL. If you attempt to add ACL entries exceeding the maximum of 32, you will get an error. An excessive number of ACL entries indicates that the requirements are better implemented by defining additional groups or users. ACLs with a very high number of entries also require additional memory and storage, and take longer to evaluate upon each permission check.

> **Important:** Ensure that all users and groups resolve on the NameNode for ACLs to work as expected.

Enabling HDFS ACLs Using Cloudera Manager

1. Go to the Cloudera Manager Admin Console and navigate to the **HDFS** service.
2. Click the **Configuration** tab.
3. Select **Scope** > ***Service_name*** **(Service-Wide)**
4. Select **Category** > **Security**
5. Locate the **Enable Access Control Lists** property and select its checkbox to enable HDFS ACLs.
6. Click **Save Changes** to commit the changes.

Enabling HDFS ACLs Using the Command Line

To enable ACLs using the command line, set the `dfs.namenode.acls.enabled` property to `true` in the NameNode's `hdfs-site.xml`.

```
<property>
<name>dfs.namenode.acls.enabled</name>
<value>true</value>
</property>
```

Commands

To set and get file access control lists (ACLs), use the file system shell commands, `setfacl` and `getfacl`.

getfacl

```
hdfs dfs -getfacl [-R] <path>

<!-- COMMAND OPTIONS
<path>: Path to the file or directory for which ACLs should be listed.
-R: Use this option to recursively list ACLs for all files and directories.
-->
```

Examples:

```
<!-- To list all ACLs for the file located at /user/hdfs/file -->
hdfs dfs -getfacl /user/hdfs/file

<!-- To recursively list ACLs for /user/hdfs/file -->
hdfs dfs -getfacl -R /user/hdfs/file
```

setfacl

```
hdfs dfs -setfacl [-R] [-b|-k -m|-x <acl_spec> <path>]|[--set <acl_spec> <path>]

<!-- COMMAND OPTIONS
<path>: Path to the file or directory for which ACLs should be set.
-R: Use this option to recursively list ACLs for all files and directories.
-b: Revoke all permissions except the base ACLs for user, groups and others.
-k: Remove the default ACL.
-m: Add new permissions to the ACL with this option. Does not affect existing permissions.
-x: Remove only the ACL specified.
<acl_spec>: Comma-separated list of ACL permissions.
--set: Use this option to completely replace the existing ACL for the path specified.
       Previous ACL entries will no longer apply.
-->
```

Examples:

```
<!-- To give user ben read & write permission over /user/hdfs/file -->
hdfs dfs -setfacl -m user:ben:rw- /user/hdfs/file

<!-- To remove user alice's ACL entry for /user/hdfs/file -->
hdfs dfs -setfacl -x user:alice /user/hdfs/file

<!-- To give user hadoop read & write access, and group or others read-only access -->
hdfs dfs -setfacl --set user::rw-,user:hadoop:rw-,group::r--,other::r-- /user/hdfs/file
```

For more information on using HDFS ACLs, see the HDFS Permissions Guide on the Apache website.

HDFS Extended ACL Example

This example demonstrates how a user ("alice"), shares folder access with colleagues from another team ("hadoopdev"), so that the hadoopdev team can collaborate on the content of that folder; this is accomplished by updating the default extended ACL of that directory:

1. Make the files and sub-directories created within the content directory readable by team "hadoopdev":

```
$ hdfs dfs -setfacl -m group:hadoopdev:r-x /project
```

2. Set the default ACL setting for the parent directory:

```
$ hdfs dfs -setfacl -m default:group:hadoopdev:r-x /project
```

3. Create a sub-directory for the content you wish to share:

```
$ hdfs dfs -mkdir /project/dev
```

4. Inspect the new sub-directory ACLs to verify that HDFS has applied the new default values:

```
$ hdfs dfs -getfacl -R /project

file: /project
owner: alice
group: appdev
user::rwx
group::r-x
other::r-x
default:user::rwx
default:group::r-x
default:group:hadoopdev:r-x
default:mask::r-x
default:other::r-x

file: /project/dev
owner: alice
group: appdev
user::rwx
group::r-x
group:hadoopdev:r-x
mask::r-x
other::r-x
default:user::rwx
default:group::r-x
default:group:hadoopdev:r-x
default:mask::r-x
default:other::r-x
```

Note: At the time it is created, the default ACL is copied from the parent directory to the child directory. Subsequent changes to the parent directory default ACL do not change the ACLs of the existing child directories.

Enabling Authorization for HDFS Web UIs

You can enforce authorization for the following HDFS web UIs: the NameNode, DataNode, and JournalNode. To do so, you must have Kerberos authentication for HTTP web consoles and Hadoop Secure Authorization enabled. When both configurations are set, only the `hdfs` user can access the HDFS web UIs by default. Any other user who attempts to access the web UI will encounter an HTTP 403 error because the user is not authorized to access the page.

For users and groups other than `hdfs` to access the web UIs, you must add them to `hdfs-site.xml` with an **HDFS Service Advanced Configuration Snippet (Safety Valve)**.

Perform the following steps to enforce authorization for the HDFS web UIs:

1. In the **Cloudera Manager Admin Console**, go to **Clusters** > ***<HDFS service>***.
2. Navigate to the **Configurations** tab and search for the following property: `HDFS Service Advanced Configuration Snippet (Safety Valve) for hdfs-site.xml`.
3. Add the following property:
   - **Name:** `dfs.cluster.administrators`
   - **Value:** *`<comma separated list of user names and/or group names>`*

   For example, a sample property might look like this:

   - **Name:** `dfs.cluster.administrators`
   - **Value:** `hdfs,admin_user_milton,HTTP,hue,admin_group`

   These values would allow the users `hdfs`, `admin_user_milton`, `HTTP`, and hue as well as the group `admin_group` to the following web UIs: NameNode, DataNode, and JournalNode.

   If you perform the steps under Additional Configuration on page 178 to restrict access to the /jmx, /stack, /conf, and /metrics servlets, you must add the `HTTP` user and the Service Monitor Kerberos Principal so that Cloudera Manager can access the /jmx and /metrics servlets.

   You can view the Service Monitor Kerberos Principal by navigating to **Cloudera Management Service** > **Configuration** and searching for `Role-Specific Kerberos Principal`. The default Service Monitor Kerberos Principal is `hue`.

4. Save the configuration.
5. Restart all stale HDFS services.

Additional Configuration

For a higher level of security, you can enforce authorization for the following HDFS web UI servlets, which may contain sensitive data: /jmx, /stack, /conf, and /metrics. When you enforce authorization for the servlets, only the users listed in the `dfs.cluster.administrators` property can access them.

Cloudera Manager requires access to the /jmx and /metrics servlets and uses the HTTP user as well as the Service Monitor Kerberos Principal to access them. Make sure to add both users to `dfs.cluster.administrators` as described in Enabling Authorization for HDFS Web UIs on page 178.

Perform the following steps to enforce authorization for the servlets:

1. In the **Cloudera Manager Admin Console**, go to **Clusters** > ***<HDFS service>***.
2. Navigate to the **Configurations** tab and search for the following property: `HDFS Service Advanced Configuration Snippet (Safety Valve) for hdfs-site.xml`.
3. Add the following property:
   - **Name:** `hadoop.security.instrumentation.requires.admin`
   - **Value:** `true`

4. Save the configuration.
5. Restart all stale HDFS services.

Configuring LDAP Group Mappings

Each host that comprises a node in a Cloudera cluster runs an operating system, such as CentOS or Oracle Linux. At the OS-level, there are `user:group` accounts created during installation that map to the services running on that specific node of the cluster. The default shell-based group mapping provider, `org.apache.hadoop.security.ShellBasedUnixGroupsMapping`, handles the mapping from the local host system (the OS) to the specific cluster service, such as HDFS. The hosts authenticate using these local OS accounts before processes are allowed to run on the node.

For clusters integrated with Kerberos for authentication, the hosts must also provide Kerberos tickets before processes can run on the node. The cluster can use the organization's LDAP directory service to provide the login credentials, including Kerberos tickets, to authenticate transparently while the system runs. That means that the local `user:group` accounts on each host must be mapped to LDAP accounts. To map local `user:group` accounts to an LDAP service:

- Use tools such as SSSD (Systems Security Services Daemon) or Centrify Server Suite (see Identity and Access management for Cloudera).
- The Hadoop `LdapGroupsMapping` group mapping mechanism. The `LdapGroupsMapping` library may not be as robust a solution needed for large organizations in terms of scalability and manageability, especially for organizations managing identity across multiple systems and not exclusively for Hadoop clusters. Support for the `LdapGroupsMapping` library is not consistent across all operating systems.
- Do not use Winbind to map Linux `user:group` accounts to Active Directory. It cannot scale, impedes cluster performance, and is not supported.
- Use the same `user:group` mappings across all cluster nodes, for ease of management.
- Use either Cloudera Manager or the command-line process detailed below.

The local `user:group` accounts must be mapped to LDAP for group mappings in Hadoop. You must create the users and groups for your Hadoop services in LDAP.

To integrate the cluster with an LDAP service, the `user:group` relationships must be contained in the LDAP directory. The admin must create the user accounts and define groups for `user:group` relationships on each host.

The user and group names listed in the table are the default `user:group` values for CDH services.

> **Note:** If the defaults have been changed for any service, use the custom values to create the users and configure the group for that service in the LDAP server, rather than the defaults listed in the table below. For example, you changed the defaults in the Cloudera Manager Admin Console to customize the **System User** or **System Group** setting for the service.

Cloudera Product or Component	User	Group
Cloudera Manager	`cloudera-scm`	`cloudera-scm`
Apache Accumulo	`accumulo`	`accumulo`
Apache Avro	(No default)	(No default)
Apache Flume (sink writing to HDFS needs write privileges)	`flume`	`flume`
Apache HBase (Master, RegionServer processes)	`hbase`	`hbase`
Apache HCatalog, WebHCat service	`hive`	`hive`
Apache Hive (HiveServer2, Hive Metastore)	`hive`	`hive`
Apache Kafka	`kafka`	`kafka`
Apache Mahout	(No default)	(No default)
Apache Oozie	`oozie`	`oozie`
Apache Pig	(No default)	(No default)

Cloudera Product or Component	User	Group
Apache Sentry	`sentry`	`sentry`
Apache Spark	`spark`	`spark`
Apache Sqoop1	`sqoop`	`sqoop`
Apache Sqoop2	`sqoop2`	`sqoop, sqoop2`
Apache Whirr	(No default)	(No default)
Apache ZooKeeper	`zookeeper`	`zookeeper`
Impala	`impala`	`impala, hive`
Cloudera Search	`solr`	`solr`
HDFS (NameNode, DataNodes)	`hdfs`	`hdfs, hadoop`
HttpFS	`httpfs`	`httpfs`
Hue	`hue`	`hue`
Hue Load Balancer (needs apache2 package)	`apache`	`apache`
Java KeyStore KMS	`kms`	`kms`
Key Trustee KMS	`kms`	`kms`
Key Trustee Server	`keytrustee`	`keytrustee`
Kudu	`kudu`	`kudu`
Llama	`llama`	`llama`
MapReduce (JobTracker, TaskTracker)	`mapred`	`mapred, hadoop`
Parquet	(No default)	(No default)
YARN	`yarn`	`yarn, hadoop`

Note: Cloudera Manager 5.3 (and later releases) can be deployed in ***single user mode***. In single user mode, Hadoop users and groups are subsumed by `cloudera-scm:cloudera-scm`. Cloudera Manager starts all Cloudera Manager Agent processes and services running on the nodes in the cluster as a unit owned by this `cloudera-scm:cloudera-scm`. Single user mode is not recommended for production clusters.

In addition:

- For Sentry with Hive, add these properties on the HiveServer2 node.
- For Sentry with Impala, add these properties to all hosts.

See Users and Groups in Sentry for more information.

Using Cloudera Manager

Minimum Required Role: Configurator (also provided by **Cluster Administrator, Full Administrator**)

Make the following changes to the HDFS service's security configuration:

1. Open the Cloudera Manager Admin Console and go to the **HDFS** service.
2. Click the **Configuration** tab.
3. Select **Scope** > **HDFS (Service Wide)**
4. Select **Category** > **Security**.
5. Modify the following configuration properties using values from the table below:

Configuration Property	Value
Hadoop User Group Mapping Implementation	`org.apache.hadoop.security.LdapGroupsMapping`
Hadoop User Group Mapping LDAP URL	`ldap://<server>`
Hadoop User Group Mapping LDAP Bind User	`Administrator@example.com`
Hadoop User Group Mapping LDAP Bind User Password	`***`
Hadoop User Group Mapping Search Base	`dc=example,dc=com`

Although the above changes are sufficient to configure group mappings for Active Directory, some changes to the remaining default configurations might be required for OpenLDAP.

Using the Command Line

Add the following properties to the `core-site.xml` on the NameNode:

```
<property>
<name>hadoop.security.group.mapping</name>
<value>org.apache.hadoop.security.LdapGroupsMapping</value>
</property>

<property>
<name>hadoop.security.group.mapping.ldap.url</name>
<value>ldap://server</value>
</property>

<property>
<name>hadoop.security.group.mapping.ldap.bind.user</name>
<value>Administrator@example.com</value>
</property>

<property>
<name>hadoop.security.group.mapping.ldap.bind.password</name>
<value>****</value>
</property>

<property>
<name>hadoop.security.group.mapping.ldap.base</name>
<value>dc=example,dc=com</value>
</property>

<property>
<name>hadoop.security.group.mapping.ldap.search.filter.user</name>
<value>(&amp;(objectClass=user)(sAMAccountName={0}))</value>
</property>

<property>
<name>hadoop.security.group.mapping.ldap.search.filter.group</name>
<value>(objectClass=group)</value>
</property>

<property>
<name>hadoop.security.group.mapping.ldap.search.attr.member</name>
<value>member</value>
</property>

<property>
<name>hadoop.security.group.mapping.ldap.search.attr.group.name</name>
<value>cn</value>
</property>
```

Authorization With Apache Sentry

Apache Sentry is a granular, role-based authorization module for Hadoop. Sentry provides the ability to control and enforce precise levels of privileges on data for authenticated users and applications on a Hadoop cluster. Sentry currently works out of the box with Apache Hive, Hive Metastore/HCatalog, Apache Solr, Impala, and HDFS (limited to Hive table data).

Sentry is designed to be a pluggable authorization engine for Hadoop components. It allows you to define authorization rules to validate a user or application's access requests for Hadoop resources. Sentry is highly modular and can support authorization for a wide variety of data models in Hadoop.

For information about configuring and managing the Sentry service, see the Apache Sentry Guide.

Continue reading:

Architecture Overview

Sentry Components

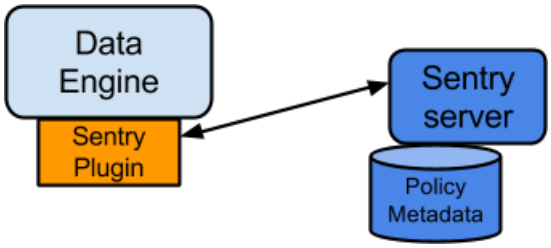

There are three components involved in the authorization process:

- **Sentry Server**

  The Sentry RPC server manages the authorization metadata. It supports interfaces to securely retrieve and manipulate the metadata. In CDH 5.13 and above, you can configure multiple Sentry Servers for high availability. See Sentry High Availability for information about how to enable high availability for Sentry.

- **Data Engine**

  This is a data processing application, such as Hive or Impala, that needs to authorize access to data or metadata resources. The data engine loads the Sentry plugin and all client requests for accessing resources are intercepted and routed to the Sentry plugin for validation.

- **Sentry Plugin**

  The Sentry plugin runs in the data engine. It offers interfaces to manipulate authorization metadata stored in the Sentry Server, and includes the authorization policy engine that evaluates access requests using the authorization metadata retrieved from the server.

Key Concepts

- Authentication - Verifying credentials to reliably identify a user
- Authorization - Limiting the user's access to a given resource
- User - Individual identified by underlying authentication system
- Group - A set of users, maintained by the authentication system
- Privilege - An instruction or rule that allows access to an object
- Role - A set of privileges; a template to combine multiple access rules
- Authorization models - Defines the objects to be subject to authorization rules and the granularity of actions allowed. For example, in the SQL model, the objects can be databases or tables, and the actions are `SELECT`, `INSERT`, and `CREATE`. For the Search model, the objects are indexes, configs, collections, documents; the access modes include query and update.

User Identity and Group Mapping

Sentry relies on underlying authentication systems, such as Kerberos or LDAP, to identify the user. It also uses the group mapping mechanism configured in Hadoop to ensure that Sentry sees the same group mapping as other components of the Hadoop ecosystem.

Consider a sample organization with users Alice and Bob who belong to an Active Directory (AD) group called `finance-department`. Bob also belongs to a group called `finance-managers`. In Sentry, you first create roles and then grant privileges to these roles. For example, you can create a role called Analyst and grant `SELECT` on tables Customer and Sales to this role.

The next step is to join these authentication entities (users and groups) to authorization entities (roles). This can be done by granting the Analyst role to the `finance-department` group. Now Bob and Alice who are members of the `finance-department` group get `SELECT` privilege to the Customer and Sales tables.

Role-Based Access Control

Role-based access control (RBAC) is a powerful mechanism to manage authorization for a large set of users and data objects in a typical enterprise. New data objects get added or removed, users join, move, or leave organisations all the time. RBAC makes managing this a lot easier. Hence, as an extension of the sample organization discussed previously, if a new employee Carol joins the Finance Department, all you need to do is add her to the `finance-department` group in AD. This will give Carol access to data from the Sales and Customer tables.

Unified Authorization

Another important aspect of Sentry is the unified authorization. The access control rules once defined, work across multiple data access tools. For example, being granted the Analyst role in the previous example will allow Bob, Alice, and others in the `finance-department` group to access table data from SQL engines such as Hive and Impala, as well as using MapReduce, Pig applications or metadata access using HCatalog.

Sentry Integration with the Hadoop Ecosystem

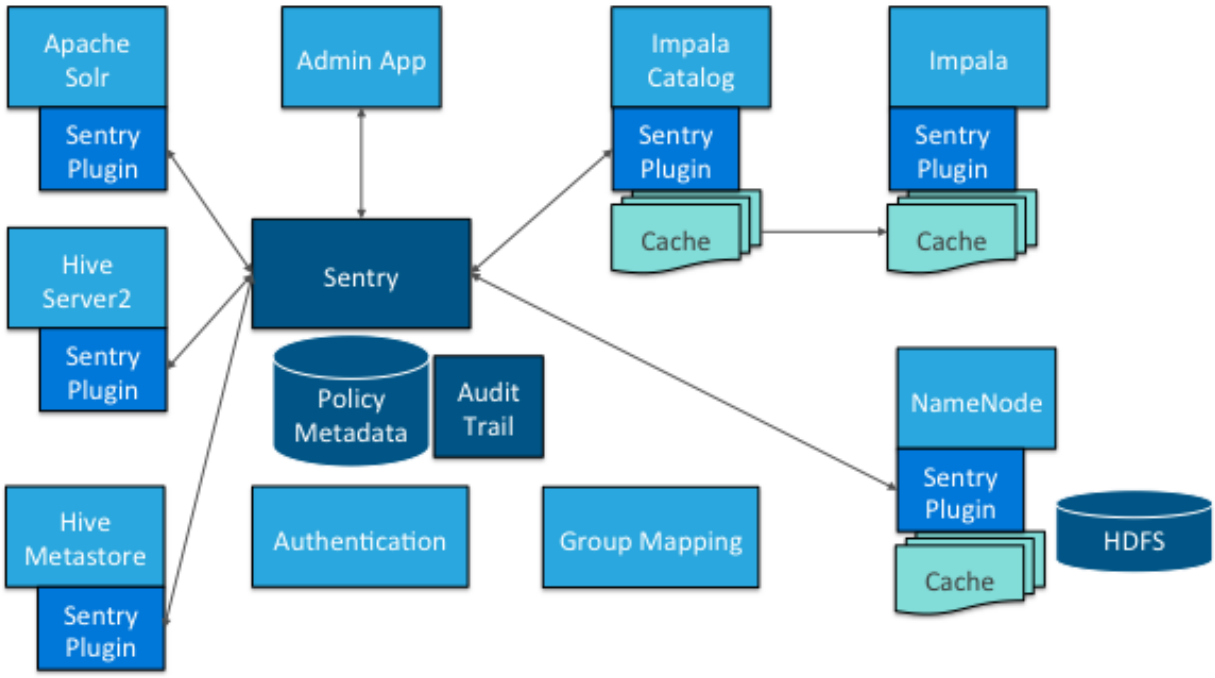

As illustrated above, Apache Sentry works with multiple Hadoop components. At the heart, you have the Sentry Server which stores authorization metadata and provides APIs for tools to retrieve and modify this metadata securely.

Note that the Sentry Server only facilitates the metadata. The actual authorization decision is made by a policy engine that runs in data processing applications such as Hive or Impala. Each component loads the Sentry plugin, which includes the service client for dealing with the Sentry service and the policy engine to validate the authorization request.

Hive and Sentry

Consider an example where Hive gets a request to access an object in a certain mode by a client. If Bob submits the following Hive query:

```
select * from production.sales
```

Hive will identify that user Bob is requesting SELECT access to the Sales table. At this point Hive will ask the Sentry plugin to validate Bob's access request. The plugin will retrieve Bob's privileges related to the Sales table and the policy engine will determine if the request is valid.

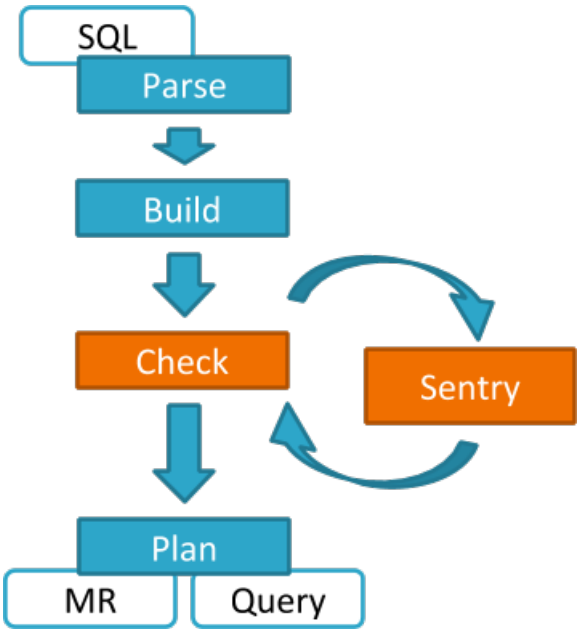

Hive works with both the Sentry service and policy files. Cloudera recommends you use the Sentry service, which makes it easier to manage user privileges. For more details and instructions, see Managing the Sentry Service.

Impala and Sentry

Authorization processing in Impala is similar to that in Hive. The main difference is caching of privileges. Impala's Catalog server manages caching schema metadata and propagating it to all Impala server nodes. This Catalog server caches Sentry metadata as well. As a result, authorization validation in Impala happens locally and much faster.

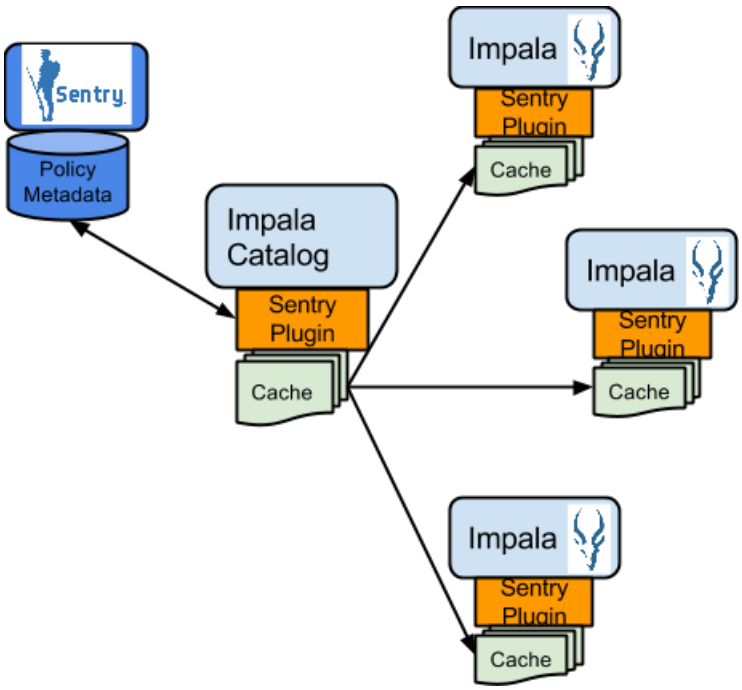

For detailed documentation, see Enabling Sentry Authorization for Impala.

Sentry-HDFS Synchronization

Sentry-HDFS authorization is focused on Hive warehouse data - that is, any data that is part of a table in Hive or Impala. The real objective of this integration is to expand the same authorization checks to Hive warehouse data being accessed from any other components such as Pig, MapReduce or Spark. At this point, this feature does not replace HDFS ACLs. Tables that are not associated with Sentry will retain their old ACLs.

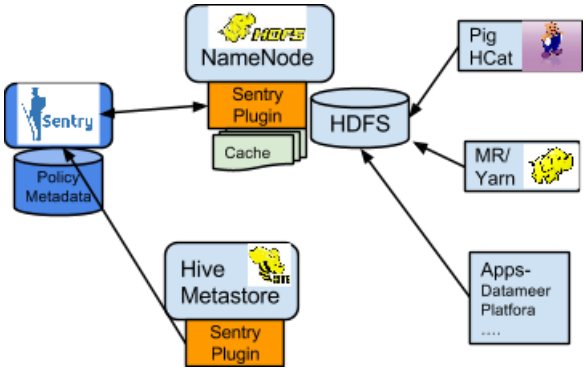

The mapping of Sentry privileges to HDFS ACL permissions is as follows:

- SELECT privilege -> Read access on the file.
- INSERT privilege -> Write access on the file.
- ALL privilege -> Read and Write access on the file.

The NameNode loads a Sentry plugin that caches Sentry privileges as well Hive metadata. This helps HDFS to keep file permissions and Hive tables privileges in sync. The Sentry plugin periodically polls Sentry to keep the metadata changes in sync.

For example, if Bob runs a Pig job that is reading from the Sales table data files, Pig will try to get the file handle from HDFS. At that point the Sentry plugin on the NameNode will figure out that the file is part of Hive data and overlay Sentry privileges on top of the file ACLs. As a result, HDFS will enforce the same privileges for this Pig client that Hive would apply for a SQL query.

For HDFS-Sentry synchronization to work, you *must* use the Sentry service, not policy file authorization. See Synchronizing HDFS ACLs and Sentry Permissions, for more details.

Search and Sentry

Sentry can apply restrictions to various Search tasks including accessing data and creating collections. These restrictions are consistently applied, regardless of the way users attempt to complete actions. For example, restricting access to data in a collection restricts that access whether queries come from the command line, from a browser, or through the admin console.

With Search, Sentry restrictions can be stored in the database-backed Sentry service or in a policy file (for example, sentry-provider.ini) which is stored in an HDFS location such as `hdfs://ha-nn-uri/user/solr/sentry/sentry-provider.ini`.

Sentry with Search does not support multiple policy files for multiple databases. If you choose to use policy files rather than database-backed Sentry service, you must use a separate policy file for each Sentry-enabled service. For example, if Hive and Search were using policy file authorization, using a combined Hive and Search policy file would result in an invalid configuration and failed authorization on both services.

Search works with both the Sentry service and policy files. Cloudera recommends you use the Sentry service, which makes it easier to manage user privileges. For more details and instructions, see Managing the Sentry Service *or* Sentry Policy File Authorization.

For detailed documentation, see Configuring Sentry Authorization for Cloudera Search.

Authorization Administration

The Sentry Server supports APIs to securely manipulate roles and privileges. Both Hive and Impala support SQL statements to manage privileges natively. Sentry assumes that HiveServer2 and Impala run as superusers, usually called `hive` and `impala`. To initiate top-level permissions for Sentry, an admin must login as a superuser. You can use either Beeline or the Impala shell to execute the following sample statement:

```
GRANT ROLE Analyst TO GROUP finance_managers
```

Disabling Hive CLI

To execute Hive queries, you must use Beeline. Hive CLI is not supported with Sentry and therefore its access to the Hive Metastore must be disabled. This is especially necessary if the Hive metastore has sensitive metadata. To do this, set the **Hive Metastore Access Control and Proxy User Groups Override** property for the Hive service in Cloudera Manager. For example, to give the `hive` user permission to impersonate only members of the `hive` and hue groups, set the property to: `hive,  hue`

If other user groups require access to the Hive Metastore, they can be added to the comma-separated list as needed. For example, setting this property to `hive,  hue` blocks the Spark shell from accessing the metastore.

Using Hue to Manage Sentry Permissions

Hue supports a Security app to manage Sentry authorization. This allows users to explore and change table permissions. Here is a video blog that demonstrates its functionality.

Configuring HBase Authorization

After configuring HBase authentication (as detailed in HBase Configuration), you must define rules on resources that is allowed to access. HBase rules can be defined for individual tables, columns, and cells within a table. Cell-level authorization is fully supported since CDH 5.2.

> **Important:** In a cluster managed by Cloudera Manager, HBase authorization is disabled by default. You have to enable HBase authorization (as detailed in Enable HBase Authorization on page 189) to use any kind of authorization method.

Understanding HBase Access Levels

HBase access levels are granted independently of each other and allow for different types of operations at a given scope.

- **Read (R)** - can read data at the given scope
- **Write (W)** - can write data at the given scope
- **Execute (X)** - can execute coprocessor endpoints at the given scope
- **Create (C)** - can create tables or drop tables (even those they did not create) at the given scope
- **Admin (A)** - can perform cluster operations such as balancing the cluster or assigning regions at the given scope

The possible scopes are:

- **Superuser** - superusers can perform any operation available in HBase, to any resource. The user who runs HBase on your cluster is a superuser, as are any principals assigned to the configuration property `hbase.superuser` in `hbase-site.xml` on the HMaster.
- **Global** - permissions granted at `global` scope allow the admin to operate on all tables of the cluster.
- **Namespace** - permissions granted at `namespace` scope apply to all tables within a given namespace.
- **Table** - permissions granted at `table` scope apply to data or metadata within a given table.
- **ColumnFamily** - permissions granted at `ColumnFamily` scope apply to cells within that ColumnFamily.
- **Cell** - permissions granted at Cell scope apply to that exact cell coordinate. This allows for policy evolution along with data. To change an ACL on a specific cell, write an updated cell with new ACL to the precise coordinates of the original. If you have a multi-versioned schema and want to update ACLs on all visible versions, you'll need to write new cells for all visible versions. The application has complete control over policy evolution. The exception is `append` and `increment` processing. `Appends` and `increments` can carry an ACL in the operation. If one is included in the operation, then it will be applied to the result of the `append` or `increment`. Otherwise, the ACL of the existing cell being appended to or incremented is preserved.

The combination of access levels and scopes creates a matrix of possible access levels that can be granted to a user. In a production environment, it is useful to think of access levels in terms of what is needed to do a specific job. The following list describes appropriate access levels for some common types of HBase users. It is important not to grant more access than is required for a given user to perform their required tasks.

- **Superusers** - In a production system, only the HBase user should have superuser access. In a development environment, an administrator might need superuser access to quickly control and manage the cluster. However, this type of administrator should usually be a `Global Admin` rather than a superuser.
- **Global Admins** - A `global admin` can perform tasks and access every table in HBase. In a typical production environment, an admin should not have `Read` or `Write` permissions to data within tables.
  - A global admin with `Admin` permissions can perform cluster-wide operations on the cluster, such as balancing, assigning or unassigning regions, or calling an explicit major compaction. This is an operations role.
  - A global admin with `Create` permissions can create or drop any table within HBase. This is more of a DBA-type role.

  In a production environment, it is likely that different users will have only one of `Admin` and `Create` permissions.

> **Warning:**
>
> In the current implementation, a `Global Admin` with `Admin` permission can grant himself `Read` and `Write` permissions on a table and gain access to that table's data. For this reason, only grant `Global Admin` permissions to trusted user who actually need them.
>
> Also be aware that a `Global Admin` with `Create` permission can perform a `Put` operation on the ACL table, simulating a `grant` or `revoke` and circumventing the authorization check for `Global Admin` permissions. This issue (but not the first one) is fixed in CDH 5.3 and higher, as well as CDH 5.2.1.
>
> Due to these issues, be cautious with granting Global Admin privileges.

- **Namespace Admin** - a namespace admin with Create permissions can create or drop tables within that namespace, and take and restore snapshots. A namespace admin with Admin permissions can perform operations such as splits or major compactions on tables within that namespace. Prior to CDH 5.4, only global admins could create namespaces. In CDH 5.4, any user with Namespace Create privileges can create namespaces.
- **Table Admins** - A table admin can perform administrative operations only on that table. A table admin with `Create` permissions can create snapshots from that table or restore that table from a snapshot. A table admin with `Admin` permissions can perform operations such as splits or major compactions on that table.
- **Users** - Users can read or write data, or both. Users can also execute coprocessor endpoints, if given `Executable` permissions.

> **Important:**
>
> If you are using Kerberos principal names when setting ACLs for users, Hadoop uses only the first part (short) of the Kerberos principal when converting it to the username. Hence, for the principal `ann/fully.qualified.domain.name@YOUR-REALM.COM`, HBase ACLs should only be set for user `ann`.

The following table shows some typical job descriptions at a hypothetical company and the permissions they might require to get their jobs done using HBase.

Table 10: Real-World Example of Access Levels

Job Title	Scope	Permissions	Description
Senior Administrator	Global	Admin, Create	Manages the cluster and gives access to Junior Administrators.
Junior Administrator	Global	Create	Creates tables and gives access to Table Administrators.
Table Administrator	Table	Admin	Maintains a table from an operations point of view.
Data Analyst	Table	Read	Creates reports from HBase data.
Web Application	Table	Read, Write	Puts data into HBase and uses HBase data to perform operations.

Further Reading

- Access Control Matrix

- Security - Apache HBase Reference Guide

Enable HBase Authorization

HBase authorization is built on top of the Coprocessors framework, specifically `AccessController Coprocessor`.

> **Note:** Once the Access Controller coprocessor is enabled, any user who uses the HBase shell will be subject to access control. Access control will also be in effect for native (Java API) client access to HBase.

Enable HBase Authorization Using Cloudera Manager

1. Go to **Clusters** and select the HBase cluster.
2. Select **Configuration**.
3. Search for **HBase Secure Authorization** and select it.
4. Search for **HBase Service Advanced Configuration Snippet (Safety Valve) for hbase-site.xml** and enter the following into it to enable `hbase.security.exec.permission.checks`. Without this option, all users will continue to have access to execute endpoint coprocessors. This option is not enabled when you enable HBase Secure Authorization for backward compatibility.

```
<property>
  <name>hbase.security.exec.permission.checks</name>
  <value>true</value>
</property>
```

5. Optionally, search for and configure **HBase Coprocessor Master Classes** and **HBase Coprocessor Region Classes**.

Enable HBase Authorization Using the Command Line

> **Important:**
> - Follow these command-line instructions on systems that do not use Cloudera Manager.
> - This information applies specifically to CDH 5.14.x. See Cloudera Documentation for information specific to other releases.

To enable HBase authorization, add the following properties to the `hbase-site.xml` file *on every HBase server host (Master or RegionServer)*:

```
<property>
     <name>hbase.security.authorization</name>
     <value>true</value>
</property>
<property>
  <name>hbase.security.exec.permission.checks</name>
  <value>true</value>
</property>
<property>
     <name>hbase.coprocessor.master.classes</name>
     <value>org.apache.hadoop.hbase.security.access.AccessController</value>
</property>
<property>
     <name>hbase.coprocessor.region.classes</name>

<value>org.apache.hadoop.hbase.security.token.TokenProvider,org.apache.hadoop.hbase.security.access.AccessController</value>
</property>
```

Configure Access Control Lists for Authorization

Now that HBase has the security coprocessor enabled, you can set ACLs using the HBase shell. Start the HBase shell as usual.

> **Important:**
>
> The host running the shell must be configured with a keytab file as described in Configuring Kerberos Authentication for HBase.

The commands that control ACLs take the following form. Group names are prefixed with the @ symbol.

```
hbase> grant <user> <permissions> [ @<namespace> [ <table>[ <column family>[ <column
qualifier> ] ] ] ]   # grants permissions

hbase> revoke <user> [ @<namespace> [ <table> [ <column family> [ <column qualifier> ]
 ] ]  # revokes permissions

hbase> user_permission <table>
                    # displays existing permissions
```

In the above commands, fields encased in <> are variables, and fields in `[]` are optional. The `permissions` variable must consist of zero or more character from the set "RWCA".

- `R` denotes read permissions, which is required to perform `Get`, `Scan`, or `Exists` calls in a given scope.
- `W` denotes write permissions, which is required to perform `Put`, `Delete`, `LockRow`, `UnlockRow`, `IncrementColumnValue`, `CheckAndDelete`, `CheckAndPut`, `Flush`, or `Compact` in a given scope.
- `X` denotes execute permissions, which is required to execute coprocessor endpoints.
- `C` denotes create permissions, which is required to perform `Create`, `Alter`, or `Drop` in a given scope.
- `A` denotes admin permissions, which is required to perform `Enable`, `Disable`, `Snapshot`, `Restore`, `Clone`, `Split`, `MajorCompact`, `Grant`, `Revoke`, and `Shutdown` in a given scope.

Access Control List Example Commands

```
grant 'user1', 'RWC'
grant 'user2', 'RW', 'tableA'
grant 'user3', 'C', '@my_namespace'
```

Be sure to review the information in Understanding HBase Access Levels on page 187 to understand the implications of the different access levels.

Configure Cell_Level Access Control Lists

If you wish to enable cell-level ACLs for HBase, then you must modify the default values for the following properties:

```
hbase.security.exec.permission.checks => true (the default value is false)
hbase.security.access.early_out => false (the default value is true)
hfile.format.version => 3 (the default value is 2)
```

Unless you modify the default properties as specified (or via the service-wide **HBase Service Advanced Configuration Snippet (Safety Valve) for hbase-site.xml**, which requires a service restart), then cell level ACLs will not work.

The following example shows how to grant (or revoke) HBase permissions (in this case, `read` permission) at the cell-level via an ACL:

```
grant 'Employee', { 'employe.name' => 'R' }, { COLUMNS => [ 'pd' ], FILTER =>
"(PrefixFilter ('T'))" }
```

Encrypting Data in Transit

Transport Layer Security (TLS) is an industry standard set of cryptographic protocols for securing communications over a network. TLS evolved from Secure Sockets Layer (SSL). Because the *SSL* terminology is still widely used, Cloudera software and documentation refer to TLS as **TLS/SSL**, but the actual protocol used is TLS. SSL is not used in Cloudera software.

In addition to TLS/SSL encryption, HDFS and HBase transfer data using remote procedure calls (RPCs). To secure this transfer, you must enable RPC encryption.

For instructions on enabling TLS/SSL and RPC encryption, see the following topics:

TLS/SSL and Its Use of Certificates

TLS/SSL provides privacy and data integrity between applications communicating over a network by encrypting the packets transmitted between endpoints (ports on a host, for example). Configuring TLS/SSL for any system typically involves creating a *private key* and *public key* for use by server and client processes to negotiate an encrypted connection at runtime. In addition, TLS/SSL can use *certificates* to verify the trustworthiness of keys presented during the negotiation to prevent spoofing and mitigate other potential security issues.

Setting up Cloudera clusters to use TLS/SSL requires creating private key, public key, and storing these securely in a keystore, among other tasks. Although adding a certificate to the keystore may be the last task in the process, the lead time required to obtain a certificate depends on the type of certificate you plan to use for the cluster.

Certificates Overview

A certificate is digitally signed, typically by a *certificate authority* (CA) that indirectly (through a chain of trust) verifies the authenticity of the public key presented during the negotiation. Certificates can be signed in one of the three different ways shown in the table:

Type	Usage Note
Public CA-signed certificates	**Recommended**. This type of certificate is signed by a public certificate authority (CA), such as Symantec or Comodo. Public CAs are trusted third-parties whose certificates can be verified through publicly accessible chains of trust. Using this type of certificate can simplify deployment because security infrastructure, such as root CAs, are already contained in the Java JDK and its default truststore. See Generate TLS Certificates on page 195 for details.
Internal CA-signed certificates	This type of certificate is signed by your organization's internal CA. Organizations using OpenSSL Certificate Authority, Microsoft Active Directory Certificate Service, or another internal CA system can use this type of certificate. See How to Configure TLS Encryption for Cloudera Manager for details about using internal CA-signed certificates for configuration.
Self-signed certificates	Not recommended for production deployments. Self-signed certificates are acceptable for use in non-production deployments, such as for proof-of-concept setups. See Using Self-signed Certificates for TLS for details.

During the process of configuring TLS/SSL for the cluster, you typically obtain a certificate for each host in the cluster, and re-use the certificate obtained in a given format (JKS, PEM) as needed for the various services (daemon roles) supported by the host. For information about converting formats, see How to Convert Certificate Encodings (DER, JKS, PEM) for TLS/SSL. As an alternative to creating discrete certificates for each host in the cluster, as of Cloudera Manager/CDH 5.9, all Cloudera cluster components support wildcard domains and SubjectAlternateName certificates.

Wildcard Domain Certificates and SAN Certificates Support

Cloudera Manager and CDH (as of release 5.9) support use of wildcard domain certificates and SAN certificates.

A wildcard certificate—a certificate with the common name `*`, as in `*.example.com`, rather than a specific host name—can be used for any number of first level sub-domains within a single domain name. For example, a wildcard certificate can be used with `host-1.example.com`, `host-2.example`, `host-3.example.com`, and so on.

Certificates obtained from public CAs are not free, so using wildcard certificates can reduce costs. Using wildcard certificates also makes it easier to enable encryption for transient clusters and for clusters that need to expand and shrink, since the same certificate and keystore can be re-used.

> **Important:** Be aware that using wildcard domain certificates has some security risks. Specifically, because all nodes use the same certificate, a breach of any one machine can result in a breach of all machines.

A SubjectAlternativeName or SAN certificate is a certificate that uses the SubjectAlternativeName extension to associate the resulting certificate with multiple specific host names. In the context of clusters, SAN certificates are used in high-availability (HA) configurations in which a load balancer targets two different specific hosts as primary and secondary nodes for fail-over purposes. See Server Certificate Requirements for HA Deployments for an example.

Wildcard Certificates	Wildcard certificates can be used by all hosts within a given domain. Using wildcard certificates for all hosts in the cluster can reduce costs but also exposes greater potential risk.
SubjectAlternativeName Certificates	SubjectAlternativeName (SAN) certificates are bound to a set of specific DNS names. A single SAN certificate can be used for all hosts or a subset of hosts in the cluster. SAN certificates are used in Cloudera Manager high-availability (HA) configurations.

Renew Certificates Before Expiration Dates

The signed certificates you obtain from a public CA (or those you obtain from an internal CA) have an expiration date, such as that shown in this excerpt:

```
$ openssl x509 -in cacert.pem -text -noout
Certificate:
    Data:
        Version: 3 (0x2)
        Serial Number: 11485830970703032316 (0x9f65de69ceef2ffc)
    Signature Algorithm: sha256WithRSAEncryption
        Issuer: C=US, ST=MD, L=Baltimore, CN=Test CA/emailAddress=test@example.com
        Validity
            Not Before: Jan 24 14:24:11 2017 GMT
            Not After : Feb 23 14:24:11 2018 GMT
        Subject: C=US, ST=MD, L=Baltimore, CN=Test CA/emailAddress=test@example.com
        Subject Public Key Info:
            Public Key Algorithm: rsaEncryption
                Public-Key: (4096 bit)
                Modulus:
                    00:b1:7f:29:be:78:02:b8:56:54:2d:2c:ec:ff:6d:
                    ...
                    39:f9:1e:52:cb:8e:bf:8b:9e:a6:93:e1:22:09:8b:
```

Expired certificates cause most cluster operations to fail. Cloudera Manager Agent hosts, for example, will not be able to validate the Cloudera Manager Server host and will fail to launch the cluster nodes. Administrators should note expiration dates of all certificates when they deploy the certificates to the cluster nodes and setup reminders to allow enough time to renew.

Tip: Use OpenSSL to check the expiration dates for certificates already deployed:

```
openssl x509 -enddate -noout -in /opt/cloudera/security/pki/$(hostname -f)-server.cert.pem
```

Understanding Keystores and Truststores

Configuring Cloudera Manager Server and cluster components to use TLS/SSL requires obtaining keys, certificates, and related security artifacts. The following provides a brief overview.

Java Keystore and Truststore

All clients in a Cloudera Manager cluster configured for TLS/SSL need access to the truststore to validate certificates presented during TLS/SSL session negotiation. The certificates assure the client or server process that the issuing authority for the certificate is part of a legitimate chain of trust.

The standard Oracle Java JDK distribution includes a default **truststore** (`cacerts`) that contains root certificates for many well-known CAs, including Symantec. Rather than using the default truststore, Cloudera recommends using the **alternative truststore**, `jssecacerts`. The alternative truststore is created by copying `cacerts` to that filename (`jssecacerts`). Certificates can be added to this truststore when needed for additional roles or services. This alternative truststore is loaded by Hadoop daemons at startup.

> **Important:** For use with Cloudera clusters, the alternative trust store—`jssecacerts`—must start as a **copy of `cacerts`** because `cacerts` contains all available default certificates needed to establish the chain of trust during the TLS/SSL handshake. After `jssecacerts` has been created, new public and private root CAs are added to it for use by the cluster. See Generate TLS Certificates on page 195 for details.

The private keys are maintained in the **keystore**.

> **Note:** For detailed information about the Java keystore and truststore, see Oracle documentation:
> - Keytool—Key and Certificate Management Tool
> - JSSE Reference Guide for Java

Although the keystore and truststore in some environments may comprise the same file, as configured for Cloudera Manager Server and CDH clusters, the keystore and truststore are distinct files. For Cloudera Manager Server clusters, each host should have its own keystore, while several hosts can share the same truststore. This table summarizes the general differences between keystore and the truststore in Cloudera Manager Server clusters.

Keystore	Truststore
Used by the server side of a TLS/SSL client-server connection.	Used by the client side of a TLS/SSL client-server connection.
Typically contains 1 private key for the host system.	Contains no keys of any kind.
Contains the certificate for the host's private key.	Contains root certificates for well-known public certificate authorities. May contain certificates for intermediary certificate authorities.
Password protected. Use the same password for the key and its keystore.	Password-protection not needed. However, if password has been used for the truststore, never use the same password as used for a key and keystore.
Password stored in a plaintext file read permissions granted to a specific group only (OS filesystem permissions set to 0440, `hadoop:hadoop`).	Password (if there is one for the truststore) stored in a plaintext file readable by all (OS filesystem permissions set to 0440).
No default. Provide a keystore name and password when you create the private key and CSR for any host system.	For Java JDK, **cacerts** is the default unless the alternative default **jssecacerts** is available.
Must be owned by hadoop user and group so that HDFS, MapReduce, YARN can access the private key.	HDFS, MapReduce, and YARN need client access to truststore.

The details in the table above are specific to the Java KeyStore (JKS) format, which is used by Java-based cluster services such as Cloudera Manager Server, Cloudera Management Service, and many (but not all) CDH components and services. See Certificate Formats (JKS, PEM) and Cluster Components on page 194 for information about certificate and key file type used various processes.

CDH Services as TLS/SSL Servers and Clients

Cluster services function as a TLS/SSL server, client, or both:

Component	Client	Server
HBase	⊘	✔
HDFS	✔	✔
Hive	✔	✔
Hue (Hue is a TLS/SSL client of HDFS, MapReduce, YARN, HBase, and Oozie.)	✔	⊘
MapReduce	✔	✔
Oozie	⊘	✔
YARN	✔	✔

Daemons that function as TLS/SSL servers load the keystores when starting up. When a client connects to an TLS/SSL server daemon, the server transmits the certificate loaded at startup time to the client, and the client then uses its truststore to validate the certificate presented by the server.

Certificate Formats (JKS, PEM) and Cluster Components

Cloudera Manager Server, Cloudera Management Service, and many other CDH services use JKS formatted keystores and certificates. Cloudera Manager Agent, Hue, Key Trustee Server, Impala, and other Python or C++ based services require PEM formatted certificates and keystores rather than Java. Specifically, PEM certificates conform to PKCS #8, which requires individual Base64-encoded text files for certificate and password-protected private key file. The table summarizes certificate types required by several components.

Component	JKS	PEM
HBase	✔	⊘
HDFS	✔	⊘
Hive (Hive clients and HiveServer 2)	✔	⊘
Hue	⊘	✔
Impala	⊘	✔
MapReduce	✔	⊘
Oozie	✔	⊘
Solr	✔	⊘
YARN	✔	⊘

For more information, see:

- How to Convert Certificate Encodings (DER, JKS, PEM) for TLS/SSL
- OpenSSL Cryptography and TLS/SSL Toolkit

Recommended Keystore and Truststore Configuration

Cloudera recommends the following for keystores and truststores for Cloudera Manager clusters:

- Create a separate keystore for each host. Each keystore should have a name that helps identify it as to the type of host—server or agent, for example. The keystore contains the private key and should be password protected.
- Create a single truststore that can be used by the entire cluster. This truststore contains the root CA and intermediate CAs used to authenticate certificates presented during TLS/SSL handshake. The truststore does not need to be password protected. (See How To Add Root and Intermediary CAs to Truststore for more information about the truststore for TLS/SSL and Cloudera clusters.)

The steps included in Generate TLS Certificates on page 195 follow this approach.

Configuring TLS Encryption for Cloudera Manager

When you configure authentication and authorization on a cluster, Cloudera Manager Server sends sensitive information over the network to cluster hosts, such as Kerberos keytabs and configuration files that contain passwords. To secure this transfer, you must configure TLS encryption between Cloudera Manager Server and all cluster hosts.

TLS encryption is also used to secure client connections to the Cloudera Manager Admin Interface, using HTTPS.

Cloudera Manager also supports TLS authentication. Without certificate authentication, a malicious user can add a host to Cloudera Manager by installing the Cloudera Manager Agent software and configuring it to communicate with Cloudera Manager Server. To prevent this, you must install certificates on each agent host and configure Cloudera Manager Server to trust those certificates.

This guide shows how to configure and enable TLS encryption and certificate authentication for Cloudera Manager. The provided examples use an internal certificate authority (CA) to sign all TLS certificates, so this guide also shows you how to establish trust with the CA. (For certificates signed by a trusted public CA, establishing trust is not necessary, because the Java Development Kit (JDK) already trusts them.)

Use this guide to enable TLS encryption and certificate authentication for Cloudera Manager:

Generate TLS Certificates

The following procedure assumes that an internal certificate authority (CA) is used, and shows how to establish trust for that internal CA. If you are using a trusted public CA (such as Symantec, GeoTrust, Comodo, and others), you do not need to explicitly establish trust for the issued certificates, unless you are using an older JDK and a newer public CA. Older JDKs might not trust newer public CAs by default.

On Each Cluster Host:

Complete the following procedure on each cluster host, including the Cloudera Manager Server host.

1. Configure your environment to set `JAVA_HOME` to the Oracle JDK. For example:

```
export JAVA_HOME=/usr/java/jdk1.8.0_162
```

   If you log out of the host before completing this procedure, make sure to set `JAVA_HOME` again when you log in to complete the steps.
2. Create the `/opt/cloudera/security/pki` directory:

```
sudo mkdir -p /opt/cloudera/security/pki
```

   If you choose to use a different directory, make sure you use the same directory on all cluster hosts to simplify management and maintenance.
3. Use the `keytool` utility to generate a Java keystore and certificate signing request (CSR). Replace the `OU`, `O`, `L`, `ST`, and `C` entries with the values for your environment. When prompted, use the same password for the `keystore`

`password` and `key password`. Cloudera Manager does not support using different passwords for the key and keystore.

```
$JAVA_HOME/bin/keytool -genkeypair -alias $(hostname -f) -keyalg RSA -keystore
/opt/cloudera/security/pki/$(hostname -f).jks -keysize 2048 -dname "CN=$(hostname
-f),OU=Engineering,O=Cloudera,L=Palo Alto,ST=California,C=US" -ext san=dns:$(hostname
-f)
```

```
$JAVA_HOME/bin/keytool -certreq -alias $(hostname -f) -keystore
/opt/cloudera/security/pki/$(hostname -f).jks -file /opt/cloudera/security/pki/$(hostname
 -f).csr -ext san=dns:$(hostname -f) -ext EKU=serverAuth,clientAuth
```

Note: You must ensure that your Issuing Authority will issue the certificates with the extensions CDH requires.

4. Submit the CSR files (for example, `cm01.example.com.csr`) to your certificate authority to obtain a server certificate.

   For security purposes, many commercial CAs ignore requested extensions in a CSR. Make sure that you inform the CA that you require certificates with both server and client authentication options.

   If possible, obtain the certificate in PEM (Base64 ASCII) format. The certificate file is in PEM format if it looks similar to this (some lines omitted):

```
-----BEGIN CERTIFICATE-----
MIIDAzCCAesCAQAwgY0xCzAJBgNVBAYTAlVTMRMwEQYDVQQIEwpDYWxpZm9ybmlh
MRIwEAYDVQQHEwlQYWxvIEFsdG8xETAPBgNVBAoTCENsb3VkZXJhMRQwEgYDVQQL
…
tudY0C32LjGjWOg5ALliN9Oy1u2xRKGAVfapbzAZ2rchtlCZc7mtaT6BXgW8S+Db
0HhuObn1/8TL4Ho9G+KlJB3MWik2oEbOvQt0rBidMr9qaNX86m0i7pouXZelZ5c5
UnDPtrhW6A==
-----END CERTIFICATE-----
```

   If your issued certificate is in binary (DER) format, convert it to PEM format.

5. After you have received the signed certificate, copy the signed certificate to the following location:

```
/opt/cloudera/security/pki/$(hostname -f).pem
```

6. Inspect the signed certificate to verify that both server and client authentication options are present, as well as the subject alternative name:

```
openssl x509 -in /opt/cloudera/security/pki/$(hostname -f).pem -noout -text
```

   Look for output similar to the following to verify the server and client authentication options:

```
            X509v3 Extended Key Usage:
                TLS Web Server Authentication, TLS Web Client Authentication
```

   Look for output similar to the following to validate the subject alternative name:

```
            X509v3 Subject Alternative Name:
                DNS:hostname.example.com
```

Important:

If the certificate does not have the DNS field, re-submit the CSR to the CA, and request that they generate a certificate that keeps the Subject Alternative Name field intact.

If the certificate does not have both `TLS Web Server Authentication` and `TLS Web Client Authentication` listed in the `X509v3 Extended Key Usage` section, re-submit the CSR to the CA, and request that they generate a certificate that can be used for both server and client authentication.

7. Copy the root and intermediate CA certificates to `/opt/cloudera/security/pki/rootca.pem` and `/opt/cloudera/security/pki/intca.pem` on each host. If you have a concatenated file containing the root CA and an intermediate CA certificate, split the file along the `END CERTIFICATE/BEGIN CERTIFICATE` boundary into individual files. If there are multiple intermediate CA certificates, use unique file names such as `intca-1.pem`, `intca-2.pem`, and so on.
8. Copy the JDK `cacerts` file to `jssecacerts` as follows:

```
sudo cp $JAVA_HOME/jre/lib/security/cacerts $JAVA_HOME/jre/lib/security/jssecacerts
```

Note: The default password for the `cacerts` file is `changeit`. The same applies to the `jssecacerts` file if you copied it from the `cacerts` before changing the password. Cloudera recommends changing these passwords by running the following commands:

```
$JAVA_HOME/bin/keytool -storepasswd -keystore
$JAVA_HOME/jre/lib/security/cacerts
```

```
$JAVA_HOME/bin/keytool -storepasswd -keystore
$JAVA_HOME/jre/lib/security/jssecacerts
```

The Oracle JDK uses the `jssecacerts` file for its default truststore if it exists. Otherwise, it uses the `cacerts` file. Creating the `jssecacerts` file allows you to trust an internal CA without modifying the `cacerts` file that is included with the JDK.

Note: If you upgrade your JDK, make sure to copy your existing `jssecacerts` file to the new JDK (under `$JAVA_HOME/jre/lib/security`).

9. Import the root CA certificate into the JDK truststore.

```
sudo $JAVA_HOME/bin/keytool -importcert -alias rootca -keystore
$JAVA_HOME/jre/lib/security/jssecacerts -file /opt/cloudera/security/pki/rootca.pem
```

If you see a message like the following, enter `yes` to continue:

```
Trust this certificate? [no]: yes
```

You *must* see the following response verifying that the certificate has been properly imported:

```
Certificate was added to keystore
```

10. Append the intermediate CA certificate to the signed host certificate, and then import it into the keystore. Make sure that you use the append operator (>>) and not the overwrite operator (>):

```
sudo cat /opt/cloudera/security/pki/intca.pem >> /opt/cloudera/security/pki/$(hostname
 -f).pem
```

```
sudo $JAVA_HOME/bin/keytool -importcert -alias $(hostname -f) -file
/opt/cloudera/security/pki/$(hostname -f).pem -keystore
/opt/cloudera/security/pki/$(hostname -f).jks
```

If you see a message like the following, enter `yes` to continue:

```
... is not trusted. Install reply anyway? [no]:  yes
```

You *must* see the following response verifying that the certificate has been properly imported:

```
Certificate reply was installed in keystore
```

If you do not see this response, contact Cloudera Support.

11. Create symbolic links (symlink) for the certificate and keystore files:

```
sudo ln -s /opt/cloudera/security/pki/$(hostname -f).pem
/opt/cloudera/security/pki/agent.pem
```

This allows you to use the same `/etc/cloudera-scm-agent/config.ini` file on all agent hosts rather than maintaining a file for each agent.

On the Cloudera Manager Server Host

On the Cloudera Manager Server host, create an additional symlink for the keystore file:

```
sudo ln -s /opt/cloudera/security/pki/$(hostname -f).jks
/opt/cloudera/security/pki/server.jks
```

Configure TLS for the Cloudera Manager Admin Console

Minimum Required Role: **Cluster Administrator** (also provided by **Full Administrator**)

Use the following procedure to enable TLS encryption for the Cloudera Manager Server admin interface. Make sure you have generated the host certificate as described in Generate TLS Certificates on page 195.

Step 1: Enable HTTPS for the Cloudera Manager Admin Console

1. Log in to the Cloudera Manager Admin Console.
2. Select **Administration** > **Settings**.
3. Select the **Security** category.
4. Configure the following TLS settings:

Property	Description
Cloudera Manager TLS/SSL Server JKS Keystore File Location	The complete path to the keystore file. For example: `/opt/cloudera/security/pki/server.jks`
Cloudera Manager TLS/SSL Server JKS Keystore File Password	The password for the `/opt/cloudera/security/jks/server.jks` keystore.

Property	Description
Use TLS Encryption for Admin Console	Check this box to enable TLS encryption for Cloudera Manager.

5. Enter a **Reason for Change**, then click **Save Changes** to save the settings.

Step 2: Specify SSL Truststore Properties for Cloudera Management Services

When enabling TLS for the Cloudera Manager Server admin interface, you must set the Java truststore location and password in the Cloudera Management Services configuration. Otherwise, roles such as Host Monitor and Service Monitor cannot connect to Cloudera Manager Server and will not start.

Configure the path and password for the `$JAVA_HOME/jre/lib/security/jssecacerts` truststore that you created earlier. Make sure that you have created this file on all hosts, including the Cloudera Management Service hosts, as instructed in Generate TLS Certificates on page 195.

1. Open the Cloudera Manager Administration Console and go to the **Cloudera Management Service** service.
2. Click the **Configuration** tab.
3. Select **Scope** > **Cloudera Management Service (Service-Wide)**.
4. Select **Category** > **Security**.
5. Edit the following TLS/SSL properties according to your cluster configuration.

Property	Description
TLS/SSL Client Truststore File Location	The path to the client truststore file used in HTTPS communication. This truststore contains certificates of trusted servers, or of Certificate Authorities trusted to identify servers. For this example, set the value to: `<JAVA_HOME>/jre/lib/security/jssecacerts` Replace <*JAVA_HOME*> with the path to the Oracle JDK.
Cloudera Manager Server TLS/SSL Certificate Trust Store Password	The password for the truststore file.

6. Click **Save Changes** to commit the changes.

Step 3: Restart Cloudera Manager and Services

You must restart both Cloudera Manager Server and the Cloudera Management Service for TLS encryption to work. Otherwise, the Cloudera Management Services (such as Host Monitor and Service Monitor) cannot communicate with Cloudera Manager Server.

1. Restart the Cloudera Manager Server by running the following command on the Cloudera Manager Server host:
   - RHEL 7 compatible:

```
sudo systemctl restart cloudera-scm-server
```

   - RHEL 6 compatible, SLES, Ubuntu:

```
sudo service cloudera-scm-server restart
```

2. After the restart completes, connect to the Cloudera Manager Admin Console using the HTTPS URL (for example: `https://cm01.example.com:7183`). If you used an internal CA-signed certificate, you must configure your browser to trust the certificate. Otherwise, you will see a warning in your browser any time you access the Cloudera Manager Administration Console. By default, certificates issued by public commercial CAs are trusted by most browsers, and no additional configuration is necessary if your certificate is signed by one of them.

3. Restart the Cloudera Management Service (**Cloudera Management Service** > **Actions** > **Restart**).

Configure TLS for Cloudera Manager Agents

Minimum Required Role: Cluster Administrator (also provided by **Full Administrator**)

Use the following procedure to encrypt the communication between Cloudera Manager Server and Cloudera Manager Agents:

Step 1: Enable TLS Encryption for Agents in Cloudera Manager

Configure the TLS properties for Cloudera Manager Agents.

1. Log in to the Cloudera Manager Admin Console.
2. Select **Administration** > **Settings**.
3. Select the **Security** category.
4. Select the **Use TLS Encryption for Agents** option.
5. Click **Save Changes** to commit the changes.

Step 2: Enable TLS on Cloudera Manager Agent Hosts

To enable TLS between the Cloudera Manager agents and Cloudera Manager, you must specify values for the TLS properties in the `/etc/cloudera-scm-agent/config.ini` configuration file on all agent hosts.

1. On each agent host (including the Cloudera Manager Server host, which also has an agent), open the `/etc/cloudera-scm-agent/config.ini` configuration file and set the `use_tls` parameter in the `[Security]` section as follows:

```
use_tls=1
```

   Alternatively, you can edit the `config.ini` file on one host, and then copy it to the other hosts because this file by default does not contain host-specific information. If you have modified properties such as `listening_hostname` or `listening_ip address` in `config.ini`, you must edit the file individually on each host.

Step 3: Restart Cloudera Manager Server and Agents

1. Restart the Cloudera Manager Server by running the following command on the Cloudera Manager Server host:
   - RHEL 7 compatible:

```
sudo systemctl restart cloudera-scm-server
```

   - RHEL 6 compatible, SLES, Ubuntu:

```
sudo service cloudera-scm-server restart
```

2. On each agent host (including the Cloudera Manager Server host), restart the Cloudera Manager agent service:
   - RHEL 7 compatible:

```
sudo systemctl restart cloudera-scm-agent
```

   - RHEL 6 compatible, SLES, Ubuntu:

```
sudo service cloudera-scm-agent restart
```

Step 4: Verify that the Cloudera Manager Server and Agents are Communicating

In the Cloudera Manager Admin Console, go to **Hosts** > **All Hosts**. If you see successful heartbeats reported in the **Last Heartbeat** column after restarting the agents, TLS encryption is working properly.

Enable Server Certificate Verification on Cloudera Manager Agents

Minimum Required Role: **Cluster Administrator** (also provided by **Full Administrator**)

If you have completed the previous sections, communication between Cloudera Manager server and the agents is encrypted, but the certificate authenticity is not verified. For full security, you must configure the agents to verify the Cloudera Manager server certificate. If you are using a server certificate signed by an internal certificate authority (CA), you must configure the agents to trust that CA:

1. On each agent host (including the Cloudera Manager Server host), open the `/etc/cloudera-scm-agent/config.ini` configuration file, and then uncomment and set the following property:

```
verify_cert_file=/opt/cloudera/security/pki/rootca.pem
```

   Alternatively, you can edit the `config.ini` file on one host, and then copy it to the other hosts because this file by default does not contain host-specific information. If you have modified properties such as `listening_hostname` or `listening_ip address` in `config.ini`, you must edit the file individually on each host.

2. Restart the Cloudera Manager agents. On each agent host (including the Cloudera Manager Server host), run the following command:
   - RHEL 7 compatible:

```
sudo systemctl restart cloudera-scm-agent
```

   - RHEL 6 compatible, SLES, Ubuntu:

```
sudo service cloudera-scm-agent restart
```

3. Verify that the Cloudera Manager server and agents are communicating. In the Cloudera Manager Admin Console, go to **Hosts** > **All Hosts**. If you see successful heartbeats reported in the **Last Heartbeat** column after restarting the agents and management service, TLS verification is working properly. If not, check the agent log (`/var/log/cloudera-scm-agent/cloudera-scm-agent.log`) for errors.

Configure Agent Certificate Authentication

> **Important:** Perform this procedure on each agent host, including the Cloudera Manager Server host, which also has an agent.

Without certificate authentication, a malicious user can add a host to Cloudera Manager by installing the Cloudera Manager agent software and configuring it to communicate with Cloudera Manager Server. To prevent this, you must configure Cloudera Manager to trust the agent certificates.

Step 1: Export the Private Key to a File

On each Cloudera Manager Agent host, use the `keytool` utility to export the private key and certificate to a PKCS12 file, which can then be split up into individual key and certificate files using the `openssl` command:

1. Export the private key and certificate:

```
sudo $JAVA_HOME/bin/keytool -importkeystore -srckeystore
/opt/cloudera/security/pki/$(hostname -f).jks -destkeystore
/opt/cloudera/security/pki/$(hostname -f)-key.p12 -deststoretype PKCS12 -srcalias
$(hostname -f)
```

2. Use the `openssl` command to export the private key into its own file:

```
sudo openssl pkcs12 -in /opt/cloudera/security/pki/$(hostname -f)-key.p12 -nocerts -out
 /opt/cloudera/security/pki/$(hostname -f).key
```

3. Create a symbolic link for the `.key` file:

```
sudo ln -s /opt/cloudera/security/pki/$(hostname -f).key
/opt/cloudera/security/pki/agent.key
```

   This allows you to use the same `/etc/cloudera-scm-agent/config.ini` file on all agent hosts rather than maintaining a file for each agent.

Step 2: Create a Password File

The Cloudera Manager agent obtains the password from a text file, not from a command line parameter or environment variable. The password file allows you to use file permissions to protect the password. For example, run the following commands on each Cloudera Manager Agent host, or run them on one host and copy the file to the other hosts:

Create and secure the file containing the password used to protect the private key of the Agent:

1. Use a text editor to create a file called `/etc/cloudera-scm-agent/agentkey.pw` that contains the password.
2. Change ownership of the file to `root`:

```
sudo chown root:root /etc/cloudera-scm-agent/agentkey.pw
```

3. Change the permissions of the file:

```
sudo chmod 440 /etc/cloudera-scm-agent/agentkey.pw
```

Step 3: Configure the Agent to Use Private Keys and Certificates

On a Cloudera Manager Agent, open the `/etc/cloudera-scm-agent/config.ini` configuration file, uncomment and edit the following properties:

Property	Example Value	Description
`client_key_file`	/opt/cloudera/security/pki/agent.key	Path to the private key file.
`client_keypw_file`	/etc/cloudera-scm-agent/agentkey.pw	Path to the private key password file.
`client_cert_file`	/opt/cloudera/security/pki/agent.pem	Path to the client certificate file.

Copy the file to all other cluster hosts. If you have modified properties such as `listening_hostname` or `listening_ip address` in `config.ini`, you must edit the file individually on each host.

Step 4: Enable Agent Certificate Authentication

1. Log in to the Cloudera Manager Admin Console.
2. Select **Administration** > **Settings**.
3. Click the **Security** category.
4. Configure the following TLS settings:

Setting	Description
Use TLS Authentication of Agents to Server	Select this option to enable TLS authentication of agents to the server.
Cloudera Manager TLS/SSL Certificate Trust Store File	Specify the full filesystem path to the `jssecacerts` file located on the Cloudera Manager Server host. For this example, set the value to: `<JAVA_HOME>/jre/lib/security/jssecacerts` Replace <*JAVA_HOME*> with the path to the Oracle JDK.
Cloudera Manager TLS/SSL Certificate Trust Store Password	Specify the password for the `jssecacerts` truststore.

5. Click **Save Changes** to commit the changes.

Step 5: Restart Cloudera Manager Server and Agents

1. On the Cloudera Manager server host, restart the Cloudera Manager server:
   - RHEL 7 compatible:

```
sudo systemctl restart cloudera-scm-server
```

   - RHEL 6 compatible, SLES, Ubuntu:

```
sudo service cloudera-scm-server restart
```

2. On every agent host, restart the Cloudera Manager agent:
   - RHEL 7 compatible:

```
sudo systemctl restart cloudera-scm-agent
```

   - RHEL 6 compatible, SLES, Ubuntu:

```
sudo service cloudera-scm-agent restart
```

Step 6: Verify that Cloudera Manager Server and Agents are Communicating

In the Cloudera Manager Admin Console, go to **Hosts** > **All Hosts**. If you see successful heartbeats reported in the **Last Heartbeat** column after restarting the agents and server, TLS certificate authentication is working properly. If not, check the agent log (`/var/log/cloudera-scm-agent/cloudera-scm-agent.log`) for errors.

For example, you might see the following error:

```
WrongHost: Peer certificate commonName does not match host, expected 192.0.2.155, got
cdh-1.example.com
[02/May/2018 15:04:15 +0000] 4655 MainThread agent        ERROR    Heartbeating to
192.0.2.155:7182 failed
```

For this scenario, make sure that your DNS and `/etc/hosts` file are configured correctly, and that your `server_host` parameter in `/etc/cloudera-scm-agent/config.ini` uses the Cloudera Manager Server hostname, and not IP address.

Configuring TLS/SSL Encryption for CDH Services

In addition to configuring Cloudera Manager cluster to use TLS/SSL (as detailed in Configuring TLS Encryption for Cloudera Manager on page 195), the various CDH services running on the cluster should also be configured to use TLS/SSL. The process of configuring TLS/SSL varies by component, so follow the steps below as needed for your system. Before trying to configure TLS/SSL, however, be sure your cluster meets prerequisites.

In general, all the roles on any given node in the cluster can use the same certificates, assuming the certificates are in the appropriate format (JKS, PEM) and that the configuration properly points to the location. If you follow the steps in How to Configure TLS Encryption for Cloudera Manager to create your CSRs and use the symbolic link for the path to the certificates, you will be setting up the certificates in the cluster for optimal reuse.

> **Note:** TLS/SSL for Hadoop core services—HDFS, MapReduce, and YARN—must be enabled as a group. TLS/SSL for other components such as HBase, Hue, and Oozie can be enabled independently.

Not all components support TLS/SSL, nor do all external engines support TLS/SSL. Unless explicitly listed in this guide, the component you want to configure may not currently support TLS/SSL. For example, Sqoop does not currently support TLS/SSL to Oracle, MySQL, or other databases.

Prerequisites

Cloudera recommends that the cluster and all services use Kerberos for authentication. If you enable TLS/SSL for a cluster that has not been configured to use Kerberos, a warning displays. You should integrate the cluster with your Kerberos deployment before proceeding.

The steps below require the cluster to have TLS enabled to ensure that Cloudera Manager Server certificate and Cloudera Manager Agent certificates are properly configured and already in place. In addition, you should have the certificates and keys needed by the specific CDH server ready.

If the cluster meets these requirements, you can configure the specific CDH service to use TLS/SSL, as detailed in this section.

Configuring TLS/SSL for HDFS, YARN and MapReduce

Required Role: Configurator, Cluster Administrator, or Full Administrator

TLS/SSL for the core Hadoop services—HDFS, MapReduce, and YARN—must be enabled as a group. Because most clusters run either MapReduce or YARN, not both, you will typically enable HDFS and YARN, or HDFS and MapReduce. Enabling TLS/SSL on HDFS is required before it can be enabled on either MapReduce or YARN.

> **Note:** If you enable TLS/SSL for HDFS, you must also enable it for MapReduce or YARN.

The steps below include enabling Kerberos authentication for HTTP Web-Consoles. Enabling TLS/SSL for the core Hadoop services on a cluster without enabling authentication displays a warning.

Before You Begin

- Before enabling TLS/SSL, keystores containing certificates bound to the appropriate domain names will need to be accessible on all hosts on which at least one HDFS, MapReduce, or YARN daemon role is running.
- Since HDFS, MapReduce, and YARN daemons act as TLS/SSL clients as well as TLS/SSL servers, they must have access to truststores. In many cases, the most practical approach is to deploy truststores to all hosts in the cluster, as it may not be desirable to determine in advance the set of hosts on which clients will run.
- Keystores for HDFS, MapReduce and YARN must be owned by the hadoop group, and have permissions `0440` (that is, readable by owner and group). Truststores must have permissions `0444` (that is, readable by all)
- Cloudera Manager supports TLS/SSL configuration for HDFS, MapReduce and YARN at the service level. For each of these services, you must specify absolute paths to the keystore and truststore files. These settings apply to all

hosts on which daemon roles of the service in question run. Therefore, the paths you choose must be valid on all hosts.

An implication of this is that the keystore file names for a given service must be the same on all hosts. If, for example, you have obtained separate certificates for HDFS daemons on hosts `node1.example.com` and `node2.example.com`, you might have chosen to store these certificates in files called `hdfs-node1.keystore` and `hdfs-node2.keystore` (respectively). When deploying these keystores, you must give them both the same name on the target host — for example, `hdfs.keystore`.

- Multiple daemons running on a host can share a certificate. For example, in case there is a DataNode and an Oozie server running on the same host, they can use the same certificate.

Configuring TLS/SSL for HDFS

1. Go to the **HDFS** service.
2. Click the **Configuration** tab.
3. Select **Scope** > **HDFS (Service-Wide)**.
4. Select **Category** > **Security**.
5. In the Search field, type **TLS/SSL** to show the TLS/SSL properties (found under the **Service-Wide** > **Security** category).
6. Edit the following properties according to your cluster configuration:

Property	Description
Hadoop TLS/SSL Server Keystore File Location	Path to the keystore file containing the server certificate and private key.
Hadoop TLS/SSL Server Keystore File Password	Password for the server keystore file.
Hadoop TLS/SSL Server Keystore Key Password	Password that protects the private key contained in the server keystore.

7. If you are not using the default truststore, configure TLS/SSL client truststore properties:

> **Important:** The HDFS properties below define a cluster-wide default truststore that can be overridden by YARN and MapReduce (see the **Configuring TLS/SSL for YARN and MapReduce** section below).

Property	Description
Cluster-Wide Default TLS/SSL Client Truststore Location	Path to the client truststore file. This truststore contains certificates of trusted servers, or of Certificate Authorities trusted to identify servers.
Cluster-Wide Default TLS/SSL Client Truststore Password	Password for the client truststore file.

8. **(Optional)** Cloudera recommends you enable web UI authentication for the HDFS service. Web UI authentication uses SPNEGO. After enabling this, you cannot access the Hadoop web consoles without a valid Kerberos ticket and proper client-side configuration. For more information, see How to Configure Browsers for Kerberos Authentication on page 372.

   To enable web UI authentication, enter **web consoles** in the Search field to bring up the **Enable Authentication for HTTP Web-Consoles** property (found under the **Service-Wide**>**Security** category). Check the property to enable web UI authentication.

Enable Authentication for HTTP Web-Consoles	Enables authentication for Hadoop HTTP web-consoles for all roles of this service. **Note:** This is effective only if security is enabled for the HDFS service.

9. Click **Save Changes**.
10. Follow the procedure described in the following **Configuring TLS/SSL for YARN and MapReduce** section, at the end of which you will be instructed to restart all the affected services (HDFS, MapReduce and YARN).

Configuring TLS/SSL for YARN or MapReduce

Perform the following steps to configure TLS/SSL for the YARN or MapReduce services:

1. Go to the **YARN** or **MapReduce** service.
2. Click the **Configuration** tab.
3. Select **Scope** > ***service name*** **(Service-Wide)**.
4. Select **Category** > **Security**.
5. Locate the **<property name>** property or search for it by typing its name in the Search box.
6. In the Search field, type **TLS/SSL** to show the TLS/SSL properties (found under the **Service-Wide** > **Security** category).
7. Edit the following properties according to your cluster configuration:

Property	Description
Hadoop TLS/SSL Server Keystore File Location	Path to the keystore file containing the server certificate and private key.
Hadoop TLS/SSL Server Keystore File Password	Password for the server keystore file.
Hadoop TLS/SSL Server Keystore Key Password	Password that protects the private key contained in the server keystore.

8. Configure the following TLS/SSL client truststore properties for MRv1 or YARN only if you want to override the cluster-wide defaults set by the HDFS properties configured above.

Property	Description
TLS/SSL Client Truststore File Location	Path to the client truststore file. This truststore contains certificates of trusted servers, or of Certificate Authorities trusted to identify servers.
TLS/SSL Client Truststore File Password	Password for the client truststore file.

9. Cloudera recommends you enable Web UI authentication for the service in question.

   Enter **web consoles** in the Search field to bring up the **Enable Authentication for HTTP Web-Consoles** property (found under the **Service-Wide**>**Security** category). Check the property to enable web UI authentication.

Enable Authentication for HTTP Web-Consoles	Enables authentication for Hadoop HTTP web-consoles for all roles of this service. **Note:** This is effective only if security is enabled for the HDFS service.

10. Click **Save Changes** to commit the changes.
11. Go to the **HDFS** service

12. Click the **Configuration** tab.
13. Type **Hadoop SSL Enabled** in the Search box.
14. Select the **Hadoop SSL Enabled** property to enable SSL communication for HDFS, MapReduce, and YARN.

Property	Description
Hadoop TLS/SSL Enabled	Enable TLS/SSL encryption for HDFS, MapReduce, and YARN web UIs, as well as encrypted shuffle for MapReduce and YARN.

15. Click **Save Changes** to commit the changes.
16. Restart all affected services (HDFS, MapReduce and YARN), as well as their dependent services.

Configuring TLS/SSL for HBase

Required Role: Configurator, Cluster Administrator, or Full Administrator

Before You Begin

- Before enabling TLS/SSL, ensure that keystores containing certificates bound to the appropriate domain names will need to be accessible on all hosts on which at least one HBase daemon role is running.
- Keystores for HBase must be owned by the `hbase` group, and have permissions `0440` (that is, readable by owner and group).
- You must specify absolute paths to the keystore and truststore files. These settings apply to all hosts on which daemon roles of the HBase service run. Therefore, the paths you choose must be valid on all hosts.
- Cloudera Manager supports the TLS/SSL configuration for HBase at the service level. Ensure you specify absolute paths to the keystore and truststore files. These settings apply to all hosts on which daemon roles of the service in question run. Therefore, the paths you choose must be valid on all hosts.

  An implication of this is that the keystore file names for a given service must be the same on all hosts. If, for example, you have obtained separate certificates for HBase daemons on hosts `node1.example.com` and `node2.example.com`, you might have chosen to store these certificates in files called `hbase-node1.keystore` and `hbase-node2.keystore` (respectively). When deploying these keystores, you must give them both the same name on the target host — for example, `hbase.keystore`.

Configuring TLS/SSL for HBase Web UIs

The steps for configuring and enabling TLS/SSL for HBase are similar to those for HDFS, YARN and MapReduce:

1. Go to the **HBase service**.
2. Click the **Configuration** tab.
3. Select **Scope** > **HBASE (Service-Wide)**.
4. Select **Category** > **Security**.
5. In the Search field, type **TLS/SSL** to show the HBase TLS/SSL properties.
6. Edit the following TLS/SSL properties according to your cluster configuration:

Table 11: HBase TLS/SSL Properties

Property	Description
HBase TLS/SSL Server JKS Keystore File Location	Path to the keystore file containing the server certificate and private key used for encrypted web UIs.
HBase TLS/SSL Server JKS Keystore File Password	Password for the server keystore file used for encrypted web UIs.
HBase TLS/SSL Server JKS Keystore Key Password	Password that protects the private key contained in the server keystore used for encrypted web UIs.

7. Check the **Web UI TLS/SSL Encryption Enabled** property.

Web UI TLS/SSL Encryption Enabled	Enable TLS/SSL encryption for the HBase Master, RegionServer, Thrift Server, and REST Server web UIs.

8. Click **Save Changes**.
9. Restart the HBase service.

Configuring TLS/SSL for HBase REST Server

1. Go to the **HBase service**
2. Click the **Configuration** tab.
3. Select **Scope** > **HBase REST Server**.
4. Select **Category** > **Security**.
5. In the Search field, type **TLS/SSL REST** to show the HBase REST TLS/SSL properties.
6. Edit the following TLS/SSL properties according to your cluster configuration:

Property	Description
Enable TLS/SSL for HBase REST Server	Encrypt communication between clients and HBase REST Server using Transport Layer Security (TLS).
HBase REST Server TLS/SSL Server JKS Keystore File Location	The path to the TLS/SSL keystore file containing the server certificate and private key used for TLS/SSL. Used when HBase REST Server is acting as a TLS/SSL server. The keystore must be in JKS format.file.
HBase REST Server TLS/SSL Server JKS Keystore File Password	The password for the HBase REST Server JKS keystore file.
HBase REST Server TLS/SSL Server JKS Keystore Key Password	The password that protects the private key contained in the JKS keystore used when HBase REST Server is acting as a TLS/SSL server.

7. Click **Save Changes**.
8. Restart the HBase service.

Configuring TLS/SSL for HBase Thrift Server

1. Go to the **HBase service**
2. Click the **Configuration** tab.
3. Select **Scope** > **HBase Thrift Server**.
4. Select **Category** > **Security**.
5. In the Search field, type **TLS/SSL Thrift** to show the HBase Thrift TLS/SSL properties.
6. Edit the following TLS/SSL properties according to your cluster configuration:

Property	Description
Enable TLS/SSL for HBase Thrift Server over HTTP	Encrypt communication between clients and HBase Thrift Server over HTTP using Transport Layer Security (TLS).
HBase Thrift Server over HTTP TLS/SSL Server JKS Keystore File Location	Path to the TLS/SSL keystore file (in JKS format) with the TLS/SSL server certificate and private key. Used when HBase Thrift Server over HTTP acts as a TLS/SSL server.
HBase Thrift Server over HTTP TLS/SSL Server JKS Keystore File Password	Password for the HBase Thrift Server JKS keystore file.
HBase Thrift Server over HTTP TLS/SSL Server JKS Keystore Key Password	Password that protects the private key contained in the JKS keystore used when HBase Thrift Server over HTTP acts as a TLS/SSL server.

7. Click **Save Changes**.

8. Restart the HBase service.

Configuring TLS/SSL for Flume Thrift Source and Sink

This topic describes how to enable TLS/SSL communication between Flume's Thrift source and sink.

The following tables list the properties that must be configured to enable TLS/SSL communication between Flume's Thrift source and sink instances.

Table 12: Thrift Source TLS/SSL Properties

Property	Description
`ssl`	Set to `true` to enable TLS/SSL encryption.
`keystore`	Path to a Java keystore file. Required for TLS/SSL.
`keystore-password`	Password for the Java keystore. Required for TLS/SSL.
`keystore-type`	The type of the Java keystore. This can be JKS or PKCS12.

Table 13: Thrift Sink TLS/SSL Properties

Property	Description
`ssl`	Set to `true` to enable TLS/SSL for this ThriftSink. When configuring TLS/SSL, you can optionally set the following `truststore`, `truststore-password` and `truststore-type` properties. If a custom truststore is not specified, Flume will use the default Java JSSE truststore (typically `jssecacerts` or `cacerts` in the Oracle JRE) to verify the remote Thrift Source's TLS/SSL credentials.
`truststore`	(Optional) The path to a custom Java truststore file.
`truststore-password`	(Optional) The password for the specified truststore.
`truststore-type`	(Optional) The type of the Java truststore. This can be JKS or any other supported Java truststore type.

Make sure you are configuring TLS/SSL for ***each*** Thrift source and sink instance. For example, to the existing `flume.conf` file, for agent `a1`, source `r1`, and sink `k1`, you would add the following properties:

```
# TLS/SSL properties for Thrift source s1
a1.sources.r1.ssl=true
a1.sources.r1.keystore=<path/to/keystore>
a1.sources.r1.keystore-password=<keystore password>
a1.sources.r1.keystore-type=<keystore type>

# TLS/SSL properties for Thrift sink k1
a1.sinks.k1.ssl=true
a1.sinks.k1.truststore=<path/to/truststore>
a1.sinks.k1.truststore-password=<truststore password>
a1.sinks.k1.truststore-type=<truststore type>
```

Configure these sets of properties for more instances of the Thrift source and sink as required. You can use either Cloudera Manager or the command line to edit the `flume.conf` file.

Using Cloudera Manager

Minimum Required Role: Configurator (also provided by **Cluster Administrator, Full Administrator**)

1. Open the Cloudera Manager Admin Console and go to the **Flume** service.
2. Click the **Configuration** tab.

3. Select **Scope** > **Agent**.
4. Select **Category** > **Main**.
5. Edit the **Configuration File** property and add the Thrift source and sink properties for each Thrift source and sink instance as described above to the configuration file.
6. Click **Save Changes** to commit the changes.
7. Restart the Flume service.

Using the Command Line

Go to the `/etc/flume-ng/conf/flume.conf` file and add the Thrift source and sink properties for each Thrift source and sink instance as described above.

Configuring Encrypted Communication Between HiveServer2 and Client Drivers

Starting with CDH 5.5, encryption between HiveServer2 and its clients has been decoupled from Kerberos authentication. (Prior to CDH 5.5, SASL QOP encryption for JDBC client drivers required connections authenticated by Kerberos.) De-coupling the authentication process from the transport-layer encryption process means that HiveServer2 can support two different approaches to encryption between the service and its clients (Beeline, JDBC/ODBC) regardless of whether Kerberos is being used for authentication, specifically:

- SASL
- TLS/SSL

Unlike TLS/SSL, SASL QOP encryption does not require certificates and is aimed at protecting core Hadoop RPC communications. However, SASL QOP may have performance issues when handling large amounts of data, so depending on your usage patterns, TLS/SSL may be a better choice. See the following topics for details about configuring HiveServer2 services and clients for TLS/SSL and SASL QOP encryption.

Configuring TLS/SSL Encryption for HiveServer2

HiveServer2 can be configured to support TLS/SSL connections from JDBC/ODBC clients using the Cloudera Manager Admin Console (for clusters that run in the context of Cloudera Manager Server), or manually using the command line.

Requirements and Assumptions

Whether you use Cloudera Manager Admin Console or manually modify the Hive configuration file for TLS/SSL encryption, the steps assume that the HiveServer2 node in the cluster has the necessary server key, certificate, keystore, and trust store set up on the host system. For details, see any of the following:

- Encrypting Data in Transit on page 191
- How to Configure TLS Encryption for Cloudera Manager
- How To Obtain and Deploy Keys and Certificates for TLS/SSL

The configuration paths and filenames shown below assume that hostname variable (`$(hostname -f)-server.jks`) was used with Java keytool commands to create keystore, as shown in this example:

```
$ sudo keytool -genkeypair -alias $(hostname -f)-server -keyalg RSA -keystore \
/opt/cloudera/security/pki/$(hostname -f)-server.jks -keysize 2048 -dname \
"CN=$(hostname
-f),OU=dept-name-optional,O=company-name,L=city,ST=state,C=two-digit-nation" \
-storepass password -keypass password
```

See the appropriate How-To guide from the above list for more information.

Using Cloudera Manager to Enable TLS/SSL

To configure TLS/SSL for Hive in clusters managed by Cloudera Manager:

1. Log in to the Cloudera Manager Admin Console.
2. Select **Clusters** > **Hive**.
3. Click the **Configuration** tab.
4. Select **Hive (Service-Wide)** for the **Scope** filter.

5. Select **Security** for the **Category** filter. The TLS/SSL configuration options display.
6. Enter values for your cluster as follows:

Property	Description
Enable TLS/SSL for HiveServer2	Click the checkbox to enable encrypted client-server communications between HiveServer2 and its clients using TLS/SSL.
HiveServer2 TLS/SSL Server JKS Keystore File Location	Enter the path to the Java keystore on the host system. For example: `/opt/cloudera/security/pki/server-name-server.jks`
HiveServer2 TLS/SSL Server JKS Keystore File Password	Enter the password for the keystore that was passed at the Java keytool command-line when the key and keystore were created. As detailed in How To Obtain and Deploy Keys and Certificates for TLS/SSL, the password for the keystore must be the same as the password for the key.
HiveServer2 TLS/SSL Certificate Trust Store File	Enter the path to the Java trust store on the host system. Cloudera clusters are typically configured to use the alternative trust store, `jssecacerts`, set up at `$JAVA_HOME/jre/lib/security/jssecacerts`.

For example:

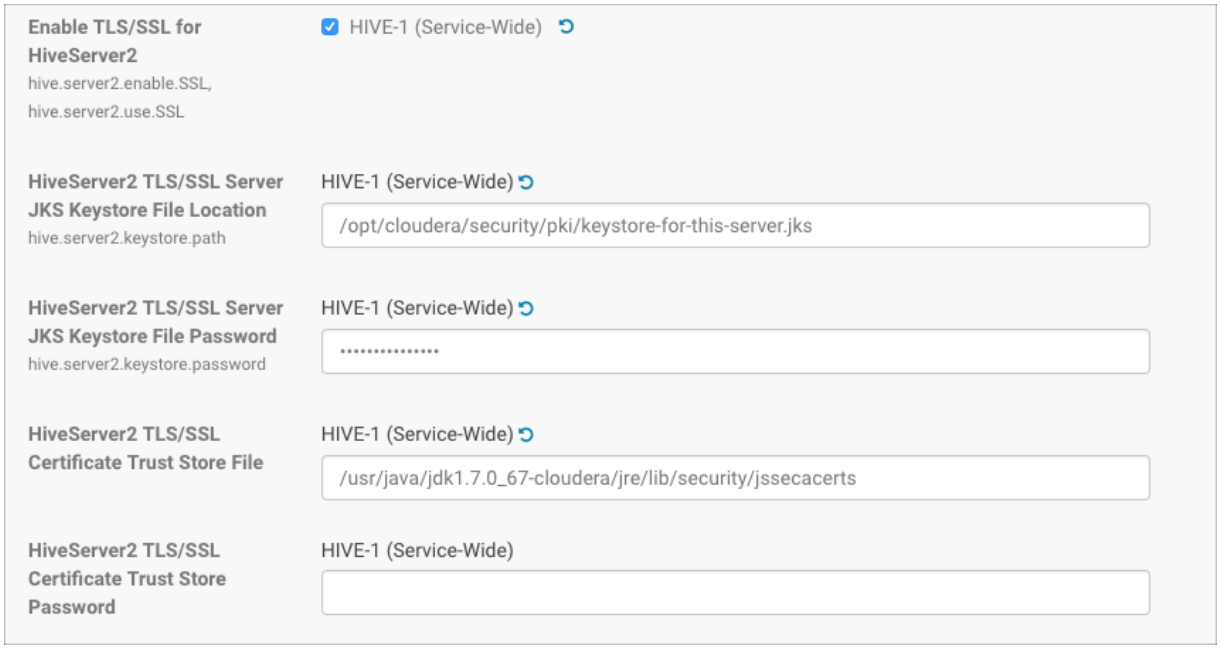

The entry field for certificate trust store password has been left empty because the trust store is typically not password protected—it contains no keys, only publicly available certificates that help establish the chain of trust during the TLS/SSL handshake. In addition, reading the trust store does not require the password.

7. Click **Save Changes**.
8. Restart the Hive service.

Using the Command Line to Enable TLS/SSL

To configure TLS/SSL for Hive in CDH clusters (without Cloudera Manager), add the properties to enable TLS/SSL and to specify the path to the keystore and the keystore password to the HiveServer2 configuration file (`hive-site.xml`) as shown below. The `keystore` must contain the server certificate and the host must meet all Requirements and Assumptions on page 210 listed above.

```
<property>
  <name>hive.server2.use.SSL</name>
  <value>true</value>
```

```
</property>

<property>
  <name>hive.server2.keystore.path</name>
  <value>/opt/cloudera/security/pki/keystore-for-this-server.jks</value>
</property>

<property>
  <name>hive.server2.keystore.password</name>
  <value>password</value>
</property>
```

Client Connections to HiveServer2 Over TLS/SSL

Clients connecting to a HiveServer2 over TLS/SSL must be able to access the trust store on the HiveServer2 host system. The trust store contains intermediate and other certificates that the client uses to establish a chain of trust and verify server certificate authenticity. The trust store is typically not password protected.

> **Note:** The trust store may have been password protected to prevent its contents from being modified. However, password protected trust stores can be read from without using the password.

The client needs the path to the trust store when attempting to connect to HiveServer2 using TLS/SSL. This can be specified using two different approaches, as follows:

- Pass the path to the trust store each time you connect to HiveServer in the JDBC connection string:

```
jdbc:hive2://fqdn.example.com:10000/default;ssl=true;\
sslTrustStore=$JAVA_HOME/jre/lib/security/jssecacerts;trustStorePassword=extraneous
```

  or,

- Set the path to the trust store one time in the Java system `javax.net.ssl.trustStore` property:

```
java
-Djavax.net.ssl.trustStore=/usr/java/jdk1.7.0_67-cloudera/jre/lib/security/jssecacerts
 \
-Djavax.net.ssl.trustStorePassword=extraneous MyClass \
jdbc:hive2://fqdn.example.com:10000/default;ssl=true
```

Configuring SASL Encryption for HiveServer2

Communications between Hive JDBC or ODBC drivers and HiveServer2 can be encrypted using SASL, a framework for authentication and data security rather than a protocol like TLS/SSL. Support for SASL (Simple Authentication and Security Layer) in HiveServer2 preceded the support for TLS/SSL. SASL offers three different Quality of Protection (QOP) levels as shown in the table:

`auth`	Default. Authentication only.
`auth-int`	Authentication with integrity protection. Signed message digests (checksums) verify the integrity of messages sent between client and server.
`auth-conf`	Authentication with confidentiality (transport-layer encryption). Use this setting for encrypted communications from clients to HiveServer2.

To support encryption for communications between client and server processes, specify the QOP `auth-conf` setting for the SASL QOP property in the HiveServer2 configuration file (`hive-site.xml`). For example,

```
<property>
  <name>hive.server2.thrift.sasl.qop</name>
  <value>auth-conf</value>
</property>
```

Client Connections to HiveServer2 Using SASL

The client connection string must match the parameter value specified for the HiveServer2 configuration. This example shows how to specify encryption for the Beeline client in the JDBC connection URL:

```
beeline> !connect jdbc:hive2://fqdn.example.com:10000/default;  \
principal=hive/_HOST@EXAMPLE.COM;sasl.qop=auth-conf
```

The `_HOST` is a wildcard placeholder that gets automatically replaced with the fully qualified domain name (FQDN) of the server running the HiveServer2 daemon process.

Configuring TLS/SSL for Hue

Hue as a TLS/SSL Client

Minimum Required Role: Configurator (also provided by **Cluster Administrator, Full Administrator**)

Hue acts as a TLS/SSL client when communicating with other services such as core Hadoop, HBase, Oozie and Amazon S3. This means Hue must authenticate HDFS, MapReduce, YARN daemons, the HBase Thrift server, and so on. To do so, Hue needs the certificate chains of their hosts in its truststore.

The Hue truststore is a single PEM file that contains the CA root, and all intermediate certificates, to authenticate the certificate installed on each TLS/SSL-enabled server. These servers host the various services with which Hue communicates.

> **Note:** A certificate is specific to a host. It is signed by a certificate authority (CA) and tells the requesting client (Hue) that "this host" is the same one as represented by the host public key. Hue uses chain of signing authorities in its truststore to validate the CA that signed the host certificate.

Creating a Hue Truststore File in PEM Format

Server certificates are stored in JKS format and must be converted to PEM. To create the Hue truststore, extract each certificate from its keystore with `keytool`, convert to PEM format with `openssl`, and add to the truststore.

1. Extract the certificate from the keystore of each TLS/SSL-enabled server with which Hue communicates.

   For example, `hadoop-server.jks` contains server certificate, `foo-1.example.com`, and password, `example123`.

```
keytool -exportcert -keystore hadoop-server.jks -alias foo-1.example.com -storepass
example123 -file foo-1.cert
```

2. Convert each certificate into a PEM file.

```
openssl x509 -inform der -in foo-1.cert > foo-1.pem
```

3. Concatenate all the PEM certificates into one PEM file.

```
cat foo-1.pem foo-2.pem foo-n.pem ... > hue_trustore.pem
```

> **Note:** Ensure the final PEM truststore is deployed in a location that is accessible by the Hue service.

Configuring Hue as a TLS/SSL Client with Cloudera Manager

1. Go to the **Hue** service and click the **Configuration** tab.
2. Filter by **Scope** > **Hue Server** and **Category** > **Security**.
3. Find the property, `Hue TLS/SSL Server CA Certificate (PEM Format)`, or `ssl_cacerts`.

4. Enter the path to `<hue_truststore>.pem` on the host running the Hue web server.
5. Click **Save Changes**.
6. Select **Actions** > **Restart** to restart the Hue service.

Configuring Hue as a TLS/SSL Client at the Command Line

For unmanaged deployments only, manually set `ssl_cacerts` in `hue.ini` to the path of the `<hue_truststore>.pem` file:

```
[desktop]
# Path to default Certificate Authority certificates.
ssl_cacerts=/etc/hue/<hue_truststore>.pem
```

Hue as a TLS/SSL Server

Hue and other Python-based services expect certificates and keys to be stored in PEM format. You can manage such services with the openssl tool. To configure Hue to use HTTPS, generate a private key and certificate as described in Configuring TLS Encryption for Cloudera Manager on page 195 and reuse a host's existing Java keystore by converting it to the PEM format. See Converting JKS Key and Certificate to PEM on page 386.

Enabling TLS/SSL for the Hue Server with Cloudera Manager

Minimum Required Role: Configurator (also provided by **Cluster Administrator, Full Administrator**)

1. Go to the **Hue service** and click **Configuration**.
2. Filter by **Scope** > **Hue Server** and **Category** > **Security**.
3. Edit the following **TLS/SSL** properties according to your cluster configuration.

Property	Description
Enable TLS/SSL for Hue	Encrypt communication between clients and Hue with TLS/SSL.
Hue TLS/SSL Server Certificate File (PEM Format) `ssl_certificate`	Path to TLS/SSL certificate on host running Hue web server.
Hue TLS/SSL Server Private Key File (PEM Format `ssl_private_key`	Path to TLS/SSL private key on host running Hue web server.
Hue TLS/SSL Private Key Password `ssl_password`	Password for private key in Hue TLS/SSL Server Certificate and Private Key file.

You can also store `ssl_password` more securely in a script and set this parameter instead:

```
ssl_password_script=<your_hue_passwords_script.sh>
```

For more, see Storing Hue Passwords in a Script on page 216.

To apply this configuration property to other role groups as needed, edit the value for the appropriate role group. See Modifying Configuration Properties Using Cloudera Manager.

4. Click **Save Changes**.
5. Select **Actions** > **Restart** to restart the Hue service.

For more details on configuring Hue with TLS/SSL, see this blog post.

Enabling TLS/SSL for the Hue Server at the Command Line

1. Enable secure session cookies in `hue.ini` under `[desktop]>[[session]]`.

```
[desktop]
  [[session]]
  secure=true
```

2. Edit the following properties in `hue.ini` under `[desktop]`.

```
[desktop]
 ssl_certificate=/path/to/server.cert
 ssl_private_key=/path/to/server.key
 ssl_password=<private_key_password>
```

You can store `ssl_password` more securely in a script and set this parameter instead:

```
ssl_password_script=<your_hue_passwords_script.sh>
```

For more, see Storing Hue Passwords in a Script on page 216.

Enabling Hue TLS/SSL Communication with HiveServer2

In CDH 5.5.x and higher, HiveServer2 is enabled for TLS/SSL communication by default.

By providing a CA certificate, private key, and public certificate, Hue can communicate with HiveServer2 over TLS/SSL. You can now configure the following properties in the `[beeswax]` section under `[[ssl]]` in the Hue configuration file, `hue.ini`.

`enabled`	Choose to enable/disable TLS/SSL communication for this server. Default: `false`
`cacerts`	Path to Certificate Authority certificates. Default: `/etc/hue/cacerts.pem`
`validate`	Choose whether Hue should validate certificates received from the server. Default: `true`

Related Information

- Configuring Encrypted Communication Between HiveServer2 and Client Drivers on page 210

Enabling Hue TLS/SSL Communication with Impala

In CDH 5.5.x and higher, Impala is enabled for TLS/SSL communication by default.

By providing a CA certificate, private key, and public certificate, Hue can communicate with Impala over TLS/SSL. You can configure the following properties in the `[impala]` section under `[[ssl]]` in the Hue configuration file, `hue.ini`.

`enabled`	Choose to enable/disable TLS/SSL communication for this server. Default: `false`
`cacerts`	Path to Certificate Authority certificates. Default: `/etc/hue/cacerts.pem`
`validate`	Choose whether Hue should validate certificates received from the server. Default: `true`

Securing Database Connections using TLS/SSL

Connections vary depending on the database. Hue uses different clients to communicate with each database internally. Client specific options, such as secure connectivity, can be passed through the interface.

For example, for MySQL you can enable TLS/SSL communication by specifying the `options` configuration property under the `desktop>[[database]]` section in `hue.ini`. Here we identify the Certificate Authority (CA) certificate:

```
[desktop]
  [[databases]]
    …
    options={"ssl":{"ca":"/tmp/ca-cert.pem"}}
```

You can also identify public and private keys, for example:

```
options='{"ssl": {"ca": "/tmp/newcerts2/ca.pem", "key": "/tmp/newcerts2/client-key.pem",
 "cert": "/tmp/newcerts2/client-cert.pem"}}'
```

Storing Hue Passwords in a Script

In CDH 5.4, Hue added the ability to store passwords in a secure script and pull passwords from `stdout`. On startup, Hue runs one or more passwords scripts and grabs each password from `stdout`.

In `hue_ini`, add the suffix, `_script`, to any password property and set it equal to the script name. In Cloudera Manager, set these properties in the configuration field, **Hue Service Advanced Configuration Snippet (Safety Valve) for hue_safety_valve.ini**. For example:

```
[desktop]
ldap_username=hueservice
ldap_password_script="/var/lib/hue/<your_hue_passwords_script.sh> ldap_password"
ssl_password_script="/var/lib/hue/<your_hue_passwords_script.sh> ssl_password"

[[ldap]]
bind_password_script="/var/lib/hue/<your_hue_passwords_script.sh> bind_password"

[[database]]
password_script="/var/lib/hue/<your_hue_passwords_script.sh> database"
```

Store the script in a directory that only the hue user can read, write, and execute. You can have one script per password or one script with parameters for all passwords. Here is an example of a script with parameters for multiple passwords:

```
#!/bin/bash

SERVICE=$1

if [[ ${SERVICE} == "ldap_password" ]]
then
   echo "password"
fi

if [[ ${SERVICE} == "ssl_password" ]]
then
   echo "password"
fi

if [[ ${SERVICE} == "bind_password" ]]
then
   echo "Password1"
fi

if [[ ${SERVICE} == "database_password" ]]
then
   echo "password"
fi
```

Note: The bind password parameter was added in CDH 5.4.6.

Configuring TLS/SSL for Impala

Impala supports TLS/SSL network encryption, between Impala and client programs, and between the Impala-related daemons running on different nodes in the cluster. This feature is important when you also use other features such as Kerberos authentication or Sentry authorization, where credentials are being transmitted back and forth.

Important:

- You can use either Cloudera Manager or the following command-line instructions to complete this configuration.
- This information applies specifically to the version of Impala shown in the HTML page header or on the PDF title page. If you use an earlier version of CDH, see the documentation for that version located at Cloudera Documentation.

Using Cloudera Manager

To configure Impala to listen for Beeswax and HiveServer2 requests on TLS/SSL-secured ports:

1. Open the Cloudera Manager Admin Console and go to the **Impala** service.
2. Click the **Configuration** tab.
3. Select **Scope** > **Impala (Service-Wide)**.
4. Select **Category** > **Security**.
5. Edit the following properties:

Table 14: Impala SSL Properties

Property	Description
Enable TLS/SSL for Impala	Encrypt communication between clients (like ODBC, JDBC, and the Impala shell) and the Impala daemon using Transport Layer Security (TLS) (formerly known as Secure Socket Layer (SSL)).
Impala TLS/SSL Server Certificate File (PEM Format)	Local path to the X509 certificate that identifies the Impala daemon to clients during TLS/SSL connections. This file must be in PEM format.
Impala TLS/SSL Server Private Key File (PEM Format)	Local path to the private key that matches the certificate specified in the Certificate for Clients. This file must be in PEM format.
Impala TLS/SSL Private Key Password	The password for the private key in the Impala TLS/SSL Server Certificate and Private Key file. If left blank, the private key is not protected by a password.
Impala TLS/SSL CA Certificate	The location on disk of the certificate, in PEM format, used to confirm the authenticity of SSL/TLS servers that the Impala daemons might connect to. Because the Impala daemons connect to each other, this should also include the CA certificate used to sign all the SSL/TLS Certificates. Without this parameter, SSL/TLS between Impala daemons will not be enabled.
SSL/TLS Certificate for *Impala component* Webserver	There are three of these configuration settings, one each for “Impala Daemon”, “Catalog Server”, and “Statestore”. Each of these Impala components has its own internal web server that powers the associated web UI with diagnostic information. The configuration setting represents the local path to the X509 certificate that identifies the web server to clients during TLS/SSL connections. This file must be in PEM format.

6. Click **Save Changes** to commit the changes.
7. Select each scope, one each for "Impala Daemon", "Catalog Server", and "Statestore", and repeat the above steps. Each of these Impala components has its own internal web server that powers the associated web UI with diagnostic information. The configuration setting represents the local path to the X509 certificate that identifies the web server to clients during TLS/SSL connections.
8. Restart the Impala service.

For information on configuring TLS/SSL communication with the `impala-shell` interpreter, see Configuring TLS/SSL Communication for the Impala Shell on page 218.

Using the Command Line

To enable SSL for when client applications connect to Impala, add both of the following flags to the `impalad` startup options:

- `--ssl_server_certificate`: the full path to the server certificate, on the local filesystem.
- `--ssl_private_key`: the full path to the server private key, on the local filesystem.

In and higher, Impala can also use SSL for its own internal communication between the `impalad`, `statestored`, and `catalogd` daemons. To enable this additional SSL encryption, set the `--ssl_server_certificate` and `--ssl_private_key` flags in the startup options for `impalad`, `catalogd`, and `statestored`, and also add the `--ssl_client_ca_certificate` flag for all three of those daemons.

> **Warning:** Prior to CDH 5.5.2 / Impala 2.3.2, you could enable Kerberos authentication between Impala internal components, or SSL encryption between Impala internal components, but not both at the same time. This restriction has now been lifted. See IMPALA-2598 to see the maintenance releases for different levels of CDH where the fix has been published.

If either of these flags are set, both must be set. In that case, Impala starts listening for Beeswax and HiveServer2 requests on SSL-secured ports only. (The port numbers stay the same; see Ports Used by Impala for details.)

Since Impala uses passphrase-less certificates in PEM format, you can reuse a host's existing Java keystore by converting it to the PEM format. For instructions, see Converting JKS Key and Certificate to PEM on page 386.

Configuring TLS/SSL Communication for the Impala Shell

Typically, a client program has corresponding configuration properties in Cloudera Manager to verify that it is connecting to the right server. For example, with SSL enabled for Impala, you use the following options when starting the `impala-shell` interpreter:

- `--ssl`: enables TLS/SSL for `impala-shell`.
- `--ca_cert`: the local pathname pointing to the third-party CA certificate, or to a copy of the server certificate for self-signed server certificates.

If `--ca_cert` is not set, `impala-shell` enables TLS/SSL, but does not validate the server certificate. This is useful for connecting to a known-good Impala that is only running over TLS/SSL, when a copy of the certificate is not available (such as when debugging customer installations).

For `impala-shell` to successfully connect to an Impala cluster that has the minimum allowed TLS/SSL version set to 1.2 (`--ssl_minimum_version=tlsv1.2`), the Python version on the cluster that `impala-shell` runs on must be 2.7.9 or higher (or a vendor-provided Python version with the required support. Some vendors patched Python 2.7.5 versions on Red Hat Enterprise Linux 7 and derivatives).

Using TLS/SSL with Business Intelligence Tools

You can use Kerberos authentication, TLS/SSL encryption, or both to secure connections from JDBC and ODBC applications to Impala. See Configuring Impala to Work with JDBC and Configuring Impala to Work with ODBC for details.

Prior to CDH 5.7 / Impala 2.5, the Hive JDBC driver did not support connections that use both Kerberos authentication and SSL encryption. If your cluster is running an older release that has this restriction, to use both of these security features with Impala through a JDBC application, use the Cloudera JDBC Connector as the JDBC driver.

Specifying TLS/SSL Minimum Allowed Version and Ciphers

Depending on your cluster configuration and the security practices in your organization, you might need to restrict the allowed versions of TLS/SSL used by Impala. Older TLS/SSL versions might have vulnerabilities or lack certain features. In , you can use startup options for the `impalad`, `catalogd`, and `statestored` daemons to specify a minimum allowed version of TLS/SSL.

Specify one of the following values for the `--ssl_minimum_version` configuration setting:

- `tlsv1`: Allow any TLS version of 1.0 or higher. This setting is the default when TLS/SSL is enabled.
- `tlsv1.1`: Allow any TLS version of 1.1 or higher.
- `tlsv1.2`: Allow any TLS version of 1.2 or higher.

> **Note:**
>
> As of , TLSv1.2 may not work for Impala on RHEL 6 or CentOS 6, even if OpenSSL 1.0.1 is available. The daemons fail to start, with a socket error stating the TLS version is not supported. The underlying cause is related to Red Hat issue 1497859. The issue applies if you build on a RHEL 6 or CentOS 6 system with OpenSSL 1.0.0, and run on a RHEL 6 or CentOS 6 system with OpenSSL 1.0.1.

Along with specifying the version, you can also specify the allowed set of TLS ciphers by using the `--ssl_cipher_list` configuration setting. The argument to this option is a list of keywords, separated by colons, commas, or spaces, and optionally including other notation. For example:

```
--ssl_cipher_list="RC4-SHA,RC4-MD5"
```

By default, the cipher list is empty, and Impala uses the default cipher list for the underlying platform. See the output of `man ciphers` for the full set of keywords and notation allowed in the argument string.

Configuring TLS/SSL for Oozie

Minimum Required Role: Configurator (also provided by **Cluster Administrator, Full Administrator**)

Before You Begin

- Keystores for Oozie must be readable by the `oozie` user. This could be a copy of the Hadoop services' keystore with permissions `0440` and owned by the `oozie` group.
- Truststores must have permissions `0444` (that is, readable by all).
- Specify absolute paths to the keystore and truststore files. These settings apply to all hosts on which daemon roles of the Oozie service run. Therefore, the paths you choose must be valid on all hosts.
- In case there is a DataNode and an Oozie server running on the same host, they can use the same certificate.

For more information on obtaining signed certificates and creating keystores, see Encrypting Data in Transit on page 191. You can also view the upstream documentation located here.

> **Important:**
>
> - This configuration process can be completed using either Cloudera Manager or the command-line instructions.
> - This information applies specifically to CDH 5.14.x. If you use an earlier version of CDH, see the documentation for that version located at Cloudera Documentation.

Using Cloudera Manager

The steps for configuring and enabling Hadoop TLS/SSL for Oozie are as follows:

1. Open the Cloudera Manager Admin Console and go to the **Oozie service**.
2. Click the **Configuration** tab.
3. Select **Scope** > **All**.
4. Select **Category** > **All**.
5. In the Search field, type **TLS/SSL** to show the Oozie TLS/SSL properties.
6. Edit the following TLS/SSL properties according to your cluster configuration.

Table 15: Oozie TLS/SSL Properties

Property	Description
Enable TLS/SSL for Oozie	Check this field to enable TLS/SSL for Oozie.
Oozie TLS/SSL Server Keystore File Location	Location of the keystore file on the local file system.
Oozie TLS/SSL Server JKS Keystore File Password	Password for the keystore.

7. Click **Save Changes**.
8. Restart the Oozie service.

Using the Command Line

To configure the Oozie server to use TLS/SSL:

1. Stop Oozie by running

```
sudo /sbin/service oozie stop
```

2. To enable TLS/SSL, set the MapReduce version that the Oozie server should work with using the `alternatives` command.

> **Note:** The `alternatives` command is only available on RHEL systems. For SLES, Ubuntu and Debian systems, the command is `update-alternatives`.

For RHEL systems, to use YARN with TLS/SSL:

```
alternatives --set oozie-tomcat-conf /etc/oozie/tomcat-conf.https
```

For RHEL systems, to use MapReduce (MRv1) with TLS/SSL:

```
alternatives --set oozie-tomcat-conf /etc/oozie/tomcat-conf.https.mr1
```

> **Important:**
>
> The `OOZIE_HTTPS_KEYSTORE_PASS` variable must be the same as the password used when creating the keystore file. If you used a password other than `password`, you'll have to change the value of the `OOZIE_HTTPS_KEYSTORE_PASS` variable in this file.

3. Start Oozie by running

```
sudo /sbin/service oozie start
```

Connect to the Oozie Web UI using TLS/SSL (HTTPS)

Use `https://oozie.server.hostname:11443/oozie` though most browsers should automatically redirect you if you use `http://oozie.server.hostname:11000/oozie`.

Additional Considerations when Configuring TLS/SSL for Oozie HA

To allow clients to talk to Oozie servers (the target servers) through the load balancer using TLS/SSL, Configure the load balancer for TLS/SSL pass-through, which means the load balancer does not perform encryption/decryption but simply passes traffic from clients and servers to the appropriate target host. See documentation for your load balancer for details.

Configuring TLS/SSL for Solr

Minimum Required Role: Configurator (also provided by **Cluster Administrator, Full Administrator**)

Before You Begin

> **Important:** Enabling TLS on a cluster breaks existing collections. Because of this, you are advised to enable TLS before creating the first collection on the cluster.

- The Solr service must be running.
- Keystores for Solr must be readable by the `solr` user. This could be a copy of the Hadoop services' keystore with permissions `0440` and owned by the `solr` group.
- Truststores must have permissions `0444` (that is, readable by all).
- Specify absolute paths to the keystore and truststore files. These settings apply to all hosts on which daemon roles of the Solr service run. Therefore, the paths you choose must be valid on all hosts.
- In case there is a DataNode and a Solr server running on the same host, they can use the same certificate.

For more information on obtaining signed certificates and creating keystores, see Encrypting Data in Transit on page 191. You can also see the *Enabling SSL* section in the Apache Solr 4.10 Reference Guide (PDF).

> **Important:**
>
> - This configuration process can be completed using either Cloudera Manager or the command-line instructions.
> - This information applies specifically to CDH 5.14.x. If you use an earlier version of CDH, see the documentation for that version located at Cloudera Documentation.

Additional Considerations when Configuring TLS/SSL for Solr HA

To allow clients to talk to Solr servers (the target servers) through the load balancer using TLS/SSL, Configure the load balancer for TLS/SSL pass-through, which means the load balancer does not perform encryption/decryption but simply passes traffic from clients and servers to the appropriate target host. See documentation for your load balancer for details.

Configuring TLS/SSL for Solr Using Cloudera Manager

The steps for configuring and enabling Hadoop TLS/SSL for Search are as follows:

1. Open the Cloudera Manager Admin Console and go to the **Solr service**.
2. Click the **Configuration** tab.
3. Select **Scope** > **All**.
4. Select **Category** > **All**.
5. In the Search field, type **TLS/SSL** to show the Solr TLS/SSL properties.
6. Edit the following properties according to your cluster configuration.

> **Note:** These values must be the same for all hosts running the Solr role.

Table 16: Solr TLS/SSL Properties

Property	Description
Enable TLS/SSL for Solr	Check this field to enable SSL for Solr.
Solr TLS/SSL Server Keystore File Location	The path to the TLS/SSL keystore file containing the server certificate and private key used for TLS/SSL. Used when Solr is acting as a TLS/SSL server. The keystore must be in JKS format.
Solr TLS/SSL Server JKS Keystore File Password	Password for the Solr JKS keystore.
Solr TLS/SSL Certificate Trust Store File	Required in case of self-signed or internal CA signed certificates. The location on disk of the truststore, in .jks format, used to confirm the authenticity of TLS/SSL servers that Solr might connect to. This is used when Solr is the client in a TLS/SSL connection. This truststore must contain the certificate(s) used to sign the service(s) being connected to. If this parameter is not provided, the default list of well-known certificate authorities is used instead.
Solr TLS/SSL Certificate Trust Store Password	The password for the Solr TLS/SSL Certificate Trust Store File. This password is not required to access the truststore: this field can be left blank. This password provides optional integrity checking of the file. The contents of truststores are certificates, and certificates are public information.

7. Click **Save Changes** to commit the changes.
8. Restart the service.

Additional Considerations When Using a Load Balancer TLS/SSL for Solr HA

To configure a load balancer:

1. Go to the Solr service.
2. Click the **Configuration** tab.
3. Select **Scope** > **Solr**.
4. Enter the hostname and port number of the load balancer in the **Solr Load Balancer** property in the format *`hostname:port number`*.

> **Note:**
>
> When you set this property, Cloudera Manager regenerates the keytabs for Solr roles. The principal in these keytabs contains the load balancer hostname.
>
> If there are services that depends on this Solr service, such as Hue, those services use the load balancer to communicate with Solr.

5. Click **Save Changes** to commit the changes.
6. Restart Solr and any dependent services or restart the entire cluster for this configuration to take effect.

Configuring TLS/SSL for Solr Using the Command Line

To configure the Search to use TLS/SSL:

1. Use `solrctl` to modify the `urlScheme` setting to specify `https`. For example:

```
solrctl --zk myZKEnsemble:2181/solr cluster --set-property urlScheme https
```

2. Stop Solr by running

```
sudo service solr-server stop
```

3. Edit `/etc/default/solr` to include the following environment variable settings:

```
SOLR_SSL_ENABLED=true
SOLR_KEYSTORE_PATH=<absolute_path_to_keystore_file>
SOLR_KEYSTORE_PASSWORD=<keystore_password>

#Following required only in case of self-signed or internal CA signed certificates
SOLR_TRUSTSTORE_PATH=<absolute_path_to_truststore_file>
SOLR_TRUSTSTORE_PASSWORD=<truststore_password>
```

4. Start Solr by running

```
sudo service solr-server start
```

Configuring TLS/SSL for the Key-Value Store Indexer Using Cloudera Manager

The steps for configuring and enabling Hadoop TLS/SSL for the Keystore Indexer are as follows:

1. Open the Cloudera Manager Admin Console and go to the **Key-Value Store Indexer**.
2. Click the **Configuration** tab.
3. Select **Scope** > **All**.
4. Select **Category** > **All**.
5. In the Search field, type **TLS/SSL** to show the Solr TLS/SSL properties.
6. Edit the following TLS/SSL properties according to your cluster configuration.

> **Note:** These values must be the same for all hosts running the Key-Value Store Indexer role.

Table 17: Key-Value Store TLS/SSL Properties

Property	Description
HBase Indexer TLS/SSL Certificate Trust Store File	The location on disk of the truststore, in .jks format, used to confirm the authenticity of TLS/SSL servers that HBase Indexer might connect to. This is used when HBase Indexer is the client in a TLS/SSL connection. This truststore must contain the certificate(s) used to sign the service(s) being connected to. If this parameter is not provided, the default list of well-known certificate authorities is used instead.
HBase Indexer TLS/SSL Certificate Trust Store Password (Optional)	The password for the HBase Indexer TLS/SSL Certificate Trust Store File. This password is not required to access the truststore: this field can be left blank. This password provides optional integrity checking of the file. The contents of truststores are certificates, and certificates are public information.

7. Restart the service.

Configuring TLS/SSL for the Key-Value Store Indexer Using the Command Line

For every host running Key-Value Store Indexer server, specify Solr Trust Store details using the `HBASE_INDEXER_OPTS` environmental variable using following Java system properties:

- `-Djavax.net.ssl.trustStore=<absolute_path_to_truststore_file>`
- `-Djavax.net.ssl.trustStorePassword=<truststore_password>` (Optional)

Restart the Key-Value Store Indexer servers to apply these changes.

Configuring TLS/SSL for Flume Using Cloudera Manager

The steps for configuring and enabling Hadoop TLS/SSL for Flume are as follows:

1. Open the Cloudera Manager Admin Console and go to **Flume**.
2. Click the **Configuration** tab.
3. Select **Scope** > **All**.
4. Select **Category** > **All**.
5. In the Search field, type **TLS/SSL** to show the properties.
6. Edit the following SSL properties according to your cluster configuration.

> **Note:** These values must be the same for all hosts running the Flume role.

Table 18: Key-Value Store SSL Properties

Property	Description
Flume TLS/SSL Certificate Trust Store File	The location on disk of the truststore, in .jks format, used to confirm the authenticity of TLS/SSL servers that Flume might connect to. This is used when Flume is the client in a TLS/SSL connection. This truststore must contain the certificate(s) used to sign the service(s) being connected to. If this parameter is not provided, the default list of well-known certificate authorities is used instead.
Flume TLS/SSL Certificate Trust Store Password (Optional)	The password for the Flume TLS/SSL Certificate Trust Store File. This password is not required to access the truststore: this field can be left blank. This password provides optional integrity checking of the file. The contents of truststores are certificates, and certificates are public information.

7. Click **Save Changes** to commit the changes.
8. Restart the service.

Configuring TLS/SSL for Flume Using the Command Line

For every host running Flume agent, specify Solr Trust Store details using the `FLUME_AGENT_JAVA_OPTS` environmental variable using following Java system properties:

- `-Djavax.net.ssl.trustStore=<absolute_path_to_truststore_file>`
- `-Djavax.net.ssl.trustStorePassword=<truststore_password>` (Optional)

Restart the Flume agents to apply these changes.

Spark Encryption

Spark supports the following means of encrypting Spark data at rest, and data in transit.

Enabling Spark Encryption Using Cloudera Manager

Minimum Required Role: Configurator (also provided by **Cluster Administrator, Full Administrator**)

1. Open the Cloudera Manager Admin Console and go to the **Spark** service.
2. Click the **Configuration** tab.
3. **(Prerequisite)** Search for the **Spark Authentication** property and make sure it has been enabled. If this property is not set, the following settings to enable encryption will not work.

4. Search for the **Enable Network Encryption** property. Use the checkbox to enable encrypted communication between Spark processes belonging to the same application.
5. Search for the **Enable I/O Encryption** property. Use the checkbox to enabled encryption for temporary shuffle and cache files stored by Spark on local disks.
6. Click **Save Changes** to commit the changes.
7. Redeploy client configuration.
8. Restart stale services (if indicated by Cloudera Manager).

Enabling Spark Encryption on an Unmanaged Cluster

> **Important:** If you are using Cloudera Manager, do not manually edit the Spark configuration files to configure the properties listed in this section. Use the steps at Enabling Spark Encryption Using Cloudera Manager on page 224 instead. To configure a property that is not yet available in Cloudera Manager, use **Advanced Configuration Snippets** (or Safety Valves).

Prerequisite - Before enabling encryption, make sure `spark.authenticate` is set to `true`. Without authentication enabled, the following settings to enable encryption will not work.

Enabling Encryption for Shuffle and Cache Files

Configure the following properties to enable encrypted shuffle for Spark on YARN.

Property	Description
`spark.shuffle.encryption.enabled`	Enable encryption of temporary shuffle and cache files.
`spark.shuffle.encryption.keySizeBits`	Shuffle file encryption key size in bits. The valid numbers include 128, 192, and 256.

Enabling Encryption for Spark RPCs

Configure the following property to enable encryption for Spark RPCs.

Property	Default Value	Description
`spark.authenticate.enableSaslEncryption`	`false`	Enable encryption for Spark RPCs.

If you are using an external shuffle service, configure the following property in the shuffle service configuration to disable unencrypted connections. Note that the external shuffle service is enabled by default in CDH 5.5 and higher.

Property	Default Value	Description
`spark.network.sasl.serverAlwaysEncrypt`	`false`	Disable unencrypted connections for the external shuffle service.

Configuring TLS/SSL for HttpFS

> **Important:**
> - This configuration process can be completed using either Cloudera Manager or the command-line instructions.
> - This information applies specifically to CDH 5.14.x. If you use an earlier version of CDH, see the documentation for that version located at Cloudera Documentation.

Using Cloudera Manager

Minimum Required Role: Configurator (also provided by **Cluster Administrator, Full Administrator**)

The steps for configuring and enabling TLS/SSL for HttpFS using Cloudera Manager are as follows:

1. Go to the **HDFS service**
2. Click the **Configuration** tab.
3. Select **Scope** > **HttpFS**.
4. **Select** > **Security**.
5. Edit the following TLS/SSL properties according to your cluster configuration:

Table 19: HttpFS TLS/SSL Properties

Property	Description
Use TLS/SSL	Use TLS/SSL for HttpFS.
HttpFS Keystore File	Location of the keystore file used by the HttpFS role for TLS/SSL. **Default:** `/var/run/hadoop-httpfs/.keystore`. Note that the default location for the keystore file is on non-persistent disk.
HttpFS Keystore Password	Password of the keystore used by the HttpFS role for TLS/SSL. If the keystore password has a percent sign, it must be escaped. For example, for a password that is `pass%word`, use `pass%%word`.
HttpFS TLS/SSL Certificate Trust Store File	The location on disk of the truststore, in `.jks` format, used to confirm the authenticity of TLS/SSL servers that HttpFS might connect to. This is used when HttpFS is the client in a TLS/SSL connection.
HttpFS TLS/SSL Certificate Trust Store Password	The password for the HttpFS TLS/SSL Certificate Trust Store File. This password is not required to access the truststore; this field can be left blank. If the truststore password has a percent sign, it must be escaped. For example, for a password that is `pass%word`, use `pass%%word`.

6. Click **Save Changes**.
7. Restart the HDFS service.

Connect to the HttpFS Web UI using TLS/SSL (HTTPS)

Use `https://<httpfs_server_hostname>:14000/webhdfs/v1/`, though most browsers should automatically redirect you if you use `http://<httpfs_server_hostname>:14000/webhdfs/v1/`

Using the Command Line

Configure the HttpFS Server to use TLS/SSL (HTTPS)

1. Stop HttpFS by running

```
sudo /sbin/service hadoop-httpfs stop
```

2. To enable TLS/SSL, change which configuration the HttpFS server should work with using the `alternatives` command.

> **Note:** The `alternatives` command is only available on RHEL systems. For SLES, Ubuntu and Debian systems, the command is `update-alternatives`.

For RHEL systems, to use TLS/SSL:

```
alternatives --set hadoop-httpfs-tomcat-conf /etc/hadoop-httpfs/tomcat-conf.https
```

> **Important:**
>
> The `HTTPFS_TLS/SSL_KEYSTORE_PASS` variable must be the same as the password used when creating the keystore file. If you used a password other than `password`, you'll have to change the value of the `HTTPFS_TLS/SSL_KEYSTORE_PASS` variable in `/etc/hadoop-httpfs/conf/httpfs-env.sh`.

3. Start HttpFS by running

```
sudo /sbin/service hadoop-httpfs start
```

Connect to the HttpFS Web UI using TLS/SSL (HTTPS)

Use `https://<httpfs_server_hostname>:14000/webhdfs/v1/`, though most browsers should automatically redirect you if you use `http://<httpfs_server_hostname>:14000/webhdfs/v1/`

> **Important:**
>
> If using a Self-Signed Certificate, your browser will warn you that it cannot verify the certificate or something similar. You will probably have to add your certificate as an exception.

Encrypted Shuffle and Encrypted Web UIs

> **Important:**
>
> - If you use Cloudera Manager, do not use these command-line instructions. For the Cloudera Manager instructions, see Configuring TLS/SSL for HDFS, YARN and MapReduce on page 204.
> - This information applies specifically to CDH 5.14.0. If you use a lower version of CDH, see the documentation for that version located at Cloudera Documentation.

CDH 5 supports encryption of the MapReduce shuffle phase for both MapReduce v1 (MRv1) and MapReduce v2 (MRv2), also known as YARN. CDH also supports enabling TLS/SSL for the MRv1 and YARN web UIs, with optional client authentication (also known as bi-directional HTTPS, or HTTPS with client certificates). The configuration properties required to enable these features have been combined. In most cases, these properties are common to both MRv1 and YARN. They include:

- `hadoop.ssl.enabled`:
  - Toggles the shuffle for MRv1 between HTTP and HTTPS.
  - Toggles the MRv1 and YARN web UIs between HTTP and HTTPS.
- `mapreduce.shuffle.ssl.enabled`: Toggles the shuffle for YARN between HTTP and HTTPS.

  By default, this property is not specified in `mapred-site.xml`, and YARN encrypted shuffle is controlled by the value of `hadoop.ssl.enabled`. If this property is set to `true`, encrypted shuffle is enabled for YARN. Note that you cannot successfully enable encrypted shuffle for YARN by only setting this property to `true`, if `hadoop.ssl.enabled` is still set to `false`.
- Configuration settings for specifying keystore and truststore properties that are used by the MapReduce shuffle service, the Reducer tasks that fetch shuffle data, and the web UIs.
- `ssl.server.truststore.reload.interval`: A configuration property to reload truststores across the cluster when a node is added or removed.

> **Important:**
>
> When the web UIs are served over HTTPS, you must specify `https://` as the protocol. There is no redirection from `http://`. If you attempt to access an HTTPS resource over HTTP, your browser will show an empty screen with no warning.

Configuring Encrypted Shuffle and Encrypted Web UIs

Configure encryption for the MapReduce shuffle, and the MRv1 and YARN web UIs, as follows:

Enable encrypted shuffle for MRv1, and encryption for the MRv1 and YARN web UIs (core-site.xml)

Set the following properties in the `core-site.xml` files of all nodes in the cluster.

`hadoop.ssl.enabled`

Default value: `false`

For MRv1, set this value to `true` to enable encryption for both the MapReduce shuffle and the web UI.

For YARN, this property enables encryption for the web UI only . Enable shuffle encryption with a property in the `mapred-site.xml` file as described here.

`hadoop.ssl.require.client.cert`

Default value: `false`

When this property is set to `true`, client certificates are required for all shuffle operations and all browsers used to access web UIs.

Cloudera recommends that this be set to `false`. This is because client certificates are easily susceptible to attacks from malicious clients or jobs. For more details, see Client Certificates on page 232.

`hadoop.ssl.hostname.verifier`

Default value: `DEFAULT`

The `SSLHostnameVerifier` interface present inside the hadoop-common security library checks if a hostname matches the name stored inside the server's X.509 certificate. The value assigned to this property determines how Hadoop verifies hostnames when it establishes new `HttpsURLConnection` instances. Valid values are:

- `DEFAULT`: The hostname must match either the first common name (CN) or any of the subjectAltNames (SAN). Wildcards can occur in either the CN or the SANs. For example, a hostname, such as `*.example.com`, will match all subdomains, including `test.cloudera.example.com`.
- `DEFAULT_AND_LOCALHOST`: This verifier mechanism works just like `DEFAULT`. However, it also allows all hostnames of the type: `localhost`, `localhost.example`, or `127.0.0.1`.
- `STRICT`: This verifier works just like `DEFAULT` with an additional restriction for hostnames with wildcards. For example, a hostname with a wildcard such as `*.example.com`, will only match subdomains at the same level. Hence, `cloudera.example.com` will match, but, unlike `DEFAULT`, `test.cloudera.example.com` will be rejected.
- `STRICT_IE6`: This verifier works just like `STRICT`, however, it will allow hostnames that match any of the common names (CN) within the server's X.509 certificate, not just the first one.
- `ALLOW_ALL`: Using this verifier will essentially turn off the hostname verifier mechanism.

`hadoop.ssl.keystores.factory.class`

Default value: `org.apache.hadoop.security.ssl.FileBasedKeyStoresFactory`

The `KeyStoresFactory` implementation to be used. Currently, `FileBasedKeyStoresFactory` is the only implementation of `KeyStoresFactory`.

`hadoop.ssl.server.conf`

Default value: `ssl-server.xml`

Resource file from which TLS/SSL server keystore information is extracted. Typically, it should be in the `/etc/hadoop/conf/` directory so that it can be looked up in the `CLASSPATH`.

`hadoop.ssl.client.conf`

Default value: `ssl-client.xml`

Resource file from which TLS/SSL client keystore information is extracted. Typically, it should be in the `/etc/hadoop/conf/` directory so that it can be looked up in the `CLASSPATH`.

Set the `<final>` field for all these properties to `true` as in the following sample configuration snippet:

```
...
    <property>
      <name>hadoop.ssl.require.client.cert</name>
      <value>false</value>
      <final>true</final>
    </property>

    <property>
      <name>hadoop.ssl.hostname.verifier</name>
      <value>DEFAULT</value>
      <final>true</final>
    </property>

    <property>
      <name>hadoop.ssl.keystores.factory.class</name>
      <value>org.apache.hadoop.security.ssl.FileBasedKeyStoresFactory</value>
      <final>true</final>
    </property>

    <property>
      <name>hadoop.ssl.server.conf</name>
      <value>ssl-server.xml</value>
      <final>true</final>
    </property>

    <property>
      <name>hadoop.ssl.client.conf</name>
      <value>ssl-client.xml</value>
      <final>true</final>
    </property>

    <property>
      <name>hadoop.ssl.enabled</name>
      <value>true</value>
    </property>
...
```

Enable encrypted shuffle for YARN (mapred-site.xml)

To enable encrypted shuffle for YARN, set the following property in the `mapred-site.xml` file on every node in the cluster:

`mapreduce.shuffle.ssl.enabled`

Default value: Not specified

By default, this property is not specified in `mapred-site.xml`, and YARN encrypted shuffle is controlled by the value of `hadoop.ssl.enabled`. If this property is set to `true`, encrypted shuffle is enabled for YARN. Note that you cannot successfully enable encrypted shuffle for YARN by only setting this property to `true`, if `hadoop.ssl.enabled` is still set to `false`.

Set the `<final>` field for this property to `true` as in the following configuration snippet:

```
...
    <property>
      <name>mapreduce.shuffle.ssl.enabled</name>
      <value>true</value>
      <final>true</final>
    </property>
...
```

Configure the keystore and truststore for the Shuffle server (ssl-server.xml)

> **Note:** To run job tasks so they are prevented from reading the server keystore and gaining access to the shuffle server certificates:
>
> - Configure the Linux Task Controller for MRv1
> - Configure the Linux Container Executor for YARN

Currently, `FileBasedKeyStoresFactory` is the only implementation of `KeyStoresFactory`. It uses properties in the `ssl-server.xml` and `ssl-client.xml` files to configure the keystores and truststores.

The `ssl-server.xml` should be owned by the `hdfs` or `mapred` Hadoop system user, belong to the `hadoop` group, and it should have 440 permissions. Regular users should not belong to the `hadoop` group.

Use the following settings to configure the keystores and truststores in the `ssl-server.xml` file.

Property	Default Value	Description
`ssl.server.keystore.type`	`jks`	Keystore file type
`ssl.server.keystore.location`	NONE	Keystore file location. The `mapred` user must own this file and have exclusive read access to it.
`ssl.server.keystore.password`	NONE	Keystore file password
`ssl.server.keystore.keypassword`	NONE	Key password
`ssl.server.truststore.type`	`jks`	Truststore file type
`ssl.server.truststore.location`	NONE	Truststore file location. The `mapred` user must own this file and have exclusive read access to it.
`ssl.server.truststore.password`	NONE	Truststore file password
`ssl.server.truststore.reload.interval`	10000	Truststore reload interval, in milliseconds

Sample `ssl-server.xml`

```
<configuration>
<!-- Server Certificate Store -->
<property>
  <name>ssl.server.keystore.type</name>
     <value>jks</value>
</property>
<property>
  <name>ssl.server.keystore.location</name>
  <value>${user.home}/keystores/server-keystore.jks</value>
</property>
<property>
  <name>ssl.server.keystore.password</name>
  <value>serverfoo</value>
```

```
</property>
<property>
  <name>ssl.server.keystore.keypassword</name>
  <value>serverfoo</value>
</property>

<!-- Server Truststore -->
<property>
  <name>ssl.server.truststore.type</name>
  <value>jks</value>
</property>
<property>
  <name>ssl.server.truststore.location</name>
  <value>${user.home}/keystores/truststore.jks</value>
</property>
<property>
  <name>ssl.server.truststore.password</name>
  <value>clientserverbar</value>
</property>
<property>
  <name>ssl.server.truststore.reload.interval</name>
  <value>10000</value>
</property>
</configuration>
```

Configure the keystore and truststore for the Reducer/Fetcher (ssl-client.xml)

Use the following settings to configure the keystore and truststore in the `ssl-client.xml` file. This file must be owned by the `mapred` user for MRv1 and by the `yarn` user for YARN. The file permissions should be 444 (read access for all users).

Property	Default Value	Description
`ssl.client.keystore.type`	`jks`	Keystore file type
`ssl.client.keystore.location`	NONE	Keystore file location. The `mapred` user must own this file and should have read access to it.
`ssl.client.keystore.password`	NONE	Keystore file password
`ssl.client.keystore.keypassword`	NONE	Key password
`ssl.client.truststore.type`	`jks`	Truststore file type
`ssl.client.truststore.location`	NONE	Truststore file location. The `mapred` user must own this file and should have read access to it.
`ssl.client.truststore.password`	NONE	Truststore file password
`ssl.client.truststore.reload.interval`	10000	Truststore reload interval, in milliseconds

Sample `ssl-client.xml`

```
<configuration>
  <!-- Client Certificate Store -->
  <property>
    <name>ssl.client.keystore.type</name>
    <value>jks</value>
  </property>
  <property>
    <name>ssl.client.keystore.location</name>
    <value>${user.home}/keystores/client-keystore.jks</value>
  </property>
  <property>
    <name>ssl.client.keystore.password</name>
```

```
    <value>clientfoo</value>
  </property>
  <property>
    <name>ssl.client.keystore.keypassword</name>
    <value>clientfoo</value>
  </property>

  <!-- Client Truststore -->
  <property>
    <name>ssl.client.truststore.type</name>
    <value>jks</value>
  </property>
  <property>
    <name>ssl.client.truststore.location</name>
    <value>${user.home}/keystores/truststore.jks</value>
  </property>
  <property>
    <name>ssl.client.truststore.password</name>
    <value>clientserverbar</value>
  </property>
  <property>
    <name>ssl.client.truststore.reload.interval</name>
    <value>10000</value>
  </property>
</configuration>
```

Activating Encrypted Shuffle

> **Important:**
>
> Encrypted shuffle has a significant performance impact. You should benchmark this before implementing it in production. In many cases, one or more additional cores are needed to maintain performance.

When you have made the configuration changes described in the previous section, activate Encrypted Shuffle by re-starting all TaskTrackers in MRv1 and all NodeManagers in YARN.

Client Certificates

Client Certificates are supported but they do not guarantee that the client is a reducer task for the job. The Client Certificate keystore file that contains the private key must be readable by all users who submit jobs to the cluster, which means that a rogue job could read those keystore files and use the client certificates in them to establish a secure connection with a Shuffle server. The JobToken mechanism that the Hadoop environment provides is a better protector of the data; each job uses its own JobToken to retrieve only the shuffle data that belongs to it. Unless the rogue job has a proper JobToken, it cannot retrieve Shuffle data from the Shuffle server.

However, if your cluster requires client certificates, ensure that browsers connecting to the web UIs are configured with appropriately signed certificates. If your certificates are signed by a certificate authority (CA), make sure you include the complete chain of CA certificates in the server's keystore.

Reloading Truststores

By default, each truststore reloads its configuration every 10 seconds. If you bring in a new truststore file to replace an old one, when the truststore is reloaded, the new certificates will be override the previous ones. If a client certificate is added to (or removed from) all the truststore files in the system, both YARN and MRv1 will pick up the new configuration without requiring that the TaskTracker or NodeManager daemons are restarted. This mechanism is useful for adding or removing nodes from the cluster, or for adding or removing trusted clients.

The reload interval is controlled by the `ssl.client.truststore.reload.interval` and `ssl.server.truststore.reload.interval` configuration properties in the `ssl-client.xml` and `ssl-server.xml` files described here.

> **Note:** The keystores are not automatically reloaded. To change a keystore for a TaskTracker in MRv1 or a NodeManager in YARN, you must restart the TaskTracker or NodeManager daemon.

Debugging

> **Important:** Enable debugging only for troubleshooting, and only for jobs running on small amounts of data. Debugging is very verbose and slows jobs down significantly. You may need to increase the value for the `mapred.task.timeout` property to prevent jobs from failing for taking too long.

To enable TLS/SSL debugging in the reducers, set the `mapred.reduce.child.java.opts` property as follows. You can do this on a per-job basis, or by means of a cluster-wide setting in `mapred-site.xml`.:

```
<configuration>
...
    <property>
        <name>mapred.reduce.child.java.opts</name>
        <value>-Xmx200m -Djavax.net.debug=all</value>
    </property>
...
</configuration>
```

To enable debugging for MRv1 TaskTrackers, edit `hadoop-env.sh` as follows:

```
HADOOP_TASKTRACKER_OPTS="-Djavax.net.debug=all $HADOOP_TASKTRACKER_OPTS"
```

To enable debugging for YARN NodeManagers for YARN, edit `yarn-env.sh` as follows:

```
YARN_OPTS="-Djavax.net.debug=all $YARN_OPTS"
```

Configuring TLS/SSL for Navigator Audit Server

Cloudera Navigator supports TLS/SSL encryption for network communications between the Navigator Audit Server and clients, such as the Cloudera Manager. Typically, TLS/SSL is configured for the entire cluster, so it is likely that the server key and certificate already exist on the specific host running the Navigator Audit Server role. See Configuring TLS Encryption for Cloudera Manager on page 195 for more information about configuring TLS/SSL for Cloudera Manager clusters.

> **Important:** When configuring Navigator to use TLS, the **Cloudera Manager TLS/SSL Client Trust Store Password** is required. If not provided, Navigator Audit Server isn't able to communicate with Cloudera Manager and a NullPointerException error appears in the Cloudera Manager server log.

1. Log in to the Cloudera Manager Admin Console.
2. Select **Clusters** > **Cloudera Management Service**.
3. Click the **Configuration** tab.
4. Select **Scope** > **Navigator Audit Server**.
5. Select **Category** > **Security**.
6. Edit the following properties according to your cluster configuration.

Property	Description
Enable TLS/SSL for Navigator Audit Server	Encrypt network communications between clients and Navigator Audit Server using TLS/SSL.
TLS/SSL Keystore File Location	The path to the keystore file containing the server private key and certificate. The keystore must be in JKS format.

Property	Description
TLS/SSL Keystore File Password	The password for the Navigator Audit Server JKS keystore file.
TLS/SSL Keystore Key Password	The password for the private key contained in the JKS keystore.
Navigator TLS/SSL Certificate Trust Store File	The path to the trust store. The trust store is used when Navigator is the client in a TLS/SSL connection. This trust store must contain the certificate(s) used to sign the service(s) connected to. If this parameter is not provided, the default list of well-known certificate authorities is used instead.
Navigator TLS/SSL Certificate Trust Store Password	The password for the Navigator TLS/SSL Certificate Trust Store File. This password is not required to access the trust store; this field can be left blank. This password provides optional integrity checking of the file.

7. Click **Save Changes**.
8. Restart the Navigator Audit Server role.

> **Note:** After TLS/SSL is enabled, Cloudera Manager links to the Cloudera Navigator console use HTTPS rather than unencrypted HTTP.

Configuring TLS/SSL for Navigator Metadata Server

Cloudera Navigator supports TLS/SSL encryption for network communications between the Navigator Metadata Server and clients, such as the web browser used for Cloudera Navigator console. Typically, TLS/SSL is configured for the entire cluster, so it is possible that the server key and certificate already exist on the specific host running the Navigator Metadata Server role. The assumption in the steps below is that the cluster is already configured for TLS/SSL and the security artifacts have already been obtained and deployed to the host running the Navigator Metadata Server role instance.

1. Validate that the JDK on the Navigator Metadata Server host has certificates configured.

   Navigator establishes trust for the TLS certificate used for network communication using the JDK on the local host. The JDK looks for certificates in three places:

   - `$JAVA_HOME/jre/lib/security/jssecacerts`
   - `$JAVA_HOME/jre/lib/security/cacerts`
   - Location specified by `java.net.ssl.truststore`. This property can be set through the **Navigator Metadata Server Advanced Configuration Snippet (Safety Valve) for cloudera-navigator.properties**.

   If the `java.net.ssl.truststore` location is set, the JDK ignores the default certificate locations.

   If there isn't a certificate in any of these locations, review Generate TLS Certificates on page 195 and make sure that the Navigator Metadata Server host was included in step 7 on page 197.

2. Validate the location of the keystore file on the host running the Navigator Metadata Server role instance.

   For example, if TLS was enabled according to the instructions Configuring TLS Encryption for Cloudera Manager on page 195, the TLS keystore would be located in `/opt/cloudera/security/pki`. If the keystore isn't already on the Navigator Metadata Server host, generate one from the instructions in Generate TLS Certificates on page 195.

3. Log in to the Cloudera Manager Admin Console.
4. Select **Clusters** > **Cloudera Management Service**.
5. Click the **Configuration** tab.
6. Select **Scope** > **Navigator Metadata Server**.
7. Select **Category** > **Security**.

8. Edit the following properties according to your cluster configuration.

Property	Description
Enable TLS/SSL for Navigator Metadata Server	Encrypt network communications between clients and Navigator Metadata Server using TLS/SSL.
TLS/SSL Keystore File Location	The path to the keystore file containing the server private key and certificate. The keystore must be in JKS format.
TLS/SSL Keystore File Password	The password for the Navigator Metadata Server JKS keystore file.
TLS/SSL Keystore Key Password	The password for the private key contained in the JKS keystore.

9. Click **Save Changes**.
10. Restart the Navigator Metadata Server role.

Note: After TLS/SSL is enabled, Cloudera Manager links to the Cloudera Navigator console use HTTPS rather than unencrypted HTTP.

Configuring TLS/SSL for Kafka (Navigator Event Broker)

To enable TLS/SSL encryption between Navigator Audit Server and Kafka for publishing audit events to Kafka:

1. Log in to the Cloudera Manager Admin Console.
2. Select **Clusters** > **Kafka**.
3. Click the **Configuration** tab.
4. Select **Kafka Broker** for the **Scope** filter.
5. Select **Security** for the **Category** filter.
6. Enter the following properties according to your cluster configuration.

Property	Description
Enable TLS/SSL for Kafka Broker	Select the checkbox to enable TLS/SSL for encrypted communication between clients and the Kafka Broker service.
Kafka Broker TLS/SSL Certificate Trust Store File	Enter the path (location on disk) to the JKS truststore. Leave this field empty to have the list of well-known CAs checked to provide a chain of proof for the Navigator Audit Server.
Kafka Broker TLS/SSL Certificate Trust Store Password	The truststore does not need password protection. Its contents are public certificates already included in the default Java truststore.

7. Click **Save Changes**.
8. Restart the Kafka service.

Configuring Encrypted Transport for HDFS

This topic describes how to configure encrypted HDFS data transport using both, Cloudera Manager, and the command line.

You must enable Kerberos before configuring encrypted HDFS data transport. See Authentication on page 40 for instructions.

Using Cloudera Manager

Minimum Required Role: Full Administrator

To enable encryption of data transferred between DataNodes and clients, and among DataNodes, perform the following steps:

1. Enable Hadoop security using Kerberos.
2. Select the HDFS service.
3. Click the **Configuration** tab.
4. Select **Scope** > **HDFS (Service Wide)**.
5. Select **Category** > **Security**.
6. Configure the following properties: (You can type the property name in the **Search** box to locate the property.)

Property	Description
Enable Data Transfer Encryption	Check this field to enable wire encryption.
Data Transfer Encryption Algorithm	Optionally configure the algorithm used to encrypt data.
Hadoop RPC Protection	Select **privacy**.

7. Click **Save Changes**.
8. Restart the HDFS service.

Using the Command Line

Important:

- This configuration process can be completed using either Cloudera Manager or the command-line instructions.
- This information applies specifically to CDH 5.14.x. If you use an earlier version of CDH, see the documentation for that version located at Cloudera Documentation.

To enable encrypted data transport using the command line, perform the following steps:

1. Enable Kerberos authentication, following these instructions.
2. Set the optional RPC encryption by setting `hadoop.rpc.protection` to "`privacy`" in the `core-site.xml` file in both client and server configurations.

   Note:

   If RPC encryption is not enabled, transmission of other HDFS data is also insecure.

3. Set `dfs.encrypt.data.transfer` to `true` in the `hdfs-site.xml` file on all server systems.
4. Restart all daemons.

Configuring Encrypted Transport for HBase

This topic describes how to configure encrypted HBase data transport using Cloudera Manager and the command line.

Configuring Encrypted HBase Data Transport Using Cloudera Manager

Minimum Required Role: Full Administrator

To enable encryption of data transferred between HBase masters and RegionServers and between RegionServers and clients:

1. Enable Hadoop security using Kerberos.
2. Configure Kerberos authentication for HBase.
3. Select the HBase service.
4. Click the **Configuration** tab.
5. Select **Scope** > **HBase (Service Wide)**.
6. Select **Category** > **Security**.
7. Search for the HBase Transport Security property and select one of the following:

Property	Description
authentication	Enables simple authentication using Kerberos.
integrity	Checks the integrity of data received to ensure it was not corrupted in transit. Selecting `integrity` also enables `authentication`.
privacy	Ensures privacy by encrypting the data in transit using TLS/SSL encryption. Selecting `privacy` also enables `authentication` and `integrity`.

8. Click **Save Changes**.
9. Restart the HBase service.

Configuring Encrypted HBase Data Transport Using the Command Line

Important:

- This configuration process can be completed using either Cloudera Manager or the command-line instructions.
- This information applies specifically to CDH 5.14.x. If you use an earlier version of CDH, see the documentation for that version located at Cloudera Documentation.

1. Enable Hadoop Security using Kerberos.
2. Enable HBase security using Kerberos.
3. Enable RPC encryption by setting `hbase.rpc.protection` in the `hbase-site.xml` file to one of the following:

Property	Description
authentication	Enables simple authentication using Kerberos.
integrity	Checks the integrity of data received to ensure it was not corrupted in transit. Selecting `integrity` also enables `authentication`.
privacy	Ensures privacy by encrypting the data in transit using TLS/SSL encryption. Selecting `privacy` also enables `authentication` and `integrity`.

4. Restart all daemons.

Encrypting Data at Rest

Cloudera clusters can use a combination of data at rest encryption mechanisms, including HDFS transparent encryption and Cloudera Navigator Encrypt. Both of these data at rest encryption mechanisms can be augmented with key management using Cloudera Navigator Key Trustee Server on page 299 and Cloudera Navigator Key HSM on page 318.

Before configuring encryption for data at rest, familiarize yourself with the Data at Rest Encryption Requirements in the Product Compatibility Matrix.

Data at Rest Encryption Reference Architecture

The following diagram illustrates the supported architecture for deploying Cloudera Navigator encryption for data at rest:

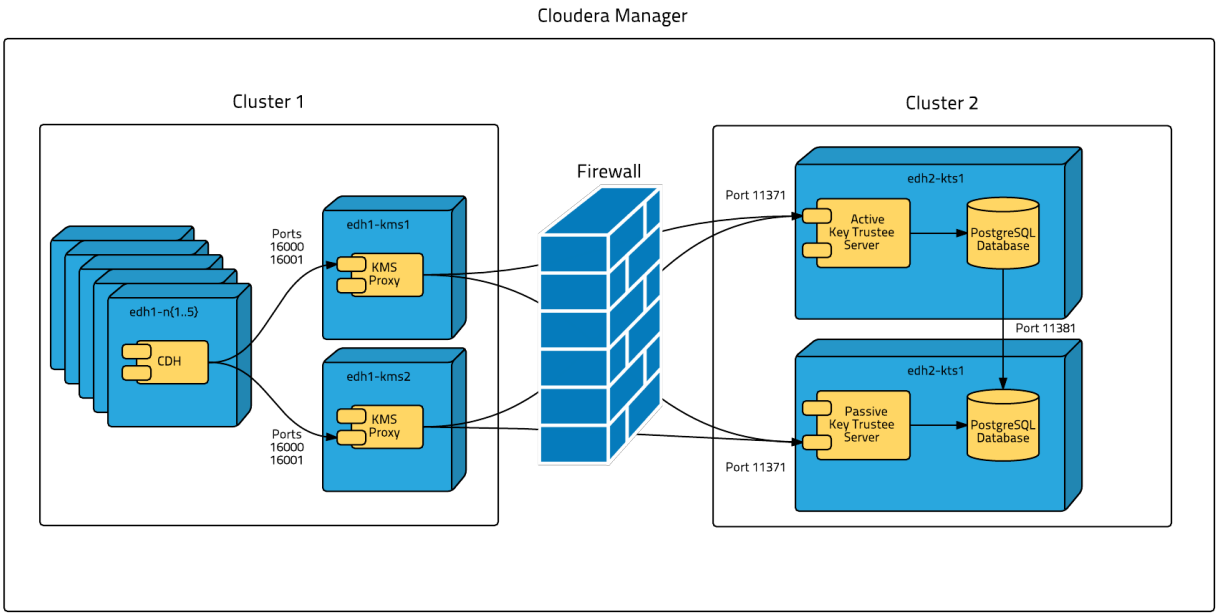

To isolate Key Trustee Server from other enterprise data hub (EDH) services, you must deploy Key Trustee Server on dedicated hosts in a separate cluster in Cloudera Manager. Deploy Key Trustee KMS on dedicated hosts in the same cluster as the EDH services that require access to Key Trustee Server. This provides the following benefits:

- You can restart your EDH cluster without restarting Key Trustee Server, avoiding interruption to other clusters or clients that use the same Key Trustee Server instance.
- You can manage the Key Trustee Server upgrade cycle independently of other cluster components.
- You can limit access to the Key Trustee Server hosts to authorized key administrators only, reducing the attack surface of the system.
- Resource contention is reduced. Running Key Trustee Server and Key Trustee KMS services on dedicated hosts prevents other cluster services from reducing available resources (such as CPU and memory) and creating bottlenecks.

If you are using virtual machines for the Key Trustee Server or Key Trustee KMS hosts, see Navigator HSM KMS HA Planning on page 239.

Data at Rest Encryption Requirements

Encryption comprises several components, each with its own requirements.

Data at rest encryption protection can be applied at a number of levels within Hadoop:

- OS filesystem-level
- Network-level
- HDFS-level (protects both data at rest and in transit)

For detailed data at rest encryption requirements, see the Data at Rest Encryption Requirements in the Product Compatibility Matrix.

For the individual compatibility matrices for each Cloudera Navigator encryption component, see Product Compatibility Matrix for Cloudera Navigator Encryption.

Resource Planning for Data at Rest Encryption

For production environments, you must configure high availability for:

- Key Trustee Server
- Key Trustee KMS
- Navigator HSM KMS

Key Trustee Server and Key Trustee KMS HA Planning

For high availability, you must provision two dedicated Key Trustee Server hosts and at least two dedicated Key Trustee KMS hosts, for a minimum of four separate hosts. Do not run multiple Key Trustee Server or Key Trustee KMS services on the same physical host, and do not run these services on hosts with other cluster services. Doing so causes resource contention with other important cluster services and defeats the purpose of high availability. See Data at Rest Encryption Reference Architecture on page 238 for more information.

The Key Trustee KMS workload is CPU intensive. Cloudera recommends using machines with capabilities equivalent to your NameNode hosts, with Intel CPUs that support AES-NI for optimum performance.

Make sure that each host is secured and audited. Only authorized key administrators should have access to them. Red Hat provides security guides for RHEL:

- RHEL 6 Security Guide
- RHEL 7 Security Guide

For hardware sizing information, see Data at Rest Encryption Requirements on page 239 for recommendations for each Cloudera Navigator encryption component.

For Cloudera Manager deployments, deploy Key Trustee Server in its own dedicated cluster. Deploy Key Trustee KMS in each cluster that uses Key Trustee Server. See Data at Rest Encryption Reference Architecture on page 238 for more information.

For information about enabling Key Trustee Server high availability, refer to Configuring Key Trustee Server High Availability Using Cloudera Manager or Configuring Key Trustee Server High Availability Using the Command Line.

For information about enabling Key Trustee KMS high availability, refer to Enabling Key Trustee KMS High Availability.

Navigator HSM KMS HA Planning

For Navigator HSM KMS high availability, you need to provision two dedicated HSM KMS hosts only.

Make sure that each host is secured and audited. Only authorized key administrators should have access to them. Red Hat provides security guides for RHEL:

- RHEL 6 Security Guide

- RHEL 7 Security Guide

For hardware sizing information, see Data at Rest Encryption Requirements on page 239 for recommendations for each Cloudera Navigator encryption component.

For information about enabling HSM KMS high availability, refer to Enabling Navigator HSM KMS High Availability.

Virtual Machine Considerations

If you are using virtual machines, make sure that the resources (such as virtual disks, CPU, and memory) for each Key Trustee Server and Key Trustee KMS host are allocated to separate physical hosts. Hosting multiple services on the same physical host defeats the purpose of high availability, because a single machine failure can take down multiple services.

To maintain the security of the cryptographic keys, make sure that all copies of the virtual disk (including any back-end storage arrays, backups, snapshots, and so on) are secured and audited with the same standards you apply to the live data.

HDFS Transparent Encryption

Data encryption is mandatory for many government, financial, and regulatory entities, worldwide, to meet privacy and other security requirements. For example, the card payment industry has adopted the Payment Card Industry Data Security Standard (PCI DSS) for information security. Other examples include requirements imposed by United States government's Federal Information Security Management Act (FISMA) and Health Insurance Portability and Accountability Act (HIPAA). Encrypting data stored in HDFS can help your organization comply with such regulations.

Introduced in CDH 5.3, transparent encryption for HDFS implements transparent, end-to-end encryption of data read from and written to HDFS blocks across your cluster. *Transparent* means that end-users are unaware of the encryption/decryption processes, and *end-to-end* means that data is encrypted at-rest and in-transit (see the Cloudera Engineering Blog post for complete details).

> **Note:** HDFS Transparent Encryption is not the same as TLS encryption. Clusters configured TLS/SSL encrypt network communications throughout the cluster. Depending on the type of services your cluster supports, you may want to configure both HDFS Transparent Encryption and TLS/SSL for the cluster. See Configuring Encryption for more information.

HDFS encryption has these capabilities:

- Only HDFS clients can encrypt or decrypt data.
- Key management is external to HDFS. HDFS cannot access unencrypted data or encryption keys. Administration of HDFS and administration of keys are separate duties encompassed by distinct user roles (HDFS administrator, Key Administrator), thus ensuring that no single user has unrestricted access to both data and keys.
- The operating system and HDFS interact using encrypted HDFS data only, mitigating threats at the OS- and file-system-level.
- HDFS uses the Advanced Encryption Standard-Counter mode (AES-CTR) encryption algorithm. AES-CTR supports a 128-bit encryption key (default), or can support a 256-bit encryption key when Java Cryptography Extension (JCE) unlimited strength JCE is installed.
- HDFS encryption has been designed to take advantage of the AES-NI instruction set, a hardware-based encryption acceleration technique, so your cluster performance should not adversely affected by configuring encryption. (The AES-NI instruction set can be an order of magnitude faster than software implementations of AES.) However, you may need to update cryptography libraries on your HDFS and MapReduce client hosts to use the acceleration mechanism. See Optimizing Performance for HDFS Transparent Encryption on page 245 for details.

Key Concepts and Architecture

Keystores and the Hadoop Key Management Server

Integrating HDFS with an *external*, enterprise-level keystore is the first step to deploying transparent encryption. This is because separation of duties between a key administrator and an HDFS administrator is a very important aspect of this feature. However, most keystores are not designed for the encrypt/decrypt request rates seen by Hadoop workloads. This led to the development of a new service, called the **Hadoop Key Management Server (KMS)**, which serves as a proxy between HDFS clients and the backing keystore. Both the keystore and Hadoop KMS must use Hadoop's KeyProvider API to interact with each other and with HDFS clients.

While HDFS encryption can be used with a local Java KeyStore for key management, Cloudera does not recommend this for production environments where a more robust and secure key management solution should be used. Cloudera offers the following two options for enterprise-grade key management:

- Cloudera Navigator Key Trustee Server is a key store for managing encryption keys. To integrate with the Navigator Key Trustee Server, Cloudera provides a custom KMS service, Key Trustee KMS.

- Hardware security modules (HSM) are third-party appliances that provide the highest level of security for keys. To integrate with a list of supported HSMs, Cloudera provides a custom KMS service, Navigator HSM KMS (see Installing Navigator HSM KMS Backed by Thales HSM and Installing Navigator HSM KMS Backed by Luna HSM).

The diagram below illustrates how HDFS clients and the NameNode interact with an enterprise keystore through the Hadoop Key Management Server. The keystore can be either the Cloudera Navigator Key Trustee Server or a support HSM.

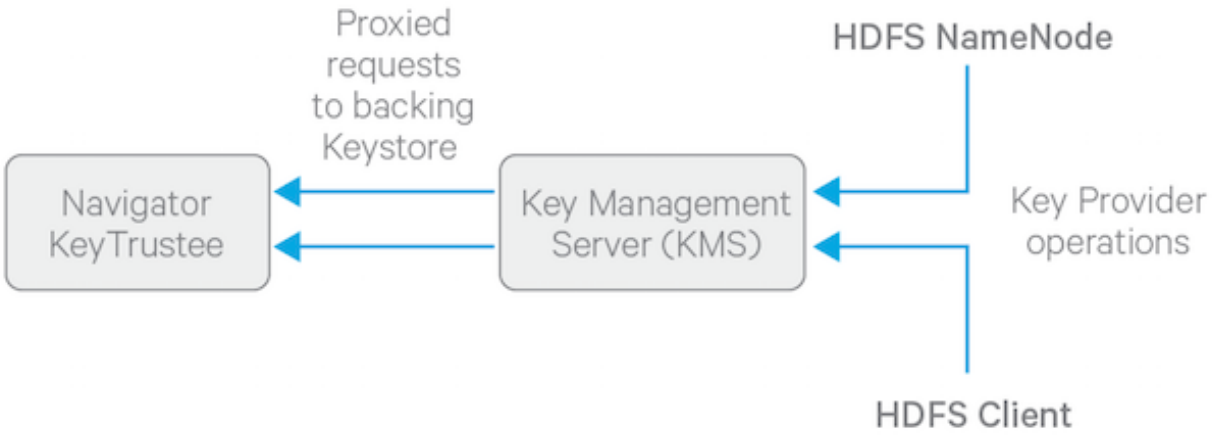

To get started with deploying the KMS and a keystore, see Enabling HDFS Encryption Using the Wizard on page 246.

For information on configuring and securing the KMS, see Configuring the Key Management Server (KMS) on page 265 and Securing the Key Management Server (KMS) on page 270.

KMS Solutions

This section describes the various KMS solutions available, and identifies criteria often used for selecting each.

Key Trustee KMS with Key Trustee Server

Choose KT KMS with Key Trustee Server if:

- Enterprise-grade key management is required
- Encryption zone key protection by an HSM is not required

Key Trustee KMS with Key Trustee Server and Key HSM

Choose KT KMS with Key Trustee Server and Key HSM if:

- Enterprise-grade key management is required
- Encryption zone key protection by an HSM (as root of trust) is required
- Performance for encryption zone key operations is critical

HSM KMS

Choose HSM KMS if:

- Enterprise-grade key management is required
- Encryption zone keys must be stored only on the HSM
- Performance for encryption zone key operations is not critical

Encryption Zones and Keys

HDFS transparent encryption introduces the concept of an ***encryption zone*** (EZ), which is a directory in HDFS whose contents will be automatically encrypted on write and decrypted on read. Encryption zones always start off as empty directories, and tools such as `distcp` with the `-skipcrccheck -update` flags can be used to add data to a zone.

(These flags are required because encryption zones are being used.) Every file and subdirectory copied to an encryption zone will be encrypted.

> **Note:** An encryption zone cannot be created on top of an existing directory.

Each encryption zone is associated with a key (EZ Key) specified by the key administrator when the zone is created. EZ keys are stored on a backing keystore external to HDFS. Each file within an encryption zone has its own encryption key, called the Data Encryption Key (DEK). These DEKs are encrypted with their respective encryption zone's EZ key, to form an Encrypted Data Encryption Key (EDEK).

The following diagram illustrates how encryption zone keys (EZ keys), data encryption keys (DEKs), and encrypted data encryption keys (EDEKs) are used to encrypt and decrypt files.

EDEKs are stored persistently on the NameNode as part of each file's metadata, using HDFS extended attributes (implemented as of Cloudera CDH 5.2). EDEKs can be safely stored and handled by the NameNode because the `hdfs` user does not have access to the EDEK's encryption keys (EZ keys). Even if HDFS is compromised (for example, by gaining unauthorized access to a superuser account), a malicious user only gains access to the encrypted text and EDEKs. EZ keys are controlled by a separate set of permissions on the KMS and the keystore.

An EZ key can have multiple key versions, where each key version has its own distinct key material (that is, the portion of the key used during encryption and decryption). Key rotation is achieved by bumping up the version for an EZ key. Per-file key rotation is then achieved by re-encrypting the file's DEK with the new version of the EZ key to create new EDEKs. HDFS clients can identify an encryption key either by its key name, which returns the latest version of the key, or by a specific key version.

For more information on creating and managing encryption zones, see Managing Encryption Keys and Zones on page 257.

Accessing Files Within an Encryption Zone

To encrypt a new file, the HDFS client requests a new EDEK from the NameNode. The NameNode then asks the KMS to decrypt it with the encryption zone's EZ key. This decryption results in a DEK, which is used to encrypt the file.

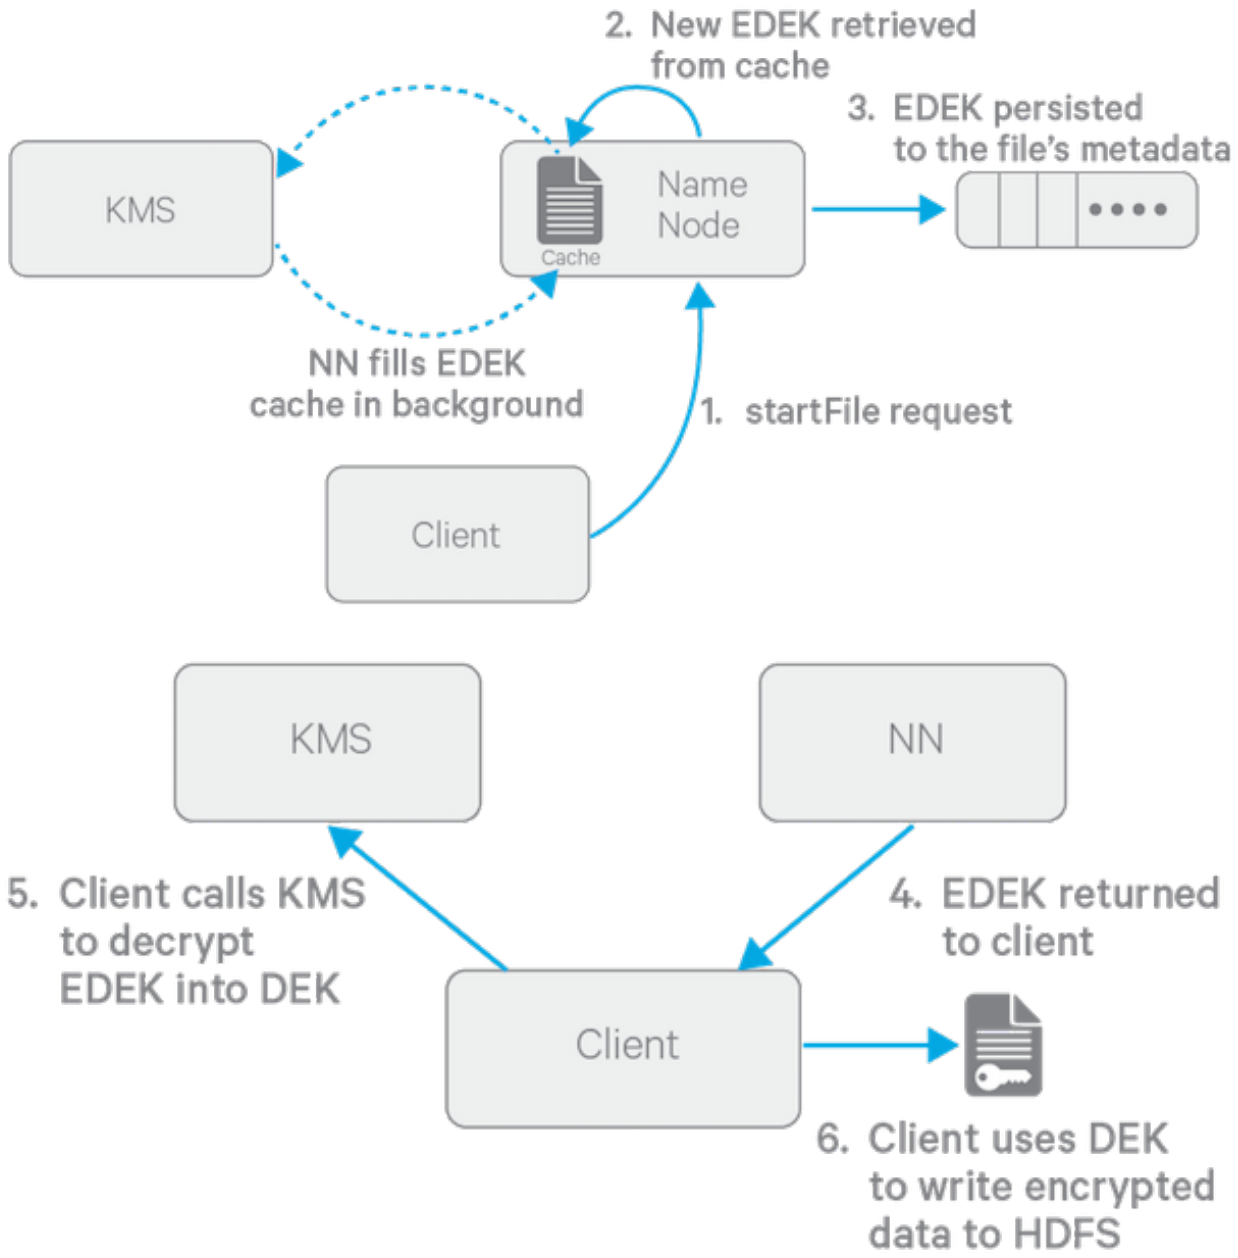

The diagram above depicts the process of writing a new encrypted file. Note that the EDEK cache on the NameNode is populated in the background. Since it is the responsibility of KMS to create EDEKs, using a cache avoids having to call the KMS for each create request. The client can request new EDEKs directly from the NameNode.

To decrypt a file, the HDFS client provides the NameNode with the file's EDEK and the version number of the EZ key that was used to generate the EDEK. The NameNode requests the KMS to decrypt the file's EDEK with the encryption zone's EZ key, which involves checking that the requesting client has permission to access that particular version of the EZ key. Assuming decryption of the EDEK is successful, the client then uses this DEK to decrypt the file.

Encryption and decryption of EDEKs takes place entirely on the KMS. More importantly, the client requesting creation or decryption of an EDEK never handles the EZ key. Only the KMS can use EZ keys to create and decrypt EDEKs as requested. It is important to note that the KMS does not store any keys, other than temporarily in its cache. It is up to the enterprise keystore to be the authoritative storage for keys, and to ensure that keys are never lost, as a lost key is equivalent to introducing a security hole. For production use, Cloudera recommends you deploy two or more redundant enterprise key stores.

Optimizing Performance for HDFS Transparent Encryption

> **Warning:** To ensure that HDFS encryption functions as expected, the steps described in this section are *mandatory for production use*.

CDH implements the ***Advanced Encryption Standard New Instructions*** (AES-NI), which provide substantial performance improvements. To get these improvements, you need a recent version of `libcrypto.so` on HDFS and MapReduce client hosts -- that is, any host from which you originate HDFS or MapReduce requests. Many OS versions have an older version of the library that does not support AES-NI. The instructions that follow tell you what you need to do for each OS version that CDH supports.

RHEL/CentOS 6.5 or later

The installed version of `libcrypto.so` supports AES-NI, but you need to install the `openssl-devel` package on all clients:

```
$ sudo yum install openssl-devel
```

RHEL/CentOS 6.4 or earlier 6.x versions, or SLES 11

Download and extract a newer version of `libcrypto.so` from a CentOS 6.5 repository and install it on all clients in `/var/lib/hadoop/extra/native/`:

1. Download the latest version of the `openssl` package. For example:

```
$ wget
http://mirror.centos.org/centos/6/os/x86_64/Packages/openssl-1.0.1e-30.el6.x86_64.rpm
```

   The `libcrypto.so` file in this package can be used on SLES 11 as well as RHEL/CentOS.

2. Decompress the files in the package, but **do not** install it:

```
$ rpm2cpio openssl-1.0.1e-30.el6.x86_64.rpm | cpio -idmv
```

3. If you are using parcels, create the `/var/lib/hadoop/extra/native/` directory:

```
$ sudo mkdir -p /var/lib/hadoop/extra/native
```

4. Copy the shared library into `/var/lib/hadoop/extra/native/`. Name the target file `libcrypto.so`, with no suffix at the end, exactly as in the command that follows.

```
$ sudo cp ./usr/lib64/libcrypto.so.1.0.1e /var/lib/hadoop/extra/native/libcrypto.so
```

RHEL/CentOS 5

In this case, you need to build `libcrypto.so` and copy it to all clients:

1. On one client, compile and install `openssl` from source:

```
$ wget http://www.openssl.org/source/openssl-1.0.1j.tar.gz
$ cd openssl-1.0.1j
$ ./config --shared --prefix=/opt/openssl-1.0.1j
$ sudo make install
```

2. If you are using parcels, create the `/var/lib/hadoop/extra/native/` directory:

```
$ sudo mkdir -p /var/lib/hadoop/extra/native
```

3. Copy the files into `/var/lib/hadoop/extra/native/`:

```
$ sudo cp /opt/openssl-1.0.1j/lib/libcrypto.so /var/lib/hadoop/extra/native
```

4. Copy the files to the remaining clients using a utility such as `rsync`

Debian Wheezy

The installed version of `libcrypto.so` supports AES-NI, but you need to install the `libssl-devel` package on all clients:

```
$ sudo apt-get install libssl-dev
```

Ubuntu Precise and Ubuntu Trusty

Install the `libssl-devel` package on all clients:

```
$ sudo apt-get install libssl-dev
```

Testing if encryption optimization works

To verify that a client host is ready to use the AES-NI instruction set optimization for HDFS encryption at rest, use the following command:

```
hadoop checknative
```

You should see a response such as the following:

```
14/12/12 13:48:39 INFO bzip2.Bzip2Factory: Successfully loaded & initialized native-bzip2
library system-native14/12/12 13:48:39 INFO zlib.ZlibFactory: Successfully loaded &
initialized native-zlib library
Native library checking:
hadoop:  true /usr/lib/hadoop/lib/native/libhadoop.so.1.0.0
zlib:    true /lib64/libz.so.1
snappy:  true /usr/lib64/libsnappy.so.1
lz4:     true revision:99
bzip2:   true /lib64/libbz2.so.1
openssl: true /usr/lib64/libcrypto.so
```

If you see `true` in the `openssl` row, Hadoop has detected the right version of `libcrypto.so` and optimization will work. If you see `false` in this row, you do not have the right version.

Enabling HDFS Encryption Using the Wizard

To accommodate the security best practice of separation of duties, enabling HDFS encryption using the wizard requires different Cloudera Manager user roles for different steps.

Launch the **Set up HDFS Data At Rest Encryption** wizard in one of the following ways:

- **Cluster** > **Set up HDFS Data At Rest Encryption**

  Minimum Required Role: Key Administrator *or* Cluster Administrator (also provided by Full Administrator)

- **Administration** > **Security** > **Set up HDFS Data At Rest Encryption**

  Minimum Required Role: Key Administrator *or* Cluster Administrator (also provided by Full Administrator)

- **HDFS service** > **Actions** > **Set up HDFS Data At Rest Encryption**

  Minimum Required Role: Cluster Administrator (also provided by **Full Administrator**)

On the first page of the wizard, select the root of trust for encryption keys:

- **Cloudera Navigator Key Trustee Server**

- **Navigator HSM KMS backed by Thales HSM**
- **Navigator HSM KMS backed by Luna HSM**
- **A file-based password-protected Java KeyStore**

Cloudera strongly recommends using Cloudera Navigator Key Trustee Server as the root of trust for production environments. The file-based Java KeyStore root of trust is insufficient to provide the security, scalability, and manageability required by most production systems. More specifically, the Java KeyStore KMS does not provide:

- Scalability, so you are limited to only one KMS, which can result in bottlenecks
- High Availability (HA)
- Recoverability, so if you lose the node where the Java KeyStore is stored, then you can lose access to all the encrypted data

Ultimately, the Java KeyStore does not satisfy the stringent security requirements of most organizations for handling master encryption keys.

Choosing a root of trust displays a list of steps required to enable HDFS encryption using that root of trust. Each step can be completed independently. The **Status** column indicates whether the step has been completed, and the **Notes** column provides additional context for the step. If your Cloudera Manager user account does not have sufficient privileges to complete a step, the **Notes** column indicates the required privileges.

Available steps contain links to wizards or documentation required to complete the step. If a step is unavailable due to insufficient privileges or a prerequisite step being incomplete, no links are present and the **Notes** column indicates the reason the step is unavailable.

Continue to the section for your selected root of trust for further instructions:

Enabling HDFS Encryption Using Navigator Key Trustee Server

Enabling HDFS encryption using Key Trustee Server as the key store involves multiple components. For an overview of the components involved in encrypting data at rest, see Cloudera Navigator Data Encryption Overview. For guidelines on deploying the Navigator Key Trustee Server in production environments, Resource Planning for Data at Rest Encryption on page 239.

Before continuing, make sure the Cloudera Manager server host has access to the internal repository hosting the Key Trustee Server software. See Setting Up an Internal Repository for more information.

After selecting **Cloudera Navigator Key Trustee Server** as the root of trust, the following steps are displayed:

1. Enable Kerberos

Minimum Required Role: Cluster Administrator (also provided by **Full Administrator**)

For more information about enabling Kerberos, see Enabling Kerberos Authentication Using the Wizard on page 51.

2. Enable TLS/SSL

Minimum Required Role: Cluster Administrator (also provided by **Full Administrator**)

For more information about enabling TLS, see Configuring TLS/SSL Encryption for CDH Services on page 204.

3. Add a dedicated cluster for the Key Trustee Server

Minimum Required Role: Cluster Administrator (also provided by **Full Administrator**)

If you haven't already done so, you *must* create an internal repository to install Cloudera Navigator before you can set up and use Navigator Key Trustee Server. For instructions on creating internal repositories (including Cloudera Manager, CDH, and Cloudera Navigator encryption components), see Creating and Using a Parcel Repository for Cloudera Manager and Creating and Using a Package Repository for Cloudera Manager.

This step creates a new cluster in Cloudera Manager for the Key Trustee Server hosts to isolate them from other enterprise data hub (EDH) services for increased security and durability. For more information, see Data at Rest Encryption Reference Architecture on page 238.

To complete this step:

1. Click **Add a dedicated cluster for the Key Trustee Server**.
2. Leave **Enable High Availability** checked to add two hosts to the cluster. For production environments, you must enable high availability for Key Trustee Server. Failure to enable high availability can result in complete data loss in the case of catastrophic failure of a standalone Key Trustee Server. Click **Continue**.
3. Search for new hosts to add to the cluster, or select the **Currently Managed Hosts** tab to add existing hosts to the cluster. After selecting the hosts, click **Continue**.
4. Select the KEYTRUSTEE_SERVER parcel to install Key Trustee Server using parcels, or select **None** if you want to use packages. If you do not see a parcel available, click **More Options** and add the repository URL to the **Remote Parcel Repository URLs** list. After selecting a parcel or **None**, click **Continue**.

   If you selected **None**, click **Continue** again, and skip to 4. Install Key Trustee Server binary using packages or parcels on page 248.

5. After the KEYTRUSTEE_SERVER parcel is successfully downloaded, distributed, unpacked, and activated, click **Continue**.
6. Click **Continue** to complete this step and return to the main page of the wizard.

4. Install Key Trustee Server binary using packages or parcels

Minimum Required Role: Cluster Administrator (also provided by **Full Administrator**)

> **Note:** If you selected **None** on the parcel selection page in step 3. Add a dedicated cluster for the Key Trustee Server on page 247, the step title is changed to **Install Parcel for Key Trustee Server**. If you are using packages, skip this step and see Installing Key Trustee Server Using the Command Line for package-based installation instructions. After installing Key Trustee Server using packages, continue to 5. Install Parcel for Key Trustee KMS on page 248.

If you haven't already done so, you *must* create an internal repository to install Cloudera Navigator before you can set up and use Navigator Key Trustee Server. For instructions on creating internal repositories (including Cloudera Manager, CDH, and Cloudera Navigator encryption components), see Creating and Using a Parcel Repository for Cloudera Manager and Creating and Using a Package Repository for Cloudera Manager.

This step is completed automatically during 3. Add a dedicated cluster for the Key Trustee Server on page 247 if you are using parcels. If the step is incomplete for any reason (such as the wizard being interrupted or a failure installing the parcel), complete it manually:

1. Click **Install Key Trustee Server binary using packages or parcels**.
2. Select the KEYTRUSTEE_SERVER parcel to install Key Trustee Server, or select **None** if you need to install Key Trustee Server manually using packages. If you do not see a parcel available, click **More Options** and add the repository URL to the **Remote Parcel Repository URLs** list. After selecting a parcel, click **Continue**.
3. After the KEYTRUSTEE_SERVER parcel is successfully downloaded, distributed, unpacked, and activated, click **Finish** to complete this step and return to the main page of the wizard.

5. Install Parcel for Key Trustee KMS

Minimum Required Role: Cluster Administrator (also provided by **Full Administrator**)

If you haven't already done so, you *must* create an internal repository to install Cloudera Navigator before you can set up and use Navigator Key Trustee Server. For instructions on creating internal repositories (including Cloudera Manager, CDH, and Cloudera Navigator encryption components), see Creating and Using a Parcel Repository for Cloudera Manager and Creating and Using a Package Repository for Cloudera Manager.

This step installs the Key Trustee KMS parcel. If you are using packages, skip this step and see Installing Key Trustee KMS Using Packages for instructions. After installing Key Trustee KMS using packages, continue to 6. Add a Key Trustee Server Service on page 249.

To complete this step for parcel-based installations:

1. Click **Install Parcel for Key Trustee KMS**.
2. Select the KEYTRUSTEE parcel to install Key Trustee KMS. If you do not see a parcel available, click **More Options** and add the repository URL to the **Remote Parcel Repository URLs** list. After selecting a parcel, click **Continue**.
3. After the KEYTRUSTEE parcel is successfully downloaded, distributed, unpacked, and activated, click **Finish** to complete this step and return to the main page of the wizard.

6. Add a Key Trustee Server Service

Minimum Required Role: Key Administrator (also provided by **Full Administrator**)

This step adds the **Key Trustee Server** service to Cloudera Manager. To complete this step:

1. Click **Add a Key Trustee Server Service**.
2. Click **Continue**.
3. On the **Customize Role Assignments for Key Trustee Server** page, select the hosts for the **Active Key Trustee Server** and **Passive Key Trustee Server** roles. Make sure that the selected hosts are not used for other services (see Resource Planning for Data at Rest Encryption on page 239 for more information), and click **Continue**.
4. The **Entropy Considerations** page provides commands to install the `rng-tools` package to increase available entropy for cryptographic operations. For more information, see Entropy requirements in Data at Rest Encryption Requirements. After completing these commands, click **Continue**.
5. The **Synchronize Active and Passive Key Trustee Server Private Keys** page provides instructions for generating and copying the Active Key Trustee Server private key to the Passive Key Trustee Server. Cloudera recommends following security best practices and transferring the private key using offline media, such as a removable USB drive. For convenience (for example, in a development or testing environment where maximum security is not required), you can copy the private key over the network using the provided `rsync` command.

   After you have synchronized the private keys, run the `ktadmin init` command on the Passive Key Trustee Server as described in the wizard. After the initialization is complete, check the box to indicate you have synchronized the keys and click **Continue** in the wizard.

6. The **Setup TLS for Key Trustee Server** page provides instructions on replacing the auto-generated self-signed certificate with a production certificate from a trusted Certificate Authority (CA). For more information, see Managing Key Trustee Server Certificates on page 315. Click **Continue** to view and modify the default certificate settings.
7. On the **Review Changes** page, you can view and modify the following settings:

   - **Database Storage Directory** (`db_root`)

     Default value: `/var/lib/keytrustee/db`

     The directory on the local filesystem where the Key Trustee Server database is stored. Modify this value to store the database in a different directory.

   - **Active Key Trustee Server TLS/SSL Server Private Key File (PEM Format)** (`ssl.privatekey.location`)

     Default value: `/var/lib/keytrustee/.keytrustee/.ssl/ssl-cert-keytrustee-pk.pem`

     The path to the Active Key Trustee Server TLS certificate private key. Accept the default setting to use the auto-generated private key. If you have a CA-signed certificate, change this path to the CA-signed certificate private key file. This file must be in PEM format.

   - **Active Key Trustee Server TLS/SSL Server Certificate File (PEM Format)** (`ssl.cert.location`)

     Default value: `/var/lib/keytrustee/.keytrustee/.ssl/ssl-cert-keytrustee.pem`

     The path to the Active Key Trustee Server TLS certificate. Accept the default setting to use the auto-generated self-signed certificate. If you have a CA-signed certificate, change this to the path to the CA-signed certificate. This file must be in PEM format.

   - **Active Key Trustee Server TLS/SSL Server CA Certificate (PEM Format)** (`ssl.cacert.location`)

     Default value: (none)

The path to the file containing the CA certificate and any intermediate certificates (if any intermediate certificates exist, then they are *required* here) used to sign the Active Key Trustee Server certificate. If you have a CA-signed certificate, set this value to the path to the CA certificate or certificate chain file. This file must be in PEM format.

- **Active Key Trustee Server TLS/SSL Private Key Password** (`ssl.privatekey.password`)

  Default value: (none)

  The password for the Active Key Trustee Server private key file. Leave this blank if the file is not password-protected.

- **Passive Key Trustee Server TLS/SSL Server Private Key File (PEM Format)** (`ssl.privatekey.location`)

  Default value: `/var/lib/keytrustee/.keytrustee/.ssl/ssl-cert-keytrustee-pk.pem`

  The path to the Passive Key Trustee Server TLS certificate private key. Accept the default setting to use the auto-generated private key. If you have a CA-signed certificate, change this path to the CA-signed certificate private key file. This file must be in PEM format.

- **Passive Key Trustee Server TLS/SSL Server Certificate File (PEM Format)** (`ssl.cert.location`)

  Default value: `/var/lib/keytrustee/.keytrustee/.ssl/ssl-cert-keytrustee.pem`

  The path to the Passive Key Trustee Server TLS certificate. Accept the default setting to use the auto-generated self-signed certificate. If you have a CA-signed certificate, change this to the path to the CA-signed certificate. This file must be in PEM format.

- **Passive Key Trustee Server TLS/SSL Server CA Certificate (PEM Format)** (`ssl.cacert.location`)

  Default value: (none)

  The path to the file containing the CA certificate and any intermediate certificates (if any intermediate certificates exist, then they are *required* here) used to sign the Passive Key Trustee Server certificate. If you have a CA-signed certificate, set this value to the path to the CA certificate or certificate chain file. This file must be in PEM format.

- **Passive Key Trustee Server TLS/SSL Private Key Password** (`ssl.privatekey.password`)

  Default value: (none)

  The password for the Passive Key Trustee Server private key file. Leave this blank if the file is not password-protected.

After reviewing the settings and making any changes, click **Continue**.

8. After all commands complete successfully, click **Continue**. If the **Generate Key Trustee Server Keyring** appears stuck, make sure that the Key Trustee Server host has enough entropy. See entropy requirements in Data at Rest Encryption Requirements for more information.
9. Click **Finish** to complete this step and return to the main page of the wizard.

For parcel-based Key Trustee Server releases 5.8 and higher, Cloudera Manager automatically backs up Key Trustee Server (using the `ktbackup.sh` script) after adding the Key Trustee Server service. It also schedules automatic backups using `cron`. For package-based installations, you must manually back up Key Trustee Server and configure a `cron` job.

Cloudera Manager configures `cron` to run the backup script hourly. The latest 10 backups are retained in `/var/lib/keytrustee` in cleartext. For information about using the backup script and configuring the `cron` job (including how to encrypt backups), see Backing Up Key Trustee Server and Key Trustee KMS Using the ktbackup.sh Script on page 300.

7. Add a Key Trustee KMS Service

Minimum Required Role: Key Administrator (also provided by **Full Administrator**)

This step adds a Key Trustee KMS service to the cluster. The Key Trustee KMS service is required to enable HDFS encryption to use Key Trustee Server for cryptographic key management. Key Trustee KMS high availability uses ZooKeeper to automatically configure load balancing. If you do not have a ZooKeeper service in your cluster, add one using the instructions in Adding a Service.

> **Note:** Before performing this step, you must have already added hosts to the environment for KMS. For details, see Adding a Host to the Cluster.

To complete this step:

1. Click **Add a Key Trustee KMS Service**.
2. Select an existing Key Trustee Server pair or specify an external Key Trustee Server pair. If you have an existing Key Trustee Server pair outside of Cloudera Manager control, select the **External Key Trustee Server** option and specify the fully-qualified domain names (FQDNs) of the Key Trustee Server pair. Click **Continue**.
3. Select cluster hosts for the Key Trustee KMS service. For production environments, select at least two hosts for high availability. If you proceed with only one host, you can enable high availability later. See Enabling Key Trustee KMS High Availability for more information.

   Make sure that the selected hosts are not used for other services (see Resource Planning for Data at Rest Encryption on page 239 for more information), and click **Continue**.

4. The **Entropy Considerations** page provides commands to install the `rng-tools` package to increase available entropy for cryptographic operations. For more information, see entropy requirements in Data at Rest Encryption Requirements. After completing these commands, click **Continue**.
5. The **Setup Organization and Auth Secret** page generates the necessary commands to create an organization in Key Trustee Server. An organization is required to be able to register the Key Trustee KMS with Key Trustee Server. See Managing Key Trustee Server Organizations on page 312 for more information.

   Enter an organization name and click **Generate Instruction**. Run the displayed commands to generate an organization and obtain the `auth_secret` value for the organization. Enter the secret in the **auth_secret** field and click **Continue**.

6. The **Setup Access Control List (ACL)** page allows you to generate ACLs for the Key Trustee KMS or to provide your own ACLs. To generate the recommended ACLs, enter the username and group responsible for managing cryptographic keys and click **Generate ACLs**. To specify your own ACLs, select the **Use Your Own kms-acls.xml File** option and enter the ACLs. For more information on the KMS Access Control List, see Configuring KMS Access Control Lists (ACLs) on page 274.

   After generating or specifying the ACL, click **Continue**.

7. The **Setup TLS for Key Trustee KMS** page provides high-level instructions for configuring TLS communication between the Key Trustee KMS and the Key Trustee Server, as well as between the EDH cluster and the Key Trustee KMS. See Configuring TLS/SSL for the KMS on page 272 for more information.

   Click **Continue**.

8. The **Review Changes** page lists all of the settings configured in this step. Click the question mark icon next to any setting for information about that setting. Review the settings and click **Continue**.
9. After the **First Run** commands have successfully completed, click **Continue**.
10. The **Synchronize Private Keys and HDFS Dependency** page provides instructions for copying the private key from one Key Management Server Proxy role to all other roles.

> **Warning:** It is *very important* that you perform this step. Failure to do so leaves Key Trustee KMS in a state where keys are intermittently inaccessible, depending on which Key Trustee KMS host a client interacts with, because cryptographic key material encrypted by one Key Trustee KMS host cannot be decrypted by another. If you are already running multiple Key Trustee KMS hosts with different private keys, immediately back up all Key Trustee KMS hosts, and contact Cloudera Support for assistance correcting the issue.
>
> To determine whether the Key Trustee KMS private keys are different, compare the MD5 hash of the private keys. On each Key Trustee KMS host, run the following command:
>
> ```
> $ md5sum /var/lib/kms-keytrustee/keytrustee/.keytrustee/secring.gpg
> ```
>
> If the outputs are different, contact Cloudera Support for assistance. Do not attempt to synchronize existing keys. If you overwrite the private key and do not have a backup, any keys encrypted by that private key are permanently inaccessible, and any data encrypted by those keys is permanently irretrievable. If you are configuring Key Trustee KMS high availability for the first time, continue synchronizing the private keys.

Cloudera recommends following security best practices and transferring the private key using offline media, such as a removable USB drive. For convenience (for example, in a development or testing environment where maximum security is not required), you can copy the private key over the network using the provided `rsync` command.

After you have synchronized the private keys, check the box to indicate you have done so and click **Continue**.

11. After the Key Trustee KMS service starts, click **Finish** to complete this step and return to the main page of the wizard.

For parcel-based Key Trustee KMS releases 5.8 and higher, Cloudera Manager automatically backs up Key Trustee KMS (using the `ktbackup.sh` script) after adding the Key Trustee KMS service. It does not schedule automatic backups using `cron`. For package-based installations, you must manually back up Key Trustee Server and configure a `cron` job.

The backup is stored in `/var/lib/kms-keytrustee` in cleartext. For more information about using the backup script and configuring the `cron` job (including how to encrypt backups), see Backing Up Key Trustee Server and Key Trustee KMS Using the ktbackup.sh Script on page 300.

8. Restart stale services and redeploy client configuration

Minimum Required Role: Cluster Administrator (also provided by **Full Administrator**)

This step restarts all services which were modified while enabling HDFS encryption. To complete this step:

1. Click **Restart stale services and redeploy client configuration**.
2. Click **Restart Stale Services**.
3. Ensure that **Re-deploy client configuration** is checked, and click **Restart Now**.
4. After all commands have completed, click **Finish**.

9. Validate Data Encryption

Minimum Required Role: Key Administrator *or* **Cluster Administrator** (also provided by Full Administrator)

This step launches a tutorial with instructions on creating an encryption zone and putting data into it to verify that HDFS encryption is enabled and working.

Enabling HDFS Encryption Using Navigator HSM KMS Backed by Thales HSM

1. Enable Kerberos

Minimum Required Role: Cluster Administrator (also provided by **Full Administrator**)

For more information about enabling Kerberos, see Enabling Kerberos Authentication Using the Wizard on page 51.

2. Enable TLS/SSL

Minimum Required Role: Cluster Administrator (also provided by **Full Administrator**)

For more information about enabling TLS, see Configuring TLS Encryption for Cloudera Manager on page 195.

3. Install the Thales HSM Client

Before installing the Navigator HSM KMS backed by Thales HSM, you *must* install the Thales HSM client on the host. Attempts to install the HSM KMS service before installing the Thales HSM client will fail.

For details about how to install the Thales HSM client, refer to the Thales HSM product documentation.

4. Install Key Trustee KMS binary using parcels

Minimum Required Role: Cluster Administrator (also provided by **Full Administrator**)

This step completes automatically when you download the parcel. If the step is incomplete for any reason (such as the wizard being interrupted or a failure installing the parcel), complete it manually:

1. Click **Install Key Trustee KMS binary using parcels**.
2. Select the KEYTRUSTEE parcel to install Key Trustee KMS, or select **None** if you need to install Key Trustee KMS manually using packages. If you do not see a parcel available, click **More Options** and add the repository URL to the **Remote Parcel Repository URLs** list. After selecting a parcel, click **Continue**.
3. After the KEYTRUSTEE parcel is successfully downloaded, distributed, unpacked, and activated, click **Finish** to complete this step and return to the main page of the wizard.

5. Add the HSM KMS backed by Thales Service

1. Click **Add Navigator HSM KMS Services backed by Thales HSM**.
2. In the **Thales HSM KMS Proxy** field, select the hosts to which you want to assign a new or existing role. Click **OK**, and then click **Continue**.
3. To set up the ACL for the cluster, specify a comma-separated list of users and groups, and then click **Generate ACLs**. Click **Continue**.
4. Click **Continue**.
5. Review your selections and specify the:
   - Thales HSM Password

     Contact your HSM administrator for the Thales HSM password.
   - Keystore Password

     Important: Specify a unique password to protect the KMS metadata. *Safely store this password for future use*. In the event that the service is deleted, you *must* provide this password to re-add the service with the same metastore.

   Note: It is a good practice to review the Thales HSM Server port at this point, because it could have changed during configuration.

   Then click **Continue**.
6. Upon notification that you have successfully added the Thales KMS Service, click **Continue** and **Finish**.

6. Restart stale services and redeploy client configuration

Minimum Required Role: Cluster Administrator (also provided by **Full Administrator**)

This step restarts all services that were modified while enabling HDFS encryption. To complete this step:

1. Click **Restart stale services and redeploy client configuration**.
2. Click **Restart Stale Services**.
3. Ensure that **Re-deploy client configuration** is checked, and click **Restart Now**.
4. After all commands have completed, click **Finish**.

7. Validate Data Encryption

Minimum Required Role: Key Administrator *or* **Cluster Administrator** (also provided by Full Administrator)

This step launches a Validate Data Encryption tutorial with instructions describing how to create an encryption zone and place data into it to verify that HDFS encryption is enabled and working.

Enabling HDFS Encryption Using Navigator HSM KMS Backed by Luna HSM

1. Enable Kerberos

Minimum Required Role: Cluster Administrator (also provided by **Full Administrator**)

For more information about enabling Kerberos, see Enabling Kerberos Authentication Using the Wizard on page 51.

2. Enable TLS/SSL

Minimum Required Role: Cluster Administrator (also provided by **Full Administrator**)

For more information about enabling TLS, see Configuring TLS Encryption for Cloudera Manager on page 195.

3. Install Luna HSM Client

Before installing the Navigator HSM KMS backed by Luna HSM, you must install the Luna HSM client on the host. Attempts to install the Navigator HSM KMS backed by Luna HSM before installing the Luna HSM client will fail.

For details about how to install the Luna HSM client, refer to the Luna HSM product documentation.

4. Install Parcel for Cloudera Key Providers

Minimum Required Role: Cluster Administrator (also provided by **Full Administrator**)

This step completed automatically when you downloaded the parcel. If the step is incomplete for any reason (such as the wizard being interrupted or a failure installing the parcel), complete it manually:

1. Click **Install Key Trustee KMS binary using packages or parcels**.
2. Select the KEYTRUSTEE parcel to install Key Trustee KMS, or select **None** if you need to install Key Trustee KMS manually using packages. If you do not see a parcel available, click **More Options** and add the repository URL to the **Remote Parcel Repository URLs** list. After selecting a parcel, click **Continue**.
3. After the KEYTRUSTEE parcel is successfully downloaded, distributed, unpacked, and activated, click **Finish** to complete this step and return to the main page of the wizard.

5. Add the Navigator HSM KMS backed by SafeNet Luna HSM

1. Click **Add Navigator HSM KMS backed by Safenet Luna HSM**.
2. In the **Luna HSM-backed KMS Proxy** field, select the hosts to which you want to assign a new or existing role. Click **OK**, and then click **Continue**.
3. To set up the ACL for the cluster, specify a comma-separated list of users and groups, and then click **Generate ACLs**. Click **Continue**.
4. Click **Continue**.
5. Review your selections and specify the:
   - Luna HSM Password

     Contact your HSM administrator for the Luna HSM Partition password.

- Keystore Password

  > **Important:** Specify a unique password to protect the KMS metadata. *Safely store this password for future use*. In the event that the service is deleted, you *must* provide this password to re-add the service with the same metastore.

- Luna HSM Server Slot

  Identification number of the Luna HSM Server slot/device to use. If you do not know what value(s) to enter here, see the Luna product documentation for instructions on configuring your Luna HSM. Alternatively, run the `/usr/safenet/lunaclient/bin/vtl verify` command on the Luna HSM client host to view the slot value.

Then click **Continue**.

6. Upon notification that you have successfully added the Navigator Safenet Luna KMS Service, click **Continue** and **Finish**.

6. Restart stale services and redeploy client configuration

Minimum Required Role: **Cluster Administrator** (also provided by **Full Administrator**)

This step restarts all services that were modified while enabling HDFS encryption. To complete this step:

1. Click **Restart stale services and redeploy client configuration**.
2. Click **Restart Stale Services**.
3. Ensure that **Re-deploy client configuration** is checked, and click **Restart Now**.
4. After all commands have completed, click **Finish**.

7. Validate Data Encryption

Minimum Required Role: **Key Administrator** *or* **Cluster Administrator** (also provided by Full Administrator)

This step launches a Validate Data Encryption tutorial with instructions describing how to create an encryption zone and place data into it to verify that HDFS encryption is enabled and working.

Enabling HDFS Encryption Using a Java KeyStore

> **Note:** Cloudera strongly recommends using Cloudera Navigator Key Trustee Server as the root of trust for production environments. The file-based Java KeyStore root of trust is insufficient to provide the security, scalability, and manageability required by most production systems.

After selecting **A file-based password-protected Java KeyStore** as the root of trust, the following steps are displayed:

1. Enable Kerberos

Minimum Required Role: **Cluster Administrator** (also provided by **Full Administrator**)

For more information on enabling Kerberos, see Enabling Kerberos Authentication Using the Wizard on page 51.

2. Enable TLS/SSL

Minimum Required Role: **Cluster Administrator** (also provided by **Full Administrator**)

For more information on enabling TLS, see Configuring TLS Encryption for Cloudera Manager on page 195.

3. Add a Java KeyStore KMS Service

Minimum Required Role: **Key Administrator** (also provided by **Full Administrator**)

This step adds the Java KeyStore KMS service to the cluster. The Java KeyStore KMS service uses a password-protected Java KeyStore for cryptographic key management. To complete this step:

1. Click **Add a Java KeyStore KMS Service**.
2. Select a cluster host for the Java KeyStore KMS service. Click **Continue**.
3. The **Setup TLS for Java KeyStore KMS** page provides high-level instructions for configuring TLS communication between the EDH cluster and the Java KeyStore KMS. See Configuring TLS/SSL for the KMS on page 272 for more information.

   Click **Continue**.
4. The **Review Changes** page lists the Java KeyStore settings. Click the question mark icon next to any setting for information about that setting. Enter the location and password for the Java KeyStore and click **Continue**.
5. Click **Continue** to automatically configure the HDFS service to depend on the Java KeyStore KMS service.
6. Click **Finish** to complete this step and return to the main page of the wizard.

4. Restart stale services and redeploy client configuration

Minimum Required Role: Cluster Administrator (also provided by **Full Administrator**)

This step restarts all services which were modified while enabling HDFS encryption. To complete this step:

1. Click **Restart stale services and redeploy client configuration**.
2. Click **Restart Stale Services**.
3. Ensure that **Re-deploy client configuration** is checked, and click **Restart Now**.
4. After all commands have completed, click **Finish**.

5. Validate Data Encryption

Minimum Required Role: Key Administrator *or* **Cluster Administrator** (also provided by Full Administrator)

This step launches a tutorial with instructions on creating an encryption zone and putting data into it to verify that HDFS encryption is enabled and working.

Hints and Tips

This section includes hints and tips that can help simplify the HSM KMS installation when using the HDFS Encryption Wizard.

Limit the Number of ZooKeeper DEBUG Messages

When setting the KMS log level to DEBUG, there can be a lot of ZooKeeper DEBUG messages that clutter the log. To prevent this, in the **Logging Advanced Configuration Snippet (Safety Valve)** field, enter:

```
log4j.category.org.apache.zookeeper=INFO
```

Limit Encryption Zone Timeouts

When creating encryption zones, there can be client timeouts due to the time it takes to fill the encrypted data encryption key (EDEK) cache. To avoid this, adjust the low watermark threshold settings as follows.

On the server side, in the field **HSM KMS Proxy Advanced Configuration Snippet (Safety Valve) for kms-site.xml**:

```
<property>
   <name>hadoop.security.kms.encrypted.key.cache.low.watermark</name>
   <value>.03</value>
</property>
```

On the client side, in the field **HDFS Cluster-wide Advanced Configuration Snippet (Safety Valve) for core-site.xml**:

```
<property>
   <name>hadoop.security.kms.client.encrypted.key.cache.low-watermark</name>
   <value>.02</value>
</property>
```

Increase KMS Client Timeout Value

Due to potential latency during installation, it is recommended that you increase the KMS client timeout value.

Change from the default of 60 seconds to a value between 100 and 120 seconds in the field **HDFS Cluster-wide Advanced Configuration Snippet (Safety Valve) for core-site.xml**:

```
<property>
   <name>hadoop.security.kms.client.timeout</name>
   <value>110</value>
</property>
```

Managing Encryption Keys and Zones

Interacting with the KMS and creating encryption zones requires the use of two new CLI commands: `hadoop key` and `hdfs crypto`. The following sections will help you get started with creating encryption keys and setting up encryption zones.

Before continuing, make sure that your KMS ACLs have been set up according to best practices. For more information, see Configuring KMS Access Control Lists (ACLs) on page 274.

Validating Hadoop Key Operations

> **Warning:** If you are using or plan to use Cloudera Navigator Key HSM in conjunction with Cloudera Navigator Key Trustee Server, ensure that:
>
> Key names begin with alphanumeric characters and do not use special characters other than hyphen (`-`), period (`.`), or underscore (_). Using other special characters can prevent you from migrating your keys to an HSM.
>
> See Integrating Key HSM with Key Trustee Server on page 323 for more information.

> **Warning:**
>
> HSM KMS key names should not exceed 28 characters. If the key name exceeds 28 characters, then it is silently truncated when passed on to the HSM; however, it is used at full length in the HSM KMS local data store. Truncation only impacts key names created on the HSM, not in those in the HSM KMS metastore.

Use `hadoop key create` to create a test key, and then use `hadoop key list` to retrieve the key list:

```
$ sudo -u <key_admin> hadoop key create keytrustee_test
$ hadoop key list
```

Creating Encryption Zones

> **Important:** Cloudera does not currently support configuring the root directory as an encryption zone. Nested encryption zones are also not supported.

> **Important:** The Java Keystore KMS default Truststore (for example, `org.apache.hadoop.crypto.key.JavaKeyStoreProvider`) does not support uppercase key names.

Once a KMS has been set up and the NameNode and HDFS clients have been correctly configured, use the `hadoop key` and `hdfs crypto` command-line tools to create encryption keys and set up new encryption zones.

- Create an encryption key for your zone as `keyadmin` for the user/group (regardless of the application that will be using the encryption zone):

```
$ sudo -u hdfs hadoop key create <key_name>
```

- Create a new empty directory and make it an encryption zone using the key created above.

```
$ sudo -u hdfs hadoop fs -mkdir /encryption_zone
$ sudo -u hdfs hdfs crypto -createZone -keyName <key_name> -path /encryption_zone
```

  You can verify creation of the new encryption zone by running the `-listZones` command. You should see the encryption zone along with its key listed as follows:

```
$ sudo -u hdfs hdfs crypto -listZones
/encryption_zone    <key_name>
```

> **Warning:** Do not delete an encryption key as long as it is still in use for an encryption zone. This results in loss of access to data in that zone. Also, do not delete the KMS service, as your encrypted HDFS data will be inaccessible. To prevent data loss, first copy the encrypted HDFS data to a non-encrypted zone using the `distcp` command.

For more information and recommendations on creating encryption zones for each CDH component, see Configuring CDH Services for HDFS Encryption on page 293.

Adding Files to an Encryption Zone

You can add files to an encryption zone by copying them to the encryption zone using `distcp`. For example:

```
sudo -u hdfs hadoop distcp /user/dir /encryption_zone
```

> **Important:** Starting with CDH 5.7.1, you can delete files or directories that are part of an HDFS encryption zone. For CDH 5.7.0 and lower, you will need to manually configure HDFS trash to allow deletions. For details on how to configure trash in HDFS, see Trash Behavior with HDFS Transparent Encryption Enabled.

DistCp Considerations

A common use case for DistCp is to replicate data between clusters for backup and disaster recovery purposes. This is typically performed by the cluster administrator, who is an HDFS superuser. To retain this workflow when using HDFS encryption, a new virtual path prefix has been introduced, `/.reserved/raw/`, that gives superusers direct access to the underlying block data in the filesystem. This allows superusers to `distcp` data without requiring access to encryption keys, and avoids the overhead of decrypting and re-encrypting data. It also means the source and destination data will be byte-for-byte identical, which would not have been true if the data was being re-encrypted with a new EDEK.

> **Warning:**
>
> When using `/.reserved/raw/` to `distcp` encrypted data, make sure you preserve extended attributes with the `-px` flag. This is because encrypted attributes such as the EDEK are exposed through extended attributes and *must* be preserved to be able to decrypt the file.
>
> This means that if the `distcp` is initiated at or above the encryption zone root, it will automatically create a new encryption zone at the destination if it does not already exist. Hence, Cloudera recommends you first create identical encryption zones on the destination cluster to avoid any potential mishaps.

Copying data from encrypted locations

By default, `distcp` compares checksums provided by the filesystem to verify that data was successfully copied to the destination. When copying from an encrypted location, the file system checksums will not match because the underlying block data is different. This is true whether or not the destination location is encrypted or unencrypted.

In this case, you can specify the `-skipcrccheck` and `-update` flags to avoid verifying checksums. When you use -skipcrccheck, `distcp` checks the file integrity by performing a file size comparison, right after the copy completes for each file.

Deleting Encryption Zones

To remove an encryption zone, delete the encrypted directory:

> **Warning:** This command deletes the entire directory and all of its contents. Ensure that the data is no longer needed before running this command.

```
$ sudo -u hdfs hadoop fs -rm -r -skipTrash /encryption_zone
```

> **Important:** The Key Trustee KMS does not directly execute a key deletion (for example, it may perform a soft delete instead, or delay the actual deletion to prevent mistakes). In these cases, errors may occur when creating or deleting a key using the same name after it has already been deleted.

Backing Up Encryption Keys

> **Warning:** It is *very important* that you regularly back up your encryption keys. Failure to do so can result in irretrievable loss of encrypted data.

If you are using the Java KeyStore KMS, make sure you regularly back up the Java KeyStore that stores the encryption keys. If you are using the Key Trustee KMS and Key Trustee Server, see Backing Up and Restoring Key Trustee Server and Clients on page 299 for instructions on backing up Key Trustee Server and Key Trustee KMS.

Rolling Encryption Keys

Before attempting to roll an encryption key (also known as an encryption zone key, or EZ key), familiarize yourself with the concepts described in Cloudera Navigator Data Encryption Overview, and HDFS Transparent Encryption, as the material in these sections presumes you are familiar with the fundamentals of HDFS transparent encryption and Cloudera data at rest encryption.

When you roll an EZ key, you are essentially creating a new version of the key (`ezKeyVersionName`). Rolling EZ keys regularly helps enterprises minimize the risk of key exposure. If a malicious attacker were to obtain the EZ key and decrypt encrypted data encryption keys (EDEKs) into DEKs, they could gain the ability to decrypt HDFS files. Rolling an EZ key ensures that all DEKs for newly-created files will be encrypted with the new version of the EZ key. The older EZ

key version that the attacker obtained cannot decrypt these EDEKs. You may want to roll the encryption key periodically, as part of your security policy or when an external security compromise is detected.

To roll an EZ key:

1. (***Optional***) Before rolling any keys, log in as HDFS Superuser and verify/identify the encryption zones to which the current key applies. This operation also helps clarify the relationship between the EZ key and encryption zones, and, if necessary, makes it easier to identify more important, high priority zones:

```
$ hdfs crypto –listZones
/ez key1
/ez2 key2
/user key1
```

   The first column identifies the encryption zone paths; the second column identifies the encryption key name.

2. (***Optional***) You can verify that the files inside an encryption zone are encrypted using the `hdfs crypto -getFileEncryptionInfo` command. Note the EZ key version name and value, which you can use for comparison and verification after rolling the EZ key.

```
$ hdfs crypto –getFileEncryptionInfo –path /ez/f
{cipherSuite: {name: AES/CTR/NoPadding, algorithmBlockSize: 16}. cryptoProtocolVersion:
CryptoProtocolVersion{description=’Encryption zones’, version=2, unknownValue=null},
edek: 373c0c2e919c27e58c1c343f54233cbd,
iv: d129c913c8a34cde6371ec95edfb7337, keyName: key1, ezKeyVersionName:
7mbvopZ0Weuvs0XtTkpGw3G92KuWc4e4xcTXl0bXCpF}
```

   Log off as HDFS Superuser.

3. Log in as Key Administrator. Because keys can be rolled, a key can have multiple key versions, where each key version has its own key material (the actual secret bytes used during DEK encryption and EDEK decryption). You can fetch an encryption key by either its key name, returning the latest version of the key, or by a specific key version.

   Roll the encryption key (previously identified/confirmed by the HDFS Superuser in step 1. Here, the `<key name>` is `key1`:

```
$ hadoop key roll key1
```

   This operation contacts the KMS and rolls the keys there. Note that this can take a considerable amount of time, depending on the number of key versions residing in the KMS.

```
Rolling key version from KeyProvider:
org.apache.hadoop.crypto.key.kms.LoadBalancingKMSClientProvider@5ea434c8
for keyName: key1
key1 has been successfully rolled.
org.apache.hadoop.crypto.key.kms.LoadBalancingKMSClientProvider@5ea434c8 has been updated.
```

4. (***Optional***) Log out as Key Administrator, and log in as HDFS Superuser. Verify that new files in the encryption zone have a new EZ key version.

   Note: For performance reasons, the NameNode caches EDEKs, so after rolling an encryption key, you may not be able to see the new version encryption key immediately, or at least until after the EDEK cache is consumed. Of course, the file decryption and encryption still works with these EDEKs. If you require that all files' DEKs in an encryption zone are encrypted using the latest version encryption key, please re-encrypt the EDEKs.

```
$ hdfs crypto –getFileEncryptionInfo –path /ez/new_file
{cipherSuite: {name: AES/CTR/NoPadding, algorithmBlockSize: 16}. cryptoProtocolVersion:
CryptoProtocolVersion{description=’Encryption zones’, version=2, unknownValue=null},
edek: 9aa13ea4a700f96287cfe1349f6ff4f2,
iv: 465c878ad9325e42fa460d2a22d12a72, keyName: key1, ezKeyVersionName:
4tuvorJ6Feeqk8WiCfdDs9K32KuEj7g2ydCAv0gNQbY}
```

Alternatively, you can use KMS Rest API to view key metadata and key versions. Elements appearing in brackets should be replaced with your actual values. So in this case, before rolling a key, you can view the key metadata and versions as follows:

```
$ curl -k --negotiate -u: "https://<KMS_HOSTNAME>:16000/kms/v1/key/<key-name>/_metadata"
{
  "name" : "<key-name>",
  "cipher" : "<cipher>",
  "length" : <length>,
  "description" : "<decription>",
  "created" : <millis-epoc>,
  "versions" : <versions> (For example, 1)
}

$ curl -k  --negotiate -u:
"https://<KMS_HOSTNAME>:16000/kms/v1/key/<key-name>/_currentversion"
{
  "material" : "<material>",
  "name" : "<key-name>",
  "versionName" : "<versionName>" (For example, version 1)
}
```

Roll the key and compare the results:

```
$ hadoop key roll key1

Rolling key version from KeyProvider:
KMSClientProvider[https://<KMS_HOSTNAME>:16000/kms/v1/]

  for key name: <key-name>

key1 has been successfully rolled.

KMSClientProvider[https://<KMS_HOSTNAME>/kms/v1/] has been updated.

$ curl -k  --negotiate -u:
"https://<KMS_HOSTNAME>:16000/kms/v1/key/<key-name>/_currentversion"

{
  "material" : "<material>", (New material)
  "name" : "<key-name>",
  "versionName" : "<versionName>" (New version name. For example, version 2)
}

$ curl -k --negotiate -u: "https://<KMS_HOSTNAME>:16000/kms/v1/key/<key-name>/_metadata"
{
  "name" : "<key-name>",
  "cipher" : "<cipher>",
  "length" : <length>,
  "description" : "<decription>",
  "created" : <millis-epoc>,
  "versions" : <versions> (For example, version 2)
}
```

Re-encrypting Encrypted Data Encryption Keys (EDEKs)

Before attempting to re-encrypt an EDEK, familiarize yourself with the concepts described in Cloudera Navigator Data Encryption Overview, and HDFS Transparent Encryption, as the material in this section presumes you are familiar with the fundamentals of HDFS transparent encryption and Cloudera data at rest encryption.

When you re-encrypt an EDEK, you are essentially decrypting the original EDEK created by the DEK, and then re-encrypting it using the new (rolled) version of the EZ key (see Rolling Encryption Keys on page 259). The file's metadata, which is stored in the NameNode, is then updated with this new EDEK. Re-encryption does not impact the data in the HDFS files or the DEK–the same DEK is still used to decrypt the file, so re-encryption is essentially transparent.

Benefits and Capabilities

In addition to minimizing security risks, re-encrypting the EDEK offers the following capabilities and benefits:

- Re-encrypting EDEKs does *not* require that the user explicitly re-encrypt HDFS files.
- In cases where there are several zones using the same key, the Key Administrator has the option of selecting which zone's EDEKs are re-encrypted first.
- The HDFS Superuser can also monitor and cancel re-encryption operations.
- Re-encryption is restarted automatically in cases where you have a NameNode failure during the re-encryption operation.

Prerequisites and Assumptions

Before attempting to re-encrypt an EDEK, familiarize yourself with the concepts and rules described in Managing Encryption Keys and Zones on page 257.

- It is recommended that you perform EDEK re-encryption at the same time that you perform regular cluster maintenance because the operation can adversely impact CPU resources on the NameNode.
- The following account credentials and settings must exist before you can re-encrypt EDEKs: CDH 5.13.0 with KT KMS 5.13.0.
- In Cloudera Manager, review the cluster's NameNode status, which must be in "Good Health". If the cluster NameNode does not have a status of "Good Health", then do not proceed with the re-encryption of the EDEK. In the Cloudera Manager WebUI menu, you can verify the status for the cluster NameNode, which must *not* be in Safe mode (in other words, the WebUI should indicate "Safemode is off").

Running the re-encryption command without successfully verifying the preceding items will result in failures with errors.

Limitations

This section identifies limitations associated with the re-encryption of EDEKs.

EDEK re-encryption doesn't change EDEKs on snapshots, due to the immutable nature HDFS snapshots. Thus, you should be aware that after EZ key exposure, the Key Administrator must delete snapshots.

Re-encrypting an EDEK

This scenario operates on the assumption that an encryption zone has already been set up for this cluster. For more details about creating an encryption zone, see Creating Encryption Zones on page 257.

1. Navigate to the cluster in which you will be rolling keys and re-encrypting the EDEK.
2. Log in as HDFS Superuser.
3. (***Optional***) To view all of the options for the `hdfs crypto` command:

```
$ hdfs crypto
[-createZone –keyName <keyName> -path <path>]
[-listZones]
[-provisionTrash –path <path>]
[-getFileEncryptionInfo –path <path>]
[-reencryptZone <action> -path <zone>]
[-listReencryptionStatus]
[-help <command-name>]
```

4. Before rolling any keys, verify/identify the encryption zones to which the current key applies. This operation also helps clarify the relationship between the EZ key and encryption zones, and, if necessary, makes it easier to identify more important, high priority zones:

```
$ hdfs crypto –listZones
/ez     key1
```

   The first column identifies the encryption zone path (`/ez`); the second column identifies the encryption key name (`key1`).
5. Exit from the HDFS Superuser account and log in as Key Administrator.

6. Roll the encryption key (previously identified/confirmed by the HDFS Superuser in step 4). Here, the <key name> is key1:

```
$ hadoop key roll key1
```

This operation contacts the KMS and rolls the keys. Note that this can take a considerable amount of time, depending on the number of key versions.

```
Rolling key version from KeyProvider:
org.apache.hadoop.crypto.key.kms.LoadBalancingKMSClientProvider@5ea434c8
for keyName: key1
key1 has been successfully rolled.
org.apache.hadoop.crypto.key.kms.LoadBalancingKMSClientProvider@5ea434c8 has been updated.
```

7. Log out as Key Administrator, and log in as HDFS Superuser.
8. (***Optional***) Before performing the re-encryption, you can verify the status of the current key version being used (keyName). Then, after re-encrypting, you can confirm that the EZ key version (`ezKeyVersionName`) and EDEK have changed:

```
$ hdfs crypto –getFileEncryptionInfo –path /ez/f
{cipherSuite: {name: AES/CTR/NoPadding, algorithmBlockSize: 16}. cryptoProtocolVersion:
CryptoProtocolVersion{description='Encryption zones', version=2, unknownValue=null},
edek: 9aa13ea4a700f96287cfe1349f6ff4f2,
iv: d129c913c8a34cde6371ec95edfb7337, keyName: key1, ezKeyVersionName:
7mbvopZ0Weuvs0XtTkpGw3G92KuWc4e4xcTXl0bXCpF}
```

9. After the EZ key has been rolled successfully, re-encrypt the EDEK by running the re-encryption command on the encryption zone:

```
$ hdfs crypto –reencryptZone –start –path /ez
```

The following information appears when the submission is complete. At this point, the NameNode is processing and re-encrypting all of the EDEKs under the `/ez` directory.

```
re-encrypt command successfully submitted for zone: /ez action:  START:
```

Depending on the number of files, the re-encryption operation can take a long time. Re-encrypting a 1M EDEK file typically takes between 2-6 minutes, depending on the NameNode hardware. To check the status of the re-encryption for the zone:

```
$ hdfs crypto –listReencryptionStatus
```

Table 20: Re-encryption Status Column Data

Column Name	Description	Sample Data
ZoneName	The encryption zone name	/ez
Status	• Submitted: the command is received, but not yet being processed by the NameNode. • Processing: the zone is being processed by the NameNode. • Completed: the NameNode has finished processing the zone, and every file in the zone has been re-encrypted.	Completed
EZKey Version Name	The encryption zone key version name, which used for re-encryption	7MHRoGKeXgf0QcO8q16Nzw2sq0WRTOHS3YCz

Column Name	Description	Sample Data
	comparison. After re-encryption is complete, all files in the encryption zone are guaranteed to have an EDEK whose encryption zone key version is at least equal to this version.	
Submission Time	The time at which the re-encryptiom operation commenced.	2017-09-07 10:01:09,262-0700
Is Canceled?	True: the encryption operation has been canceled. False: the encryption operation has not been canceled.	False
Completion Time	The time at which the re-encryption operation completed.	2017-09-07 10:01:10,441-0700
Number of files re-encrypted	The number that appears in this column reflects *only* the files whose EDEKs have been updated. If a file is created after the key is rolled, then it will already have an EDEK that has been encrypted by the new key version, so the re-encryption operation will skip that file. In other words, it's possible for a "Completed" re-encryption to reflect a number of re-encrypted files that is less than the number of files actually in the encryption zone. **Note:** In cases when you re-encrypt an EZ key that has already been re-encrypted and there are no new files, the number of files re-encrypted will be 0.	1
Number of failures	When 0, no errors occurred during the re-encryption operation. If larger than 0, then investigate the NameNode log, and re-encrypt.	0
Last file Checkpointed	Identifies the current position of the re-encryption process in the encryption zone--in other words, the file that was most recently re-encrypted.	0

10. (***Optional***) After the re-encryption completes, you can confirm that the EDEK and EZ Key Version Name values have changed:

```
$ hdfs crypto –getFileEncryptionInfo –path /ez/f
{cipherSuite: {name: AES/CTR/NoPadding, algorithmBlockSize: 16}. cryptoProtocolVersion:
CryptoProtocolVersion{description=’Encryption zones’, version=2, unknownValue=null},
edek: 373c0c2e919c27e58c1c343f54233cbd,
iv: d129c913c8a34cde6371ec95edfb7337, keyName: key1, ezKeyVersionName:
ZMHfRoGKeXXgf0QzCX8q16NczIw2sq0rWRTOHS3YjCz }
```

Managing Re-encryption Operations

This section includes information that can help you manage various facets of the EDEK re-encryption process.

Cancelling Re-encryption

Only users with the HDFS Superuser privilege can cancel the EDEK re-encryption after the operation has started.

To cancel a re-encryption:

```
$ hadoop crypto -reencryptZone cancel -path <zone>
```

Rolling Keys During a Re-encryption Operation

While it is *not* recommended, it is possible to roll the encryption zone key version on the KMS while a re-encryption of that encryption zone is already in progress in the NameNode. The re-encryption is guaranteed to complete with all DEKs re-encrypted, with a key version equal to or later than the encryption zone key version when the re-encryption command was submitted. This means that, if initially the key version is rolled from v0 to v1, then a re-encryption command was submitted. If later on the KMS the key version is rolled again to v2, then all EDEKs will be at least re-encrypted to v1. To ensure that all EDEKs are re-encrypted to v2, submit another re-encryption command for the encryption zone.

Rolling keys during re-encryption is not recommended because of the potential negative impact on key management operations. Due to the asynchronous nature of re-encryption, there is no guarantee of when, exactly, the rolled encryption keys will take effect. Re-encryption can only guarantee that all EDEKs are re-encrypted at least on the EZ key version that existed when the re-encryption command is issued.

Throttling Re-encryption Operations

With the default operation settings, you will not typically need to throttle re-encryption operations. However, in cases of excessive performance impact due to the re-encryption of large numbers of files, advanced users have the option of throttling the operation so that the impact on the HDFS NameNode and KT KMS are minimized.

Specifically, you can throttle the operation to control the rate of the following:

- The number of EDEKs that the NameNode should send to the KMS to re-encrypt in a batch (dfs.namenode.reencrypt.batch.size)
- The number of threads in the NameNode that can run concurrently to contact the KMS. (dfs.namenode.reencrypt.edek.threads)
- Percentage of time the NameNode read-lock should be held by the re-encryption thread (dfs.namenode.reencrypt.throttle.limit.handler.ratio)
- Percentage of time the NameNode write-lock should be held by the re-encryption thread (dfs.namenode.reencrypt.throttle.limit.updater.ratio)

You can monitor the HDFS NameNode heap and CPU usage from Cloudera Manager.

Configuring the Key Management Server (KMS)

Hadoop Key Management Server (KMS) is a cryptographic key management server based on the Hadoop **KeyProvider** API. It provides a KeyProvider implementation client that interacts with the KMS using the HTTP REST API. Both the KMS and its client support HTTP SPNEGO Kerberos authentication and TLS/SSL-secured communication. The KMS is a Java-based web application that uses a preconfigured Tomcat server bundled with the Hadoop distribution.

For instructions on securing the KMS, see Securing the Key Management Server (KMS) on page 270.

Cloudera provides the following implementations of the Hadoop KMS:

- **Java KeyStore KMS** - The default Hadoop KMS included in CDH that uses a file-based Java KeyStore (JKS) for its backing keystore. For parcel-based installations, no additional action is required to install or upgrade the KMS. For package-based installations, you must install additional packages. For more information, see Installing and

Upgrading Java KeyStore KMS. Cloudera strongly recommends not using Java Keystore KMS in production environments.

- **Key Trustee KMS** - A custom KMS that uses Cloudera Navigator Key Trustee Server for its backing keystore instead of the file-based Java KeyStore (JKS) used by the default Hadoop KMS. Cloudera strongly recommends using Key Trustee KMS in production environments to improve the security, durability, and scalability of your cryptographic key management. For more information about the architecture and components involved in encrypting data at rest for production environments, see Cloudera Navigator Data Encryption Overview and Data at Rest Encryption Reference Architecture on page 238. For instructions on installing and upgrading Key Trustee KMS, see:
  - Installing Key Trustee KMS
  - Upgrading Key Trustee KMS

  Also, integrating Key Trustee Server with Cloudera Navigator Key HSM provides an additional layer of protection.
- **Navigator KMS Services backed by Thales HSM** - A custom KMS that uses a supported Thales Hardware Security Module (HSM) as its backing keystore. This KMS service provides the highest level of key isolation to customers who require it.

  For installation information about Navigator KMS Services backed by Thales HSM, see Installing Navigator HSM KMS Backed by Thales HSM.
- **Navigator KMS Services backed by Luna HSM** - A custom KMS that uses a supported Luna Hardware Security Module (HSM) as its backing keystore. This KMS provides the highest level of key isolation to customers who require it.

  For installation information about Navigator KMS Services backed by Luna HSM, see Installing Navigator HSM KMS Backed by Luna HSM.

Configuring KMS High Availability

For Key Trustee KMS high availability, see Enabling Key Trustee KMS High Availability. Java KeyStore KMS does not support high availability.

Configuring the KMS Using Cloudera Manager

If you are using Cloudera Manager, you can view and edit the KMS configuration by navigating to the following pages, depending on the KMS implementation you are using:

- **Key Trustee KMS service** > **Configuration**
- **Java KeyStore KMS service** > **Configuration**

For more information on using Cloudera Manager to find and change configuration parameters, see Modifying Configuration Properties Using Cloudera Manager.

For instructions about configuring the KMS and its clients using the command line for package-based installations, continue reading:

Configuring the KMS Cache Using Cloudera Manager

By default, the KMS caches keys to reduce the number of interactions with the key provider. You can disable the cache by setting the `hadoop.kms.cache.enable` property to `false`.

The cache is only used with the `getCurrentKey()`, `getKeyVersion()` and `getMetadata()` methods.

For the `getCurrentKey()` method, entries are cached for a maximum of `30000` milliseconds to prevent stale keys.

For the `getKeyVersion()` method, entries are cached with a default inactivity timeout of `600000` milliseconds (10 minutes).

You can configure the cache and its timeout values by adding the following properties to **KMS service** > **Configuration** > **Advanced** > **Key Management Server Proxy Advanced Configuration Snippet (Safety Valve) for kms-site.xml**:

```
<property>
    <name>hadoop.kms.cache.enable</name>
    <value>true</value>
</property>

<property>
    <name>hadoop.kms.cache.timeout.ms</name>
    <value>600000</value>
</property>

<property>
    <name>hadoop.kms.current.key.cache.timeout.ms</name>
    <value>30000</value>
</property>
```

See Custom Configuration for more information on adding custom properties using the **Advanced Configuration Snippet (Safety Valve)** feature.

Configuring the Audit Log Aggregation Interval Using the Command Line

Audit logs are generated for `GET_KEY_VERSION`, `GET_CURRENT_KEY`, `DECRYPT_EEK`, and `GENERATE_EEK` operations.

Entries are aggregated by user, key, and operation for a configurable interval, after which the number of aggregated operations by the user for a given key is written to the audit log.

The interval is configured in milliseconds by adding the `hadoop.kms.aggregation.delay.ms` property to **KMS service** > **Configuration** > **Advanced** > **Key Management Server Proxy Advanced Configuration Snippet (Safety Valve) for kms-site.xml**:

```
<property>
    <name>hadoop.kms.aggregation.delay.ms</name>
    <value>10000</value>
</property>
```

For more information about adding custom properties using the **Advanced Configuration Snippet (Safety Valve)** feature, see Custom Configuration.

Configuring the Java KeyStore KMS Using the Command Line

> **Note:** Because Key Trustee KMS is supported only in Cloudera Manager deployments, the following command line instructions apply only to Java KeyStore KMS. For instructions about configuring Key Trustee KMS, see Configuring the KMS Using Cloudera Manager on page 266.

For instructions about configuring the Java KeyStore KMS and its clients using the command line for package-based installations, continue reading:

Configuring the Java KeyStore KMS KeyProvider Using the Command Line

> **Important:**
> - Follow these command-line instructions on systems that do not use Cloudera Manager.
> - This information applies specifically to CDH 5.14.x. See Cloudera Documentation for information specific to other releases.

Configure the KMS backing KeyProvider properties in the `/etc/hadoop-kms/conf/kms-site.xml` configuration file:

```
<property>
    <name>hadoop.kms.key.provider.uri</name>
    <value>jceks://file@/${user.home}/kms.keystore</value>
</property>

<property>
    <name>hadoop.security.keystore.java-keystore-provider.password-file</name>
    <value>keystore_password_file</value>
</property>
```

If you do not specify the absolute path to the password file, you must include it in the Hadoop `CLASSPATH`.

Restart the KMS for configuration changes to take effect. See Starting and Stopping the Java KeyStore KMS Using the Command Line on page 269 for instructions.

Configuring the Java KeyStore KMS Cache Using the Command Line

Important:

- Follow these command-line instructions on systems that do not use Cloudera Manager.
- This information applies specifically to CDH 5.14.x. See Cloudera Documentation for information specific to other releases.

By default, the KMS caches keys to reduce the number of interactions with the key provider. You can disable the cache by setting the `hadoop.kms.cache.enable` property to `false`.

The cache is only used with the `getCurrentKey()`, `getKeyVersion()` and `getMetadata()` methods.

For the `getCurrentKey()` method, entries are cached for a maximum of `30000` milliseconds to prevent stale keys.

For the `getKeyVersion()` method, entries are cached with a default inactivity timeout of `600000` milliseconds (10 minutes).

The cache and its timeout values are configured using the following properties in the `/etc/hadoop-kms/conf/kms-site.xml` configuration file:

```
<property>
    <name>hadoop.kms.cache.enable</name>
    <value>true</value>
</property>

<property>
    <name>hadoop.kms.cache.timeout.ms</name>
    <value>600000</value>
</property>

<property>
    <name>hadoop.kms.current.key.cache.timeout.ms</name>
    <value>30000</value>
</property>
```

Configuring KMS Clients Using the Command Line

Important:

- Follow these command-line instructions on systems that do not use Cloudera Manager.
- This information applies specifically to CDH 5.14.x. See Cloudera Documentation for information specific to other releases.

To configure KMS clients, set the `hadoop.security.key.provider.path` property in `core-site.xml` or `hdfs-site.xml`. Specify the value in the format `kms://<scheme>@<kms_hosts>:<port>/kms`. Replace *<scheme>* with `http` or `https`, depending on whether you have configured TLS. Replace *<kms_hosts>* with a semicolon-separated list of the KMS hosts. Replace *<port>* with the port number on which the KMS is running (`16000` by default).

For example, for a KMS running on `http://localhost:16000/kms`, the KeyProvider URI is `kms://http@localhost:16000/kms`. For high availability KMS (Key Trustee KMS only) running on `https://kms01.example.com:16000/kms` and `https://kms02.example.com:16000/kms`, the KeyProvider URI is `kms://https@kms01.example.com;kms02.example.com:16000/kms`.

See the following for an excerpt from `core-site.xml`:

```
<property>
  <name>hadoop.security.key.provider.path</name>
  <value>kms://https@kms01.example.com;kms02.example.com:16000/kms</value>
</property>
```

Starting and Stopping the Java KeyStore KMS Using the Command Line

Important:

- Follow these command-line instructions on systems that do not use Cloudera Manager.
- This information applies specifically to CDH 5.14.x. See Cloudera Documentation for information specific to other releases.

To start or stop KMS use the `kms.sh` script. For example, to start the KMS:

```
$ sudo /usr/lib/hadoop-kms/sbin/kms.sh start
```

Running the script without parameters lists all possible parameters.

To use an `init` script to manage the KMS service, use your package manager to install the `hadoop-kms-server` package from the CDH repository. For example, for RHEL 6:

```
$ sudo yum install hadoop-kms-server
```

After installation, use the `service hadoop-kms-server` command to manage the KMS service.

Configuring the Audit Log Aggregation Interval Using the Command Line

Important:

- Follow these command-line instructions on systems that do not use Cloudera Manager.
- This information applies specifically to CDH 5.14.x. See Cloudera Documentation for information specific to other releases.

Audit logs are generated for `GET_KEY_VERSION`, `GET_CURRENT_KEY`, `DECRYPT_EEK`, and `GENERATE_EEK` operations.

Entries are aggregated by user, key, and operation for a configurable interval, after which the number of aggregated operations by the user for a given key is written to the audit log.

The interval is configured in milliseconds using the `hadoop.kms.aggregation.delay.ms` property:

```
<property>
    <name>hadoop.kms.aggregation.delay.ms</name>
    <value>10000</value>
</property>
```

Configuring the Embedded Tomcat Server Using the Command Line

> **Important:**
> - Follow these command-line instructions on systems that do not use Cloudera Manager.
> - This information applies specifically to CDH 5.14.x. See Cloudera Documentation for information specific to other releases.

You can configure the embedded Tomcat server by using the `/etc/hadoop-kms/tomcat-conf/conf/server.xml.conf` file.

The following environment variables can be set in KMS `/etc/hadoop-kms/conf/kms-env.sh` script and can be used to alter the default ports and log directory:

- `KMS_HTTP_PORT`
- `KMS_ADMIN_PORT`
- `KMS_LOG`

Restart the KMS for the configuration changes to take effect.

Securing the Key Management Server (KMS)

Cloudera provides the following implementations of the Hadoop KMS:

- **Java KeyStore KMS** - The default Hadoop KMS included in CDH that uses a file-based Java KeyStore (JKS) for its backing keystore. For parcel-based installations, no additional action is required to install or upgrade the KMS. For package-based installations, you must install additional packages. For more information, see Installing and Upgrading Java KeyStore KMS. Cloudera strongly recommends not using Java Keystore KMS in production environments.
- **Key Trustee KMS** - A custom KMS that uses Cloudera Navigator Key Trustee Server for its backing keystore instead of the file-based Java KeyStore (JKS) used by the default Hadoop KMS. Cloudera strongly recommends using Key Trustee KMS in production environments to improve the security, durability, and scalability of your cryptographic key management. For more information about the architecture and components involved in encrypting data at rest for production environments, see Cloudera Navigator Data Encryption Overview and Data at Rest Encryption Reference Architecture on page 238. For instructions on installing and upgrading Key Trustee KMS, see:
  - Installing Key Trustee KMS
  - Upgrading Key Trustee KMS

  Also, integrating Key Trustee Server with Cloudera Navigator Key HSM provides an additional layer of protection.
- **Navigator KMS Services backed by Thales HSM** - A custom KMS that uses a supported Thales Hardware Security Module (HSM) as its backing keystore. This KMS service provides the highest level of key isolation to customers who require it.

  For installation information about Navigator KMS Services backed by Thales HSM, see Installing Navigator HSM KMS Backed by Thales HSM.
- **Navigator KMS Services backed by Luna HSM** - A custom KMS that uses a supported Luna Hardware Security Module (HSM) as its backing keystore. This KMS provides the highest level of key isolation to customers who require it.

  For installation information about Navigator KMS Services backed by Luna HSM, see Installing Navigator HSM KMS Backed by Luna HSM.

This topic contains information on securing the KMS using Kerberos, TLS/SSL communication, and access control lists (ACLs) for operations on encryption keys. Cloudera Manager instructions can be performed for both Key Trustee KMS and Java KeyStore KMS deployments. Command-line instructions apply only to Java KeyStore KMS deployments. Key Trustee KMS is not supported outside of Cloudera Manager. For more information, see Installing Key Trustee KMS.

Enabling Kerberos Authentication for the KMS

Enabling Kerberos Authentication for the KMS Using Cloudera Manager

Minimum Required Role: Full Administrator

To enable Kerberos for the KMS using Cloudera Manager:

1. Open the Cloudera Manager Admin Console and go to the KMS service.
2. Click **Configuration**.
3. Set the **Authentication Type** property to `kerberos`.
4. Click **Save Changes**.
5. Because Cloudera Manager does not automatically create the principal and keytab file for the KMS, you must run the **Generate Credentials** command manually. On the top navigation bar, go to **Administration** > **Security** > **Kerberos Credentials**and click **Generate Missing Credentials**.

   > **Note:** This does not create a new Kerberos principal if an existing HTTP principal exists for the KMS host.

6. Return to the Home page by clicking the Cloudera Manager logo.
7. Click the icon next to any stale services to invoke the cluster restart wizard.
8. Click **Restart Stale Services**.
9. Click **Restart Now**.
10. Click **Finish**.

Enabling Kerberos Authentication for the Java KeyStore KMS Using the Command Line

> **Important:**
> - Follow these command-line instructions on systems that do not use Cloudera Manager.
> - This information applies specifically to CDH 5.14.x. See Cloudera Documentation for information specific to other releases.

Configure `/etc/krb5.conf` with information for your KDC server. Create an HTTP principal and keytab file for the KMS.

Configure `/etc/hadoop-kms/conf/kms-site.xml` with the following properties:

```
<property>
    <name>hadoop.kms.authentication.type</name>
    <value>kerberos</value>
</property>

<property>
    <name>hadoop.kms.authentication.kerberos.keytab</name>
    <value>${user.home}/kms.keytab</value>
</property>

<property>
    <name>hadoop.kms.authentication.kerberos.principal</name>
    <value>HTTP/localhost</value>
</property>

<property>
    <name>hadoop.kms.authentication.kerberos.name.rules</name>
    <value>DEFAULT</value>
</property>
```

Restart the KMS service for the configuration changes to take effect.

Configuring the Java KeyStore KMS Proxyuser Using the Command Line

> **Important:**
> - Follow these command-line instructions on systems that do not use Cloudera Manager.
> - This information applies specifically to CDH 5.14.x. See Cloudera Documentation for information specific to other releases.

Each proxyuser must be configured in `/etc/hadoop-kms/conf/kms-site.xml` using the following properties:

```
<property>
    <name>hadoop.kms.proxyuser.#USER#.users</name>
    <value>*</value>
</property>

<property>
    <name>hadoop.kms.proxyuser.#USER#.groups</name>
    <value>*</value>
</property>

<property>
    <name>hadoop.kms.proxyuser.#USER#.hosts</name>
    <value>*</value>
</property>
```

where #USER# is the username of the proxyuser to be configured.

The `hadoop.kms.proxyuser.#USER#.users` property indicates the users that can be impersonated. The `hadoop.kms.proxyuser.#USER#.groups` property indicates the groups to which the users being impersonated must belong. At least one of these properties must be defined. If both are defined, the configured proxyuser can impersonate any user in the `users` list and any user belonging to a group listed in the `groups` list.

The `hadoop.kms.proxyuser.#USER#.hosts` property indicates the host from which the proxyuser can make impersonation requests. "*" means there are no restrictions for the #USER# regarding users, groups, or hosts.

Configuring TLS/SSL for the KMS

Configuring TLS/SSL for the KMS Using Cloudera Manager

Minimum Required Role: Configurator (also provided by **Cluster Administrator, Full Administrator**)

The steps for configuring and enabling Hadoop TLS/SSL for the KMS are as follows:

1. Go to the KMS service.
2. Click **Configuration**.
3. In the Search field, type **TLS/SSL** to show the KMS TLS/SSL properties (in the **Key Management Server Default Group** > **Security** category).
4. Edit the following TLS/SSL properties according to your cluster configuration.

Table 21: KMS TLS/SSL Properties

Property	Description
Enable TLS/SSL for Key Management Server	Encrypt communication between clients and Key Management Server using Transport Layer Security (TLS) (formerly known as Secure Socket Layer (TLS/SSL)).
Key Management Server TLS/SSL Server JKS Keystore File Location	The path to the TLS/SSL keystore file containing the server certificate and private key used for TLS/SSL. Used when Key Management Server is acting as a TLS/SSL server. The keystore must be in JKS format.

Property	Description
Key Management Server TLS/SSL Server JKS Keystore File Password	The password for the Key Management Server JKS keystore file.
Key Management Server Proxy TLS/SSL Certificate Trust Store File	The location on disk of the truststore, in .jks format, used to confirm the authenticity of TLS/SSL servers that Key Management Server Proxy might connect to. This is used when Key Management Server Proxy is the client in a TLS/SSL connection. This truststore must contain the certificates used to sign the services connected to. If this parameter is not provided, the default list of well-known certificate authorities is used instead.
Key Management Server Proxy TLS/SSL Certificate Trust Store Password	The password for the Key Management Server Proxy TLS/SSL Certificate Trust Store File. This password is not required to access the truststore; this field can be left blank. This password provides optional integrity checking of the file. The contents of truststores are certificates, and certificates are public information.

5. Click **Save Changes**.
6. Return to the Home page by clicking the Cloudera Manager logo.
7. Click the icon next to any stale services to invoke the cluster restart wizard.
8. Click **Restart Stale Services**.
9. Click **Restart Now**.
10. Click **Finish**.

For help troubleshooting TLS/SSL for KMS configuration issues, see Troubleshooting TLS/SSL Issues in Cloudera Manager on page 429.

Configuring TLS/SSL for the Java KeyStore KMS Using the Command Line

Important:

- Follow these command-line instructions on systems that do not use Cloudera Manager.
- This information applies specifically to CDH 5.14.x. See Cloudera Documentation for information specific to other releases.

To configure KMS to work over HTTPS, set the following properties in the `/etc/hadoop-kms/conf/kms_env.sh` script:

- `KMS_SSL_KEYSTORE_FILE`
- `KMS_SSL_KEYSTORE_PASS`
- `KMS_SSL_TRUSTSTORE_FILE`
- `KMS_SSL_TRUSTSTORE_PASS`

In the `/etc/hadoop-kms/tomcat-conf/conf/` directory, replace the `server.xml` file with the provided `ssl-server.xml` file.

Create a TLS/SSL certificate for the KMS. As the `kms` user, use the Java `keytool` command to create the TLS/SSL certificate:

```
$ keytool -genkey -alias tomcat -keyalg RSA
```

You are asked a series of questions in an interactive prompt. It creates the keystore file, which is named **.keystore** and located in the `kms` user home directory. The password you enter for the keystore must match the value of the `KMS_SSL_KEYSTORE_PASS` environment variable set in the `kms-env.sh` script in the configuration directory.

The answer to "What is your first and last name?" (CN) must be the hostname of the machine where the KMS will be running.

> **Note:** Restart the KMS for the configuration changes to take effect.

Configuring KMS Access Control Lists (ACLs)

An Access Control List (ACL) is a list of specific permissions or controls that allow individual users, groups, a host, or applications to perform specific actions upon specific objects. The Hadoop KMS supports a range of ACLs that control access to encryption keys and key operations on a granular basis.

KMS ACLs indirectly impact data access by controlling key access, and are decoupled from HDFS file permissions and ACLs. KMS ACLs alone do not directly control data access. Instead, KMS ACLs control whether or not an authorized client can perform a specific operation on a named encryption key.

While KMS ACLs play a primary role in controlling encryption key security, it is important to understand that they are not the only mechanism by which access is controlled. A user's role also factors into the level of access.

Proper configuration of KMS ACLs depends on a variety of variables such as workload, CDH components in use, and how your clusters are configured. This documentation does not take into consideration or describe these outside variables. For details about your specific component's ACL behavior and requirements, refer to the product documentation for the CDH components in your configuration. See the following for additional CDH component ACL information:

- Synchronizing HDFS ACLs and Sentry Permissions
- HDFS Extended ACLs on page 175
- Navigator Encrypt Access Control List on page 340

KMS ACLs and Roles

Cloudera's framework for key management is based on enforcing a secure-by-default configuration based upon the KMS ACLs and roles described here.

Table 22: KMS ACLs and Roles

Role	Description	Allowed To:	Not Allowed To:
Key Administrators	The sole purpose of a Key Administrator is to create and manage keys. This user is whitelisted in a number of areas so that they can handle defined and undefined keys within the context of the KMS.	• Create and manage encryption zone keys • Add, update, or otherwise modify ACLs that protect all encryption zone keys	
HDFS Superusers	Responsible for HDFS administration, HDFS Superusers are not granted rights to decrypt data within encryption zones. Rather, they are authorized to only create zones and attach keys to those zones for the data sets that they manage. HDFS Superusers are usually also HDFS Superusers.	• Create encryption zones	• Manage, create, or read keys • Add, update, or otherwise modify ACLs that protect all keys • Decrypt EEKs

Role	Description	Allowed To:	Not Allowed To:
	Important: To maintain the separation of duties, the HDFS Encryption Wizard automatically blacklists all of the specified HDFS Superuser Groups set in `dfs.permissions.supergroup`. Do not authorize HDFS Superusers to create, manage, or read keys, or to decrypt EEKs.		
HDFS Service User	There is only one HDFS Service User; this is the user the HDFS service runs as within the Hadoop framework. HDFS Service Users are granted special permissions to generate keys (EEKs) that populate per encryption zone key caches.	• Generate keys that are made available for use through the per encryption zone key caches	
End Users	Producers and consumers of data who store or retrieve information within specific encryption zones.		• Read and write in encrypted data spaces

KMS ACL Classes

ACLs are implemented in the upstream Hadoop KMS and apply to any variant of the service, whether it is through Key Trustee Server (KTS), the Java Keystore (JKS), or the HSM KMS. Understanding how these ACLs are defined and work enables you to more accurately create ACLs that meet both your needs and the needs of your customers.

The KMS ACL classes provide fine-grain control of key access at different decision points (for example, preventing a user from accessing the KMS altogether), and are applied either KMS-wide or key-specific:

- **KMS-wide**

  You can use KMS-wide ACLs to specify the types of operations a user can perform. Configure KMS-wide operations using the classes:

  - `hadoop.kms.acl.<OPERATION>`
  - `hadoop.kms.blacklist.<OPERATION>`

  KMS-wide ACLs precede key-specific ACLs. In other words, the permission check on a key-specific ACL only occurs if permission is already granted at the KMS-wide level; if it is, then the permission check against key-specific ACLs can proceed.

  Users accessing a KMS-wide ACL are first checked for inclusion in the ACL for the requested operation, and then checked for exclusion in the blacklist before the operation is performed or access is granted.

- **Key-specific**

  You can use key-specific ACLs, set on a per-encryption zone key basis, to specify the types of operations a user can perform. Configure key-specific operations using the classes:

  - `default.key.acl.<OPERATION>`
  - `whitelist.key.acl.<OPERATION>`
  - `key.acl.<key_name>.<OPERATION>`

> **Important:**
> The KMS supports both whitelist and blacklist ACLs. Blacklist entries apply and are implemented KMS-wide, while whitelist entries apply and are implemented at the key-specific level.

There are five distinct classes of ACLs, each defined by the operations defined within.

The hadoop.kms.acl.<OPERATION> Class

This class controls permissions to perform KMS-wide operations. The only ACLs that fall back to a value of `'*'` when the values are empty are those within the `hadoop.kms.acl` definition.

You can define fine-grain access for the following operations in this ACL class:

Table 23: hadoop.kms.acl.<OPERATION> Class

hadoop.kms.acl. <OPERATION>	<OPERATION> Name	Description
CREATE	Create Key	Creates an encryption zone key. **Note:** Encryption zone key material is returned only if the user also has GET permission.
DELETE	Delete Key	Deletes an encryption zone key.
ROLLOVER	Rollover Key	Rolls an encryption zone key to the new version, using different material. The encrypted encryption keys (also known as EEKs or EDEKs) are always generated using the latest version of the encryption key. **Note:** Encryption zone key material is returned only if the user also has GET permission.
	Invalidate Key Cache	Invalidates all cached EEKs.
GET	Get Current Key	Gets the latest version of the encryption zone key, including its material.
	Get Key Version	Gets a specified version of the encryption zone key, including its material.
	Get Key Versions	Gets all versions of the encryption zone key, including their materials.
GET_KEYS	Get Key Names	Lists all the encryption zone key names on the Key Trustee KMS. No additional key material is returned.
GET_METADATA	Get Key Metadata	Gets the metadata of an encryption zone key, including the cipher, key length, description, time created, and number of versions and other attributes.
	Get Keys Metadata	Gets the metadata for a collection of encryption zone keys.

hadoop.kms.acl. <OPERATION>	<OPERATION> Name	Description
SET_KEY_MATERIAL		For ROLLOVER and/or CREATE encryption key operations, determines whether or not the user can provide key material. If encryption zone key material is not provided, then the Key Trustee KMS randomly generates the material.
GENERATE_EEK	Generate Encrypted Key for Current KeyVersion	Generates an encrypted encryption zone key from a given encrypted key name, using its current key version.
	Re-encrypt Encrypted Key With The Latest KeyVersion	Takes a previously generated EEK, and re-encrypts it using the latest KeyVersion of the same encrypted key. If the latest KeyVersion is the same as the one used to generate the EEK, then the same EEK is returned.
	Batch Re-encrypt Encrypted Keys With The Latest KeyVersion	Takes a batch of previously generated EEKs, and re-encrypts them using the latest KeyVersion of the same encrypted key. If the latest KeyVersion is the same as the one used to generate the EEK, then the same EEK is returned.
DECRYPT_EEK	Decrypt Encrypted Key	Decrypts an EEK and returns the decrypted encryption zone key.

The hadoop.kms.blacklist.<OPERATION> Class

This class controls permissions to perform KMS-wide operations. Users and groups listed here will be prevented from performing any other listed OPERATION on any encryption keys.

You can define fine-grain access for the following operations in this ACL category:

Table 24: hadoop.kms.blacklist.<OPERATION> Class

hadoop.kms.blacklist. <OPERATION>	<OPERATION> Name	Description
CREATE	Create Key	Creates an encryption zone key. **Note:** Encryption zone key material is returned only if the user also has GET permission.
DELETE	Delete Key	Deletes an encryption zone key.
ROLLOVER	Rollover Key	Rolls an encryption zone key to the new version, using different material. The encrypted encryption keys (EEKs) are always generated using the latest version of the encryption key. **Note:** Encryption zone key material is returned only if the user also has GET permission.
	Invalidate Key Cache	Invalidates all cached EEKs.
GET	Get Current Key	Gets the latest version of the encryption zone key, including its material.

hadoop.kms.blacklist. <OPERATION>	<OPERATION> Name	Description
	Get Key Version	Gets a specified version of the encryption zone key, including its material.
	Get Key Versions	Gets all versions of the encryption zone key, including their materials.
GET_KEYS	Get Key Names	Lists all the encryption zone key names on the Key Trustee KMS. No additional key material is returned.
GET_METADATA	Get Key Metadata	Gets the metadata of an encryption zone key, including the cipher, key length, description, time created, and number of versions and other attributes.
	Get Keys Metadata	Gets the metadata for a collection of encryption zone keys.
SET_KEY_MATERIAL		For ROLLOVER and/or CREATE encryption key operations, determines whether or not the user can provide key material. If encryption zone key material is not provided, then the Key Trustee KMS randomly generates the material.
GENERATE_EEK	Generate Encrypted Key for Current KeyVersion	Generates an encrypted encryption zone key (EEK) from a given encrypted key name, using its current key version.
	Re-encrypt Encrypted Key With The Latest KeyVersion	Takes a previously generated EEK, and re-encrypts it using the latest KeyVersion of the same encrypted key. If the latest KeyVersion is the same as the one used to generate the EEK, then the same EEK is returned.
	Batch Re-encrypt Encrypted Keys With The Latest KeyVersion	Takes a batch of previously generated EEKs, and re-encrypts them using the latest KeyVersion of the same encrypted key. If the latest KeyVersion is the same as the one used to generate the EEK, then the same EEK is returned.
DECRYPT_EEK	Decrypt Encrypted Key	Decrypts an EEK and returns the decrypted encryption zone key.

The key.acl.<key-name>.<OPERATION> Class

This class controls permissions to perform operations for a specific key, and applies to key-specific ACLs.

> **Important:**
>
> In the process of evaluating key access, `key.acl.<OPERATION>` overrides default operations provided in `default.key.acl.<OPERATION>` or `whitelist.key.acl.<OPERATION>`.
>
> Be aware that the `key.acl` does not fall back to a value of `'*'` when you use empty values.

You can define fine-grain access for the following operations in this ACL class:

Table 25: key.acl<key-name>.<OPERATION> Class

key.acl.key-name. <OPERATION>	<OPERATION> Name	Description
MANAGEMENT	createKey, deleteKey, rolloverNewVersion	Use for the following operations: • `createKey`

key.acl.key-name. <OPERATION>	<OPERATION> Name	Description
		• deleteKey • rolloverNewKeyVersion
GENERATE_EEK	Generate Encryption Key	Use for the following operations: • generateEncryptedKey • reencryptEncryptedKey • reencryptEncryptedKeys • warmUpEncryptedKeys
DECRYPT_EEK	Decrypt Encrypted Key	Use for the decryptEncryptedKey operation.
READ	Read	Use for the following operations: • getKeyVersion • getKeyVersions • getMetadata • getKeysMetadata • getCurrentKey
ALL	All	Use to specify all operations in this class.

The default.key.acl.<OPERATION> Class

This class controls permission to perform operations for keys that are not otherwise specified by key.acl.<key-name>.<OPERATION>, and applies to key-specific ACLs.

The default.key.acl.<OPERATION> applies to all keys for which an ACL has not been explicitly configured. Be aware that if all of the following conditions exist, key access is denied:

- There is no key-specific ACL configured
- There is no KMS ACL configured for the requested operation
- There is no whitelist key ACL configured for the requested operation

Also note that the default.key.acl does not fall back to a value of '*' when you use empty values. All default.key.acls are (by default) empty, which means that you must create the required key definition entries for each key you wish to use.

Note: The default.key.acl class does *not* support the ALL operation. If specified, it will be ignored.

Table 26: default.key.acl.<OPERATION> Class

default.key.acl. <OPERATION>	<OPERATION> Name	Description
MANAGEMENT	Manage the Zone Key	Use for the following operations: • createKey • deleteKey • rolloverNewKeyVersion
GENERATE_EEK	Create Encryption Key	Use for the decryptEncryptedKey operation.
DECRYPT_EEK	Decrypt Encryption Key	Use for the following operations:

default.key.acl. <OPERATION>	<OPERATION> Name	Description
		• `getKeyVersion` • `getKeyVersions` • `getMetadata` • `getKeysMetadata` • `getCurrentKey`
READ	Read	Use for the following operations: • `getKeyVersion` • `getKeyVersions` • `getMetadata` • `getKeysMetadata` • `getCurrentKey`

The whitelist.key.acl Class

This class controls permissions to perform key operations across all keys, and applies to key-specific ACLs.

Table 27: whitelist.key.acl.<OPERATION> Class

whitelist.key.acl.<OPERATION>	<OPERATION> Name	Description
`MANAGEMENT`	Manage the Zone Key	Use for the following operations: • `createKey` • `deleteKey` • `rolloverNewKeyVersion`
`GENERATE_EEK`	Create Encryption Key	Use for the `decryptEncryptedKey` operation.
`DECRYPT_EEK`	Decrypt Encryption Key	Use for the following operations: • `getKeyVersion` • `getKeyVersions` • `getMetadata` • `getKeysMetadata` • `getCurrentKey`
READ	Read	Use for the following operations: • `getKeyVersion` • `getKeyVersions` • `getMetadata` • `getKeysMetadata` • `getCurrentKey`

> **Important:**
>
> To successfully create an encryption zone, the HDFS Superuser must include the following two entries in the `whitelist.key.acl.<OPERATION>`:
>
> - `GENERATE_EEK`
> - `READ`
>
> These whitelist entries enable the NameNode to build per encryption zone key EEK caches when an encryption zone is created.

KMS ACL Evaluation Flow

Before diving into the details of how KMS ACLs are evaluated, it is critical that you understand the key rules that the KMS uses in performing this evaluation.

KMS ACL Flow Rules:

- The whitelist class bypasses `key.acl` and `default.key.acl` controls.
- The `key.acl` definitions override all default definitions.

The better you understand these rules, the more likely it is that you will be successful creating and managing KMS ACL sets to achieve your desired goals.

Encryption key access is evaluated as follows:

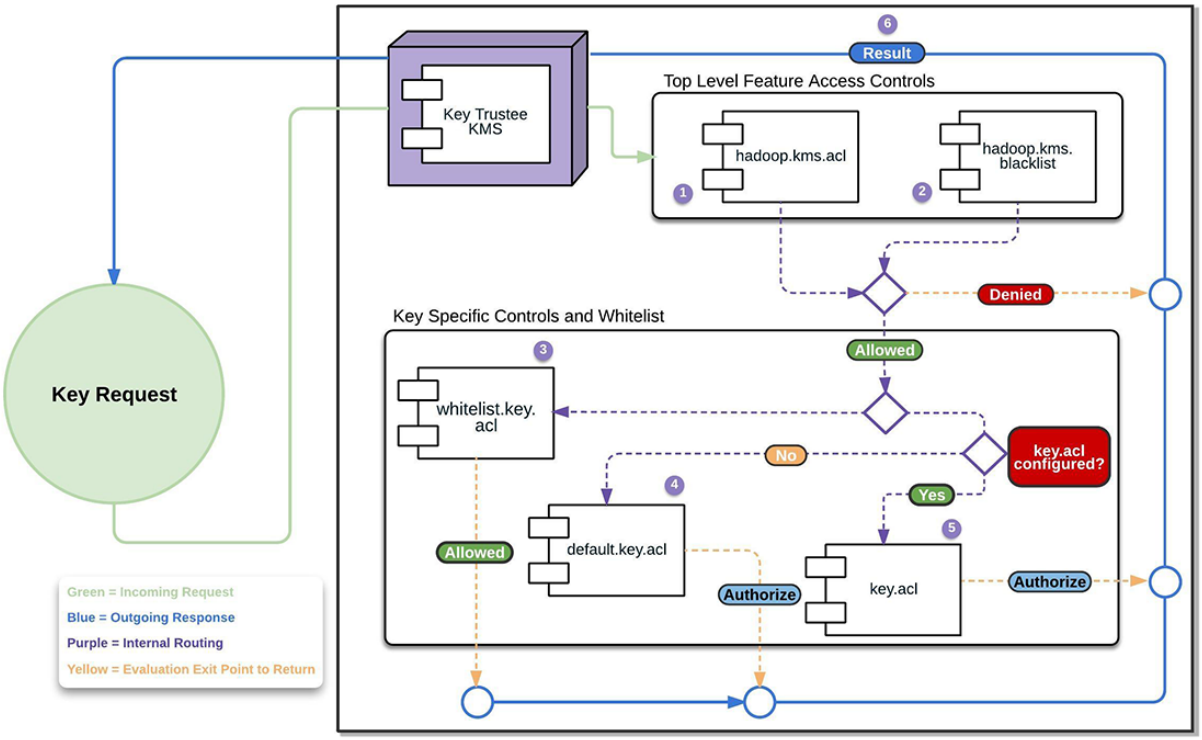

The KMS evaluates `hadoop.kms.acl.<OPERATION>` and `hadoop.kms.blacklist.<OPERATION>` classes to determine whether or not access to a specific KMS feature or function is authorized.

In other words, a user must be allowed by `hadoop.kms.acl.<OPERATION>`, and not be disallowed by `hadoop.kms.blacklist.<OPERATION>`.

If a user is blacklisted or otherwise not allowed access to a KMS-wide operation, then the flow halts and returns the result "Denied".

If a user is allowed access to a KMS-wide operation, then the evaluation flow proceeds.

3

The KMS evaluates `whitelist.key.acl` class.

The KMS ACL workflow evaluates the `whitelist.key.acl.<OPERATION>`, and if the user is allowed access, then it is granted ("Allowed") . If not, then the flow continues with the evaluation.

4 5

The KMS evaluates `default.key.acl.<OPERATION>`, and `key.acl.<OPERATION>` classes.

The KMS evaluates whether or not there is a `key.acl.KEY.<OPERATION>` that matches the action the user is attempting to perform. If there is, it then evaluates that value to determine whether or not the user can perform the requested operation.

> **Note:** Before evaluating the `default.key.acl.<OPERATION>` and `key.acl.<OPERATION>`, the flow logic determines which classes exist. Only one of these can exist and be used at any time (for example, `key.acl.prodkey.READ` would override `default.key.acl.READ` for `prodkey`, so it will be configured with its own READ ACLs).

6

Depending on the result of the KMS ACL evaluation, controls are applied to the key and results (Allowed or Denied).

KMS ACL Syntax and Tips

Blacklist and Whitelist Syntax

The ACL syntax for both blacklist and whitelist entries is as follows:

- Users only: `user1,user2,userN`

  There are no spaces following the commas separating the users in the list.

- Groups only: `nobody group1,group2,groupN`

  There is a space between nobody and the comma-separated group list. The nobody user, if it exists, must not have privileges to log in to or interact with the system. If you are uncertain about its access privileges, specify a different nonexistent user in its place.

- Users and Groups: `user1,user2,userN group1,group2,groupN`

  The comma-separated user list is separated from the comma-separated group list by a space.

Blocking User Access

If you wish to block access to an operation entirely, use the value of an empty space, or some non-existent values (for example,'NOUSERS NOGROUPS'). By doing this, you ensure that no user maps to a particular operation by default. Alternatively, you can restrict those features to Key Administrators only by setting the value to Keyadmin users and/or groups.

Group Membership in KMS ACLs

The group membership used by ACL entries depends upon the configured group mapping mechanism for HDFS. By default, group membership is determined on the local Linux system running the KMS service. If you have configured HDFS to use LDAP for group mapping, then group membership for the ACL entries is determined using the configured LDAP settings. For more information about LDAP-based group membership, see Configuring LDAP Group Mappings on page 179.

Configuring KMS ACLs Using Cloudera Manager

Minimum Required Role: Key Administrator (also provided by **Full Administrator**)

The KMS installation wizard includes an option to generate the recommended ACLs. To view or edit the ACLs:

1. Go to the KMS service.
2. Click **Configuration**.
3. In the Search field, type **acl** to show the **Key Management Server Advanced Configuration Snippet (Safety Valve) for kms-acls.xml** (in the **Key Management Server Default Group** category).
4. Add or edit the ACL properties according to your cluster configuration. See "Recommended KMS Access Control List" for example ACL entries.
5. Click **Save Changes**.
6. Return to the Home page by clicking the Cloudera Manager logo, and return to the KMS service.
7. Click (**Refresh Needed**) . Then click **Refresh Cluster**.
8. Click **Finish**.

Configuring KMS ACLs Using the Command Line

KMS ACLs are defined in the `/etc/hadoop-kms/conf/kms-acls.xml` configuration file. This file is hot-reloaded when it changes. See Recommended KMS ACL on page 283 for recommended ACL entries.

Recommended KMS ACL

Cloudera recommends the following ACL definition for secure production settings. Replace `keyadmin` and `keyadmingroup` with the user and group responsible for maintaining encryption keys.

> **Note:** If you are entering the ACL using Cloudera Manager, omit the surrounding `<configuration>` and `</configuration>` tags; Cloudera Manager adds this automatically.

```
<configuration>

<!--
KMS ACLs control which users can perform various actions on the KMS,and which
users and groups have access to which keys.

This file includes the following sections:
  * ACLs for KMS operations
  ** Access to specific KMS operations
  ** Blacklists for specific operations
  * ACLs for keys
  ** Default ACLs for keys
  ** Whitelist ACLs for keys
  ** Key-specific ACLs
-->

<!--
KMS ACLs that govern access to specific key operations. If access is not granted
for an operation here, then the operation is forbidden, even if a key ACL allows it.

The ACL value should be either a username or a username and groupname separated
by whitespace.

A value of "*" (for the username or groupname) indicates that all users are
granted access to that operation. Any operation for which there is no ACL or an
empty (zero-length) ACL is treated as having an ACL with a value of "*".
To disallow all users, add an ACL with a value of " " (a single space).
```

```
Note: This convention applies only to the KMS-wide ACLs beginning with 'hadoop.kms.acl'.
-->
<property>
  <name>hadoop.kms.acl.CREATE</name>
  <value>keyadmin keyadmingroup</value>
  <description>
    ACL for create-key operations.
    If the user is not in the GET ACL, the key material is not returned
    as part of the response.
  </description>
</property>

<property>
  <name>hadoop.kms.acl.DELETE</name>
  <value>keyadmin keyadmingroup</value>
  <description>
   ACL for delete-key operations.
  </description>
</property>

<property>
  <name>hadoop.kms.acl.ROLLOVER</name>
  <value>keyadmin keyadmingroup</value>
  <description>
    ACL for rollover-key operations.
    If the user does is not in the GET ACL, the key material is not returned as part of
 the response.
  </description>
</property>

<property>
  <name>hadoop.kms.acl.GET</name>
  <value></value>
  <description>
    ACL for get-key-version and get-current-key operations.
  </description>
</property>

<property>
  <name>hadoop.kms.acl.GET_KEYS</name>
  <value>keyadmin keyadmingroup</value>
  <description>
    ACL for get-keys operations.
  </description>
</property>

<property>
  <name>hadoop.kms.acl.SET_KEY_MATERIAL</name>
  <value></value>
  <description>
    Complementary ACL for CREATE and ROLLOVER operations to allow the client to provide
 the
    key material when creating or rolling a key.
  </description>
</property>

<property>
  <name>hadoop.kms.acl.GENERATE_EEK</name>
  <value>hdfs supergroup</value>
  <description>
    ACL for generateEncryptedKey CryptoExtension operations.
  </description>
</property>

<!--
  KMS blacklists to prevent access to operations. These settings override the permissions
 granted by the KMS ACLs listed above.

  The blacklist value should be either a username or a username and groupname separated
 by whitespace.

  A blank value indicates that no user is blacklisted from the operation. A value of
```

```
"*" (for either the username or
  groupname) indicates that all users are blacklisted from the operation. Any operation
 for which there is no blacklist
  will be treated as having a blacklist with an empty value.
-->

<!--
  In this template the hdfs user is blacklisted for everything except GET_METADATA,
GET_KEYS, and GENERATE_EEK. The
  GET and SET_KEY_MATERIAL operations are blacklisted for all users because Hadoop users
 should not need to perform
  those operations, and access to the key material should be as restricted as possible.
-->

<property>
  <name>hadoop.kms.blacklist.CREATE</name>
  <value>hdfs supergroup</value>
</property>

<property>
  <name>hadoop.kms.blacklist.DELETE</name>
  <value>hdfs supergroup</value>
</property>

<property>
  <name>hadoop.kms.blacklist.ROLLOVER</name>
  <value>hdfs supergroup</value>
</property>

<property>
  <name>hadoop.kms.blacklist.GET</name>
  <value>*</value>
</property>

<property>
  <name>hadoop.kms.blacklist.GET_KEYS</name>
  <value></value>
</property>

<property>
  <name>hadoop.kms.blacklist.SET_KEY_MATERIAL</name>
  <value>*</value>
</property>

<property>
  <name>hadoop.kms.blacklist.DECRYPT_EEK</name>
  <value>hdfs supergroup</value>
</property>

<property>
  <name>keytrustee.kms.acl.UNDELETE</name>
  <value></value>
  <description>
    ACL that grants access to the UNDELETE operation on all keys.
    Only used by Key Trustee KMS.
  </description>
</property>

<property>
  <name>keytrustee.kms.acl.PURGE</name>
  <value></value>
  <description>
    ACL that grants access to the PURGE operation on all keys.
    Only used by Key Trustee KMS.
  </description>
</property>

<!--
  Default key ACLs that govern access to key operations for key-operation pairs that do
 not have a
  key-specific ACL already. Key-specific ACLs override the default key ACLs.

  The ACL value should be either a username or a username and group name separated by
```

```
whitespace.

  An empty value for an ACL indicates that no user is granted access to that operation.
 A value
  of "*" (for the username or groupname) indicates that all users are granted access to
 that operation.
  Any operation for which there is no ACL will be treated as having an ACL with an empty
 value.
-->

<property>
  <name>default.key.acl.MANAGEMENT</name>
  <value></value>
  <description>
    Default ACL that grants access to the MANAGEMENT operation on all keys.
  </description>
</property>

<property>
  <name>default.key.acl.GENERATE_EEK</name>
  <value></value>
  <description>
    Default ACL that grants access to the GENERATE_EEK operation on all keys.
  </description>
</property>

<property>
  <name>default.key.acl.DECRYPT_EEK</name>
  <value></value>
  <description>
    Default ACL that grants access to the DECRYPT_EEK operation on all keys.
  </description>
</property>

<property>
  <name>default.key.acl.READ</name>
  <value></value>
  <description>
    Default ACL that grants access to the READ operation on all keys.
  </description>
</property>

<!--
  Whitelist key ACLs that grant access to key-specific operations. Any permissions
granted here
  will be added to whatever permissions are granted by the specific key ACL or the
default key ACL.
  Note that these whitelist ACLs grant access to operations on specific keys. If the
operations
  are not allowed because of the KMS ACLs/blacklists, then the they will not be permitted,

  regardless of the whitelist settings.

  The ACL value should be either a username or a username and group name separated by
whitespace.

  An empty value for an ACL indicates that no user is granted access to that operation.
 A value
  of "*" (for the username or groupname) indicates that all users are granted access to
 that
  operation. Any operation for which there is no ACL will
  be treated as having an ACL with an empty value.
-->

<property>
  <name>whitelist.key.acl.MANAGEMENT</name>
  <value>keyadmin keyadmingroup</value>
  <description>
    Whitelist ACL for MANAGEMENT operations for all keys.
  </description>
</property>

<property>
```

```
  <name>whitelist.key.acl.READ</name>
  <value>hdfs supergroup</value>
  <description>
    Whitelist ACL for READ operations for all keys.
  </description>
</property>

<property>
  <name>whitelist.key.acl.GENERATE_EEK</name>
  <value>hdfs supergroup</value>
  <description>
    Whitelist ACL for GENERATE_EEK operations for all keys.
  </description>
</property>

<property>
  <name>whitelist.key.acl.DECRYPT_EEK</name>
  <value>keyadmin keyadmingroup</value>
  <description>
    Whitelist ACL for DECRYPT_EEK operations for all keys.
  </description>
</property>

<!--
  Key ACLs that grant access to specific key operations. Any permissions granted here
are added
  to whatever permissions are granted by the whitelists.
  The key ACL name should be key.acl.<keyname>.<OPERATION>.

  The ACL value should be either a username or a username and group name separated by
whitespace.

  An empty value for an ACL indicates that no user is granted access to that operation.
 A value
  of "*" (for the username or group name) indicates that all users are granted access
to that operation.
  Any key operation for which there is no ACL will default to the default ACL for the
operation.

  Normally adding users or groups for a specific key and DECRYPT_EEK is sufficient to
allow access
  to data protected with HDFS data at rest encryption.
-->

<!--
  The following ACLs are required for proper functioning of services. Cloudera Manager
 does not create keys or
  encryption zones; however, our best practices recommend encryption zones on certain
directories.
  An assumption in these ACLs is that the user has followed the recommended naming scheme
 and named the keys
  according to documented best practices: "hive-key" for the Hive service,
  "hbase-key" for the Hbase service, etc. If the key names are different, then none of
 this will work
  out of the box, and you will need to edit these ACLs to match your key names.
-->

<property>
  <name>key.acl.hive-key.DECRYPT_EEK</name>
  <value>hive hive</value>
  <description>
    Gives the Hive user and the Hive group access to the key named "hive-key".
    This allows the Hive service to read and write files in /user/hive/.
    Also note that the Impala user ought to be a member of the Hive group to enjoy this
 same access.
  </description>
</property>

<property>
  <name>key.acl.hive-key.READ</name>
  <value>hive hive</value>
  <description>
    Required because Hive compares key strengths when joining tables.
```

```
  </description>
</property>

<property>
  <name>key.acl.hbase-key.DECRYPT_EEK</name>
  <value>hbase hbase</value>
  <description>
    Gives the hbase user and hbase group access to the key named "hbase-key".
    This allows the hbase service to read and write files in /hbase.
  </description>
</property>

<property>
  <name>key.acl.solr-key.DECRYPT_EEK</name>
  <value>solr solr</value>
  <description>
    Gives the solr user and solr group access to the key named "solr-key".
    This allows the solr service to read and write files in /solr.
  </description>
</property>

<property>
  <name>key.acl.mapred-key.DECRYPT_EEK</name>
  <value>mapred,yarn hadoop</value>
  <description>
    Gives the mapred user and mapred group access to the key named "mapred-key".
    This allows mapreduce to read and write files in /user/history.
    This is required by YARN.
  </description>
</property>

<property>
  <name>key.acl.hue-key.DECRYPT_EEK</name>
  <value>oozie,hue oozie,hue</value>
  <description>
    Gives the appropriate users and groups access to the key named "hue-key".
    This allows Hue and Oozie to read and write files in /user/hue.
    Oozie is required here because it will attempt to access workflows in
/user/hue/oozie/workspaces.
  </description>
</property>



<!-- This example is required if there are encryption zones on user’s home directories.
 -->
<!--
<property>
  <name>key.acl.username-key.DECRYPT_EEK</name>
  <value>username username,hive,hbase,solr,oozie,hue,yarn</value>
  <description>
    Designed to be placed on a key that protects the EZ /user/username, and assumes that

    the key name is also "username-key"; this shows that a number of services can reach
 in
    to access data. Remove those that are not needed for your use case.
  </description>
</property>
-->

</configuration>
```

Migrating from a Key Trustee KMS to an HSM KMS

In cases where your enterprise requires the highest level of encryption zone key (EZ key) security, or, it must adhere to strict security certification compliance standards, you have the option of migrating from a Key Trustee KMS (KT KMS) to a Hardware Security Module KMS (HSM KMS). Migrating from KT KMS to HSM KMS essentially means that you are moving EZ key metadata from one KMS to another.

You should be aware of the following important points associated with the migration of key metadata from a KT KMS to an HSM KMS:

- Only EZ key metadata (name, cipher, length, description, created, and versions) is migrated. Encryption zone key material is *not* migrated.
- EZ key versions on the destination KMS are distinct and disjoint from the EZ key versions on the source KMS.

  The key isolation standards of the HSM KMS require that keys originate on and never leave the HSM. However, importing key material (which includes the corresponding unique identifier and key version) would violate this requirement. If you want to verify metadata after migrating to an HSM KMS, note that the key versions will be different than before. With the Java Keystore KMS and KT KMS, this is not the case.

Before attempting a KT KMS to HSM KMS migration, you should be familiar with basic concepts about EZ keys and the HSM KMS. For more details, see:

- Managing Encryption Keys and Zones on page 257
- Cloudera Navigator HSM KMS Overview

Requirements and Limitations

The following requirements and limitations apply to the KT KMS to HSM KMS migration of EZ key metadata:

- The KT KMS and HSM KMS must be on separate hosts. Running each KMS on a distinct node ensures that there will be no conflicts, especially when both are running at the same time.
- Key Trustee Server must be configured with Key HSM using either the SafeNet Luna or Thales HSM.
- The KT KMS must remain active until all encrypted data encryption keys (EDEKs) are re-encrypted using an EZ key version from the destination KMS. This is necessary so that you can continue to use EDEKs encrypted with an EZ key version from the source KMS.
- It is recommended that the HSM KMS is installed, but not running as the default KMS for HDFS during the actual migration.
- The latency between the physical HSM and HSM KMS must be low. It is recommended that the measurable latency be no higher than 30 milliseconds. Note that this latency can result in timeout conditions, as the EDEK must be sent to the HSM to be encrypted and/or decrypted.

Migrating from a KT KMS to an HSM KMS

1. Copy the configuration file (`kts-site.xml`) from the KT KMS node to all HSM KMS nodes and, if not already the owner, change the owner to 'kms':

```
root@kms-1.example.com> scp
/var/run/cloudera-scm-agent/process/##--keytrustee-KMS_KEYTRUSTEE/kts-site.xml
root@hsm-1.example.com:/tmp/kts-site.xml
# Where ## is the latest number that is present for keytrustee-KMS_KEYTRUSTEE in the
process folder.
root@hsm-1.example.com> chown kms:kms /tmp/kts-site.xml
```

   The `kts-site` configuration file includes required connection information about the Key Trustee servers to which you will need to connect.

2. Copy the `/var/lib/kms-keytrustee` directory from the KT KMS node to all HSM KMS nodes:

```
root@kms-1.example.com> scp -r /var/lib/kms-keytrustee
root@hsm-1.example.com:/var/lib/kms-keytrustee/
root@hsm-1.example.com> chown -R kms:kms /var/lib/kms-keytrustee
```

   If you plan to re-encrypt with new keys, then you will require a way to decrypt the old keys first, before you can re-encrypt them. Copying this directory ensures that you will have all the GPG keys necessary to contact, read, and decrypt the EZ keys from the KT KMS.

3. *If TLS is enabled*, then import the Key Trustee Server certificates into the HSM KMS truststore on both nodes:

```
hsm-1.example.com> echo -n | openssl s_client -connect kts-1.example.com>11371 | sed
-ne '/-BEGIN CERTIFICATE-/,/-END CERTIFICATE-/p' > /tmp/keytrustee_certificate.pem
# Use the default Oracle Java keytool to import the certificate:
/usr/java/jdk1.7.0_67/bin/keytool -importcert -keystore
/path_to_truststore_file/truststore.jks -file /tmp/keytrustee_certificate.pem -alias
kts1
# If Key Trustee Server HA is configured, then import certificates from both nodes.
```

4. Set the following Advanced Configuration Snippets for `kms-site.xml` (also known as "safety valves") on *all* HSM KMS nodes, unless otherwise specified. Refer to Custom Configuration for details about safety valves.

   1. Configure the re-encryption source URI file. The configuration information provided here instructs the HSM to check the KT KMS for a key version (if it cannot be found in the HSM KMS).

      Note: If TLS is enabled, use HTTPS. Otherwise, use HTTP.

```
<property>
<name>hadoop.kms.key.provider.reencryption.source.uri</name>
<value>kms://https@kms-1.example.com:16000/kms</value>
</property>

<property>
<name>hadoop.kms.key.provider.reencryption.source.uri</name>
<value>kms://http@kms-1.example.com:16000/kms</value>
</property>
```

      This new HSM KMS URI includes the old KT KMS configuration as a secondary source. If the HSM KMS does not have that key version, then it reaches out to the KT KMS to try and locate it there. In cases where you have several KMS nodes, specify them using a semi-colon (`;`).

   2. Enable re-encryption:

```
<property>
<name>hadoop.kms.reencryption.crypto.enable</name>
<value>true</value>
</property>
```

   3. Configure the location of the `kts-site.xml` file:

```
<property>
<name>cloudera.hsmkp.ktkms.config.dir</name>
<value>/path/to/kts-site.xml</value>
</property>
```

   4. Configure HSM-specific settings for the `hadoop.kms.key.provider.uri`:

```
For Luna, use the luna.keystore:
<property>
<name>hadoop.kms.key.provider.uri</name>
<value>cloudera.hsmkp://file@/${user.home}/luna.keystore,keytrustee://file@/var/lib/kms-keytrustee/keytrustee/.keytrustee</value>
</property>
```

```
For Thales, use the ncipher.keystore:
<property>
<name>hadoop.kms.key.provider.uri</name>
<value>cloudera.hsmkp://file@/${user.home}/ncipher.keystore,keytrustee://file@/var/lib/kms-keytrustee/keytrustee/.keytrustee</value>
</property>
```

5. Configure the following safety valve in `kms-site.xml` for a single Navigator HSM KMS Proxy role (for details, refer to Overriding Configuration Properties):

```
<property>
<name>cloudera.hsmkp.initiatemigration</name>
<value>true</value>
</property>
```

6. (*Optional*) If the HSM KMS is operating in a topology where there is significant latency between the HSM KMS and HSM (for example, an HSM KMS in the cloud with an HSM on premise), then you can configure cache sizes to prevent timeouts (the values here are suggestions; you may need to fine tune them for your specific topology):

```
In hdfs-site.xml on HDFS:
<property>
<name>hadoop.security.kms.client.encrypted.key.cache.low-watermark</name>
<value>.05</value>
</property>
<property>
<name>hadoop.security.kms.client.encrypted.key.cache.size</name>
<value>40</value>
</property>
```

```
In kms-site.xml on all HSM KMS nodes:
<property>
<name>hadoop.security.kms.encrypted.key.cache.low-watermark</name>
<value>.05</value>
</property>
<property>
<name>hadoop.security.kms.encrypted.key.cache.size</name>
<value>40</value>
</property>
```

7. *If Kerberos is enabled*, in the HSM KMS service configuration, set **Zookeeper Authentication Type for Secret Signer** to `sasl`.
8. In the HDFS configuration, replace `KT KMS` with `HSM KMS`.
9. Start the HSM KMS, and then restart HDFS. Upon startup, key migration automatically occurs.
10. Post migration, it is recommended that you set `cloudera.hsmkp.initiatemigration` to `false` after key migration completes (set in Step 5), and then restart the HSM KMS.

    It is highly recommended that you re-encrypt encryption zones after migration is complete. For details about re-encrypting EDEKs, refer to Re-encrypting Encrypted Data Encryption Keys (EDEKs) on page 261.

11. Before operating the new HSM KMS, it is recommended that you restore the properties (`hadoop.kms.key.provider.uri`), and also remove and clean up the safety valves that you set previously, as they are no longer necessary:
    - `cloudera.hsmkp.ktkms.config.dir`
    - `hadoop.kms.key.provider.reencryption.source.uri`
    - `hadoop.kms.reencryption.crypto.enable`
12. (For Luna HSM) To restore `hadoop.kms.key.provider.uri`: `cloudera.hsmkp://file@/${user.home}/luna.keystore`
13. (For Thales HSM) To restore `hadoop.kms.key.provider.uri`: `cloudera.hsmkp://file@/${user.home}/ncipher.keystore`

Migrating Keys from a Java KeyStore to Cloudera Navigator Key Trustee Server

You can migrate keys from an existing Java KeyStore (JKS) to Key Trustee Server to improve security, durability, and scalability. If you are using the Java KeyStore KMS service, and want to use Key Trustee Server as the backing key store for HDFS Transparent Encryption on page 241, use the following procedure.

This procedure assumes that the Java KeyStore (JKS) is on the same host as the new Key Trustee KMS service.

1. Stop the Java KeyStore KMS service.
2. Add and configure the Key Trustee KMS service, and configure HDFS to use it for its **KMS Service** setting. For more information about how to install Key Trustee KMS, see Installing Key Trustee KMS.
3. Restart the HDFS service and redeploy client configuration for this to take effect:
   a. **Home** > **Cluster-wide** > **Deploy Client Configuration**
4. Add the following to the **Key Management Server Proxy Advanced Configuration Snippet (Safety Valve) for kms-site.xml** (**Key Trustee KMS Service** > **Configuration** > **Category** > **Advanced**):

```
<property>
  <name>hadoop.kms.key.provider.uri</name>

<value>keytrustee://file@/var/lib/kms-keytrustee/keytrustee/.keytrustee/,jceks://file@/path/to/kms.keystore</value>

  <description>URI of the backing KeyProvider for the KMS</description>
</property>

<property>
  <name>hadoop.security.keystore.java-keystore-provider.password-file</name>
  <value>/tmp/password.txt</value>
  <description>Java KeyStore password file</description>
</property>
```

   If the Java KeyStore is *not* password protected, omit the `hadoop.security.keystore.java-keystore-provider.password-file` property.

5. Click **Save Changes** and restart the Key Trustee KMS service. If the Java KeyStore is *not* password protected, skip to step 7.
6. Create the file `/var/lib/keytrustee-kms/tomcat-deployment/webapps/kms/WEB-INF/classes/tmp/password.txt` and add the Java KeyStore password to it.
7. Change the ownership of `/var/lib/keytrustee-kms/tomcat-deployment/webapps/kms/WEB-INF/classes/tmp/password.txt` to `kms:kms`:

```
$ sudo chown kms:kms
/var/lib/keytrustee-kms/tomcat-deployment/webapps/kms/WEB-INF/classes/tmp/password.txt
```

8. From the host running the Key Trustee KMS service, if you have not configured Kerberos and TLS/SSL, run the following command:

```
$ curl -L -d "trusteeOp=migrate"
"http://kms01.example.com:16000/kms/v1/trustee/key/migrate?user.name=username&trusteeOp=migrate"
```

   If you have configured Kerberos and TLS/SSL, use the following command instead:

```
$ curl --negotiate -u : -L -d "trusteeOp=migrate"
"https://kms01.example.com:16000/kms/v1/trustee/key/migrate?user.name=username&trusteeOp=migrate"
 --cacert /path/to/kms/cert
```

9. Monitor `/var/log/kms-keytrustee/kms.log` and `/var/log/kms-keytrustee/kms-catalina.<date>.log` to verify that the migration is successful. You can also run `sudo -u <key_admin> hadoop key list` to verify that the keys are listed.
10. After you have verified that the migration is successful, remove the safety valve entry used in step 3 and restart the Key Trustee KMS service.

Configuring CDH Services for HDFS Encryption

This page contains recommendations for setting up HDFS Transparent Encryption on page 241 with various CDH services.

> **Important:** HDFS encryption does not support file transfer (reading, writing files) between zones through WebHDFS. For web-based file transfer between encryption zones managed by HDFS, use HttpFS with a load balancer instead.

> **Important:** Encrypting `/tmp` using HDFS encryption is *not* supported.

HBase

Recommendations

Make `/hbase` an encryption zone. Do not create encryption zones as subdirectories under `/hbase`, because HBase may need to rename files across those subdirectories. When you create the encryption zone, name the key `hbase-key` to take advantage of auto-generatedKMS ACLs .

Steps

On a cluster without HBase currently installed, create the `/hbase` directory and make that an encryption zone.

On a cluster with HBase already installed, perform the following steps:

1. Stop the HBase service.
2. Move data from the `/hbase` directory to `/hbase-tmp`.
3. Create an empty `/hbase` directory and make it an encryption zone.
4. Distcp all data from `/hbase-tmp` to `/hbase`, preserving user-group permissions and extended attributes.
5. Start the HBase service and verify that it is working as expected.
6. Remove the `/hbase-tmp` directory.

KMS ACL Configuration for HBase

In theKMS ACL, grant the `hbase` user and group `DECRYPT_EEK` permission for the HBase key:

```
<property>
  <name>key.acl.hbase-key.DECRYPT_EEK</name>
  <value>hbase hbase</value>
  </description>
</property>
```

Hive

HDFS encryption has been designed so that files cannot be moved from one encryption zone to another or from encryption zones to unencrypted directories. Therefore, the landing zone for data when using the `LOAD DATA INPATH` command must always be inside the destination encryption zone.

To use HDFS encryption with Hive, ensure you are using *one* of the following configurations:

Single Encryption Zone

With this configuration, you can use HDFS encryption by having all Hive data inside the same encryption zone. In Cloudera Manager, configure the Hive Scratch Directory (`hive.exec.scratchdir`) to be inside the encryption zone.

Recommended HDFS Path: `/user/hive`

To use the auto-generated KMS ACL, make sure you name the encryption key `hive-key`.

For example, to configure a single encryption zone for the entire Hive warehouse, you can rename `/user/hive` to `/user/hive-old`, create an encryption zone at `/user/hive`, and then `distcp` all the data from `/user/hive-old` to `/user/hive`.

In Cloudera Manager, configure the Hive Scratch Directory (`hive.exec.scratchdir`) to be inside the encryption zone by setting it to `/user/hive/tmp`, ensuring that permissions are `1777` on `/user/hive/tmp`.

Multiple Encryption Zones

With this configuration, you can use encrypted databases or tables with different encryption keys. To read data from read-only encrypted tables, users must have access to a temporary directory that is encrypted at least as strongly as the table.

For example:

1. Configure two encrypted tables, `ezTbl1` and `ezTbl2`.
2. Create two new encryption zones, `/data/ezTbl1` and `/data/ezTbl2`.
3. Load data to the tables in Hive using `LOAD` statements.

For more information, see Changed Behavior after HDFS Encryption is Enabled on page 294.

Other Encrypted Directories

- **`LOCALSCRATCHDIR`**: The MapJoin optimization in Hive writes HDFS tables to a local directory and then uploads them to the distributed cache. To ensure these files are encrypted, either disable MapJoin by setting `hive.auto.convert.join` to `false`, or encrypt the *local* Hive Scratch directory (`hive.exec.local.scratchdir`) using Cloudera Navigator Encrypt.
- **`DOWNLOADED_RESOURCES_DIR`**: JARs that are added to a user session and stored in HDFS are downloaded to `hive.downloaded.resources.dir` on the HiveServer2 local filesystem. To encrypt these JAR files, configure Cloudera Navigator Encrypt to encrypt the directory specified by `hive.downloaded.resources.dir`.
- **NodeManager Local Directory List:** Hive stores JARs and MapJoin files in the distributed cache. To use MapJoin or encrypt JARs and other resource files, the `yarn.nodemanager.local-dirs` YARN configuration property must be configured to a set of encrypted local directories on all nodes.

Changed Behavior after HDFS Encryption is Enabled

- Loading data from one encryption zone to another results in a copy of the data. Distcp is used to speed up the process if the size of the files being copied is higher than the value specified by `HIVE_EXEC_COPYFILE_MAXSIZE`. The minimum size limit for `HIVE_EXEC_COPYFILE_MAXSIZE` is 32 MB, which you can modify by changing the value for the `hive.exec.copyfile.maxsize` configuration property.
- When loading data to encrypted tables, Cloudera strongly recommends using a landing zone inside the same encryption zone as the table.
  - **Example 1: Loading unencrypted data to an encrypted table** - Use one of the following methods:
    - If you are loading new unencrypted data to an encrypted table, use the `LOAD DATA ...` statement. Because the source data is not inside the encryption zone, the `LOAD` statement results in a copy. For this reason, Cloudera recommends landing data that you need to encrypt inside the destination encryption zone. You can use `distcp` to speed up the copying process if your data is inside HDFS.
    - If the data to be loaded is already inside a Hive table, you can create a new table with a `LOCATION` inside an encryption zone as follows:

```
CREATE TABLE encrypted_table [STORED AS] LOCATION ... AS SELECT * FROM <unencrypted_table>
```

The location specified in the `CREATE TABLE` statement must be inside an encryption zone. Creating a table pointing `LOCATION` to an unencrypted directory does not encrypt your source data. You must copy your data to an encryption zone, and then point `LOCATION` to that zone.

– **Example 2: Loading encrypted data to an encrypted table** - If the data is already encrypted, use the `CREATE TABLE` statement pointing `LOCATION` to the encrypted source directory containing the data. This is the fastest way to create encrypted tables.

```
CREATE TABLE encrypted_table [STORED AS] LOCATION ... AS SELECT * FROM
<encrypted_source_directory>
```

- Users reading data from encrypted tables that are read-only must have access to a temporary directory which is encrypted with at least as strong encryption as the table.
- Temporary data is now written to a directory named `.hive-staging` in each table or partition
- Previously, an `INSERT OVERWRITE` on a partitioned table inherited permissions for new data from the existing partition directory. With encryption enabled, permissions are inherited from the table.

KMS ACL Configuration for Hive

When Hive joins tables, it compares the encryption key strength for each table. For this operation to succeed, you must configure the KMS ACL to allow the `hive` user and group `READ` access to the Hive key:

```
<property>
  <name>key.acl.hive-key.READ</name>
  <value>hive hive</value>
</property>
```

If you have restricted access to the `GET_METADATA` operation, you must grant permission for it to the `hive` user or group:

```
<property>
  <name>hadoop.kms.acl.GET_METADATA</name>
  <value>hive hive</value>
</property>
```

If you have disabled HiveServer2 Impersonation on page 130 (for example, to use Apache Sentry), you must configure the KMS ACLs to grant `DECRYPT_EEK` permissions to the `hive` user, as well as any user accessing data in the Hive warehouse.

Cloudera recommends creating a group containing all Hive users, and granting `DECRYPT_EEK` access to that group.

For example, suppose user `jdoe` (home directory `/user/jdoe`) is a Hive user and a member of the group `hive-users`. The encryption zone (EZ) key for `/user/jdoe` is named `jdoe-key`, and the EZ key for `/user/hive` is `hive-key`. The following ACL example demonstrates the required permissions:

```
<property>
  <name>key.acl.hive-key.DECRYPT_EEK</name>
  <value>hive hive-users</value>
</property>

<property>
  <name>key.acl.jdoe-key.DECRYPT_EEK</name>
  <value>jdoe,hive</value>
</property>
```

If you have enabled HiveServer2 impersonation, data is accessed by the user submitting the query or job, and the user account (`jdoe` in this example) may still need to access data in their home directory. In this scenario, the required permissions are as follows:

```
<property>
  <name>key.acl.hive-key.DECRYPT_EEK</name>
  <value>nobody hive-users</value>
</property>

<property>
  <name>key.acl.jdoe-key.DECRYPT_EEK</name>
```

```
  <value>jdoe</value>
</property>
```

Hue

Recommendations

Make `/user/hue` an encryption zone because Oozie workflows and other Hue-specific data are stored there by default. When you create the encryption zone, name the key `hue-key` to take advantage of auto-generated KMS ACLs .

Steps

On a cluster without Hue currently installed, create the `/user/hue` directory and make it an encryption zone.

On a cluster with Hue already installed:

1. Create an empty `/user/hue-tmp` directory.
2. Make `/user/hue-tmp` an encryption zone.
3. DistCp all data from `/user/hue` into `/user/hue-tmp`.
4. Remove `/user/hue` and rename `/user/hue-tmp` to `/user/hue`.

KMS ACL Configuration for Hue

In the KMS ACLs, grant the hue and `oozie` users and groups `DECRYPT_EEK` permission for the Hue key:

```
<property>
  <name>key.acl.hue-key.DECRYPT_EEK</name>
  <value>oozie,hue oozie,hue</value>
</property>
```

Impala

Recommendations

- If HDFS encryption is enabled, configure Impala to encrypt data spilled to local disk.
- In releases lower than Impala 2.2.0 / CDH 5.4.0, Impala does not support the `LOAD DATA` statement when the source and destination are in different encryption zones. If you are running an affected release and need to use `LOAD DATA` with HDFS encryption enabled, copy the data to the table's encryption zone prior to running the statement.
- Use Cloudera Navigator to lock down the local directory where Impala UDFs are copied during execution. By default, Impala copies UDFs into `/tmp`, and you can configure this location through the `--local_library_dir` startup flag for the `impalad` daemon.
- Limit the rename operations for internal tables once encryption zones are set up. Impala cannot do an `ALTER TABLE RENAME` operation to move an internal table from one database to another, if the root directories for those databases are in different encryption zones. If the encryption zone covers a table directory but not the parent directory associated with the database, Impala cannot do an `ALTER TABLE RENAME` operation to rename an internal table, even within the same database.
- Avoid structuring partitioned tables where different partitions reside in different encryption zones, or where any partitions reside in an encryption zone that is different from the root directory for the table. Impala cannot do an `INSERT` operation into any partition that is not in the same encryption zone as the root directory of the overall table.
- If the data files for a table or partition are in a different encryption zone than the HDFS trashcan, use the `PURGE` keyword at the end of the `DROP TABLE` or `ALTER TABLE DROP PARTITION` statement to delete the HDFS data

files immediately. Otherwise, the data files are left behind if they cannot be moved to the trashcan because of differing encryption zones. This syntax is available in Impala 2.3 / CDH 5.5 and higher.

Steps

Start every `impalad` process with the `--disk_spill_encryption=true` flag set. This encrypts all spilled data using AES-256-CFB. Set this flag by selecting the **Disk Spill Encryption** checkbox in the Impala configuration (**Impala service** > **Configuration** > **Category** > **Security**).

> **Important:** Impala does not selectively encrypt data based on whether the source data is already encrypted in HDFS. This results in at most 15 percent performance degradation when data is spilled.

KMS ACL Configuration for Impala

Cloudera recommends making the `impala` user a member of the `hive` group, and following the ACL recommendations in KMS ACL Configuration for Hive on page 295.

MapReduce and YARN

MapReduce v1

Recommendations

MRv1 stores both history and logs on local disks by default. Even if you do configure history to be stored on HDFS, the files are not renamed. Hence, no special configuration is required.

MapReduce v2 (YARN)

Recommendations

Make `/user/history` a single encryption zone, because history files are moved between the `intermediate` and done directories, and HDFS encryption does not allow moving encrypted files across encryption zones. When you create the encryption zone, name the key `mapred-key` to take advantage of auto-generated KMS ACLs.

Steps

On a cluster with MRv2 (YARN) installed, create the `/user/history` directory and make that an encryption zone.

If `/user/history` already exists and is not empty:

1. Create an empty `/user/history-tmp` directory.
2. Make `/user/history-tmp` an encryption zone.
3. DistCp all data from `/user/history` into `/user/history-tmp`.
4. Remove `/user/history` and rename `/user/history-tmp` to `/user/history`.

KMS ACL Configuration for MapReduce

In the KMS ACLs, grant DECRYPT_EEK permission for the MapReduce key to the `mapred` and `yarn` users and the `hadoop` group:

```
<property>
  <name>key.acl.mapred-key.DECRYPT_EEK</name>
  <value>mapred,yarn hadoop</value>
  </description>
</property>
```

Search

Recommendations

Make `/solr` an encryption zone. When you create the encryption zone, name the key `solr-key` to take advantage of auto-generated Configuring KMS Access Control Lists (ACLs) on page 274KMS ACLs.

Steps

On a cluster without Solr currently installed, create the `/solr` directory and make that an encryption zone.

On a cluster with Solr already installed:

1. Create an empty `/solr-tmp` directory.
2. Make `/solr-tmp` an encryption zone.
3. DistCp all data from `/solr` into `/solr-tmp`.
4. Remove `/solr`, and rename `/solr-tmp` to `/solr`.

KMS ACL Configuration for Search

In the KMS ACL, grant the `solr` user and group `DECRYPT_EEK` permission for the Solr key:

```
<property>
  <name>key.acl.solr-key.DECRYPT_EEK</name>
  <value>solr solr</value>
  </description>
</property>
```

Spark

Recommendations

- By default, application event logs are stored at `/user/spark/applicationHistory`, which can be made into an encryption zone.
- Spark also optionally caches its JAR file at `/user/spark/share/lib` (by default), but encrypting this directory is not required.
- Spark does not encrypt shuffle data. To do so, configure the Spark local directory, `spark.local.dir` (in Standalone mode), to reside on an encrypted disk. For YARN mode, make the corresponding YARN configuration changes.

KMS ACL Configuration for Spark

In the KMS ACLs, grant `DECRYPT_EEK` permission for the Spark key to the `spark` user and any groups that can submit Spark jobs:

```
<property>
  <name>key.acl.spark-key.DECRYPT_EEK</name>
  <value>spark spark-users</value>
</property>
```

Sqoop

Recommendations

- **For Hive support:** Ensure that you are using Sqoop with the `--target-dir` parameter set to a directory that is inside the Hive encryption zone. For more details, see Hive on page 293.
- **For append/incremental support:** Make sure that the `sqoop.test.import.rootDir` property points to the same encryption zone as the `--target-dir` argument.
- **For HCatalog support:** No special configuration is required.

Cloudera Navigator Key Trustee Server

Cloudera Navigator Key Trustee Server is an enterprise-grade cryptographic key storage and management system used by Cloudera Navigator Encrypt on page 326 that separates encryption keys from the data, thus ensuring data is protected even if unauthorized users gain access to the storage media. It enables your cluster to meet the strictest data security regulations. Furthermore, Navigator Key Trustee Server can be integrated with a hardware security module (HSM) to provide the highest level of security for your keys. See Cloudera Navigator Key HSM on page 318 for details.

In addition, Navigator Key Trustee Server can be used by other cluster components. For example, HDFS Transparent Encryption on page 241 can use Navigator Key Trustee Server (KTS) as its backing key store (for Hadoop KMS, instead of the default Java KeyStore) for better security and scalability. See Migrating Keys from a Java KeyStore to Cloudera Navigator Key Trustee Server on page 291 for more information about using Navigator KTS with HDFS encryption.

> **Important:** Cloudera recommends that each cluster use its own KTS instance. Although sharing a single KTS across clusters is technically possible, it is neither approved nor supported for security reasons—specifically, the increased security risks associated with single point of failure for encryption keys used by multiple clusters.

After Installing Cloudera Navigator Key Trustee Server, follow the steps below to manage the system:

Backing Up and Restoring Key Trustee Server and Clients

Key Trustee Server high availability applies to read operations only. If either Key Trustee Server fails, the client automatically retries fetching keys from the functioning server. New write operations (for example, creating new encryption keys) are not allowed unless both Key Trustee Servers are operational.

If a Key Trustee Server fails catastrophically, you must restore it from backup to a new host with the same hostname and IP address as the failed host. Cloudera does not support PostgreSQL promotion to convert a passive Key Trustee Server to an active Key Trustee Server.

Cloudera strongly recommends regularly backing up Key Trustee Server databases and configuration files. Because these backups contain encryption keys and encrypted deposits, you must ensure that your backup repository is as secure as the Key Trustee Server.

You must also back up client configuration files and keys for Key Trustee Server clients, such as Key Trustee KMS and Navigator Encrypt clients.

> **Note:** In an HA configuration, the backup need only be performed on one of the hosts for Key Trustee Server and the Key Trustee KMS. For Key Trustee Server, run the backup on the *active* server. For Key Trustee KMS, you can run the backup on any instance.

Backing Up Key Trustee Server and Key Trustee KMS Using Cloudera Manager

Cloudera Manager versions 5.8 and higher, when used with Key Trustee Server and Key Trustee KMS versions 5.7 and higher, allow for backups of the KT Server and KT KMS configurations.

The actions executed in this procedure are equivalent to running the `ktbackup.sh` script on the node in question (see Backing Up Key Trustee Server and Key Trustee KMS Using the ktbackup.sh Script on page 300 for additional details.

In addition, when using the HDFS Encryption Wizard in Cloudera Manager 5.8 or higher to install and configure Key Trustee Server and Key Trustee KMS versions 5.7 and higher, a `cron` job is automatically set up to back up the Key Trustee Server on an ongoing basis. See Initializing Standalone Key Trustee Server on page 309 for more detail.

To back up the KT Server or KT KMS service configuration using Cloudera Manager:

1. Select the KT Server or KMS service configuration that you wish to back up.

2. For a KT Server backup, select **Create Backup on Active Server** (or **Create Backup on Passive Server**) from the Actions menu. For a KMS backup, select **Create Backup**.

A successfully completed backup of the KT Server is indicated by the message "Command Create Backup on Active Server finished successfully on service keytrustee_server".

Backing Up Key Trustee Server and Key Trustee KMS Using the ktbackup.sh Script

Key Trustee Server releases 5.7 and higher include a script, `ktbackup.sh`, to simplify and automate backing up Key Trustee Server. Key Trustee KMS releases 5.7 and higher include the same script for backing up Key Trustee KMS.

When run on a Key Trustee Server host, the script creates a tarball containing the Key Trustee Server private GPG keys and the PostgreSQL database. When run on a Key Trustee KMS host, the script creates a tarball containing the Key Trustee KMS private GPG keys and configuration file.

To preserve the security of the backup, you must specify a GPG recipient. Because this recipient is the only entity that can decrypt the backup, the recipient must be someone authorized to access the Key Trustee Server database, such as a key administrator.

Creating and Importing a GPG Key for Encrypting and Decrypting Backups

If the key administrator responsible for backing up and restoring Key Trustee Server and Key Trustee KMS does not already have a GPG key pair, they can create one using the `gpg --gen-key` command. The following example demonstrates this procedure:

> **Note:** By default, `gpg --gen-key` fails at the password prompt if you have logged in to your user account with the su command. You must log in to the SSH session with the user account for which you want to generate the GPG key pair.

```
[john.doe@backup-host ~]$ gpg --gen-key
gpg (GnuPG) 2.0.14; Copyright (C) 2009 Free Software Foundation, Inc.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

Please select what kind of key you want:
   (1) RSA and RSA (default)
   (2) DSA and Elgamal
   (3) DSA (sign only)
   (4) RSA (sign only)
Your selection? 1
RSA keys may be between 1024 and 4096 bits long.
What keysize do you want? (2048)
Requested keysize is 2048 bits
Please specify how long the key should be valid.
         0 = key does not expire
      <n>  = key expires in n days
      <n>w = key expires in n weeks
      <n>m = key expires in n months
      <n>y = key expires in n years
Key is valid for? (0)
Key does not expire at all
Is this correct? (y/N) y

GnuPG needs to construct a user ID to identify your key.

Real name: John Doe
Email address: john.doe@example.com
Comment: Key Trustee Backup
You selected this USER-ID:
    "John Doe (Key Trustee Backup) <john.doe@example.com>"

Change (N)ame, (C)omment, (E)mail or (O)kay/(Q)uit? O
You need a Passphrase to protect your secret key.

can't connect to `/home/john.doe/.gnupg/S.gpg-agent': No such file or directory
gpg-agent[10638]: directory `/home/john.doe/.gnupg/private-keys-v1.d' created
```

```
We need to generate a lot of random bytes. It is a good idea to perform
some other action (type on the keyboard, move the mouse, utilize the
disks) during the prime generation; this gives the random number
generator a better chance to gain enough entropy.
We need to generate a lot of random bytes. It is a good idea to perform
some other action (type on the keyboard, move the mouse, utilize the
disks) during the prime generation; this gives the random number
generator a better chance to gain enough entropy.
gpg: /home/john.doe/.gnupg/trustdb.gpg: trustdb created
gpg: key 0936CB67 marked as ultimately trusted
public and secret key created and signed.

gpg: checking the trustdb
gpg: 3 marginal(s) needed, 1 complete(s) needed, PGP trust model
gpg: depth: 0  valid:   1  signed:   0  trust: 0-, 0q, 0n, 0m, 0f, 1u
pub   2048R/0936CB67 2016-02-10
      Key fingerprint = CE57 FDED 3AFE E67D 2041  9EBF E64B 7D00 0936 CB67
uid                  John Doe (Key Trustee Backup) <john.doe@example.com>
sub   2048R/52A6FC5C 2016-02-10
```

After the GPG key pair is generated, you can export the public key:

```
[john.doe@backup-host ~]$ gpg --armor --output /path/to/johndoe.pub --export 'John Doe'
```

Copy the public key (`johndoe.pub` in this example) to the Key Trustee Server or Key Trustee KMS host, and import it into the service account keyring (`keytrustee` for Key Trustee Server and `kms` for Key Trustee KMS):

- On the Key Trustee Server host:

```
$ sudo -u keytrustee gpg --import /path/to/johndoe.pub
```

- On the Key Trustee KMS host:

```
$ sudo -u kms gpg --import /path/to/johndoe.pub
```

Running the ktbackup.sh Script

You must run `ktbackup.sh` as the service account. The location of the script depends on the service and installation method. See the following table for the script location and default service account for package- and parcel-based installations for Key Trustee Server and Key Trustee KMS.

Table 28: Backup Script Locations

Service	Service Account	Parcel-Based Installation	Package-Based Installation
Key Trustee Server	`keytrustee`	/opt/cloudera/parcels/KEYTRUSTEE-SERVER/bin/ktbackup.sh	`/usr/bin/ktbackup.sh`
Key Trustee KMS	`kms`	/opt/cloudera/parcels/KEYTRUSTEE/bin/ktbackup.sh	/usr/share/keytrustee-provider/bin/ktbackup.sh

The following table lists the command options for `ktbackup.sh`.

Table 29: Command Options for ktbackup.sh

Command Option	Description
`-c, --confdir=CONFDIR`	Specifies the Key Trustee configuration directory. Defaults to `/var/lib/keytrustee/.keytrustee` for parcel-based Key Trustee Server. For Key Trustee KMS and package-based Key Trustee Server, you must specify this option.

Command Option	Description
`--database-port=PORT`	Specifies the Key Trustee Server database port. Defaults to `11381` for parcel-based installations. For package-based Key Trustee Server installations, you must specify this option.
`--gpg-recipient=GPG_RECIPIENT`	Specifies the GPG recipient. The backup is encrypted with the public key of the specified recipient. The GPG recipient public key must be imported into the service account keyring before running the script. See Creating and Importing a GPG Key for Encrypting and Decrypting Backups on page 300 for more information.
`--cleartext`	Outputs an unencrypted tarball. To preserve the security of the cryptographic keys, *do not* use this option in production environments.
`--output=DIR`	Specifies the output directory for the tarball. Defaults to `/var/lib/keytrustee` for parcel-based Key Trustee Server. For Key Trustee KMS and package-based Key Trustee Server, you must specify this option.
`--roll=`*n*	Deletes backups older than the last *n* backups from the directory specified by the `--output` parameter. For example, if you have 10 backups, specifying `--roll=10` creates a new backup (11 backups total) and then delete the oldest backup. Specifying `--roll=1` creates a new backup and then deletes all other backups. **Note:** This option works for Key Trustee Server only.
`-q, --quiet`	Suppresses console log messages and, if successful, returns only the backup tarball file path. This is useful for automating backups.
`--verbose`	Outputs additional log messages to the console for debugging.

Important: Running the `ktbackup.sh` script on the KMS server creates a backed up copy of the GPG private keys currently being used by the KMS. If this backup is not performed, and the GPG private keys are not saved, then the system cannot be fully restored in the event of a catastrophic failure, and *access to both keys and data may be lost*. It is therefore imperative that you perform a manual backup after the KMS starts successfully for the first time, and on all KMS instances whenever a new Key Trustee High Availability instance is added. Running this backup ensures that the most recent backup contains a copy of the GPG private keys currently in use.

The following examples demonstrate the command usage for different scenarios:

- To back up a parcel-based Key Trustee Server, specifying the GPG recipient by name:

```
$ sudo -u keytrustee /opt/cloudera/parcels/KEYTRUSTEE_SERVER/bin/ktbackup.sh
--gpg-recipient='John Doe'
```

- To back up a parcel-based Key Trustee KMS, specifying the GPG recipient by email:

```
$ sudo -u kms /opt/cloudera/parcels/KEYTRUSTEE/bin/ktbackup.sh -c
/var/lib/kms-keytrustee/keytrustee/.keytrustee --output=/var/lib/kms-keytrustee
--gpg-recipient=john.doe@example.com
```

- To back up a package-based Key Trustee Server with the database running on a non-default port (12345 in this example):

```
$ sudo -u keytrustee ktbackup.sh --database-port=12345
--gpg-recipient=john.doe@example.com
```

- To back up a package-based Key Trustee KMS, specifying the GPG recipient by email:

```
$ sudo -u kms /usr/share/keytrustee-keyprovider/bin/ktbackup.sh -c
/var/lib/kms-keytrustee/keytrustee/.keytrustee --output=/var/lib/kms-keytrustee
--gpg-recipient=john.doe@example.com
```

Automating Backups Using cron

You can schedule automatic backups of Key Trustee Server using the `cron` scheduling utility.

Create a `crontab` entry using the following commands:

1. Edit the `crontab` by running the following command:

```
$ sudo -u keytrustee crontab -e
```

2. Add the following entry to run the backup script every 30 minutes. This example is for a parcel-based installation of Key Trustee Server. See the Backup Script Locations table for the package-based script location.

```
*/30 * * * * /opt/cloudera/parcels/KEYTRUSTEE_SERVER/bin/ktbackup.sh --gpg-recipient='John
 Doe' --quiet --output=/tmp/backups --roll=10
```

   Run `man 5 crontab` to see the `crontab` man page for details on using `cron` to schedule backups at different intervals.

Backing Up Key Trustee Server Manually

Use this procedure for both parcel-based and package-based installations.

If you have deployed Cloudera Navigator Key Trustee Server High Availability, perform these steps on both the active and passive Key Trustee Servers. The following procedure references the default database port and location; if you modified these settings during installation, replace the database and port with your values.

1. Back up the Key Trustee Server database:

   - For Key Trustee Server 3.8:

```
$ su - postgres
$ pg_dump -c -p 5432 keytrustee | zip --encrypt keytrustee-db.zip -
```

   - For Key Trustee Server 5.4 and higher:

```
$ su - keytrustee
$ pg_dump -c -p 11381 keytrustee | zip --encrypt keytrustee-db.zip -
```

   The `--encrypt` option prompts you to create a password used to encrypt the zip file. This password is required to decrypt the file.

For parcel-based installations, you must set environment variables after switching to the `keytrustee` user:

```
$ su - keytrustee
$ export PATH=$PATH:/opt/cloudera/parcels/KEYTRUSTEE_SERVER/PG_DB/opt/postgres/9.3/bin
$ export
LD_LIBRARY_PATH=/opt/cloudera/parcels/KEYTRUSTEE_SERVER/PG_DB/opt/postgres/9.3/lib
$ pg_dump -c -p 11381 keytrustee | zip --encrypt keytrustee-db.zip -
```

2. Back up the Key Trustee Server configuration directory (`/var/lib/keytrustee/.keytrustee`):

```
$ zip -r --encrypt keytrustee-conf.zip /var/lib/keytrustee/.keytrustee
```

The `--encrypt` option prompts you to create a password used to encrypt the zip file. This password is required to decrypt the file.

3. Move the backup files (*keytrustee-db.zip* and *keytrustee-conf.zip*) to a secure location.

Backing Up Key Trustee Server Clients

Cryptographic keys stored in Key Trustee Server are encrypted by clients before they are sent to Key Trustee Server. The primary clients for Key Trustee Server are Key Trustee KMS and Navigator Encrypt. Cloudera strongly recommends backing up regularly the configuration files and GPG keys for Key Trustee Server clients. See Backing Up Key Trustee Server and Key Trustee KMS Using the ktbackup.sh Script on page 300 for instructions on backing up Key Trustee KMS using the provided backup script.

> **Warning:** Failure to back up these files can result in irretrievable data loss. For example, encryption zone keys used for HDFS Transparent Encryption on page 241 are encrypted by the KMS before being stored in Key Trustee Server. A catastrophic failure of the KMS with no backup causes all HDFS data stored in encryption zones to become permanently irretrievable.

To prevent permanent data loss, regularly back up the following directories on each client that stores objects in Key Trustee Server:

Table 30: Key Trustee Server Client Configuration Directories

Key Trustee Server Client	Directories to Back Up
Key Trustee KMS	`/var/lib/kms-keytrustee`
Navigator Encrypt	`/etc/navencrypt`

Restoring Key Trustee Server

When restoring the Key Trustee Server database from backup, keep in mind that any keys or deposits created after the backup are not restored. If you are using Key Trustee Server high availability, you can restore the Active Key Trustee Server from the Passive Key Trustee Server. This restores all keys that were successfully written to the Passive Key Trustee Server before the failure.

The procedure to restore Key Trustee Server is different for parcel-based than for package-based installations. For more information about parcels, see Parcels.

Restoring Key Trustee Server in Parcel-Based Installations

> **Note:** These instructions apply to Key Trustee Servers deployed using parcels. For package-based deployments, skip to the Restoring Key Trustee Server in Package-Based Installations on page 305 section.

If you have deployed Cloudera Navigator Key Trustee Server High Availability, perform these steps on both the active and passive Key Trustee Servers. The following procedures assume the default database port and location; if you modified these settings during installation, replace the database and port with your custom values.

If the Key Trustee Server host has failed completely, remove the host from the cluster and add a new host using Cloudera Manager:

1. Remove the failed host from the cluster. See Deleting Hosts for instructions.
2. Add a new host with the same hostname and IP address as the failed host to the cluster. See Adding a Host to the Cluster for instructions.

   Important: Make sure that the replacement host uses the same operating system version as the failed host.

3. Install Key Trustee Server on the new host. See Installing Cloudera Navigator Key Trustee Server for instructions. Make sure to install the same Key Trustee Server version as the failed host.

After you have provisioned a new host and installed Key Trustee Server (or if you are restoring the database or configuration on the original host), restore the database and configuration directory. If your backups were created using the `ktbackup.sh` script, skip to Restoring Key Trustee Server and Key Trustee KMS from ktbackup.sh Backups on page 306. If you need to restore the Active Key Trustee Server from the Passive Key Trustee Server, skip to Restoring Active Key Trustee Server from Passive Key Trustee Server on page 308.

If your backups were created manually using the `pg_dump` command, do the following:

1. Copy or move the backup files (*keytrustee-db.zip* and *keytrustee-conf.zip*) to the Key Trustee Server host.
2. Start the PostgreSQL server:

```
$ sudo ktadmin db --start --pg-rootdir /var/lib/keytrustee/db --background
```

3. Restore the Key Trustee Server database:

```
$ su - keytrustee
$ export PATH=$PATH:/opt/cloudera/parcels/KEYTRUSTEE_SERVER/PG_DB/opt/postgres/9.3/bin
$ export
LD_LIBRARY_PATH=/opt/cloudera/parcels/KEYTRUSTEE_SERVER/PG_DB/opt/postgres/9.3/lib
$ unzip -p /path/to/keytrustee-db.zip | psql -p 11381 -d keytrustee
```

   If the zip file is encrypted, you are prompted for the password to decrypt the file.

4. Restore the Key Trustee Server configuration directory:

```
$ su - keytrustee
$ cd /var/lib/keytrustee
$ unzip /path/to/keytrustee-conf.zip
```

   If the zip file is encrypted, you are prompted for the password to decrypt the file.

5. Stop the PostgreSQL server:

```
$ sudo ktadmin db --stop --pg-rootdir /var/lib/keytrustee/db
```

6. Start the Key Trustee Server service in Cloudera Manager (**Key Trustee Server service** > **Actions** > **Start**).
7. Restart the Key Trustee KMS service in Cloudera Manager (**Key Trustee KMS service** > **Actions** > **Restart**).
8. Remove the backup files (*keytrustee-db.zip* and *keytrustee-conf.zip*) from the Key Trustee Server host.

Restoring Key Trustee Server in Package-Based Installations

If you have deployed Cloudera Navigator Key Trustee Server High Availability, perform these steps on both the active and passive Key Trustee Servers. The following procedures assume the default database port and location; if you modified these settings during installation, replace the database and port with your custom values.

If the Key Trustee Server host has failed completely, provision a new host with the same hostname and IP address as the failed host, and re-install Key Trustee Server. See Installing Cloudera Navigator Key Trustee Server for instructions.

> **Important:** Make sure to install the same operating system and Key Trustee Server versions as the failed host.

After you have provisioned a new host and installed Key Trustee Server (or if you are restoring the database or configuration on the original host), restore the database and configuration directory. If your backups were created using the `ktbackup.sh` script, skip to Restoring Key Trustee Server and Key Trustee KMS from ktbackup.sh Backups on page 306. If you need to restore the Active Key Trustee Server from the Passive Key Trustee Server, skip to Restoring Active Key Trustee Server from Passive Key Trustee Server on page 308.

If your backups were created manually using the `pg_dump` command, do the following:

1. Copy or move the backup files (*keytrustee-db.zip* and *keytrustee-conf.zip*) to the Key Trustee Server host.
2. Change the file ownership on the backup files to `keytrustee:keytrustee`:

```
$ sudo chown keytrustee:keytrustee /path/to/keytrustee*.zip
```

3. Restore the Key Trustee Server database:

```
$ su - keytrustee
$ unzip -p /path/to/keytrustee-db.zip | psql -p 11381 -d keytrustee
```

   If the zip file is encrypted, you are prompted for the password to decrypt the file.

4. Restore the Key Trustee Server configuration directory:

```
$ cd /var/lib/keytrustee
$ unzip /path/to/keytrustee-conf.zip
```

   If the zip file is encrypted, you are prompted for the password to decrypt the file.

5. Start the Key Trustee Server service:
   - RHEL 6-compatible: `$ sudo service keytrusteed start`
   - RHEL 7-compatible: `$ sudo systemctl start keytrusteed`
6. Remove the backup files (*keytrustee-db.zip* and *keytrustee-conf.zip*) from the Key Trustee Server host.

Restoring Key Trustee Server and Key Trustee KMS from ktbackup.sh Backups

After installing Key Trustee Server or Key Trustee KMS on a new host after a failure, or if you need to restore accidentally deleted keys on the same host, use the following procedure to restore Key Trustee Server or Key Trustee KMS from backups generated by the `ktbackup.sh` script.

1. Decrypt the backup tarball using the private key of the GPG recipient specified in the backup command by running the following command as the GPG recipient user account. The GPG recipient private key must be available on the Key Trustee Server or Key Trustee KMS host on which you run this command.

```
$ gpg -d -o /path/to/decrypted/backup.tar /path/to/encrypted/tarball
```

2. Verify the decrypted tarball using the `tar tvf /path/to/decrypted/backup.tar` command. For example:

```
$ tar tvf kts_bak_kts01_example_com_2016-02-10_11-14-37.tar
drwx------ keytrustee/keytrustee 0 2016-02-09 16:43 var/lib/keytrustee/.keytrustee/
-rw------- keytrustee/keytrustee 434 2016-02-09 16:43
var/lib/keytrustee/.keytrustee/keytrustee.conf
-rw------- keytrustee/keytrustee 1280 2016-02-09 16:43
var/lib/keytrustee/.keytrustee/trustdb.gpg
-rw------- keytrustee/keytrustee 4845 2016-02-09 16:43
var/lib/keytrustee/.keytrustee/secring.gpg
```

```
-rw------- keytrustee/keytrustee  600 2016-02-09 16:43
var/lib/keytrustee/.keytrustee/random_seed
drwx------ keytrustee/keytrustee    0 2016-02-09 16:40
var/lib/keytrustee/.keytrustee/.ssl/
-rw------- keytrustee/keytrustee 1708 2016-02-09 16:40
var/lib/keytrustee/.keytrustee/.ssl/ssl-cert-keytrustee-pk.pem
-rw------- keytrustee/keytrustee 1277 2016-02-09 16:40
var/lib/keytrustee/.keytrustee/.ssl/ssl-cert-keytrustee.pem
-rw------- keytrustee/keytrustee 2263 2016-02-09 16:43
var/lib/keytrustee/.keytrustee/pubring.gpg
-rw-r--r-- keytrustee/keytrustee  457 2016-02-09 16:43
var/lib/keytrustee/.keytrustee/logging.conf
-rw------- keytrustee/keytrustee 2263 2016-02-09 16:43
var/lib/keytrustee/.keytrustee/pubring.gpg~
-rw------- keytrustee/keytrustee  157 2016-02-09 16:40
var/lib/keytrustee/.keytrustee/gpg.conf
-rw-r--r-- keytrustee/keytrustee 47752 2016-02-10 11:14
var/lib/keytrustee/kts_bak_kts01_example_com_2016-02-10_11-14-37.sql
```

3. Restore the files to their original locations, using this command for both Key Trustee Server and Key Trustee KMS backups:

```
$ tar xvf /path/to/decrypted/backup.tar -C /
```

4. **(Key Trustee Server Only)** Drop and re-create the `keytrustee` PostgreSQL database, and restore the database from the backup.

   - For parcel-based installations:

```
$ su - keytrustee
$ source /opt/cloudera/parcels/KEYTRUSTEE_SERVER/meta/keytrustee_env.sh
$ /opt/cloudera/parcels/KEYTRUSTEE_SERVER/PG_DB/opt/postgres/9.3/bin/psql -p 11381
psql (9.3.6)
Type "help" for help.

keytrustee=# \list
                                      List of databases
    Name    |   Owner    | Encoding |   Collate   |    Ctype    |     Access privileges
------------+------------+----------+-------------+-------------+-----------------------

 keytrustee | keytrustee | UTF8     | en_US.UTF-8 | en_US.UTF-8 |
 postgres   | keytrustee | UTF8     | en_US.UTF-8 | en_US.UTF-8 |
 template0  | keytrustee | UTF8     | en_US.UTF-8 | en_US.UTF-8 | =c/keytrustee
     +
            |            |          |             |             |
keytrustee=CTc/keytrustee
 template1  | keytrustee | UTF8     | en_US.UTF-8 | en_US.UTF-8 | =c/keytrustee
     +
            |            |          |             |             |
keytrustee=CTc/keytrustee
(4 rows)

keytrustee=# \c postgres;
You are now connected to database "postgres" as user "keytrustee".
postgres=# drop database keytrustee;
DROP DATABASE
postgres=# create database keytrustee;
CREATE DATABASE
postgres=# \q
$ sudo -u keytrustee
/opt/cloudera/parcels/KEYTRUSTEE_SERVER/PG_DB/opt/postgres/9.3/bin/psql -p 11381 -f
/var/lib/keytrustee/kts_bak_kts01_example_com_2016-02-10_11-14-37.sql
```

   - For package-based installations:

```
$ su - keytrustee
$ psql -p 11381
psql (9.3.6)
Type "help" for help.
```

```
keytrustee=# \list
                                   List of databases
    Name    |   Owner    | Encoding |   Collate   |    Ctype    |     Access privileges
------------+------------+----------+-------------+-------------+-----------------------

 keytrustee | keytrustee | UTF8     | en_US.UTF-8 | en_US.UTF-8 |
 postgres   | keytrustee | UTF8     | en_US.UTF-8 | en_US.UTF-8 |
 template0  | keytrustee | UTF8     | en_US.UTF-8 | en_US.UTF-8 | =c/keytrustee
     +
            |            |          |             |             |
keytrustee=CTc/keytrustee
 template1  | keytrustee | UTF8     | en_US.UTF-8 | en_US.UTF-8 | =c/keytrustee
     +
            |            |          |             |             |
keytrustee=CTc/keytrustee
(4 rows)

keytrustee=# \c postgres;
You are now connected to database "postgres" as user "keytrustee".
postgres=# drop database keytrustee;
DROP DATABASE
postgres=# create database keytrustee;
CREATE DATABASE
postgres=# \q
$ sudo -u keytrustee psql -p 11381 -f
/var/lib/keytrustee/kts_bak_kts01_example_com_2016-02-10_11-14-37.sql
```

5. Restart Key Trustee Server.
   - **Using Cloudera Manager: Key Trustee Server service** > **Actions** > **Restart**
   - **Using the Command Line:** Run the following command on the Key Trustee Server hosts:

```
$ sudo service keytrusteed restart      #RHEL 6-compatible
$ sudo systemctl restart keytrusteed    #RHEL 7-compatible
```

6. Restart the Key Trustee KMS service in Cloudera Manager (**Key Trustee KMS service** > **Actions** > **Restart**).

Restoring Active Key Trustee Server from Passive Key Trustee Server

If the Active Key Trustee Server fails, and you do not have a backup, you can restore it from the Passive Key Trustee Server using the following procedure. You can also use this procedure if you need to restore keys that were successfully written to the Passive Key Trustee Server, but are not included in the most recent backup.

The following procedure assumes you have installed Key Trustee Server on the replacement host and (if you are using Cloudera Manager) added the Key Trustee Server service. For instructions on installing Key Trustee Server, see Installing Cloudera Navigator Key Trustee Server

1. Copy the Key Trustee Server database from the Passive Key Trustee Server host to the new Active Key Trustee Server host. Run the following command on the Passive Key Trustee Server host:

```
$ sudo rsync --exclude recovery.conf -a /var/lib/keytrustee/db
root@kts01.example.com:/var/lib/keytrustee/
```

   Replace *kts01.example.com* with the hostname of the new Active Key Trustee Server.

2. Make sure that the `recovery.conf` file did not get copied to the Active Key Trustee Server (for example, if there was a typo in your `rsync` command). Run the following command on the Active Key Trustee Server host:

```
$ sudo ls -l /var/lib/keytrustee/db/recovery.conf
```

   If the file exists on the Active Key Trustee Server host, delete it. Make sure you are on the Active Key Trustee Server host before deleting the file. Do not delete the `recovery.conf` file on the Passive Key Trustee Server host.

3. Copy the configuration directory from the Passive Key Trustee Server host to the new Active Key Trustee Server host. Run the following command on the Passive Key Trustee Server host:

```
$ sudo rsync --exclude .ssl --exclude '*.pid' -a /var/lib/keytrustee/.keytrustee
root@kts01.example.com:/var/lib/keytrustee/
```

   Replace *kts01.example.com* with the hostname of the new Active Key Trustee Server.

4. Create the `logs` directory and make sure it is owned by the `keytrustee` user and group:

```
$ sudo mkdir /var/lib/keytrustee/logs
$ sudo chown keytrustee:keytrustee /var/lib/keytrustee/logs
```

5. **(Cloudera Manager only)** Generate the Key Trustee Server keyring: **Key Trustee Server service** > **Actions** > **Generate Key Trustee Server Keyring**
6. Set up the database on the Active Key Trustee Server host.
   - **Using Cloudera Manager: Key Trustee Server service** > **Actions** > **Set Up Key Trustee Server Database**
   - **Using the Command Line:**

```
$ sudo ktadmin --confdir /var/lib/keytrustee db --port 11381 --pg-rootdir
/var/lib/keytrustee/db --bootstrap --slave kts02.example.com
```

   Replace *kts02.example.com* with the hostname of the Passive Key Trustee Server.

7. Start the database.
   - **Using Cloudera Manager: Key Trustee Server service** > **Instances** > **Active Database** > **Actions** > **Start this Active Database**
   - **Using the Command Line:** Run the following command on the Active Key Trustee Server host:

```
$ sudo ktadmin --confdir /var/lib/keytrustee db --port 11381 --pg-rootdir
/var/lib/keytrustee/db --bootstrap --slave kts02.example.com
```

   Replace *kts02.example.com* with the hostname of the Passive Key Trustee Server.

8. Enable synchronous replication.
   - **Using Cloudera Manager: Key Trustee Server service** > **Actions** > **Setup Enable Synchronous Replication in HA mode**
   - **Using the Command Line:** Run the following command on the Active Key Trustee Server host:

```
$ sudo ktadmin --confdir /var/lib/keytrustee enable-synchronous-replication
```

9. Restart the Active Key Trustee Server.
   - **Using Cloudera Manager: Key Trustee Server service** > **Actions** > **Restart**
   - **Using the Command Line:** Run the following command on the Active Key Trustee Server host:

```
$ sudo service keytrusteed restart      #RHEL 6-compatible
$ sudo systemctl restart keytrusteed    #RHEL 7-compatible
```

10. Restart the Key Trustee KMS service in Cloudera Manager (**Key Trustee KMS service** > **Actions** > **Restart**).

Initializing Standalone Key Trustee Server

If you are configuring high availability Key Trustee Server, skip this step and proceed to Cloudera Navigator Key Trustee Server High Availability. Cloudera strongly recommends configuring high availability for Key Trustee Server.

Initializing Standalone Key Trustee Server Using Cloudera Manager

> **Important:** If you are using SSH software other than OpenSSH, the initialization fails. To prevent this, pre-create the SSH key before continuing:
>
> ```
> $ sudo -u keytrustee ssh-keygen -t rsa -f /var/lib/keytrustee/.ssh/id_rsa
> ```

For new installations, use the **Set up HDFS Data At Rest Encryption** wizard and follow the instructions in Enabling HDFS Encryption Using the Wizard on page 246. When prompted, deselect the **Enable High Availability** option to proceed in standalone mode.

To set up Key Trustee Server manually, add the Key Trustee Server service to your cluster, following the instructions in Adding a Service. When customizing role assignments, assign only the Active Key Trustee Server and Active Database roles.

> **Important:** You *must* assign the Key Trustee Server and Database roles to the same host. Key Trustee Server does not support running the database on a different host.

For parcel-based Key Trustee Server releases 5.8 and higher, Cloudera Manager automatically backs up Key Trustee Server (using the `ktbackup.sh` script) after adding the Key Trustee Server service. It also schedules automatic backups using `cron`. For package-based installations, you must manually back up Key Trustee Server and configure a `cron` job.

Cloudera Manager configures `cron` to run the backup script hourly. The latest 10 backups are retained in `/var/lib/keytrustee` in cleartext. For information about using the backup script and configuring the `cron` job (including how to encrypt backups), see Backing Up Key Trustee Server and Key Trustee KMS Using the ktbackup.sh Script on page 300.

Initializing Standalone Key Trustee Server Using the Command Line

To initialize a standalone Key Trustee Server, run the following commands on the Key Trustee Server:

> **Important:** For Key Trustee Server 5.4.0 and higher, the `ktadmin init-master` command is deprecated. Use the `ktadmin init` command instead. If you are using SSH software other than OpenSSH, the initialization fails. To prevent this, pre-create the SSH key before continuing:
>
> ```
> $ sudo -u keytrustee ssh-keygen -t rsa /var/lib/keytrustee/.ssh/id_rsa
> ```

```
$ sudo ktadmin init --external-address keytrustee.example.com
$ sudo ktadmin db --bootstrap --port 11381 --pg-rootdir /var/lib/keytrustee/db
## For RHEL/CentOS 7, use 'sudo systemctl [stop|start] <service_name>' instead of 'sudo
 service <service_name> [stop|start]' ##
$ sudo service keytrustee-db stop
$ sudo service keytrustee-db start
$ sudo service keytrusteed start
$ sudo chkconfig keytrustee-db on
$ sudo chkconfig keytrusteed on
```

Replace *`keytrustee.example.com`* with the fully qualified domain name (FQDN) of the Key Trustee Server. Cloudera recommends using the default `/var/lib/keytrustee/db` directory for the PostgreSQL database.

To use a different port for the database, modify the `ktadmin init` and `ktadmin db` commands as follows:

```
$ sudo ktadmin init --external-address keytrustee.example.com --db-connect
postgresql://localhost:<port>/keytrustee?host=/tmp
$ sudo ktadmin db --bootstrap --port <port> --pg-rootdir /var/lib/keytrustee/db
```

If you specify a database directory other than `/var/lib/keytrustee/db`, create or edit the `/etc/sysconfig/keytrustee-db` file and add the following line:

```
ARGS="--pg-rootdir /path/to/db"
```

The `ktadmin init` command initializes the Key Trustee configuration directory (`/var/lib/keytrustee/.keytrustee` by default) and generates a self-signed certificate that Key Trustee Server uses for HTTPS communication.

The `ktadmin db --bootstrap` command initializes the database in the directory specified by the `--pg-rootdir` parameter.

The `sudo service keytrustee-db stop` and `sudo service keytrustee-db start` commands restart the Key Trustee Server database.

The `sudo service keytrusteed start` command starts Key Trustee Server.

> **Note:** The `/etc/init.d/postgresql` script does not work when the PostgreSQL database is started by Key Trustee Server, and cannot be used to monitor the status of the database. Use `/etc/init.d/keytrustee-db` instead.

(Optional) Replace Self-Signed Certificate with CA-Signed Certificate

> **Important:** Key Trustee Server certificates must be issued to the fully qualified domain name (FQDN) of the Key Trustee Server host. If you are using CA-signed certificates, ensure that the generated certificates use the FQDN, and not the short name.

If you have a CA-signed certificate for Key Trustee Server, see Managing Key Trustee Server Certificates on page 315 for instructions on how to replace the self-signed certificate.

Specifying TLS/SSL Minimum Allowed Version and Ciphers

Depending on your cluster configuration and the security practices in your organization, you might need to restrict the allowed versions of TLS/SSL used by Key Trustee Server. Older TLS/SSL versions might have vulnerabilities or lack certain features.

Specify one of the following values using the Minimum TLS Support configuration setting:

- `tlsv1`: Allow any TLS version of 1.0 or higher. This setting is the default when TLS/SSL is enabled.
- `tlsv1.1`: Allow any TLS version of 1.1 or higher.
- `tlsv1.2`: Allow any TLS version of 1.2 or higher.

> **Note:** The `pyOpenSSL` version on the Key Trustee Server cluster should be updated to 16.2 before changing the TLS version to 1.2. If `pyOpenSSL` is not updated, then the following error appears when the Key Trustee Server service attempts to restart:
>
> ```
> keytrustee-server: Error in setting the protocol to the value TLSv1.2.
> This usually means there is no support for the entered value.
> Python Error: 'module' object has no attribute 'OP_NO_TLSv1_1"
> ```

Along with specifying the version, you can also specify the allowed set of TLS ciphers using the Supported Cipher Configuration for SSL configuration setting. The argument to this option is a list of keywords, separated by colons, commas, or spaces, and optionally including other notation.

```
AES256:CAMELLIA256-SHA
```

By default, the cipher list is empty, and Key Trustee Server uses the default cipher list for the underlying platform. See the output of `man ciphers` for the full set of keywords and notation allowed in the argument string.

Configuring a Mail Transfer Agent for Key Trustee Server

The Key Trustee Server requires a mail transfer agent (MTA) to send email. Cloudera recommends Postfix, but you can use any MTA that meets your needs.

To configure Postfix for local delivery, run the following commands:

```
export
KEYTRUSTEE_SERVER_PK="/var/lib/keytrustee/.keytrustee/.ssl/ssl-cert-keytrustee-pk.pem"
export
KEYTRUSTEE_SERVER_CERT="/var/lib/keytrustee/.keytrustee/.ssl/ssl-cert-keytrustee.pem"
export
KEYTRUSTEE_SERVER_CA="/var/lib/keytrustee/.keytrustee/.ssl/ssl-cert-keytrustee-ca.pem"
export KEYTRUSTEE_SERVER_HOSTNAME="$(hostname -f)"  # or adjust as required
postconf -e 'mailbox_command ='
postconf -e 'smtpd_sasl_local_domain ='
postconf -e 'smtpd_sasl_auth_enable = yes'
postconf -e 'smtpd_sasl_security_options = noanonymous'
postconf -e 'broken_sasl_auth_clients = yes'
postconf -e 'smtpd_recipient_restrictions =
permit_sasl_authenticated,permit_mynetworks,reject_unauth_destination'
postconf -e 'inet_interfaces = localhost'
postconf -e 'smtp_tls_security_level = may'
postconf -e 'smtpd_tls_security_level = may'
```

Start the Postfix service and ensure that it starts at boot:

```
$ service postfix restart
$ sudo chkconfig --level 235 postfix on
```

For information on installing Postfix or configuring a relay host, see the Postfix documentation.

Verifying Cloudera Navigator Key Trustee Server Operations

Verify that the installation was successful by running the following command on all Key Trustee Servers.

```
curl -k https://keytrustee.example.com:11371/?a=fingerprint
```

Replace `keytrustee.example.com` with the fully qualified domain name (FQDN) of each Key Trustee Server you are validating. This command outputs a fingerprint similar to the following:

```
4096R/4EDC46882386C827E20DEEA2D850ACA33BEDB0D1
```

If high availability is enabled, the output should be identical for all Key Trustee Servers.

Managing Key Trustee Server Organizations

Organizations allow you to configure Key Trustee for use in a multi-tenant environment. Using the `keytrustee-orgtool` utility, you can create organizations and administrators for multiple organizations. Organization administrators can then approve or deny the registration of clients, depending on the registration method.

The `keytrustee-orgtool` Utility

`keytrustee-orgtool` is a command-line utility for administering organizations. The `keytrustee-orgtool` command must be run as the root user.

The following table explains the various `keytrustee-orgtool` commands and parameters. Run `keytrustee-orgtool --help` to view this information at the command line.

Table 31: Usage for `keytrustee-orgtool`

Operation	Usage	Description
Add	`keytrustee-orgtool add [-h] –n` *`name`* `–c` *`contacts`*	Adds a new organization and administrators for the organization.
List	`keytrustee-orgtool list`	Lists current organizations, including the authorization secret, all administrators, the organization creation date, and the organization expiration date.
Disable client	`keytrustee-orgtool disable-client [-h] --fingerprint` *`fingerprint`*	Disables a client that has already been activated by the organization administrator.
Enable client	`keytrustee-orgtool enable-client [-h] --fingerprint` *`fingerprint`*	Enables a client that has requested activation but has not yet been approved by the organization administrator.
Set authorization Code	`keytrustee-orgtool set-auth [-h] –n` *`name`* `–s` *`secret`*	Sets the authorization code to a new string, or to blank to allow automatic approvals without the code.

Create Organizations

Each new Key Trustee tenant needs its own organization. You can create new organizations using the `keytrustee-orgtool add` command. For example, to create a new organization for the Disaster Recovery group and add two administrators, Finn and Jake:

```
$ keytrustee-orgtool add -n disaster-recov -c finn@example.com,jake@example.com
```

When adding organizations:

- Do not use spaces or special characters in the organization name. Use hyphens or underscores instead.
- Do not use spaces between email addresses (when adding multiple administrators to an organization). Use a comma to separate email addresses, without any space (as shown in the example above).

Each contact email address added when creating the organization receives a notification email, as detailed below.

Once an organization exists, use the `keytrustee-orgtool add` command to add new administrators to the organization. For example, to add an administrator to the `disaster-recov` organization:

```
keytrustee-orgtool add -n disaster-recov -c marceline@example.com
```

Note: You cannot remove contacts from an organization with the `keytrustee-orgtool` utility.

List Organizations

After creating an organization, verify its existence with the `keytrustee-orgtool list` command. This command lists details for all existing organizations. The following is the entry for the `disaster-recov` organization created in the example:

```
"disaster-recov": {
           "auth_secret": "/qFiICsyYqMLhdTznNY3Nw==",
           "contacts": [
            "finn@example.com",
            "jake@example.com"
        ],
        "creation": "2013-12-02T09:55:21",
        "expiration": "9999-12-31T15:59:59",
        "key_info": null,
        "name": "disaster-recov",
        "state": 0,
        "uuid": "xY3Z8xCwMuKZMiTYJa0mZOdhMVdxhyCUOc6vSNc9I8X"
    }
```

Change the Authorization Code

When an organization is created, an authorization code is automatically generated. When you run the `keytrustee-orgtool list` command, the code is displayed in the `auth_secret` field. To register with a Key Trustee Server, the client must have the authorization code along with the organization name. To set a new `auth_secret`, run the following command:

```
$ keytrustee-orgtool set-auth -n disaster-recov -s ThisISAs3cr3t!
```

Run the `keytrustee-orgtool list` command again, and confirm the updated `auth_secret` field:

```
"disaster-recov": {
          "auth_secret": "ThisISAs3cr3t!",
          "contacts": [
            "finn@example.com",
            "jake@example.com"
        ],
        "creation": "2013-12-02T09:55:21",
        "expiration": "9999-12-31T15:59:59",
        "key_info": null,
        "name": "disaster-recov",
        "state": 0,
        "uuid": "xY3Z8xCwMuKZMiTYJa0mZOdhMVdxhyCUOc6vSNc9I8X"
    }
```

If you do not want to use an authorization code, set the `auth_secret` field to an empty string:

```
$ keytrustee-orgtool set-auth -n disaster-recov -s ""
```

Cloudera recommends requiring an authorization code.

Notification Email and GPG Keys

Whenever an administrator is added to an organization, the Key Trustee Server sends an automated email message (subject: “KeyTrustee Contact Registration”) to the newly added administrator:

```
Hello, this is an automated message from your Cloudera keytrustee Server.

Welcome to Cloudera keytrustee! You have been listed as an administrator contact
for keytrustee services at your organization [test-org]. As an administrator,
you may be contacted to authorize the activation of new keytrustee clients.

We recommend that you register a GPG public key for secure administration of
your clients. To do so, visit the link below and follow the instructions.
```

```
https://keytrustee01.example.com:11371/?q=CnRV6u0nbm7zB07BQEpXCXsN0QJFBz684uC0lcHMoWL

This link will expire in 12 hours, at Thu Sep  3 00:08:25 2015 UTC.
```

Cloudera recommends that an organization's administrators:

- Register the GPG public key by following the link contained in the notification email. Registering the GPG public key is optional, but if you choose to register your public key:
  - Complete the process within 12 hours, before the link expires.
  - Upload the entire key, including the BEGIN and END strings, as shown here:

```
-----BEGIN PGP PUBLIC KEY BLOCK-----
Version: GnuPG v1.2.1 (GNU/Linux)
Comment: For info see http://www.gnupg.org

mQGiBDkHP3URBACkWGsYh43pkXU9wj/X1G67K8/DSrl85r7dNtHNfLL/ewil10k2
q8saWJn26QZPsDVqdUJMOdHfJ6kQTAt9NzQbgcVrxLYNfgeBsvkHF/POtnYcZRgL
tZ6syBBWs8JB4xt5V09iJSGAMPUQE8Jpdn2aRXPApdoDw179LM8Rq6r+gwCg5ZZa
. . .
-----END PGP PUBLIC KEY BLOCK-----
```

- Import the Key Trustee Server's public GPG key to verify that the server is the sender.

The organization's administrators are notified by email when new clients are registered to the Key Trustee Server using the mail transfer agent (as discussed in Configuring a Mail Transfer Agent for Key Trustee Server on page 312). However, if the server does not have access to email, you can use a local system mail address, such as `username@hostname`, where `hostname` is the system hostname and `username` is a valid system user. If you use a local system email address, be sure to regularly monitor the email box.

Managing Key Trustee Server Certificates

Transport Layer Security (TLS) certificates are used to secure communication with Key Trustee Server. By default, Key Trustee Server generates self-signed certificates when it is first initialized. Cloudera strongly recommends using certificates signed by a trusted Certificate Authority (CA).

Generating a New Certificate

1. Generate a new certificate signing request (CSR):

```
$ openssl req -new -key keytrustee_private_key.pem -out new.csr
```

   Replace *keytrustee_private_key.pem* with the filename of the private key. You can reuse the existing private key or generate a new private key in accordance with your company policies. For existing auto-generated self-signed certificates, the private key file is located at `/var/lib/keytrustee/.keytrustee/.ssl/ssl-cert-keytrustee-pk.pem`.

2. Generate a new certificate from the CSR:
   - For a CA-signed certificate, submit the CSR to the CA, and they will provide a signed certificate.
   - To generate a new self-signed certificate, run the following command:

```
$ openssl x509 -req -days 365 -in new.csr -signkey keytrustee_private_key.pem \
-out new_keytrustee_certificate.pem
```

Replacing Key Trustee Server Certificates

Use the following procedure if you need to replace an existing certificate for the Key Trustee Server. For example, you can use this procedure to replace the auto-generated self-signed certificate with a CA-signed certificate, or to replace an expired certificate.

> **Note:** Key Trustee Server supports password-protected private keys, but not password-protected certificates.

1. After Generating a New Certificate on page 315, back up the original certificate and key files:

```
$ sudo cp -r /var/lib/keytrustee/.keytrustee/.ssl /var/lib/keytrustee/.keytrustee/.ssl.bak
```

2. **(CA-Signed Certificates Only)** Provide either the root or intermediate CA certificate:

> **Important:** If you have separate root CA and intermediate CA certificate files, then you must concatenate them into a single file. If you need to convert from a JKS with combined certificate and private key file to a PEM file and separate private key file, refer to Converting JKS Key and Certificate to PEM on page 386.

```
$ sudo mv /path/to/rootca.pem
/var/lib/keytrustee/.keytrustee/.ssl/ssl-cert-keytrustee-ca.pem
```

3. Make sure that the certificate and key files are owned by the `keytrustee` user and group, with file permissions set to `600`:

```
$ sudo chown keytrustee:keytrustee /path/to/new/certificate.pem
/path/to/new/private_key.pem
$ sudo chmod 600 /path/to/new/certificate.pem /path/to/new/private_key.pem
```

4. Update the Key Trustee Server configuration with the location of the new certificate and key files:
   - Using Cloudera Manager:
     1. Go to the Key Trustee Server service.
     2. Click the **Configuration** tab.
     3. Select **Category** > **Security**.
     4. Edit the following properties to specify the location of the new certificate and key files. If the private keys are not password protected, leave the password fields empty.
        - **Active Key Trustee Server TLS/SSL Server Private Key File (PEM Format)**
        - **Active Key Trustee Server TLS/SSL Server Certificate File (PEM Format)**
        - **Active Key Trustee Server TLS/SSL Private Key Password**
        - **Passive Key Trustee Server TLS/SSL Server Private Key File (PEM Format)**
        - **Passive Key Trustee Server TLS/SSL Server Certificate File (PEM Format)**
        - **Passive Key Trustee Server TLS/SSL Private Key Password**
     5. Click **Save Changes** to commit the changes.
   - Using the command line:
     1. Edit `/var/lib/keytrustee/.keytrustee/keytrustee.conf` on the active and passive Key Trustee Server hosts and modify the `SSL_CERTIFICATE` and `SSL_PRIVATE_KEY` parameters as follows:

```
    "SSL_CERTIFICATE": "/path/to/new/certificate.pem",
    "SSL_PRIVATE_KEY": "/path/to/new/private_key.pem"
```

If the private key is password protected, add the following entry:

```
"SSL_PRIVATE_KEY_PASSWORD_SCRIPT": "/path/to/password_script [arguments]"
```

Replace */path/to/password_script [arguments]* with the path to a script (and any necessary command arguments) that returns the password for the private key file. If you do not want to create a script, you can use a simple command, such as `echo -n password`. For example:

```
"SSL_PRIVATE_KEY_PASSWORD_SCRIPT": "/bin/echo -n password"
```

Keep in mind that this method can expose the private key password in plain text to anyone who can view the `/var/lib/keytrustee/.keytrustee/keytrustee.conf` file.

5. Restart Key Trustee Server:
   - **Using Cloudera Manager:** Restart the Key Trustee Server service (**Key Trustee Server service** > **Actions** > **Restart**).
   - **Using the Command Line:** Restart the Key Trustee Server daemon:
     - RHEL 6-compatible: `$ sudo service keytrusteed restart`
     - RHEL 7-compatible: `$ sudo systemctl restart keytrusteed`
6. If you are using the Key Trustee KMS service in Cloudera Manager for HDFS Transparent Encryption on page 241, update the Java KeyStore (JKS) used on the Key Trustee KMS host:
   a. Download the new certificate to the Key Trustee KMS host:

```
$ echo -n | openssl s_client -connect keytrustee01.example.com:11371 \
| sed -ne '/-BEGIN CERTIFICATE-/,/-END CERTIFICATE-/p' > /tmp/keytrustee_certificate.pem
```

   b. Delete the existing keystore entry for *keytrustee01.example.com*:

```
$ keytool -delete -alias key_trustee_alias_name -keystore /path/to/truststore -v
```

   c. Add the new keystore entry for *keytrustee01.example.com*:

```
$ keytool -import -trustcacerts -alias keytrustee01.example.com \
-file /tmp/keytrustee_certificate.pem -keystore /path/to/truststore
```

   d. Restart the Key Trustee KMS service in Cloudera Manager.
7. If you are using Key HSM, update the Key Trustee Server and Key HSM configuration:
   a. Run the `keyhsm trust` command, using the path to the new certificate:

```
$ sudo keyhsm trust /path/to/new/key_trustee_server/cert
```

   b. Run the `ktadmin keyhsm` command, using the `--client-certfile` and `--client-keyfile` options to specify the location of the new certificate file and private key:

```
$ sudo ktadmin keyhsm --server https://keyhsm01.example.com:9090 --client-certfile
/path/to/new/key_trustee_server/cert --client-keyfile
/path/to/new/key_trustee_server/private_key
```

Cloudera Navigator Key HSM

Cloudera Navigator Key HSM allows Cloudera Navigator Key Trustee Server to seamlessly integrate with a hardware security module (HSM). Key HSM enables Key Trustee Server to use an HSM as a root of trust for cryptographic keys, taking advantage of Key Trustee Server's policy-based key and security asset management capabilities while at the same time satisfying existing, internal security requirements regarding treatment of cryptographic materials.

For instructions on installing Key HSM, see Installing Cloudera Navigator Key HSM.

For instructions on configuring Key HSM, continue reading:

Initializing Navigator Key HSM

Before initializing Navigator Key HSM, verify that the HSM is properly configured and accessible from the Key HSM host, and that the HSM client libraries are installed on the Key HSM host:

- **SafeNet Luna**

  Install the SafeNet Luna client. No additional configuration is needed.

- **SafeNet KeySecure**

  Extract the KeySecure client tarball in the Key HSM library directory (`/usr/share/keytrustee-server-keyhsm/`).

- **Thales**

  Install the Thales client service. Copy `nCipherKM.jar`, `jcetools.jar`, and `rsaprivenc.jar` from the installation media (usually located in `opt/nfast/java/classes` relative to the installation media mount point) to the Key HSM library directory (`/usr/share/keytrustee-server-keyhsm/`).

See your HSM product documentation for more information on installing and configuring your HSM and client libraries.

> **Note:** When using an HSM with Key Trustee Server and Navigator Encrypt, encrypting many block devices may exceed the capacity of the HSM. A key is created in the HSM for each encrypted block device, so be sure that your HSM can support your encryption requirements.

To initialize Key HSM, use the `service keyhsm setup` command in conjunction with the name of the target HSM distribution:

```
$ sudo service keyhsm setup [keysecure|thales|luna]
```

For all HSM distributions, this first prompts for the IP address and port number that Key HSM listens on.

> **Important:** If you have implemented Key Trustee Server high availability, initialize Key HSM on each Key Trustee Server.

Cloudera recommends using the loopback address (`127.0.0.1`) for the listener IP address and `9090` as the port number:

```
-- Configuring keyHsm General Setup --
Cloudera Recommends to use 127.0.0.1 as the listener port for Key HSM
Please enter Key HSM SSL listener IP address: [127.0.0.1]127.0.0.1
Will attempt to setup listener on 127.0.0.1
Please enter Key HSM SSL listener PORT number: 9090

validate Port:                                    :[ Successful ]
```

If the setup utility successfully validates the listener IP address and port, you are prompted for additional information specific to your HSM. For HSM-specific instructions, continue to the HSM-Specific Setup for Cloudera Navigator Key HSM on page 319 section for your HSM.

The Key HSM keystore defaults to a strong, randomly-generated password. However, you can change the keystore password in the `application.properties` file:

```
keyhsm.keystore.password.set=yes
```

Then, run the `service keyhsm setup` command with the name of the HSM to which the keystore password applies. You will be prompted to enter the new keystore password, which must be a minimum of six characters in length:

```
$ sudo service keyhsm setup [keysecure|thales|luna]
```

After initial setup, the configuration is stored in the `/usr/share/keytrustee-server-keyhsm/application.properties` file, which contains human-readable configuration information for the Navigator Key HSM server.

> **Important:** The truststore file created during Key HSM initialization must be stored at `/usr/share/keytrustee-server-keyhsm/`. There is no way to change the default location.

For additional details about keystores and truststores, see Understanding Keystores and Truststores on page 193.

HSM-Specific Setup for Cloudera Navigator Key HSM

SafeNet KeySecure

> **Note:** KeySecure was previously named DataSecure, but the Key HSM configuration process is the same for both.

Prerequisites

Before setting up SafeNet KeySecure, be sure to:

- Set the protocol to `NAE-XML`
- Set **Allow Key and Policy Configuration Operations** to `enabled`
- Set **Password Authentication** to `required`
- Disable **Client Certificate Authentication**
- Set **KeySecure Crypto Operations** to `activated`

For additional details about SafeNet KeySecure settings, see the SafeNet KeySecure product documentation.

After entering the Key HSM listener IP address and port, the HSM setup for SafeNet KeySecure prompts for login credentials, the IP address of the KeySecure HSM, and the port number:

```
-- Ingrian HSM Credential Configuration --
Please enter HSM login USERNAME: keyhsm
Please enter HSM login PASSWORD: ******

Please enter HSM IP Address or Hostname: 172.19.1.135
Please enter HSM Port number: 9020
```

If the connection is successful, the following message is displayed:

```
Valid address:                                    :[ Successful ]
```

The KeySecure setup utility then prompts you whether to use SSL:

```
Use SSL? [Y/n] Y
```

If you choose to use SSL, Key HSM attempts to resolve the server certificate, and prompts you to trust the certificate:

```
  [0]         Version: 3
         SerialNumber: 0
             IssuerDN: C=US,ST=TX,L=Austin,O=ACME,OU=Dev,
CN=172.19.1.135,E=webadmin@example.com
           Start Date: Thu Jan 29 09:55:57 EST 2015
           Final Date: Sat Jan 30 09:55:57 EST 2016
            SubjectDN: C=US,ST=TX,L=Austin,O=ACME,OU=Dev,
CN=172.19.1.135,E=webadmin@example.com
           Public Key: RSA Public Key
            modulus: abe4a8dcef92e145984309bd466b33b35562c7f875
                     1d1c406b1140e0584890272090424eb347647ba04b
                     34757cacc79652791427d0d8580a652c106bd26945
                     384b30b8107f8e15d2deba8a4e868bf17bb0207383
                     7cffef0ef16d5b5da5cfb4d3625c0affbda6320daf
                     7c6b6d8adfcb563960fcd1207c059300feb6513408
                     79dd2d929a5b986517531be93f113c8db780c92ddf
                     30f5c8bf2b0bea60359b67be306c520358cc0c3fc3
                     65500d8abeeac99e53cc2b369b2031174e72e6fca1
                     f9a4639e09240ed6d4a73073885868e814839b09d5
                     6aa98a5a1e230b46cdb4818321f546ac15567c5968
                     33be47ef156a73e537fd09605482790714f4a276e5
                     f126f935
    public exponent: 10001

  Signature Algorithm: SHA256WithRSAEncryption
            Signature: 235168c68567b27a30b14ab443388039ff12357f
                       99ba439c6214e4529120d6ccb4a9b95ab25f81b4
                       7deb9354608df45525184e75e80eb0948eae3e15
                       c25c1d58c4f86cb9616dc5c68dfe35f718a0b6b5
                       56f520317eb5b96b30cd9d027a0e42f60de6dd24
                       5598d1fcea262b405266f484143a74274922884e
                       362192c4f6417643da2df6dd1a538d6d5921e78e
                       20a14e29ca1bb82b57c02000fa4907bd9f3c890a
                       bdae380c0b4dc68710deeaef41576c0f767879a7
                       90f30a4b64a6afb3a1ace0f3ced17ae142ee6f18
                       5eff64e8b710606b28563dd99e8367a0d3cbab33
                       2e59c03cadce3a5f4e0aaa9d9165e96d062018f3
                       6a7e8e3075c40a95d61ebc8db43d77e7
       Extensions:
                       critical(false) BasicConstraints: isCa(true)
                       critical(false) NetscapeCertType: 0xc0

Trust this server? [y/N] Y

Trusted server:                                    :[ Successful ]
```

Thales HSM

By default, the Thales HSM client process listens on ports 9000 and 9001. The Cloudera Manager agent also listens on port 9000. To prevent a port conflict, you must change the Thales client ports. Cloudera recommends using ports 11500 and 11501.

To change the Thales client ports, run the following commands:

```
$ sudo /opt/nfast/bin/config-serverstartup --enable-tcp --enable-privileged-tcp
--port=11500 --privport=11501
$ sudo /opt/nfast/sbin/init.d-ncipher restart
```

To configure Key HSM to use the modified port, edit the `/usr/share/keytrustee-server-keyhsm/start.sh` file and add the `-DNFAST_SERVER_PORT=11500` Java system property. For example:

```
java -classpath "*:/usr/safenet/lunaclient/jsp/lib/*:/opt/nfast/java/classes/*"
-Djava.library.path=/usr/safenet/lunaclient/jsp/lib/ -DNFAST_SERVER_PORT=11500
com.cloudera.app.run.Program $@
```

Before completing the Thales HSM setup, run the `nfkminfo` command to verify that the Thales HSM is properly configured:

```
$ sudo /opt/nfast/bin/nfkminfo
  World generation  2
  state       0x17270000 Initialised Usable Recovery !PINRecovery !ExistingClient
              RTC  NVRAM FTO !AlwaysUseStrongPrimes SEEDebug
```

If `state` reports `!Usable` instead of `Usable`, configure the Thales HSM before continuing. See the Thales product documentation for instructions.

After entering the Key HSM listener IP address and port, the HSM setup for Thales prompts for the OCS card password:

```
Please enter the OCS Card Password (input suppressed):

Configuration saved in 'application.properties' file
Configuration stored in: 'application.properties'. (Note: You can also use service keyHsm
 settings to quickly view your current configuration)
```

Luna HSM

> **Important:** If you have implemented Key Trustee Server high availability, ensure that the Luna client on each Key Trustee Server is configured with access to the same partition. See the Luna product documentation for instructions on configuring the Luna client.

Before completing the Luna HSM setup, run the `vtl verify` command (usually located at `/usr/safenet/lunaclient/bin/vtl`) to verify that the Luna HSM is properly configured.

After entering the Key HSM listener IP address and port, the HSM setup for Luna prompts for the slot number and password:

```
-- Configuring SafeNet Luna HSM --
Please enter SafeNetHSM Slot Number: 1
Please enter SafeNet HSM password (input suppressed):
Configuration stored in: 'application.properties'. (Note: You can also use service keyHsm
 settings to quickly view your current configuration)
Configuration saved in 'application.properties' file
```

See the Luna product documentation for instructions on configuring your Luna HSM if you do not know what values to enter here.

Validating Key HSM Settings

After the setup completes, the Key HSM configuration is stored in `/usr/share/keytrustee-server-keyhsm/application.properties`.

You can view these settings using the `service keyhsm settings` command:

```
$ sudo service keyhsm settings

  # keyHsm Server Configuration information:
  keyhsm.management.address : 172.19.1.2
  keyhsm.server.port : 9090
  keyhsm.management.port : 9899
```

```
keyhsm.service.port : 19791
keyhsm.hardware : ncipher

# Module OCS Password
thales.ocs_password :
  GIqhXDuZsj1Oet137Lb+f+tqkYvKYDm/8StefpNqZWwlB+LfSYlB4eHd
  endtYJio8qLjjbT+e7j2th5xf8O9t8FwfVguuyFW+6wdD
  uNGvse1LY/itCwqF0ScMlB1Mnz4010xqC6ylPW7l+0JjjkkqqM5gJJbl8lsQFFaIGVM/pY=
```

These settings can be manually configured by modifying the `application.properties` file, with the exception of any passwords. These are encrypted by design, and can only be changed by re-running the setup utility.

Verifying Key HSM Connectivity to HSM

To verify Hardware Security Module (HSM) operations using Key HSM, run the following command on the Key Trustee Server host (which should also be the Key HSM host as described in Installing Cloudera Navigator Key HSM):

```
$ curl -k https://keytrustee01.example.com:11371/test_hsm
```

If Key HSM operations to the HSM are successful, the command returns output similar to the following:

```
"Sample Key TEST_HELLO_DEPOSIT2016-06-03-072718 has been created"
```

You must run this command from the Key Trustee Server host. If you run it from a different host, the command returns an HTTP `403` error code.

If the command returns an HTTP `405` error code, restart Key Trustee Server and try again.

> **Note:** If you are using the `test_hsm` script to verify that the Key Hardware Security Module (Key HSM) has successfully integrated with the Key Trustee Server, or to verify that the Key HSM is connected to HSM, *and* the Key Trustee Server private key file is password-protected, then the verification may fail. This can occur even if the integration is successful or connected.
>
> If this occurs, then create a key through Hadoop for the test.

Managing the Navigator Key HSM Service

Use the `keyhsm` service for all basic server operations:

```
$  sudo service keyhsm
keyHsm service usage:
    setup <hsm name> - set up a new connection to an HSM
    trust <path>     - add a trusted client certificate
    validate         - validate that Key HSM is properly configured
    settings         - display the current server configuration
    start            - start the Key HSM proxy server
    status           - show the current Key HSM server status
    shutdown         - force Key HSM server to shut down
    debug            - start the Key HSM proxy server in debug mode
    reload           - reload the server (without shutdown)
```

The `reload` command causes the application to restart internal services without ending the process itself. If you want to stop and start the process, use the `restart` command.

Logging and Audits

The Navigator Key HSM logs contain all log and audit information, and by default are stored in the `/var/log/keyhsm` directory.

You can configure the maximum log size (in bytes) and maximum number of log files to retain by adding or editing the following entries in `/usr/share/keytrustee-server-keyhsm/application.properties`:

```
keyhsm.log.size = 100000000
keyhsm.roll.size = 3
```

The values used in this example are the default values, and are used if these parameters are not set.

To enable debug logging, add the debug parameter to the start command:

```
$ sudo service keyhsm start debug
```

> **Note:** You cannot start Key HSM in debug mode using the `systemctl` command on RHEL 7-compatible OS. You must use the `service` command.

This enables debug logging until the service is restarted without the debug parameter.

Key Naming Convention

To ensure you can manage keys (for example, delete a key), you need to understand the naming convention for keys. Keys adhere to the following naming convention: `handle name-uuid-date`, which means if you know the key name and date, you can make modifications to it.

The following example shows the key nomenclature in the context of a key list query on Luna HSM:

```
[root@user 64]# cmu list
Please enter password for token in slot 1 : **********
handle=220
label=key1-3T17-YYdn-2015-07-23
handle=806
label=key2-CMYZ-8Sym-2015-07-23
handle=108
label=key3-qo62-XQfx-2015-07-23
handle=908
label=key2-CMYZ-8Sym-2015-07-23--cert0
handle=614
label=key3-qo62-RWz0-2015-07-23--cert0
handle=825
label=key1-3T11-YYdz-2015-07-23--cert0
```

Integrating Key HSM with Key Trustee Server

Using a hardware security module with Navigator Key Trustee Server requires Key HSM. This service functions as a driver to support interactions between Navigator Key Trustee Server and the hardware security module, and it must be installed on the same host system as Key Trustee Server. The steps below assume that both Key HSM and Key Trustee Server are set up and running. See Installing Cloudera Navigator Key HSM for details. Integrating Key HSM and Key Trustee Server involves the following steps:

1. Check Existing Key Names (for existing Key Trustee Server users only)
2. Establish Trust from Key HSM to Key Trustee Server
3. Integrate Key HSM and Key Trustee Server

Check Existing Key Names

During the process detailed below, you are prompted to migrate any existing keys from the Key Trustee Server to the HSM.

> **Warning:** Migration fails if any existing keys do not adhere to the constraints.

Successful migration depends on the existing keys conforming to the following constraints:

- Key names can begin with alpha-numeric characters only
- Key names can include only these special characters:
  - Hyphen -
  - Period .
  - Underscore _

If any existing key names in Key Trustee Server do not meet the requirements listed above, the migration fails. To prepare for migration, check your key names and do the following if any of them are non-conforming:

- Decrypt any data using the non-conforming key
- Create a new key, named per the requirements
- Re-encrypt the data using the new key

After this, the migration from Key Trustee Server to the HSM will succeed during the process below.

> **Important:** Keys are not available during migration, so you should perform these tasks during a maintenance window.

Establish Trust from Key HSM to Key Trustee Server

This step assumes that Key Trustee Server has a certificate for TLS (wire) encryption as detailed in Managing Key Trustee Server Certificates on page 315. Before you can run the commands in the steps below, Key HSM service must explicitly trust the Key Trustee Server certificate (presented during TLS handshake). To establish this trust, run the following command:

```
$ sudo keyhsm trust /path/to/key_trustee_server/cert
```

The `/path/to/key_trustee_server/cert` in this command (and in the commands below) depends on whether the Key Trustee Server uses the **default certificate** (created by default during install), or uses a **custom certificate** (obtained from a commercial or internal CA). The two alternate paths are shown in the table below. The custom path is a common example but may differ from that shown.

Default	Custom
/var/lib/keytrustee/.keytrustee/.ssl/ssl-cert-keytrustee.pem	`/etc/pki/cloudera/certs/cert-file.crt`
/var/lib/keytrustee/.keytrustee/.ssl/ssl-cert-keytrustee-pk.pem	`/etc/pki/cloudera/private/private-key.key`

> **Note:** The system requires TLS and Kerberos authentication throughout the system for security reasons. Connections attempted over SSL (1 through 3) and connections from untrusted clients are immediately terminated to prevent POODLE (Padding Oracle On Downgraded Legacy Encryption) exploits. See the Cloudera Security Bulletin for more information.

Integrate Key HSM and Key Trustee Server

The steps below assume that both Key HSM and the Key Trustee Server are on the same host system, as detailed in Installing Cloudera Navigator Key HSM. These steps invoke commands on the Key HSM service and the Key Trustee Server, and they must be run on the host—they cannot be run remotely from another host.

1. Ensure the Key HSM service is running:

```
$ sudo service keyhsm start
```

2. Establish trust from Key Trustee Server to Key HSM specifying the path to the private key and certificate (Key Trustee Server is a client to Key HSM). This example shows how to use the `--client-certfile` and `--client-keyfile` options to specify the path to non-default certificate and key:

```
$ sudo ktadmin keyhsm --server https://keyhsm01.example.com:9090 \
--client-certfile /etc/pki/cloudera/certs/mycert.crt \
--client-keyfile /etc/pki/cloudera/certs/mykey.key --trust
```

For a password-protected Key Trustee Server private key, add the `--passphrase` argument to the command and enter the password when prompted:

```
$ sudo ktadmin keyhsm --passphrase \
--server https://keyhsm01.example.com:9090 \
--client-certfile /etc/pki/cloudera/certs/mycert.crt \
--client-keyfile /etc/pki/cloudera/certs/mykey.key --trust
```

> **Note:** The preceding commands also create a certificate file for the Key HSM that is used by the Key Trustee Server. This certificate file is stored in `/var/lib/keytrustee/.keytrustee/.ssl/ssl-cert-keyhsm.pem`.

3. Restart Key Trustee Server:
   - **Using Cloudera Manager:** Restart the Key Trustee Server service (**Key Trustee Server service** > **Actions** > **Restart**).
   - **Using the Command Line:** Restart the Key Trustee Server daemon:
     - RHEL 6-compatible: `$ sudo service keytrusteed restart`
     - RHEL 7-compatible: `$ sudo systemctl restart keytrusteed`

4. Verify connectivity between the Key HSM service and the HSM:

```
$ curl -k https://keytrustee01.example.com:11371/test_hsm
```

> **Important:** You must perform the connection verification between Key HSM and the HSM for all Key Trustee Server hosts.

Successful connection and test of operations returns output like the following:

```
"Sample Key TEST_HELLO_DEPOSIT2016-06-03-072718 has been created"
```

> **Note:** If you are using the `test_hsm` script to verify that the Key Hardware Security Module (Key HSM) has successfully integrated with the Key Trustee Server, or to verify that the Key HSM is connected to HSM, *and* the Key Trustee Server private key file is password-protected, then the verification may fail. This can occur even if the integration is successful or connected.
>
> If this occurs, then create a key through Hadoop for the test.

See Verifying Key HSM Connectivity to HSM on page 322 for more information about the validation process.

Cloudera Navigator Encrypt

Cloudera Navigator Encrypt transparently encrypts and secures data at rest without requiring changes to your applications and ensures there is minimal performance lag in the encryption or decryption process. Advanced key management with Cloudera Navigator Key Trustee Server on page 299 and process-based access controls in Navigator Encrypt enable organizations to meet compliance regulations and ensure unauthorized parties or malicious actors never gain access to encrypted data.

For instructions on installing Navigator Encrypt, see Installing Cloudera Navigator Encrypt.

For instructions on configuring Navigator Encrypt, continue reading:

Registering Cloudera Navigator Encrypt with Key Trustee Server

Prerequisites

Functioning Navigator Key Trustee Server

After Installing Cloudera Navigator Encrypt on a host, you must register the host with Navigator Key Trustee Server. If you have not yet installed Navigator Key Trustee Server, see Installing Cloudera Navigator Key Trustee Server for instructions.

Key Trustee Server Organization

To register with Key Trustee Server, you must have an existing organization. See Managing Key Trustee Server Organizations on page 312 for information on creating and viewing organizations on a Key Trustee Server.

Master Password

The Master Key is the primary Navigator Encrypt administrator access code and is configured by the Navigator Encrypt administrator during installation. The Master Key can take any one of three different forms:

- If you choose a passphrase (single), it must be between 15 and 32 characters long.
- If you choose passphrase (dual), both must be between 15 and 32 characters long.
- If you choose the RSA option, enter a path to the RSA key file, and if it has RSA passphrase, enter it for this private key.

> **Warning:** It is ***extremely*** important that you keep your master password secret and safe. In the event that you lose your master password, **you will never be able to recover it**, leaving your encrypted data **irretrievably locked away**.

Registering with Key Trustee Server

After Installing Cloudera Navigator Encrypt on a host, you must register the host with Navigator Key Trustee Server to be able to encrypt and decrypt data. The following section lists the command options for registering your Navigator Encrypt client.

> **Note:** Do not run Navigator Encrypt commands with the `screen` utility.

If the TLS certificate is signed by an internal CA that is not publicly recognized, then you must add the root certificate to the host certificate truststore of each Navigator Encrypt client. For details, see Setting Up TLS for Navigator Encrypt Clients.

Example command:

```
$ sudo navencrypt register --server=https://keytrustee01.example.com:11371
--passive-server=https://keytrustee02.example.com:11371 --org=your_keytrustee_org
--auth=org_auth_token
```

Table 32: Registration Options

Command Option	Explanation
`--clientname=my_client_name`	User-defined unique name for this client to be used for administration and reports. You can verify your client name in the `/etc/navencrypt/keytrustee/clientname` file.
`--server=https://keytrustee01.example.com:11371`	Target Active Key Trustee Server for key storage. Replace *keytrustee01.example.com:11371* with the hostname and port of the Active Key Trustee Server. The default port is `11371`.
`--passive-server=https://keytrustee02.example.com:11371`	Target Passive Key Trustee Server for key storage. Replace *keytrustee02.example.com:11371* with the hostname and port of the Passive Key Trustee Server. The default port is `11371`.
`--org=your_keytrustee_org`	Key Trustee organization name configured by the Key Trustee Server administrator
`--auth=org_auth_token`	Organization authorization token, a pre-shared secret by the Navigator Key Trustee Server administrator
`--skip-ssl-check`	Skip SSL certificate verification. Use with self-signed certificates on the Navigator Key Trustee Server
`--trustee`	Add trustees for retrieval of the master key
`--votes`	Configure voting policy for trustees
`--recoverable`	Master Key will be uploaded without encrypting it with your local GPG Navigator Key Trustee
`--scheme "<scheme>"`	Key Trustee Server scheme that Navigator Encrypt uses for public key operations. Specify `"http"` or `"https"`.
`--port`	Key Trustee Server port that Navigator Encrypt uses for public key operations.

Registering with Previous Versions of Key Trustee Server

By default, new installations of Navigator Key Trustee Server 5.4.0 use a single HTTPS port for key storage and public key operations. Previous versions and upgrades use separate ports for key storage and public key operations. For backward compatibility, Navigator Encrypt 3.7.0 introduces the `--scheme` and `--port` parameters for the `navencrypt register` command.

For example, to register a version 3.7.0 Navigator Encrypt client with a version 3.8.0 Key Trustee Server using HTTPS over port 443 for key storage and HTTP over port 80 for public key operations, run the following command:

```
$ sudo navencrypt register --server=https://keytrustee.example.com:443
--org=key_trustee_org --auth=auth_token --scheme "http" --port 80
```

Navigator Encrypt versions lower than 3.7.0 do not support the `--scheme` and `--port` parameters. For these versions of Navigator Encrypt, you must ensure that the Key Trustee Server is configured to use port 443 (HTTPS) for key storage and port 80 (HTTP) for public key operations.

Navigator Encrypt versions lower than 3.8.0 do not support the `--passive-server` parameter.

Updating Key Trustee Server Ports

The `navencrypt register` command does not provide the ability to change the ports for existing registrations. If the Key Trustee Server ports are changed, you must update `/etc/navencrypt/keytrustee/ztrustee.conf` with the new port and scheme parameters (HKP_PORT and HKP_SCHEME, respectively).

For example, see the following `ztrustee.conf` excerpt from a registered client that has been upgraded to Navigator Encrypt 3.7.0:

```
{
        "LOCAL_FINGERPRINT":    "2048R/182AAA838DC300AC334258D8E7F299BFB68A6F6F",
        "REMOTES":      {
                "kts01.example.com": {
                        "REMOTE_FINGERPRINT":
"4096R/AF6400E12DC149799CA8CE6BF1604C34D830DE20",
                        "REMOTE_SERVER":        "https://kts01.example.com",
                        "DEFAULT":      true,
                        "SSL_INSECURE": false,
                        "PROTOCOL":     "json-encrypt"
                }
        }
}
```

In this example, the Key Trustee Server (`keytrustee.example.com`) is using the default configuration of port 443 (HTTPS) for key storage and port 80 (HTTP) for public key operations.

If the Key Trustee Server is then updated to use port 11371 (HTTPS) for both key storage and public key operations, you must update `ztrustee.conf` as follows (changes in **bold**):

```
{
        "LOCAL_FINGERPRINT":    "2048R/182AAA838DC300AC334258D8E7F299BFB68A6F6F",
        "REMOTES":      {
                "kts01.example.com": {
                        "REMOTE_FINGERPRINT":
"4096R/AF6400E12DC149799CA8CE6BF1604C34D830DE20",
                        "REMOTE_SERVER":        "https://kts01.example.com:11371",
                        "HKP_PORT":     11371,
                        "HKP_SCHEME":   "https",
                        "DEFAULT":      true,
                        "SSL_INSECURE": false,
                        "PROTOCOL":     "json-encrypt"
                }
        }
}
```

Updating Navigator Encrypt for High Availability Key Trustee Server

If you registered a Navigator Encrypt client with a standalone Key Trustee Server, and then configured high availability for Key Trustee Server, you can edit `/etc/navencrypt/keytrustee/ztrustee.conf` to enable the client to take advantage of the high availability features. The following example shows the contents of `ztrustee.conf` after adding the required REMOTE_SERVERS entry (changes in **bold**):

```
{
        "LOCAL_FINGERPRINT":    "2048R/182AAA838DC300AC334258D8E7F299BFB68A6F6F",
        "REMOTES":      {
                "kts01.example.com": {
                        "REMOTE_FINGERPRINT":
"4096R/AF6400E12DC149799CA8CE6BF1604C34D830DE20",
                        "REMOTE_SERVER":        "https://kts01.example.com:11371",
                        "HKP_PORT":     11371,
```

```
                        "HKP_SCHEME":   "https",
                        "DEFAULT":      true,
                        "REMOTE_SERVERS":       ["https://kts01.example.com:11371",
"https://kts02.example.com:11371"],
                        "SSL_INSECURE": true,
                        "PROTOCOL":     "json-encrypt"
                }
        }
}
```

Configuration Files

The installer creates the `/etc/navencrypt` directory. All configuration settings are saved in this directory. **Do not** delete any file from `/etc/navencrypt`. These files provide the necessary information for the Navigator Encrypt application to function properly.

> **Warning:** Perform backups of encrypted data, mount-points, and Navigator Encrypt configuration directories on a regular basis. To do this, ensure you have a backup of `/etc/navencrypt`. **Failure to backup this directory will make your backed up encrypted data unrecoverable in the event of data loss.**

Change Master Key by UUID

It is possible to re-use a previously used Master Key by its UUID. For example, if you currently have a Master key with a single passphrase, you can see the corresponding Navigator Key Trustee UUID in the `/etc/navencrypt/control` file:

```
$ cat /etc/navencrypt/control
{
 "app": {
  "name": "navencrypt",
  "version": "3.5"
 },
 "keys": {
  "master": {
   "type": "single-passphrase",
   "uuid": "qMAKRMdk4HVbhfzR79cp9w92YBmNHJ5nSLhfd8ZVo6L"
  },
  "targets": []
 }
}
```

> **Note:** If the control file is accidentally deleted, you can restore it using the `navencrypt control --restore-control-file` command.

You can copy the UUID (`qMAKRMdk4HVbhfzR79cp9w92YBmNHJ5nSLhfd8ZVo6L` in this example) and run `navencrypt key --change` with option `--new-master-key-uuid` to change a Master Key by using its UUID only:

```
$ sudo navencrypt key --change
--new-master-key-uuid=qMAKRMdk4HVbhfzR79cp9w92YBmNHJ5nSLhfd8ZVo6L
>> Type your OLD Master key
Type MASTER passphrase 1:
Type MASTER passphrase 2:
Verifying Master Key against Navigator Key Trustee (wait a moment)...
OK
Changing Master key (wait a moment)...
 * Setting up EXISTING MASTER key...
 * Uploading CONTROL content...
 * Re-encrypting local keys...
Master key successfully changed.
```

> **Note:** The `navencrypt key` command fails if no volumes are encrypted or the kernel module is not loaded.

Preparing for Encryption Using Cloudera Navigator Encrypt

Before you can encrypt data, you must prepare a storage repository to hold the encrypted data and a mount point through which to access the encrypted data. The storage repository and mount point must exist before encrypting data using the `navencrypt-move` command.

Data stored and retrieved from the repository is encrypted and decrypted transparently.

Cloudera Navigator Encrypt *does not* support:

- Encrypting a directory that contains or is contained within a mount point for another service (including Navigator Encrypt and NFS). See Encrypting Data on page 337 for more information.
- Encrypting immutable files or directories containing immutable files.
- Installation or use in `chroot` environments, including creating `chroot` environments within an encrypted directory.
- Encrypting HDFS data files.

Navigator Encrypt Commands

> **Note:** Do not run Navigator Encrypt commands with the `screen` utility.

The following table lists the commands used to encrypt data:

Table 33: Navigator Encrypt Commands

Command	Description
`navencrypt`	Manage, update, and verify your data.
`navencrypt-prepare`	Prepare your system for encryption by creating mount-points and specifying storage.
`navencrypt-prepare --undo`	Remove a mountpoint that is no longer in use.
`navencrypt-move`	Encrypt/decrypt your data to/from the encrypted filesystem.
`navencrypt-profile`	Generate process profile information in JSON format.
`navencrypt-module-setup`	Build or rebuild the Navigator Encrypt kernel module.

Preparing for Encryption

> **Note:** When using an HSM with Key Trustee Server and Navigator Encrypt, encrypting many block devices may exceed the capacity of the HSM. A key is created in the HSM for each encrypted block device, so be sure that your HSM can support your encryption requirements.

To get an in-depth look at the details behind the `navencrypt-prepare` command, or to use a unique configuration, use the interactive prompt by executing `navencrypt-prepare` with no options. This launches an interactive console that guides you through the following operations:

- Creating internal encryption keys
- Registering internal keys in Navigator Key Trustee

- Registering mount point in `/etc/navencrypt/ztab`
- Mounting current mount point
- Establishing encryption method (`dm-crypt` for devices, `ecryptfs` for directories)

Using the console, you can choose how you want your data stored and accessed. Navigator Encrypt offers two different types of encryption:

- Block-level encryption with `dm-crypt`: Protect your data by encrypting the entire device. This option enables full disk encryption and is optimized for some system configurations. Block-level encryption can be used with logical devices such as a loop device.
- File-level encryption with `ecryptfs`: Protect your data by mounting an encrypted filesystem on top of an existing one. Enables transparent access to encrypted data without modifying your storage.

> **Note:** As of August 2015, Filesystem-level encryption using eCryptfs is deprecated. For new installations, use block-level encryption. For existing installations using eCryptfs, see Migrating eCryptfs-Encrypted Data to dm-crypt on page 338 for instructions on migrating data encrypted using eCryptfs to use dm-crypt.

See Block-Level Encryption with dm-crypt on page 331 and Filesystem-Level Encryption with eCryptfs on page 334 for more information.

To prepare for encryption, you must set a location to store the encrypted data and a mount point through which to access the data. The storage location and mount point must be created before encrypting data.

> **Note:** If you are performing a file system check as part of your preparation work, then note that the crypto device must be mapped and active. Also, be aware that if you execute `fsck` in force mode (`-f`) , there is a risk of data loss. If the force operation fails, it could cause file system corruption at the device header.

Block-Level Encryption with dm-crypt

> **Note:** For best performance, Cloudera strongly recommends using block encryption with dm-crypt. See Migrating eCryptfs-Encrypted Data to dm-crypt on page 338 for instructions on migrating data encrypted using eCryptfs to use dm-crypt.

When choosing block-level encryption during the interactive console, you must specify two parameters:

1. The first parameter is the block device that you want to store the encrypted file system in. Because this device stores all of the encrypted data, it must be as large as or larger than the target data. The device must exist and be empty. Supported storage devices are:
   - Physical block devices (for example, a disk device)
   - Virtual block devices (for example, a block device created by LVM)
   - Loop devices (see Block-Level Encryption with a Loop Device on page 332 for instructions on creating a loop device)

   > **Important:** If the block device to be used for encryption was previously used by the host, entries for it must be removed from the file `/etc/fstab` before running the `navencrypt-prepare` command.

2. The second parameter is the mount point for the encrypted file system. This is the location where you can access the encrypted data stored in the first parameter. The mount point must already exist. It is not created by the `navencrypt-prepare` command.

The entire device in the first parameter is used for encrypted data.

> **Note:** Do not manually unmount the encryption mount point (for example, using `umount`). If you do so, you must manually close the `dm-crypt` device using the following procedure:
>
> 1. Run `dmsetup table` to list the `dm-crypt` devices.
> 2. Run `cryptsetup luksClose <device_name>` to close the device for the unmounted mount point.

After choosing these two parameters and following the interactive console (discussed further in Preparing for Encryption on page 330), you are ready to encrypt your data. The following example shows successful output from a `navencrypt-prepare` command using dm-crypt for block-level encryption:

```
$ sudo /usr/sbin/navencrypt-prepare /dev/sda1 /mnt/dm_encrypted
Type MASTER passphrase:
Encryption Type:  dmCrypt (LUKS)
Cipher:        aes
Key Size:  256
Random Interface: /dev/urandom
Filesystem:    ext4
Verifying MASTER key against Navigator Key Trustee (wait a moment) ... OK
Generation Encryption Keys with /dev/urandom       ... OK
Preparing dmCrypt device (--use-urandom)               ... OK
Creating ext4 filesystem                               ... OK
Registering Encryption Keys (wait a moment)        ... OK
Mounting /dev/sda1                                    ... OK
```

Block-Level Encryption with a Loop Device

A block-level encrypted device can be a physical device or a storage space treated as a device. See Migrating eCryptfs-Encrypted Data to dm-crypt on page 338 for instructions on migrating data encrypted using eCryptfs to use dm-crypt with a loop device.

To configure a loop device, use the dd command to create a storage space:

> **Warning:** The space for the loop device is pre-allocated. After the loop device is created, the size cannot be changed. Make sure you are allocating enough storage for the current encrypted data as well as any future data.
>
> If your disks are mounted individually with a filesystem on each disk, and your storage needs exceed the capacity of a single disk, you can create a loop device on each disk for which you want to allocate space for encrypted data. If you have consolidated storage (for example, using LVM), you can create a single loop device or multiple devices.

```
$ sudo dd if=/dev/zero of=/dmcrypt/storage bs=1G count=500
```

The dd command above creates a 500 GB file. Modify the `bs` and `count` values to generate the required file size.

After generating the file, run `losetup -f` to view unused loop devices. Use the available loop device with the `navencrypt-prepare -d` command, demonstrated below.

Specifically for loop devices, the `-d` parameter enables Navigator Encrypt to manage the loop device association. You no longer need to use the `losetup` command to associate the file with the loop device, and the loop device is automatically prepared at boot. For RHEL 7-compatible OS, you must run the following commands to ensure that a loop device is available at boot:

```
$ sudo bash -c 'echo "loop" > /etc/modules-load.d/loop.conf'
$ sudo bash -c 'echo "options loop max_loop=8" > /etc/modprobe.d/loop_options.conf'
```

> **Warning:** Loop devices are not created by Navigator Encrypt. Instead, Navigator Encrypt assigns a datastore to a loop device when the `navencrypt-prepare --datastore` option is used. So, it is up to the system administrator to create persistent /dev/loopX devices, which are required to prepare a virtual block device. If the loop device being prepared is not persistent, then Navigator Encrypt will not mount the device upon a reboot.

The data storage directory name (`/dmcrypt/storage` in the previous example) must contain only alphanumeric characters, spaces, hyphens (`-`), or underscores (_). Other special characters are not supported.

The following example shows the output from a successful command:

```
$ losetup -f
/dev/loop0
$ sudo navencrypt-prepare -d /dmcrypt/storage /dev/loop0 /dmcrypt/mountpoint
Type MASTER passphrase:

Encryption Type:  dmCrypt (LUKS)
Cipher:           aes
Key Size:         256
Random Interface: OpenSSL
Filesystem:       ext4
Options:

Verifying MASTER key against KeyTrustee (wait a moment)   ... OK
Generation Encryption Keys with OpenSSL                   ... OK
Assigning '/dev/loop0'->'/dmcrypt/storage'                ... OK
Preparing dmCrypt device                                  ... OK
Creating ext4 filesystem                                  ... OK
Registering Encryption Keys (wait a moment)               ... OK
Mounting /dev/loop0                                       ... OK
```

For upgraded Navigator Encrypt clients that already use loop devices, you can enable Navigator Encrypt to manage the loop device file association (instead of configuring the system to run the `losetup` command at boot) by adding the `nav_datastore` option to the entry in `/etc/navencrypt/ztab`. For example:

```
# <target mount-dir>        <source device>      <type>    <options>
/dmcrypt/mountpoint     /dev/loop0      luks
key=keytrustee,nav_datastore='/dmcrypt/storage'
```

> **Important:** Use caution when editing the `/etc/navencrypt/ztab` file. Entries are tab-separated (not space-separated). The `ztab` file must not contain empty lines.

After you have created the loop device, continue with the instructions in Block-Level Encryption with dm-crypt on page 331.

Navigator Encrypt and Device UUIDs

Navigator Encrypt has always prepared and identified devices simply using a device name, such as `/dev/sdb1` or `/dev/loop0`. However, we know that using a device name or label could lead to a conflict and impact system operations.

Navigator Encrypt also supports preparing devices using a UUID, in addition to device name. This UUID is simply a symbolic link to the actual device, and is created when preparing a device with Navigator Encrypt during a `navencrypt-prepare` operation.

The advantage of using a device UUID is that if a device's name changes, the UUID associated with that device does *not* change. To ensure that Navigator Encrypt recognizes devices even when the device name changes, enter the command:

```
navencrypt-prepare --use-uuid /dev/sda1 /mountpoint
```

To unprepare (ensure the device UUID is included), enter either of the following commands:

```
navencrypt-prepare --undo-force /dev/disk/by-uuid/3a602a15-11f7-46ac-ae98-0a51e1b25cf9
navencrypt-prepare --undo /dev/disk/by-uuid/3a602a15-11f7-46ac-ae98-0a51e1b25cf9
```

Note: While the device name is still used in the `navencrypt-prepare` statement, rest assured that Navigator Encrypt handles the device by the UUID, which is calculated and used for mount during boot. The device name is used for convenience so that you do not have to explicitly input the UUID in the command. Ultimately, Navigator Encrypt handles the device via the device UUID rather than the device name.

Important: UUID device support does *not* include loop devices; rather, it only applies to physical devices. When preparing loop devices with Navigator Encrypt, always use the device name.

Filesystem-Level Encryption with eCryptfs

Note: As of August 2015, Filesystem-level encryption using eCryptfs is deprecated. For best performance, Cloudera strongly recommends using Block-Level Encryption with dm-crypt on page 331 where possible. See Migrating eCryptfs-Encrypted Data to dm-crypt on page 338 for instructions on migrating data encrypted using eCryptfs to use dm-crypt.

RHEL 7 does not support eCryptfs. For new installations on RHEL 7, you must use Block-Level Encryption with dm-crypt on page 331. If you are planning on upgrading to RHEL 7 and are currently using eCryptfs, migrate to dm-crypt before upgrading.

When choosing file-level encryption during the interactive console, you must specify two parameters:

1. The first parameter is the storage directory you want to store the encrypted file system in. Because this directory will hold all of the encrypted data, it must be as large as or larger than the target data.
2. The second parameter is the mount point for the encrypted file system. This is the location where you can access the encrypted data stored in the location identified by the first parameter.

While the data is technically stored at the location identified by the first parameter, you can only access the data from the mount point identified by the second parameter. Consider this when choosing where to mount your data.

After choosing these two parameters and following the interactive console (discussed further in Preparing for Encryption on page 330), you are ready to encrypt your data.

In the following example, we will use the directory `/navencrypt/encrypted-storage` for the encrypted storage and `/navencrypt/mount-point` for the mount point. If you have specific space/partition requirements, you can select a different directory, although Cloudera highly recommends that you place the encrypted directory on the same partition as the data you are planning to encrypt.

The syntax for the prepare command is as follows:

```
$ sudo navencrypt-prepare <data_storage_directory> <mount_point>
```

When specifying the storage path and the mount point path, *do not* use a trailing / in the path names. Both directories must exist prior to running the `navencrypt-prepare` command. They are not automatically created.

To create the encrypted partition, create the mount point and storage directories, and then use the `navencrypt-prepare` utility:

```
$ sudo mkdir -p /navencrypt/encrypted-storage /navencrypt/mount-point
$ sudo navencrypt-prepare /navencrypt/encrypted-storage /navencrypt/mount-point
```

For RHEL 7, run the following command after the `navencrypt-prepare` command completes:

```
$ sudo systemctl start navencrypt-mount
```

To demonstrate the difference between the two directories, this example uses different directories for the encrypted storage and the mount point. It is also possible to store and access the data using the same directory.

To see the effects of these commands, run `df -h`. This command displays the partition information about your system. You should see an `ecryptfs` partition located at `/navencrypt/encrypted-storage`, and mounted at `/navencrypt/mount-point`.

After you have successfully prepared a client for encryption, you can encrypt and decrypt data using the commands described in Encrypting and Decrypting Data Using Cloudera Navigator Encrypt on page 336.

Undo Operation

Navigator Encrypt 3.5 and higher supports a new command option, `navencrypt-prepare --undo`. This command reverses the operations from the regular `navencrypt-prepare` command by removing the device from Navigator Encrypt control and removing registered encryption keys.

The only parameter of the undo operation is the storage device used to store the encrypted file system (not the mount point). Here is an example showing `navencrypt-prepare` and `navencrypt-prepare --undo` operations:

```
$ sudo navencrypt-prepare /path/to/storage /path/to/mountpoint
Type MASTER passphrase:

Encryption Type:  eCryptfs
Cipher:           aes
Key Size:         256
Random Interface: OpenSSL
Filesystem:       ext4
Options:

Verifying MASTER key against Navigator Key Trustee (wait a moment)     ... OK
Generation Encryption Keys with OpenSSL                   ... OK
Registering Encryption Keys (wait a moment)               ... OK
Mounting /path/to/mountpoint                              ... OK
$ sudo navencrypt-prepare --undo /path/to/storage
Type MASTER passphrase:
Verifying MASTER key against Navigator Key Trustee (wait a moment)     ... OK
Umounting /path/to/mountpoint                             ... OK
```

Pass-through Mount Options for `navencrypt-prepare`

Navigator Encrypt 3.5 and higher provides the ability to specify options to pass to the `mount` command that is executed during `/etc/init.d/navencrypt-mount start` (`systemctl start navencrypt-mount` on RHEL 7). These options are specified with the `-o` option when preparing a mountpoint with the `navencrypt-prepare` command.

The following shows an example `navencrypt-prepare` command output when passing mount options with the `-o` option:

```
$ sudo navencrypt-prepare -o discard,resize /mnt/t2 /mnt/t2
Type MASTER passphrase:

Encryption Type:  eCryptfs
Cipher:           aes
Key Size:         256
Random Interface: OpenSSL
Filesystem:       ext4
Options:          discard,resize

Verifying MASTER key against Navigator Key Trustee(wait a moment)     ... OK
Generation Encryption Keys with OpenSSL                               ... OK
Registering Encryption Keys (wait a moment)                           ... OK
Mounting /mnt/t2                                                      ... OK
```

You can verify the results by viewing the `/etc/navencrypt/ztab` file:

```
$ cat /etc/navencrypt/ztab
/mnt/t2 /mnt/t2 ecryptfs key=keytrustee,cipher=aes,keysize=256,discard,resize
```

Options can be added or removed to existing mount points prepared with versions of Navigator Encrypt prior to 3.5 by editing the `/etc/navencrypt/ztab` file and adding the comma-separated options (no spaces) to the end of each line as seen in the previous example above.

> **Important:** Use caution when editing the `/etc/navencrypt/ztab` file. Entries are tab-separated (not space-separated). The `ztab` file must not contain empty lines.

To see the mounted filesystems and options, run `mount`:

```
$ mount
/mnt/t2 on /mnt/t2 type ecryptfs
(rw,ecryptfs_sig=6de3db1e87077adb,ecryptfs_unlink_sigs,noauto,\
ecryptfs_cipher=aes,ecryptfs_key_bytes=32,discard,resize)
```

Pass-through mount options work for both dm-crypt and eCryptfs. For a list of available mount options, see the `man` pages for `cryptsetup` and `ecryptfs` respectively.

Encrypting and Decrypting Data Using Cloudera Navigator Encrypt

> **Warning:** Before encrypting or decrypting any data, stop all processes (for example, MySQL, MongoDB, PostgreSQL, and so on) that have access to the target data. **Failure to do so could lead to data corruption.**

After the encrypted file system is created and initialized, it is ready to hold data. All encryption and decryption functionality is performed with a single command: `navencrypt-move`.

Do not manually create directories or files under a Cloudera Navigator Encrypt mount point; use only the `navencrypt-move` command to encrypt and decrypt data. See Preparing for Encryption Using Cloudera Navigator Encrypt on page 330 for more information about mount points.

After encrypting a file or directory, all data written and read through the mount point is transparently encrypted and decrypted.

Before You Begin

Navigator Encrypt does not support encrypting data in certain environments, including the following:

- *Do not* attempt to encrypt a directory that contains or is contained within a mount point for another service (including Navigator Encrypt and NFS). For example:
  - If your encryption mount point is `/var/lib/navencrypt/mount`, do not attempt to encrypt `/var`, `/var/lib`, `/var/lib/navencrypt`, `/var/lib/navencrypt/mount`, or anything under `/var/lib/navencrypt/mount/`.
  - If you have mounted an NFS filesystem at `/mnt/home`, do not attempt to encrypt `/mnt`, `/mnt/home`, or anything under `/mnt/home`.
- Do not attempt to encrypt immutable files or directories containing immutable files.
- Do not use Navigator Encrypt within a `chroot` environment, or create a `chroot` environment within an encrypted directory.
- If your Key Trustee Server is managed by Cloudera Manager, *do not* encrypt the Cloudera Manager database with Navigator Encrypt; doing so prevents Cloudera Manager from starting.

Encrypting Data

Do not manually create directories or files under a Navigator Encrypt mount point; use only the `navencrypt-move` command to encrypt data.

Here is an example command to encrypt data, with an explanation for each option:

```
$ sudo navencrypt-move encrypt @<category> <directory_or_file_to_encrypt>
<encrypted_mount_point>
```

Important: *Do not* run `navencrypt-move` commands simultaneously in multiple terminals. Doing so results in failure to encrypt or decrypt all of the specified data. No data is lost, as the source data is not removed, but you must re-run the failed operations sequentially.

Table 34: `navencrypt-move` Command Options

Command Option	Explanation
`navencrypt-move`	Main command interface for all actions that require moving data either to or from the encrypted file system. For more information see the `navencrypt-move` man page (`man navencrypt-move`).
`encrypt`	Identifies the cryptographic operation, in this case, encrypting data. The `decrypt` option is described later in Decrypting Data on page 338. **Note:** By default, all Navigator Encrypt encryption commands require free space equal to twice the size of the encrypted data. If your environment does not have enough free space, add `--per-file` to the end of the command. This moves each file individually. Per-file encryption only requires free space equal to twice the size of the largest individual file, but is a slower operation.
`@<category>`	The access category that is applied to the data being encrypted. Encrypted data is protected by process-based access controls that restrict access to only the processes that you allow. You can use any naming convention you want (the `@` symbol is required), but Cloudera recommends keeping it simple and memorable. For example, you can use a name referencing the data type being encrypted, such as `@mysql` for a MySQL deployment. See Listing Categories on page 347 for instructions on viewing existing categories.
`<directory_or_file_to_encrypt>`	The data that you want to encrypt. This can be a single file or an entire directory. Navigator Encrypt starts after the system boots, so do not encrypt required system files and directories (such as the root partition, `/var`, and so on). Some examples of recommended data directories to encrypt are `/var/lib/mysql/data`, `/db/data`, and so on.

Command Option	Explanation
`<encrypted_mount_point>`	Where you want to store the data. This is the path to the mount point specified during the `navencrypt-prepare` command.

When a file is encrypted, a symbolic link (symlink) is created which points to a mount point `@<category>` directory. The `navencrypt-move` command moves all specified data to the encrypted filesystem and replaces it with a symlink to the mount point for that encrypted filesystem.

Encrypting a directory is similar to encrypting a file. The following command encrypts a directory:

```
$ sudo /usr/sbin/navencrypt-move encrypt @mycategory /path/to/directory_to_encrypt/
/path/to/mount
```

In this command, a directory is specified instead of a filename, and a symlink is created for that particular directory. To see the effects of this command, run:

```
$ ls -l <directory_to_encrypt>
$ du -h <encrypted_storage_directory>
```

The output demonstrates the new filesystem layout. Everything that was in the target directory is now securely stored in the encrypted filesystem.

Decrypting Data

The decryption command requires only the path to the original data, which is now a symbolic link, as an argument. The following example demonstrates how to decrypt a file using the `navencrypt-move` command:

```
$ sudo /usr/sbin/navencrypt-move decrypt /path/to/encrypted/directory_or_file
```

> **Important:** *Do not* run `navencrypt-move` commands simultaneously in multiple terminals. Doing so results in failure to encrypt or decrypt all of the specified data. No data is lost, as the source data is not removed, but you must re-run the failed operations sequentially.

As with encryption, you can specify a directory instead of a file:

```
$ sudo /usr/sbin/navencrypt-move decrypt /path/to/encrypted/directory
```

Migrating eCryptfs-Encrypted Data to dm-crypt

As of August 2015, Filesystem-level encryption using eCryptfs is deprecated. Use this procedure to migrate to dm-crypt.

RHEL 7 does not support eCryptfs. For new installations on RHEL 7, you must use Block-Level Encryption with dm-crypt on page 331. If you are planning on upgrading to RHEL 7 and are currently using eCryptfs, migrate to dm-crypt before upgrading.

> **Warning:** Before encrypting or decrypting any data, stop all processes (for example, MySQL, MongoDB, PostgreSQL, and so on) that have access to the target data. **Failure to do so could lead to data corruption.**

1. Prepare an empty block device. This can be a physical block device (such as an unused disk) or a virtual block device (for example, a logical block device created by LVM, or a loop device). For instructions on creating a loop device, see Block-Level Encryption with a Loop Device on page 332.
2. Stop any services which depend on the encrypted data to be moved.

3. Prepare a block-level encrypted mount point. See Preparing for Encryption Using Cloudera Navigator Encrypt on page 330 for details about the procedure.
4. Add ACL rules for the new encrypted mount point that match the ACL rules for the mount point you are migrating from. To view existing ACL rules, run `sudo navencrypt acl --print`.
5. Add an ACL rule for your preferred shell (for example, `/bin/bash`) to enable command-line utilities such as `mv` and `cp`:

```
$ sudo navencrypt acl --add --rule="ALLOW @category * /bin/bash"
```

6. Copy the encrypted data from the eCryptfs mount point to the dm-crypt mount point:

```
$ sudo cp -rp /ecryptfs/mountpoint/path/to/data /dmcrypt/mountpoint/path/to/data
```

7. Update any symbolic links referencing the encrypted data. The following example demonstrates updating a symbolic link for a PostgreSQL database that was originally encrypted using eCryptfs, but has been migrated to dm-crypt:

```
$ sudo ls -l /var/lib/db/data/base/16385
lrwxrwxrwx 1 root root 72 Jul 22 15:33 /var/lib/db/data/base/16385 ->
/ecryptfs/mountpoint/postgres/var/lib/db/data/base/16385
$ sudo ln -sif /dmcrypt/mountpoint/postgres/var/lib/db/data/base/16385
/var/lib/db/data/base/16385
$ sudo ls -l /var/lib/db/data/base/16385
lrwxrwxrwx 1 root root 72 Jul 22 15:33 /var/lib/db/data/base/16385 ->
/dmcrypt/mountpoint/postgres/var/lib/db/data/base/16385
```

8. Remove the ACL rule enabling command-line utilities:

```
$ sudo navencrypt acl --del --rule="ALLOW @category * /bin/bash"
```

9. Restart any services which depend on the encrypted data.
10. Verify that the data was successfully copied, then delete the original eCryptfs-encrypted data. *Do not* delete any data until you are certain that you no longer need the original data.
    a. Stop the `navencrypt-mount` service:

```
$ sudo service navencrypt-mount stop
```

    b. Remove the original mountpoint directory and the storage directory with the original encrypted data.
    c. Edit `/etc/navencrypt/ztab` and remove entries for the original encrypted directory where `eCryptfs` is listed as the `<type>`.

> **Important:** Use caution when editing the `/etc/navencrypt/ztab` file. Entries are tab-separated (not space-separated). The `ztab` file must not contain empty lines.

    d. Start the `navencrypt-mount` service:

```
$ sudo service navencrypt-mount start
```

Navigator Encrypt Access Control List

Managing the Access Control List

Cloudera Navigator Encrypt manages file system permissions with an access control list (ACL). This ACL is a security access control created by Cloudera that enables a predefined Linux process to access a file or directory managed by Navigator Encrypt.

The ACL uses rules to control process access to files. The rules specify whether a Linux process has access permissions to read from or write to a specific Navigator Encrypt path.

A rule is defined in the following order:

```
# TYPE @CATEGORY PATH PROCESS PARAMETERS
```

The following table defines the ACL rule components:

Table 35: ACL Rule Components

Component	Description
TYPE	Specifies whether to allow or deny a process. It can have either of the following values: `ALLOW` or `DENY`.
@CATEGORY	This is a user-defined shorthand, or container, for the encrypted dataset that the process will have access to. For example, if you are encrypting the directory `/var/lib/mysql`, you could use the category `@mysql` to indicate that this rule is granting access to a process on the MySQL data. See Listing Categories on page 347 for instructions on viewing existing categories.
PATH	Specifies the rights permissions of a specific path. For example: `*`, `www/*.htaccess`. Omit the leading slash (`/`).
PROCESS	Specifies the process or command name for the rule.
PARAMETERS	Tells the process the parent-child process to be executed: `--shell` defines the script for Navigator Encrypt to allow for executable process. Supported shells are `/usr/bin/bash`, `/bin/bash`, `/usr/bin/dash`, and `/bin/bash`. `--children` defines for Navigator Encrypt which child processes to allow that are executed by a process/script. Example: `--shell=/bin/bash`, `--children=/bin/df,/bin/ls`

All rules are stored in an encrypted policy file together with a set of process signatures that are used by Navigator Encrypt to authenticate each Linux process. This file is encrypted with the Navigator Encrypt key you defined during installation.

Cloudera recommends using `permissive` mode to assist with the initial ACL rule creation for your environment. In `permissive` mode, Navigator Encrypt allows full access to the encrypted data by all processes, but logs them in `dmesg` as `action="denied"` messages. Consult these messages to identify required ACL rules. To set Navigator Encrypt to permissive mode, use the following command:

```
$ sudo /usr/sbin/navencrypt set --mode=permissive
```

To view the current mode, run `navencrypt status -d`. For more information on access modes, see Access Modes.

`deny2allow`

After you generate the `action="denied"` messages, use the `navencrypt deny2allow` command to show which ACL rules are required, based on the `action="denied"` messages in `dmesg`. To show which ACL rules are required, perform the following steps:

1. Save the `dmesg` content to a file:

```
$ sudo dmesg > /tmp/dmesg.txt
```

2. Use the `dmesg.txt` file content as input to the `deny2allow` command to analyze the `action="denied"` messages and display a list of required ACL rules based on the `action="denied"` messages. Here is an example command and output:

```
$ sudo /usr/sbin/navencrypt deny2allow /tmp/dmesg.txt
ALLOW @mysql employees/* /usr/sbin/mysqld
ALLOW @mysql * /bin/bash
ALLOW @mysql * /bin/ls
```

If you need to clear the `dmesg` log and start fresh, run `dmesg -c`.

If a rule is displayed in the output from the command, it does not automatically mean the ACL rule must be added. You must determine which rules are actually needed. For example, the rule for `ls` would not typically be added as an ACL rule.

After the initial ACL rules are created, disable `permissive` mode with the following command:

```
$ sudo /usr/sbin/navencrypt set --mode=enforcing
```

Adding ACL Rules

Rules can be added one at a time using the command line or by specifying a policy file containing multiple rules. The following example shows how to add a single rule using the `navencrypt acl --add` command:

```
$ sudo /usr/sbin/navencrypt acl --add --rule="ALLOW @mysql * /usr/sbin/mysqld"
```

See Listing Categories on page 347 for instructions on viewing existing categories.

The following example shows how to add multiple rules using a policy file:

```
$ sudo /usr/sbin/navencrypt acl --add --file=/mnt/private/acl_rules
```

The contents of the policy file should contain one rule per line. For example:

```
ALLOW @mysql * /usr/sbin/mysqld
ALLOW @log * /usr/sbin/mysqld
ALLOW @apache * /usr/lib/apache2/mpm-prefork/apache2
```

Navigator Encrypt releases 3.10 and higher support comments in the policy file. Comments begin with the hash (#) symbol. You can use comments to annotate the policy file, or to temporarily disable a rule for testing. For example:

```
# Cloudera Navigator Encrypt policy file
# Allow mysqld to access all database files
ALLOW @mysql * /usr/sbin/mysqld
# Allow mysqld to write logs
ALLOW @log * /usr/sbin/mysqld
# ALLOW @apache * /usr/lib/apache2/mpm-prefork/apache2
```

Using a policy file is the fastest way to add multiple rules because it only requires the security key one time.

It is also possible to overwrite the entire current rules set with the option `--overwrite`. When this command is executed, all current rules are replaced by the ones specified in the file that contains the new set of rules. Cloudera recommends to save a copy of your current set of rules by printing it with the option `--print`.

Here is an example command using the `--overwrite` option:

```
$ sudo /usr/sbin/navencrypt acl --overwrite --file=/mnt/private/acl_rules
```

Adding ACL Rules by Profile

If your environment requires more granular controls on the processes that can access the data, you can add extra controls by using profiles. Profiles set requirements on a process other than just having the correct fingerprint. They can include such things as process owner and group, required open files, and the current working directory. To see more about adding rules by profile, see ACL Profile Rules on page 343. For details about fingerprints, see Process-Based Access Control List.

Deleting ACL Rules

Rules can be deleted in one of two ways:

1. Manually specifying the rule to delete using the command line.
2. Specifying the line number of the rule to delete.

The following example shows how to delete a rule by passing it as a parameter:

```
$ sudo /usr/sbin/navencrypt acl --del --rule="ALLOW @mysql * /usr/sbin/mysqld "
```

If you remove a MySQL `ALLOW` rule, the MySQL cache must be cleaned by executing the `FLUSH TABLES;` MySQL statement. Otherwise, it will still be possible to view data from encrypted table.

The following example shows how to delete a rule by specifying a line number:

```
$ sudo /usr/sbin/navencrypt acl --del --line 3
```

It is also possible to delete multiple ACL rules in a single command:

```
$ sudo /usr/sbin/navencrypt acl --del --line=1,3
```

See Printing ACL Rules on page 342 for information on determining line numbers.

Deleting ACL Rules by Profile

See ACL Profile Rules on page 343 for instructions on deleting rules by profile.

Printing ACL Rules

You can print the current Access Control List using the following command:

```
$ sudo /usr/sbin/navencrypt acl --print
```

Save the ACL to a file with the `--file` option:

```
$ sudo /usr/sbin/navencrypt acl --print --file=policy-backup
```

To display additional information about the organization of the policy file, use the `--list` option:

```
$ sudo /usr/sbin/navencrypt acl --list
```

Universal ACL Rules

Universal ACLs will allow or deny a process access to all files or directories encrypted with Navigator Encrypt.

The rule `ALLOW @* * /process` allows the designated process to access anything from all encrypted categories.

The rule `ALLOW @data * *` allows all processes access to any path under the `@data` category.

The rule `ALLOW @* * *` allows all processes access to all encrypted categories. Cloudera does not recommend using this rule. Use it only in test environments.

Here is an example adding a universal ACL rule and then displaying it:

```
$ sudo /usr/sbin/navencrypt acl --add --rule="ALLOW @* * /usr/sbin/mysqld"
Type MASTER passphrase:
1 rule(s) were added
# navencrypt acl --listType MASTER passphrase:
# -  Type     Category       Path                     Profile  Process
1    ALLOW    @*             *                                 /usr/sbin/mysqld
```

Enabling Shell Scripts to Be Detected by ACL

All of the previous rules work for binary files. There may be times a script, such as a shell script, must be allowed to access the encrypted directory.

You can add the script as a rule by indicating the executable binary process of this script using the `--shell` option, for example:

```
ALLOW @scripts * /root/script.sh --shell=/bin/bash
```

The `--shell` option identifies which executable process is used to execute the script. Supported shells are `/usr/bin/bash`, `/bin/bash`, `/usr/bin/dash`, and `/bin/bash`

If the script is altered, it will no longer be trusted by the ACL because the fingerprint has changed. If you edit the script you must invoke the update option to update the ACL with the new fingerprint.

In some cases, it may be necessary to grant permissions to sub-processes invoked by scripts. For example, it may be necessary to grant permissions to `/bin/bash` that also allow running the `/bin/df` command to allow the system administrator to check disk capacity through a script run using a `crontab` entry. By using the `--children` option, you can specify these permissions. For example:

```
ALLOW @scripts * /root/script.sh --shell=/bin/bash --children=/bin/df
```

The `--children` option tells Navigator Encrypt to allow the `/bin/df` binary process if it is executed by `/root/script.sh`.

To allow more than one sub-process, identify them with the `--children` option as comma-separated values. For example:

```
ALLOW @scripts * /root/script.sh --shell=/bin/bash --children=/bin/df,/bin/ls
```

To add shell-children sub-processes, execute the `navencrypt acl --add` command, for example:

```
$ sudo /usr/sbin/navencrypt acl --add --rule="ALLOW @mysql * /usr/bin/mysqld_safe \
--shell=/bin/bash --children=/bin/df,/bin/ls"
```

ACL Profile Rules

If your environment requires more granular controls on the processes that can access the data, you can add extra controls by using profiles. Profiles set requirements on a process other than just having the correct fingerprint. They can include such things as process owner and group, required open files, and the current working directory.

A profile is generated by using the following command:

```
$ usr/sbin/navencrypt-profile --pid=<pid>
```

The output, by default, will be displayed on the screen. You can redirect the output to a file using the > or >> redirect operators. You can then edit the JSON output in the file to remove lines you do not want. By default, the profile includes the UID, the short name of the binary or script (identified as comm), and the full command line of the running process (including any parameters passed). You can generate information by using one of these flags:

- `-c`, `--with-cwd`

  Output the current working directory

- `-e`, `--with-egid`

  Output the egid

- `-g`, `--with-gid`

  Output the gid

- `-u`, `--with-euid`

  Output the euid

Example output from the `navencrypt-profile` command:

```
{
"uid":"0",
"comm":"NetworkManager",
"cmdline":"NetworkManager –pid-file=/var/run/NetwrkManager/NetworkManager.pid",
"gid":"0"
"cwd":"/",
"fd0":"/dev/null",
"fd1":"/dev/null",
"fd2":"/dev/null"
}
```

Some distributions do not support `euid` and `guid`. Make sure that your profile file is correct by executing the following command to verify the expected IDs:

```
$ ps -p <pid_of_process> -o euid,egid
```

If `cmdline` parameters are variable, such as appending a process start timestamp to a filename, then the process profile will not match on subsequent restarts of the process because the current profile will have an updated timestamp and access will be denied by the ACL. You can mark those parameters as variable inside the profile file. For example, if the `cmdline` of a process is something like this:

```
"cmdline":"NetworkManager –pid-file=/var/run/NetworkManager/NetworkManager.pid \
-logfile=/var/log/NetworkManager/log-20130808152300.log"
```

Where `log-20130505122300.log` is a variable `cmdline` parameter, before adding the process profile to the ACL, edit the process profile file and use `##` to specify that a particular parameter is variable:

```
"cmdline":"NetworkManager –pid-file=/var/run/NetworkManager/NetworkManager.pid
-logfile=##"
```

With the above configuration, the ACL will allow any value for the `-logfile` `cmdline` parameter.

To enable a profile in the ACL, use the additional parameter `--profile-file=<filename>` when adding the rule to the ACL:

```
$ sudo /usr/sbin/navencrypt acl --add --rule="ALLOW @mysql * /usr/sbin/mysqld" \
–-profile-file=/path/to/profile/file
```

To display the profile portion of the rules, use the `--all` parameter with `navencrypt acl --list`:

```
$ sudo /usr/sbin/navencrypt acl --list --all
Type MASTER passphrase:
# - Type Category Path Profile Process
1 ALLOW @mysql * YES /usr/sbin/mysqld
PROFILE:
{"uid":"120","comm":"mysqld","cmdline":"mysqld"}
```

Maintaining Cloudera Navigator Encrypt

Manually Backing Up Navigator Encrypt

It is recommended that you back up Navigator Encrypt configuration directory after installation, and again after any configuration updates.

1. To manually back up the Navigator Encrypt configuration directory (`/etc/navencrypt`):

```
$ zip -r --encrypt nav-encrypt-conf.zip /etc/navencrypt
```

   The `--encrypt` option prompts you to create a password used to encrypt the zip file. This password is also required to decrypt the file. Ensure that you protect the password by storing it in a secure location.

2. Move the backup file (*nav-encrypt-conf.zip*) to a secure location.

> **Warning:** Failure to back up the configuration directory makes your backed-up encrypted data unrecoverable in the event of data loss.

Validating Navigator Encrypt Configuration

To validate the Navigator Encrypt deployment, run the following command:

```
$ sudo navencrypt status --integrity
```

This command verifies that:

- The mount encryption key (MEK) exists for each mount point.
- Each mount point in `/etc/navencrypt/ztab` has a corresponding entry in the control file (`/etc/navencrypt/control`).
- Each mount point directory exists.
- For loop devices, the file used for encrypted storage exists.

The output is similar to the following:

```
$ sudo navencrypt status --integrity
Checking MEKs integrity

        Mountpoint: /dev/loop0
                MEK file exist: ......................   [YES]
        Mountpoint: /dev/loop1
                MEK file exist: ......................   [YES]

Checking Ztab Mountpoints integrity

         Mountpoint: /dev/loop0
                 ztab vs control correspondence: .....   [YES]
                 Mountpoint directory exists: ........   [YES]
         Mountpoint: /dev/loop1
                 ztab vs control correspondence: .....   [YES]
                 Mountpoint directory exists: ........   [YES]
```

```
Checking Datastore backend files

         Datastore: '/root/my_storage_test'
                 Backend file exist: ................   [YES]
```

Restoring Mount Encryption Keys (MEKs) and Control File

Navigator Encrypt deposits its mount encryption keys (MEKs) and control file (`/etc/navencrypt/control`) in Cloudera Navigator Key Trustee Server. If these files are accidentally deleted, they can be restored from Key Trustee Server using the following commands:

- To restore MEKs:

```
$ sudo navencrypt key --restore-keys
```

- To restore the control file:

```
$ sudo navencrypt control --restore-control-file
```

Access Modes

Navigator Encrypt provides three different access modes:

- **Enforcing** is the default mode in which Navigator Encrypt validates access from all processes against the ACL. To protect your data, enforcing mode must be enabled.
- **Permissive** mode causes `action="denied"` messages to be logged in `dmesg`. It **does not** prevent access to the encrypted data. This mode is a dry-run feature to run and build ACL rules.
- **Admin** mode, as well as permissive mode, **does not** prevent access to the encrypted data. It allows any process to access the information because the ACL rules are not validated against the process. Admin mode does not cause `action="denied"` messages to be logged in `dmesg`.

To view the current access mode, run the following command:

```
$ sudo /usr/sbin/navencrypt status -d
```

> **Note:** The `navencrypt status` command reports that the `navencrypt` module is not running if no volumes are encrypted or the kernel module is not loaded.

To change the access mode, use the following command:

```
$ sudo /usr/sbin/navencrypt set --mode={enforcing|permissive|admin}
```

You cannot change the Navigator Encrypt access mode unless the Navigator Encrypt module is running. To view the status of the Navigator Encrypt module, run `navencrypt status --module`.

To start the Navigator Encrypt module there must be at least one active mount-point. To verify the mount-points status, run the following command:

```
$ sudo /etc/init.d/navencrypt-mount status
```

For RHEL 7, use `systemctl` instead:

```
$ sudo systemctl status navencrypt-mount
```

Changing and Verifying the Master Key

You can perform two operations with the `navencrypt key` command: change and `verify`.

You can verify a key against the Navigator Encrypt module, the Navigator Key Trustee server, or both. For example:

```
$ sudo /usr/sbin/navencrypt key --verify
$ sudo /usr/sbin/navencrypt key --verify --only-module
$ sudo /usr/sbin/navencrypt key --verify --only-keytrustee
```

Note: The `navencrypt key` command fails if no volumes are encrypted or the kernel module is not loaded.

The master key can be changed in the event that another key-type authentication mechanism or a new master key is required. Valid master key types are single-passphrase, dual-passphrase, and RSA key files. To change the master key type, issue the following command and follow the interactive console:

```
$ sudo /usr/sbin/navencrypt key –-change
```

You can use the `--trustees`, `--votes`, and `--recoverable` options for the new key as described in Table 32: Registration Options on page 327.

Listing Categories

To list the existing categories for each mount point, run the command `navencrypt-move --list-categories`. For example:

```
$ sudo navencrypt-move --list-categories
Navigator Encrypt Categories found per Mountpoint:

  /dmcrypt-storage
     @mysql
     @keytabs

  /home/jdoe/secrets
     @moms_recipes
     @world_domination_plan
```

Updating ACL Fingerprints

All rules reference a process fingerprint (a SHA256 digest) that is used to authenticate the process into the file system. If the filesystem detects a fingerprint that is different from the one stored in the ACL, the Linux process is denied access and treated as an untrusted process.

Occasionally this process fingerprint must be updated, such as when software is upgraded. When the fingerprint must be updated, the Navigator Encrypt administrator re-authenticates the process on the ACL by executing the command `navencrypt acl --update`.

The following example demonstrates how to determine when a process fingerprint has been changed and must be updated:

```
$ sudo /usr/sbin/navencrypt acl --list
Type MASTER passphrase:
# -  Type Category    Path         Profile  Process
1 !! ALLOW @mysql *            /usr/sbin/mysqld
2 ALLOW @log * /usr/sbin/mysqld
3 !! ALLOW @apache * /usr/lib/apache2/mpm-prefork/
```

In the example above, the double exclamation (`!!`) characters indicate that a process fingerprint has changed and must be updated. Similarly, double E (EE) characters indicate a process read error. This error can be caused by a process that does not exist or that has permission issues.

> **Note:**
>
> For RHEL-compatible OSes, the `prelink` application may also be responsible for ACL fingerprint issues. Prelinking is intended to speed up a system by reducing the time a program needs to begin. Cloudera highly recommends disabling any automated `prelink` jobs, which are enabled by default in some systems. As an alternative, Cloudera recommends that you integrate a manual `prelink` run into your existing change control policies to ensure minimal downtime for applications accessing encrypted data.
>
> To disable prelinking, modify the `/etc/sysconfig/prelink` file and change `PRELINKING=yes` to `PRELINKING=no`. Then, run the `/etc/cron.daily/prelink` script as root. Once finished, automatic prelinking is disabled.
>
> For more information about how prelinking affects your system, see prelink.

Managing Mount Points

The `/etc/init.d/navencrypt-mount` command mounts all mount points that were registered with the `navencrypt-prepare` command and are listed in the `/etc/navencrypt/ztab` file. The possible operations are:

- `start`
- `stop`
- `status`
- `restart`

For RHEL 7, use `systemctl [start|stop|status|restart] navencrypt-mount`.

> **Note:** Do not manually unmount the encryption mount point (for example, using `umount`). If you do so, you must manually close the `dm-crypt` device using the following procedure:
>
> 1. Run `dmsetup table` to list the `dm-crypt` devices.
> 2. Run `cryptsetup luksClose` *<device_name>* to close the device for the unmounted mount point.

When executing the `stop` operation, the encrypted mount point is unmounted, and your data becomes inaccessible.

The following example shows how to execute `navencrypt-mount status` with some inactive mount points:

```
$ sudo /etc/init.d/navencrypt-mount status
```

The following example shows how to execute the `navencrypt-mount stop` command:

```
$ sudo /etc/init.d/navencrypt-mount stop
```

The following example shows how to execute the `navencrypt-mount start` command:

```
$ sudo /etc/init.d/navencrypt-mount start
```

Here is an example command used to manually mount a directory:

```
$ sudo /usr/sbin/mount.navencrypt /path/to_encrypted_data/ /path/to/mountpoint
```

This command can be executed only if the `navencrypt-prepare` command was previously executed.

Navigator Encrypt Kernel Module Setup

If the kernel headers were not installed on your host, or if the wrong version of the kernel headers were installed, the Navigator Encrypt module was not built at installation time. To avoid reinstalling the system, install the correct headers and execute the `navencrypt-module-setup` command. This attempts to build the module and install it.

This method is also an efficient way to install any new Navigator Encrypt module feature or fix without otherwise modifying your current Navigator Encrypt environment.

Navigator Encrypt Configuration Directory Structure

The file and directory structure of `/etc/navencrypt` is as follows:

```
$ tree /etc/navencrypt/
/etc/navencrypt/
    control -> /etc/navencrypt/jSpi9SM65xUIIhrau1Nn8ZXmQhrrQ9e363EUz8HKiRs
    jSpi9SM65xUIIhrau1Nn8ZXmQhrrQ9e363EUz8HKiRs
    rules
    ztab
locust
    keytrustee
     clientname
     deposits
         dev.loop0
        media.31E5-79B9locustlocust[system ~]# . /etc/*release[system ~]# . /etc/*release

         mnt.a
         mnt.encrypted
         mnt.tomount
     pubring.gpg
     pubring.gpg~
     random_seed
     secring.gpg
     trustdb.gpg
     ztrustee.conf
```

The following files and folders are part of the created file structure:

- `control`

  File that saves information about the mount points and corresponding Navigator Key Trustee keys. If this file is accidentally deleted, you can restore it using the `navencrypt control --restore-control-file` command.

- `rules`

  File that contains the ACL rules. It is encrypted with the user-provided master key.

- `ztab`

  File that contains information about all the mount points and their encryption type.

  > **Important:** Use caution when editing the `/etc/navencrypt/ztab` file. Entries are tab-separated (not space-separated). The `ztab` file must not contain empty lines.

- `keytrustee`

  Directory where Navigator Key Trustee GPG keys are stored. These are generated during `navencrypt register` operations.

- `keytrustee/deposits`

  Directory where the Navigator Encrypt mount encryption keys (MEKs) are saved. These are encrypted with the user-provided master key. If these are accidentally deleted, you can can restore them from Key Trustee Server using the `navencrypt key --restore-keys` command.

Every mount point has an internal randomly-generated encryption passphrase.

Collecting Navigator Encrypt Environment Information

When troubleshooting problems with Navigator Encrypt, it is helpful to gather information about the installation and environment. Navigator Encrypt provides a command to facilitate this:

```
$ sudo navencrypt-collect
```

This command collects and outputs to the console the following information:

- Information about the system on which Navigator Encrypt is installed
- Entries from `/etc/navencrypt/ztab`
- The contents of the `keytrustee.conf` file
- Recent entries from the Navigator Encrypt log file
- Configured software repositories
- Checksums of all `/usr/src/navencrypt*` and `/usr/sbin/navencrypt*` files

You can use this information to compare Navigator Encrypt installations and to provide to Cloudera Support for troubleshooting. The `navencrypt-collect` command only outputs this information on the console, and does not generate any files or upload to Cloudera.

To save the information to a file, use the redirect operator (>). For example:

```
$ sudo navencrypt-collect > navencrypt.info
```

Upgrading Navigator Encrypt Hosts

See Best Practices for Upgrading Navigator Encrypt Hosts for considerations when upgrading operating systems (OS) and kernels on hosts that have Navigator Encrypt installed.

Configuring Encryption for Data Spills

Some CDH services can encrypt data stored temporarily on the local filesystem outside HDFS. For example, data may *spill* to disk during memory-intensive operations, or when a service exceeds its allotted memory on a host. You can enable on-disk spill encryption for the following services.

MapReduce v2 (YARN)

MapReduce v2 can encrypt intermediate files generated during encrypted shuffle and data spilled to disk during the map and reduce stages. To enable encryption on these intermediate files, use Cloudera Manager to specify settings for the **MapReduce Client Advanced Configuration Snippet (Safety Valve)** for the `mapred-site.xml` associated with a gateway node.

From the Cloudera Manager Admin Console:

1. Select **Clusters** > **YARN**.
2. Click the **Configuration** tab.
3. In the Search field, enter **MapReduce Client Advanced Configuration Snippet (Safety Valve)** to find the safety valve on one of YARN's gateway nodes:

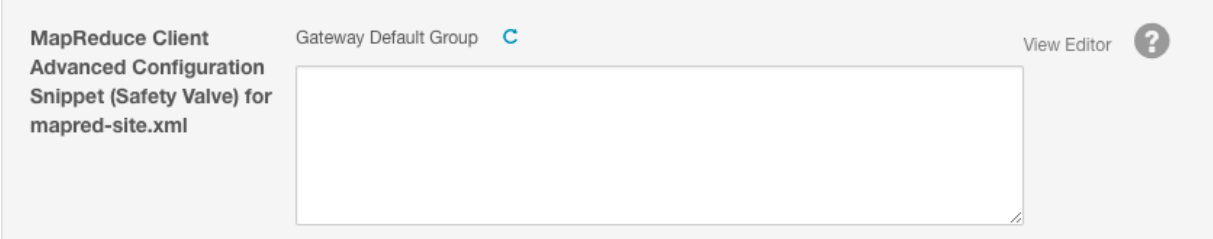

4. Enter the XML in the field (or switch to Editor mode and enter each property and its value in the fields provided). A complete example of the XML is shown here:

```
<property>
    <name>mapreduce.job.encrypted-intermediate-data</name>
    <value>true</value>
</property>
<property>
    <name>mapreduce.job.encrypted-intermediate-data-key-size-bits</name>
    <value>128</value>
</property>
<property>
    <name>mapreduce.job.encrypted-intermediate-data.buffer.kb</name>
    <value>128</value>
</property>
```

5. Click **Save Changes**.

The table provides descriptions of the properties used for data-spill encryption:

Property	Default	Description
`mapreduce.job.encrypted-intermediate-data`	`false`	Enables (true) and disables (false) encryption for intermediate MapReduce spills.
`mapreduce.job.encrypted-intermediate-data-key-size-bits`	`128`	Length of key (bits) used for encryption.
`mapreduce.job.encrypted-intermediate-data.buffer.kb`	`128`	Buffer size (Kb) for the stream written to disk after encryption.

> **Note:** Enabling encryption for intermediate data spills limits the number of attempts for a job to 1.

HBase

HBase does not write data outside HDFS, and does not require spill encryption.

Impala

Impala allows certain memory-intensive operations to be able to write temporary data to disk in case these operations approach their memory limit on a host. For details, read SQL Operations that Spill to Disk. To enable disk spill encryption in Impala:

1. Go to the Cloudera Manager Admin Console.
2. Click the **Configuration** tab.
3. Select **Scope** > **Impala Daemon**.
4. Select **Category** > **Security**.
5. Check the checkbox for the **Disk Spill Encryption** property.
6. Click **Save Changes.**

Hive

Hive jobs occasionally write data temporarily to local directories. If you enable HDFS encryption, you must ensure that the following intermediate local directories are also protected:

- **`LOCALSCRATCHDIR`**: The MapJoin optimization in Hive writes HDFS tables to a local directory and then uploads them to the distributed cache. To ensure these files are encrypted, either disable MapJoin by setting `hive.auto.convert.join` to `false`, or encrypt the *local* Hive Scratch directory (`hive.exec.local.scratchdir`) using Cloudera Navigator Encrypt.
- **`DOWNLOADED_RESOURCES_DIR`**: JARs that are added to a user session and stored in HDFS are downloaded to `hive.downloaded.resources.dir` on the HiveServer2 local filesystem. To encrypt these JAR files, configure Cloudera Navigator Encrypt to encrypt the directory specified by `hive.downloaded.resources.dir`.
- **NodeManager Local Directory List:** Hive stores JARs and MapJoin files in the distributed cache. To use MapJoin or encrypt JARs and other resource files, the `yarn.nodemanager.local-dirs` YARN configuration property must be configured to a set of encrypted local directories on all nodes.

For more information on Hive behavior with HDFS encryption enabled, see Using HDFS Encryption with Hive.

Flume

Flume supports on-disk encryption for log files written by the Flume file channels. See Configuring Encrypted On-disk File Channels for Flume on page 352.

Configuring Encrypted On-disk File Channels for Flume

Flume supports on-disk encryption of data on the local disk. To implement this:

- Generate an encryption key to use for the Flume Encrypted File Channel
- Configure on-disk encryption by setting parameters in the `flume.conf` file

Important:

Flume on-disk encryption operates with a maximum strength of 128-bit AES encryption unless the JCE unlimited encryption cryptography policy files are installed. Please see this Oracle document for information about enabling strong cryptography:
http://www.oracle.com/technetwork/java/javase/downloads/jce-6-download-429243.html

Consult your security organization for guidance on the acceptable strength of your encryption keys. Cloudera has tested with AES-128, AES-192, and AES-256.

Generating Encryption Keys

Use the `keytool` program included with the Oracle JDK to create the AES encryption keys for use with Flume.

The command to generate a 128-bit key that uses the same password as the key store password is:

```
keytool -genseckey -alias key-1 -keyalg AES -keysize 128 -validity 9000 \
-keystore test.keystore -storetype jceks \
-storepass keyStorePassword
```

The command to generate a 128-bit key that uses a different password from that used by the key store is:

```
keytool -genseckey -alias key-0 -keypass keyPassword -keyalg AES \
-keysize 128 -validity 9000 -keystore test.keystore \
-storetype jceks -storepass keyStorePassword
```

The key store and password files can be stored anywhere on the file system; both files should have `flume` as the owner and `0600` permissions.

Please note that `-keysize` controls the strength of the AES encryption key, in bits; 128, 192, and 256 are the allowed values.

Configuration

Flume on-disk encryption is enabled by setting parameters in the `/etc/flume-ng/conf/flume.conf` file.

Basic Configuration

The first example is a basic configuration with an alias called `key-0` that uses the same password as the key store:

```
agent.channels.ch-0.type = file
agent.channels.ch-0.capacity = 10000
agent.channels.ch-0.encryption.cipherProvider = AESCTRNOPADDING
agent.channels.ch-0.encryption.activeKey = key-0
agent.channels.ch-0.encryption.keyProvider = JCEKSFILE
agent.channels.ch-0.encryption.keyProvider.keyStoreFile = /path/to/my.keystore
agent.channels.ch-0.encryption.keyProvider.keyStorePasswordFile =
/path/to/my.keystore.password
agent.channels.ch-0.encryption.keyProvider.keys = key-0
```

In the next example, `key-0` uses its own password which may be different from the key store password:

```
agent.channels.ch-0.type = file
agent.channels.ch-0.capacity = 10000
agent.channels.ch-0.encryption.cipherProvider = AESCTRNOPADDING
agent.channels.ch-0.encryption.activeKey = key-0
agent.channels.ch-0.encryption.keyProvider = JCEKSFILE
agent.channels.ch-0.encryption.keyProvider.keyStoreFile = /path/to/my.keystore
agent.channels.ch-0.encryption.keyProvider.keyStorePasswordFile =
/path/to/my.keystore.password
agent.channels.ch-0.encryption.keyProvider.keys = key-0
agent.channels.ch-0.encryption.keyProvider.keys.key-0.passwordFile =
/path/to/key-0.password
```

Changing Encryption Keys Over Time

To modify the key, modify the configuration as shown below. This example shows how to change the configuration to use `key-1` instead of `key-0`:

```
agent.channels.ch-0.type = file
agent.channels.ch-0.capacity = 10000
agent.channels.ch-0.encryption.cipherProvider = AESCTRNOPADDING
agent.channels.ch-0.encryption.activeKey = key-1
agent.channels.ch-0.encryption.keyProvider = JCEKSFILE
agent.channels.ch-0.encryption.keyProvider.keyStoreFile = /path/to/my.keystore
agent.channels.ch-0.encryption.keyProvider.keyStorePasswordFile =
/path/to/my.keystore.password
agent.channels.ch-0.encryption.keyProvider.keys = key-0 key-1
```

The same scenario except that `key-0` and `key-1` have their own passwords is shown here:

```
agent.channels.ch-0.type = file
agent.channels.ch-0.capacity = 10000
agent.channels.ch-0.encryption.cipherProvider = AESCTRNOPADDING
agent.channels.ch-0.encryption.activeKey = key-1
agent.channels.ch-0.encryption.keyProvider = JCEKSFILE
agent.channels.ch-0.encryption.keyProvider.keyStoreFile = /path/to/my.keystore
agent.channels.ch-0.encryption.keyProvider.keyStorePasswordFile =
/path/to/my.keystore.password
agent.channels.ch-0.encryption.keyProvider.keys = key-0 key-1
agent.channels.ch-0.encryption.keyProvider.keys.key-0.passwordFile =
/path/to/key-0.password
agent.channels.ch-0.encryption.keyProvider.keys.key-1.passwordFile =
/path/to/key-1.password
```

Troubleshooting

If the unlimited strength JCE policy files are not installed, an error similar to the following is printed in the `flume.log`:

```
07 Sep 2012 23:22:42,232 ERROR [lifecycleSupervisor-1-0]
(org.apache.flume.channel.file.encryption.AESCTRNoPaddingProvider.getCipher:137) - Unable
 to load key using transformation: AES/CTR/NoPadding; Warning: Maximum allowed key length
 = 128 with the available JCE security policy files. Have you installed the JCE unlimited
 strength jurisdiction policy files?
java.security.InvalidKeyException: Illegal key size
at javax.crypto.Cipher.a(DashoA13*..)
at javax.crypto.Cipher.a(DashoA13*..)
at javax.crypto.Cipher.a(DashoA13*..)
at javax.crypto.Cipher.init(DashoA13*..)
at javax.crypto.Cipher.init(DashoA13*..)
at
org.apache.flume.channel.file.encryption.AESCTRNoPaddingProvider.getCipher(AESCTRNoPaddingProvider.java:120)
at
org.apache.flume.channel.file.encryption.AESCTRNoPaddingProvider.access$200(AESCTRNoPaddingProvider.java:35)
at
org.apache.flume.channel.file.encryption.AESCTRNoPaddingProvider$AESCTRNoPaddingDecryptor.<init>(AESCTRNoPaddingProvider.java:94)
at
org.apache.flume.channel.file.encryption.AESCTRNoPaddingProvider$AESCTRNoPaddingDecryptor.<init>(AESCTRNoPaddingProvider.java:91)
at
org.apache.flume.channel.file.encryption.AESCTRNoPaddingProvider$DecryptorBuilder.build(AESCTRNoPaddingProvider.java:66)
at
org.apache.flume.channel.file.encryption.AESCTRNoPaddingProvider$DecryptorBuilder.build(AESCTRNoPaddingProvider.java:62)
at
org.apache.flume.channel.file.encryption.CipherProviderFactory.getDecrypter(CipherProviderFactory.java:47)
at org.apache.flume.channel.file.LogFileV3$SequentialReader.<init>(LogFileV3.java:257)
at
org.apache.flume.channel.file.LogFileFactory.getSequentialReader(LogFileFactory.java:110)
at org.apache.flume.channel.file.ReplayHandler.replayLog(ReplayHandler.java:258)
at org.apache.flume.channel.file.Log.replay(Log.java:339)
at org.apache.flume.channel.file.FileChannel.start(FileChannel.java:260)
at
org.apache.flume.lifecycle.LifecycleSupervisor$MonitorRunnable.run(LifecycleSupervisor.java:236)
at java.util.concurrent.Executors$RunnableAdapter.call(Executors.java:441)
at java.util.concurrent.FutureTask$Sync.innerRunAndReset(FutureTask.java:317)
```

```
at java.util.concurrent.FutureTask.runAndReset(FutureTask.java:150)
at
java.util.concurrent.ScheduledThreadPoolExecutor$ScheduledFutureTask.access$101(ScheduledThreadPoolExecutor.java:98)
at
java.util.concurrent.ScheduledThreadPoolExecutor$ScheduledFutureTask.runPeriodic(ScheduledThreadPoolExecutor.java:180)
at
java.util.concurrent.ScheduledThreadPoolExecutor$ScheduledFutureTask.run(ScheduledThreadPoolExecutor.java:204)
at java.util.concurrent.ThreadPoolExecutor$Worker.runTask(ThreadPoolExecutor.java:886)
at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:908)
at java.lang.Thread.run(Thread.java:662)
```

Impala Security Overview

Impala includes a fine-grained authorization framework for Hadoop, based on the Sentry open source project. Sentry authorization was added in Impala 1.1.0. Together with the Kerberos authentication framework, Sentry takes Hadoop security to a new level needed for the requirements of highly regulated industries such as healthcare, financial services, and government. Impala also includes an auditing capability; Impala generates the audit data, the Cloudera Navigator product consolidates the audit data from all nodes in the cluster, and Cloudera Manager lets you filter, visualize, and produce reports. The auditing feature was added in Impala 1.1.1.

The Impala security features have several objectives. At the most basic level, security prevents accidents or mistakes that could disrupt application processing, delete or corrupt data, or reveal data to unauthorized users. More advanced security features and practices can harden the system against malicious users trying to gain unauthorized access or perform other disallowed operations. The auditing feature provides a way to confirm that no unauthorized access occurred, and detect whether any such attempts were made. This is a critical set of features for production deployments in large organizations that handle important or sensitive data. It sets the stage for multi-tenancy, where multiple applications run concurrently and are prevented from interfering with each other.

The material in this section presumes that you are already familiar with administering secure Linux systems. That is, you should know the general security practices for Linux and Hadoop, and their associated commands and configuration files. For example, you should know how to create Linux users and groups, manage Linux group membership, set Linux and HDFS file permissions and ownership, and designate the default permissions and ownership for new files. You should be familiar with the configuration of the nodes in your Hadoop cluster, and know how to apply configuration changes or run a set of commands across all the nodes.

The security features are divided into these broad categories:

authorization

Which users are allowed to access which resources, and what operations are they allowed to perform? Impala relies on the open source Sentry project for authorization. By default (when authorization is not enabled), Impala does all read and write operations with the privileges of the `impala` user, which is suitable for a development/test environment but not for a secure production environment. When authorization is enabled, Impala uses the OS user ID of the user who runs `impala-shell` or other client program, and associates various privileges with each user. See Enabling Sentry Authorization for Impala for details about setting up and managing authorization.

authentication

How does Impala verify the identity of the user to confirm that they really are allowed to exercise the privileges assigned to that user? Impala relies on the Kerberos subsystem for authentication. See Enabling Kerberos Authentication for Impala on page 141 for details about setting up and managing authentication.

auditing

What operations were attempted, and did they succeed or not? This feature provides a way to look back and diagnose whether attempts were made to perform unauthorized operations. You use this information to track down suspicious activity, and to see where changes are needed in authorization policies. The audit data produced by this feature is collected by the Cloudera Manager product and then presented in a user-friendly form by the Cloudera Manager product. See Auditing Impala Operations for details about setting up and managing auditing.

These other topics in the *Security Guide* cover how Impala integrates with security frameworks such as Kerberos, LDAP, and Sentry:

- Impala Authentication on page 140
- Enabling Sentry Authorization for Impala

Security Guidelines for Impala

The following are the major steps to harden a cluster running Impala against accidents and mistakes, or malicious attackers trying to access sensitive data:

- Secure the `root` account. The `root` user can tamper with the `impalad` daemon, read and write the data files in HDFS, log into other user accounts, and access other system services that are beyond the control of Impala.
- Restrict membership in the `sudoers` list (in the `/etc/sudoers` file). The users who can run the `sudo` command can do many of the same things as the `root` user.
- Ensure the Hadoop ownership and permissions for Impala data files are restricted.
- Ensure the Hadoop ownership and permissions for Impala log files are restricted.
- Ensure that the Impala web UI (available by default on port 25000 on each Impala node) is password-protected. See Impala Web User Interface for Debugging for details.
- Create a policy file that specifies which Impala privileges are available to users in particular Hadoop groups (which by default map to Linux OS groups). Create the associated Linux groups using the `groupadd` command if necessary.
- The Impala authorization feature makes use of the HDFS file ownership and permissions mechanism; for background information, see the HDFS Permissions Guide. Set up users and assign them to groups at the OS level, corresponding to the different categories of users with different access levels for various databases, tables, and HDFS locations (URIs). Create the associated Linux users using the `useradd` command if necessary, and add them to the appropriate groups with the `usermod` command.
- Design your databases, tables, and views with database and table structure to allow policy rules to specify simple, consistent rules. For example, if all tables related to an application are inside a single database, you can assign privileges for that database and use the `*` wildcard for the table name. If you are creating views with different privileges than the underlying base tables, you might put the views in a separate database so that you can use the `*` wildcard for the database containing the base tables, while specifying the precise names of the individual views. (For specifying table or database names, you either specify the exact name or `*` to mean all the databases on a server, or all the tables and views in a database.)
- Enable authorization by running the `impalad` daemons with the `-server_name` and `-authorization_policy_file` options on all nodes. (The authorization feature does not apply to the `statestored` daemon, which has no access to schema objects or data files.)
- Set up authentication using Kerberos, to make sure users really are who they say they are.

Securing Impala Data and Log Files

One aspect of security is to protect files from unauthorized access at the filesystem level. For example, if you store sensitive data in HDFS, you specify permissions on the associated files and directories in HDFS to restrict read and write permissions to the appropriate users and groups.

If you issue queries containing sensitive values in the `WHERE` clause, such as financial account numbers, those values are stored in Impala log files in the Linux filesystem and you must secure those files also. For the locations of Impala log files, see Using Impala Logging.

All Impala read and write operations are performed under the filesystem privileges of the `impala` user. The `impala` user must be able to read all directories and data files that you query, and write into all the directories and data files for `INSERT` and `LOAD DATA` statements. At a minimum, make sure the `impala` user is in the `hive` group so that it can access files and directories shared between Impala and Hive. See User Account Requirements for more details.

Setting file permissions is necessary for Impala to function correctly, but is not an effective security practice by itself:

- The way to ensure that only authorized users can submit requests for databases and tables they are allowed to access is to set up Sentry authorization, as explained in Enabling Sentry Authorization for Impala. With authorization enabled, the checking of the user ID and group is done by Impala, and unauthorized access is blocked by Impala itself. The actual low-level read and write requests are still done by the `impala` user, so you must have appropriate file and directory permissions for that user ID.

- You must also set up Kerberos authentication, as described in Enabling Kerberos Authentication for Impala on page 141, so that users can only connect from trusted hosts. With Kerberos enabled, if someone connects a new host to the network and creates user IDs that match your privileged IDs, they will be blocked from connecting to Impala at all from that host.

Installation Considerations for Impala Security

Impala 1.1 comes set up with all the software and settings needed to enable security when you run the `impalad` daemon with the new security-related options (`-server_name` and `-authorization_policy_file`). You do not need to change any environment variables or install any additional JAR files. In a cluster managed by Cloudera Manager, you do not need to change any settings in Cloudera Manager.

Securing the Hive Metastore Database

It is important to secure the Hive metastore, so that users cannot access the names or other information about databases and tables the through the Hive client or by querying the metastore database. Do this by turning on Hive metastore security, using the instructions in the CDH 5 Security Guide for securing different Hive components:

- Secure the Hive Metastore.
- In addition, allow access to the metastore only from the HiveServer2 server, and then disable local access to the HiveServer2 server.

Securing the Impala Web User Interface

The instructions in this section presume you are familiar with the `.htpasswd` mechanism commonly used to password-protect pages on web servers.

Password-protect the Impala web UI that listens on port 25000 by default. Set up a `.htpasswd` file in the `$IMPALA_HOME` directory, or start both the `impalad` and `statestored` daemons with the `--webserver_password_file` option to specify a different location (including the filename).

This file should only be readable by the Impala process and machine administrators, because it contains (hashed) versions of passwords. The username / password pairs are not derived from Unix usernames, Kerberos users, or any other system. The `domain` field in the password file must match the domain supplied to Impala by the new command-line option `--webserver_authentication_domain`. The default is `mydomain.com`.

Impala also supports using HTTPS for secure web traffic. To do so, set `--webserver_certificate_file` to refer to a valid `.pem` TLS/SSL certificate file. Impala will automatically start using HTTPS once the TLS/SSL certificate has been read and validated. A `.pem` file is basically a private key, followed by a signed TLS/SSL certificate; make sure to concatenate both parts when constructing the `.pem` file.

If Impala cannot find or parse the `.pem` file, it prints an error message and quits.

Note:

If the private key is encrypted using a passphrase, Impala will ask for that passphrase on startup, which is not useful for a large cluster. In that case, remove the passphrase and make the `.pem` file readable only by Impala and administrators.

When you turn on TLS/SSL for the Impala web UI, the associated URLs change from `http://` prefixes to `https://`. Adjust any bookmarks or application code that refers to those URLs.

Configuring Secure Access for Impala Web Servers

Cloudera Manager supports two methods of authentication for secure access to the Impala Catalog Server, Daemon, and StateStoreweb servers: password-based authentication and TLS/SSL certificate authentication.

Authentication for the three types of daemons can be configured independently.

Configuring Password Authentication

1. Navigate to **Clusters** > **Impala Service** > **Configuration**.
2. Search for "password" using the Search box in the **Configuration** tab. This should display the password-related properties (Username and Password properties) for the Impala Daemon, StateStore, and Catalog Server. If there are multiple role groups configured for Impala Daemon instances, the search should display all of them.
3. Enter a username and password into these fields.
4. Click **Save Changes**, and restart the Impala service.

Now when you access the Web UI for the Impala Daemon, StateStore, or Catalog Server, you are asked to log in before access is granted.

Configuring TLS/SSL Certificate Authentication

1. Create or obtain an TLS/SSL certificate.
2. Place the certificate, in `.pem` format, on the hosts where the Impala Catalog Server and StateStore are running, and on each host where an Impala Daemon is running. It can be placed in any location (path) you choose. If all the Impala Daemons are members of the same role group, then the `.pem` file must have the same path on every host.
3. Navigate to **Clusters** > **Impala Service** > **Configuration**.
4. Search for "certificate" using the Search box in the **Configuration** tab. This should display the certificate file location properties for the Impala Catalog Server, Impala Daemon, and StateStore. If there are multiple role groups configured for Impala Daemon instances, the search should display all of them.
5. In the property fields, enter the full path name to the certificate file.
6. Click **Save Changes**, and restart the Impala service.

> **Important:** If Cloudera Manager cannot find the `.pem` file on the host for a specific role instance, that role will fail to start.

When you access the Web UI for the Impala Catalog Server, Impala Daemon, and StateStore, `https` will be used.

Kudu Security Overview

Kudu includes security features that allow Kudu clusters to be hardened against access from unauthorized users. Kudu uses strong authentication with Kerberos, while communication between Kudu clients and servers can now be encrypted with TLS. Kudu also allows you to use HTTPS encryption to connect to the web UI.

The rest of this topic describes the security capabilities of Apache Kudu and how to configure a secure Kudu cluster. Currently, there are a few known limitations in Kudu security that might impact your cluster. For the list, see Security Limitations.

Kudu Authentication with Kerberos

Kudu can be configured to enforce secure authentication among servers, and between clients and servers. Authentication prevents untrusted actors from gaining access to Kudu, and securely identifies connecting users or services for authorization checks. Authentication in Kudu is designed to interoperate with other secure Hadoop components by utilizing Kerberos.

Configure authentication on Kudu servers using the `--rpc-authentication` flag, which can be set to one of the following options:

- `required` - Kudu will reject connections from clients and servers who lack authentication credentials.
- `optional` - Kudu will attempt to use strong authentication, but will allow unauthenticated connections.
- `disabled` - Kudu will only allow unauthenticated connections.

By default, the flag is set to `optional`. To secure your cluster, set `--rpc-authentication` to `required`.

Internal Private Key Infrastructure (PKI)

Kudu uses an internal PKI to issue X.509 certificates to servers in the cluster. Connections between peers who have both obtained certificates will use TLS for authentication. In such cases, neither peer needs to contact the Kerberos KDC.

X.509 certificates are only used for internal communication among Kudu servers, and between Kudu clients and servers. These certificates are never presented in a public facing protocol. By using internally-issued certificates, Kudu offers strong authentication which scales to huge clusters, and allows TLS encryption to be used without requiring you to manually deploy certificates on every node.

Authentication Tokens

After authenticating to a secure cluster, the Kudu client will automatically request an authentication token from the Kudu master. An authentication token encapsulates the identity of the authenticated user and carries the Kudu master's RSA signature so that its authenticity can be verified. This token will be used to authenticate subsequent connections. By default, authentication tokens are only valid for seven days, so that even if a token were compromised, it cannot be used indefinitely. For the most part, authentication tokens should be completely transparent to users. By using authentication tokens, Kudu is able to take advantage of strong authentication, without paying the scalability cost of communicating with a central authority for every connection.

When used with distributed compute frameworks such as Apache Spark, authentication tokens can simplify configuration and improve security. For example, the Kudu Spark connector will automatically retrieve an authentication token during the planning stage, and distribute the token to tasks. This allows Spark to work against a secure Kudu cluster where only the planner node has Kerberos credentials.

Client Authentication to Secure Kudu Clusters

Users running client Kudu applications must first run the `kinit` command to obtain a Kerberos ticket-granting ticket. For example:

```
$ kinit admin@EXAMPLE-REALM.COM
```

Once authenticated, you use the same client code to read from and write to Kudu servers with and without the Kerberos configuration.

Scalability

Kudu authentication is designed to scale to thousands of nodes, which means it must avoid unnecessary coordination with a central authentication authority (such as the Kerberos KDC) for each connection. Instead, Kudu servers and clients use Kerberos to establish initial trust with the Kudu master, and then use alternate credentials for subsequent connections. As described previously, the Kudu master issues internal X.509 certificates to tablet servers on startup, and temporary authentication tokens to clients on first contact.

Encryption

Kudu allows you to use TLS to encrypt all communications among servers, and between clients and servers. Configure TLS encryption on Kudu servers using the `--rpc-encryption` flag, which can be set to one of the following options:

- `required` - Kudu will reject unencrypted connections.
- `optional` - Kudu will attempt to use encryption, but will allow unencrypted connections.
- `disabled` - Kudu will not use encryption.

By default, the flag is set to `optional`. To secure your cluster, set `--rpc-encryption` to `required`.

> **Note:** Kudu will automatically turn off encryption on local loopback connections, since traffic from these connections is never exposed externally. This allows locality-aware compute frameworks, such as Spark and Impala, to avoid encryption overhead, while still ensuring data confidentiality.

Coarse-grained Authorization

Kudu supports coarse-grained authorization checks for client requests based on the client's authenticated Kerberos principal (user or service). Access levels are granted based on whitelist-style Access Control Lists (ACLs), one for each level. Each ACL specifies a comma-separated list of users, or may be set to `'*'` to indicate that all authenticated users have access rights at the specified level.

The two levels of access which can be configured are:

- **Superuser** - Principals authorized as a superuser can perform certain administrative functions such as using the `kudu` command line tool to diagnose and repair cluster issues.
- **User** - Principals authorized as a user are able to access and modify all data in the Kudu cluster. This includes the ability to create, drop, and alter tables, as well as read, insert, update, and delete data. The default value for the User ACL is `'*'`, which allows all users access to the cluster. However, if authentication is enabled, this will restrict access to only those users who are able to successfully authenticate using Kerberos. Unauthenticated users on the same network as the Kudu servers will be unable to access the cluster.

> **Note:** Internally, Kudu has a third access level for the daemons themselves called **Service**. This is used to ensure that users cannot connect to the cluster and pose as tablet servers.

Web UI Encryption

The Kudu web UI can be configured to use secure HTTPS encryption by providing each server with TLS certificates. Use the `--webserver-certificate-file` and `--webserver-private-key-file` properties to specify the certificate and private key to be used for communication.

Alternatively, you can choose to completely disable the web UI by setting `--webserver-enabled` flag to `false` on the Kudu servers.

Web UI Redaction

To prevent sensitive data from being included in the web UI, all row data is redacted. Table metadata, such as table names, column names, and partitioning information is not redacted. Alternatively, you can choose to completely disable the web UI by setting the `--webserver-enabled` flag to `false` on the Kudu servers.

> **Note:** Disabling the web UI will also disable REST endpoints such as `/metrics`. Monitoring systems rely on these endpoints to gather metrics data.

Log Redaction

To prevent sensitive data from being included in Kudu server logs, all row data will be redacted. You can turn off log redaction using the `--redact` flag.

Configuring a Secure Kudu Cluster using Cloudera Manager

> **Warning:** If you are upgrading from Kudu 1.2.0 / CDH 5.10.x, you must upgrade both Kudu and CDH parcels (or packages) at the same time. If you upgrade Kudu but do not upgrade CDH, new Kudu features such as Security will not be available. Note that even though you might be able to see the updated configuration options for Kudu security in Cloudera Manager, configuring them will have no effect.

Use the following set of instructions to secure a Kudu cluster using Cloudera Manager:

Enabling Kerberos Authentication and RPC Encryption

> **Important:** The following instructions assume you already have a secure Cloudera Manager cluster with Kerberos authentication enabled. If this is not the case, first secure your cluster using the steps described at Enabling Kerberos Authentication Using the Cloudera Manager Wizard.

To enable Kerberos authentication for Kudu:

1. Go to the **Kudu** service.
2. Click the **Configuration** tab.
3. Select **Category** > **Main**.
4. In the Search field, type **Kerberos** to show the relevant properties.
5. Edit the following properties according to your cluster configuration:

Field	Usage Notes
Kerberos Principal	Set to the default principal, kudu. Currently, Kudu does not support configuring a custom service principal for Kudu processes.

Field	Usage Notes
Enable Secure Authentication And Encryption	Select this checkbox to enable authentication and RPC encryption between all Kudu clients and servers, as well as between individual servers. Only enable this property after you have configured Kerberos.

6. Click **Save Changes**.
7. You will see an error message that tells you the Kudu keytab is missing. To generate the keytab, go to the top navigation bar and click **Administration** > **Security**.
8. Go to the **Kerberos Credentials** tab. On this page you will see a list of the existing Kerberos principals for services running on the cluster.
9. Click **Generate Missing Credentials**. Once the Generate Missing Credentials command has finished running, you will see the Kudu principal added to the list.

Configuring Coarse-grained Authorization with ACLs

1. Go to the **Kudu** service.
2. Click the **Configuration** tab.
3. Select **Category** > **Security**.
4. In the Search field, type **ACL** to show the relevant properties.
5. Edit the following properties according to your cluster configuration:

Field	Usage Notes
***Superuser* Access Control List**	Add a comma-separated list of superusers who can access the cluster. By default, this property is left blank. '*' indicates that all authenticated users will be given superuser access.
***User* Access Control List**	Add a comma-separated list of users who can access the cluster. By default, this property is set to '*'. The default value of `'*'` allows all users access to the cluster. However, if authentication is enabled, this will restrict access to only those users who are able to successfully authenticate using Kerberos. Unauthenticated users on the same network as the Kudu servers will be unable to access the cluster. Add the `impala` user to this list to allow Impala to query data in Kudu. You might choose to add any other relevant usernames if you want to give access to Spark Streaming jobs.

6. Click **Save Changes**.

Configuring HTTPS Encryption for the Kudu Master and Tablet Server Web UIs

Use the following steps to enable HTTPS for encrypted connections to the Kudu master and tablet server web UIs.

1. Go to the **Kudu** service.
2. Click the **Configuration** tab.
3. Select **Category** > **Security**.
4. In the Search field, type **TLS/SSL** to show the relevant properties.
5. Edit the following properties according to your cluster configuration:

Field	Usage Notes
Master TLS/SSL Server Private Key File (PEM Format)	Set to the path containing the Kudu master host's private key (PEM-format). This is used to enable TLS/SSL encryption (over HTTPS) for browser-based connections to the Kudu master web UI.

Field	Usage Notes
Tablet Server TLS/SSL Server Private Key File (PEM Format)	Set to the path containing the Kudu tablet server host's private key (PEM-format). This is used to enable TLS/SSL encryption (over HTTPS) for browser-based connections to Kudu tablet server web UIs.
Master TLS/SSL Server Certificate File (PEM Format)	Set to the path containing the signed certificate (PEM-format) for the Kudu master host's private key (set in Master TLS/SSL Server Private Key File). The certificate file can be created by concatenating all the appropriate root and intermediate certificates required to verify trust.
Tablet Server TLS/SSL Server Certificate File (PEM Format)	Set to the path containing the signed certificate (PEM-format) for the Kudu tablet server host's private key (set in Tablet Server TLS/SSL Server Private Key File). The certificate file can be created by concatenating all the appropriate root and intermediate certificates required to verify trust.
Enable TLS/SSL for Master Server	Enables HTTPS encryption on the Kudu master web UI.
Enable TLS/SSL for Tablet Server	Enables HTTPS encryption on the Kudu tablet server Web UIs.

6. Click **Save Changes**.

Configuring a Secure Kudu Cluster using the Command Line

> **Important:** Follow these command-line instructions on systems that do not use Cloudera Manager. If you are using Cloudera Manager, see Configuring a Secure Kudu Cluster using Cloudera Manager on page 362.

The following configuration parameters should be set on all servers (master and tablet servers) to ensure that a Kudu cluster is secure:

```
# Connection Security
#--------------------
--rpc_authentication=required
--rpc_encryption=required
--keytab_file=<path-to-kerberos-keytab>

# Web UI Security
#--------------------
--webserver_certificate_file=<path-to-cert-pem>
--webserver_private_key_file=<path-to-key-pem>
# optional
--webserver_private_key_password_cmd=<password-cmd>

# If you prefer to disable the web UI entirely:
--webserver_enabled=false

# Coarse-grained authorization
#--------------------------------

# This example ACL setup allows the 'impala' user as well as the
# 'etl_service_account' principal access to all data in the
# Kudu cluster. The 'hadoopadmin' user is allowed to use administrative
# tooling. Note that by granting access to 'impala', other users
# may access data in Kudu via the Impala service subject to its own
# authorization rules.
--user_acl=impala,etl_service_account
--admin_acl=hadoopadmin
```

More information about these flags can be found in the configuration reference documentation.

Security How-To Guides

Configuring security for clusters can be complex and time-consuming. Here are some start-to-finish guides and single-focused instruction sets aimed at simplifying various tasks.

Administrative Basics

- Check Cluster Security Settings
- Configure Antivirus Software on CDH Hosts
- Log a Security Support Case

Authentication

- Authenticate Kerberos Principals Using Java
- Configure Browser-based Interfaces to Require Kerberos Authentication (SPNEGO)
- Configure Clusters for Kerberos Authentication (Start to finish)

Cloud Security

Amazon Web Services (AWS) and Amazon S3

- Amazon Web Services (AWS) Security on page 367
- How to Configure AWS Credentials on page 393
- Configure Authentication for Amazon S3
- Configure Encryption for Amazon S3

Microsoft Azure

- How To Set Up Access to Cloudera EDH or Cloudera Director (Microsoft Azure Marketplace) on page 409

Client Access

- Configure Browsers for Kerberos Authentication (SPNEGO)
- Set Up a Gateway Node to Restrict Access to the Cluster

Data Privacy

- Enable Sensitive Data Redaction

Data in Transit (TLS/SSL) Encryption

- Add Root and Intermediate CAs to Truststore for TLS/SSL
- Configure Encrypted Transport for HBase Data
- Configure Encrypted Transport for HDFS Data
- Configuring TLS Encryption for Cloudera Manager on page 195
- Convert DER, JKS, PEM for TLS/SSL Clients and Services
- Obtain and Deploy Keys and Certificates for TLS/SSL

- Renew and Redistribute Certificates Before Expiration
- Use Self-Signed Certificates for TLS

How to Add Root and Intermediate CAs to Truststore for TLS/SSL

If a signed certificate is from a certificate authority (CA) that does not have certificates in the truststore for whatever reason (internal CA or a public CA not included in the Java truststore, for example), you must explicitly establish trust for the CA, as detailed below. The content of the truststore for the Cloudera Manager Server cluster can be modified (certificates added, for example) without restarting the server. Changes to the truststore are adopted by the system within 10 seconds.

Explicit Trust for Certificates

Before importing the certificate into the keystore of the host system, you must load the root CAs and any intermediate CAs into the truststore.

1. Copy the root and intermediate CA certificates to these locations on the Cloudera Manager Server host:

```
/opt/cloudera/security/pki/rootca.cert.pem
/opt/cloudera/security/pki/intca.cert.pem
```

   a. For concatenated files containing root CA and intermediate CA certificates, split the file between the `END CERTIFICATE` and `BEGIN CERTIFICATE` boundaries that separate each certificate in the file and make individual files instead.
   b. When extracting multiple intermediate CA certificates from a concatenated file, use unique file names such as `intca-1.cert.pem`, `intca-1.cert.pem`, and so on.

2. Import the root CA certificate into the JDK truststore. If you do not have the `$JAVA_HOME` variable set, replace it with the path to the Oracle JDK.

```
$ sudo keytool -importcert -alias rootca -keystore $JAVA_HOME/jre/lib/security/jssecacerts
 \
-file /opt/cloudera/security/pki/rootca.cert.pem -storepass changeit
```

   The default password for the `cacerts` file is `changeit` (as shown in the above command). Cloudera recommends changing this password by running the keytool command:

```
keytool -storepasswd -keystore $JAVA_HOME/jre/lib/security/cacerts
```

3. Copy the `jssecacerts` file from the Cloudera Manager Server host to all other cluster hosts. Copy the file to the same location on each host using the path required by Oracle JDK, which is as follows:

```
$JAVA_HOME/jre/lib/security/jssecacerts
```

4. On the Cloudera Manager Server host, append the intermediate CA certificate to the signed server certificate. Be sure to use the **append** (>>) operator—not overwrite (>)—when executing the statement:

```
$ sudo cat /opt/cloudera/security/pki/intca.cert.pem >> \
/opt/cloudera/security/pki/$(hostname -f)-server.cert.pem
```

Amazon Web Services (AWS) Security

Amazon Web Services (AWS) is Amazon's cloud solution that offers compute, storage, networking, and other infrastructure services that can be used for Cloudera cluster deployments, whether completely cloud-based or in combination with on-premises clusters.

For example, Amazon Elastic Compute Cloud (EC2) can be used for the instances that make-up the nodes of a Cloudera cluster deployed to the AWS cloud. Amazon's cloud-based storage solution, Amazon Simple Storage Service (Amazon S3), can be used by both on-premises and AWS-cloud-based clusters in various ways, including as storage for Impala tables for direct use by Hue and Hive, and other CDH components such as HDFS client, Hive, Impala, MapReduce.

As of release 5.11, Cloudera Manager supports Amazon's IAM-role based access to Amazon S3 storage, in addition to its prior support of AWS access key and secret key. See How to Configure AWS Credentials on page 393 for details.

For any AWS service, including Amazon S3, you must obtain an Amazon Web Services account and have appropriate access to the AWS Management Console to set up the various services you want, including Amazon S3. Assuming you have an account for AWS, to provide access from your Cloudera cluster to Amazon S3 storage you must configure AWS credentials.

- Configuring Authentication
- Configuring Encryption

Getting Started with Amazon Web Services

To get started with AWS, including Amazon S3, you must have:

1. An Amazon Web Services account. Both Amazon and Cloudera recommend that you do not use your primary Amazon account—known as the root account—for working with Amazon S3 and other AWS services. See the AWS IAM documentation for details about how to set up your AWS account.
2. Access to the AWS Management Console and appropriate permissions to create and configure the AWS services needed for your use case, such as the following:
   a. AWS Elastic Compute Cloud (EC2) to deploy your cluster to the AWS cloud.
   b. AWS Identity and Access Management (IAM) to set up users and groups, or to set up an IAM role.
   c. Amazon S3 and the specific storage bucket (or buckets) for use with your cluster.
   d. Amazon DynamoDB to enable the database needed by Cloudera S3Guard, if you plan to enable S3Guard for your cluster. Cloudera S3Guard augments Amazon S3 with a database to track metadata changes so that the 'eventual consistency' model inherent to Amazon S3 does not pose a problem for transactions or other use cases in which changes may not be apparent to each other in real time. See Configuring and Managing S3Guard in Cloudera Administration for details. To use S3Guard, you will also need to set up the appropriate access policy (create table, read, write) to DynamoDB for the same AWS identity that owns the Amazon S3 storage.
   e. AWS Key Management Services (KMS) (AWS KMS) to create encryption keys for your Amazon S3 bucket if you plan to use SSE-KMS for server-side encryption (not necessary for SSE-S3 encryption. See How to Configure Encryption for Amazon S3 for details).

Configuration Properties Reference

This table provides reference documentation for the `core-site.xml` properties relevant for use with AWS and Amazon S3.

Property	Description
fs.s3a. server-side-encryption-algorithm	Enable server-side encryption for the Amazon S3 storage bucket associated with the cluster. Allowable values: • **`AES256`** Specifies SSE-S3 server-side encryption for Amazon S3. • **`SSE-KMS`** Specifies SSE-KMS server-side encryption for Amazon S3. Requires adding the `fs.s3a.server-side-encryption-key` property with a valid value.

Property	Description
fs.s3a.server-side-encryption-key	Specify the ARN, ARN plus alias, Alias, or globally unique ID of the key created in AWS Key Management Service for use with SSE-KMS.
fs.s3a.awsAccessKeyId	Specify the AWS access key ID. This property is irrelevant and not used to access Amazon S3 storage from a cluster launched using an IAM role.
fs.s3a.awsSecretAccessKey	Specify the AWS secret key provided by Amazon. This property is irrelevant and not used to access Amazon S3 storage from a cluster launched using an IAM role.
fs.s3a.endpoint	Use this property only if the endpoint is outside the standard region (s3.amazonaws.com), such as regions and endpoints in China or in the US GovCloud. See AWS regions and endpoints documentation for details.
fs.s3a.connection.ssl.enabled	Enables (`true`) and disables (`false`) TLS/SSL connections to Amazon S3. Default is `true`.

Connecting to Amazon S3 Using TLS

The boolean parameter `fs.s3a.connection.ssl.enabled` in `core-site.xml` controls whether the `hadoop-aws` connector uses TLS when communicating with Amazon S3. Because this parameter is set to `true` by default, you do not need to configure anything to enable TLS. If you are not using TLS on Amazon S3, the connector will automatically fall back to a plaintext connection.

The root Certificate Authority (CA) certificate that signed the Amazon S3 certificate is trusted by default. If you are using custom truststores, make sure that the configured truststore for each service trusts the root CA certificate.

To import the root CA certificate into your custom truststore, run the following command:

```
$JAVA_HOME/bin/keytool -importkeystore -srckeystore $JAVA_HOME/jre/lib/security/cacerts
 -destkeystore /path/to/custom/truststore -srcalias baltimorecybertrustca
```

If you are using S3Guard, you must import an additional CA certificate:

```
$JAVA_HOME/bin/keytool -importkeystore -srckeystore $JAVA_HOME/jre/lib/security/cacerts
 -destkeystore /path/to/custom/truststore -srcalias verisignclass3g5ca
```

If you are using JDK 1.7.0_131 or higher, or JDK 1.8.0_131 or higher, the aliases are appended with `[jdk]` as follows:

```
$JAVA_HOME/bin/keytool -importkeystore -srckeystore $JAVA_HOME/jre/lib/security/cacerts
 -destkeystore /path/to/custom/truststore -srcalias "baltimorecybertrustca [jdk]"
$JAVA_HOME/bin/keytool -importkeystore -srckeystore $JAVA_HOME/jre/lib/security/cacerts
 -destkeystore /path/to/custom/truststore -srcalias "verisignclass3g5ca [jdk]"
```

If you do not have the `$JAVA_HOME` variable set, replace it with the path to the Oracle JDK (for example, `/usr/java/jdk1.7.0_67-cloudera/`). When prompted, enter the password for the destination and source truststores. The default password for the Oracle JDK `cacerts` truststore is `changeit`.

The truststore configurations for each service that accesses S3 are as follows:

hadoop-aws Connector

All components that can use Amazon S3 storage rely on the `hadoop-aws` connector, which uses the built-in Java truststore (`$JAVA_HOME/jre/lib/security/cacerts`). To override this truststore, create a truststore named `jssecacerts` in the same directory (`$JAVA_HOME/jre/lib/security/jssecacerts`) on all cluster nodes. If you are using the `jssecacerts` truststore, make sure that it includes the root CA certificate that signed the Amazon S3 certificate.

Hive/Beeline CLI

The Hive and Beeline command line interfaces (CLI) rely on the HiveServer2 truststore. To view or modify the truststore configuration:

1. Go to the **Hive** service in the Cloudera Manager Admin Interface.
2. Select the **Configuration** tab.
3. Select **Scope** > ***HIVE-1* (Service-Wide)**.
4. Select **Category** > **Security**.
5. Locate the **HiveServer2 TLS/SSL Certificate Trust Store File** and **HiveServer2 TLS/SSL Certificate Trust Store Password** properties or search for them by typing `Trust` in the **Search** box.

Impala Shell

The Impala shell uses the `hadoop-aws` connector truststore. To override it, create the `$JAVA_HOME/jre/lib/security/jssecacerts` file, as described in hadoop-aws Connector on page 368.

Hue S3 File Browser

For instructions on enabling the S3 file browser in Hue, see How to Enable S3 Cloud Storage in Hue. The S3 file browser uses TLS if it is enabled, and the S3 File Browser trusts the S3 certificate by default. No additional configuration is necessary.

Impala Query Editor (Hue)

The Impala query editor in Hue uses the `hadoop-aws` connector truststore. To override it, create the `$JAVA_HOME/jre/lib/security/jssecacerts` file, as described in hadoop-aws Connector on page 368.

Hive Query Editor (Hue)

The Hive query editor in Hue uses the HiveServer2 truststore. For instructions on viewing and modifying the HiveServer2 truststore, see Hive/Beeline CLI on page 368.

How to Authenticate Kerberos Principals Using Java

To authenticate Kerberos principals in custom applications using Java, import the UserGroupInformation class from the Hadoop security library into your application and use it to pass principals and keytabs to the Kerberos service. This basic example shows hard-coded passing of principal `cloudera` and a `cloudera.keytab` file:

```
// Authenticating Kerberos principal
System.out.println("Principal Authentication: ");
final String user = "cloudera@CLOUDERA.COM";
final String keyPath = "cloudera.keytab";
UserGroupInformation.loginUserFromKeytab(user, keyPath);
```

The UserGroupInformation class (org.apache.hadoop.security.UserGroupInformation) has methods to handle tokens, proxy users, and many other tasks required for your Java application to work with Kerberos. See the Apache API documentation for details.

How to Check Security Settings on a Cluster

Quickly perform a high level check of your cluster's security configuration by doing one of the following:

Check Security for Cloudera Manager Clusters

Use Cloudera Manager to verify security mechanisms for your cluster by simply examining the properties for the cluster.

For clusters not managed by Cloudera Manager Server, see Check Security for CDH Clusters on page 371.

To check Kerberos and HDFS encryption:

1. Log into the Cloudera Manager Admin Console.
2. Select **Security** from the **Administration** drop-down selector to display a list of managed clusters:

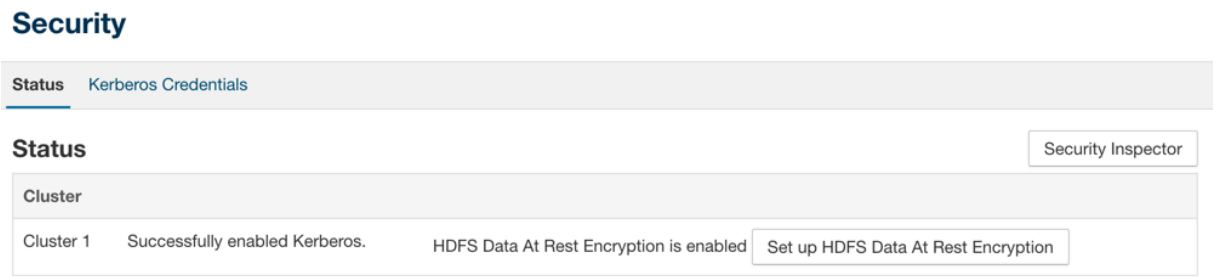

This shows at a glance that both Kerberos and HDFS transparent encryption have been configured for this cluster.

To check TLS settings:

1. Select **Settings** from the **Administration** drop-down selector to open a search field.
2. Enter **TLS** in the search field to display all TLS related configuration settings.
3. Scroll through the displayed results, looking for "Use TLS..." for various services and processes. For example, the test system shown below is not using TLS for the Cloudera Manager Admin Console:

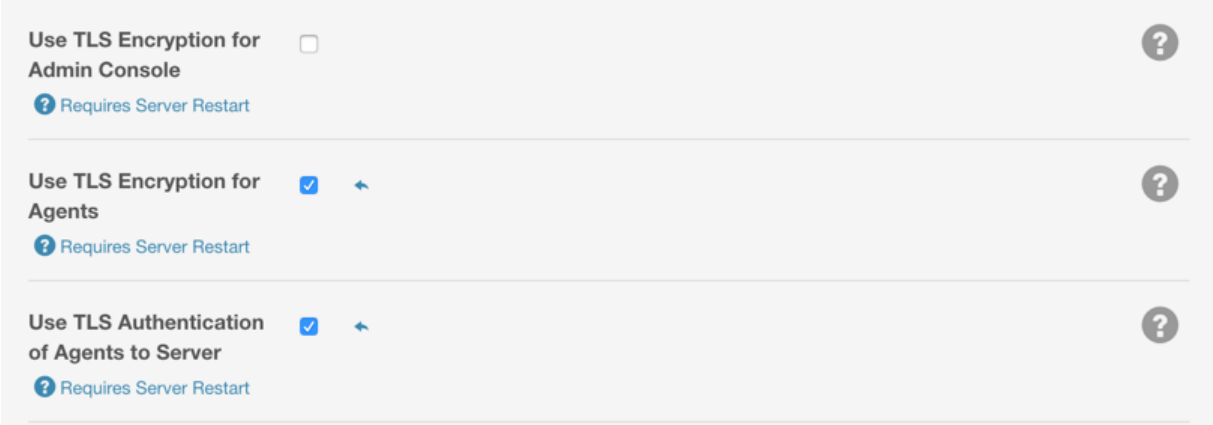

See How to Configure TLS Encryption for Cloudera Manager for complete information about configuring TLS for the cluster.

To find all TLS settings, cluster-wide, enter **"TLS enabled"** (or simply, "TLS") in the top-most search field on the Cloudera Manager Admin Console. Then you can easily select from among the display list to examine the actual setting.

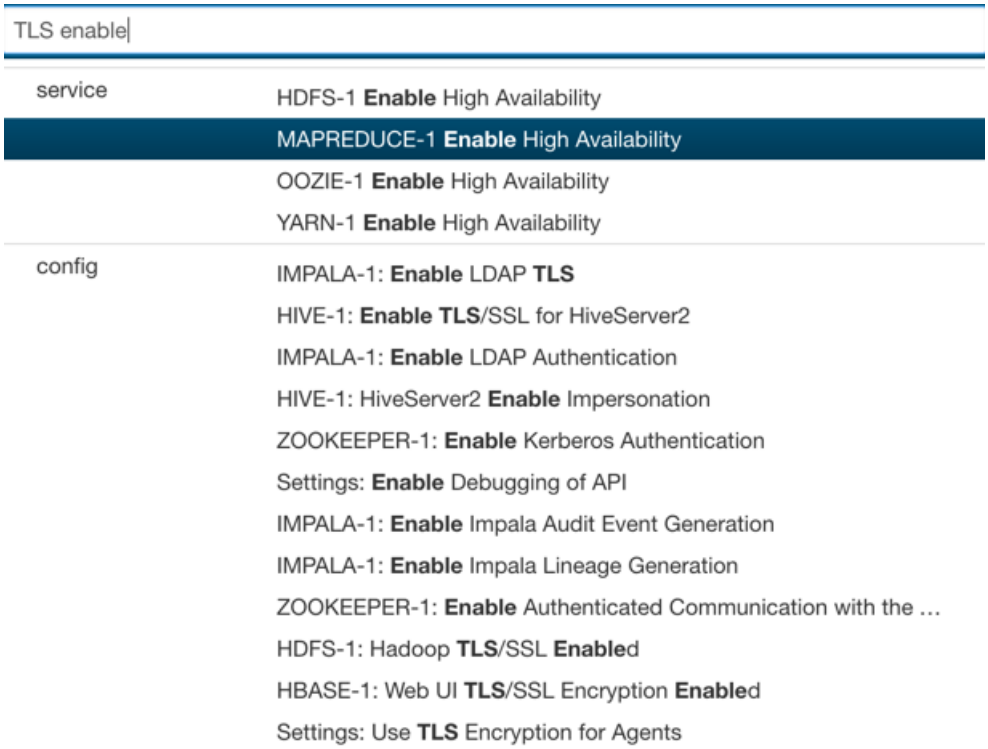

Check Security for CDH Clusters

To check security settings for CDH components not managed by Cloudera Manager, open the configuration file (`core-site.xml`) in a text editor and examine the property values shown below:

Functionality	Property	Value
TLS	hadoop.ssl.enabled	true
Kerberos	hadoop.security.authentication	kerberos
	hadoop.security.authorization	true

See Configuring Authentication in CDH Using the Command Line and Configuring TLS/SSL Encryption for CDH Services for more information.

How to Use Antivirus Software on CDH Hosts

If you use antivirus software on your servers, consider configuring it to skip scans on certain types of Hadoop-specific resources. It can take a long time to scan large files or directories with a large number of files. In addition, if your antivirus software locks files or directories as it scans them, those resources will be unavailable to your Hadoop processes during the scan, and can cause latency or unavailability of resources in your cluster. Consider skipping scans on the following types of resources:

- Scratch directories used by services such as Impala
- Log directories used by various Hadoop services
- Data directories which can grow to petabytes in size

The specific directory names and locations depend on the services your cluster uses and your configuration. In general, avoid scanning very large directories and filesystems. Instead, limit write access to these locations using security mechanisms such as access controls at the level of the operating system, HDFS, or at the service level.

How to Configure Browser-based Interfaces to Require Kerberos Authentication

Required Role: Configurator, Cluster Administrator, or Full Administrator

Access to the web (browser-based) consoles for HTTP services, such as HDFS, MapReduce, and YARN roles can be restricted to only those users with valid Kerberos credentials. After configuring the cluster services as detailed below, the web service and the browser negotiate the authentication process using SPNEGO. After following the steps detailed below for any specific service, all browsers that connect to the service over HTTP must have their configurations modified to handle SPNEGO. See How to Configure Browsers for Kerberos Authentication

To require authentication for HTTP services on the cluster, including web-based user interfaces, such as the Cloudera Manager Admin Console, among others:

1. From the **Clusters** tab, select the service (HDFS, MapReduce, or YARN) for which you want to enable authentication.
2. Click the **Configuration** tab.
3. Under the **Scope** filter, click ***service_name* (Service-Wide)**.
4. Under the **Category** filter, click **Security** to display the security configuration options.
5. In the **Enable Kerberos Authentication for HTTP Web-Consoles** setting, click the box to activate authentication requirement for the selected *service_name* (Service-Wide).
6. Click **Save Changes** to save the change.

Cloudera Manager generates new credentials for the service. When the process finishes, restart all roles for that service.

Repeat this process for the other services for which you want to require Kerberos authentication.

Configure your browser to support authentication to the HTTP services by following the appropriate steps for your browser in How to Configure Browsers for Kerberos Authentication on page 372.

How to Configure Browsers for Kerberos Authentication

The browser configurations below are required only for those browsers used to connect to component web interfaces with the Require Authentication for HTTP Web Consoles configuration property enabled. The settings below enable the respective browser to use SPNEGO to negotiate Kerberos authentication for the browser. The host running the browser must have a valid TGT to authenticate to Kerberos Web Consoles.

Mozilla Firefox

1. Open the low level Firefox configuration page by loading the `about:config` page.
2. In the **Search** text box, enter: `network.negotiate-auth.trusted-uris`
3. Double-click the `network.negotiate-auth.trusted-uris` preference and enter the hostname or the domain of the web server that is protected by Kerberos HTTP SPNEGO. Separate multiple domains and hostnames with a comma.
4. Click **OK**.

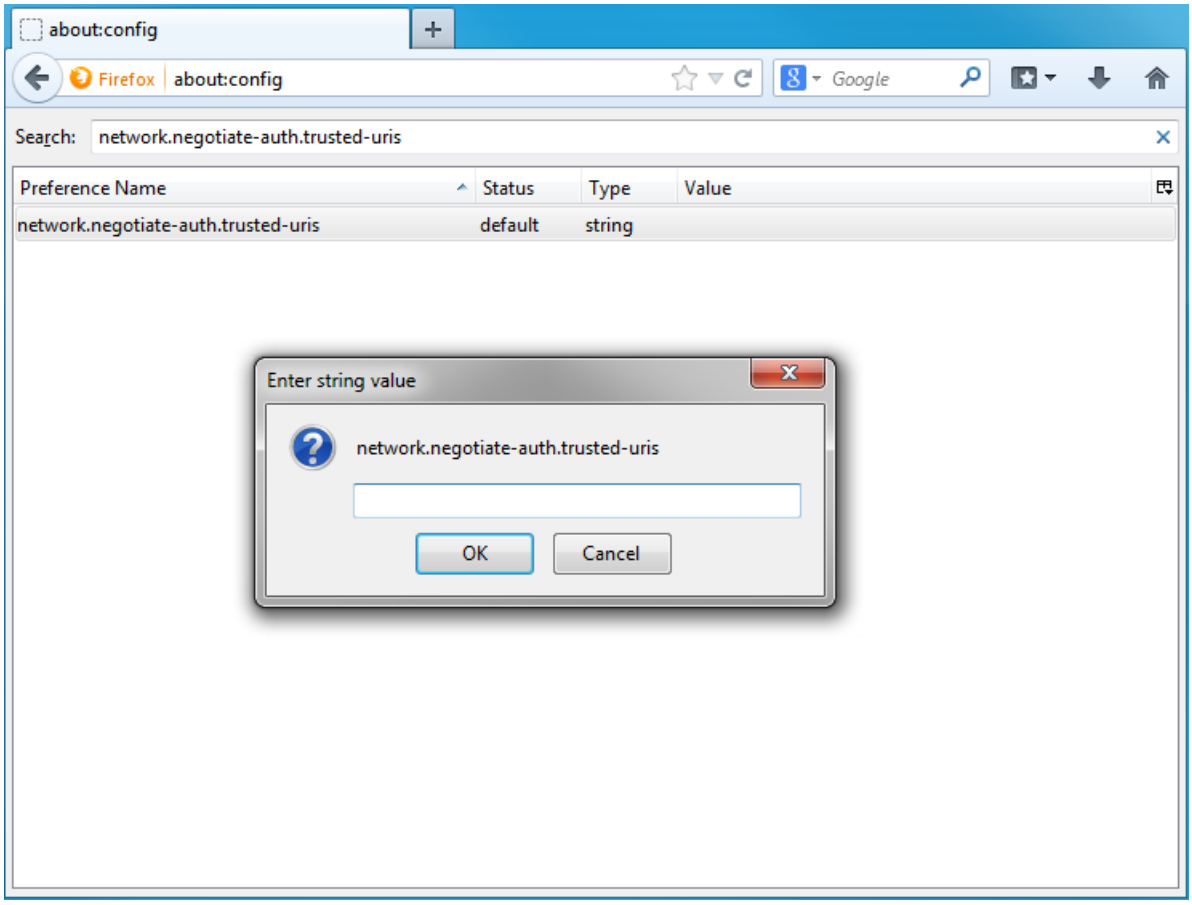

Internet Explorer

Follow the steps below to configure Internet Explorer.

Configuring the Local Intranet Domain

1. Open Internet Explorer and click the Settings gear icon in the top-right corner. Select **Internet options**.
2. Select the **Security** tab.
3. Select the **Local Intranet** zone and click the **Sites** button.
4. Make sure that the first two options, **Include all local (intranet) sites not listed in other zones** and **Include all sites that bypass the proxy server** are checked.
5. Click **Advanced** and add the names of the domains that are protected by Kerberos HTTP SPNEGO, one at a time, to the list of websites. For example, `myhost.example.com`. Click **Close**.
6. Click **OK** to save your configuration changes.

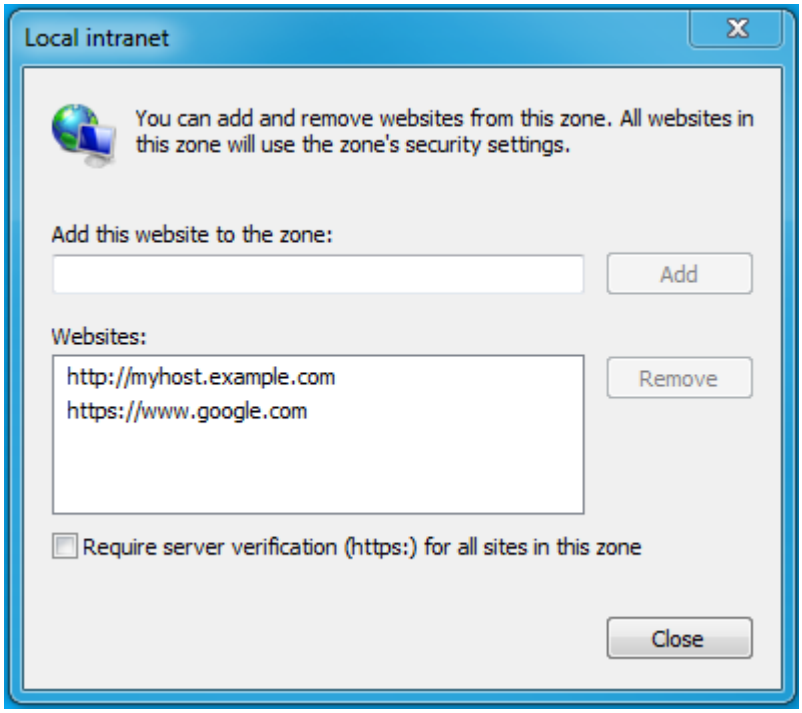

Configuring Intranet Authentication

1. Click the Settings gear icon in the top-right corner. Select **Internet options**.
2. Select the **Security** tab.
3. Select the **Local Intranet** zone and click the **Custom level...** button to open the **Security Settings - Local Intranet Zone** dialog box.
4. Scroll down to the **User Authentication** options and select **Automatic logon only in Intranet zone**.
5. Click **OK** to save these changes.

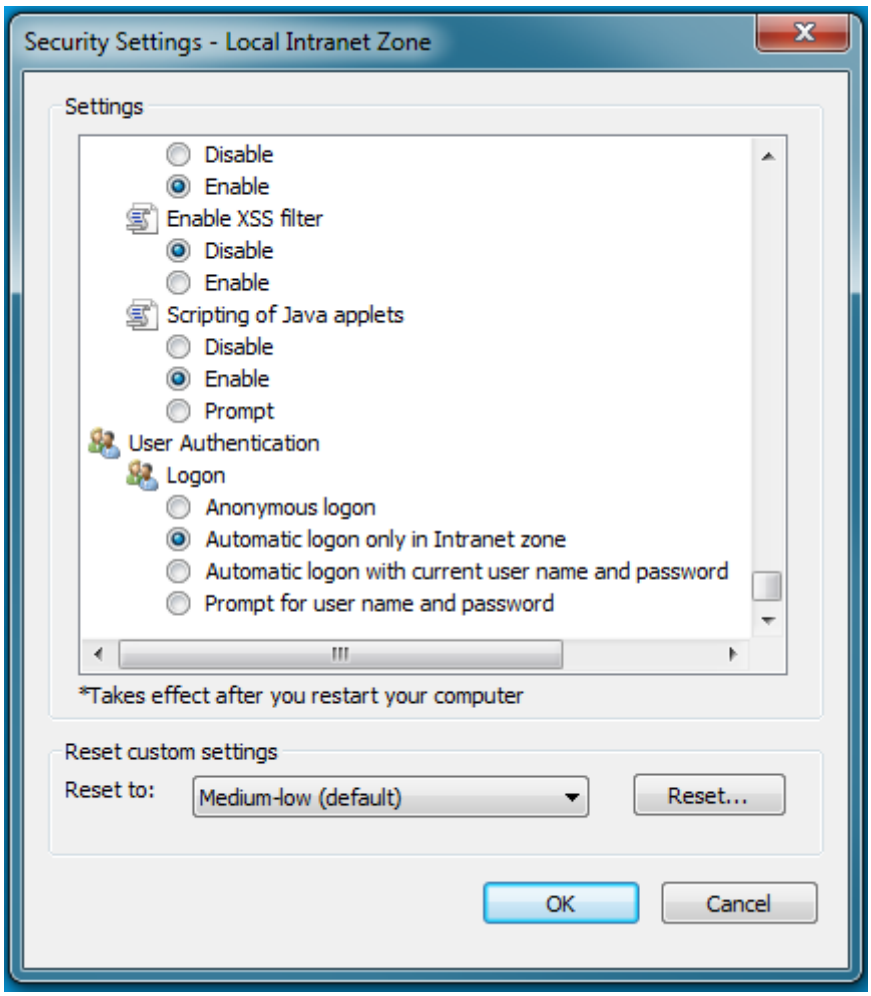

Verifying Proxy Settings

Perform these steps only if you have a proxy server already enabled.

1. Click the Settings gear icon in the top-right corner. Select **Internet options**.
2. Select the **Connections** tab and click **LAN Settings**.
3. Verify that the proxy server **Address** and **Port** number settings are correct.
4. Click **Advanced** to open the **Proxy Settings** dialog box.
5. Add the Kerberos-protected domains to the **Exceptions** field.
6. Click **OK** to save any changes.

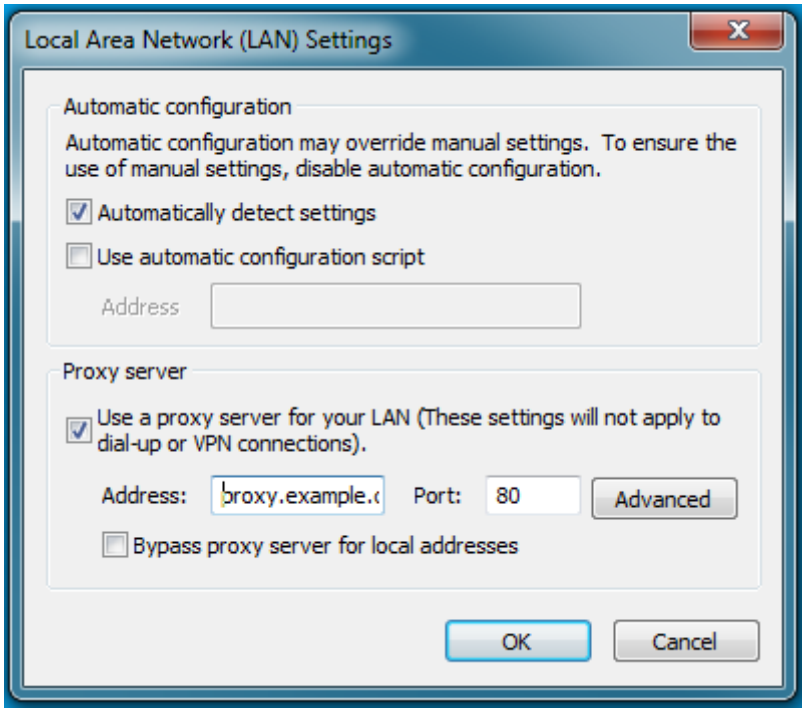

Google Chrome

For Windows:

- Open the Control Panel to access the **Internet Options** dialog. Use the same configuration as detailed in Configuration changes required are the same as those described above for Internet Explorer.

For Linux or MacOS:

- Add the `--auth-server-whitelist` parameter to the `google-chrome` command. For example, to run Chrome from a Linux prompt, run the `google-chrome` command as follows:

```
> google-chrome --auth-server-whitelist = "hostname/domain"
```

How to Configure Clusters to Use Kerberos for Authentication

Cloudera clusters can use Kerberos to authenticate services running on the cluster and the users who need access to those services. This How To guide provides the requirements, pre-requisites, and high-level summary of the steps needed to integrate clusters with Kerberos for authentication.

> **Note:** For clusters deployed using Cloudera Manager Server, Cloudera recommends using the Kerberos configuration wizard available through the Cloudera Manager Admin Console. See Enabling Kerberos Authentication Using the Wizard for details.

The following are the general steps for integrating Kerberos with Cloudera clusters without using the Cloudera Manager configuration wizard.

Step 1: Verify Requirements and Assumptions

The steps outlined below assume that:

- The Kerberos instance has been setup, is running, and is available during the configuration process.

- The Cloudera cluster has been installed and is operational, with all services fully-functional—Cloudera Manager Server, CDH, and Cloudera Manager Agent processes on all cluster nodes.

Hosts Configured for AES-256 Encryption

By default, CentOS and RHEL 5.5 (and higher) use AES-256 encryption for Kerberos tickets. If you use either of these platforms for your cluster, the Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction Policy File must be installed on all cluster hosts.

To install the JCE Policy file on the host system at the OS layer:

1. Download the `jce_policy-x.zip`
2. Unzip the file
3. Follow the steps in the `README.txt` to install it.

> **Note:** The AES-256 encryption can also be configured on a running cluster by using Cloudera Manager Admin Console. See To use Cloudera Manager to install the JCE policy file for details.

Required Administrator Privileges

Setting up the Cloudera cluster to use Kerberos for authentication requires complete administrator access to the cluster and administrator privileges on the Kerberos instance:

- **Cluster Administrator** or **Full Administrator**
- Kerberos administrator privileges:

```
someone/admin@YOUR-DOMAIN.FQDN.EXAMPLE.COM
```

If you do not have administrator privileges on the Kerberos instance, you will need help from the Kerberos administrator before you can complete the process.

Required User (Service Account) Directories

During installation, the `cloudera-scm` account is created on the host system. When Cloudera Manager and CDH services are installed at the same time, Cloudera Manager creates other accounts as needed to support the service role daemons. However, if the CDH services and Cloudera Manager are installed separately, you may need to specifically set directory permissions for certain Hadoop user (service daemon) accounts for successful integration with Kerberos. The following table shows the accounts used for core service roles. Note that `hdfs` acts as superuser for the system.

User	Service Roles
`hdfs`	NameNode, DataNodes, Secondary NameNode (and HDFS superuser)
`mapred`	JobTracker, TaskTrackers (MR1), Job History Server (YARN)
`yarn`	ResourceManager, NodeManager (YARN)
`oozie`	Oozie Server
`hue`	Hue Server, Beeswax Server, Authorization Manager, Job Designer

These accounts require ownership control over specific directories.

- For newly installed Cloudera clusters (Cloudera Manager and CDH installed at the same time)—The Cloudera Manager Agent process on each cluster host automatically configures the appropriate directory ownership when the cluster launches.
- For existing CDH clusters using HDFS and running MapReduce jobs prior to Cloudera Manager installation—The directory ownership must be manually configured, as shown in the table below. The directory owners cannot differ from those shown in the table to ensure that the service daemons can set permissions as needed on each directory.

Directory Specified in this Property	Owner
`dfs.name.dir`	`hdfs:hadoop`
`dfs.data.dir`	`hdfs:hadoop`
`mapred.local.dir`	`mapred:hadoop`
`mapred.system.dir` in HDFS	`mapred:hadoop`
`yarn.nodemanager.local-dirs`	`yarn:yarn`
`yarn.nodemanager.log-dirs`	`yarn:yarn`
`oozie.service.StoreService.jdbc.url` (if using Derby)	`oozie:oozie`
`[[database]] name`	`hue:hue`
`javax.jdo.option.ConnectionURL`	`hue:hue`

Step 2. Create Principal for Cloudera Manager Server in the Kerberos KDC

Cloudera Manager Server has its own principal to connect to the Kerberos KDC and import user and service principals for use by the cluster.

The steps below summarize the process of adding a principal specifically for Cloudera Manager Server to an MIT KDC and an Active Directory KDC. See documentation from MIT, Microsoft, or the appropriate vendor for more detailed information.

> **Note:** If an administrator principal to act on behalf of Cloudera Manager cannot be created on the Kerberos KDC for whatever reason, Cloudera Manager will not be able to create or manage principals and keytabs for CDH services. That means these principals must be created manually on the Kerberos KDC and then imported (retrieved) by Cloudera Manager. See Using a Custom Kerberos Keytab Retrieval Script on page 67 for details about this process.

Creating a Principal in Active Directory

Check your Microsoft documentation for specific details for your Active Directory KDC. The general process is as follows:

1. Create an Organizational Unit (OU) in your Active Directory KDC service that will contain the principals for use by the CDH cluster.
2. Add a new user account to Active Directory, for example, *`username`*`@YOUR-REALM.EXAMPLE.COM`. Set the password for the user to never expire.
3. Use the Delegate Control wizard of Active Directory and grant this new user permission to **Create, Delete, and Manage User Accounts**.

Creating a Principal in an MIT KDC

For MIT Kerberos, user principals that include the instance name `admin` designate a user account with administrator privileges. For example:

```
username/admin@YOUR-REALM.EXAMPLE.COM
```

Create the Cloudera Manager Server principal as shown in one of the examples below, appropriate for the location of the Kerberos instance and using the correct REALM name for your setup.

For MIT Kerberos KDC on a remote host:

```
kadmin: addprinc -pw password cloudera-scm/admin@YOUR-REALM.EXAMPLE.COM
```

For MIT Kerberos KDC on a local host:

```
kadmin.local: addprinc -pw password cloudera-scm/admin@YOUR-REALM.EXAMPLE.COM
```

Step 3: Add the Credentials for the Principal to the Cluster

Assuming the principal was successfully added to the Kerberos KDC, it can be added to the cluster as follows:

1. Log in to the Cloudera Manager Admin Console.
2. Select **Administration** > **Security**.
3. Click the **Kerberos Credentials** tab.
4. Click the **Import Kerberos Account Manager Credentials** button.
5. Enter the credentials for the principal added to the Kerberos KDC in the previous step:
   - For **Username**, enter the primary and realm portions of the Kerberos principal. Enter the realm name in all upper-case only (*`YOUR-REALM.EXAMPLE.COM`*) as shown below.
   - Enter the **Password** for the principal.

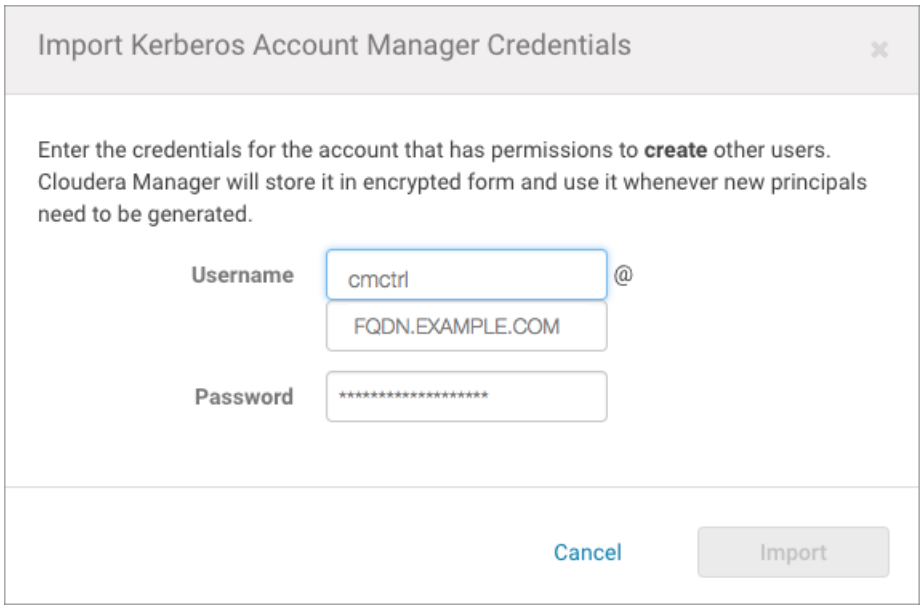

6. Click **Import**.

Cloudera Manager encrypts the username and password into a keytab and uses it to create new principals in the KDC as needed.

Click **Close** when complete.

Step 4: Identify Default Kerberos Realm for the Cluster

Each host in the cluster must have the default realm property (`default_realm`) specified in the `libdefaults` section of its Kerberos configuration file (`/etc/krb5.conf`).

```
[libdefaults]
default_realm = FQDN.EXAMPLE.COM
```

After adding the default realm to the configuration file for all hosts in the cluster, configure the same default realm for Cloudera Manager Server.

In the Cloudera Manager Admin Console:

1. Select **Administration** > **Settings**.

2. Select **Kerberos** for the Category filter.
3. In the **Kerberos Security Realm** field, enter the default realm name set in the Kerberos configuration file (`/etc/krb5.conf`) of each host in the cluster. For example:

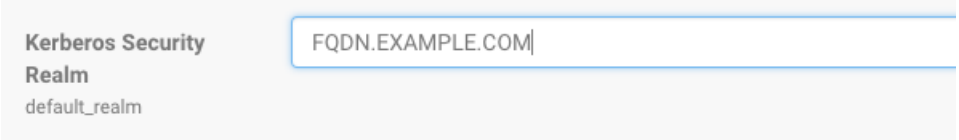

4. Click **Save Changes**.

Step 5: Stop all Services

All service daemons in the cluster must be stopped so that they can be restarted at the same time and start as authenticated services in the cluster. Service daemons running without authenticating to Kerberos first will not be able to communicate with other daemons in the cluster that have authenticated to Kerberos, so they must be shut down and restarted at the end of the configuration process, as a unit.

> **Note:** The requirement to stop all daemons prevents using the rolling upgrade process to enable Kerberos integration on the cluster.

Stop all running services and the Cloudera Management Service as follows:

In the Cloudera Manager Admin Console:

1. Select **Clusters** > **Cluster-*n***.
2. Click the **Actions** drop-down menu and select **Stop** to stop all services on the cluster.
3. Click **Stop** on the warning message to stop all services on the cluster. The Command Details window displays the progress as each service shuts down. When the message **All services successfully stopped** displays, close the Command Details window.
4. Select > **Clusters** > **Cloudera Management Service**.
5. Click the **Actions** drop-down menu and select **Stop** to stop the Cloudera Management Service. The Command Details window displays the progress as each role instance running on the Cloudera Management Service shuts down. The process is completed when the message **Command completed with n/n successful subcommands** displays.
6. Click **Close**.

Step 6. Specify Kerberos for Security

Kerberos must be specified as the security mechanism for Hadoop infrastructure, starting with the HDFS service. Enable Cloudera Manager Server security for the cluster on an HDFS service. After you do so, the Cloudera Manager Server automatically enables Hadoop security on the MapReduce and YARN services associated with that HDFS service.

In the Cloudera Manager Admin Console:

1. Select **Clusters** > **HDFS-*n***.
2. Click the **Configuration** tab.
3. Select **HDFS-*n*** for the **Scope** filter.
4. Select **Security** for the **Category** filter.
5. Scroll (or search) to find the **Hadoop Secure Authentication** property.
6. Click the **kerberos** button to select Kerberos:

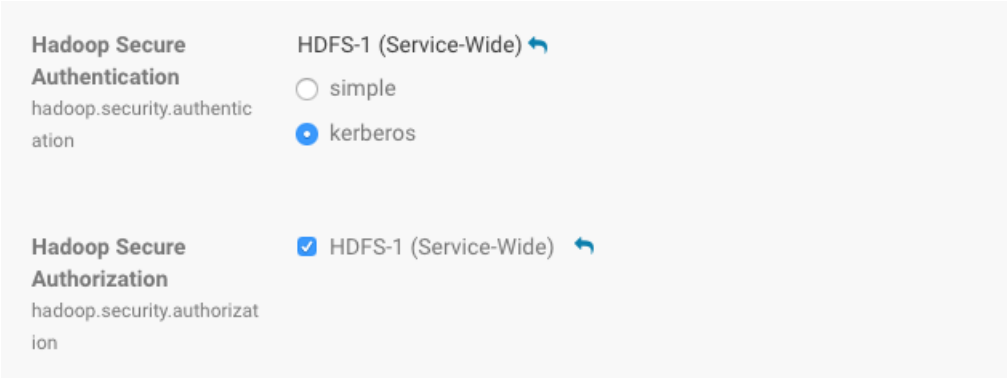

7. Click the value for the **Hadoop Secure Authorization** property and select the checkbox to enable service-level authorization on the selected HDFS service. You can specify comma-separated lists of users and groups authorized to use Hadoop services or perform admin operations using the following properties under the **Service-Wide Security** section:
   - **Authorized Users**: Comma-separated list of users authorized to use Hadoop services.
   - **Authorized Groups**: Comma-separated list of groups authorized to use Hadoop services.
   - **Authorized Admin Users**: Comma-separated list of users authorized to perform admin operations on Hadoop.
   - **Authorized Admin Groups**: Comma-separated list of groups authorized to perform admin operations on Hadoop.

   > **Important:** For Cloudera Manager's Monitoring services to work, the hue user should always be added as an authorized user.

8. In the Search field, type **DataNode Transceiver** to find the **DataNode Transceiver Port** property.
9. Click the value for the **DataNode Transceiver Port** property and specify a privileged port number (below 1024). Cloudera recommends 1004.

   > **Note:** If there is more than one DataNode Role Group, you must specify a privileged port number for each DataNode Transceiver Port property.

10. In the Search field, type **DataNode HTTP** to find the **DataNode HTTP Web UI Port** property and specify a privileged port number (below 1024). Cloudera recommends 1006.

    > **Note:** The port numbers for the two DataNode properties must be below 1024 to provide part of the security mechanism to make it impossible for a user to run a MapReduce task that impersonates a DataNode. The port numbers for the NameNode and Secondary NameNode can be anything you want, but the default port numbers are good ones to use.

11. In the Search field type **Data Directory Permissions** to find the **DataNode Data Directory Permissions** property.
12. Reset the value for the **DataNode Data Directory Permissions** property to the default value of `700` if not already set to that.
13. Make sure you have changed the **DataNode Transceiver Port**, **DataNode Data Directory Permissions** and **DataNode HTTP Web UI Port** properties for every DataNode role group.
14. Click **Save Changes** to save the configuration settings.

To enable ZooKeeper security:

1. Go to the **ZooKeeper Service Configuration** tab and click **View and Edit**.
2. Click the value for **Enable Kerberos Authentication** property.
3. Click **Save Changes** to save the configuration settings.

To enable HBase security:

1. Go to the **HBase Service** > **Configuration** tab and click **View and Edit**.
2. In the Search field, type **HBase Secure** to show the Hadoop security properties (found under the **Service-Wide** > **Security** category).
3. Click the value for the **HBase Secure Authorization** property and select the checkbox to enable authorization on the selected HBase service.
4. Click the value for the **HBase Secure Authentication** property and select `kerberos` to enable authorization on the selected HBase service.
5. Click **Save Changes** to save the configuration settings.

To enable Solr security:

1. Go to the **Solr Service** > **Configuration** tab and click **View and Edit**.
2. In the Search field, type **Solr Secure** to show the Solr security properties (found under the **Service-WideSecurity** category).
3. Click the value for the **Solr Secure Authentication** property and select `kerberos` to enable authorization on the selected Solr service.
4. Click **Save Changes** to save the configuration settings.

> **Note:** Using Cloudera Manager Admin Console to generate client configuration files **after** enabling Kerberos does not provide the Kerberos principals and keytabs that users need to authenticate to the cluster. Users must obtain their Kerberos principals from the Kerberos administrator and then run the `kinit` command themselves.

Credentials Generated

After you enable security for any of the services in Cloudera Manager, a command called Generate Credentials will be triggered automatically. You can watch the progress of the command on the top right corner of the screen that shows the running commands. Wait for this command to finish (indicated by a grey box containing "0" in it).

Step 7: Restart All Services

Start all services on the cluster using the Cloudera Manager Admin Console:

1. Select **Clusters** > **Cluster-*n***.
2. Click the **Actions** drop-down button menu and select **Start**. The confirmation prompt displays.
3. Click **Start** to confirm and continue. The **Command Details** window displays progress. When **All services successfully started** displays, close the **Command Details**window.
4. Select > **Clusters** > **Cloudera Management Service**.
5. Click the **Actions** drop-down menu and select **Start**. The Command Details window displays the progress as each role instance running on the Cloudera Management Service starts up. The process is completed when the message **Command completed with *n*/*n* successful subcommands** displays.

Step 8: Deploy Client Configurations

Deploy client configurations for services supported on the cluster using the Cloudera Manager Admin Console:

1. Select **Clusters** > **Cluster-*n***.
2. Click the **Actions** drop-down button menu and select **Deploy Client Configuration**.

Step 9: Create the HDFS Superuser Principal

To be able to create home directories for users, you will need access to the HDFS superuser account. (CDH automatically created the HDFS superuser account on each cluster host during CDH installation.) When you enabled Kerberos for the HDFS service, you lost access to the default HDFS superuser account using `sudo -u hdfs` commands. Cloudera recommends you use a different user account as the superuser, not the default `hdfs` account.

Designating a Non-Default Superuser Group

To designate a different group of superusers instead of using the default `hdfs` account, follow these steps:

1. Go to the Cloudera Manager Admin Console and navigate to the HDFS service.
2. Click the **Configuration** tab.
3. Select **Scope** > **HDFS (Service-Wide)**.
4. Select **Category** > **Security**.
5. Locate the **Superuser Group** property and change the value to the appropriate group name for your environment. For example, *superuser*.
6. Click **Save Changes**.
7. Restart the HDFS service.

   To enable your access to the superuser account now that Kerberos is enabled, you must now create a Kerberos principal or an Active Directory user whose first component is *superuser*:

For Active Directory

Add a new user account to Active Directory, *`superuser@YOUR-REALM.EXAMPLE.COM`* . The password for this account should be set to never expire.

For MIT KDC

1. In the `kadmin.local` or `kadmin` shell, type the following command to create a Kerberos principal called *superuser*:

```
kadmin:  addprinc superuser@YOUR-REALM.EXAMPLE.COM
```

   This command prompts you to create a password for the *superuser* principal. You should use a strong password because having access to this principal provides superuser access to all of the files in HDFS.
2. To run commands as the HDFS superuser, you must obtain Kerberos credentials for the *superuser* principal. To do so, run the following command and provide the appropriate password when prompted.

```
$ kinit superuser@YOUR-REALM.EXAMPLE.COM
```

Step 10: Get or Create a Kerberos Principal for Each User Account

Now that Kerberos is configured and enabled on your cluster, you and every other Hadoop user must have a Kerberos principal or keytab to obtain Kerberos credentials to be allowed to access the cluster and use the Hadoop services. In the next step of this procedure, you need to create your own Kerberos principals to verify that Kerberos security is working on your cluster. If you and the other Hadoop users already have a Kerberos principal or keytab, or if your Kerberos administrator can provide them, you can skip ahead to the next step.

To create Kerberos principals for all users:

Active Directory

Add a new AD user account, `<username>@EXAMPLE.COM` for each Cloudera Manager service that should use Kerberos authentication. The password for these service accounts should be set to never expire.

MIT KDC

1. In the kadmin.local or kadmin shell, use the following command to create a principal for your account by replacing `EXAMPLE.COM` with the name of your realm, and replacing `username` with a username:

```
kadmin:  addprinc username@EXAMPLE.COM
```

2. When prompted, enter the password twice.

Step 11: Prepare the Cluster for Each User

Before you and other business users can access the cluster, the hosts must be prepared for each user. Perform the following tasks to prepare the hosts for each user:

1. Each host in the cluster must have a Unix user account with the same name as primary component of the user's principal name. For example, the principal `jcarlos@YOUR-REALM.EXAMPLE.COM` needs the Linux account `jcarlos` on each host system. Use LDAP (OpenLDAP, Microsoft Active Directory) for this step if possible.

   > **Note:** Each account must have a user ID that is greater than or equal to 1000. In the `/etc/hadoop/conf/taskcontroller.cfg` file, the default setting for the `banned.users` property is `mapred`, `hdfs`, and `bin` to prevent jobs from being submitted from those user accounts. The default setting for the `min.user.id` property is `1000` to prevent jobs from being submitted with a user ID less than 1000, which are conventionally Unix super users.

2. Create a subdirectory under `/user` on HDFS for each user account (for example, `/user/jcarlos`). Change the owner and group of that directory to be the user.

```
$ hadoop fs -mkdir /user/jcarlos
$ hadoop fs -chown jcarlos /user/jcarlos
```

> **Note:** The commands above do not include `sudo -u hdfs` because it is not required with Kerberos configured for the cluster (assuming you created the Kerberos credentials for the HDFS super user as detailed in Step 9: Create the HDFS Superuser Principal).

Step 12: Verify Successful Kerberos Integration

To verify that Kerberos has been successfully integrated for the cluster, try running one of the sample MapReduce jobs (`sleep` or `pi`, for example) provided at:

```
/usr/lib/hadoop/hadoop-examples.jar
```

This assumes you have the fully-functional cluster as recommended in Step 1: Verify Requirements and Assumptions and that the client configure files have been generated as detailed in Step 8: Deploying Client Configuration Files.

To verify that Kerberos security is working:

1. Obtain Kerberos credentials for your user account.

```
$ kinit
            youruserid@YOUR-REALM.EXAMPLE.COM
```

2. Enter a password when prompted.
3. Submit a sample pi calculation as a test MapReduce job. Use the following command if you use a package-based setup for Cloudera Manager:

```
$ hadoop jar /usr/lib/hadoop-0.20/hadoop-0.20.2*examples.jar pi 10 10000
Number of Maps = 10
Samples per Map = 10000
...
Job Finished in 38.572 seconds
Estimated value of Pi is 3.14120000000000000000
```

If you have a parcel-based setup, use the following command instead:

```
$ hadoop jar /opt/cloudera/parcels/CDH/lib/hadoop-0.20-mapreduce/hadoop-examples.jar pi
 10 10000
Number of Maps = 10
Samples per Map = 10000
```

```
...
Job Finished in 30.958 seconds
Estimated value of Pi is 3.14120000000000000000
```

You have now verified that Kerberos security is working on your cluster.

Important:

Running a MapReduce job will fail if you do not have a valid Kerberos ticket in your credentials cache. You can examine the Kerberos tickets currently in your credentials cache by running the `klist` command. You can obtain a ticket by running the `kinit` command and either specifying a keytab file containing credentials, or entering the password for your principal. If you do not have a valid ticket, you will receive an error such as:

```
11/01/04 12:08:12 WARN ipc.Client:
Exception encountered while connecting to the server :
javax.security.sasl.SaslException:GSS initiate failed
[Caused by GSSException: No valid credentials provided (Mechanism level:
 Failed to find any
Kerberos tgt)]
Bad connection to FS. command aborted. exception: Call to
nn-host/10.0.0.2:8020 failed on local exception:
java.io.IOException:javax.security.sasl.SaslException: GSS initiate
failed
[Caused by GSSException: No valid credentials provided
(Mechanism level: Failed to find any Kerberos tgt)]
```

How to Convert File Encodings (DER, JKS, PEM) for TLS/SSL Certificates and Keys

Client and server processes require specific file formats for certificates, keys, and other digital artifacts used for TLS/SSL encryption. For example, when TLS is enabled, Cloudera Manager Server presents Java KeyStore (JKS) formatted key and certificate to requesting Cloudera Manager Agent hosts. The Hue client also connects to Cloudera Manager Server, but Hue requires a PEM-formatted key and certificate, rather than JKS. The PEM format used by Cloudera Manager is PKCS #8, which handles certificates and keys as individual Base64-encoded text files.

If you receive binary DER files from your certificate authority, you must convert them to the appropriate format. Since neither Java Keytool nor OpenSSL work directly with PKCS format, many of the configuration tasks detailed in Configuring TLS Encryption for Cloudera Manager on page 195 involve converting formats, or extracting keys or certificates from an artifact in one format to another.

Certificates issued by a CA in one format (encoding) can be used to create certificates in a different format using Java Keytool and OpenSSL as detailed below:

Converting DER Encoded Certificates to PEM

OpenSSL can be used to convert a DER-encoded certificate to an ASCII (Base64) encoded certificate. Typically, DER-encoded certificates may have file extension of .DER, .CRT, or .CER, but regardless of the extension, a DER encoded certificate is not readable as plain text (unlike PEM encoded certificate).

A PEM-encoded certificate may also have file extension of .CRT or .CER, in which case, you can simply copy the file to a new name using the .PEM extension:

```
$ cp hostname.cer hostname.pem
```

To convert a DER-encoded certificate to PEM encoding using OpenSSL:

```
$ openssl x509 -inform der -in hostname.cer -out hostname.pem
```

For example:

```
$ openssl x509 -inform der -in /opt/cloudera/security/pki/hostname.cer -out
/tmp/hostname.pem
```

Converting JKS Key and Certificate to PEM

This process uses both Java keytool and OpenSSL (`keytool` and `openssl`, respectively, in the commands below) to export the composite private key and certificate from a Java keystore and then extract each element into its own file.

The `PKCS12` file created below is an interim file used to obtain the individual key and certificate files.

Replace *hostname-keystore*, *cmhost*, *hostname*, and *password* with values from your system.

1. Export the private key and certificate command line:

```
$ keytool -importkeystore -srckeystore /opt/cloudera/security/jks/hostname-keystore.jks
 \
-srcstorepass password -srckeypass password -destkeystore /tmp/hostname-keystore.p12 \
-deststoretype PKCS12 -srcalias hostname -deststorepass password -destkeypass password
```

2. Extract the certificate file from the resulting `PKCS12` file:

```
$ openssl pkcs12 -in /tmp/hostname-keystore.p12 -passin pass:password  -nokeys \
-out /opt/cloudera/security/pki/hostname.pem
```

This extracted certificate can be used, as is, but you can additionally extract the private key from the keystore as detailed in the next section.

Extracting the Private Key from PKCS Keystore

Use OpenSSL to extract the private key from the PKCS keystore when needed. The command shown below extracts the key and saves it to a keystore that is protected using the password you provide:

```
$ openssl pkcs12 -in /tmp/hostname-keystore.p12 -passin pass:password \
-nocerts -out /opt/cloudera/security/pki/hostname.key -passout pass:password
```

To generate a key without a password, use this version of the command:

```
$ openssl rsa -in /tmp/hostname-keystore.p12 -passin pass:password \
 -nocerts -out /opt/cloudera/security/pki/hostname.pem
```

Converting PEM Key and Certificate to JKS

Replace `hostname` in the commands below with the FQDN of the host whose certificate is being imported.

1. Convert the `openssl` private key and certificate files into a `PKCS12` file.

```
$ openssl pkcs12 -export -in /opt/cloudera/security/pki/hostname.pem \
-inkey /opt/cloudera/security/pki/hostname.key -out /tmp/hostname.p12 \
-name hostname -passin pass:password -passout pass:password
```

2. Import the PKCS12 file into the Java keystore.

```
$ keytool -importkeystore -srckeystore /tmp/hostname.p12 -srcstoretype PKCS12 \
-srcstorepass password -alias hostname -deststorepass password \
-destkeypass password -destkeystore /opt/cloudera/security/jks/hostname-keystore.jks
```

How To Configure Authentication for Amazon S3

There are several ways to integrate Amazon S3 storage with Cloudera clusters, depending on your use case and other factors, including whether the cluster has been deployed using Amazon EC2 instances and if those instances were deployed using an IAM role, such as might be the case for clusters that have a single-user or small-team with comparable privileges. Clusters deployed to support many different users with various privilege levels to the Amazon S3 need to use AWS Credentials and have privileges to target data set up in Sentry. See How to Configure AWS Credentials on page 393 for details.

Authentication through the S3 Connector Service

Starting with CDH/Cloudera Manager 5.10, integration with Amazon S3 from Cloudera clusters has been simplified. Specifically, the **S3 Connector Service** automates the authentication process to Amazon S3 for Impala, Hive, and Hue, the components used for business-analytical use cases designed to run on persistent multi-tenant clusters.

The S3 Connector Service transparently and securely distributes AWS credentials needed by the cluster for the Amazon S3 storage. Access to the underlying Impala tables is controlled by Sentry role-based permissions. The S3 Connector Service runs on a secure cluster only, that is, a cluster configured to use:

- **Kerberos** for authentication, and
- **Sentry** for role-based authorization.

> **Note:** Of the items listed in the screenshot below, only the Sentry service and Kerberos enabled messages are actual requirements. The other messages are for informational purposes only.

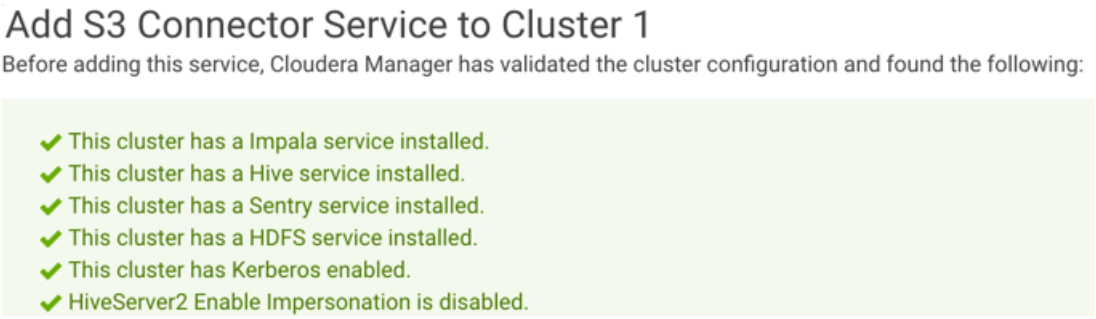

In Cloudera Manager 5.11, the S3 Connector Service setup wizard is launched automatically during the AWS Credential setup process when you select the path to add the S3 Connector Service.

See Configuring the Amazon S3 Connector for more information about the S3 Connector Service.

Authentication through Advanced Configuration Snippets

Before release 5.10 and the introduction of the S3 Connector Service, using Amazon S3 storage with the cluster involved adding the credentials to the `core-site.xml` configuration file (through Cloudera Manager's Advanced Configuration Snippet mechanism). This approach is not recommended. AWS credentials provide read and write access to data stored on Amazon S3, so they should be kept secure at all times.

- Never share the credentials with other cluster users or services.

- Do not store in cleartext in any configuration files. When possible, use Hadoop's credential provider to encrypt and store credentials in the local JCEK (Java Cryptography Extension Keystore).
- Enable Cloudera sensitive data redaction to ensure that passwords and other sensitive information does not appear in log files.

> **Important:** Cloudera recommends using this approach for **single-user clusters** on **secure networks** only—networks that allow access only to authorized users, all of whom are also authorized to use the S3 credentials.

To enable CDH services to access Amazon S3, AWS credentials can be specified using the `fs.s3a.access.key` and `fs.s3a.secret.key` properties:

```
<property>
    <name>fs.s3a.access.key</name>
    <value>your_access_key</value>
</property>

<property>
    <name>fs.s3a.secret.key</name>
    <value>your_secret_key</value>
</property>
```

The process of adding AWS credentials is generally the same as that detailed in configuring server-side encryption for Amazon S3, that is, using the Cloudera Manager Admin Console to add the properties and values to the `core-site.xml` configuration file (Advanced Configuration Snippet). However, Cloudera strongly discourages this approach: in general, adding AWS credentials to the `core-site.xml` is not recommended. A somewhat more secure approach is to use temporary credentials, which include a session token that limits the viability of the credentials to a shorter time-frame within which a key can be stolen and used.

Using Temporary Credentials for Amazon S3

The AWS Security Token Service (STS) issues temporary credentials to access AWS services such as Amazon S3. These temporary credentials include an access key, a secret key, and a session token that expires within a configurable amount of time. Temporary credentials are not currently handled transparently by CDH, so administrators must obtain them directly from Amazon STS. For details, see Temporary Security Credentials in the AWS Identity and Access Management Guide.

> **Important:** Cloudera recommends using this approach only for **single-user clusters** on **secure networks**—networks that allow access only to authorized users, all of whom are also authorized to use the S3 credentials.

To connect to Amazon S3 using temporary credentials obtained from STS, submit them as command-line arguments with the Hadoop job. For example:

```
-Dfs.s3a.access.key=your_temp_access_key
-Dfs.s3a.secret.key=your_temp_secret_key
-Dfs.s3a.session.token=your_session_token_from_AmazonSTS
-Dfs.s3a.aws.credentials.provider=org.apache.hadoop.fs.s3a.TemporaryAWSCredentialsProvider
```

Creating a Table in a Bucket

To create a table in a bucket, a user must have Sentry permissions on the S3 database and bucket URI. Cloudera recommends that you create a database specifically for the purpose of creating tables. Then, grant the user's role ALL permissions on the database and the URI.

It is possible to give the user ALL permissions on the server to allow the user to create external tables, but that approach is not recommended because it is not secure.

To allow a user to create tables in a bucket, complete the following steps:

1. Create a database where you want to create the tables.
2. Grant ALL on the database to the user's role.
3. Grant ALL on the bucket URI to the user's role.
4. The user must create the table with a reference to the database. For example:

```
CREATE EXTERNAL TABLE
sentrys3demo.your_external_s3_table
(name STRING, type STRING) ROW FORMAT
DELIMITED FIELDS TERMINATED BY ','
LOCATION 's3a://your-bucket-name/';
```

How to Configure Encryption for Amazon S3

Amazon offers several server-side encryption mechanisms for use with Amazon S3 storage. Cloudera clusters support server-side encryption for Amazon S3 data using either SSE-S3 (CDH 5.10 and later) or SSE-KMS (CDH 5.11 and later). With SSE-S3, keys are completely under the control of Amazon. With SSE-KMS, you have more control over the encryption keys, and can upload your own key material to use for encrypting Amazon S3. With either mechanism, encryption is applied transparently to the Amazon S3 bucket objects after you configure your cluster to use it.

Amazon S3 also supports TLS/SSL ('wire' or data-in-transit) encryption by default. Configuring data-at-rest encryption for Amazon S3 for use with your cluster involves some configuration both on Amazon S3 and for the cluster, using the Cloudera Manager Admin Console, as detailed below:

Requirements

Using Amazon S3 assumes that you have an Amazon Web Services account and the appropriate privileges on the AWS Management Console to set up and configure Amazon S3 buckets.

In addition, to configure Amazon S3 storage for use with a Cloudera cluster, you must have privileges as the User Administrator or Full Administrator on the Cloudera Manager Admin Console. See How to Configure AWS Credentials for details.

Amazon S3 and TLS/SSL Encryption

Amazon S3 uses TLS/SSL by default. Cloudera clusters (release 5.9 and later) include in their default configuration file the boolean property, `fs.s3a.connection.ssl.enabled` set to `true`, which activates TLS/SSL. This means that if the cluster has been configured to use TLS/SSL, connections from the cluster to Amazon S3 automatically use TLS wire encryption for the communication. The value of the `fs.s3a.connection.ssl.enabled` property can be confirmed by running `hadoop org.apache.hadoop.conf.Configuration`.

> **Important:** Cloudera recommends that clusters always be configured for TLS/SSL. See How to Configure TLS Encryption for Cloudera Manager for details.

If the cluster is not configured to use TLS, the connection to Amazon S3 silently reverts to an unencrypted connection.

Amazon S3 and Data at Rest Encryption

Cloudera clusters can use either of these two Amazon S3 mechanisms for data-at-rest encryption:

- Server-side Encryption with AWS KMS-Managed Keys (SSE-KMS), which requires using Amazon Key Management Server (AWS KMS) in conjunction with your Amazon S3. You can have Amazon generate and manage the keys in AWS KMS for you, or you can provide your own key material, but you must configure AWS KMS and create a key before you can use it with your cluster. See Prerequisites for Using SSE-KMS on page 390 for details.

- Server-side Encryption with S3-Managed Encryption Keys (SSE-S3) , which is simplest to set up because it uses Amazon-provided and -managed keys and has no requirements beyond setting a single property. See Configuring the Cluster to Use Server-Side Encryption on Amazon S3 on page 391 for details.

Enabling the cluster to use Amazon S3 server-side encryption involves using the Cloudera Manager Admin Console to configure the Advanced Configuration Snippet (Safety Valve) as detailed in Configuring the Cluster to Use Server-Side Encryption on Amazon S3 on page 391, below.

The steps assume that your cluster has been set up and that you have set up AWS credentials.

> **Note:** Cloudera clusters can be configured one encryption mode and key at a time for all objects in a given Amazon S3 bucket. However, you can change encryption modes or keys. See Changing Encryption Modes or Keys on page 392 for details.

Prerequisites for Using SSE-KMS

To use SSE-KMS with your Amazon S3 bucket, you must log in to the AWS Management Console using the account you set up in step 1 of Getting Started with Amazon Web Services on page 367. For example, the account *lab-iam* has an IAM user named *etl-workload* that has been granted permissions on the Amazon S3 storage bucket to be configured using SSE-KMS.

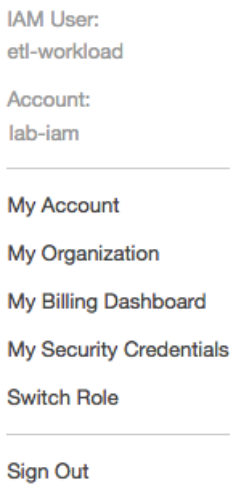

- Select **My Security Credentials** from the menu.
- Click **Encryption keys** (bottom left-hand on the AWS Management Console that displays at step 1, above).
- Click the **Create key** button to start the 5-step key-creation wizard that leads you through entry pages for giving the key an alias and description; adding tags, defining administrator permissions to the key, and defining usage permissions. The last page of the wizard shows you the policy that will be applied to the key before creating the key.

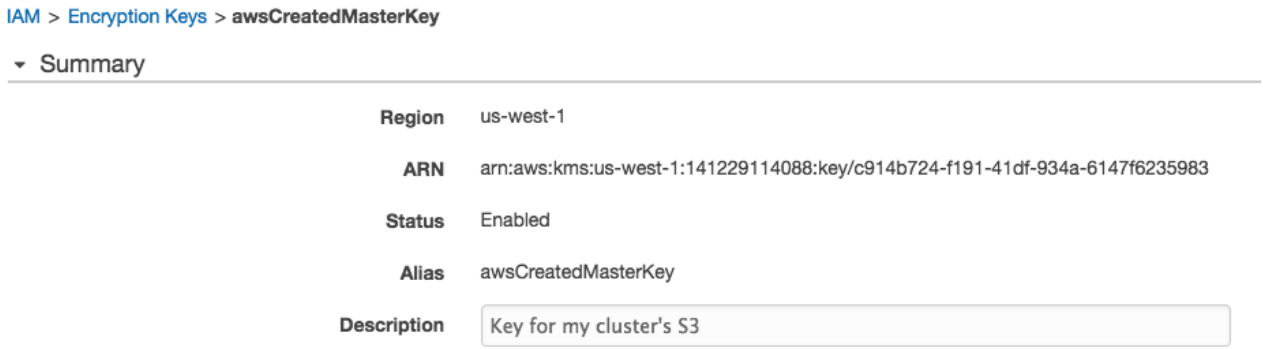

Configuring the Cluster to Use Server-Side Encryption on Amazon S3

Follow the steps below to enable server-side encryption on Amazon S3. To use SSE-KMS encryption, you will need your KMS key ID at step 7. Using SSE-S3 has no pre-requisites—Amazon generates and manages the keys transparently.

To configure the cluster to encrypt data stored on Amazon S3:

1. Log into the Cloudera Manager Admin Console.
2. Select **Clusters** > **HDFS**.
3. Click the **Configuration** tab.
4. Select **Scope** > **HDFS (Service-Wide)**.
5. Select **Category** > **Advanced**.
6. Locate the **Cluster-wide Advanced Configuration Snippet (Safety Valve) for core-site.xml** property.
7. In the text field, define the properties and values appropriate for the type of encryption you want to use.

   a. **To use SSE-S3 encryption:**

```
<property>
    <name>fs.s3a.server-side-encryption-algorithm</name>
    <value>AES256</value>
</property>
```

   b. **To use SSE-KMS encryption:**

```
<property>
    <name>fs.s3a.server-side-encryption-algorithm</name>
    <value>SSE-KMS</value>
</property>
<property>
    <name>fs.s3a.server-side-encryption-key</name>
    <value>your_kms_key_id</value>
</property>
```

8. Click **Save Changes**.
9. Restart the HDFS service.

For the value of *your_kms_key_id* (step 7b., above), you can use any of Amazon's four different key ID formats. Here are some examples:

Format	Example
Key ARN	arn:aws:kms:us-west-1:141229114088:key/c914b724-f191- 41df-934a-6147f6235983
Alias ARN	arn:aws:kms:us-west-1: 141229114088:key/c914b724-f191-41df-934a-6147f6235983: alias/awsCreatedMasterKey

Format	Example
Globally Unique Key ID	141229114088:key/c914b724-f191-41df-934a- 6147f6235983
Alias Name	alias/awsCreatedMasterKey

Changing Encryption Modes or Keys

Cloudera clusters can be configured to use only one type of server-side encryption for Amazon S3 data at a time.

However, Amazon S3 *does* support using different encryption keys for objects on an Amazon S3 bucket, which means that Amazon S3 continues to use whichever key was initially used to encrypt data (to decrypt on reads and re-encrypt on writes), even after a new mechanism and key exists. Your new key is used on new data written to the Amazon S3 bucket, while your 'old' key is used on any existing data that was encrypted with it. To summarize the behavior:

- Changing encryption mechanisms or keys on Amazon S3 has no effect on existing encrypted or unencrypted data.
- Data stored on Amazon S3 without encryption remains unencrypted even after you configure encryption for Amazon S3.
- Any existing encrypted data continues using the original mechanism and key to decrypt data (on reads) and re-encrypt data (on writes).
- After changing encryption mode or key, new objects stored on Amazon S3 from the cluster use the new mode and key.

Effect of Changing Encryption Mode or Key

This table shows the effect on existing encrypted or unencrypted data on Amazon S3 (far left column labeled "Data starts as...," reading down) and the result of "New" and "Existing" data and the keys that would be used after changing encryption-key configuration on the cluster. After changing encryption mode or key, existing data (Existing) and new data (New) use the mode and keys shown in columns 2 ("Unencrypted") through 5 ("Non-SSE Key"):

Data starts as...↓	Data results after modifying encryption mode or keys...			
	Unencrypted	SSE-S3	SSE-KMS	Non-SSE Key
Unencrypted	Existing	New	New	~
SSE-S3 encrypted	~	Existing	New	~
SSE-KMS [key1]	~	New	Existing [key1] New [key2]	~
Non-SSE key	~	~	~	Existing

Migrating Encrypted Data to New Encryption Mode or Key

To use a new key with existing data (data stored on Amazon S3 using a different mechanism or key), you must first decrypt the data, and then re-encrypt it using the new key. The process is as follows:

1. Create an Amazon S3 bucket as temporary storage for the unencrypted files.
2. Decrypt the data on the Amazon S3 bucket using the mechanism and key used for encryption (legacy encryption mode or key), moving the unencrypted data to the temporary bucket created in step 1.
3. Configure the Amazon S3 bucket to use the new encryption mechanism and key of your choice (SSE-S3, SSE-KMS).
4. Move the unencrypted data from the temporary bucket back to the Amazon S3 bucket that is now configured using the new mechanism and key.

Deleting an Encryption Key

If you change encryption modes or keys on Amazon S3, do not delete the key. To replace the old key and mode with a completely new mode or key, you must manually migrate the data.

When you delete an encryption key, Amazon puts the key in a **Pending Deletion** state (as shown in the Status column of the screenshot below) for at least 7 days, allowing you to reinstate a key if you change your mind or realize an error.

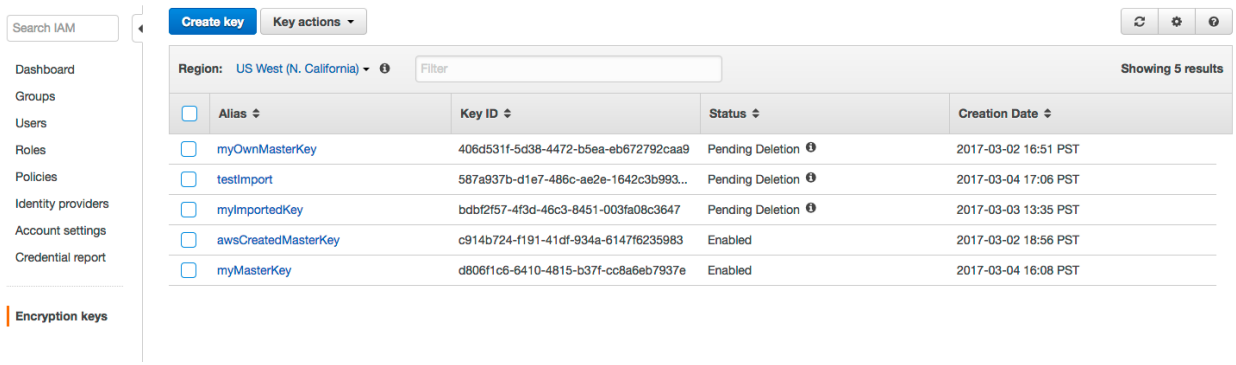

The pending time frame is configurable, from 7 up to 30 days. See AWS Key Management Service Documentation for complete details.

How to Configure AWS Credentials

Minimum Required Role: User Administrator (also provided by **Full Administrator**)

Amazon S3 (Simple Storage Service) can be used in a CDH cluster managed by Cloudera Manager in the following ways:

- As storage for Impala tables
- As a source or destination for HDFS and Hive/Impala replication and for cluster storage
- To enable Cloudera Navigator to extract metadata from Amazon S3 storage
- To browse S3 data using Hue

You can use the **S3Guard** feature to address possible issues with the "eventual consistency" guarantee provided by Amazon for data stored in S3. To use the S3Guard feature, you provision an Amazon DynamoDB that CDH uses as an additional metadata store to improve performance and guarantee that your queries return the most current data. See Configuring and Managing S3Guard.

To provide access to Amazon S3, you configure **AWS Credentials** that specify the authentication type (role-based, for example) and the access and secret keys. Amazon offers two types of authentication you can use with Amazon S3:

IAM Role-based Authentication

Amazon Identity and Access Management (IAM) can be used to create users, groups, and roles for use with Amazon Web Services, such as EC2 and Amazon S3. IAM role-based access provides the same level of access to all clients that use the role. All jobs on the cluster will have the same level of access to Amazon S3, so this is better suited for single-user clusters, or where all users of a cluster should have the same privileges to data in Amazon S3.

If you are setting up a peer to copy data to and from Amazon S3, using Cloudera Manager Hive or HDFS replication, select this option.

If you are configuring Amazon S3 access for a cluster deployed to Amazon Elastic Compute Cloud (EC2) instances using the IAM role for the EC2 instance profile, you do not need configure IAM role-based authentication for services such as Impala, Hive, or Spark.

> **Note:** IAM role-based authentication does not provide access to Amazon S3 using Cloudera Navigator.

Access Key Credentials

This type of authentication requires an AWS Access Key and an AWS Secret key that you obtain from Amazon and is better suited for environments where you have multiple users or multi-tenancy. You must enable the Sentry service and Kerberos when using the **S3 Connector** service. Enabling these services allows you to configure selective access for different data paths. (The Sentry service is not required for BDR replication or access by Cloudera Navigator.)

Cloudera Manager stores these values securely and does not store them in world-readable locations. The credentials are masked in the Cloudera Manager Admin console, encrypted in the configurations passed to processes managed by Cloudera Manager, and redacted from the logs.

For more information about Amazon S3, see the Amazon S3 documentation.

The client configuration files generated by Cloudera Manager based on configured services do not include AWS credentials. These clients must manage access to these credentials outside of Cloudera Manager. Cloudera Manager uses credentials stored in Cloudera Manager for trusted clients such as the Impala daemon and Hue. For access from YARN, MapReduce or Spark, see Using S3 Credentials with YARN, MapReduce, or Spark.

> **Note:** Isilon storage cannot be used with AWS credentials.

Adding AWS Credentials

Minimum Required Role: User Administrator (also provided by **Full Administrator**)

To add **AWS Credentials** for Amazon S3:

1. Open the Cloudera Manager Admin Console.
2. Click **Administration** > **External Accounts**.
3. Select the **AWS Credentials** tab.
4. Select one of the following:
   - **Add Access Key Credentials**

     This authentication mechanism requires you to obtain AWS credentials from Amazon.

     1. Enter a **Name** of your choosing for this account.
     2. Enter the **AWS Access Key ID**.
     3. Enter the **AWS Secret Key**.
   - **Add IAM Role-Based Authentication**

     1. Enter a name for your IAM Role-based authentication.

     > **Note:** You cannot use IAM Role-based authentication for Cloudera Navigator access.
5. Click **Add**.

   The **Edit S3Guard** dialog box displays.

   S3Guard enables a consistent view of data stored in Amazon S3 and requires that you provision a DynamoDB database from Amazon Web Services. S3Guard is optional but can help improve performance and accuracy for certain types of workflows. To configure S3Guard, see Configuring and Managing S3Guard and return to these steps after completing the configuration.

   If you do not want to enable S3Guard, click **Save** to finish adding the AWS Credential.

   The **Connect to Amazon Web Services** dialog box displays.
6. Choose one of the following options:
   - **Cloud Backup and Restore**

     To configure Amazon S3 as the source or destination of a replication schedule (to back up and restore data, for example), click the **Replication Schedules** link. See Data Replication for details.
   - **Cluster Access to S3**

To enable cluster access to S3 using the S3 Connector Service, click the **Enable for *Cluster Name*** link, which launches a wizard for adding the S3 Connector service. See Adding the S3 Connector Service for details.

- **Cloudera Navigator Access to S3**

  To give Cloudera Navigator access to Amazon S3, click the **Enable for Cloudera Navigator** link. Restart the Cloudera Navigator Metadata Server to enable access.

  > **Note:** You cannot enable **Cloudera Navigator Access to S3** when using **IAM Role-based Authentication**.

Managing AWS Credentials

To remove AWS credentials:

> **Note:** You cannot remove AWS credentials if they are in use by a service in the cluster. You might need to edit the connectivity of the credential before removing it.

1. Open the Cloudera Manager Admin Console.
2. Click **Administration** > **External Accounts**.
3. Select the **AWS Credentials** tab.
4. Locate the row with the credentials you want to delete and click **Actions** > **Remove**.

To edit AWS Access Key credentials:

1. Open the Cloudera Manager Admin Console.
2. Click **Administration** > **External Accounts**.
3. Select the **AWS Credentials** tab.
4. Locate the row with the Access Key Credentials you want to delete and click **Actions** > **Edit Credential**.

   The **Edit Credential** dialog box displays.

5. Edit the account fields.
6. Click **Save**.
7. Restart cluster services that use these credentials. If connectivity is for Cloudera Navigator, restart the Cloudera Navigator Metadata server.

To rename the IAM Role-Based Authentication:

1. Open the Cloudera Manager Admin Console.
2. Click **Administration** > **External Accounts**.
3. Select the **AWS Credentials** tab.
4. Locate the row with the IAM Role-Based Authentication you want to rename and click **Actions** > **Rename**.
5. Enter a new name.
6. Click **Save**.

   The **Connect to Amazon Web Services** screen displays.

7. Click the links to change any service connections or click **Close** to leave them unchanged.

To edit the services connected to an AWS Credentials account:

1. Open the Cloudera Manager Admin Console.
2. Click **Administration** > **External Accounts**.
3. Select the **AWS Credentials** tab.
4. Locate the row with the credentials you want to edit and click **Actions** > **Edit Connectivity**.

   The **Connect to Amazon Web Services** screen displays.

5. Click one of the following options:
   - **Cloud Backup and Restore**

     To configure Amazon S3 as the source or destination of a replication schedule (to back up and restore data, for example), click the **Replication Schedules** link. See Data Replication for details.
   - **Cluster Access to S3**

     To enable cluster access to S3 using the S3 Connector Service, click the **Enable for *Cluster Name*** link, which launches a wizard for adding the S3 Connector service. See Adding the S3 Connector Service for details.
   - **Cloudera Navigator Access to S3**

     To give Cloudera Navigator access to Amazon S3, click the **Enable for Cloudera Navigator** link. Restart the Cloudera Navigator Metadata Server to enable access.

     > **Note:** You cannot enable **Cloudera Navigator Access to S3** when using **IAM Role-based Authentication**.

How to Enable Sensitive Data Redaction

Redaction is a process that obscures data. It helps organizations comply with government and industry regulations, such as PCI (Payment Card Industry) and HIPAA, by making personally identifiable information (PII) unreadable except to those whose jobs require such access. For example, in simple terms, HIPAA legislation requires that patient PII is available only to appropriate medical professionals (and the patient), and that any medical or personal information exposed outside the appropriate context cannot be used to associate an individual's identity with any medical information. Data redaction can help ensure this privacy, by transforming PII to meaningless patterns—for example, transforming U.S. social security numbers to `XXX-XX-XXXX` strings.

Data redaction works separately from Cloudera data encryption techniques. Data encryption alone does not preclude administrators with full access to the cluster from viewing sensitive user data. Redaction ensures that cluster administrators, data analysts, and others cannot see PII or other sensitive data that is not within their job domain. At the same time, it does not prevent users with appropriate permissions from accessing data to which they have privileges.

Cloudera clusters implement some redaction features by default, while some features are configurable and require administrators to specifically enable them. The details are covered below:

Cloudera Manager and Passwords

As of Cloudera Manager 5.5 (and later releases) passwords are no longer stored in cleartext, neither through the Cloudera Manager Admin Console nor in the configuration files on disk. Passwords managed by Cloudera Manager and Cloudera Navigator are redacted internally, with the following results:

- In the Cloudera Manager Admin Console:
  - In the **Processes** page for a given role instance, passwords in the linked configuration files are replaced by `*******`.
  - **Advanced Configuration Snippet (Safety Valve)** parameters, such as passwords and secret keys, are visible to users (such as admins) who have edit permissions on the parameter, while those with read-only access see redacted data. However, the parameter name is visible to anyone. (Data to be redacted from these snippets is identified by a fixed list of key words: *password*, *key*, *aws*, and *secret*.)
- On all Cloudera Manager Server and Cloudera Manager Agent hosts:
  - Passwords in the configuration files in `/var/run/cloudera-scm-agent/process` are replaced by `********`.

Cloudera Manager Server Database Password Handling

Unlike the other passwords that are redacted or encrypted by Cloudera Manager, the password used for the Cloudera Manager Server database is stored in plaintext in the configuration file, `/etc/cloudera-scm-server/db.properties`, as shown in this example:

```
# Auto-generated by scm_prepare_database.sh on Mon Jan 30 05:02:18 PST 2017
#
# For information describing how to configure the Cloudera Manager Server
# to connect to databases, see the "Cloudera Manager Installation Guide."
#
com.cloudera.cmf.db.type=mysql
com.cloudera.cmf.db.host=localhost
com.cloudera.cmf.db.name=cm
com.cloudera.cmf.db.user=cm
com.cloudera.cmf.db.setupType=EXTERNAL
com.cloudera.cmf.db.password=password
```

However, as of Cloudera Manager 5.10 (and higher), rather than using a cleartext password you can use a script or other executable that uses `stdout` to return a password for use by the system.

During installation of the database, you can pass the script name to the `scm_prepare_database.sh` script with the `--scm-password-script` parameter. See Setting up the Cloudera Manager Server Database and Syntax for scm_prepare_database.sh for details.

You can also replace an existing cleartext password in `/etc/cloudera-scm-server/db.properties` by replacing the `com.cloudera.cmf.db.password` setting with `com.cloudera.cmf.db.password_script` and setting the name of the script or executable:

Cleartext Password (5.9 and prior)	Script (5.10 and higher)
com.cloudera.cmf.db.password=*password*	com.cloudera.cmf.db.password_script=*script_name_here*

At runtime, if `/etc/cloudera-scm-server/db.properties` does not include the script identified by `com.cloudera.cmf.db.password_script`, the system looks for the value of `com.cloudera.cmf.db.password`.

Cloudera Manager API Redaction

Cloudera Manager API does not have redaction enabled by default. You can configure redaction of the sensitive items by specifying a JVM parameter for Cloudera Manager. When you set this parameter, API calls to Cloudera Manager for configuration data do not include the sensitive information. For more information, see Redacting Sensitive Information from the Exported Configuration.

Log and Query Redaction

Cloudera Manager provides a configurable log and query redaction feature that lets you redact sensitive data in the CDH cluster as it's being written to the log files (see the Cloudera Engineering Blog "Sensitive Data Redaction" post for a technical overview), to prevent leakage of sensitive data. Redaction works only on data, not metadata—that is, sensitive data inside files is redacted, but the name, owner, and other metadata about the file is not.

Redaction is enabled for the entire cluster through the Cloudera Manager Admin Console, which also lets you define rules to target sensitive data in SQL data and log files. After enabling data redaction, the following contain replacement strings (such as a series of Xs) for the sensitive data items you define in your rules:

- Logs in HDFS and any dependent cluster services.
- Audit data sent to Cloudera Navigator.
- SQL query strings displayed by Hue, Hive, and Impala.

Note: Log redaction is not available in Isilon-based clusters.

See Enabling Log and Query Redaction Using Cloudera Manager on page 398 (below) for information about how to enable and define rules for sensitive data redaction for your cluster's logs and SQL queries (Hive, Hue, Impala).

How Redaction Rules Work

Cloudera's redaction process (redactor) uses regular expressions to target data for redaction. Common regular expression patterns for sensitive data include social security numbers, credit card numbers, email addresses, and dates, for example. The redaction rules are specified using the following elements:

- **Search -** Regular expression to compare against the data. For example, the regular expression `\d{4}[^\w]\d{4}[^\w]\d{4}[^\w]\d{4}` searches for a credit card number pattern. Segments of data that match the regular expression are redacted using the Replace string.
- **Replace -** String used to redact (obfuscate) data, such as a pattern of Xs to replace digits of a credit card number: `XXXX-XXXX-XXXX-XXXX`.
- **Trigger -** Optional simple string to be searched before applying the regular expression. If the string is found, the redactor searches for matches using the Search regular expression. Using the Trigger field improves performance: simple string matching is faster than regular expression matching.

You can use the following preconfigured redaction rules on your cluster. Rules are applied in the order listed in the table.

Rule	Regex Pattern	Replacement
Credit Card numbers (with separator)	`\d{4}[^\w]\d{4}[^\w]\d{4}[^\w]\d{4}`	`XXXX-XXXX-XXXX-XXXX`
Social Security numbers (with separator)	`\d{3}[^\w]\d{2}[^\w]\d{4}`	`XXX-XX-XXXX`
Email addresses	`\b([A-Za-z0-9]|[A-Za-z0-9][A-Za-z0-9\-\._] \ *[A-Za-z0-9])@(([A-Za-z0-9]|[A-Za-z] \ [A-Za-z0-9\-]*[A-Za-z0-9])\.)+([A-Za-z0-9] \ |[A-Za-z0-9][A-Za-z0-9\-]*[A-Za-z0-9])\b`	`email@redacted.host`
Hostnames	`\b(([A-Za-z]|[A-Za-z][A-Za-z0-9\-] \ *[A-Za-z0-9])\.)+([A-Za-z0-9] \ |[A-Za-z0-9][A-Za-z0-9\-]*[A-Za-z0-9])\b`	`HOSTNAME.REDACTED`

Enabling Log and Query Redaction Using Cloudera Manager

To enable redaction, you must use Cloudera Manager's new layout rather than the classic layout (the Switch to the new layout, Switch to the classic layout links toggle between these two variations of the UI). To enable log and query redaction in Cloudera Manager:

1. Login to the Cloudera Manager Admin Console.
2. Select **Clusters** > **HDFS**.
3. Click the **Configuration** tab.
4. In the **Search** box, type `redaction` to find the redaction property settings:
   - Enable Log and Query Redaction
   - Log and Query Redaction Policy List of rules for redacting sensitive information from log files and query strings. Choose a preconfigured rule or add a custom rule. See How Redaction Rules Work on page 398 for more information about rule pattern definitions.

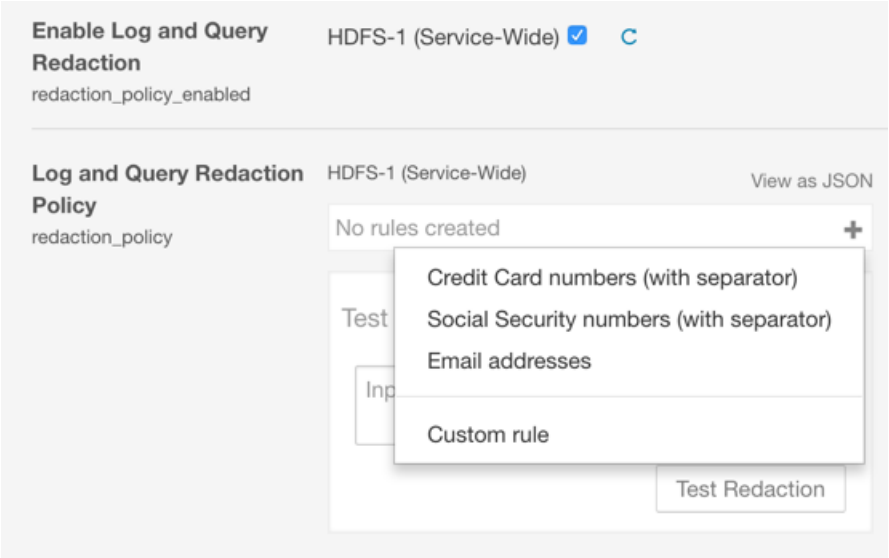

Test your rules:

- Enter sample text into the Test Redaction Rules text box
- Click Test Redaction.

5. Click **Save Changes**.
6. Restart the cluster.

If no rules match, the text you entered displays in the Results field, unchanged.

Using Cloudera Navigator Data Management for Data Redaction

> **Important:** This approach has been supplanted by Cloudera Manager's cluster-wide log and query redaction feature, and is not recommended.

You can specify credit card number patterns and other PII to be masked in audit events, in the properties of entities displayed in lineage diagrams, and in information retrieved from the Audit Server database and the Metadata Server persistent storage. Redacting data other than credit card numbers is not supported by default with the Cloudera Navigator property. You can use regular expressions to redact social security numbers or other PII. Masking applies only to audit events and lineage entities generated after enabling a mask.

Required Role: Navigator Administrator or Full Administrator

1. Log into **Cloudera Manager Admin Console**.
2. Select **Clusters** > **Cloudera Management Service**.
3. Click the **Configuration** tab.
4. Expand the **Navigator Audit Server Default Group** category.
5. Click the **Advanced** category.
6. Configure the **PII Masking Regular Expression** property with a regular expression that matches the credit card number formats to be masked. The default expression is:

```
(4[0-9]{12}(?:[0-9]{3})?)|(5[1-5][0-9]{14})|
(3[47][0-9]{13})|(3(?:0[0-5]|[68][0-9])[0-9]{11})|
(6(?:011|5[0-9]{2})[0-9]{12})|((?:2131|1800|35\\d{3})\\d{11})
```

which consolidates these regular expressions:

- Visa - `(4[0-9]{12}(?:[0-9]{3})?)`
- MasterCard - `(5[1-5][0-9]{14})`
- American Express - `(3[47][0-9]{13})`

- Diners Club - `(3(?:0[0-5]|[68][0-9])[0-9]{11})`
- Discover - `(6(?:011|5[0-9]{2})[0-9]{12})`
- JCB - `((?:2131|1800|35\\d{3})\\d{11})`

If the property is left blank, PII information is not masked.

7. Click **Save Changes**.

How to Log a Security Support Case

Before logging a support case, collect as much information about your issue as possible, as detailed below:

- If possible, gather a diagnostic bundle.
- Note specific details about the issue, including these:
  - Is this a new install, or did you upgrade a working cluster? For upgrades, identify release numbers ("trying to upgrade from Cloudera Manager 5.5.6 to Cloudera Manager 5.8.1," for example).
  - Is this a production deployment, development environment, or sandbox environment?
  - Note the severity of the impact and whether it is causing a production outage.
  - Capture error messages (screenshots, command-line capture, and so on) and attach to the case.
  - Note the commands executed and the results, captured to a file if possible.
  - Itemize all changes made to your cluster configuration after the issue arose (possibly, as attempts to resolve the issue).

The following are some additional details to gather when contacting support:

- Kerberos Issues on page 400
- LDAP Issues on page 400
- TLS/SSL Issues on page 400

Kerberos Issues

- For Kerberos issues, your `krb5.conf` and `kdc.conf` files are valuable for support to be able to understand your configuration.
- If you are having trouble with client access to the cluster, provide the output for `klist -ef` after kiniting as the user account on the client host in question. Additionally, confirm that your ticket is renewable by running `kinit -R` after successfully kiniting.
- Specify if you are authenticating (kiniting) with a user outside of the Hadoop cluster's realm (such as Active Directory, or another MIT Kerberos realm).
- If using AES-256 encryption, ensure you have the Unlimited Strength JCE Policy Files deployed on all cluster and client nodes.

LDAP Issues

- Specify the LDAP service in use (Active Directory, OpenLDAP, one of Oracle Directory Server offerings, OpenDJ, etc)
- Provide a screenshot of the LDAP configuration screen you are working with if you are troubleshooting setup issues.
- Be prepared to troubleshoot using the `ldapsearch` command (requires the `openldap-clients` package) on the host where LDAP authentication or authorization issues are being seen.

TLS/SSL Issues

- Specify whether you are using a private/commercial CA for your certificates, or if they are self-signed. Note that Cloudera strongly recommends against using self-signed certificates in production clusters.
- Clarify what services you are attempting to setup TLS/SSL for in your description.

- When troubleshooting TLS/SSL trust issues, provide the output of the following `openssl` command:

```
openssl s_client -connect host.fqdn.name:port
```

How To Obtain and Deploy Keys and Certificates for TLS/SSL

Cloudera recommends obtaining certificates from one of the trusted public certificate authorities (CA) such as Symantec or Comodo for TLS/SSL encryption for the cluster. This How To provides a brief overview of the tools used to create the various artifacts followed by the steps in the process. Plan ahead to allow time to receive signed certificates from whichever CA you use.

> **Note:** Always check with your selected public CA before creating any CSRs and follow the CA-specific processes.

If your organization uses its own internal CA, follow your internal process for creating and submitting the CSRs.

Tools Overview

Java Keytool and OpenSSL are key management tools that let you create the security artifacts needed for TLS/SSL. See How to Convert PEM to JKS and JKS to PEM for TLS/SSL Services and Clients for more information beyond the two short overviews below.

Java Keytool

Oracle Java `keytool` is a utility included with the Oracle JDK for creating and managing cryptographic keys and certificates. During configuring the Cloudera Manager Cluster for TLS/SSL, you create the private key pairs, keystore, certificate signing requests, and create a truststore for specific use by the cluster using this software tool, as detailed in the steps throughout this guide.

Java Keytool Requirements for Cloudera Manager TLS/SSL Configuration

For any steps using the Java Keytool, be sure to:

- Use the Oracle Java `keytool` rather than tools such as OpenJDK.
- Use the JDK downloaded from Oracle or the Cloudera-provided Oracle JDK located in this default path on a Cloudera Manager Server host:

```
/usr/java/jdk1.7.0_67-cloudera/bin/jre/lib/security
```

- Use the same version of the Java keytool for all steps. If the host has multiple JDKs installed, set the `PATH` variable so that the Oracle JDK is invoked first, as in this example:

```
$ export JAVA_HOME=/usr/java/jdk1.7.0_67-cloudera
$ export PATH=$JAVA_HOME/bin:$PATH
```

- Use the same *password* for the `-keypass` and `-storepass` in any commands that invoke these two options. Cloudera Manager requires the same password for a key and its keystore.

OpenSSL

OpenSSL is an open source cryptography and TLS/SSL toolkit that has been widely used since its inception ~ 1999. Just as with Java Keytool, OpenSSL lets you create private keys, certificate requests, and keystores, and it provides options for verifying certificates.

Cloudera Manager Agent hosts act as clients of a Cloudera Manager Server host during RPC client and server communications. The Agent hosts, Hue, Impala and other Python-based services require PEM-formatted keys and certificates (PKCS #8), which is why the steps below include converting some of the JKS artifacts using this tool. See How to Convert PEM to JKS and JKS to PEM for TLS/SSL Services and Clients for more information.

Step 1: Create Directory for Security Artifacts

Distributing the certificates, keys, truststore—in short, all security artifacts used for TLS/SSL intra-cluster communications—is part of this and some subsequent processes. To keep things organized, Cloudera recommends that you create the directory and distribute and store artifacts when you receive them, even though they may not be immediately needed.

The following table shows an example directory structure for the security artifacts created in the steps in this and subsequent sections. Use different names if you like, but for ease of deployment, use the same directory names across all hosts in the cluster.

Example	Description
`cmhost.sec.example.com`	FQDN of an example Cloudera Manager Server host.
`/opt/cloudera/security`	Base path for security-related files.
`/opt/cloudera/security/pki`	Path for all security artifacts associated with TLS/SSL, including keys, keystores (*`keystore`*`.jks`), CSR, and root- and intermediate-CA certificates.
`/usr/java/jdk1.7.0_67-cloudera/jre/lib/security/jssecacerts`	Path to the default alternative Java truststore on a Cloudera Manager Server host system.

1. On the Cloudera Manager Server host, create the `/opt/cloudera/security/pki` directory:

```
$ sudo mkdir -p /opt/cloudera/security/pki
```

2. This directory must be owned by `cloudera-scm:cloudera-scm` and have its executable bit set correctly so that Cloudera Manager can access the keystore at runtime:

```
$ sudo chown -R cloudera-scm:cloudera-scm /opt/cloudera/security/jks
$ umask 022
$ cd /opt/cloudera/security/jks
```

Directories and artifacts persist during system upgrades. For security purposes, for any host you remove from a cluster, remove any directories you create and more importantly, remove any security artifacts (keys, certificates, and so on) they may contain.

Step 2: Create the Java Truststore

On the Cloudera Manager Server host, copy the JDK `cacerts` file to `jssecacerts`:

```
$ sudo cp $JAVA_HOME/jre/lib/security/cacerts $JAVA_HOME/jre/lib/security/jssecacerts
```

If you do not have the `$JAVA_HOME` variable set, replace it with the path to the Oracle JDK. For example, the default path to the Java JDK on a Cloudera Manager Server host is:

`/usr/java/jdk1.7.0_67-cloudera/`

Step 3: Generate Server Key and CSR

1. On the Cloudera Manager Server host, use the `keytool` utility to generate a keypair and a keystore named for the server. Replace the OU, O, L, ST, and C entries with the values for your organization name, location, and country code (US, in the example):

```
$ sudo keytool -genkeypair -alias $(hostname -f)-server -keyalg RSA -keystore \
/opt/cloudera/security/pki/$(hostname -f)-server.jks -keysize 2048 -dname \
"CN=$(hostname -f),OU=Dept,O=Example.com,L=City,ST=State,C=US" \
-storepass password -keypass password
```

Use the same *password* for the key and its keystore (`-keypass` and `-storepass`, respectively): Cloudera Manager does not support using different passwords for the key and keystore.

2. Generate a certificate signing request (CSR) for the public key (contained in the keystore as a result of the command above). In this command below, enter the password that you set for the key and the keystore in the command above:

```
$ sudo keytool -certreq -alias $(hostname -f)-server \
-keystore /opt/cloudera/security/pki/$(hostname -f)-server.jks \
-file /opt/cloudera/security/pki/$(hostname -f)-server.csr -storepass password \
-keypass password
```

Step 4: Submit the CSR to the CA

1. Submit the CSR file to your certificate authority using the process and means required by the CA, for example, email or web submission. For the certificate format, specify PEM (base64 ASCII) (see Step 5 below for an example of PEM formatted certificate heading and closing structure).
2. The public CA will request specific details from you, to verify that you own the domain name contained in the CSR, before they issue the certificate.
3. When you receive the signed certificate from the CA, you can proceed with Step 5.

Step 5: Verify the Certificate

From the public (or internal) CA, you receive the certificate for the server signed by the CA and several other digital artifacts, including root CA and possibly one (or more) intermediate CA certificates. Before distributing these to the hosts, make sure the certificate is in PEM format. A PEM formatted certificate file looks like this:

```
-----BEGIN CERTIFICATE-----
<The encoded certificate is represented by multiple lines of exactly 64 characters,
except
for the last line, which can contain 64 characters or fewer.>
-----END CERTIFICATE-----
```

If your issued certificate is in binary (DER) format, convert it to PEM format before continuing. See How to Convert Formats (PEM, JKS) for TLS/SSL Clients and Services for details.

To modify the truststore (jssecacerts) to explicitly trust certificates or certificate chain (as you might for certificates signed by an internal CA), follow the steps in How to Add Root and Intermediate CAs to Truststore for TLS/SSL.

Step 6: Import the Certificate into the Keystore

Copy the signed certificate you receive from the CA to the Cloudera Manager Server host. To identity the certificate's functionality, include a suffix such as "-server.cert.pem" in the target name for the copy command, as shown in this example:

```
cp cert-file-recd /opt/cloudera/security/pki/$(hostname -f)-server.cert.pem
```

Import the certificate into the keystore.

```
$ sudo keytool -importcert -alias $(hostname -f)-server \
-file /opt/cloudera/security/pki/$(hostname -f)-server.cert.pem \
-keystore /opt/cloudera/security/pki/$(hostname -f)-server.jks
```

Assuming the certificate was obtained from a public CA, you can safely disregard this message about trust, and enter `yes` to continue:

```
... is not trusted. Install reply anyway? [no]:  yes
```

You *must* see the following response confirming that the certificate has been properly imported and can verify the private key that it certifies.

```
Certificate reply was installed in keystore
```

If you do not see this response, double-check all your steps up to this point: are you working in the correct path? Do you have the proper certificate? and so on. See Getting Support for information about how to contact Cloudera Support and to find out about other sources of help if you cannot successfully import the certificates.

How To Renew and Redistribute Certificates

Certificates installed throughout TLS/SSL-configured clusters have a defined lifetime. That lifetime begins on a certain date and ends on another date, typically a year or two later as specified in the certificate's **validity** attribute. For example, the following certificate is valid between January 24, 2017 and February 23, 2018 only:

```
$ openssl x509 -in cacert.pem -text -noout
Certificate:
    Data:
        Version: 3 (0x2)
        Serial Number: 11485830970703032316 (0x9f65de69ceef2ffc)
    Signature Algorithm: sha256WithRSAEncryption
        Issuer: C=US, ST=MD, L=Baltimore, CN=Test CA/emailAddress=test@example.com
        Validity
            Not Before: Jan 24 14:24:11 2017 GMT
            Not After : Feb 23 14:24:11 2018 GMT
        Subject: C=US, ST=MD, L=Baltimore, CN=Test CA/emailAddress=test@example.com
        Subject Public Key Info:
            Public Key Algorithm: rsaEncryption
                Public-Key: (4096 bit)
                Modulus:
                    00:b1:7f:29:be:78:02:b8:56:54:2d:2c:ec:ff:6d:
                    ...
                    39:f9:1e:52:cb:8e:bf:8b:9e:a6:93:e1:22:09:8b:
```

If any certificates used for TLS/SSL in the Cloudera cluster expire, the cluster can fail completely or can exhibit various issues related to connectivity between components. For example, an expired server certificate on the Cloudera Manager Server host will be rejected by the Cloudera Manager Agent host and prevent the cluster node from launching.

> **Warning:** Replace certificates **before** they expire to avoid issues. If certificates expire and the cluster or any nodes fail to start as a result, see Replacing Certificates After a Certificate Expires on page 407 for steps to resolve.

Obtaining certificates with new expiration dates takes just as much time as it did to obtain them in the first place. This guide steps you through the process.

> **Note:** Keep track of the expiration dates for all certificates installed in Cloudera clusters. For example, add certificate expiration dates to your calendar with notifications set to alter you at least 1 month in advance of the expiration date.

Assumptions and Requirements

The steps below assume that the cluster has been configured for TLS/SSL using certificates obtained previously as detailed in How to Configure TLS Encryption for Cloudera Manager, and that the cluster has been operational using those certificates. This table recaps recommended paths to various security artifacts:

Example	Description
`/opt/cloudera/security`	Base path for security-related files.

Example	Description
`/opt/cloudera/security/pki`	Path for all security artifacts associated with TLS/SSL, including keys, keystores (*`keystore`*`.jks`), CSR, and root- and intermediate-CA certificates.
`/usr/java/jdk1.7.0_67-cloudera/jre/lib/security/jssecacerts`	Path to the default alternative Java truststore on a Cloudera Manager Server host system.

This guide assumes that the Cloudera Manager Server host uses the **`jssecacerts`** truststore and includes all CA certs from **`cacerts`** and any intermediate CA certificates needed to enable successful chain of trust traversal during handshake.

```
$ sudo cp $JAVA_HOME/jre/lib/security/cacerts $JAVA_HOME/jre/lib/security/jssecacerts
```

Use Cloudera Manager Admin Console to check TLS/SSL configuration details for the cluster and services, and to verify paths to keys and keystores, certificates, and trust stores configured for each service. These do not need to be re-enabled or changed (unless you replace existing keys with new ones as part of this process), but you can note all paths and names of all TLS-related security artifacts before you begin.

Certificates and keys may have been converted from one format to another (as detailed in How to Convert File Encodings (DER, JKS, PEM) for TLS/SSL Certificates and Keys). That means that a CSR may have been used to obtain a JKS formatted certificate for one service that was then converted to PEM for use by another service (or services) running on the same node of the cluster as needed.

> **Note:** Use the same names for certificates to avoid having to reconfigure the TLS/SSL settings in Cloudera Manager Admin Console.

Check Certificate Expiration Dates

If you do not know the expiration dates for certificates installed on the cluster, use OpenSSL (for PEM-formatted certificates) or use Java Keytool (for JKS-formatted certificates) to determine certificate expiration dates.

PEM-formatted certificates (PKCS #8) are used by Cloudera Manager Agent hosts, Hue, Impala and other Python-based services, while JKS-formatted certificates are used by HDFS, MapReduce, and YARN, for example. See Certificate Formats (JKS, PEM) and Cluster Components on page 194 for more information.

Using OpenSSL to Obtain Certificate Details

To check the expiration dates of PEM-formatted certificates using `openssl` and the certificate by name:

```
openssl x509 -enddate -noout -in /opt/cloudera/security/pki/$(hostname -f)-server.cert.pem
```

To check expiration dates by querying the active listener ports for any TLS-enabled services from the command line, use OpenSSL as in this example of querying the Cloudera Manager TLS listener port (7183):

```
echo | openssl s_client -connect fqdn.example.com:7183 2>/dev/null | openssl x509 -noout
 -subject -dates
```

The subject and dates should display, as shown in this example:

```
subject=/C=US/ST=California/L=Los Angeles/O=Internet Corporation for Assigned Names
and Numbers/OU=Technology/CN=www.example.org
notBefore=Nov 3 00:00:00 2015 GMT
notAfter=Nov 28 12:00:00 2018 GMT
```

Using Java Keytool to Obtain Certificate Details

To obtain all information about a certificate:

```
keytool -list -v -keystore keystore_name.jks
```

Renewing and Replacing Certificates Before Expiration

There are many different ways to proceed depending on your specific needs. Renewing a certificate really means obtaining a new certificate, one with a new start and end date, and updating the key in the keystore with the certificate so it can being presented during the TLS handshake. That means you can replace the key at the same time, or re-use the existing key but apply a new certificate to it. If you have the original Certificate Signing Request (CSR), you can resubmit that to the commercial Certificate Authority (CA) as was done when TLS/SSL was configured for the cluster. Here is a summary of the various approaches:

- Use the existing key to generate a new CSR for a certificate to replace the one currently associated with the key, as detailed in Step 1 below. This requires knowing the key/keystore password.
- Create a new private key and public key to replace those already in use and generate a new CSR to obtain a completely new certificate to use with the public key. To use this approach, follow steps in How To Obtain and Deploy Keys and Certificates for TLS/SSL.
- Submit the original CSR to the CA and obtain a new certificate. If you have the CSR that matches the key used by the service, you can skip Step 1 and start with Step 2 below.

Step 1: Generate a New CSR from the Key

You must know the password for the key and keystore to generate a new signing request from it. Use the same alias used for the key when it was created.

1. Use `keytool` to generate the CSR as follows:

```
$ sudo keytool -certreq -keystore server-key.jks -file newcert.csr
Enter keystore password:
```

Step 2: Submit the CSR to the CA

1. Submit the CSR file to your certificate authority using the process and means required by the CA, for example, email or web submission. For the certificate format, specify PEM (base64 ASCII) (see Step 3 below for an example of PEM formatted certificate heading and closing structure).
2. The public CA will request specific details from you, to verify that you own the domain name contained in the CSR, before they issue the certificate.
3. When you receive the signed certificate from the CA, you can proceed with Step 3.

Step 3: Verify the Certificate

From any Certificate Authority (CA), you should receive a certificate signed by the CA attesting to the validity of the key. Before distributing these to the hosts, make sure the certificate is in PEM format. A PEM formatted certificate file looks like this:

```
-----BEGIN CERTIFICATE-----
<The encoded certificate is represented by multiple lines of exactly 64 characters,
except
for the last line, which can contain 64 characters or fewer.>
-----END CERTIFICATE-----
```

If your issued certificate is in binary (DER) format, convert it to PEM format before continuing. See How to Convert Formats (PEM, JKS) for TLS/SSL Clients and Services for details.

Along with the certificate, the CA also provides several other digital artifacts, including root CA and possibly one (or more) intermediate CA certificates. The intermediate certificates may need to be added to the trust store in the following situations:

For certificates obtained from a public (commercial) CA—Intermediate or root certificates do not need to be installed in the trust stores for the cluster. The JDK trust store already includes the certificates needed to establish chain of trust for certificates obtained from commercial CAs.

For certificates obtained from an internal CA—To modify the truststore (jssecacerts) to explicitly trust certificates or certificate chain (as you might for certificates signed by an internal CA), follow the steps in How to Add Root and Intermediate CAs to Truststore for TLS/SSL.

Step 4: Import the Certificate into the Keystore

Copy the signed certificate you receive from the CA to the Cloudera Manager Server host. To identity the certificate's functionality, include a suffix such as "-server.cert.pem" in the target name for the copy command, as shown in this example:

```
cp cert-file-recd /opt/cloudera/security/pki/$(hostname -f)-server.cert.pem
```

Import the certificate into the keystore.

```
$ sudo keytool -importcert -alias $(hostname -f)-server \
-file /opt/cloudera/security/pki/$(hostname -f)-server.cert.pem \
-keystore /opt/cloudera/security/pki/$(hostname -f)-server.jks
```

Assuming the certificate was obtained from a public CA, you can safely disregard this message about trust, and enter `yes` to continue:

```
... is not trusted. Install reply anyway? [no]:  yes
```

You *must* see the following response confirming that the certificate has been properly imported and can verify the private key that it certifies.

```
Certificate reply was installed in keystore
```

If you do not see this response, double-check all your steps up to this point: are you working in the correct path? Do you have the proper certificate? and so on. See Getting Support for information about how to contact Cloudera Support and to find out about other sources of help if you cannot successfully import the certificates.

Replacing Certificates After a Certificate Expires

If the cluster cannot start up because of an expired certificate, then perform the steps in Renewing and Replacing Certificates Before Expiration on page 406 to resolve the issue.

If the replacement certificate has not been obtained yet, you can use a self-signed certificate for the short term, until you receive the CA signed certificate.

How to Set Up a Gateway Node to Restrict Access to the Cluster

The steps below configure a firewall-protected Hadoop cluster that allows access only through the node configured as the gateway. Clients access the cluster through the gateway using the REST API, for example, using HttpFS (which provides REST access to HDFS) or using Oozie, which allows REST access for submitting and monitoring jobs.

Installing and Configuring the Firewall and Gateway

Follow these steps:

1. Choose a cluster node to be the gateway machine.
2. Install and configure the Oozie server by following the standard directions starting here: Installing Oozie.
3. Install HttpFS.

4. Start the Oozie server:

```
$ sudo service oozie start
```

5. Start the HttpFS server:

```
$ sudo service hadoop-httpfs start
```

6. Configure firewalls.

   Block all access from outside the cluster.

   - The gateway node should have ports 11000 (oozie) and 14000 (hadoop-httpfs) open.
   - Optionally, to maintain access to the Web UIs for the cluster's JobTrackers and NameNode, open their HTTP ports: see Ports Used by Components of CDH 5.

7. Optionally configure authentication in simple mode (default) or using Kerberos. See HttpFS Authentication on page 133 to configure Kerberos for HttpFS and Oozie Authentication on page 149 to configure Kerberos for Oozie.
8. Optionally encrypt communication using HTTPS for Oozie by following these directions.

Accessing HDFS

With the Hadoop client:

All of the standard `hadoop fs` commands work; just make sure to specify `-fs webhdfs://HOSTNAME:14000`. For example (where *GATEWAYHOST* is the hostname of the gateway machine):

```
$ hadoop fs -fs webhdfs://GATEWAYHOST:14000 -cat /user/me/myfile.txt
Hello World!
```

Without the Hadoop client:

You can run all of the standard `hadoop fs` commands by using the WebHDFS REST API and any program that can do `GET`, `PUT`, `POST`, and `DELETE` requests; for example:

```
$ curl "http://GATEWAYHOST:14000/webhdfs/v1/user/me/myfile.txt?op=OPEN&user.name=me"
Hello World!
```

> **Important:** The user.name parameter is valid only if security is disabled. In a secure cluster, you must a initiate a valid Kerberos session.

In general, the command will look like this:

```
$ curl "http://GATEWAYHOST/webhdfs/v1/PATH?[user.name=USER&]op=…"
```

You can find a full explanation of the commands in the WebHDFS REST API documentation.

Submitting and Monitoring Jobs

The Oozie REST API supports the direct submission of jobs for MapReduce, Pig, and Hive; Oozie automatically creates a workflow with a single action. For any other action types, or to execute anything more complicated than a single job, you must create an actual workflow. Required files (JAR files, input data) must already exist on HDFS; if they do not, you can use HttpFS to upload the files.

With the Oozie client:

All of the standard Oozie commands will work. You can find a full explanation of the commands in the documentation for the command-line utilities.

Without the Oozie client:

You can run all of the standard Oozie commands by using the REST API and any program that can do GET, PUT, and POST requests. You can find a full explanation of the commands in the Oozie Web Services API documentation.

How To Set Up Access to Cloudera EDH or Cloudera Director (Microsoft Azure Marketplace)

The Cloudera Enterprise Data Hub (EDH) and Cloudera Director bundles that are available on the Microsoft Azure Marketplace only support SSH access (port 22). To access Cloudera Manager (port 7180), Cloudera Director (port 7189), or other services, you can:

- Set up a SOCKS (Sockets Secure protocol) proxy on your client machine. Cloudera recommends that you use this option.
- Add inbound rules to the Network Security Group in the Azure instance after you deploy EDH or Director to Azure.

Configure the SOCKS Proxy

The SOCKS5 protocol is implemented as a client and server process that enables traversal of IP network firewalls. After you configure the SOCKS proxy, your browser resolves DNS lookups using the Microsoft Azure network (through the proxy server), and lets you connect to services using internal FQDNs or private IP address.

With this approach, you complete the following tasks:

- Set up a single SSH tunnel to one of the hosts on the network and create a SOCKS proxy on the host.
- Change the browser configuration to perform all lookups through the SOCKS proxy host.

Network Prerequisites

Verify the following prerequisites before you connect to your cluster with a SOCKS proxy:

- You must be able to reach the host that you want to proxy to from the public internet or from the network that you're connecting from.
- The host that you proxy to must be on the same network as the Cloudera services that you're connecting to. For example, if you’re using the Cloudera EDH offering, tunnel to the Cloudera Manager host. If you're using the Cloudera Director offering, tunnel to the Cloudera Director host.

Find the Public IP of the Host

Only one VM is created for the Cloudera Director. Use the public IP of that VM.

For the Cloudera EDH offering, use the public IP of the 0th master node VM: [dnsName]-mn0.

Start the SOCKS Proxy

On Linux

To start a SOCKS proxy over SSH, run the following command:

```
ssh -i your-key-file.pem -CND 1080
the_username_you_specified@publicIP_of_VM
```

The command uses the following parameters:

- `-i your-key-file.pem` Specifies the path to the private key needed to SSH to the Cloudera Director server. Omit if using SSH passwords.
- `C` Sets up compression.
- `N` Suppresses any command execution once established.
- `D` Sets up the SOCKS proxy on a port.
- `1080` The port to set the SOCKS proxy locally.

On Windows

Follow the instructions on Microsoft's website.

Configure Google Chrome to Use the Proxy

By default, Google Chrome uses system-wide proxy settings on a per-profile basis. To start Chrome without these settings, open Chrome through the command line and specify the following:

- The SOCKS proxy port. This must be the same port that you used when you started the proxy.
- The profile. This examples below create a new profile.

Use one of the following commands to create a profile and launch a new instance of Chrome that does not conflict with any currently running Chrome instances.

Linux

```
/usr/bin/google-chrome \
--user-data-dir="$HOME/chrome-with-proxy" \
--proxy-server="socks5://localhost:1080"
```

Mac OS X

```
"/Applications/Google Chrome.app/Contents/MacOS/Google Chrome" \
--user-data-dir="$HOME/chrome-with-proxy" \
--proxy-server="socks5://localhost:1080"
```

Microsoft Windows

```
"C:\Program Files (x86)\Google\Chrome\Application\chrome.exe" ^
--user-data-dir="%USERPROFILE%\chrome-with-proxy" ^
--proxy-server="socks5://localhost:1080"
```

In this Chrome session, you can use the private IP address or internal FQDN to connect to any host that is accessible by Cloudera Director. For example, if you proxy to the Cloudera Director server, you can connect to Cloudera Director as if it were local by entering `localhost:7189` in the Chrome URL bar.

Network Security Group

Warning: This method is not recommended for any purpose other than a Proof of Concept. If the data is not carefully locked down, it will be accessible to hackers and malicious entities.

On portal.azure.com, find the Network Security Group(s) and add inbound rules for the various services. You may have to create these rules for the services. Refer to Cloudera documentation for more information on ports used by Cloudera Manager, CDH components, managed services, and third-party components.

How to Use Self-Signed Certificates for TLS

Self-signed certificates should not be used for production deployments. Self-signed certificates are created and stored in the keystore specified during the key-generation process, and should be replaced by a signed certificate. Using self-signed certificates requires generating and distributing the certificates and establishing explicit trust for the certificate.

However, using self-signed certificates lets you easily obtain certificates for TLS/SSL configuration and may be appropriate for non-production or test setups. See Encrypting Data in Transit on page 191 and Configuring TLS Encryption for Cloudera Manager on page 195 for more information.

Replace paths, file names, aliases, and other examples in the commands below for your system.

1. Create the directory for the certificates:

```
$ mkdir -p /opt/cloudera/security/x509/ /opt/cloudera/security/jks/
```

Give Cloudera Manager access to the directory, set the correct permissions, and then change to the directory:

```
$ sudo chown -R cloudera-scm:cloudera-scm /opt/cloudera/security/jks
$ sudo umask 0700
$ cd /opt/cloudera/security/jks
```

2. Generate the key pair and self-signed certificate, storing everything in the keystore with the same password for keystore and storepass, as shown below. Use the FQDN of the current host for the CN to avoid raising a `java.io.IOException: HTTPS hostname wrong` exception. Replace values for OU, O, L, ST, and C with entries appropriate for your environment:

```
keytool -genkeypair -alias cmhost -keyalg RSA -keysize 2048 -dname "cn=cm01.example.com,
 ou=Department,
o=Company, l=City, st=State, c=US" -keypass password -keystore example.jks -storepass
password
```

3. Copy the default Java truststore (`cacerts`) to the alternate system truststore (`jssecacerts`):

```
$ sudo cp $JAVA_HOME/jre/lib/security/cacerts $JAVA_HOME/jre/lib/security/jssecacerts
```

4. Export the certificate from the keystore (`example.jks`).

```
$ keytool -export -alias cmhost -keystore example.jks -rfc -file selfsigned.cer
```

5. Copy the self-signed certificate (`selfsigned.cer`) to the `/opt/cloudera/security/x509/` directory.

```
$ cp selfsigned.cer /opt/cloudera/security/x509/cmhost.pem
```

6. Import the public key into the alternate system truststore (`jssecacerts`), so that any process that runs with Java on this machine will trust the key. The default password for the Java truststore is `changeit`. Do not use the password created for the keystore in step 2.

```
$ keytool -import -alias cmhost -file /opt/cloudera/security/jks/selfsigned.cer
-keystore $JAVA_HOME/jre/lib/security/jssecacerts -storepass changeit
```

> **Important:** Repeat this process on each host in the cluster.

7. Rename the keystore:

```
$ mv /opt/cloudera/security/jks/example.jks /opt/cloudera/security/jks/cmhost-keystore.jks
```

You can also delete the certificate because it was copied to the appropriate path in step 5.

```
$ rm /opt/cloudera/security/jks/selfsigned.cer
```

The self-signed certificate set up is complete.

Troubleshooting Security Issues

Security-related issues can manifest in various ways and can be associated with various sources, including Kerberos authentication, HDFS encryption, and TLS/SSL (data in transit encryption). This section includes error messages and troubleshooting for common issues and some architectural details about some of the underlying components.

For information about Sentry troubleshooting, see Troubleshooting Sentry.

Error Messages and Various Failures

These error messages may indicate an issue with Kerberos authentication or other issues.

Cluster cannot run jobs after Kerberos enabled

Symptom: Cluster previously configured without Kerberos authentication may fail to run jobs for certain users on certain TaskTrackers (MRv1) or NodeManagers (YARN) after enabling Kerberos for the cluster. Errors may display in the TaskTracker or NodeManager logs. The following example errors are from TaskTracker on MRv1:

```
10/11/03 01:29:55 INFO mapred.JobClient: Task Id : attempt_201011021321_0004_m_000011_0,
 Status : FAILED
Error initializing attempt_201011021321_0004_m_000011_0:
java.io.IOException: org.apache.hadoop.util.Shell$ExitCodeException:
at org.apache.hadoop.mapred.LinuxTaskController.runCommand(LinuxTaskController.java:212)

at
org.apache.hadoop.mapred.LinuxTaskController.initializeUser(LinuxTaskController.java:442)

at
org.apache.hadoop.mapreduce.server.tasktracker.Localizer.initializeUserDirs(Localizer.java:272)

at org.apache.hadoop.mapred.TaskTracker.localizeJob(TaskTracker.java:963)
at org.apache.hadoop.mapred.TaskTracker.startNewTask(TaskTracker.java:2209)
at org.apache.hadoop.mapred.TaskTracker$TaskLauncher.run(TaskTracker.java:2174)
Caused by: org.apache.hadoop.util.Shell$ExitCodeException:
at org.apache.hadoop.util.Shell.runCommand(Shell.java:250)
at org.apache.hadoop.util.Shell.run(Shell.java:177)
at org.apache.hadoop.util.Shell$ShellCommandExecutor.execute(Shell.java:370)
at org.apache.hadoop.mapred.LinuxTaskController.runCommand(LinuxTaskController.java:203)

... 5 more
```

Possible cause: The issue is caused by legacy content in directories for TaskTracker and NodeManager that may exist after configuring Kerberos authentication for a cluster that previously did not use Kerberos. The sequence of events leading to this issue is as follows:

1. Cluster that had not been configured for Kerberos authentication was used to run jobs, which created local user directory (or directories) on each TaskTracker or NodeManager host.
2. Cluster was then configured to use Kerberos authentication.
3. Users try running jobs on the newly secured cluster but local user directories on TaskTrackers or NodeManagers are owned by the wrong user or have overly-permissive permissions.

These directories should have been cleaned up at the time Kerberos authentication was enabled for the cluster.

Steps to resolve: Delete `mapred.local.dir` or `yarn.nodemanager.local-dirs` directories across the cluster for affected users.

NameNode fails to start

Symptom: Login failure occurs when attempting to start the NameNode. With debugging enabled on the cluster, the following KrbException messages may display:

```
Caused by: KrbException: Integrity check on decrypted field failed (31) - PREAUTH_FAILED}}
```

```
Caused by: KrbException: Identifier does not match expected value (906)
```

Possible cause: This issue may be due to incorrect configuration for AES-256. By default, certain operating systems—CentOS/Red Hat Enterprise Linux 5.6 (and higher), Ubuntu—use AES-256 encryption which requires installing the Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction Policy File on all hosts (see JCE Policy File for AES-256 Encryption for details), or disabling AES-256 support in the `kdc.conf` or `krb5.com` (see disable AES-256 encryption from the Kerberos instance for details).

Steps to resolve: KrbException 31 and KrbException 906 can be caused by various issues, but the most likely cause is incorrectly configured AES-256 encryption. Resolving the issue should start by determining the type of encryption configured for the cluster.

To verify the type of encryption configured for the cluster:

1. On the local KDC host, type this command to create a test principal:

```
$ kadmin -q "addprinc test"
```

2. On a cluster host, type this command to start a Kerberos session as test:

```
$ kinit test
```

3. On a cluster host, type this command to view the encryption type in use:

```
$ klist -e
```

If AES-256 is being used, output such as the following displays:

```
Ticket cache: FILE:/tmp/krb5cc_0
Default principal: test@SCM
Valid starting     Expires            Service principal
05/19/11 13:25:04  05/20/11 13:25:04  krbtgt/SCM@SCM
    Etype (skey, tkt): AES-256 CTS mode with 96-bit SHA-1 HMAC, AES-256 CTS mode with
96-bit SHA-1 HMAC
```

To remove AES-256 encryption from the Kerberos configuration files:

- Remove `aes256-cts:normal` from the `supported_enctypes` field of the `kdc.conf` or `krb5.conf` file.
- After changing the configuration, restart the KDC and the kadmin servers.
- Change TGT principal (krbtgt/REALM@REALM) and other principal passwords as needed.
- In the [realms] section of the `kdc.conf` file, for the realm associated with `HADOOP.LOCALDOMAIN`, add (or replace if it exists already) the following variable:

```
supported_enctypes = des3-hmac-sha1:normal arcfour-hmac:normal des-hmac-sha1:normal
des-cbc-md5:normal des-cbc-crc:normal des-cbc-crc:v4 des-cbc-crc:afs3
```

- Recreate the `hdfs` keytab file and `mapred` keytab file using the `-norandkey` option in the `xst` command (see Step 4: Create and Deploy the Kerberos Principals and Keytab Files on page 84 for details.

```
kadmin.local: xst -norandkey -k hdfs.keytab hdfs/fully.qualified.domain.name
HTTP/fully.qualified.domain.name
kadmin.local: xst -norandkey -k mapred.keytab mapred/fully.qualified.domain.name
HTTP/fully.qualified.domain.name
```

Clients cannot connect to NameNode

Symptom: The NameNode keytab file does not have an AES-256 entry but the client tickets do. The NameNode starts but clients cannot connect to it. The error message does not specify "AES256" but rather contains enctype code "18."

Possible cause: Issue related to AES-256 encryption and the JCE library.

Steps to resolve: Verify that Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction Policy File is installed or remove `aes256-cts:normal` from the `supported_enctypes` field of the `kdc.conf` or `krb5.conf` file, as detailed above.

Hadoop commands run in local realm but not in remote realm

Symptom: After enabling cross-realm trust, authenticating as a principal in the local realm lets you successfully run Hadoop commands, but authenticating as a principal in the remote realm does not.

Possible cause: This issue is often due to principals in the two realms having different encryption types or different passwords for the cross-realm principal in each realm. Because the local and remote realm each issue their own TGTs, the local commands run but the service ticket needed for the local and remote realms to communicate cannot be granted.

Steps to resolve: Add the cross-realm `krbtgt` principal and its encryption types to the MIT KDC server using `kadmin.local` or `kadmin` as appropriate (local access vs. remote):

```
kadmin: addprinc -e "enc_type_list" krbtgt/LOCAL-REALM.EXAMPLE.COM@MAIN-REALM.COMPANY.COM
```

Specify the types of encryption supported by the cross-realm principal (`krbtgt`), for example, AES, DES, or RC4. Multiple encryption types can be specified in the *enc_type_list* as long as one of the encryption types matches that of the tickets granted by the KDC in the remote realm. For example:

```
kadmin:  addprinc -e "aes256-cts:normal rc4-hmac:normal des3-hmac-sha1:normal"
krbtgt/LOCAL-REALM.EXAMPLE.COM@MAIN-REALM.COMPANY.COM
```

Users cannot obtain credentials when running Hadoop jobs or commands

Symptom: Users attempt to authenticate but message such as the following displays:

```
13/01/15 17:44:48 DEBUG ipc.Client: Exception encountered while connecting to the server
 : javax.security.sasl.SaslException:
GSS initiate failed [Caused by GSSException: No valid credentials provided (Mechanism
level: Fail to create credential.
(63) - No service creds)]
```

Possible cause: Ticket message may be too large for the UDP protocol (which is used by SASL by default).

Steps to resolve: Force Kerberos to use TCP instead of UDP by adding the following parameter to `libdefaults` in the `krb5.conf` file on the clients where the problem is occurring:

```
[libdefaults]
udp_preference_limit = 1
```

Configure `krb5.conf` through Cloudera Manager, this will automatically get added to `krb5.conf`.

> **Note:**
>
> When sending a message to the KDC, the library will try using TCP before UDP if the size of the ticket message is larger than the setting specified for the `udp_preference_limit` property. If the ticket message is smaller than `udp_preference_limit` setting, then UDP will be tried before TCP. Regardless of the size, both protocols will be tried if the first attempt fails.

Bogus replay exceptions in service logs

Symptom: Multiple valid requests to Kerberos protected services are identified as ***replay*** attempts when they are not. The following exception shows up in the logs for one or more of the Hadoop daemons:

```
2013-02-28 22:49:03,152 INFO  ipc.Server (Server.java:doRead(571)) - IPC Server listener
 on 8020: readAndProcess threw exception javax.security.sasl.SaslException: GSS initiate
 failed [Caused by GSSException: Failure unspecified at GSS-API level (Mechanism l
javax.security.sasl.SaslException: GSS initiate failed [Caused by GSSException: Failure
 unspecified at GSS-API level (Mechanism level: Request is a replay (34))]
        at
com.sun.security.sasl.gsskerb.GssKrb5Server.evaluateResponse(GssKrb5Server.java:159)
       at org.apache.hadoop.ipc.Server$Connection.saslReadAndProcess(Server.java:1040)

        at org.apache.hadoop.ipc.Server$Connection.readAndProcess(Server.java:1213)
        at org.apache.hadoop.ipc.Server$Listener.doRead(Server.java:566)
        at org.apache.hadoop.ipc.Server$Listener$Reader.run(Server.java:363)
Caused by: GSSException: Failure unspecified at GSS-API level (Mechanism level: Request
 is a replay (34))
        at sun.security.jgss.krb5.Krb5Context.acceptSecContext(Krb5Context.java:741)
        at sun.security.jgss.GSSContextImpl.acceptSecContext(GSSContextImpl.java:323)
        at sun.security.jgss.GSSContextImpl.acceptSecContext(GSSContextImpl.java:267)
        at
com.sun.security.sasl.gsskerb.GssKrb5Server.evaluateResponse(GssKrb5Server.java:137)
        ... 4 more
Caused by: KrbException: Request is a replay (34)
        at sun.security.krb5.KrbApReq.authenticate(KrbApReq.java:300)
        at sun.security.krb5.KrbApReq.<init>(KrbApReq.java:134)
       at sun.security.jgss.krb5.InitSecContextToken.<init>(InitSecContextToken.java:79)

        at sun.security.jgss.krb5.Krb5Context.acceptSecContext(Krb5Context.java:724)
        ... 7 more
```

This issue can also manifest as poor performance for clients of the cluster, including dropped connections, timeouts attempting to make RPC calls, and so on.

Possible cause: Kerberos uses a second-resolution timestamp to protect against replay attacks (where an attacker can record network traffic, and play back recorded requests later to gain elevated privileges). That is, incoming requests are cached by Kerberos for a little while, and if there are similar requests within a few seconds, Kerberos will be able to detect them as replay attack attempts (see MIT Kerberos replay cache for more information).

However, if there are multiple valid Kerberos requests coming in at the same time, these may also be misjudged as attacks for the following reasons:

- **Multiple services in the cluster are using the same Kerberos principal.** All secure clients that run on multiple machines should use unique Kerberos principals for each machine. For example, rather than connecting as a service principal `myservice@EXAMPLE.COM`, services should have per-host principals such as `myservice/host123.example.com@EXAMPLE.COM`.
- **Clocks not synchronized:** All hosts should run NTP so that clocks are kept in sync between clients and servers.

Steps to resolve:

While having different principals for each service, and clocks in sync helps mitigate the issue, there are, however, cases where even if all of the above are implemented, the problem still persists. In such a case, disabling the cache (*and the replay protection as a consequence*), will allow parallel requests to succeed. This compromise between usability and security can be applied by setting the `KRB5RCACHETYPE` environment variable to none.

Note that the `KRB5RCACHETYPE` is not automatically detected by Java applications. For Java-based components:

- Ensure that the cluster runs on JDK 8.
- To disable the replay cache, add `-Dsun.security.krb5.rcache=none` to the Java Opts/Arguments of the targeted JVM. For example, HiveServer2 or the Sentry service.

Cloudera Manager cluster services fail to start

Symptom: One or more of the cluster services fails to start. For example, the DataNode fails to start. Error messages in the log files may display the following error messages:

```
Exception in secureMain
java.lang.ExceptionInInitializerError
at javax.crypto.KeyGenerator.nextSpi(KeyGenerator.java:324)
at javax.crypto.KeyGenerator.<init>(KeyGenerator.java:157)
...
Caused by: java.lang.SecurityException: The jurisdiction policy files are not signed by
 a trusted signer!
at javax.crypto.JarVerifier.verifyPolicySigned(JarVerifier.java:289)
at javax.crypto.JceSecurity.loadPolicies(JceSecurity.java:316)
at javax.crypto.JceSecurity.setupJurisdictionPolicies(JceSecurity.java:261)
...
```

Possible cause: This is another example of a mismatch for AES-256 encryption. Services cannot start when the version of the JCE policy file does not match the version of Java installed on a node because the cryptographic signatures for the JCE policy files cannot be verified, resulting in the message shown above.

- Check that the encryption types are matched between your KDC and `krb5.conf` on all hosts.

  Solution: If you are using AES-256, follow the instructions at Step 2: Installing JCE Policy File for AES-256 Encryption on page 53 to deploy the JCE policy file on all hosts.

- Services cannot start

Solution: Download the correct JCE policy files for the version of Java you are running:

- Java 8
- Java 7
- Java 6 [Legacy information]

Download and unpack the zip file. Copy the two JAR files to the `$JAVA_HOME/jre/lib/security` directory on each node within the cluster.

Error Messages

Incorrect permission Java exception (java.io.IOException)

Symptom: An incorrect permission error displays when trying to run a job:

```
java.io.IOException: Incorrect permission for
/var/folders/B3/B3d2vCm4F+mmWzVPB89W6E+++TI/-Tmp-/tmpYTil84/dfs/data/data1,
expected: rwxr-xr-x, while actual: rwxrwxr-x
       at org.apache.hadoop.util.DiskChecker.checkPermission(DiskChecker.java:107)
       at
org.apache.hadoop.util.DiskChecker.mkdirsWithExistsAndPermissionCheck(DiskChecker.java:144)

       at org.apache.hadoop.util.DiskChecker.checkDir(DiskChecker.java:160)
      at org.apache.hadoop.hdfs.server.datanode.DataNode.makeInstance(DataNode.java:1484)

       at
org.apache.hadoop.hdfs.server.datanode.DataNode.instantiateDataNode(DataNode.java:1432)

       at
org.apache.hadoop.hdfs.server.datanode.DataNode.instantiateDataNode(DataNode.java:1408)

      at org.apache.hadoop.hdfs.MiniDFSCluster.startDataNodes(MiniDFSCluster.java:418)

      ...
```

Possible cause: The daemon has umask 0002 rather than 0022.

Steps to resolve: Make sure that the `umask` for `hdfs` and `mapred` is `0022`.

MapReduce (MRv1) Errors

These error messages are associated with MapReduce only (not YARN).

Jobs won't run and cannot access files in mapred.local.dir

Symptom: The TaskTracker log contains the following error message:

```
WARN org.apache.hadoop.mapred.TaskTracker: Exception while localization
java.io.IOException: Job initialization failed (1)
```

Possible cause:

Steps to resolve:

1. Add the `mapred` user to the `mapred` and `hadoop` groups on all hosts.
2. Restart all TaskTrackers.

Jobs cannot run and TaskTracker cannot create local mapred directory

Symptom: The TaskTracker log contains the following error message:

```
11/08/17 14:44:06 INFO mapred.TaskController: main : user is atm
11/08/17 14:44:06 INFO mapred.TaskController: Failed to create directory
/var/log/hadoop/cache/mapred/mapred/local1/taskTracker/atm - No such file or directory
11/08/17 14:44:06 WARN mapred.TaskTracker: Exception while localization
java.io.IOException: Job initialization failed (20)
        at
org.apache.hadoop.mapred.LinuxTaskController.initializeJob(LinuxTaskController.java:191)

        at org.apache.hadoop.mapred.TaskTracker$4.run(TaskTracker.java:1199)
        at java.security.AccessController.doPrivileged(Native Method)
        at javax.security.auth.Subject.doAs(Subject.java:396)
        at
org.apache.hadoop.security.UserGroupInformation.doAs(UserGroupInformation.java:1127)
        at org.apache.hadoop.mapred.TaskTracker.initializeJob(TaskTracker.java:1174)
        at org.apache.hadoop.mapred.TaskTracker.localizeJob(TaskTracker.java:1089)
        at org.apache.hadoop.mapred.TaskTracker.startNewTask(TaskTracker.java:2257)
       at org.apache.hadoop.mapred.TaskTracker$TaskLauncher.run(TaskTracker.java:2221)
Caused by: org.apache.hadoop.util.Shell$ExitCodeException:
        at org.apache.hadoop.util.Shell.runCommand(Shell.java:255)
        at org.apache.hadoop.util.Shell.run(Shell.java:182)
        at org.apache.hadoop.util.Shell$ShellCommandExecutor.execute(Shell.java:375)
        at
org.apache.hadoop.mapred.LinuxTaskController.initializeJob(LinuxTaskController.java:184)

        ... 8 more
```

Possible cause: Mismatch of `mapred.local.dir` values specified in `mapred-site.xml` and `taskcontroller.cfg`. These values should be the same.

Steps to resolve: Verify that the setting for `mapred.local.dir` is the same in both `mapred-site.xml` and `taskcontroller.cfg`, and reconfigure if necessary.

Jobs cannot run and TaskTracker cannot create Hadoop logs directory

Symptom: The TaskTracker log contains an error message similar to the following:

```
11/08/17 14:48:23 INFO mapred.TaskController: Failed to create directory
/home/atm/src/cloudera/hadoop/build/hadoop-0.23.2-cdh3u1-SNAPSHOT/logs1/userlogs/job_201108171441_0004
 - No such file or directory
11/08/17 14:48:23 WARN mapred.TaskTracker: Exception while localization
java.io.IOException: Job initialization failed (255)
        at
org.apache.hadoop.mapred.LinuxTaskController.initializeJob(LinuxTaskController.java:191)

        at org.apache.hadoop.mapred.TaskTracker$4.run(TaskTracker.java:1199)
        at java.security.AccessController.doPrivileged(Native Method)
        at javax.security.auth.Subject.doAs(Subject.java:396)
        at
```

```
org.apache.hadoop.security.UserGroupInformation.doAs(UserGroupInformation.java:1127)
        at org.apache.hadoop.mapred.TaskTracker.initializeJob(TaskTracker.java:1174)
        at org.apache.hadoop.mapred.TaskTracker.localizeJob(TaskTracker.java:1089)
        at org.apache.hadoop.mapred.TaskTracker.startNewTask(TaskTracker.java:2257)
       at org.apache.hadoop.mapred.TaskTracker$TaskLauncher.run(TaskTracker.java:2221)
Caused by: org.apache.hadoop.util.Shell$ExitCodeException:
        at org.apache.hadoop.util.Shell.runCommand(Shell.java:255)
        at org.apache.hadoop.util.Shell.run(Shell.java:182)
        at org.apache.hadoop.util.Shell$ShellCommandExecutor.execute(Shell.java:375)
        at
org.apache.hadoop.mapred.LinuxTaskController.initializeJob(LinuxTaskController.java:184)

        ... 8 more
```

Possible cause: Misconfiguration issue.

Steps to resolve: In MRv1, the default value specified for `hadoop.log.dir` in `mapred-site.xml` is `/var/log/hadoop-0.20-mapreduce`. The path must be owned and be writable by the `mapred` user. If you change the default value specified for `hadoop.log.dir`, make sure the value is identical in `mapred-site.xml` and `taskcontroller.cfg`. If the values are different, the error message above is returned.

Authentication and Kerberos Issues

Clusters that use Kerberos for authentication have several possible sources of potential issues, including:

- Failure of the Key Distribution Center (KDC)
- Missing Kerberos or OS packages or libraries
- Incorrect mapping of Kerberos REALMs for cross-realm authentication

These are just some examples, but they can prevent users and services from authenticating and can interfere with the cluster's ability to run and process workloads. The first step whenever an issue emerges is to try to isolate the source of the actual issue, by answering basic questions such as these:

- Is the issue a local issue or a global issue? That is, are all users failing to authenticate, or is the issue specific to a single user?
- Is the issue specific to a single service, or are all services problematic? and so on.

If all users and multiple services are affected—and if the cluster has not worked at all after integrating with Kerberos for authentication—step through all settings for the Kerberos configuration files, as outlined in Auditing the Kerberos Configuration.

> **Note:** All nodes in any given cluster configured for Kerberos must use the same configuration settings in the files that are distributed throughout the cluster.

Auditing the Kerberos Configuration

Cloudera recommends verifying the Kerberos configuration whenever issues arise, especially after intially completing the integration process.

- Verify that all `/etc/hosts` files conform to Cloudera Manager's installation requirements.
- Verify forward and reverse name resolution for all cluster hosts and for the MIT KDC or Active Directory KDC hosts.
- Verify that all required Kerberos server and workstation packages have been installed and are the correct versions for the OS running on the host systems.
- Verify that the `hadoop.security.auth_to_local` property in the `core-site.xml` has proper mappings for *all* trusted Kerberos realms, including HDFS trusted realms, for all services on the cluster that use Kerberos.
- Verify your Kerberos configuration by comparing to the Sample Kerberos Configuration Files on page 425 shown below (see `krb5.conf` and `kdc.conf`).
- Review the configuration of all the KDC, `REALM`, and domain hosts referenced in the `krb5.conf` and `kdc.conf` files. The KDC host in particular, is a common point-of-failure and you may have to begin troubleshooting there.

Ensure that the REALM set in `krb5.conf` has the correct hostname listed for the KDC. For cross-realm authentication, see Reviewing Service Ticket Credentials in Cross-Realm Deployments on page 425.

- Use whether the services using Kerberos are running and responding properly with `kinit/klist`.
- Attempt to authenticate to Cloudera Manager using cluster service credentials specific to the issue or affected service. Examine the issued credentials if you are able to successfully authenticate with the service keytab.
- Use `klist` to list the principals present within a service keytab to ensure each service has one.
- Enabling debugging using either the command line or Cloudera Manager.

Kerberos Command-Line Tools

User Authentication with and Without Keytab

The `kinit` command line tool is used to authenticate a user, service, system, or device to a KDC. The most basic example is a user authenticating to Kerberos with a username (principal) and password. In the following example, the first attempt uses a wrong password, followed by a second successful attempt.

```
[alice@host1 ~]$ kinit alice@TEST.ORG.LAB
Password for alice@TEST.ORG.LAB: (wrong password)
kinit: Preauthentication failed while getting initial credentials

[alice@host1 ~]$ kinit alice@TEST.ORG.LAB
Password for alice@TEST.ORG.LAB: (correct password)
(note silent return on successful auth)
[alice@host1 ~]$ klist
Ticket cache: FILE:/tmp/krb5cc_10001
Default principal: alice@TEST.ORG.LAB

Valid starting     Expires            Service principal
03/11/14 11:55:39  03/11/14 21:54:55  krbtgt/TEST.ORG.LAB@TEST.ORG.LAB
renew until 03/18/14 11:55:39
```

Another method of authentication is using keytabs with the `kinit` command. You can verify whether authentication was successful by using the `klist` command to show the credentials issued by the KDC. The following example attempts to authenticate the `hdfs` service to the KDC by using the `hdfs` keytab file.

```
[root@host1 312-hdfs-DATANODE]# kinit -kt hdfs.keytab hdfs/host1.test.lab@TEST.LAB
[root@host1 312-hdfs-DATANODE]# klist
Ticket cache: FILE:/tmp/krb5cc_0
Default principal: hdfs/host1.test.lab@TEST.LAB

Valid starting      Expires      Service principal
03/11/14 11:18:34 03/12/14 11:18:34 krbtgt/TEST.LAB@TEST.LAB
renew until 03/18/14 11:18:34
```

Examining Kerberos credentials with `klist`

So far we've only seen basic usage examples of the `klist` command to list the contents of a keytab file, or to examine a user's credentials. To get more information from the `klist` command, such as the encryption types being negotiated, or the flags being set for credentials being issued by the KDC, use the `klist -ef` command. The output for this command will show you the negotiated encryption types for a user or service principal. This is useful information because you may troubleshoot errors caused (especially in cross-realm trust deployments) because an AD or MIT KDC server may not support a particular encryption type. Look for the encryption types under the "Etype" section of the output.

Flags indicate options supported by Kerberos that extend the features of a set of issued credentials. As discussed previously, CDH requires renewable as well as forwardable tickets for successful authentication, especially in cross realm environments. Look for these settings in the "Flags:" section of the `klist -ef` output shown below where, F = Forwardable, and, R = renewable.

For example, if you use the `klist -ef` command in an ongoing user session:

```
[alice@host1 ~]$ klist -ef
Ticket cache: FILE:/tmp/krb5cc_10001
```

```
Default principal: alice@TEST.ORG.LAB
Valid starting Expires Service principal
03/11/14 11:55:39 03/11/14 21:54:55 krbtgt/TEST.ORG.LAB@TEST.ORG.LAB
renew until 03/18/14 11:55:39, Flags: FRIA
Etype (skey, tkt): aes256-cts-hmac-sha1-96,
    aes256-cts-hmac-sha1-96
```

Enabling Debugging for Authentication Issues

Using Cloudera Manager for Debugging

To obtain additional information in the logs and facilitate troubleshooting, administrators can set debug levels for any of the services running on Cloudera Manager Server. Typically, the settings are added using the Advanced Configuration Snippet (Safety Valve) settings for the specific service, the names are specific to the service.

as for HDFS as detailed below:

1. Log in to the Cloudera Manager Admin Console.
2. Select **Clusters** > **HDFS-*n***.
3. Click the **Configuration** tab.
4. Search for properties specific to the different role types for which you want to enable debugging. For example, if you want to enable debugging for the HDFS NameNode, search for the **NameNode Logging Threshold** property and select at least DEBUG level logging.
5. Enable Kerberos debugging by using the HDFS service's Advanced Configuration Snippet. Once again, this may be different for each specific role type or service. For the HDFS NameNode, add the following properties to the HDFS Service Environment Safety Valve:

```
HADOOP_JAAS_DEBUG=true
HADOOP_OPTS="-Dsun.security.krb5.debug=true"
```

6. Click **Save Changes**.
7. Restart the HDFS service.

The output will be seen in the process logs: `stdout.log` and `stderr.log`. These can be found in the runtime path of the instance:

```
/var/run/cloudera-scm-agent/process/###-service-ROLE
```

After restarting Cloudera Manager Service, the most recent instance of the `###-service-ROLE` directory will have debug logs. Use `ls -ltr` in the `/var/run/cloudera-scm-agent/process` path to determine the most current path.

Enabling Debugging for the Authentication Process

Set the following properties on the cluster to obtain debugging information from the Kerberos authentication process.

```
# export HADOOP_ROOT_LOGGER=TRACE,console;
# export HADOOP_JAAS_DEBUG=true;
# export HADOOP_OPTS="-Dsun.security.krb5.debug=true"
```

You can then use the following command to copy the console output to the user (with debugging messages), along with all output from STDOUT and STDERR to a file.

```
# hadoop fs -ls / > >(tee fsls-logfile.txt) 2>&1
```

Kerberos Credential-Generation Issues

Cloudera Manager creates accounts needed by CDH services using an internal command (**Generate Credentials**) that is triggered automatically by the Kerberos configuration wizard or when changes are made to the cluster that require new Kerberos principals.

After configuring the cluster for Kerberos authentication or making changes that require generation of new principals, you can verify that the command ran successfully by using the Cloudera Manager Admin Console, as follows:

1. Log in to the Cloudera Manager Admin Console. Any error messages display on the **Home** page, in the **Status** area near the top of the page. The following Status message indicates that the Generate Credentials command failed:

```
Role is missing Kerberos
            keytab
```

2. To display the output of the command, go to the **Home** > **Status** tab and click the **All Recent Commands** tab.

Active Directory Credential-Generation Errors

Error: `ldap_sasl_interactive_bind_s: Can't contact LDAP server (-1)`

Possible cause: The Domain Controller specified is incorrect or LDAPS has not been enabled for it.

Steps to resolve: Verify the configuration for Active Directory Active Directory KDC, as follows:

1. Log in to Cloudera Manager Admin Console.
2. Select **Administration** > **Settings** .
3. Select **Kerberos** for the **Category** filter.

Verify all settings. Also check that LDAPS is enabled for Active Directory.

Error: `ldap_add: Insufficient access (50)`

Possible cause: The Active Directory account you are using for Cloudera Manager does not have permissions to create other accounts.

Steps to resolve: Use the Delegate Control wizard to grant permission to the Cloudera Manager account to create other accounts. You can also login to Active Directory as the Cloudera Manager user to check that it can create other accounts in your Organizational Unit.

MIT Kerberos Credential-Generation Errors

Error: `kadmin: Cannot resolve network address for admin server in requested realm while initializing kadmin interface.`

Possible cause: The hostname for the KDC server is incorrect.

Steps to resolve: Check the `kdc` field for your default realm in `krb5.conf` and make sure the hostname is correct.

Hadoop commands fail after enabling Kerberos security

Users need to obtain valid Kerberos tickets to interact with a secure cluster, that is, a cluster that has been configured to use Kerberos for authentication. Running any Hadoop command (such as `hadoop fs -ls`) will fail if you do not have a valid Kerberos ticket in your credentials cache. If you do not have a valid ticket, you will receive an error such as:

```
11/01/04 12:08:12 WARN ipc.Client: Exception encountered while connecting to the server
 : javax.security.sasl.SaslException:
GSS initiate failed [Caused by GSSException: No valid credentials provided (Mechanism
level: Failed to find any Kerberos tgt)]
Bad connection to FS. command aborted. exception: Call to nn-host/10.0.0.2:8020 failed
 on local exception: java.io.IOException:
javax.security.sasl.SaslException: GSS initiate failed [Caused by GSSException: No valid
 credentials provided (Mechanism level: Failed to find any Kerberos tgt)]
```

Steps to resolve: Examine the Kerberos tickets currently in your credentials cache by running the `klist` command. You can obtain a ticket by running the `kinit` command and either specifying a keytab file containing credentials, or entering the password for your principal.

Using the UserGroupInformation class to authenticate Oozie

Secured CDH services mainly use Kerberos to authenticate RPC communication. RPCs are one of the primary means of communication between nodes in a Hadoop cluster. For example, RPCs are used by the YARN NodeManager to communicate with the ResourceManager, or by the HDFS client to communicate with the NameNode.

CDH services handle Kerberos authentication by calling the UserGroupInformation (UGI) login method, `loginUserFromKeytab()`, once every time the service starts up. Since Kerberos ticket expiration times are typically short, repeated logins are required to keep the application secure. Long-running CDH applications have to be implemented accordingly to accommodate these repeated logins. If an application is only going to communicate with HDFS, YARN, MRv1, and HBase, then you only need to call the `UserGroupInformation.loginUserFromKeytab()` method at startup, before any actual API activity occurs. The HDFS, YARN, MRv1 and HBase services' RPC clients have their own built-in mechanisms to automatically re-login when a keytab's Ticket-Granting Ticket (TGT) expires. Therefore, such applications do not need to include calls to the UGI re-login method because their RPC client layer performs the re-login task for them.

However, some applications may include other service clients that do not involve the generic Hadoop RPC framework, such as Hive or Oozie clients. Such applications must explicitly call the `UserGroupInformation.getLoginUser().checkTGTAndReloginFromKeytab()` method before every attempt to connect with a Hive or Oozie client. This is because these clients do not have the logic required for automatic re-logins.

This is an example of an infinitely polling Oozie client application:

```
// App startup
UserGroupInformation.loginFromKeytab(KEYTAB_PATH, PRINCIPAL_STRING);
OozieClient client = loginUser.doAs(new PrivilegedAction<OozieClient>() {
    public OozieClient run() {
        try {
            returnnew OozieClient(OOZIE_SERVER_URI);
        } catch (Exception e) {
            e.printStackTrace();
            returnnull;
        }
    }
});

while (true && client != null) {
    // Application's long-running loop

    // Every time, complete the TGT check first
    UserGroupInformation loginUser = UserGroupInformation.getLoginUser();
    loginUser.checkTGTAndReloginFromKeytab();

    // Perform Oozie client work within the context of the login user object
    loginUser.doAs(new PrivilegedAction<Void>() {
        publicVoid run() {
            try {
                List<WorkflowJob> list = client.getJobsInfo("");
                for (WorkflowJob wfJob : list) {
                    System.out.println(wfJob.getId());
                }
            } catch (Exception e) {
                e.printStackTrace();
            }
        } // End of function block
    }); // End of doAs
} // End of loop
```

Certain Java versions cannot read credentials cache

Symptom: For MIT Kerberos 1.8.1 (or higher), the following error will occur when you attempt to interact with the Hadoop cluster, even after successfully obtaining a Kerberos ticket using `kinit`:

```
11/01/04 12:08:12 WARN ipc.Client: Exception encountered while connecting to the server
 : javax.security.sasl.SaslException:
GSS initiate failed [Caused by GSSException: No valid credentials provided (Mechanism
level: Failed to find any Kerberos tgt)]
Bad connection to FS. command aborted. exception: Call to nn-host/10.0.0.2:8020 failed
 on local exception: java.io.IOException:
javax.security.sasl.SaslException: GSS initiate failed [Caused by GSSException: No valid
 credentials provided (Mechanism level: Failed to find any Kerberos tgt)]
```

Possible cause:

At release 1.8.1 of MIT Kerberos, a change was made to the credentials cache format that conflicts with Oracle JDK 6 Update 26 (and earlier JDKs) rendering Java incapable of reading Kerberos credentials cache created by versions of MIT Kerberos 1.8.1 (or higher). Kerberos 1.8.1 is the default in Ubuntu Lucid and higher releases and Debian Squeeze and higher releases. On RHEL and CentOS, an older version of MIT Kerberos which does not have this issue, is the default.

Workaround: Use the -R (renew) option with `kinit` after initially obtaining credentials with `kinit`. This sequence causes the ticket to be renewed and credentials are cached using a format that Java can read. However, the initial ticket must be renewable.

For example:

```
$ klist
klist: No credentials cache found (ticket cache FILE:/tmp/krb5cc_1000)
$ hadoop fs -ls
11/01/04 13:15:51 WARN ipc.Client: Exception encountered while connecting to the server
 : javax.security.sasl.SaslException:
GSS initiate failed [Caused by GSSException: No valid credentials provided (Mechanism
level: Failed to find any Kerberos tgt)]
Bad connection to FS. command aborted. exception: Call to nn-host/10.0.0.2:8020 failed
 on local exception: java.io.IOException:
javax.security.sasl.SaslException: GSS initiate failed [Caused by GSSException: No valid
 credentials provided (Mechanism level: Failed to find any Kerberos tgt)]
$ kinit
Password for username@REALM-NAME.EXAMPLE.COM:
$ klist
Ticket cache: FILE:/tmp/krb5cc_1000
Default principal: username@REALM-NAME.EXAMPLE.COM

Valid starting     Expires            Service principal
01/04/11 13:19:31  01/04/11 23:19:31  krbtgt/REALM-NAME.EXAMPLE.COM@REALM-NAME.EXAMPLE.COM

 renew until 01/05/11 13:19:30
$ hadoop fs -ls
11/01/04 13:15:59 WARN ipc.Client: Exception encountered while connecting to the server
 : javax.security.sasl.SaslException:
GSS initiate failed [Caused by GSSException: No valid credentials provided (Mechanism
level: Failed to find any Kerberos tgt)]
Bad connection to FS. command aborted. exception: Call to nn-host/10.0.0.2:8020 failed
 on local exception: java.io.IOException:
javax.security.sasl.SaslException: GSS initiate failed [Caused by GSSException: No valid
 credentials provided (Mechanism level: Failed to find any Kerberos tgt)]
$ kinit -R
$ hadoop fs -ls
Found 6 items
drwx------   - user user     0 2011-01-02
          16:16 /user/user/.staging
```

Non-renewable tickets display this error message when the command

```
kinit: Ticket expired while renewing credentials
```

Resolving Cloudera Manager Service keytab Issues

Every service managed by Cloudera Manager has a keytab file that is provided at startup by the Cloudera Manager Agent. The most recent keytab files can be examined by navigating to the path, `/var/run/cloudera-scm-agent/process`, with an `ls -ltr` command.

As you can see in the example below, Cloudera Manager service directory names have the form: `###-service-ROLE`. Therefore, if you are troubleshooting the HDFS service, the service directory may be called, `326-hdfs-NAMENODE`.

```
[root@cehd1 ~]# cd /var/run/cloudera-scm-agent/process/
[root@cehd1 process]# ls -ltr | grep NAMENODE | tail -3
drwxr-x--x 3 hdfs         hdfs         4096 Mar  3 23:43 313-hdfs-NAMENODE
drwxr-x--x 3 hdfs         hdfs         4096 Mar  4 00:07 326-hdfs-NAMENODE
drwxr-x--x 3 hdfs         hdfs         4096 Mar  4 00:07 328-hdfs-NAMENODE-nnRpcWait

[root@cehd1 process]# cd 326-hdfs-NAMENODE

[root@cehd1 326-hdfs-NAMENODE]# ls
cloudera_manager_agent_fencer.py              dfs_hosts_allow.txt         hdfs.keytab
                log4j.properties             topology.py
cloudera_manager_agent_fencer_secret_key.txt  dfs_hosts_exclude.txt      hdfs-site.xml
               logs
cloudera-monitor.properties                   event-filter-rules.json
http-auth-signature-secret  navigator.client.properties
core-site.xml                                 hadoop-metrics2.properties  krb5cc_494
                topology.map
```

If you have root access to the `/var/run/cloudera-scm-agent/process` path, you can use any service's keytab file to log in as root or a sudo user to verify whether basic Kerberos authentication is working.

After locating a keytab file, examine its contents using the `klist` command to view the credentials stored in the file. For example, to list the credentials stored in the `hdfs.keytab` file:

```
[root@host1 326-hdfs-DATANODE]# klist -kt hdfs.keytab

Keytab name: WRFILE:hdfs.keytab
KVNO Timestamp         Principal
---- ----------------- --------------------------------------------------------
   4 02/17/14 19:09:17 HTTP/host1.test.lab@TEST.LAB
   4 02/17/14 19:09:17 HTTP/host1.test.lab@TEST.LAB
   4 02/17/14 19:09:17 HTTP/host1.test.lab@TEST.LAB
   4 02/17/14 19:09:17 HTTP/host1.test.lab@TEST.LAB
   4 02/17/14 19:09:17 HTTP/host1.test.lab@TEST.LAB
   4 02/17/14 19:09:17 HTTP/host1.test.lab@TEST.LAB
   4 02/17/14 19:09:17 hdfs/host1.test.lab@TEST.LAB
   4 02/17/14 19:09:17 hdfs/host1.test.lab@TEST.LAB
   4 02/17/14 19:09:17 hdfs/host1.test.lab@TEST.LAB
   4 02/17/14 19:09:17 hdfs/host1.test.lab@TEST.LAB
   4 02/17/14 19:09:17 hdfs/host1.test.lab@TEST.LAB
   4 02/17/14 19:09:17 hdfs/host1.test.lab@TEST.LAB
```

Now, attempt to authenticate using the keytab file and a principal within it. In this case, we use the `hdfs.keytab` file with the `hdfs/host1.test.lab@TEST.LAB` principal. Then use the `klist` command without any arguments to see the current user session's credentials.

```
root@host1 312-hdfs-DATANODE]# kinit -kt hdfs.keytab hdfs/host1.test.lab@TEST.LAB
[root@host1 312-hdfs-DATANODE]# klist

Ticket cache: FILE:/tmp/krb5cc_0
Default principal: hdfs/host1.test.lab@TEST.LAB

Valid starting     Expires            Service principal
03/11/14 11:18:34  03/12/14 11:18:34  krbtgt/TEST.LAB@TEST.LAB
renew until 03/18/14 11:18:34
```

Note that Kerberos credentials have an expiry date and time. This means, to make sure Kerberos credentials are valid uniformly over a cluster, all hosts and clients within the cluster should be using NTP and must never drift more than 5 minutes apart from each other. Kerberos session tickets have a limited lifespan, but can be renewed (as indicated in

the sample `krb5.conf` and `kdc.conf`). CDH requires renewable tickets for cluster principals. Check whether renewable tickets have been enabled by using a klist command with the `-e` (list key encryption types) and `-f` (list flags set) switches when examining Kerberos sessions and credentials.

Reviewing Service Ticket Credentials in Cross-Realm Deployments

When you examine your cluster configuration, make sure you haven't violated any of the following integration rules:

- When negotiating encryption types, follow the realm with the most specific limitations on supported encryption types.
- All realms should be known to one another through the `/etc/krb5.conf` file deployed on the cluster.
- When you make configuration decisions for Active Directory environments, you must evaluate the Domain Functional Level or Forrest Functional Level that is present.

Kerberos typically negotiates and uses the strongest form of encryption possible between a client and server for authentication into the realm. However, the encryption types for TGTs may sometimes end up being negotiated downward towards the weaker encryption types, which is not desirable. To investigate such issues, check the `kvno` of the cross-realm trust principal (`krbtgt`) as described in the following steps. Replace `CLUSTER.REALM` and `AD.REALM` (or `MIT.REALM`) with the appropriate values for your configured realm. This scenario assumes cross-realm authentication with Active Directory.

- Once trust has been configured (see sample files in previous section), `kinit` as a system user by authenticating to the AD Kerberos realm.
- From the command line, perform a `kvno` check of the local and cross-realm `krbtgt` entry. The local representation of this special `REALM` service principal is in the form, `krbtgt/CLUSTER.REALM@CLUSTER.REALM`. The cross-realm principal is named after the trusted realm in the form, `krbtgt/AD.REALM`.
- Failure of the `kvno` check indicates incorrect cross-realm trust configuration. Review encryption types again, looking for incompatibilities or unsupported encryption types configured between realms.

Sample Kerberos Configuration Files

This section contains several example Kerberos configuration files.

/etc/krb5.conf

The `/etc/krb5.conf` file is the configuration a client uses to access a realm through its configured KDC. The `krb5.conf` maps the realm to the available servers supporting those realms. It also defines the host-specific configuration rules for how tickets are requested and granted.

```
[logging]
 default = FILE:/var/log/krb5libs.log
 kdc = FILE:/var/log/krb5kdc.log
 admin_server = FILE:/var/log/kadmind.log

[libdefaults]
 default_realm = EXAMPLE.COM
 dns_lookup_realm = false
 dns_lookup_kdc = false
 ticket_lifetime = 24h
 renew_lifetime = 7d
 forwardable = true
# udp_preference_limit = 1

# set udp_preference_limit = 1 when TCP only should be
# used. Consider using in complex network environments when
# troubleshooting or when dealing with inconsistent
# client behavior or GSS (63) messages.

# uncomment the following if AD cross realm auth is ONLY providing DES encrypted tickets
# allow-weak-crypto = true

[realms]
 AD-REALM.EXAMPLE.COM = {
  kdc = AD1.ad-realm.example.com:88
```

```
  kdc = AD2.ad-realm.example.com:88
  admin_server = AD1.ad-realm.example.com:749
  admin_server = AD2.ad-realm.example.com:749
  default_domain = ad-realm.example.com
 }
 EXAMPLE.COM = {
  kdc = kdc1.example.com:88
  admin_server = kdc1.example.com:749
  default_domain = example.com
 }

# The domain_realm is critical for mapping your host domain names to the kerberos realms
# that are servicing them. Make sure the lowercase left hand portion indicates any
domains or subdomains
# that will be related to the kerberos REALM on the right hand side of the expression.
 REALMs will
# always be UPPERCASE. For example, if your actual DNS domain was test.com but your
kerberos REALM is
# EXAMPLE.COM then you would have,

[domain_realm]
test.com = EXAMPLE.COM
#AD domains and realms are usually the same
ad-domain.example.com = AD-REALM.EXAMPLE.COM
ad-realm.example.com = AD-REALM.EXAMPLE.COM
```

/var/kerberos/krb5kdc.conf

The `kdc.conf` file only needs to be configured on the actual cluster-dedicated KDC, and should be located at `/var/kerberos/krb5kdc`. Only primary and secondary KDCs need access to this configuration file. The contents of this file establish the configuration rules which are enforced for all client hosts in the `REALM`.

```
[kdcdefaults]
 kdc_ports = 88
 kdc_tcp_ports = 88

[realms]
  EXAMPLE.COM = {
  #master_key_type = aes256-cts
  max_renewable_life = 7d 0h 0m 0s
  acl_file = /var/kerberos/krb5kdc/kadm5.acl
  dict_file = /usr/share/dict/words
  admin_keytab = /var/kerberos/krb5kdc/kadm5.keytab
# note that aes256 is ONLY supported in Active Directory in a domain / forrest operating
 at a 2008 or greater functional level.
# aes256 requires that you download and deploy the JCE Policy files for your JDK release
 level to provide
# strong java encryption extension levels like AES256. Make sure to match based on the
 encryption configured within AD for
# cross realm auth, note that RC4 = arcfour when comparing windows and linux enctypes
  supported_enctypes = aes256-cts:normal aes128-cts:normal arcfour-hmac:normal
  default_principal_flags = +renewable, +forwardable
 }
```

kadm5.acl

```
*/admin@HADOOP.COM *
cloudera-scm@HADOOP.COM * flume/*@HADOOP.COM
cloudera-scm@HADOOP.COM * hbase/*@HADOOP.COM
cloudera-scm@HADOOP.COM * hdfs/*@HADOOP.COM
cloudera-scm@HADOOP.COM * hive/*@HADOOP.COM
cloudera-scm@HADOOP.COM * httpfs/*@HADOOP.COM
cloudera-scm@HADOOP.COM * HTTP/*@HADOOP.COM
cloudera-scm@HADOOP.COM * hue/*@HADOOP.COM
cloudera-scm@HADOOP.COM * impala/*@HADOOP.COM
cloudera-scm@HADOOP.COM * mapred/*@HADOOP.COM
cloudera-scm@HADOOP.COM * oozie/*@HADOOP.COM
cloudera-scm@HADOOP.COM * solr/*@HADOOP.COM
cloudera-scm@HADOOP.COM * sqoop/*@HADOOP.COM
```

```
cloudera-scm@HADOOP.COM * yarn/*@HADOOP.COM
cloudera-scm@HADOOP.COM * zookeeper/*@HADOOP.COM
```

HDFS Encryption Issues

The following are possible workarounds for issues that may arise when encrypting HDFS directories and files. HDFS encryption is sometimes referred to in the documentation as HDFS Transparent Encryption or as HDFS Data at Rest Encryption.

KMS server jute buffer exception

Description

You see the following error when the KMS (for example, as a ZooKeeper client) jute buffer size is insufficient to hold all the tokens:

```
2017-01-31 21:23:56,416 WARN org.apache.zookeeper.ClientCnxn: Session 0x259f5fb3c1000fb
 for server example.cloudera.com/10.172.0.1:2181, unexpected error, closing socket
connection and attempting reconnect
java.io.IOException: Packet len4196356 is out of range!
```

Solution

Increase the jute buffer size and restart the KMS. In Cloudera Manager, go to the KMS **Configuration** page, and in the **Additional Java Configuration Options for KMS** (`kms_java_opts`) field, enter `-Djute.maxbuffer=<number_of_bytes>`. Restart the KMS.

Retrieval of encryption keys fails

Description

You see the following error when trying to list encryption keys

```
user1@example-sles-4:~> hadoop key list
Cannot list keys for KeyProvider: KMSClientProvider[https:
//example-sles-2.example.com:16000/kms/v1/]: Retrieval of all keys failed.
```

Solution

Make sure your truststore has been updated with the relevant certificate(s), such as the Key Trustee server certificate.

DistCp between unencrypted and encrypted locations fails

Description

By default, DistCp compares checksums provided by the filesystem to verify that data was successfully copied to the destination. However, when copying between unencrypted and encrypted locations, the filesystem checksums will not match since the underlying block data is different.

Solution

Specify the `-skipcrccheck` and `-update` distcp flags to avoid verifying checksums.

(CDH 5.6 and lower) Cannot move encrypted files to trash

> **Note:** Starting with CDH 5.7, you can delete files or directories that are part of an HDFS encryption zone. For details, see Trash Behavior with HDFS Transparent Encryption Enabled.

Description

In CDH 5.6 and lower, with HDFS encryption enabled, you cannot move encrypted files or directories to the trash directory.

Solution

To remove encrypted files/directories, use the following command with the `-skipTrash` flag specified to bypass trash.

```
rm -r -skipTrash /testdir
```

NameNode - KMS communication fails after long periods of inactivity

Description

Encrypted files and encryption zones cannot be created if a long period of time (by default, 20 hours) has passed since the last time the KMS and NameNode communicated.

Solution

> **Important:** Upgrading your cluster to the latest CDH 5 release will fix this problem. For instructions, see Upgrading from an Earlier CDH 5 Release to the Latest Release.

For lower CDH 5 releases, there are two possible workarounds to this issue :

- You can increase the KMS authentication token validity period to a very high number. Since the default value is 10 hours, this bug will only be encountered after 20 hours of no communication between the NameNode and the KMS. Add the following property to the `kms-site.xml`Safety Valve:

  ```
  <property>
  <name>hadoop.kms.authentication.token.validity</name>
  <value>SOME VERY HIGH NUMBER</value>
  </property>
  ```

- You can switch the KMS signature secret provider to the string secret provider by adding the following property to the `kms-site.xml` Safety Valve:

  ```
  <property>
  <name>hadoop.kms.authentication.signature.secret</name>
  <value>SOME VERY SECRET STRING</value>
  </property>
  ```

HDFS Trash Behaviour with Transparent Encryption Enabled

The Hadoop trash feature helps prevent accidental deletion of files and directories. When you delete a file in HDFS, the file is not immediately expelled from HDFS. Deleted files are first moved to the `/user/<username>/.Trash/Current` directory, with their original filesystem path being preserved. After a user-configurable period of time (`fs.trash.interval`), a process known as trash checkpointing renames the `Current` directory to the current timestamp, that is, `/user/<username>/.Trash/<timestamp>`. The checkpointing process also checks the rest of the `.Trash` directory for any existing timestamp directories and removes them from HDFS permanently. You can restore files and directories in the trash simply by moving them to a location outside the `.Trash` directory.

Trash Behaviour with HDFS Transparent Encryption Enabled

Starting with CDH 5.7, you can delete files or directories that are part of an HDFS encryption zone. As is evident from the procedure described above, moving and renaming files or directories is an important part of trash handling in HDFS. However, currently HDFS transparent encryption only supports renames *within* an encryption zone. To accommodate this, HDFS creates a local `.Trash` directory every time a new encryption zone is created. For example, when you create an encryption zone, `enc_zone`, HDFS will also create the `/enc_zone/.Trash/` subdirectory. Files

deleted from enc_zone are moved to /enc_zone/.Trash/<username>/Current/. After the checkpoint, the Current directory is renamed to the current timestamp, /enc_zone/.Trash/<username>/<timestamp>.

If you delete the entire encryption zone, it will be moved to the .Trash directory under the user's home directory, /users/<username>/.Trash/Current/enc_zone. Trash checkpointing will occur only after the entire zone has been moved to /users/<username>/.Trash. However, if the user's home directory is already part of an encryption zone, then attempting to delete an encryption zone will fail because you cannot move or rename directories across encryption zones.

Troubleshooting TLS/SSL Issues in Cloudera Manager

To diagnose and resolve issues related to TLS/SSL configuration, verify configuration tasks appropriate for the cluster by verifying the steps in How to Configure TLS for Cloudera Manager.

After checking your settings and finding no obvious misconfiguration, try some of the troubleshooting steps below.

Test Connectivity with OpenSSL

From the host that has connectivity issues, run openssl as shown below. You can also check that the certificate used by the host is recognized by a trusted CA during the TLS/SSL negotiation.

To check the connection:

```
$ openssl s_client -connect [host.fqdn.name]:[port]
```

For example:

```
$ openssl s_client -connect test1.sec.cloudera.com:7183
```

A return code 0 means openssl was able to follow trust server chain of trust through its library of trusted public CAs.

For certificates signed by your organization's internal CA or self-signed certificates, you may need to add the certificate to the truststore using the openssl command. Use the -CAfile option to specify the path to the root CA so openssl can verify the self-signed or internal-CA-signed certificate as follows:

```
$ openssl s_client -connect test1.sec.cloudera.com:7183 -CAfile \
/opt/cloudera/security/CAcerts/RootCA.pem
```

Only the Root CA certificate is needed to establish trust for this test. The result from the command is successful when you see the return code 0 as follows:

```
...
 Verify return code: 0 (ok)
---
```

By default, Cloudera Manager Server writes logs to the /etc/cloudera-scm-server/cloudera-scm-server.log file on startup. Successful server start-up using the certificate looks similar to the following log example:

```
2014-10-06 21:33:47,515 INFO WebServerImpl:org.mortbay.log: jetty-6.1.26.cloudera.2
2014-10-06 21:33:47,572 INFO WebServerImpl:org.mortbay.log: Started
SslSelectChannelConnector@0.0.0.0:7183
2014-10-06 21:33:47,573 INFO WebServerImpl:org.mortbay.log: Started
SelectChannelConnector@0.0.0.0:7180
2014-10-06 21:33:47,573 INFO WebServerImpl:com.cloudera.server.cmf.WebServerImpl: Started
 Jetty server.
```

Upload Diagnostic Bundles to Cloudera Support

By default, Cloudera Manager uploads diagnostic bundles over HTTPS to the Cloudera Support server at cops.cloudera.com. However, the upload can fail if the Cloudera Manager truststore cannot verify the authenticity

of the Cloudera Support server certificate, and that verification process can fail due to Cloudera Manager truststore configuration issues.

To ensure the Cloudera Manager Server truststore contains the public CAs needed to verify Cloudera Support's certificate, you can explicitly establish trust by importing Cloudera Support's certificate into Cloudera Manager's truststore.

> **Note:** Cloudera Support servers use certificates signed by a commercial CA, so this step is typically not needed, unless the default truststore has been altered. Before downloading or adding any certificates, test the connection and verify that the certificate is the source of the connection issue.

Importing Cloudera Support's Certificate into the Cloudera Manager Server Truststore

To obtain Cloudera's public key certificate from the Cloudera Support server:

```
$ openssl s_client -connect cops.cloudera.com:443 | openssl x509 -text -out
/path/to/cloudera-cert.pem
```

To import this certificate into the Cloudera Manager truststore (use paths for your own system):

```
$ keytool -import -keystore /path/to/cm/truststore.jks -file /path/to/cloudera-cert.pem
```

After importing the certificate, confirm that Cloudera Manager is configured for this truststore file, as detailed in Configuring Cloudera Manager Truststore Properties on page 430.

> **Note:** Alternatively, you can use the default Java truststore for your Cloudera Manager cluster deployment, as described in Configuring TLS Encryption for Cloudera Manager on page 195.

Configuring Cloudera Manager Truststore Properties

After installing the Cloudera Support server certificate into the Cloudera Manager truststore, you must configure Cloudera Manager to use the truststore, as follows:

1. Log into the Cloudera Manager Admin Console.
2. Select **Administration** > **Settings**.
3. Click the **Security** category.
4. Enter the path to the truststore and the password (if necessary):

Setting	Description
Cloudera Manager TLS/SSL Certificate Trust Store File	Enter the complete Cloudera Manager Server host filesystem path to the truststore (the trust `.jks`). Cloudera Manager Server invokes JVM with `-Djavax.net.ssl.trustStore` to access the specified truststore.
Cloudera Manager TLS/SSL Certificate Trust Store Password	Specify the password (if there is one) for the truststore file. Password is not required to access the truststore, so you can typically leave this field blank. Cloudera Manager Server invokes JVM with `-Djavax.net.ssl.trustStore.password` if this field has an entry.

5. Click **Save Changes** to save the settings.

> **Note:** See Oracle's JSSE Reference Guide for more information about the JSSE trust mechanism.

YARN, MRv1, and Linux OS Security

Several subsystems are fundamental to Hadoop clusters, specifically, the Jsvc, TaskController, and Container Executor Programs, documented below.

MRv1 and YARN: The jsvc Program

The `jsvc` is part of the Hadoop Apache Commons libraries designed to make Java applications run better on Linux. The `jsvc`program is part of the `bigtop-jsvc` package and installed in `/usr/lib/bigtop-utils/jsvc` or `/usr/libexec/bigtop-utils/jsvc` depending on version of Linux.

In particular, `jsvc` is used to start the DataNode listening on low port numbers. Its entry point is the `SecureDataNodeStarter` class, which implements the Daemon interface that `jsvc` expects. `jsvc` is run as root, and calls the `SecureDataNodeStarter.init(...)` method while running as root. Once the `SecureDataNodeStarter` class has finished initializing, `jsvc` sets the effective UID to be the `hdfs` user, and then calls `SecureDataNodeStarter.start(...)`. `SecureDataNodeStarter` then calls the regular DataNode entry point, passing in a reference to the privileged resources it previously obtained.

MRv1 Only: The Linux TaskController

A setuid binary called `task-controller` is part of the `hadoop-0.20-mapreduce` package and is installed in either `/usr/lib/hadoop-0.20-mapreduce/sbin/Linux-amd64-64/task-controller` or `/usr/lib/hadoop-0.20-mapreduce/sbin/Linux-i386-32/task-controller`.

This `task-controller` program, which is used on MRv1 only, allows the TaskTracker to run tasks under the Unix account of the user who submitted the job in the first place. It is a setuid binary that must have a very specific set of permissions and ownership to function correctly. In particular, it must:

1. Be owned by root
2. Be owned by a group that contains only the user running the MapReduce daemons
3. Be setuid
4. Be group readable and executable

This corresponds to the ownership `root:mapred` and the permissions `4754`.

Here is the output of `ls` on a correctly-configured Task-controller:

```
-rwsr-xr-- 1 root mapred 30888 Mar 18 13:03 task-controller
```

The TaskTracker will check for this configuration on start up, and fail to start if the Task-controller is not configured correctly.

YARN Only: The Linux Container Executor

A setuid binary called `container-executor` is part of the `hadoop-yarn` package and is installed in `/usr/lib/hadoop-yarn/bin/container-executor`.

This `container-executor` program, which is used on YARN only and supported on GNU/Linux only, runs the containers as the user who submitted the application. It requires all user accounts to be created on the cluster hosts where the containers are launched. It uses a setuid executable that is included in the Hadoop distribution. The NodeManager uses this executable to launch and kill containers. The setuid executable switches to the user who has submitted the application and launches or kills the containers. For maximum security, this executor sets up restricted permissions and user/group ownership of local files and directories used by the containers such as the shared objects, jars, intermediate files, and log files. As a result, only the application owner and NodeManager can access any of the local files/directories including those localized as part of the distributed cache.

Parcel Deployments

In a parcel deployment the `container-executor` file is located inside the parcel at `/opt/cloudera/parcels/CDH/lib/hadoop-yarn/bin/container-executor`. For the `/usr/lib` mount point,

setuid should not be a problem. However, the parcel could easily be located on a different mount point. If you are using a parcel, make sure the mount point for the parcel directory is without the nosuid option.

The `container-executor` program must have a very specific set of permissions and ownership to function correctly. In particular, it must:

1. Be owned by root
2. Be owned by a group that contains only the user running the YARN daemons
3. Be setuid
4. Be group readable and executable. This corresponds to the ownership `root:yarn` and the permissions `6050`.

```
---Sr-s--- 1 root yarn 91886 2012-04-01 19:54 container-executor
```

Important: Configuration changes to the Linux container executor could result in local NodeManager directories (such as `usercache`) being left with incorrect permissions. To avoid this, when making changes using either Cloudera Manager or the command line, first manually remove the existing NodeManager local directories from all configured local directories (`yarn.nodemanager.local-dirs`), and let the NodeManager recreate the directory structure.

Troubleshooting

When you set up a secure cluster for the first time and debug problems with it, the `task-controller` or `container-executor` may encounter errors. These programs communicate these errors to the TaskTracker or NodeManager daemon via numeric error codes that appear in the TaskTracker or NodeManager logs respectively (`/var/log/hadoop-mapreduce` or `/var/log/hadoop-yarn`). The following sections list the possible numeric error codes with descriptions of what they mean:

- TaskController Error Codes (MRv1) on page 432
- ContainerExecutor Error Codes (YARN) on page 434

TaskController Error Codes (MRv1)

This table shows some of the error codes and messages generated by TaskController in MapReduce (MRv1).

Numeric Code	Name	Description
1	INVALID_ARGUMENT_NUMBER	• Incorrect number of arguments provided for the given task-controller command • Failure to initialize the job localizer
2	INVALID_USER_NAME	The user passed to the task-controller does not exist.
3	INVALID_COMMAND_PROVIDED	The task-controller does not recognize the command it was asked to execute.
4	SUPER_USER_NOT_ALLOWED_TO_RUN_TASKS	The user passed to the task-controller was the super user.
5	INVALID_TT_ROOT	The passed TaskTracker root does not match the configured TaskTracker root (`mapred.local.dir`), or does not exist.
6	SETUID_OPER_FAILED	Either could not read the local groups database, or could not set UID or GID

Numeric Code	Name	Description
7	UNABLE_TO_EXECUTE_TASK_SCRIPT	The task-controller could not execute the task launcher script.
8	UNABLE_TO_KILL_TASK	The task-controller could not kill the task it was passed.
9	INVALID_TASK_PID	The PID passed to the task-controller was negative or 0.
10	ERROR_RESOLVING_FILE_PATH	The task-controller could not resolve the path of the task launcher script file.
11	RELATIVE_PATH_COMPONENTS_IN_FILE_PATH	The path to the task launcher script file contains relative components (for example, "...").
12	UNABLE_TO_STAT_FILE	The task-controller did not have permission to stat a file it needed to check the ownership of.
13	FILE_NOT_OWNED_BY_TASKTRACKER	A file which the task-controller must change the ownership of has the wrong the ownership.
14	PREPARE_ATTEMPT_DIRECTORIES_FAILED	The `mapred.local.dir` is not configured, could not be read by the task-controller, or could not have its ownership secured.
15	INITIALIZE_JOB_FAILED	The task-controller could not get, stat, or secure the job directory or job working working directory.
16	PREPARE_TASK_LOGS_FAILED	The task-controller could not find or could not change the ownership of the task log directory to the passed user.
17	INVALID_TT_LOG_DIR	The `hadoop.log.dir` is not configured.
18	OUT_OF_MEMORY	The task-controller could not determine the job directory path or the task launcher script path.
19	INITIALIZE_DISTCACHEFILE_FAILED	Could not get a unique value for, stat, or the local distributed cache directory.
20	INITIALIZE_USER_FAILED	Could not get, stat, or secure the per-user task tracker directory.
21	UNABLE_TO_BUILD_PATH	The task-controller could not concatenate two paths, most likely because it ran out of memory.
22	INVALID_TASKCONTROLLER_PERMISSIONS	The task-controller binary does not have the correct permissions set. See Information about Other Hadoop Security Programs.
23	PREPARE_JOB_LOGS_FAILED	The task-controller could not find or could not change the ownership of the job log directory to the passed user.
24	INVALID_CONFIG_FILE	The taskcontroller.cfg file is missing, malformed, or has incorrect permissions.

Numeric Code	Name	Description
255	Unknown Error	Several possible causes for this error, including: • User accounts on the cluster with an ID less than the value specified for the `min.user.id` property in the `taskcontroller.cfg` file. Default value of 1000 is good for Ubuntu but may not be valid for other OSs. To set the `min.user.id` in the `taskcontroller.cfg` file, see Configure Secure MapReduce. • Jobs do not run and the TaskTracker is unable to create a Hadoop logs directory (see Troubleshooting Error Messages). • May result from previous errors. Check older entries in the log file for possibilities.

ContainerExecutor Error Codes (YARN)

The codes in the table apply to the container-executor in YARN, but are used by the LinuxContainerExecutor only.

Numeric Code	Name	Description
1	INVALID_ARGUMENT_NUMBER	• Incorrect number of arguments provided for the given container-executor command • Failure to initialize the container localizer
2	INVALID_USER_NAME	The user passed to the container-executor does not exist.
3	INVALID_COMMAND_PROVIDED	The container-executor does not recognize the command it was asked to run.
5	INVALID_NM_ROOT	The passed NodeManager root does not match the configured NodeManager root (`yarn.nodemanager.local-dirs`), or does not exist.
6	SETUID_OPER_FAILED	Either could not read the local groups database, or could not set UID or GID
7	UNABLE_TO_EXECUTE_CONTAINER_SCRIPT	The container-executor could not run the container launcher script.
8	UNABLE_TO_SIGNAL_CONTAINER	The container-executor could not signal the container it was passed.
9	INVALID_CONTAINER_PID	The PID passed to the container-executor was negative or 0.
18	OUT_OF_MEMORY	The container-executor couldn't allocate enough memory while reading the container-executor.cfg file, or while getting the paths for the container launcher script or credentials files.

Numeric Code	Name	Description
20	INITIALIZE_USER_FAILED	Couldn't get, stat, or secure the per-user NodeManager directory.
21	UNABLE_TO_BUILD_PATH	The container-executor couldn't concatenate two paths, most likely because it ran out of memory.
22	INVALID_CONTAINER_EXEC_PERMISSIONS	The container-executor binary does not have the correct permissions set. See Information about Other Hadoop Security Programs.
24	INVALID_CONFIG_FILE	The container-executor.cfg file is missing, malformed, or has incorrect permissions.
25	SETSID_OPER_FAILED	Could not set the session ID of the forked container.
26	WRITE_PIDFILE_FAILED	Failed to write the value of the PID of the launched container to the PID file of the container.
255	Unknown Error	This error has several possible causes. Some common causes are: • User accounts on your cluster have a user ID less than the value specified for the `min.user.id` property in the `container-executor.cfg` file. The default value is 1000 which is appropriate on Ubuntu systems, but may not be valid for your operating system. For information about setting `min.user.id` in the `container-executor.cfg` file, see this step. • This error is often caused by previous errors; look earlier in the log file for possible causes.

The following exit status codes apply to all containers in YARN. These exit status codes are part of the YARN framework and are in addition to application specific exit codes that can be set:

Numeric Code	Name	Description
0	SUCCESS	Container has finished succesfully.
-1000	INVALID	Initial value of the container exit code. A container that does not have a COMPLETED state will always return this status.
-100	ABORTED	Containers killed by the framework, either due to being released by the application or being 'lost' due to node failures, for example.
-101	DISKS_FAILED	Container exited due to local disks issues in the NodeManager node. This occurs when the number of good nodemanager-local-directories or nodemanager-log-directories drops below the health threshold.

Numeric Code	Name	Description
-102	PREEMPTED	Containers preempted by the framework. This does not count towards a container failure in most applications.
-103	KILLED_EXCEEDED_VMEM	Container terminated because of exceeding allocated virtual memory limit.
-104	KILLED_EXCEEDED_PMEM	Container terminated because of exceeding allocated physical memory limit.
-105	KILLED_BY_APPMASTER	Container was terminated on request of the application master.
-106	KILLED_BY_RESOURCEMANAGER	Container was terminated by the resource manager.
-107	KILLED_AFTER_APP_COMPLETION	Container was terminated after the application finished.

Appendix: Apache License, Version 2.0

SPDX short identifier: Apache-2.0

Apache License
Version 2.0, January 2004
http://www.apache.org/licenses/

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License.

Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.

3. Grant of Patent License.

Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims

licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.

4. Redistribution.

You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:

1. You must give any other recipients of the Work or Derivative Works a copy of this License; and
2. You must cause any modified files to carry prominent notices stating that You changed the files; and
3. You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and
4. If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. Submission of Contributions.

Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.

6. Trademarks.

This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.

7. Disclaimer of Warranty.

Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.

8. Limitation of Liability.

In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.

9. Accepting Warranty or Additional Liability.

While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

APPENDIX: How to apply the Apache License to your work

To apply the Apache License to your work, attach the following boilerplate notice, with the fields enclosed by brackets "[]" replaced with your own identifying information. (Don't include the brackets!) The text should be enclosed in the appropriate comment syntax for the file format. We also recommend that a file or class name and description of purpose be included on the same "printed page" as the copyright notice for easier identification within third-party archives.

```
Copyright [yyyy] [name of copyright owner]

Licensed under the Apache License, Version 2.0 (the "License");
you may not use this file except in compliance with the License.
You may obtain a copy of the License at

  http://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software
distributed under the License is distributed on an "AS IS" BASIS,
WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
See the License for the specific language governing permissions and
limitations under the License.
```