

WANdisco Fusion on Oracle Big Data Cloud Service

ORACLE WHITE PAPER | JULY 2017

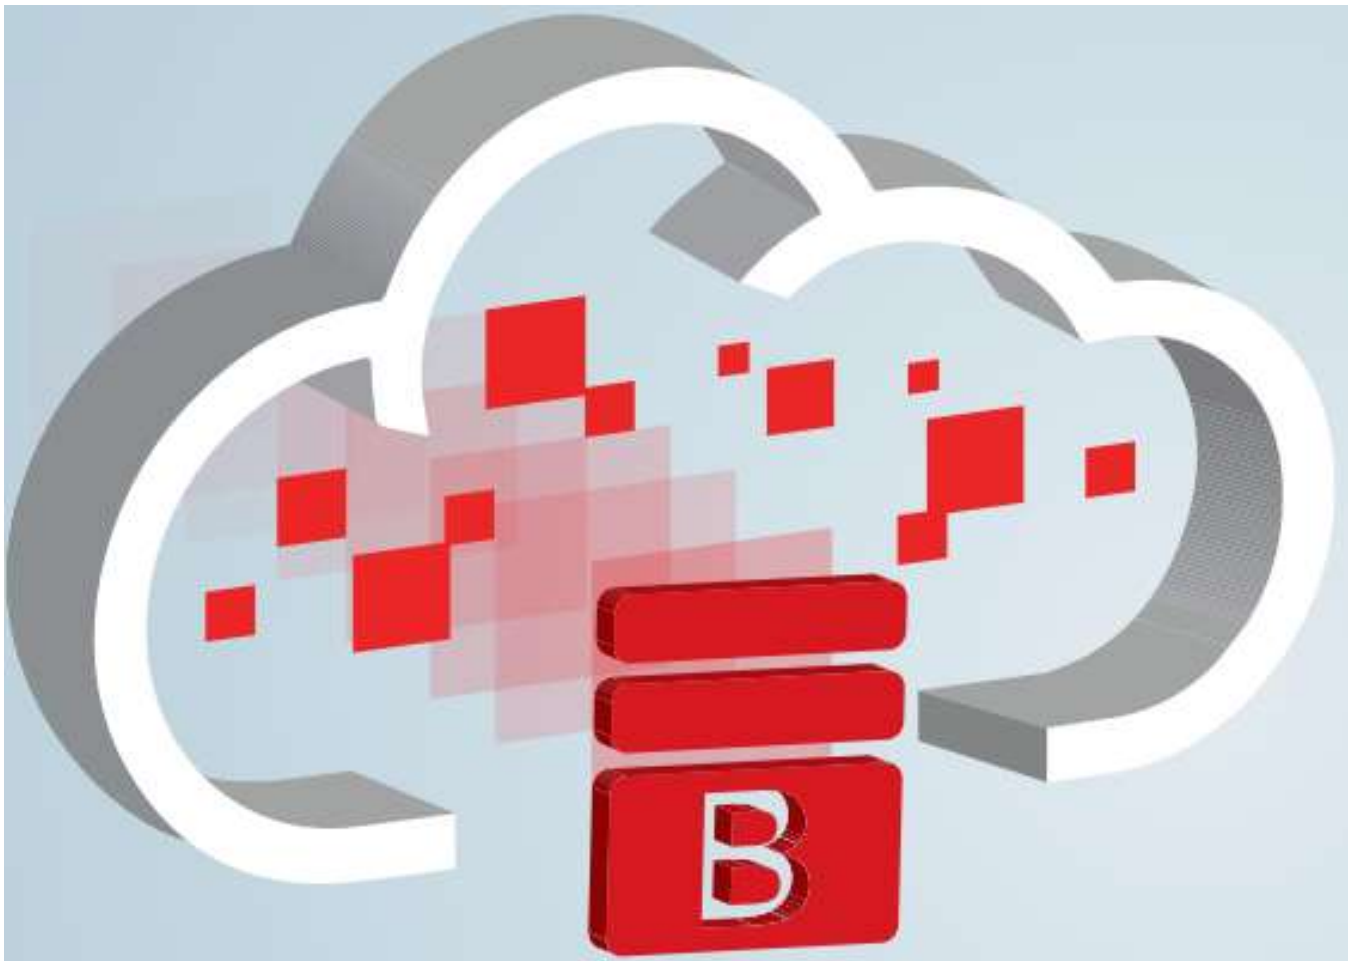


Table of Contents

What's out of the box from Oracle Big Data Cloud Services	1
WANdisco Fusion on Oracle Big Data Cloud Services.....	4
WANdisco Fusion – BDCS Deployment Considerations	7
WANdisco Installation Guidelines	7
Configuration used for the Validation Test	23
Summary	24

What's out of the box from Oracle Big Data Cloud Services

Oracle Big Data Cloud Service is a Platform as a Service (PaaS) offering from Oracle. Although the service information is detailed here, the service offering is subject to change. More up to date information can be found at [Oracle BDCS Features](#) and [Oracle BDCS Hardware and Pricing](#).

Oracle Big Data Cloud Service is a pre-configured Big Data environment which includes a complete installation of the Cloudera Distribution including Hadoop and Spark.

When you subscribe to the Oracle Big Data Cloud Service, you can select between 3 and 60 dedicated nodes. You can scale up or down the number of nodes in increments of one node at a time.

Oracle manages the complete hardware and networking infrastructure as well as the initial setup while you have complete administrative control of the software. This includes root access to the servers. The H/W and software configurations are as follows. Note that the specification can change in the future.

Each node is a DEDICATED virtualized environment with the following hardware configuration:

- » 32 OCPUs (Oracle Compute Unit: CPU capacity equivalent of one physical core of an Intel Xeon processor with hyper threading enabled. Each OCPU corresponds to two hardware execution threads, known as vCPUs)
- » 256GB Memory
- » 48TB of pre-configured HDFS storage
- » 40 Gb/s InfiniBand between each node.
The advantage of dedicated H/W is that you will not have noisy neighbors and will get predictable performance for the workload.
- » Integrated Software
- » Oracle Linux Operating System 6 update 8
- » Cloudera Distribution including Apache Hadoop (CDH 5.x)
- » File System: HDFS
- » MapReduce Engine: The MapReduce engine provides a platform for massively parallel execution of algorithms written in Java. Oracle Big Data Cloud Service runs YARN by default
- » Administrative Framework: Cloudera Manager is a comprehensive administrative tool for CDH
- » Apache Projects: CDH includes Apache projects for MapReduce and HDFS such as Hive, Pig, Oozie, Zookeeper, HBase, Sqoop and Spark to name a few
- » Cloudera Applications: Oracle Big Data Cloud Service installs all products included in the Cloudera Enterprise Data Hub Edition including Impala, Search and Navigator
- » Two Command-line utilities for managing data and resources (These utilities are for use only with Big Data Cloud Service and are not included with the on-premises engineered system - Oracle Big Data Appliance); bda-oss-admin to configure command shell access to Hadoop cluster and odcp to copy data between HDFS and Storage Cloud.
- » Single Command patching and upgrade: The advantage of cloud solutions is that the upgrades and patching procedures are thoroughly tested and implemented to minimize issues. The patching and upgrades can be performed in a rolling fashion or all at once with a down time. Oracle uses Mammoth command line utility to upgrade or install.
- » Oracle Big Data Connectors, which facilitate access to data stored in Apache Hadoop Cluster; Oracle SQL Connector for HDFS, Oracle Loader for Hadoop, Oracle XQuery for Hadoop, Oracle R Advanced Analytics for Hadoop and Oracle Data Integrator Enterprise Edition
- » Oracle Big Data Spatial and Graph, which provides spatial and graph analytic capabilities to supported Apache Hadoop and Oracle NoSQL Database Big Data Platforms
- » If you have a subscription to Oracle Storage Cloud Service, you can integrate the storage with Oracle Big Data Cloud service to stage data to load into the HDFS cluster. Additional tools like odcp helps copy data between HDFS and an Oracle Storage Cloud Service Container efficiently.
- » MySql in active passive mode for Hive metastore.
- » Security: Oracle Big Data Cloud Service allows for a 'Secure Setup' option. MIT Kerberos, Apache Sentry, HDFS Encryption (Key Trustee Servers are installed on the Master Nodes), Network Encryption and Auditing are enabled with a simple drop down menu during cluster creation. This allows you to secure your cluster right from its inception.
- » Client Network Access: Oracle Big Data Cloud Service provides a utility that uses Linux iptables and configuration files called whitelists to filter network traffic to services in the cluster. The whitelist configuration specifies whether network requests from specified clients will be accepted or denied for specified services at specified ports. When a connection tries to establish itself, iptables looks for a matching client IP address or range of IP addresses in the whitelist. If it

doesn't find one, it uses the default action. By default, the ports used for accessing Cloudera Manager and Hue are blocked. To control access to these ports, you must use the `bdacli` utility to add the ip address. It is also recommended to setup a IP Whitelist for each port. To open ports required by an application, additional security whitelists need to be created. The IP Whitelist can be configured for the following services/ports.

Services That Can Be Added to the Whitelist

The following table shows the services that can be added to the whitelist, the files that store the configurations for those services, and their default behaviors; that is, when their configuration files are empty:

Service	Whitelist Configuration File	Default (Empty File)
Cloudera Manager	<code>/opt/oracle/bda/cloudera_manager_whitelist</code>	Deny access
Hue	<code>/opt/oracle/bda/hue_whitelist</code>	Deny access
Secure Shell (SSH)	<code>/opt/oracle/bda/ssh_whitelist</code>	Allow access

Example:

```
# bdacli bdc_s Whitelist allow cloudera_manager 198.51.100.48
```

There are 3 types of compute nodes in the Oracle BDCS:

- » **Cluster nodes** are persistent data nodes that are part of the Hadoop cluster.
- » **Edge nodes** are persistent compute nodes with similar resources (Memory, CPU, Storage) that do not participate in the Hadoop cluster but are otherwise equal network citizens. They are all fully integrated in the InfiniBand network fabric.
- » **Burst Nodes** are transient cluster compute nodes. These nodes are light weight nodes that only include OCPUs and memory. For the Hadoop cluster, these nodes will not run data node roles. Burst Nodes provide additional compute power to the Cluster Nodes when necessary. Burst Nodes provide elastic compute capability on an hourly basis when required.

WANdisco Fusion on Oracle Big Data Cloud Services

WANdisco Fusion is a software application that supports continuous availability and performance with guaranteed data consistency across clusters any distance apart. Clusters can be deployed on any combination of Hadoop distributions, Hadoop compatible storage systems, or cloud environments. Oracle BDCS is a pre-configured Cloudera Hadoop distribution on Oracle Elastic Cloud, WANdisco Fusion and Oracle Big Data Cloud Services are perfectly suited to provide maximum availability of all your Hadoop data.

For additional information on WANdisco Fusion on Oracle BDCS, please refer to <https://www.wandisco.com/product/oracle-bda-oracle-bdcs>

Oracle Big Data Cloud Service and WANdisco Fusion allow you to:

- » Quickly provision a Secured Enterprise Class high performance Hadoop Cluster that is highly available and runs Spark on Yam in the Oracle Elastic Cloud with only a few clicks
- » Configure replication of data between BDCS instances in multiple regions, between BDCS and on-premise Hadoop clusters, or even between BDCS and other Hadoop environments in the Cloud, such as Oracle's Big Data Cloud Service – Compute Edition.
- » Enable true Active-Active replication of data between clusters such that updates to one environment are coordinated in real-time with the others, and delivering single-copy consistent hdfs data across all your Hadoop environments.
- » Provide a high performance single point of truth for all your data

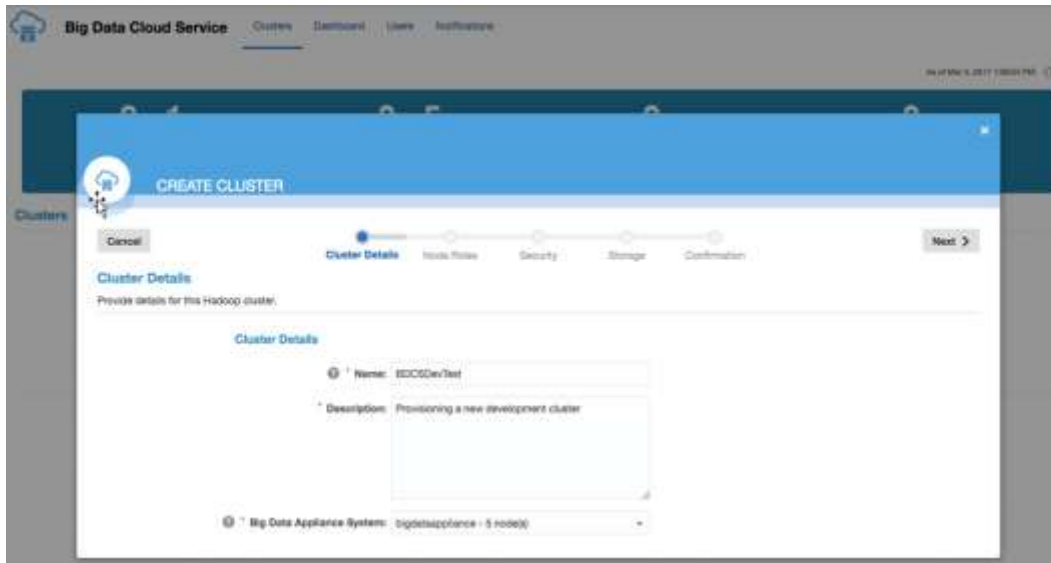
WANdisco Fusion Configuration Requirements for Oracle Big Data Cloud Service

Once you have an Oracle BDCS account, you create a Hadoop cluster, choose the number of Cluster nodes and edge nodes, navigate through a 5 step wizard which includes which includes the following:

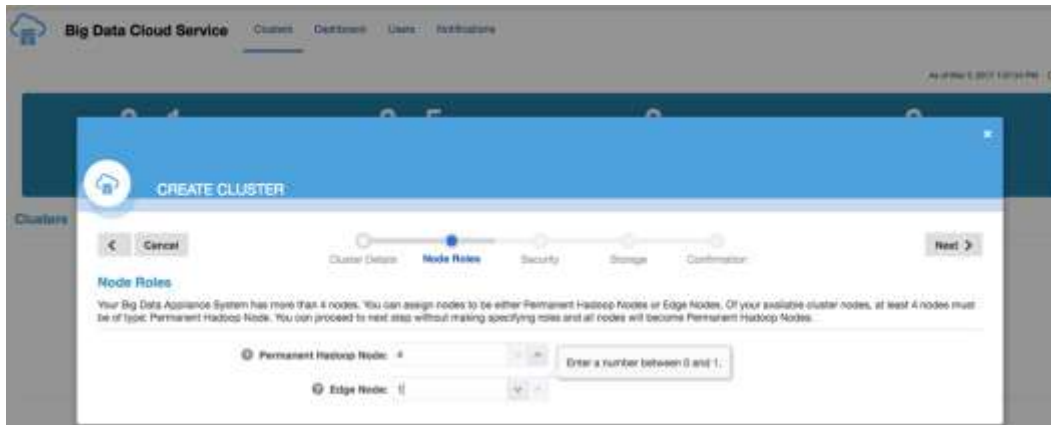
- » Name your cluster
- » Assign the node roles (Hadoop / Edge nodes)
- » Secure the cluster with a public ssh key, set the master password and optionally select the 'Secured Setup' option
- » Attaching an optional Storage Cloud with the cluster
- » Confirm your selections
- » Standard Cloudera CDH services will then be provisioned.
- » One Edge node is sufficient for WANdisco Fusion installation in a non-production configuration. For production use, installation of WANdisco Fusion on multiple edge nodes is recommended for purposes of High Availability.



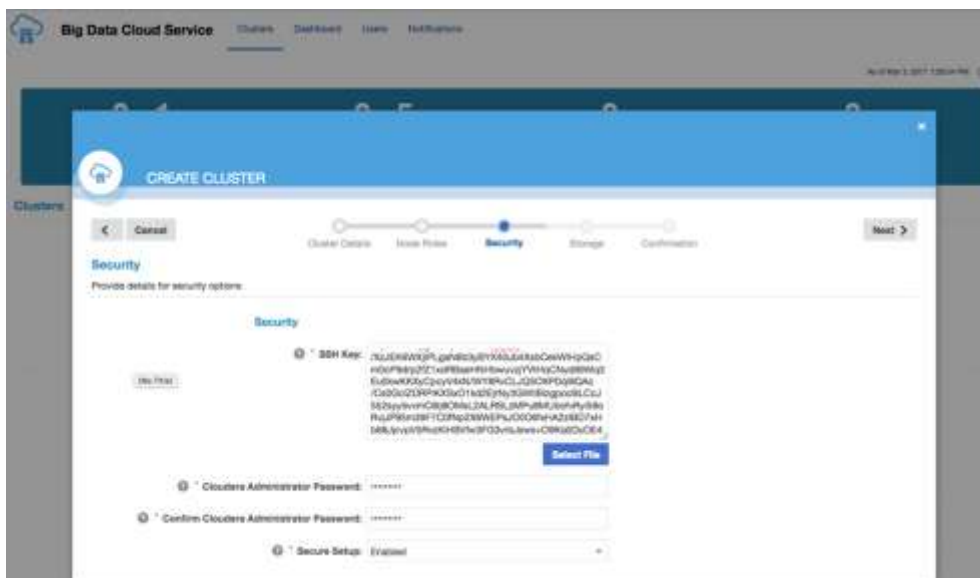
Press the + icon to create cluster



Enter the Hadoop cluster name



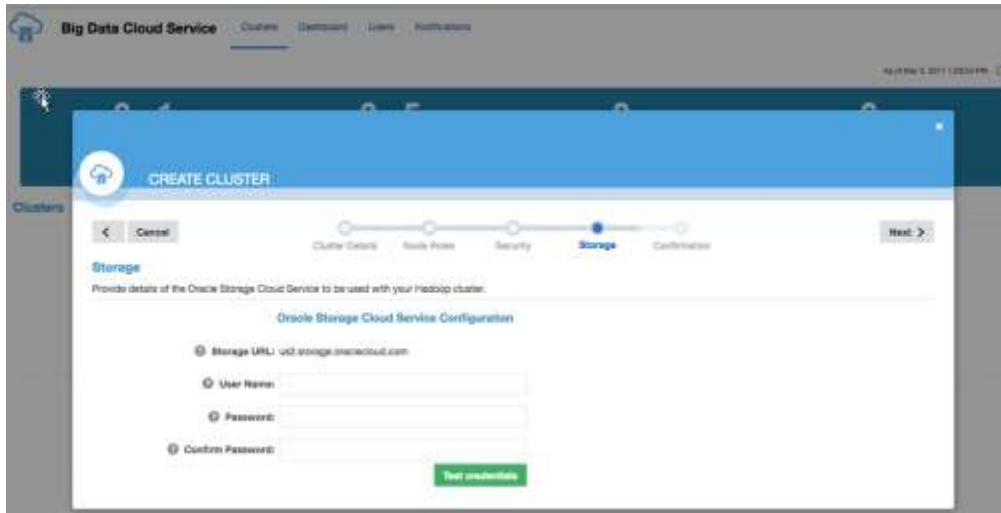
Enter the initial setup with Number of Hadoop cluster and Edge Nodes.



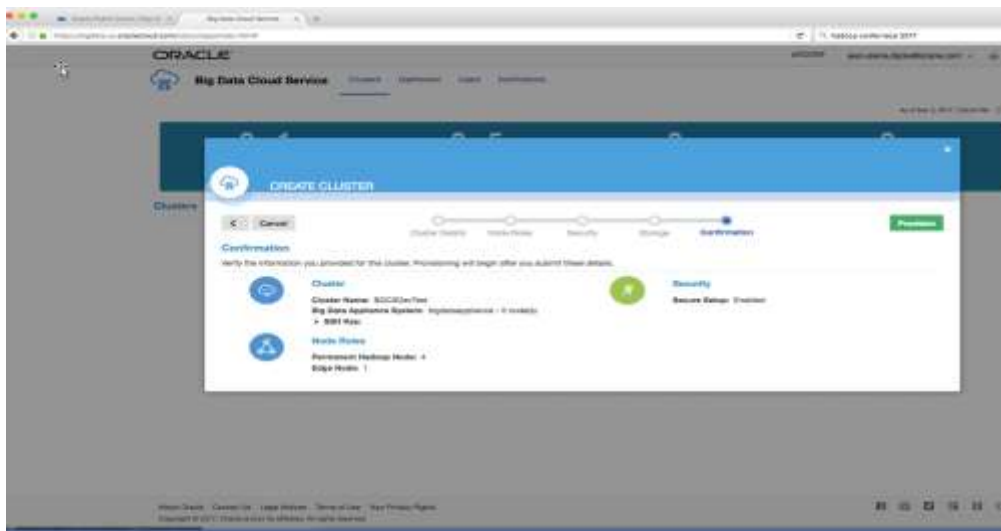
Paste a SSH public key and optionally enable security setup

Select **Enabled** to enable the included security services. When you select this option, you enable:

- » MIT Kerberos
- » Apache Sentry
- » HDFS Encryption—Key Trustee Servers are installed on the master nodes
- » Network Encryption
- » Data Encryption at rest (zones)
- » Auditing



Configure Oracle Storage Cloud Service container (optional) odcps utility is configured to copy data between HDFS and Oracle Storage Cloud Service. New to Oracle Big Data Cloud Service is the Oracle Big Data Manager which is a GUI tool to move data from Storage Cloud to HDFS and vice-versa.



That's it. You are done!

Once the cluster is created, you can now log onto the edge node with your private key to install the WANdisco Fusion application.

WANdisco Fusion – BDCS Deployment Considerations

Overview

Installation of WANdisco Fusion to an Oracle BDCS environment requires additional network configuration to allow the Fusion services to expose access to the management user interface for an administrator. This document details the steps required to make that adjustment to a stock BDCS environment

Configuration

The `/opt/oracle/bda/bin/bdacli` script includes details of the services that can be managed by an administrator of the BDCS environment. By default, it does not include entries that allow the WANdisco Fusion user interface to be accessed.

Modification of this script is required to introduce that mechanism.

Prior to the installation of WANdisco Fusion, provide access from external hosts to the port exposed by the WANdisco Fusion user interface by taking the following actions:

1. Add an entry to the `@ValidBDCSServices` definition: `my @ValidBDCSServices = ({ "name" => "wandisco_fusion", "display_name" => "WANdisco Fusion", "port" => "8083", "default" => "drop", "filename" => "wandisco_fusion_whitelist" },`
2. Create an empty file `/opt/oracle/bda/wandisco_fusion_whitelist`: `touch /opt/oracle/bda/wandisco_fusion_whitelist`
3. Allow access with that configuration: `bdacli bdcscs_whitelist allow wandisco_fusion <IP address from which access is required>`
- 4.

Additional Considerations

It should also be noted that network access between the BDCS environment and external clusters for which WANdisco Fusion provides replication is required. In particular, you should ensure that connectivity is available from other environments to the ports detailed in the product documentation at <https://docs.wandisco.com/bigdata/wdfusion/2.10/#ports>.

WANdisco Fusion Installation Guidelines

On premises installation

The following section covers the installation of WANdisco Fusion into a cluster that is based in your organization's own premises.

Starting the installation

Use the following steps to complete an installation using the installer file. This requires an administrator to enter details throughout the procedure. Once the initial settings are entered through the terminal session, the installation is then completed through a browser or alternatively, using a Silent Installation option to handle configuration programmatically.

Open a terminal session on your first installation server. Copy the WD Fusion installer script into a suitable directory.

1. Make the script executable, e.g.
`chmod +x fusion-ui-server-<version>_rpm_installer.sh`
2. Execute the file with root permissions, e.g.
`sudo ./fusion-ui-server-<version>_rpm_installer.sh`
3. The installer will now start.

Verifying archive integrity... All good.

Uncompressing WANdisco Fusion.....

```
  ::  ::  ::  #  #  ##  ###  #####  #  #####  #####  #####
  :::  :::  :::  #  #  #  #  ##  ##  #  #  #  #  #  #  #  #
  :::::::::::  :::  #  #  #  #  #  #  #  #  #  #  #  #  #  #
  :::::::::::  :::  #  #  #  #  #  #  #  #  #  #  #  #  #  #
  :::::::::::  :::  #  #  #  #  #  #  #  #  #  #  #  #  #  #
  :::  :::  :::  ##  ##  #  ##  #  #  #  #  #  #  #  #  #
  ::  ::  ::  #  #  ##  ##  #  #####  #  #####  #####  #####

Welcome to the WANdisco Fusion installation

You are about to install WANdisco Fusion version 2.4-206

Do you want to continue with the installation? (Y/n) y
```

The installer will perform an integrity check, confirm the product version that will be installed, then invite you to continue. Enter "Y" to continue the installation. The installer performs some basic checks and lets you modify the Java heap settings. The heap settings apply only to the WD Fusion UI.

```
Checking prerequisites:
Checking for perl: OK
Checking for java: OK

INFO: Using the following Memory settings:

INFO: -Xms128m -Xmx512m

Do you want to use these settings for the installation? (Y/n) y
```

The installer checks for Perl and Java. See the Installation Checklist [Java Requirements](#) for more information about these requirements. Enter "Y" to continue the installation.

1. Next, confirm the port that will be used to access WD Fusion through a browser.

```
Which port should the UI Server listen on? [8083]:
```

2. Select the Hadoop version and type from the list of supported platforms:

```
Please specify the appropriate backend from the list below:
[0] cdh-5.2.0
[1] cdh-5.3.0
[2] cdh-5.4.0
[3] cdh-5.5.0
[4] cdh-5.6.0
[5] cdh-5.7.0
[6] cdh-5.8.0

Which fusion backend do you wish to use? 5
You chose cdh-5.7.0:5.7.0.2.3.0.0-2557
```

3. The installer now confirms which system user/group will be applied to WD Fusion.

```
We strongly advise against running Fusion as the root user. For default HDFS setups, this is set to 'hdfs'.
However, you should choose a user appropriate for running HDFS commands on your system.

Which user should Fusion run as? [hdfs]
Checking 'hdfs' ... .. 'hdfs' found.

Please choose an appropriate group for your system.
```

By default CDH uses the 'hadoop' group.
Which group should Fusion run as? [hadoop]
Checking 'hadoop' 'hadoop' found.

4. The installer does a search for the commonly used account and group, assigning these by default. Check the summary to confirm that your chosen settings are appropriate: Installing with the following settings:

```
User and Group:                hdfs:hadoop
Hostname:                      node04-example.host.com
Fusion Admin UI Listening on:    0.0.0.0:8083
Fusion Admin UI Minimum Memory: 128
Fusion Admin UI Maximum memory: 512
Platform:                      cdh-5.7.0 (5.7.0.2.3.0.0-2557)
Manager Type                   CLOUDERA
Manager Host and Port:         :
Fusion Server Hostname and Port: node04-example.host.com:8082
SSL Enabled:                   false
Do you want to continue with the installation? (Y/n) y
```

You are now given a summary of all the settings provided so far. If these settings are correct then enter "Y" to complete the installation of the WD Fusion server.

5. The package will now install.

```
Installing cdh-5.7.0 packages:
  fusion-cdh-5.7.0-server-2.9_SNAPSHOT-1130.noarch.rpm ...
      Done
  fusion-cdh-5.7.0-ihc-server-2.9_SNAPSHOT-1130.noarch.rpm ...
      Done
Installing fusion-ui-server package Starting fusion-ui-server:[ OK ]
Checking if the GUI is listening on port 8083: .....Done
```

6. The WD Fusion server will now start up:

```
Please visit http://<YOUR-SERVER-ADDRESS>.com:8083/ to access the WANdisco Fusion
If 'http://<YOUR-SERVER-ADDRESS>.com' is internal or not available from your browser, replace this with
an externally available address to access it.
Installation Complete [root@node05 opt]#
```

At this point the WD Fusion server and corresponding IHC server will be installed. The next step is to configure the WD Fusion UI through a browser or using the silent installation script.

Configure WD Fusion through a browser

Follow this section to complete the installation by configuring WD Fusion using a browser-based graphical user interface.

1. Open a web browser and point it at the provided URL. e.g.
`http://<YOUR-SERVER-ADDRESS>.com:8083/`
2. In the first "Welcome" screen you're asked to choose between **Create a new Zone** and **Add to an existing Zone**.



Figure 1. Welcome - Installer screen

Make your selection as follows: **Adding a new WD Fusion cluster** Select **Add Zone**. **Adding additional WD Fusion servers to an existing WD Fusion cluster** Select **Add to an existing Zone**.

High Availability for WD Fusion / IHC Servers

It's possible to enable High Availability in your WD Fusion cluster by adding additional WD Fusion/IHC servers to a zone. These additional nodes ensure that in the event of a system outage, there will remain sufficient WD Fusion/IHC servers running to maintain replication.

Add HA nodes to the cluster using the installer and choosing to Add to an existing Zone, using a new node name.

Configuration for High Availability

When setting up the configuration for a High Availability cluster, ensure that `fs.defaultFS`, located in the `core-site.xml` is not duplicated between zones. This property is used to determine if an operation is being executed locally or remotely, if two separate zones have the same default file system address, then problems will occur. **WD Fusion should never see the same URI (Scheme + authority) for two different clusters.**

- Run through the installer's detailed Environment checks.

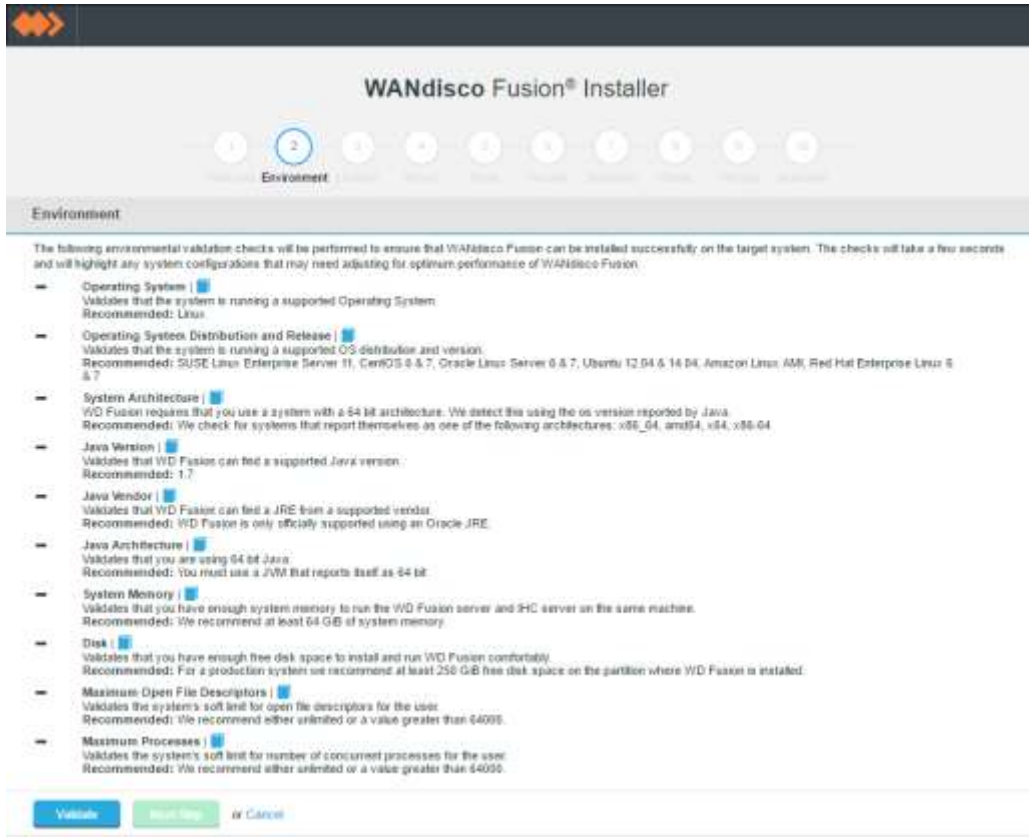


Figure 2. Installer screen

- On clicking validate the installer will run through a series of checks of your system's hardware and software setup and warn you if any of WD Fusion's prerequisites are missing.

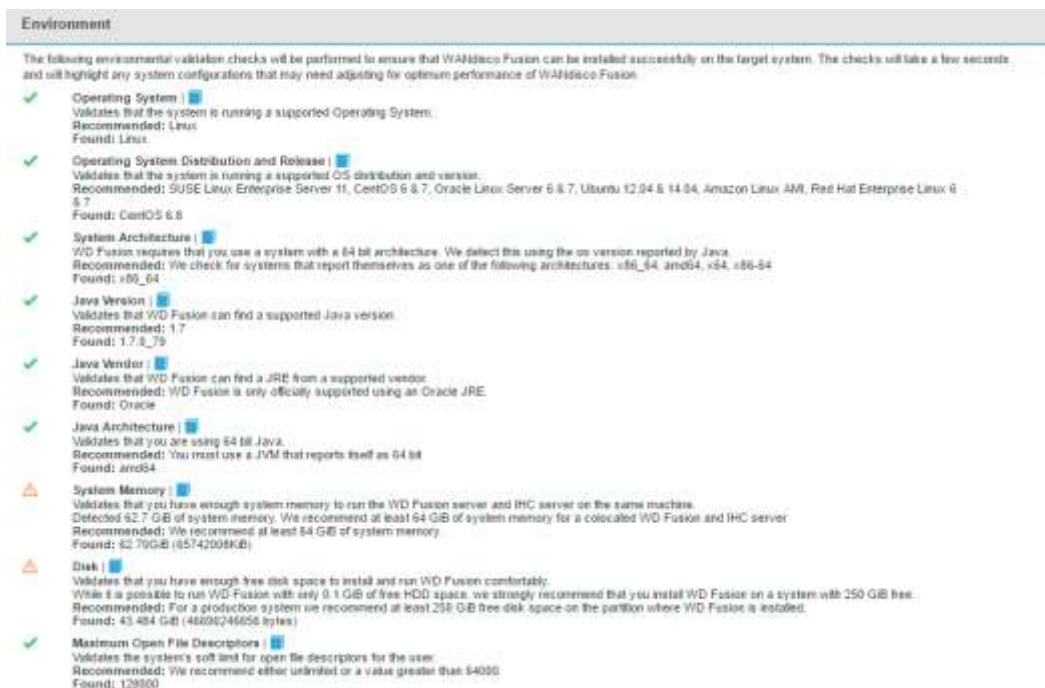


Figure 3. Validation results

Any element that fails the check should be addressed before you continue the installation. Warnings may be ignored for the purposes of completing the installation, especially if the installation is only for evaluation purposes and not for production. However, when installing for production, you should also address all warnings, or at least take note of them and exercise due care if you continue the installation without resolving and revalidating.

1. Upload the license file.

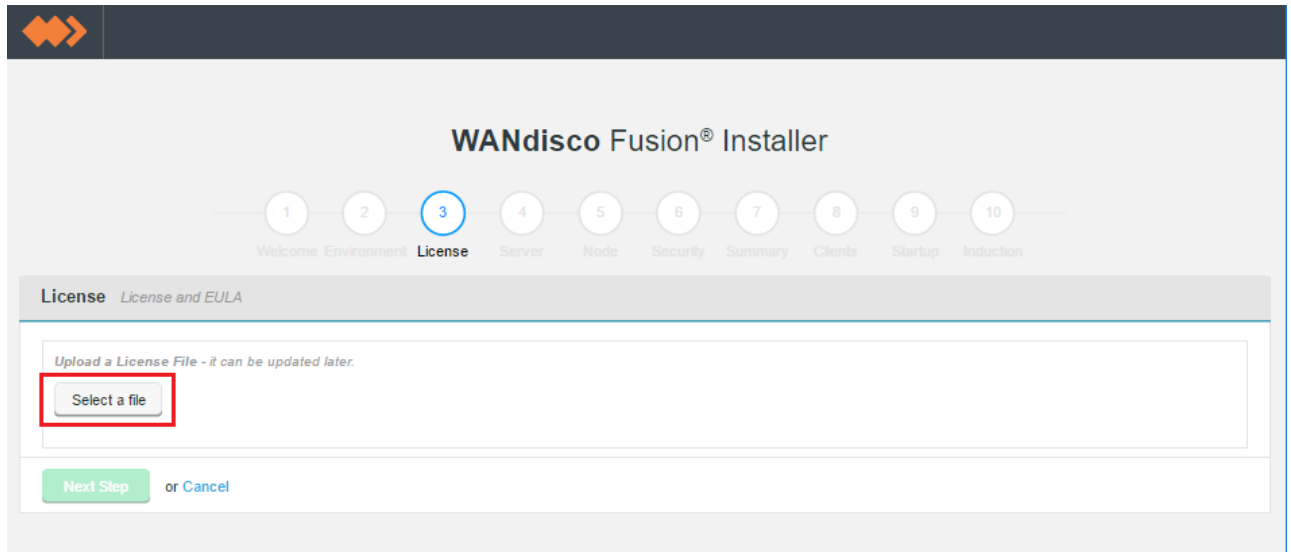


Figure 4. Installer screen

5. The conditions of your license agreement will be presented in the top panel, including License Type, Expiry data, Name Node Limit and Data Node Limit.

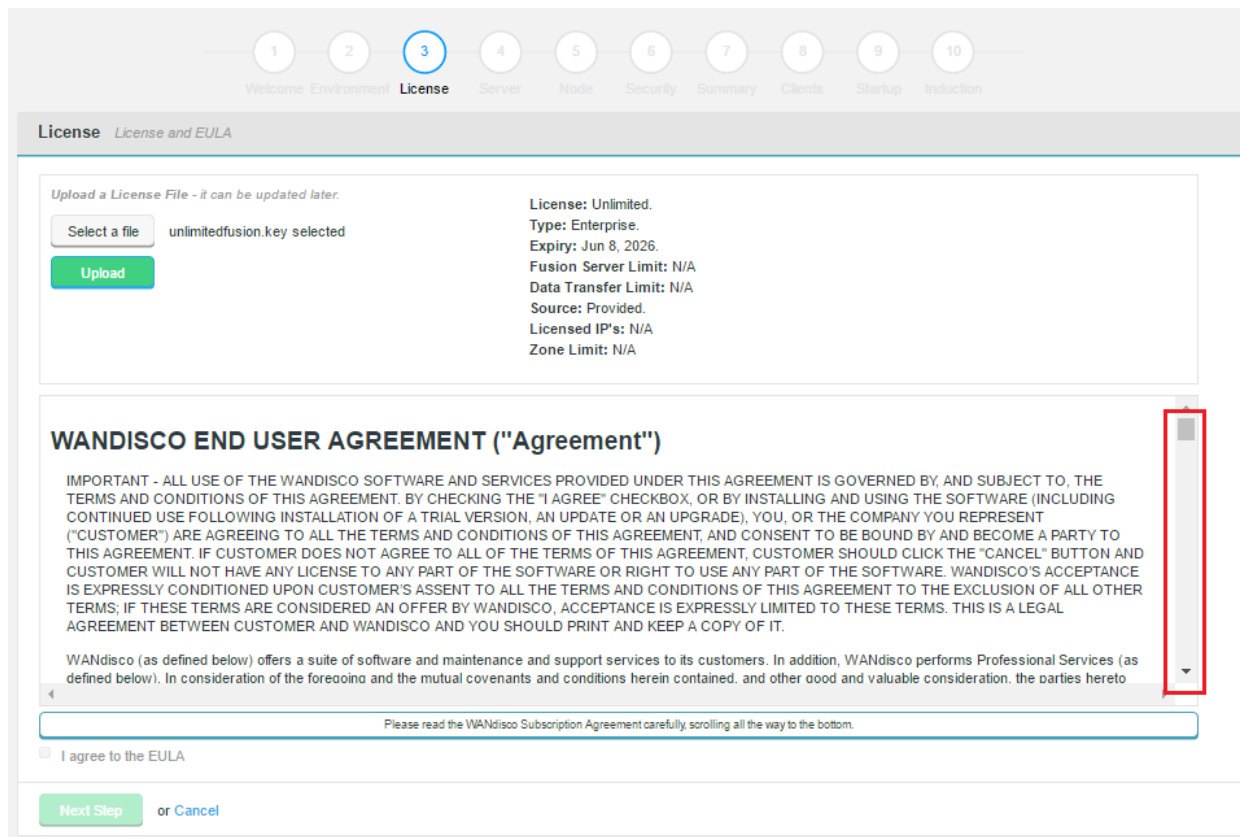


Figure 5. Check screen

6. Read through the EULA. When the scroll bar reaches the bottom you can click on the “*I agree to the EULA*” to continue, then click **Next Step**.



Figure 6. License screen

7. Enter settings for the WD Fusion server.

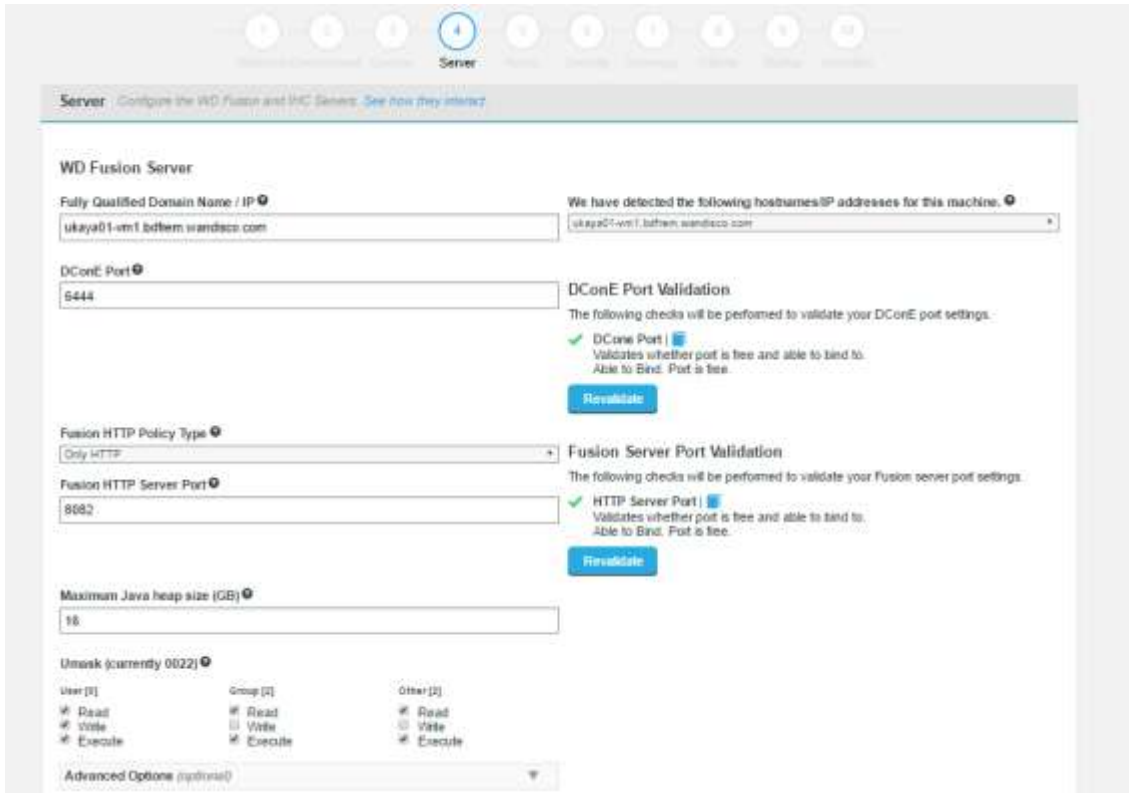


Figure 7. Server settings

WD Fusion Server

Fully Qualified Domain Name / IP

The full hostname for the server.

We have detected the following hostname/IP addresses for this machine.

The installer will try to detect the server's hostname from its network settings. Additional hostnames will be listed on a dropdown selector.

DConE Port

TCP port used by WD Fusion for replicated traffic. Validation will check that the port is free and that it can be bound to.

Fusion HTTP Policy Type

Sets the policy for communication with the WD Fusion Core Server API.

Select from one of the following policies:

Only HTTP - WD Fusion will not use SSL encryption on its API traffic.

Only HTTPS - WD Fusion will only use SSL encryption for API traffic.

Use HTTP and HTTPS - WD Fusion will use both encrypted and un-encrypted traffic.

Known Issue

Currently, the HTTP policy and SSL settings both independently alter how WD Fusion uses SSL, when they should be linked. You need to make sure that your HTTP policy selection and the use of SSL (enabled in the next section of the Installer) are in sync. If you choose either to the policies that use HTTPS, then you must enable SSL. If you stick with "Only HTTP" then you must ensure that you do not enable SSL. In a future release these two settings will be linked so it won't be possible to have contradictory settings.

Fusion HTTP Server Port

The TCP port used for standard HTTP traffic. Validation checks whether the port is free and that it can be bound.

Maximum Java heap size (GB)

Enter the maximum Java Heap value for the WD inter-Hadoop Communication server.

Umask (currently 0022)

Set the default permissions applied to newly created files. The value 022 results in default directory permissions 755 and default file permissions 644. This ensures that the installation will be able to start up/restart.

IHC Server

IHC Server

Maximum Java heap size (GB) ⓘ
16

IHC network interface ⓘ
ukaya01-vm1.bdfrem.wandisco.com

We have detected the following hostnames/IP addresses for this machine. ⓘ
ukaya01-vm1.bdfrem.wandisco.com

WARNING: Multiple local IP addresses detected. Please specify the Fully Qualified Domain Name or IP address, which other nodes will be able to use to communicate with this IHC Server.

Advanced Options (optional) ▲

Advanced options should be used by experienced users only.

IHC server binding address ⓘ
0.0.0.0

WARNING: Please enter the Fully Qualified Domain Name or IP address which the IHC server will bind to, or select from one of the options below. If left empty, this value will default to 0.0.0.0.
0.0.0.0

Next step or go back

Figure 8. IHC Server details

Maximum Java heap size (GB)

Enter the maximum Java Heap value for the WD Inter-Hadoop Communication server.

IHC network interface

The hostname for the IHC server.

Advanced Options (optional)

IHC server binding address

In the advanced settings you can decide which address the IHC server will bind to. The address is optional, by default the IHC server binds to all interfaces (0.0.0.0), using the port specified in the `ihc.server` field. In all cases the port should be identical to the port used in the `ihc.server` address. i.e.

```
_/etc/wandisco/fusion/ihc/server/cdh-5.4.0/2.6.0-cdh5.4.0.ihc_ or _/etc/wandisco/fusion/ihc/server/localfs-2.7.0/2.7.0.ihc_
```

Once all settings have been entered, click **Next step**.

» Next, you will enter the settings for your new Zone.

Zone Information

Entry fields for zone properties:

Fully Qualified Domain Name

The full hostname for the server.

Node ID

A unique identifier that will be used by WD Fusion UI to identify the server.

Induction failure

If induction fails, attempting a fresh installation may be the most straight forward cure, however, it is possible to push through an induction manually, using the REST API.

Known issue with Node IDs

You must use different Node IDs for each zone. If you use the same name for multiple zones, then you will not be able to complete the induction between those nodes.

DConE Port

TCP port used by WD Fusion for replicated traffic.

Zone Name

The name used to identify the zone in which the server operates.

Management Endpoint

Select the Hadoop manager that you are using, i.e. Cloudera Manager, Ambari or Pivotal HD. The selection will display the entry fields for your selected manager.

URI Selection

The default behavior for WD Fusion is to fix all replication to the Hadoop Distributed File System / `hdfs://` URI. Setting the `hdfs-scheme` provides the widest support for Hadoop client applications, since some applications can't support the available "`fusion://`" URI they can only use the HDFS protocol. Each option is explained below:

Use HDFS URI with HDFS file system

The element appears in a radio button selector:



Figure 9. URI option A

This option is available for deployments where the Hadoop applications support neither the WD Fusion URI nor the HCFS standards. WD Fusion operates entirely within HDFS.

This configuration will not allow paths with the `fusion:///` uri to be used; only paths starting with `hdfs:///` or no scheme that correspond to a mapped path will be replicated. The underlying file system will be an instance of the HDFS DistributedFileSystem, which will support applications that aren't written to the HCFS specification.

Use WD Fusion URI with HCFS file system



Figure 10. URI option B

This is the default option that applies if you don't enable *Advanced Options*, and was the only option in WD Fusion prior to version 2.6. When selected, you need to use `fusion://` for all data that must be replicated over an instance of the Hadoop Compatible File System. If your deployment includes Hadoop applications that are either unable to support the Fusion URI or are not written to the HCFS specification, this option will not work.

Use Fusion URI with HDFS file system




Figure 11. URI option C

This differs from the default in that while the WD Fusion URI is used to identify data to be replicated, the replication is performed using HDFS itself. This option should be used if you are deploying applications that can support the WD Fusion URI but not the Hadoop Compatible File System.

Benefits of HDFS.

The following advanced options provide a number of low level configuration settings that may be required for installation into certain environments. The incorrect application of some of these settings could cause serious problems, so for this reason we strongly recommend that you discuss their use with WANdisco's support team before enabling them.

Use Fusion URI and HDFS URI with HDFS file system

This "mixed mode" supports all the replication schemes (`fusion://`, `hdfs://` and no scheme) and uses HDFS for the underlying file system, to support applications that aren't written to the HCFS specification.




Figure 12. URI option D

Cloudera / Ambari Configuration

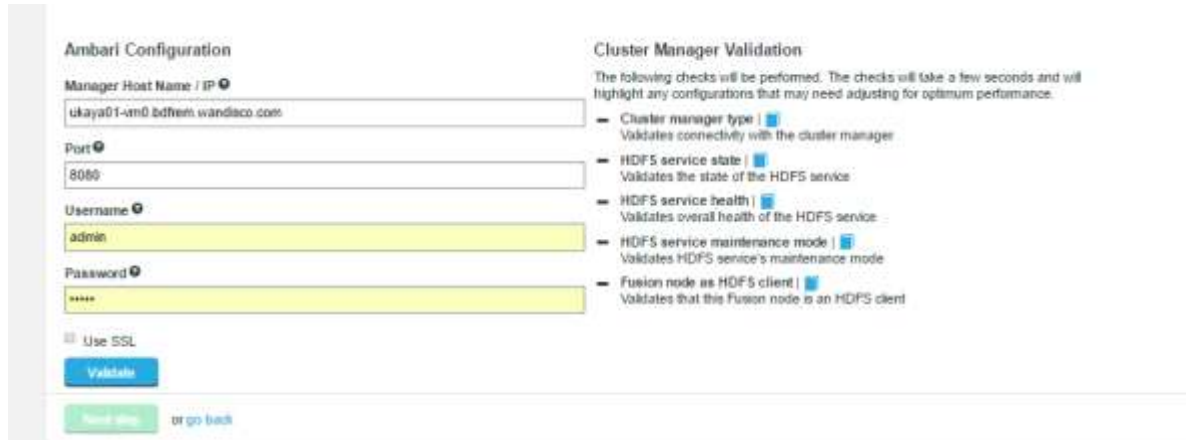


Figure 13. Manager Configuration

Cluster manager type

Validates connectivity with the cluster manager.

HDFS service state

Validates the state of the HDFS service.

HDFS service health

Validates the overall health of the HDFS service.

HDFS service health

Validates HDFS service's maintenance mode.

Fusion node as HDFS client

Validates that this Fusion node is a HDFS client.

Induction

Induction is the process used to incorporate new nodes into WANdisco's replication system. The process is run at the end of a node installation, although it is also possible to delay the process, then use the **+ Induct** link on the **Fusion Nodes** tab.

Use this procedure if you have installed a new node but did not complete its induction into your replication system at the end of the installation process.

1. Login to one of the active nodes, clicking on the Fusion **Nodes** tab. Click the **+ Induct** button.

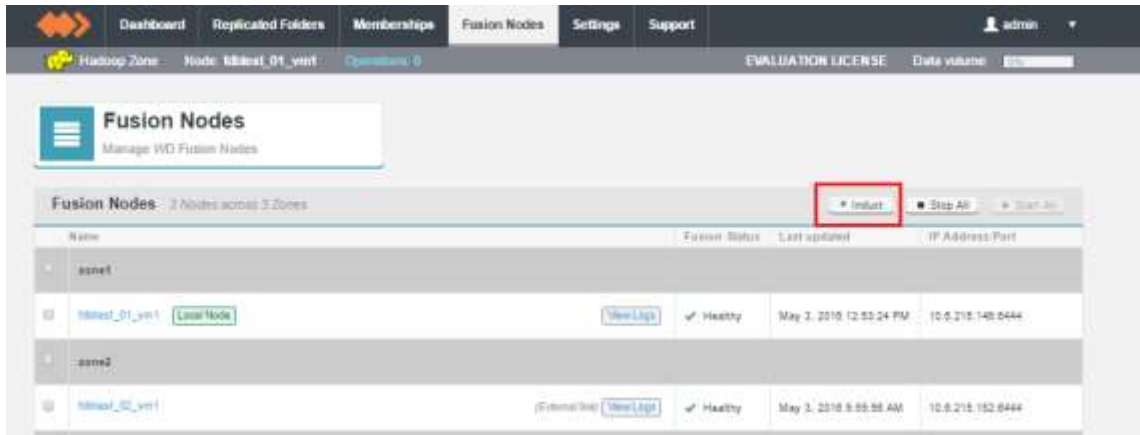


Figure 14. +Induct

2. Enter the fully qualified domain name of the new node that you wish to induct into your replication system.

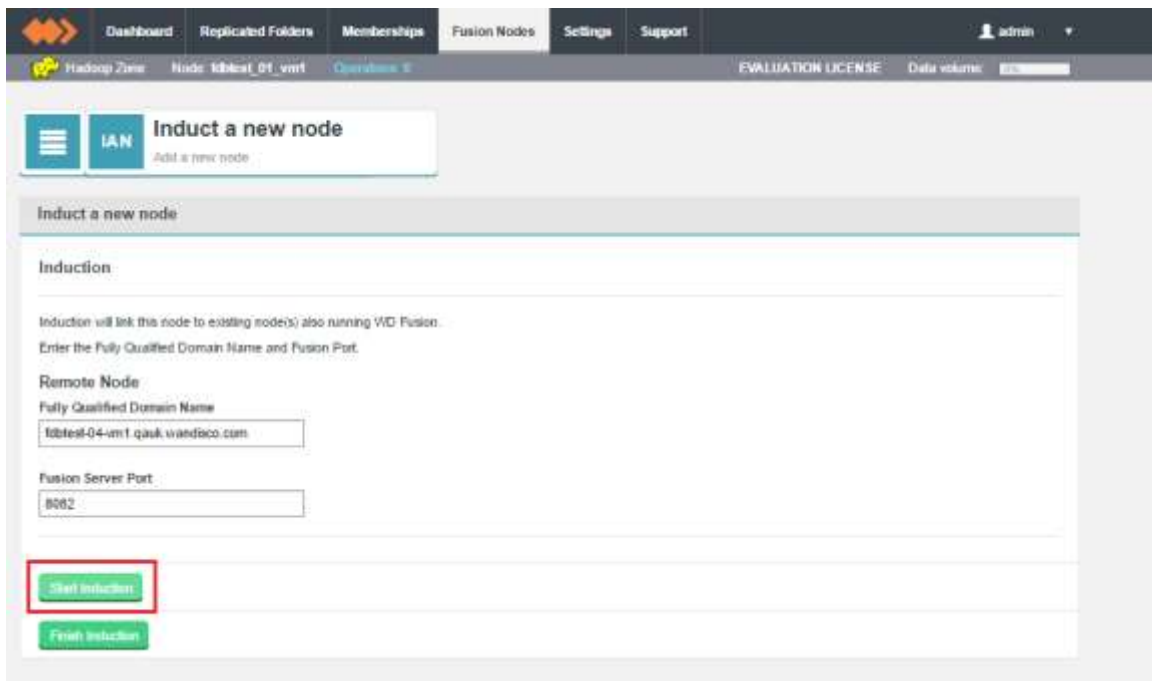


Figure 15. Start Induction

Fully Qualified Domain Name

The full domain name for the new node that you will induct into your replication system.

Fusion Server Port

The TCP port is used by the WD Fusion application for configuration and reporting, both internally and via REST API. The port needs to be open between all WD Fusion nodes and any systems or scripts that interface with WD Fusion through the REST API.

Click **Start Induction**.

When the induction process completes, the **Fusion Node** tab will refresh with the new node added to the list.

Managing Replication

WD Fusion is built on WANdisco's patented DConE active-active replication technology. DConE sets a requirement that all replicating nodes that synchronize data with each other are joined in a "membership". Memberships are coordinated groups of

Create a membership

Login to the WD Fusion UI. Click on the **Membership** tab. Click on the **Create New** tab. The "New Membership" window will open that will display the WD Fusion nodes organized by zone.

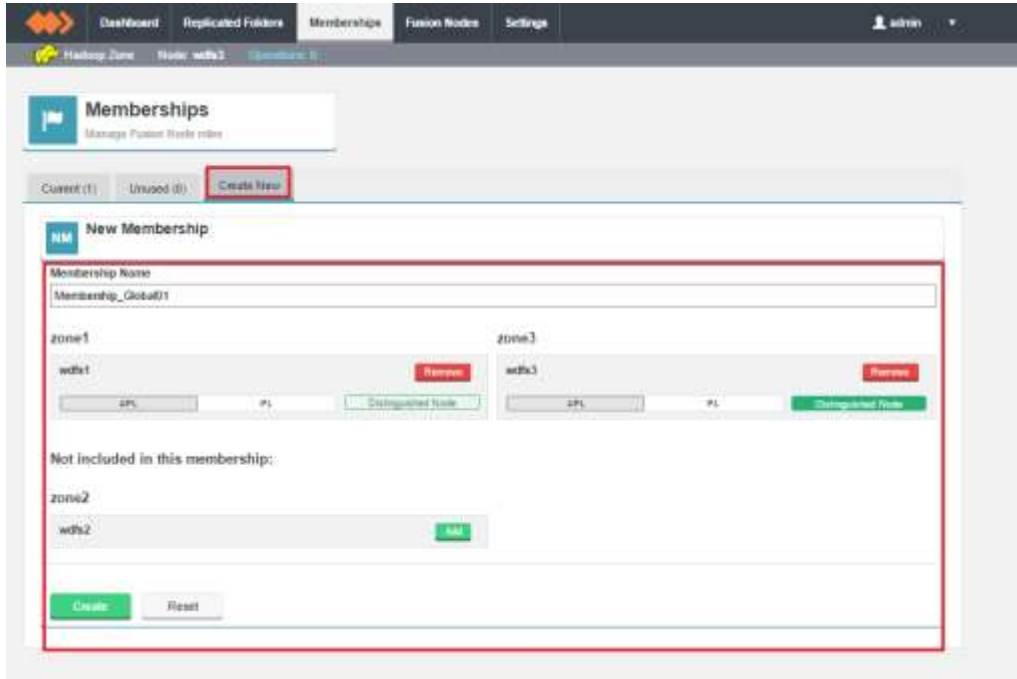


Figure 16. Example membership 1

Configure the membership by selecting which nodes should be acceptors. Acceptors vote on the ordering of changes.

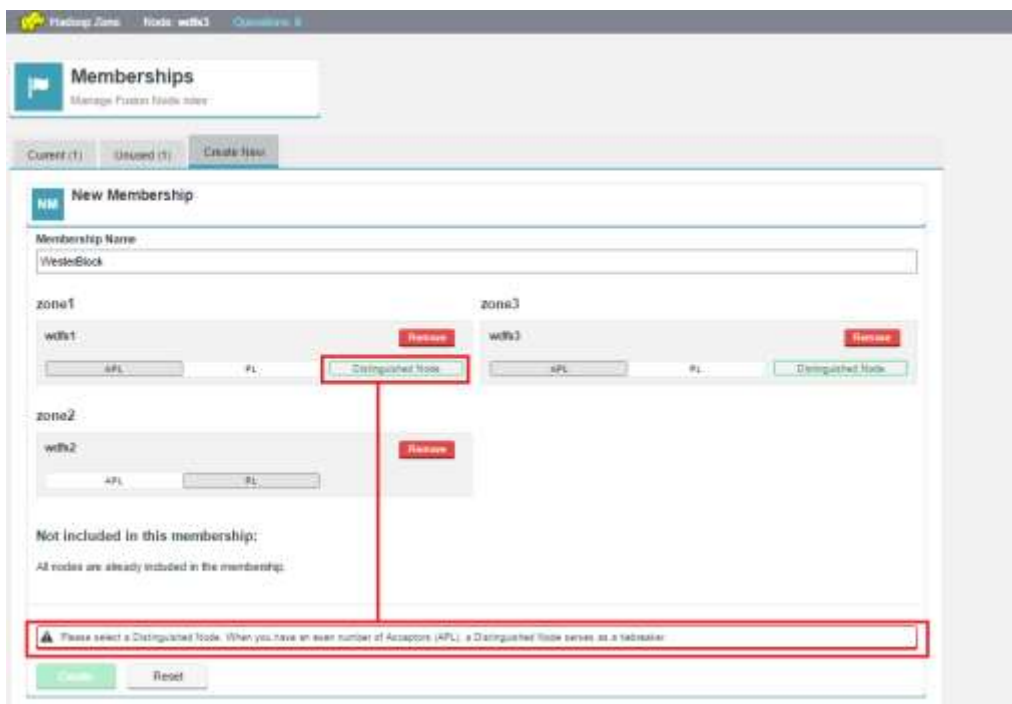


Figure 17. Example membership 2

Note how a two-node membership requires that one of the nodes be upgraded to a Distinguished Node.

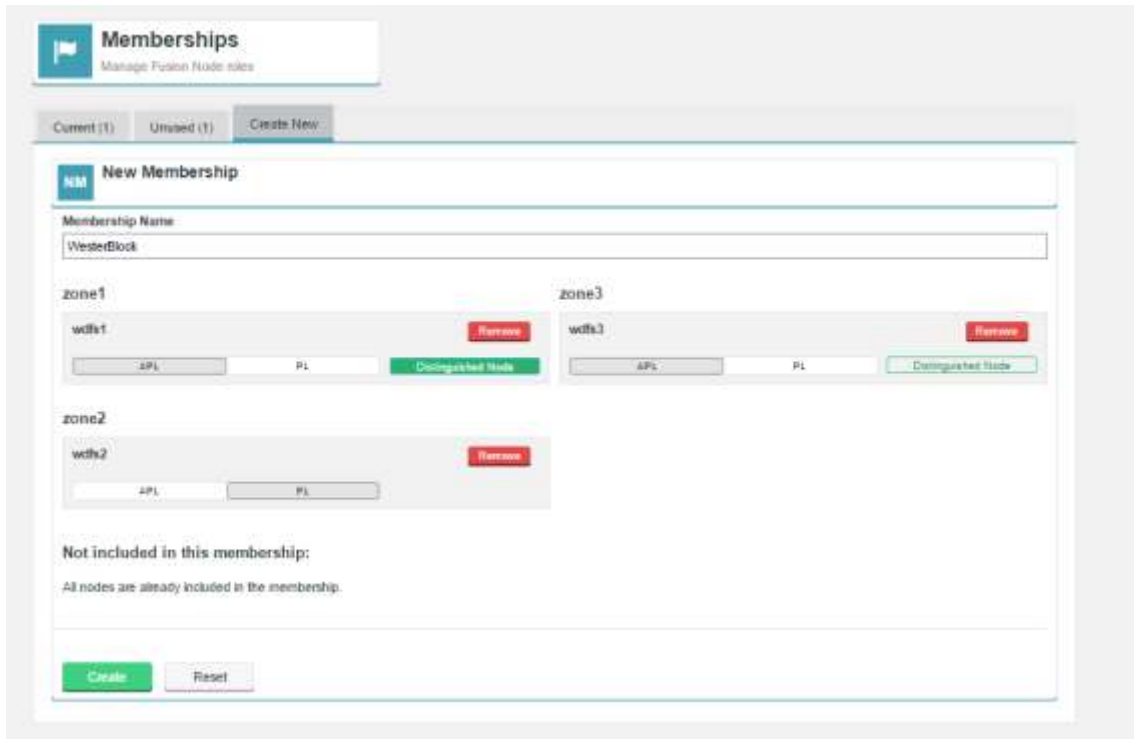


Figure 18. Example membership2

Click **Create** to complete the operation. Click **Cancel** to discard the changes.

Identical memberships are not allowed

You will be prevented from creating more than 1 membership with a particular configuration.

Guide to node types

APL

Acceptor - the node will vote on the order in which replicated changes will play out.

Proposer - the node will create proposals for changes that can be applied to the other nodes.

Learner - the node will receive replication traffic that will synchronize its data with other nodes.

PL

Proposer - the node will create proposals for changes that can be applied to the other nodes.

Learner - the node will receive replication traffic that will synchronize its data with other nodes.

Distinguished Node

Acceptor + - the distinguished node is used in situations where there is an even number of nodes, a configuration that introduces the risk of a tied vote. The Distinguished Node's bigger vote ensures that it is not possible for a vote to become tied.

Replicated Folders

WD Fusion allows selected folders within your hdfs file system to replicated to other data centers in your cluster. This section covers the set up and management of replicated folders.

Create a replicated folder

The first step in setting up a replicated folder is the creation of a target folder:

1. In each zone, create a directory in the hdfs file space. To avoid permission problems, ensure that the owning user/group are identical across the zones. Use Hadoop's filesystem command to complete the tasks:

```
hadoop fs -mkdir /user/hiver  hadoop fs -chown -R hiver:groupname /user/hiver
```

2. As user `hdfs`, run the following commands on each data center:

```
hadoop fs -mkdir /user/hiver/warehouse-replicated  
hadoop fs -chown hiver:hiver /user/hiver/warehouse-replicated
```

This ensures that the a universal system user has read/write access to the HDFS directory `warehouse-replicated` that will be replicated through WD Fusion.

Create Rule

1. Once the folder is in place on all nodes, login to WD Fusion's UI on one of the WD Fusion nodes and click on the **Replicated Folders** tab.
2. Click on the **+ Create** button.



Figure 19. Create rule 1

1. The replicated folder entry form screen will appear.

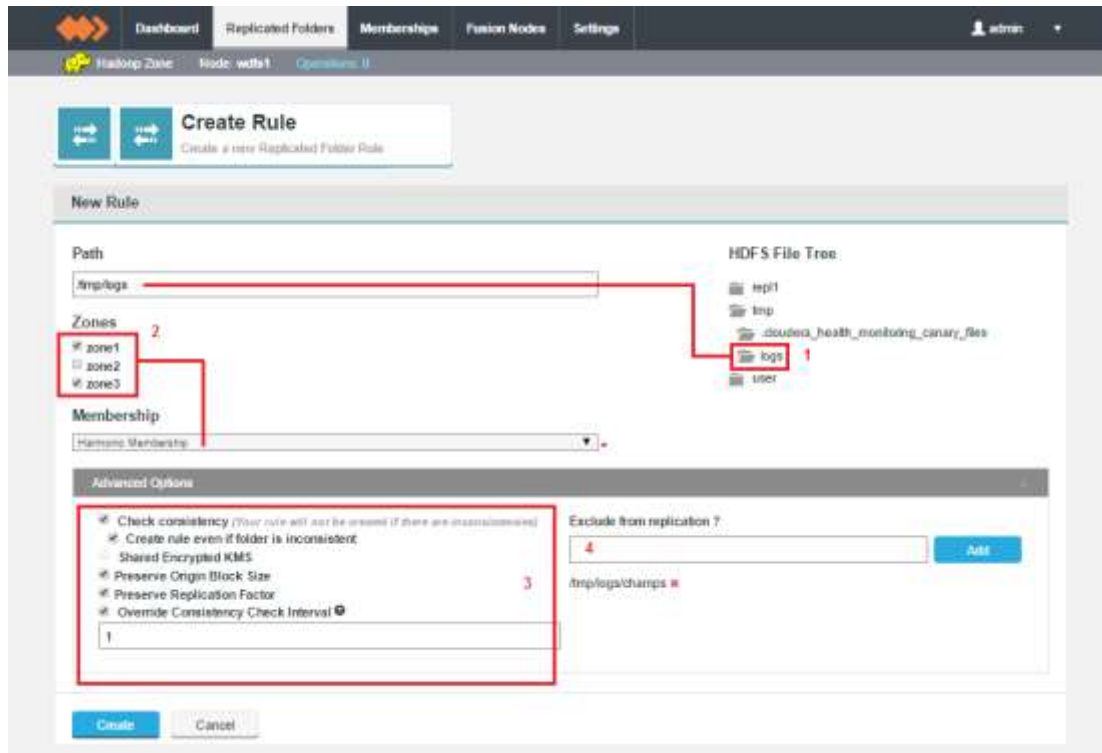


Figure 20. Create rule 2

Navigate the **HDFS File Tree (1)**, on the right-hand side of the **New Rule** panel to select your target folder, created in the previous section. The selected folder will appear in the **Path** entry field. You can, instead, enter the full path to the folder in the **Path** directory.

Next, select two or more zones from the **Zones list (2)**. You then select a **Membership** from the dropdown selector. If there's no existing membership with the combination of Zones that you selected, then you will see the message:

There are no memberships available matching your criteria.

In this case you can create a new membership, see *Create a membership* and restart the Create Replicated Folder process.

- » You can now complete the creation of the **Replicated folder** by clicking on the **Create** button. However, there are some additional options available on the **Advanced Options** panel. Consider if you need to apply any [Advanced Options](#) for the folder.

Note that the allocated writer for this zone is listed under the **Advanced Options** panel. This can be useful information in case you need to troubleshoot replication problems.



Figure 21. Advanced Options

These include **Preserve Origin Block Size** which is used for columnar storage formats such as parquet, **Preserve Replication Factor** which is used when want replica data to continue to use the replication factor that is set on its originating cluster, rather than the use the factor that applies on the new cluster.

Exclude from replication? lets you set an "exclude pattern" to indicate files and folders in your replicated folder that you don't want to be replicated. If you apply any Advanced Options you need to click the **Update** button to make sure that they are applied.

The option **Override Consistency Check Interval** allows administrators to set a consistency check interval that is specific to the replicated folder space and different from the default value that is set in the **Consistency Check** section of the **Settings** tab.

Path interpretation

If the path contains a leading slash "/", we assume it is an absolute path, if it contains no leading slash then we assume it is a relative path and the root directory will be added to the beginning of the exclusion.

1. If you didn't complete a consistency check on the selected folder, you may do so now.

Replication Rules						
Path	Target Zones	Current Membership	Consistency	File Transfers		
/repl1	dc1, dc2	simpleMembership	Unknown Last Check: Never Next Check: Apr 10, 2015 10:52:45 AM	View No files in transfer		
/tmp	dc1, dc2	simpleMembership	Unknown Last Check: Never Next Check: Apr 10, 2015 10:52:45 AM	View No files in transfer		

Figure 22. Replicate to Zones

- After the completion of a consistency check, the **Consistency** column will report the consistency status.

Path	Target Zones	Current Membership	Consistency	File Transfers
reg1	dc1, dc2	simpleMembership	Consistent Last Check: Apr 9, 2015 4:45:34 PM Next Check: Apr 10, 2015 4:45:34 PM	No files in transfer.
tmp	dc1, dc2	simpleMembership	Consistent Last Check: Apr 9, 2015 4:20:38 PM Next Check: Apr 10, 2015 4:20:38 PM	No files in transfer.

Figure 23. Replicated folder status

Configuration Used for Validation Tests

Once installed and configured as defined above, WANdisco Fusion is able to replicate data between multiple BDCS environments. Validation of operation against Oracle Big Data Cloud Services has been performed in environments made available by Oracle, including sample data sets provided by Oracle along with large datasets generated by standard Hadoop tooling (TeraGen). This validation included replication across on-premises deployments of Hadoop and Oracle BDCS, and between third-party public cloud vendor Hadoop deployments and Oracle BDCS.

BDCS Environment

A 5-node BDCS environment was used for test purposes, with no Kerberos configured (although WANdisco Fusion will interoperate with a secured BDCS deployment also).

Connectivity

A single “opc” user available with SSH access and sudo privileges for the purpose of installing WANdisco Fusion in the BDCS environment and configuring network connectivity as explained in this document.

A “wandisco” user provisioned on the BDCS environment with access privileges to test data.

Test data

A variety of public data sets were made available under a /user/oracle directory. The wandisco user was used to create further test datasets in HDFS.

In each test case, replication was performed entirely by WANdisco Fusion without the need to leverage additional replication technologies or protocols. With network connectivity available into the BDCS environment, arbitrary external systems can participate in replication regardless of their own external connectivity, demonstrating the ability for WANdisco Fusion to support active data replication even when one of the participating systems cannot accept inbound TCP connections.

Several use cases were performed to demonstrate the active data replication capabilities of WANdisco Fusion. Over a terabyte of sample data was introduced and replicated using the TeraGen Hadoop examples. Multiple data sets were added into the HDFS file system in a BDCS environment using TeraGen to demonstrate bidirectional replication between the two BDCS instances. Validation includes operation of TeraGen, TeraSort and TeraValidate in each of the BDCS environments across directories replicated by WANdisco Fusion. This validation demonstrated that single-copy consistency of data was maintained by WANdisco Fusion in conjunction with Oracle BDCS regardless of which of the two test BDCS environments was operating against the datasets.

Success Criteria

Success criteria for the tests covered the standard set of post-install validation checks that apply to a regular WANdisco Fusion installation, demonstrating equivalent functionality and correctness from its application in an Oracle Big Data Cloud Services deployment.

Additionally, a base level of performance tests has been performed to ensure that replication performance with Oracle BDCS is equivalent to that which can be achieved in similar deployments. Please refer to the Oracle MAA Whitepaper for WANdisco Fusion for details.

Summary

- » Oracle Big Data Cloud Service provisions a Secured Enterprise Class highly available Hadoop cluster which runs Spark on Yarn in the Oracle Elastic Cloud with only a few simple clicks.
- » Oracle Big Data Cloud Service provides high performance Edge Nodes for installation of any client application that requires high performance connectivity with the cluster. Oracle Big Data Cloud Service also provides the 'Bursting' option which allows for a transient increase of CPU and Memory resources on demand on an hourly basis.
- » WANdisco Fusion supports active replication of data across Oracle Big Data Cloud Service deployments for continuous data availability, cross-location consistency, disaster recovery requirements, migration and upgrade scenarios.
- » WANdisco Fusion can be deployed across on-premise and cloud deployments of Hadoop, including with Oracle Big Data Cloud Services. This enhances greatly the ability to use Oracle BDCS for cloud bursting and specialized operations.



Oracle Corporation, World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries
Phone: +1.650.506.7000
Fax: +1.650.506.7200

CONNECT WITH US



blogs.oracle.com/oracle



facebook.com/oracle



twitter.com/oracle



oracle.com

Integrated Cloud Applications & Platform Services

Copyright © 2017, Oracle and/or its affiliates. All rights reserved. This document is provided *for* information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

WANdisco Fusion on Oracle Big Data Cloud Service
July 2017
Authors: Rick Pandya, Paul Scott-Murphy, Steve Kilgore



Oracle is committed to developing practices and products that help protect the environment

Integrated Cloud Applications & Platform Services