

# Dell EMC PowerScale and Cloudera Replication Manager QATS Certification Report and Process

January 2022

H19036

## White Paper

### Abstract

This document provides the test summary report of the Cloudera Quality Assurance Test Suite (QATS) certification of Cloudera CDH 6.3.4 to Cloudera CDP 7.1.7 cluster replication using Cloudera Replication Manager on Dell EMC PowerScale powered by Dell EMC PowerScale OneFS 8.2.2.

Dell Technologies

**CLOUDERA**

## Copyright

The information in this publication is provided as is. Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2022 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Intel, the Intel logo, the Intel Inside logo and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. Other trademarks may be trademarks of their respective owners. Published in the USA January 2022 H19036.

Dell Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

# Contents

- Executive summary.....4**
- QATS test environment.....4**
- Cloudera Replication Manager features certified on Dell EMC PowerScale .....5**
- Setup, issues, and workarounds.....6**
- Technical support and resources .....9**
- Appendix: HDFS replication policy reference .....10**

## Executive summary

### Overview

This document provides details for functional testing of Cloudera Replication Manager to replicate Cloudera CDH 6.3.4 to Cloudera Data Platform (CDP) Private Cloud Base 7.1.7 with Dell EMC PowerScale through the HDFS protocol as a primary file system. Dell Technologies and Cloudera field and sales teams can use the information when engaging with customers who want to use Replication Manager to replicate Hadoop clusters running on PowerScale with HDFS as a primary file system.

### Revisions

Date	Description
January 2022	Initial release

### We value your feedback

Dell Technologies and the authors of this document welcome your feedback on this document. Contact the Dell Technologies team by [email](#).

**Authors:** Kirankumar Bhusanurmath, Analytics Solutions Architect, Dell Technologies; Nishan Shetty, Software Engineer, Cloudera

**Note:** For links to other documentation for this topic, see the [PowerScale Info Hub](#).

## QATS test environment

The following table details the QATS test environment:

**Table 1. QATS test environment**

Component	Version/resource
Source cluster: Cloudera CDH	Cloudera Manager 6.3.4: See download information <a href="#">here</a>
	Cloudera Runtime 6.3.4: See download information <a href="#">here</a>
CDP Private Cloud Base version	Cloudera Manager 7.4.4-17486246: See download information <a href="#">here</a>
	Cloudera Runtime 7.1.7: See download information <a href="#">here</a>
Cloudera Replication Manager	CDP 7.1.7 Replication Manager: See more information <a href="#">here</a>
Dell EMC PowerScale CSD	PowerScale-1.1.0.jar
Dell EMC Isilon node	A200-4U-Single-16GB-1x1GE-2X10GE SFP+-30TB-400GB SSD
OneFS operating system	Isilon OneFS v8.2.2.0 B_8_2_2_007(RELEASE)
Compute and operating system	Virtual machines
	CentOS Linux release 7.7.1908 (Core)
	Memory: 64 GB
	CPUs: 32

---

**Note:** CDP Ranger had an issue described at [RANGER-2626](#). A OneFS custom patch was developed to fix this Ranger issue for the QATS testing, and the same fix is incorporated in the “[PSP-1091]” patch made available as a Roll Up Patch (RUP) in May 2021.

---

PowerScale Custom Service Descriptors (CSD) are built for Dell EMC PowerScale. You can download CSD from the [product download](#) page and extract the contents on to the Cloudera Manager server.

The Cloudera full release documentation, including the [Cloudera Manager Release Notes](#) and [Cloudera Runtime Release Notes](#), is available [here](#).

## Cloudera Replication Manager features certified on Dell EMC PowerScale

**Introduction** This section discusses the Cloudera Replication Manager features that are supported and those that are not supported on PowerScale OneFS 8.2.2.

### Certified and supported Replication Manager features

The following Cloudera Replication Manager features successfully completed all the functional test cases of the Cloudera QATS Appliance. This feature-specific test success determines the support for HDFS replication of the components on PowerScale OneFS 8.2.2.

#### HBase snapshots

HBase snapshots allow you to create point-in-time backups of tables without making data copies and with minimal impact on RegionServers. For detailed information, see [Snapshots](#) under Replication Manager.

#### HDFS replication

HDFS replication enables you to copy (replicate) your HDFS data from one HDFS service to another, synchronizing the data set on the destination service with the data set on the source service, based on a specified replication policy. The destination service must be managed by the Cloudera Manager Server where the replication is being set up, and the source service can be managed by that same server or by a peer Cloudera Manager Server. You can also replicate HDFS data within a cluster by specifying different source and destination directories.

Remote Replication Manager automatically copies HDFS metadata to the destination cluster as it copies files. HDFS metadata need only be backed up locally. For detailed information, see [HDFS Replication](#) under Replication Manager.

#### Hive/Impala replication

Hive/Impala replication enables you to copy (replicate) your Hive metastore and data from one cluster to another and synchronize the Hive metastore and data set on the destination cluster with the source, based on a specified replication policy. For detailed information, see [Hive/Impala replication](#) under Replication Manager.

## Sentry to Ranger replication

When you create or edit a Hive replication policy, you can choose to migrate the Sentry policies for Hive objects, Impala data, and URLs that are being replicated. The Replication Manager converts the Sentry policies to Ranger policies for the migrated data in the target cluster. The minimum supported Cloudera Manager version 6.3.1 and above is required to replicate Sentry policies to Ranger. For detailed information, see [Sentry to Ranger replication for Hive replication policies](#).

## Unsupported Replication Manager features

Any Replication Manager features not listed in [Certified and supported Replication Manager features](#) are to be considered as not supported features of the Cloudera Replication Manager for a Cloudera cluster deployed on PowerScale OneFS 8.2.2.

## Setup, issues, and workarounds

### Environment setup

#### Cloudera CDP 7.1.7 installation on PowerScale OneFS 8.2.2

The following steps outline the standard process that is used for setting up the Cloudera Data Platform Private Cloud Base 7.1.7, including any service-specific workaround employed during the Cloudera QATS testing:

1. Set up the CDP 7.1.7 cluster using the instructions in the [CDP Private Cloud Base Installation Guide](#).
2. Set up CDP on PowerScale. Dell Technologies used the following installation guide, which can be downloaded from the [product download](#) page: *Dell EMC PowerScale OneFS 8.2.2 and CDP Private Base 7.1.6 Install Guide*

#### Replication Manager setup on PowerScale

The following steps outline the process used for setting up Replication Manager on PowerScale:

1. Source cluster: Cloudera CDH 6.3.4 cluster health check  
Ensure that the source cluster is healthy and that any errors or warnings are addressed before proceeding with the cluster replication process.
2. Source cluster: Cloudera Manager upgrade  
Upgrade the source Cloudera Manager to the target cluster Cloudera Manager. We recommend upgrading it to 7.4.4 (7.4.4-17486246). For detailed information, see [Upgrading Cloudera Manager 6](#).
3. Target cluster: Cloudera CDP 7.1.7 cluster health check  
Ensure that the target cluster is healthy and that any errors or warnings are addressed before proceeding with the cluster replication process.
4. Enabling replication between clusters with Kerberos authentication  
To enable replication between clusters, additional setup steps are required to ensure that the source and destination clusters can communicate. For detailed information, see [Enabling replication between clusters with Kerberos Authentication](#).

## 5. Configuring replication with Kerberos and PowerScale

If you plan to use replication between clusters that have enabled Kerberos, then you need to create a Custom Kerberos Keytab and principal for replication jobs to use to authenticate. For detailed information, see [Configuring Replication with Kerberos and Dell EMC PowerScale](#).

## 6. Target cluster: Adding “Peers” in the CDP cluster and creating replication policies

### Peers

#### Connect Multiple Instances of Cloudera Manager

For HDFS or Hive replication, add as a peer the Cloudera Manager Server that should be the **source** of replicated data. Data from the peer cluster can then be replicated to an HDFS or Hive service managed by the Cloudera Manager Server you are currently signed into.

Add Peer

Peer Name	Peer URL	Status
SRC	<a href="http://ctr-e168-1619641223258-13464-01-000007.hwx.site:7180/">http://ctr-e168-1619641223258-13464-01-000007.hwx.site:7180/</a>	✓ Connected

1 - 1 of 1

**NOTE:** Ensure that the Cloudera Manager version of the source cluster is 7.4.4 (7.4.4-17486246) and not 6.3.x.

## Issues and workaround

Note the following known issues and workarounds:

- Hive Replication issue:** When copying (distCp), the data is copied to the target cluster, but the job fails to get the job status at the **Hive Table Data Replication** step.

<p>▼ Hive Table Data Replication</p> <p>HDFS replication command failed.</p>	<p>DellEMC PowerScale</p>	<p>Aug 10, 12:57:06 PM</p>	<p>47.27s</p>
<p>▼ Trigger a HDFS replication job on one of the available HDFS roles.</p> <p>HDFS replication command failed.</p>	<p>cdp7rm3-abjain-5.cdp7rm3-abjain.root.hwx.site</p>	<p>Aug 10, 12:57:06 PM</p>	<p>47.27s</p>

```
S> dr/distcp.sh ["-bandwidth","100","-i","-m","20","-prbugpa","-skipAclErr","-update","-proxyuser","bdr@ROOT.HWX.SITE","-log","/user/PROXY_USER_PLACEHOLDER/.cm/distcp/2021-08-10_702","-stagingDirPath","/user/PROXY_USER_PLACEHOLDER/.cm/distcp-job-staging/2021-08-10_702","-sourceconf","source-client-conf","-sourceprincipal","bdr@ROOT.HWX.SITE","-sourcektcache","source.tgt","-rebase","-skipOwnershipCheck","-skipDefaultJobStagingDir","-diff","-replaceNameservice","-strategy","dynamic","-filters","exclusion-filter.list","-scheduleId","10","-scheduleName","hive-sample1","/tmp/extdata1","warehouse/tablespace/external/hive"]
```

stdout stderr

```
21/08/10 12:57:16 INFO mapreduce.Job: The url to track the job: http://cdp7rm3-abjain-1.cdp7rm3-abjain.root.hwx.site:8088/proxy/application_1628599913890_0002/
21/08/10 12:57:16 INFO distcp.DistCp: DistCp job-id: job_1628599913890_0002
21/08/10 12:57:16 INFO distcp.DistCp: map 0% reduce 0% files 0% bytes 0% throughput 0.0 (KB/s) running mappers 0
21/08/10 12:57:36 INFO distcp.DistCp: map 100% reduce 100% files 100% bytes 100% throughput 32.7 (KB/s) running mappers 0
21/08/10 12:57:47 INFO mapred.ClientServiceDelegate: Application state is completed. FinalApplicationStatus=SUCCEEDED. Redirecting to job history server
21/08/10 12:57:47 ERROR util.DistCpUtil: Exception encountered
java.io.IOException: Job status not available
    at org.apache.hadoop.mapreduce.Job.updateStatus(Job.java:341)
    at org.apache.hadoop.mapreduce.Job.isComplete(Job.java:614)
    at com.cloudera.enterprise.distcp.DistCp.checkProgress(DistCp.java:475)
    at com.cloudera.enterprise.distcp.DistCp.execute(DistCp.java:465)
    at com.cloudera.enterprise.distcp.DistCp$1.run(DistCp.java:152)
    at com.cloudera.enterprise.distcp.DistCp$1.run(DistCp.java:149)
    at java.security.AccessController.doPrivileged(Native Method)
    at javax.security.auth.Subject.doAs(Subject.java:422)
    at org.apache.hadoop.security.UserGroupInformation.doAs(UserGroupInformation.java:1898)
    at com.cloudera.enterprise.distcp.DistCp.run(DistCp.java:149)
    at org.apache.hadoop.util.ToolRunner.run(ToolRunner.java:76)
    at com.cloudera.enterprise.distcp.DistCp.main(DistCp.java:847)
    at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)
    at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62)
    at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43)
    at java.lang.reflect.Method.invoke(Method.java:498)
    at org.apache.hadoop.util.RunJar.run(RunJar.java:318)
    at org.apache.hadoop.util.RunJar.main(RunJar.java:232)
21/08/10 12:57:47 WARN distcp.DistCp: Killing submitted job job_1628599913890_0002
```

**Solution:** CM > Yarn > Configuration > HDFS Replication Advanced Configuration Snippet (Safety Valve) for yarn-site.xml

```
<property><name>mapreduce.jobhistory.intermediate-done-dir</name><value>/user/history/done_intermediate</value></property>
```

- **Hive Replication issue:** Replication fails at the last step, **Rename Snapshots on Target**, with `NullPointerException`.

### Hive Replication Command

Status ❌ Failed Context [Hive](#) 📅 Aug 11, 10:30:25 AM ⌚ 73.91s [Download](#)

java.lang.NullPointerException

▼ **Completed 5 of 5 step(s).**

Show All Steps    Show Only Failed Steps    Show Only Running Steps

✔	Export remote Hive Metastore
> ✔	Transfer Metadata Files
> ✔	Hive Table Data Replication
✔	Import Hive Metastore
❌	Rename Snapshots on Target Failed due to java.lang.NullPointerException:

**Solution:** Set `distcp.copy.diff` to `false` in `hdfs_client_config_safety_valve`

- **Impersonation error:** To avoid an impersonation error, add `root`, `bdr` (Run As User) to the `hive proxyusers` in `PowerScale`.

```
isi hdfs proxyusers modify hive --zone=System --add-user=root
isi hdfs proxyusers view hive --zone=System
```

- **Error in getting credentials for local/remote HDFS service**

**Solution:** Add the group name of the user (for example, `bdr`) who will be used as Run As Username of the replication schedule as “Authorized Admin Group” in target cluster HDFS/PowerScale service's configuration in CM.

- **Change permission not working after distcp job is run**

**Solution:** Add a user mapping rule in `PowerScale`. For example:



Apply order	Rule type	Description	On rule match
1	Replace	Replace user "hdfs" with "root".	Continue
2	Replace	Replace user "bdr" with "root".	Continue

- **No option to change default “Log Expiration” for BDR**

**Workaround:** By default (if the Log Path is not provided under the replication policy in CM), the logs are kept under “/user/<username>/.cm/distcp”; otherwise, “<userProvidedLogPath>/.cm/distcp” is used to keep the logs for distcp jobs. Currently, the user cannot change the default log expiration time (internal Cloudera Jira). Thus, if the space is the constraint, these logs have to be cleaned up manually.

## Technical support and resources

[Dell.com/support](https://dell.com/support) is focused on meeting customer needs with proven services and support.

[Storage technical documents and videos](#) provide expertise that helps to ensure customer success on Dell EMC storage platforms.

# Appendix: HDFS replication policy reference

Edit HDFS Replication Policy ✕

**General** Resources Advanced

Name:

Source:  Use [AWS Credentials](#) or [Azure Credentials](#) to add more cloud replication destinations. Use [Add Peers](#) to add more replication sources.

Source Path:

Destination:  Use [AWS Credentials](#) or [Azure Credentials](#) to add more cloud replication destinations.

Destination Path:

Schedule:

Start Time:  UTC

Run As Username:  ⚠ The user `bdr@IPA.ENG.HORTONWORKS.COM` must be a superuser because **Preserve Permissions** is enabled.

Run on Peer as Username:

[Cancel](#) [Save Policy](#)

Edit HDFS Replication Policy ✕

General Resources **Advanced**

Path Exclusion ⓘ [Add Exclusion](#)

MapReduce Service:

Log Path:

Description:

Error Handling

- Skip Checksum Checks
- Skip Listing Checksum Checks
- Abort on Error
- Abort on Snapshot Diff Failures

Preserve

- Block Size
- Replication Count
- Permissions
- Extended Attributes

Delete Policy

- Keep Deleted Files
- Delete to Trash

[Cancel](#) [Save Policy](#)

Replication Policies

Replication History

Name **repltest** Type **HDFS** Source **isilon (Cluster 1 @ SRC)** Destination **PowerScale (Cluster1)** Next Run **None scheduled.**

Start Time	Duration	Outcome	Files Expected	Files Copied	Files Failed	Files Deleted	Files Skipped
> May 31, 2021 6:04 PM	1 min	Successful	2 (80 B)	2 (80 B)	0 (0 B)	0 (0 B)	0 (0 B)
> May 31, 2021 5:56 PM	1 min	Successful	2 (80 B)	0 (0 B)	0 (0 B)	0 (0 B)	2 (80 B)
> May 31, 2021 5:54 PM	0 min	Failed Dry Run	0 (0 B)	0 (0 B)	0 (0 B)	0 (0 B)	0 (0 B)
> May 31, 2021 5:52 PM	1 min	Successful	2 (80 B)	0 (0 B)	0 (0 B)	0 (0 B)	2 (80 B)
> May 31, 2021 5:36 PM	1 min	Successful	2 (80 B)	0 (0 B)	0 (0 B)	0 (0 B)	2 (80 B)
> May 31, 2021 5:32 PM	1 min	Failed	2 (80 B)	0 (0 B)	2 (80 B)	0 (0 B)	2 (80 B)
> May 26, 2021 12:23 PM	1 min	Failed	2 (80 B)	2 (80 B)	2 (80 B)	0 (0 B)	0 (0 B)
> May 26, 2021 12:20 PM	0 min	Failed	0 (0 B)	0 (0 B)	0 (0 B)	0 (0 B)	0 (0 B)

Edit Snapshot Policy



Service:

Name:

Description:

Tables:

Schedule:  Hourly  
 Take snapshots every hour at  minute(s)  
 Hourly snapshots to keep:

Daily  
 Weekly  
 Monthly  
 Yearly

Alerts:  On failure

Snapshot Policies

Snapshots History

Policy Name **snaptest1** Type **HBASE** Cluster **Cluster1** Service **HBase** Snapshot Prefix **cm-auto-21e161f6-7d73-401a-bbdb-36a0fa197b9c**

Start Time	Outcome	Tables Processed	Tables Unprocessed	Snapshots Created	Snapshots Deleted	Errors During Creation	Errors During Deletion
> May 31, 2021 1:05 PM	Successful	1	0	1	1	0	0
> May 31, 2021 12:05 PM	Successful	1	0	1	1	0	0
> May 31, 2021 11:05 AM	Successful	1	0	1	1	0	0
> May 31, 2021 10:05 AM	Successful	1	0	1	1	0	0
> May 31, 2021 9:05 AM	Successful	1	0	1	1	0	0
> May 31, 2021 8:05 AM	Successful	1	0	1	1	0	0
> May 31, 2021 7:05 AM	Successful	1	0	1	1	0	0
> May 31, 2021 6:05 AM	Successful	1	0	1	1	0	0