

Hortonworks Data Platform

Using WebHDFS REST API

(Jul 2, 2014)

Hortonworks Data Platform : Using WebHDFS REST API

Copyright © 2012-2014 Hortonworks, Inc. Some rights reserved.

The Hortonworks Data Platform, powered by Apache Hadoop, is a massively scalable and 100% open source platform for storing, processing and analyzing large volumes of data. It is designed to deal with data from many sources and formats in a very quick, easy and cost-effective manner. The Hortonworks Data Platform consists of the essential set of Apache Hadoop projects including MapReduce, Hadoop Distributed File System (HDFS), HCatalog, Pig, Hive, HBase, Zookeeper and Ambari. Hortonworks is the major contributor of code and patches to many of these projects. These projects have been integrated and tested as part of the Hortonworks Data Platform release process and installation and configuration tools have also been included.

Unlike other providers of platforms built using Apache Hadoop, Hortonworks contributes 100% of our code back to the Apache Software Foundation. The Hortonworks Data Platform is Apache-licensed and completely open source. We sell only expert technical support, [training](#) and partner-enablement services. All of our technology is, and will remain free and open source.

Please visit the [Hortonworks Data Platform](#) page for more information on Hortonworks technology. For more information on Hortonworks services, please visit either the [Support](#) or [Training](#) page. Feel free to [Contact Us](#) directly to discuss your specific needs.



Except where otherwise noted, this document is licensed under
Creative Commons Attribution ShareAlike 3.0 License.
<http://creativecommons.org/licenses/by-sa/3.0/legalcode>

Table of Contents

1. Using WebHDFS REST API	1
1.1. WebHDFS User Guide	1
1.1.1. Authentication	2
1.2. WebHDFS Resources	3

1. Using WebHDFS REST API

Apache Hadoop provides native libraries for accessing HDFS. However, users prefer to use HDFS remotely over the heavy client side native libraries. For example, some applications need to load data in and out of the cluster, or to externally interact with the HDFS data. WebHDFS addresses these issues by providing a fully functional HTTP REST API to access HDFS.

WebHDFS provides the following features:

- Provides read and write access. Supports all HDFS operations (like granting permissions, configuring replication factor, accessing block location, etc.).
- Supports all HDFS parameters with defaults.
- Permits clients to access Hadoop from multiple languages without actually installing Hadoop. You can also use common tools like curl/wget to access HDFS.
- Uses the full bandwidth of the Hadoop cluster for streaming data: The file read and file write calls are redirected to the corresponding datanodes.
- Uses Kerberos (SPNEGO) and Hadoop delegation tokens for authentication.
- WebHDFS is completely Apache open source. Hortonworks contributed the code to Apache Hadoop as a first class built-in Hadoop component.
- Requires no additional servers. However, a proxy WebHDFS (for example: Httpfs is useful in certain cases and is complementary to WebHDFS).

For more information, see: [WebHDFS – HTTP REST Access to HDFS](#).

In this section:

- [User Guide](#)
- [Resources](#)

For information on WebHDFS Administration guide, see: [WebHDFS Administrator Guide](#).

1.1. WebHDFS User Guide

The following examples use the `curl` command tool to access HDFS via WebHDFS REST API.

- To read a file (for example: `/foo/bar`):

```
curl -i -L "http://$<Host_Name>:$<Port>/webhdfs/v1/foo/bar?op=OPEN"
```

- To list a directory (for example: `/foo`):

```
curl -i "http://$<Host_Name>:$<Port>/webhdfs/v1/foo/?op=LISTSTATUS"
```

- To list the status of a file (for example: `/foo/bar`) or a directory:

```
curl -i "http://$<Host_Name>:$<Port>/webhdfs/v1/foo/bar?op=GETFILESTATUS"
```

- To write a file into a `/foo/new` file:

```
curl -i -X PUT -L "http://$<Host_Name>:$<Port>/webhdfs/v1/foo/newFile?op=CREATE" -T newFile
```

- To rename the `/foo/bar` file to `/foo/bar3`:

```
curl -i -X PUT "http://$<Host_Name>:$<Port>/webhdfs/v1/foo/bar?op=RENAME&destination=/foo/bar2"
```

- Make new directory `/foo2`:

```
curl -i -X PUT "http://$<Host_Name>:$<Port>/webhdfs/v1/foo2?op=MKDIRS&permission=711"
```

1.1.1. Authentication

When security is enabled, authentication is performed by either Hadoop delegation token or Kerberos SPNEGO. If a token is set in the delegation query parameter, the authenticated user is the user encoded in the token. If the delegation parameter is not set, the user is authenticated by Kerberos SPNEGO.

Below are examples using the `curl` command tool.

- Login to the Key Distribution Center (KDC).

```
kinit
```

- Provide any arbitrary user name and a null password.
- Execute the following commands:

```
curl --negotiate -u:anyUser "http://$<Host_Name>:$<Port>/webhdfs/v1/foo/bar?op=OPEN"
curl --negotiate -u:anyUser -b ~/cookies.txt -c ~/cookies.txt http://$<Host_Name>:$<Port>/webhdfs/v1/foo/bar?op=OPEN
```

where:

- `--negotiate` option enables SPNEGO in `curl`.

- `-u:anyUser` option is mandatory when the user name is not specified instead, a Kerberos established user (via kinit) is used. (Ensure that you provide any user name and enter a null password when prompted.)
- `-b` and `-c` options are used for storing and sending HTTP cookies.

1.2. WebHDFS Resources

The HTTP REST API supports the complete FileSystem interface for HDFS. For more information, see the following sections in the [WebHDFS REST API](#) documentation:

- Introduction
 - [Operations](#)
 - [FileSystem URIs vs HTTP URLs](#)
 - [HDFS Configuration Options](#)
- [Authentication](#)
- [Proxy Users](#)
- [File and Directory Operations](#)
- [Other File System Operations](#)
- [Delegation Token Operations](#)
- [Error Responses](#)
- [JSON Schemas](#)
- [HTTP Query Parameter Dictionary](#)