

Oracle® Big Data SQL Installation Guide



Release 3 (3.2)

F38383-01

January 2021

The Oracle logo, consisting of a solid red square with the word "ORACLE" in white, uppercase, sans-serif font centered within it.

ORACLE®

Copyright © 2012, 2021, Oracle and/or its affiliates.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software" or "commercial computer software documentation" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

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

Intel and Intel Inside are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Epyc, and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Contents

Preface

Audience	vi
Related Documents	vi
Conventions	vi
Backus-Naur Form Syntax	vii
Changes in Oracle Big Data SQL 3.2	vii

1 Introduction

1.1 Supported System Combinations	1-1
1.2 Oracle Big Data SQL Master Compatibility Matrix	1-2
1.3 Prerequisites for Installation on the Hadoop Cluster	1-2
1.4 Prerequisites for Installation on Oracle Database Nodes	1-6
1.5 Downloading Oracle Big Data SQL	1-8
1.6 Upgrading From a Prior Release of Oracle Big Data SQL	1-9
1.7 Important Terms and Concepts	1-10
1.8 Installation Overview	1-12
1.9 Post-Installation Checks	1-17
1.10 Using the Installation Quick Reference	1-18

2 Installing or Upgrading the Hadoop Side of Oracle Big Data SQL

2.1 Before You Start	2-1
2.2 About the Jaguar Utility	2-3
2.2.1 Jaguar Configuration Parameter and Command Reference	2-4
2.3 Steps for Installing on the Hadoop Cluster	2-14
2.4 Special Installation Procedures for Oracle Big Data Appliance 4.10	2-16

3 Installing or Upgrading the Oracle Database Side of Oracle Big Data SQL

3.1 Before You Start the Database-Side Installation	3-1
3.1.1 Potential Requirement to Restart Grid Infrastructure	3-2

3.1.1.1	Understanding When Grid or Database Restart is Required	3-3
3.1.2	Special Considerations When a System Under Grid Infrastructure has Multiple Network Interfaces of the Same Type	3-3
3.2	About the Database-Side Installation Directory	3-5
3.3	Steps for Installing on Oracle Database Nodes	3-6
3.3.1	Command Line Parameter Reference for bds-database-install.sh	3-9
3.4	Granting User Access	3-11
4	Expanding or Shrinking an Installation	
4.1	Adding or Removing Oracle Big Data SQL on Hadoop Cluster Nodes	4-1
4.2	Adding or Removing Oracle Big Data SQL on Oracle Database Nodes	4-2
5	Reconfiguring an Installation	
5.1	Reconfiguring an Existing Oracle Big Data SQL Installation	5-3
6	Uninstalling Oracle Big Data SQL	
6.1	General Guidelines for Removing the Software	6-1
6.2	Uninstalling From an Oracle Database Server	6-1
6.3	Uninstalling From a Hadoop Cluster	6-2
7	Securing Big Data SQL	
7.1	Security Overview	7-1
7.2	Big Data SQL Communications and Secure Hadoop Clusters	7-2
7.3	Configuring Oracle Big Data SQL in a Kerberos-Secured Environment	7-2
7.3.1	Enabling Oracle Big Data SQL Access to a Kerberized Cluster	7-2
7.3.2	Installing a Kerberos Client on the Oracle Database Nodes	7-4
7.4	Using Oracle Secure External Password Store to Manage Database access for Oracle Big Data SQL	7-5
7.5	About Data Security on Oracle Big Data Appliance	7-5
7.6	Authentication Between Oracle Database and Oracle Big Data SQL Offload Cell Server Processes	7-6
7.7	The Multi-User Authorization Model	7-6
8	Additional Tools Installed	
8.1	Copy to Hadoop and OSH	8-1
8.1.1	Completing the OSH Configuration on Oracle Database Nodes	8-2
8.1.2	Completing the OSH Configuration on the Hadoop Cluster	8-4

A	Installation Quick Reference	
B	bds-config.json Configuration Example	
C	Oracle Big Data SQL Installation Examples	
D	Determining the Correct Software Version and Composing the Download Paths for Hadoop Clients	
E	Oracle Big Data SQL Licensing	
E.1	ANTLR 4.7	E-1
E.2	Apache Commons Exec 1.3	E-2
E.3	Apache Licensed Code	E-2
E.4	Apache License	E-2
F	Change History for Previous Releases	
F.1	Changes in Oracle Big Data SQL 3.1	F-1
F.2	Changes in Oracle Big Data SQL 3.0.1	F-3
	Index	

Preface

The *Oracle Big Data SQL User's Guide* describes how to use and manage the Oracle Big Data SQL product.

Audience

This guide is intended for administrators and users of Oracle Big Data SQL, including:

- Application developers
- Data analysts
- Data scientists
- Database administrators
- System administrators

The guide assumes that the reader has basic knowledge of Oracle Database single-node and multinode systems, the Hadoop framework, the Linux operating system, and networking concepts.

Related Documents

See the *Oracle Big Data SQL User's Guide* for instructions on using the product.

The following publications provide information about the use of Oracle Big Data SQL with the Oracle Big Data Appliance and Oracle Big Data Connectors:

- *Oracle Big Data Appliance Owner's Guide*
- *Oracle Big Data Appliance Software User's Guide*
- *Oracle Big Data Connectors User's Guide*

You can find more information about Oracle's Big Data solutions and Oracle Database at the [Oracle Help Center](#)

For more information on Hortonworks HDP and Ambari, refer to the Hortonworks documentation site at <http://docs.hortonworks.com/index.html>.

For more information on Cloudera CDH and Configuration Manager, see <http://www.cloudera.com/documentation.html>

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.
# prompt	The pound (#) prompt indicates a command that is run as the Linux root user.

Backus-Naur Form Syntax

The syntax in this reference is presented in a simple variation of Backus-Naur Form (BNF) that uses the following symbols and conventions:

Symbol or Convention	Description
[]	Brackets enclose optional items.
{ }	Braces enclose a choice of items, only one of which is required.
	A vertical bar separates alternatives within brackets or braces.
...	Ellipses indicate that the preceding syntactic element can be repeated.
delimiters	Delimiters other than brackets, braces, and vertical bars must be entered as shown.
boldface	Words appearing in boldface are keywords. They must be typed as shown. (Keywords are case-sensitive in some, but not all, operating systems.) Words that are not in boldface are placeholders for which you must substitute a name or value.

Changes in Oracle Big Data SQL 3.2

Oracle Big Data SQL Release 3.2 includes major improvements in performance, secure network connectivity, authentication, and user administration, as well as installation and configuration.

JSON CLOB Predicate Pushdown

Much improved filtering and parsing of JSON CLOB data in Hadoop enables Oracle Big Data SQL to push more processing for these large objects down to the Hadoop cluster. JSON Data can now be filtered on the Oracle Big Data SQL cells in Hadoop for CLOB columns up to 1 MB, depending on character set of the input document. The eligible JSON filter expressions for storage layer evaluation include simplified syntax, JSON_VALUE, and JSON_QUERY. In addition, Oracle Big Data SQL can project up to 32 KB of CLOB data from select list expression evaluation in Hadoop to Oracle Database. Processing falls back to Oracle Database only when column sizes exceed these two values.

Customers can disable or re-enable this functionality to suit their own needs.

In Release 3.2, this enhancement currently applies only to JSON expressions returning CLOB data. The same support will be provided for other CLOB types (such as `substr` and `instr`) as well as for BLOB data in a future release.

 **Note:**

The new JSON CLOB predicate pushdown functionality requires Oracle Database version 12.1.0.2.180417 or greater, as well as the following patches:

- The April 2018 Proactive DBBP (Database Bundle Patch). This is patch 27486326.
- The one-off patch 27767148.

Install the one-off patch on all database compute nodes.

The one-off patch 26170659, which is required on top of earlier DBBPs, is not required on top of the April DBBP.

This functionality is not available through the January 2018 and August 2017 Proactive DBBPs

See the [Oracle Big Data SQL Master Compatibility Matrix](#) (Doc ID 2119369.1 in [My Oracle Support](#)) for the most up-to-date information on software version and patch requirements.

Support for Querying Kafka Topics

Release 3.2 provides Hive and Oracle Big Data SQL the ability to query Kafka topics via a new Hive storage handler. You can use this storage handler to create external Hive tables backed by data residing in Kafka. Oracle Big Data SQL or Hive can then query the Kafka data through the external tables. The Kafka key, value, offset, topic name, and partition id are mapped to Hive columns. You can explicitly designate the offset for each topic/partition pair, otherwise the offset will start from the earliest offset in the topic and end with the latest offset in the topic for each partition.

Improved Processing of Parquet Files

Oracle has introduced its own Parquet reader for processing data in Parquet format. This new reader provides significant performance and resource utilization improvements over the existing Hive Parquet driver. These include:

- More intelligent column processing retrieval. The reader uses “lazy materialization” to process only columns with rows that satisfy the filter, thereby improving I/O.
- Leveraging of dictionaries during filter predicate processing to improve CPU usage.
- Streamlined data conversion, which also contributes to more efficient CPU usage.

The Big Data SQL installation enables the Oracle's Parquet reader by default. You have the option to disable it and revert to the generic Parquet reader.

Multi-User Authorization

In previous releases of Oracle Big Data SQL, all queries against Hadoop and Hive data are executed as the `oracle` user and there is no option to change users. Although

`oracle` is still the underlying user in all cases, Oracle Big Data SQL 3.2 now uses Hadoop Secure Impersonation to direct the `oracle` account to execute tasks on behalf of other designated users. This enables HDFS data access based on the user that is currently executing the query, rather than the singular `oracle` user.

Administrators set up the rules for identifying the query user. They can provide rules for identifying the currently connected user and mapping the connected user to the user that is impersonated. Because there are numerous ways in which users can connect to Oracle Database, this user may be a database user, a user sourced from LDAP, from Kerberos, or a user from other sources. Authorization rules on the files apply for that user and HDFS auditing identifies the actual user running the query.

See Also:

Administration for Multi-User Authorization is done through the `DBMS_BDSQL` PL/SQL Package, which is documented in the *Oracle Big Data SQL User's Guide*.

Authentication Between Oracle Database and Oracle Big Data SQL Cells

This authentication is between Oracle Database and the Big Data SQL cells on the Hadoop cluster, facilitating secure communication. The Database Authentication enhancement provides a safeguard against impersonation attacks, in which a rogue service attempts to connect to the Oracle Big Data offload server process running on a cluster node.

Kerberos Ticket Renewal Automation

On a Kerberos-secured network you can configure the installation to set up automated Kerberos ticket renewal for the `oracle` account used by Oracle Big Data SQL. This is done for both the Hadoop cluster and Oracle Database sides of the installation. You must provide the principal name and the path to the keytab file in the `bds-config.json` configuration file. A template is provided in the configuration file:

```
"kerberos" : {
  "principal" : "oracle/mycluster@MY.DOMAIN.COM",
  "keytab" : "/home/oracle/security/oracle.keytab"
}
```

If you provide the Kerberos parameters in the configuration file, then Oracle Big Data SQL installation sets up cron jobs on both the Hadoop cluster and Oracle Database servers. These jobs renew the Kerberos tickets for the principal once per day.

The principal and keytab file must already exist.

Automatic Upgrade

The current release can now be installed over an earlier release with no need to remove the older software on either the Hadoop or Oracle Database side. The previous installation is upgraded to the current release level.

Common Installation Bundle for all Platforms

In previous releases, customers needed to unpack the Oracle Big Data SQL installation bundle and choose the correct package for their Hadoop system (CDH or HDP). Now the bundle contains a single installation package that works for all supported Hadoop systems.

Simpler and Faster Installation with the new “Jaguar” Installer

The Jaguar installer replaces `setup-bds.sh`, the installer in previous releases. Jaguar includes these changes:

- **Automatic Check for Installation Prerequisites on Hadoop Nodes**
Jaguar checks for installation readiness on each Hadoop DataNode and reports any missing prerequisites.
- **No Need to Manually Generate the Database-Side Installation Bundle**
The database-side installation bundle that previously was manually generated by the customer can now be generated automatically. You still need to copy the bundle to the Oracle Database nodes and install it.
- **Faster Overall Installation Time on the Hadoop Side**
Installation time will vary, but on the Hadoop Side the installation may take approximately eight minutes if all resources are local, possibly 20 minutes if Hadoop clients must be downloaded from the Internet, depending on download speed.
- **Prerequisite Apache Services on CDH can now be Installed as Either Packages or Parcels**
Previously on CDH systems, the Oracle Big Data SQL installation required that the HDFS, YARN, and HIVE components had been installed as parcels. These components can now be installed on CDH as either packages or parcels. There is no change for HDP, where they must be installed as stacks.

Note:

On CDH systems, if the Hadoop services required by Oracle Big Data SQL are installed as packages, be sure that they are installed from within Cloudera Manager. Otherwise, Cloudera Manager will not be able to manage these services. This is not an issue with parcels.

- **In the CLI, the Jaguar utility Replaces `./setup-bds`**
The Jaguar utility is now the primary tool for Hadoop-side installation, de-installation, and configuration changes, as in these examples:


```
# ./jaguar install bds-config.json
# ./jaguar reconfigure bds-config.json
# ./jaguar uninstall bds-config.json
```
- **The Default Configuration File Name is `bds-config.json`, but Alternate File Names are Also Accepted**

You can now drop the explicit `bds-config.json` argument and allow the installer default to `bds-config.json`, as in the first example below. You can also specify an alternate configuration file of any name, though it must adhere to the same internal format as `bds-config.json` and should be given the `.json` file type.

```
# ./jaguar install
# ./jaguar install cluster2-config.json
```

You can create configurations files with settings that are tailored to the requirements of each cluster. For example, you may want to apply different security parameters to Oracle Big Data SQL installations on test and production clusters.

- **Configuration Parameters Have Changed Significantly**

Users of previous releases will see that the Jaguar configuration file includes a number of new parameters. Most of them are “optional” in the sense that they are not uniformly required, although your particular installation may require some of them. See the Related Links section below for links to the table of installer parameters as well as an example of a configuration file that uses all available parameters.

- **New `updatenodes` Command for Easier Maintenance**

Oracle Big Data SQL must be installed on each Hadoop cluster node that is provisioned with the DataNode role. It has no function on nodes where DataNode is not present. The new Jaguar utility includes the `updatenodes` command which scans the cluster for instances of the DataNode within the cluster. If the DataNode role has been removed or relocated, or if nodes provisioned with the DataNode have been added or removed, then the script installs or uninstalls Oracle Big Data SQL components from nodes as needed.

- **An Extra Installation Step is Required to Enable Some Security Features**

If you choose to enable Database Authentication between Oracle Database and Oracle Big Data SQL cells in the Hadoop cluster, or, Hadoop Secure Impersonation, then an additional “Database Acknowledge” step is required. In this process, the installation on the database server generates a ZIP file of configuration information that you must copy back to the Hadoop cluster management server for processing.

- **On the Database Side, Connections to Clusters are no Longer Classified as Primary and Secondary.**

An Oracle Database system can have Oracle Big Data SQL connections to multiple Hadoop clusters. In previous releases, the first these connections was considered the primary (and had to be uninstalled last) and the others were secondary. In the current release, management of multiple installation is simpler and `--uninstall-as-primary` and `--uninstall-as-secondary` parameters of the database-side installer are obsolete. However there is now a default cluster. The *Important Terms and Concepts* section of this guide explains the significance of the default cluster.

Support for Oracle Tablespaces in HDFS Extended to Include All Non-System Permanent Tablespaces

Previous releases supported the move of permanent online tablespaces only to HDFS. This functionality now supports online, read-only, as well as offline permanent tablespaces.

Important Change in Behavior of the “mtactl start” Command

Oracle Big Data SQL 3.1 introduced the option to install Oracle Big Data SQL on servers where Oracle Grid Infrastructure is not present. In these environments, you can use the start subcommand of the mtactl utility (`mtactl start`) to start the MTA (Multi-Threaded Agent) extproc.

Note that in the current release, the `mtactl start` command works differently from the original Release 3.1 implementation.

- Current behavior: `mtactl start` starts an MTA extproc using the init parameter values that are stored in the repository. It uses the default values only if the repository does not exist.
- Previous behavior (Oracle Big Data SQL 3.1): `mtactl start` always uses the default init parameters regardless of whether or not init parameter values are stored in the repository.

 **See Also:**

Resource Requirements

8 CPU cores and 12 GB of RAM are now recommended for each node of the Hadoop cluster.

There are some `sysctl` settings related to kernel, swap, core memory, and socket buffer size that are strongly recommended for optimal performance. These are part of the installation prerequisites explained in Chapter 2 of the installation guide.

Related Topics

- [bds-config.json Configuration Example](#)
The following is an example of a fully-populated `bds-config.json` file, which includes all available configuration parameters.
- [Oracle Big Data SQL Installation Examples](#)
The following are samples of the console output for the Oracle Big Data SQL installation.

1

Introduction

This guide describes how to install Oracle Big Data SQL, how to reconfigure or extend the installation to accommodate changes in the environment, and, if necessary, how to uninstall the software.

This installation is done in phases. The first two phases are:

- Installation on the node of the Hadoop cluster where the cluster management server is running.
- Installation on each node of the Oracle Database system.

If you choose to enable new security features available in Release 3.2, then there is an additional third phase in which you activate the security features.

The two systems must be networked together via Ethernet or InfiniBand. (Connectivity to Oracle SuperCluster is InfiniBand only).

Note:

For Ethernet connections between Oracle Database and the Hadoop cluster, Oracle recommends 10 Gb/s Ethernet.

The installation process starts on the Hadoop system, where you install the software manually on one node only (the node running the cluster management software). Oracle Big Data SQL leverages the administration facilities of the cluster management software to automatically propagate the installation to all DataNodes in the cluster.

The package that you install on the Hadoop side also generates an Oracle Big Data SQL installation package for your Oracle Database system. After the Hadoop-side installation is complete, copy this package to all nodes of the Oracle Database system, unpack it, and install it using the instructions in this guide. If you have enabled Database Authentication or Hadoop Secure Impersonation, you then perform the third installation step.

1.1 Supported System Combinations

Oracle Big Data SQL supports connectivity between a number of Oracle Engineered Systems and commodity servers.

The current release supports Oracle Big Data SQL connectivity for the following Oracle Database platforms/Hadoop system combinations:

- Oracle Database on commodity servers with Oracle Big Data Appliance.
- Oracle Database on commodity servers with commodity Hadoop systems.
- Oracle Exadata Database Machine with Oracle Big Data Appliance.
- Oracle Exadata Database Machine with commodity Hadoop systems.

Oracle SPARC SuperCluster support is not available for Oracle Big Data SQL 3.2 at this time. Release 3.1 does support this platform.

 **Note:**

The phrase “Oracle Database on commodity systems” refers to Oracle Database hosts that are not the Oracle Exadata Database Machine. Commodity database systems may be either Oracle Linux or RHEL-based. “Commodity Hadoop systems” refers to Hortonworks HDP systems and to Cloudera CDH-based systems other than Oracle Big Data Appliance.

1.2 Oracle Big Data SQL Master Compatibility Matrix

See the [Oracle Big Data SQL Master Compatibility Matrix](#) (Doc ID 2119369.1 in [My Oracle Support](#)) for up-to-date information on Big Data SQL compatibility with the following:

- Oracle Engineered Systems.
- Other systems.
- Linux OS distributions and versions.
- Hadoop distributions.
- Oracle Database releases, including required patches.

1.3 Prerequisites for Installation on the Hadoop Cluster

The following active services, installed packages, and available system tools are prerequisites to the Oracle Big Data SQL installation. These prerequisites apply to all DataNodes of the cluster.

The Oracle Big Data SQL installer checks all prerequisites before beginning the installation and reports any missing requirements on each node.

Platform requirements, such as supported Linux distributions and versions, as well as supported Oracle Database releases and required patches and are not listed here. See the [Oracle Big Data SQL Master Compatibility Matrix](#) (Doc ID 2119369.1 in [My Oracle Support](#)) for this information.

 **Important:**

Oracle Big Data SQL 3.2 does not support single user mode for Cloudera clusters.

Services Running

These Apache Hadoop services must be running on the cluster.

- HDFS
- YARN

- Hive

You do not need to take any extra steps to ensure that the correct HDFS and Hive clients URLs are specified in the database-side installation bundle.

The Apache Hadoop services listed above may be installed as parcels or packages on Cloudera CDH and as stacks on Hortonworks HDP.

! Important:

On CDH, if you install the Hadoop services required by Oracle Big Data SQL as packages, be sure that they are installed from within CM. Otherwise, CM will not be able to manage them. This is not an issue with parcel-based installation.

Packages

The following packages must be pre-installed on all Hadoop cluster nodes before installing Oracle Big Data SQL. These packages are already installed on versions of Oracle Big Data Appliance supported by Oracle Big Data SQL.

- Oracle JDK version 1.7 or later is required on Oracle Big Data Appliance. Non-Oracle commodity Hadoop servers must also use the Oracle JDK.
- dmidecode
- net-snmp, net-snmp-utils
- perl
PERL LibXML – 1.7.0 or higher, e.g. perl-XML-LibXML-1.70-5.el6.x86_64.rpm
perl-libwww-perl, perl-libxml-perl, perl-Time-HiRes, perl-libs, perl-XML-SAX

The `yum` utility is the recommended method for installing these packages:

```
# yum -y install dmidecode net-snmp net-snmp-utils perl perl-libs perl-  
Time-HiRes perl-libwww-perl perl-libxml-perl perl-XML-LibXML perl-XML-  
SAX
```

Conditional Requirements

- `perl-Env` is required for systems running Oracle Linux 7 or RHEL 7 only.

```
# yum -y install perl-Env
```

System Tools

- curl
- gcc
- libaio
- rpm
- scp

- tar
- unzip
- wget
- yum
- zip

The libaio libraries must be installed on each Hadoop cluster node:

```
# yum install -y libaio gcc
```

Environment Settings

The following environment settings are required prior to the installation.

- NTP enabled
- The path to the Java binaries must exist in `/usr/java/latest`.
- The path `/usr/java/default` must exist and must point to `/usr/java/latest`.
- Check that these system settings meet the requirements indicated. All of these settings can be temporarily set using the `sysctl c` ommand. To set them permanently, add or update them in `/etc/sysctl.conf`.

- `kernel.shmmax` and `kernel.shmall` must each be greater than physical memory size.
- `kernel.shmall` and `kernel.shmmax` values should fit this formula:

```
kernel.shmmax = kernel.shmall * PAGE_SIZE
```

(You can determine `PAGE_SIZE` with `# getconf PAGE_SIZE`.)

- `vm.swappiness=10`

If cell startup fails with an error indicating that the `SHMALL` limit has been exceeded, then increase the memory allocation and restart Oracle Big Data SQL.

- socket buffer size:

```
net.core.rmem_default >= 4194304
net.core.rmem_max >= 4194304
net.core.wmem_default >= 4194304
net.core.wmem_max >= 4194304
```

Proxy-Related Requirements:

- The installation process requires Internet access in order to download some packages from Cloudera or Hortonworks sites. If a proxy is needed for this access, then either ensure that the following are set as Linux environment variables, or, enable the equivalent parameters in the installer configuration file, `bds-config.json`)
 - `http_proxy` and `https_proxy`
 - `no_proxy`

Set `no_proxy` to include the following: "localhost,127.0.0.1,<Comma-separated list of the hostnames in the cluster (in FQDN format).>".

- On Cloudera CDH, clear any proxy settings in Cloudera Manager administration before running the installation. You can restore them after running the script that creates the database-side installation bundle (`bds-database-create-bundle.sh`).

See Also:

[Configuration Parameters for `bds-config.json`](#) describes the use of `http_proxy`, `https_proxy`, and other parameters in the installer configuration file.

Hardware Requirements

- Physical memory \geq 12GB
- Number of cores \geq 8

Networking

If Hadoop traffic is over VLANs, all DataNodes must be on the same VLAN.

Python 2.7 < 3.0 (for the Oracle Big Data SQL Installer)

The Oracle Big Data SQL installer requires Python (version 2.7 or greater, but less than 3.0) locally on the node where you run the installer. This should be the same node where the cluster management service (CM or Ambari) is running.

If any installation of Python in the version ranged described above is already present, you can use it to run the Oracle Big Data installer.

If an earlier version than Python 2.7.0 is already installed on the cluster management server and you need to avoid overwriting this existing installation, you can add Python 2.7.x as a secondary installation.

Restriction:

On Oracle Big Data Appliance do not overwrite or update the pre-installed Python release. This restriction may also apply other supported Hadoop platforms. Consult the documentation for the CDH or HDP platform you are using.

On Oracle Linux 5, add Python 2.7 as a secondary installation. On Oracle Linux 6, both Python 2.6 and 2.7 are pre-installed and you should use the provided version of Python 2.7 for the installer. Check whether the default interpreter is Python 2.6 or 2.7. To run the Oracle Big Data SQL installer, you may need to invoke Python 2.7 explicitly. On Oracle Big Data Appliance, SCL is installed so you can use to enable version 2.7 for the shell as in this example:

```
[root@myclusteradminserver:BDSjaguar] # scl enable python27 "./jaguar
install bds-config.json"
```

Below is a procedure for adding the Python 2.7.5 as a secondary installation.

 **Tip:**

If you manually install Python, first ensure that the `openssl-devel` package is installed:

```
# yum install -y openssl-devel
```

If you create a secondary installation of Python, it is strongly recommended that you apply Python update regularly to include new security fixes. On Oracle Big Data Appliance, do not update the mammoth-installed Python unless directed to do so by Oracle.

```
# pyversion=2.7.5
# cd /tmp/
# mkdir py_install
# cd py_install
# wget https://www.python.org/static/files/pubkeys.txt
# gpg --import pubkeys.txt
# wget https://www.python.org/ftp/python/$pyversion/
Python-$pyversion.tgz.asc
# wget https://www.python.org/ftp/python/$pyversion/
Python-$pyversion.tgz
# gpg --verify Python-$pyversion.tgz.asc Python-$pyversion.tgz
# tar xfvz Python-$pyversion.tgz
# cd Python-$pyversion
# ./configure --prefix=/usr/local/python/2.7
# make
# mkdir -p /usr/local/python/2.7
# make install
# export PATH=/usr/local/python/2.7/bin:$PATH
```

1.4 Prerequisites for Installation on Oracle Database Nodes

Installation prerequisites vary, depending on type of Hadoop system and Oracle Database system where Oracle Big Data SQL will be installed.

See the *Oracle Big Data SQL Master Compatibility Matrix* (Doc ID 2119369.1) in [My Oracle Support](#) for supported Linux distributions, Oracle Database release levels, and required patches.

 **Note:**

Be sure that the correct Bundle Patch and any one-off patches identified in the Compatibility Matrix have been pre-applied before starting this installation.

Before you begin the installation, review the additional environmental and user access requirements described below.

Packages Required for Kerberos

If you are installing on a Kerberos-enabled Oracle Database System, these package must be pre-installed:

- krb5-workstation
- krb5-libs

Packages for the “Oracle Tablespaces in HDFS” Feature

Oracle Big Data SQL provides a method to store Oracle Database tablespaces in the Hadoop HDFS file system. The following RPMs must be installed:

- fuse
- fuse-libs

```
# yum -y install fuse fuse-libs
```

Required Environment Variables

The following are always required. Be sure that these environment variables are set correctly.

- ORACLE_SID
- ORACLE_HOME

 **Note:**

GI_HOME (which was required in Oracle Big Data SQL 3.1 and earlier) is no longer required.

Required Credentials

- Oracle Database owner credentials (The owner is usually the `oracle` Linux account.)

Big Data SQL 3.2 is installed as an add-on to Oracle Database. Tasks related directly to database instance are performed through database owner account (`oracle` or other).

- Grid user credentials

In some cases where Grid infrastructure is present, it must be restarted. If the system uses Grid then you should have the Grid user credentials on hand in case a restart is required.

The Linux users `grid` and `oracle` (or other database owner) must both in the same group (usually `oinstall`). This user requires permission to read all files owned by the `grid` user and vice versa.

All Oracle Big Data SQL files and directories are owned by the `oracle:oinstall` user and group.

1.5 Downloading Oracle Big Data SQL

You can download Oracle Big Data SQL from the [Oracle Software Delivery Cloud](#) (also known as “eDelivery”).

Note to Oracle Big Data Appliance Customers:

On Oracle Big Data Appliance 4.10, do not enable the Oracle Big Data SQL 3.2 package that is included with the Oracle Big Data Appliance bundle. Download the “standalone” version from the Oracle Software Delivery Cloud instead. The Oracle Big Data Appliance 4.10 installation bundle includes an early-release version of Oracle Big Data SQL 3.2. The Oracle Big Data SQL 3.2 package available for download on the Oracle Software Delivery Cloud is a newer image and includes important updates.

1. Log in as root to the Hadoop node that hosts the cluster management server (CDH or Ambari). Create a new directory or choose an existing one to be the installation source directory.

Note:

Throughout the instructions, this guide uses the placeholder *<Big Data SQL Install Directory>* to refer to the location where you unzip the bundle. This is the working directory from which you configure and install Oracle Big Data SQL. You should secure this directory against accidental or unauthorized modification or deletion. The primary file to protect is your installation configuration file (by default, `bds-config.json`). As you customize the configuration to your needs, this file becomes the record of the state of the installation. It is useful for recovery purposes and as a basis for further changes.

2. Log on to the Oracle Software Delivery Cloud.
3. Search for “Oracle Big Data SQL”.
4. Select Oracle Big Data SQL 3.2.0.0 for Linux x86-64.

The actual version available may be greater than 3.2 (as in 3.2.n.n). The same bundle is compatible for with all supported systems.

5. Read and agree to the Oracle Standard Terms and Restrictions.
6. Download and unzip the bundle.

When you unzip the bundle, the *<Big Data SQL Install Directory>/BDSjaguar* directory will include the content listed the table below, with the exception of the `db-bundles` and `dbkeys` directories. These do not appear until after you have run an installation that produces the output for these directories.

Table 1-1 Oracle Big Data SQL Product Extracted Bundle Inventory

Directory or File	Description
jaguar	Hadoop cluster-side installation and configuration utility.
example-bda-config.json example-cdh-config.json example-ethernet- config.json example-infiniband- config.json example-kerberos- config.json example-localrepos- config.json example-named- config.json example-subnetwork- config.json example-unsecure- config.json	These configuration file examples show how to apply the various parameter options. The file <code>example-bda-config.json</code> includes all available configuration parameters.
deployment_manager/*	Python webserver for deployment utilities).
deployment_server/*	Deployment server Python package utility.
bdsrepo/*	Resources to build Cloudera parcels or Ambari stacks for the installation.
db-bundles/	When you run <code>jaguar</code> to install Oracle Big Data SQL on the Hadoop cluster, the installation bundle (or bundles) that you must copy over to the database side are generated and stored here.
dbkeys/	If you choose to use the Database Authentication security feature, then as part of the process, the Jaguar install operation generates a <code>.reqkey</code> file (a "request key") for setup of authentication with a specific database. All such request keys are stored in this directory. Database Authentication is enabled in the configuration by default, but there are further actions you need to take to in order to complete the Database Authentication setup process. These are described later in this document.

1.6 Upgrading From a Prior Release of Oracle Big Data SQL

On the Oracle Database side, Oracle Big Data SQL can now be installed over a previous release with no need to remove the older software. The install script automatically detects and upgrades an older version of the software.

Upgrading the Oracle Database Side of the Installation

On the database side, you need to perform the installation only once to upgrade the database side for any clusters connected to that particular database. This is because the installations on the database side are not entirely separate. They share the same

set of Oracle Big Data SQL binaries. This results in a convenience for you – if you upgrade one installation on a database instance then you have effectively upgraded the database side of all installations on that database instance.

Upgrading the Hadoop Cluster Side of the Installation

If existing Oracle Big Data SQL installations on the Hadoop side are not upgraded, these installations will continue to work with the new Oracle Big Data SQL binaries on the database side, but will not have access to the new features in this release.

1.7 Important Terms and Concepts

These are special terms and concepts in the Oracle Big Data SQL installation.

Oracle Big Data SQL Installation Directory

On both the Hadoop side and database side of the installation, the directory where you unpack the installation bundle is not a temporary directory which you can delete after running the installer. These directories are staging areas for any future changes to the configuration. You should not delete them and may want to secure them against accidental deletion.

Database Authentication Keys

Database Authentication uses a key that must be identical on both sides of the installation (the Hadoop cluster and Oracle Database). The first part of the key is created on the cluster side and stored in the `.reqkey` file. This file is consumed only once on the database side, to connect the first Hadoop cluster to the database. Subsequent cluster installations use the configured key and the `.reqkey` file is no longer required. The full key (which is completed on the database side) is stored in an `.ackkey` file. This key is included in the part of the ZIP file created by the database-side installation and must be copied back to the Hadoop cluster by the user.

Request Key

By default, the Database Authentication feature is enabled in the configuration. (You can disable it by setting the parameter `database_auth_enabled` to “false” in the configuration file.) When this setting is true, then the `Jaguar install`, `reconfigure` and `updatenodes` operations can all generate a `request key` (stored in a file with the extension `.reqkey`). This key is part of a unique GUID-key pair used for Database Authorization. This GUID-key pair is generated during the database side of the installation. The `Jaguar` operation creates a request key if the command line includes the `--requestdb` command line parameter along with a single database name (or a comma separated list of names). In this example, the `install` operation creates three keys, one for each of three different databases:

```
# ./jaguar --requestdb orcl,testdb,proddb install
```

The operation creates the request keys files in the directory `<Oracle Big Data SQL install directory>/BDSJaguar/dbkeys`. In this example, `Jaguar install` would generate these request key files:

```
orcl.reqkey  
testdb.reqkey  
proddb.reqkey
```

Prior to the database side of the installation, you copy request key to the database node and into the path of the database-side installer, which at runtime generates the GUID-key pair.

Acknowledge Key

After you copy a request key into the database-side installation directory, then when you run the database-side Oracle Big Data SQL installer it generates a corresponding *acknowledge key*. The acknowledge key is the original request key, paired with a GUID. This key is stored in a file that is included in a ZIP archive along with other information that must be returned to the Hadoop cluster by the user.

Database Request Operation (*databasereq*)

The Jaguar *databasereq* operation is “standalone” way to generate a request key. It lets you create one or more request keys without performing an *install*, *reconfigure*, or *updatenodes* operation:

```
# ./jaguar --requestdb <database name list> databasereq {configuration  
file | null}
```

Database Acknowledge ZIP File

If Database Authentication, or Hadoop Secure Impersonation is enabled for the configuration, then the database-side installer creates a ZIP bundle configuration information. If Database Authentication is enabled, this bundle includes the acknowledge key file. Information required for Hadoop Secure Impersonation is also included if that option was enabled. Copy this ZIP file back to `/opt/oracle/DM/databases/conf` on the Hadoop cluster management server for processing.

Database Acknowledge is a third phase of the installation and is performed only when any of the three security features cited above are enabled.

Database Acknowledge Operation (*databaseack*)

If you have opted to enable any or all of three new security features (Database Authentication, or Hadoop Secure Impersonation), then after copying the Database Acknowledge ZIP file back to the Hadoop cluster, run the Jaguar *Database Acknowledge* operation.

The setup process for these features is a “round trip” that starts on the Hadoop cluster management server, where you set the security directives in the configuration file and run Jaguar, to the Oracle Database system where you run the database-side installation, and back to the Hadoop cluster management server where you return a copy of the ZIP file generated by the database-side installation. The last step is when you run *databaseack*, the Database Acknowledge operation described in the outline below. Database Acknowledge completes the setup of these security features.

Default Cluster

The *default cluster* is the first Oracle Big Data SQL connection installed on an Oracle Database. In this context, the term *default cluster* refers to the installation directory on the database node where the connection to the Hadoop cluster is established. It does not literally refer to the Hadoop cluster itself. Each connection between a Hadoop cluster and a database has its own installation directory on the database node.

An important aspect of the default cluster is that the setting for Hadoop Secure Impersonation in the default cluster determines that setting for all other cluster

connections to a given database. If you run a Jaguar `reconfigure` operation some time after installation and use it to turn Hadoop Secure Impersonation in the default cluster on or off, this change is effective for *all* clusters associated with the database.

If you perform installations to add additional clusters, the first cluster remains the default. If the default cluster is uninstalled, then next one (in chronological order of installation) becomes the default.

1.8 Installation Overview

The Oracle Big Data SQL software must be installed on all Hadoop cluster DataNodes and all Oracle Database compute nodes.

Important: About Service Restarts

On the Hadoop-side installation, the following restarts may occur.

- Cloudera Configuration Manager (or Ambari) may be restarted. This in itself does not interrupt any services.
- Hive, YARN, and any other services that have a dependency on Hive or YARN (such as Impala) are restarted.

The Hive libraries parameter is updated in order to include Oracle Big Data SQL JARs. On Cloudera installations, if the YARN Resource Manager is enabled, then it is restarted in order to set cgroup memory limit for Oracle Big Data SQL and the other Hadoop services. On Oracle Big Data Appliance, the YARN Resource Manager is always enabled and therefore always restarted.

On the Oracle Database server(s), the installation may require a database and/or Oracle Grid infrastructure restart in environments where updates are required to Oracle Big Data SQL cell settings on the Grid nodes. See [Potential Requirement to Restart Grid Infrastructure](#) for details.

If A Previous Version of Oracle Big Data SQL is Already Installed

On commodity Hadoop systems (those other than Oracle Big Data Appliance) the installer automatically uninstalls any previous release from the Hadoop cluster.

You can install Oracle Big Data SQL on all supported Oracle Database systems without uninstalling a previous version.

Before installing this Oracle Big Data SQL release on Oracle Big Data Appliance, you must use `bdacli` to manually uninstall the older version if it had been enabled via `bdacli` or Mammoth. If you are not sure, try `bdacli bds disable`. If the disable comment fails, then the installation was likely done with the `setup-bds` installer. In that case, you can install the new version Oracle Big Data SQL without disabling the old version.

How Long Does It Take?

The table below estimates the time required for each phase of the installation. Actual times will vary.

Table 1-2 Installation Time Estimates

Installation on the Hadoop Cluster	Installation on Oracle Database Nodes
<p>Eight minutes to 28 minutes</p> <p>The Hadoop side installation may take eight minute if all resources are locally available. An additional 20 minutes or more may be required if resources must be downloaded from the Internet.</p>	<p>The average installation time for the database side can be estimated as follows:</p> <ul style="list-style-type: none"> • 15 minutes for a single node database if a restart is not required. If a restart is required, the time will vary, depending on the size of the database. • On a RAC database, multiply the factors above by the number of nodes. • If an Oracle Grid restart is required, factor that in as well.

The installation process on Hadoop side includes installation on the Hadoop cluster as well as generation of the bundle for the second phase of the installation on the Oracle Database side. The database bundle includes Hadoop and Hive clients and other software. The Hadoop and Hive client software enable Oracle Database to communicate with HDFS and the Hive Metastore. The client software is specific to the version of the Hadoop distribution (i.e. Cloudera or Hortonworks). As explained later in this guide, you can download these packages prior to the installation, set up an URL or repository within your network, and make that target available to the installation script. If instead you let the installer download them from the Internet, the extra time for the installation depends upon your Internet download speed.

Pre-installation Steps

- Check to be sure that the Hadoop cluster and the Oracle Database system both meet all of the prerequisites for installation. On the database side, this includes confirming that all of the required patches are in installed. Check against these sources:
 - [Oracle Big Data SQL Master Compatibility Matrix](#) (Doc ID 2119369.1 in [My Oracle Support](#))
 - Sections 2.1 in this guide, which identifies the prerequisites for installing on the Hadoop cluster. Also see Section 3.1, which describes the prerequisites for installing the Oracle Database system component of Oracle Big Data SQL.

Oracle Big Data Appliance already meets all prerequisites.

- Have these login credentials available:
 - `root` credentials for both the Hadoop cluster and all Oracle Grid nodes.
On the grid nodes you have the option of using passwordless SSH with the `root` user instead.
 - `oracle` Linux user (or other, if the database owner is not `oracle`)
 - The Oracle Grid user (if this is not the same as the database owner).
 - The Hadoop configuration management service (CM or Ambari) admin password.
- On the cluster management server (where CM or Ambari is running), download the Oracle Big Data SQL installation bundle and unzip it into a [permanent](#) location of your choice. (See [Downloading Oracle Big Data SQL](#).)

- On the cluster management server, check to see if Python (at least 2.7, but less Python 3.0) is installed, as described in [Prerequisites for Installation on the Hadoop Cluster](#).

Outline of the Installation Steps

This is an overview to familiarize you with the process. Complete installation instructions are provided in Chapters 2 and 3.

The installation always has two phases – the installation on the Hadoop cluster and the subsequent installation on the Oracle Database system. It may also include the third, “Database Acknowledge,” phase, depending on your configuration choices.

Note to Oracle Big Data Appliance 4.10 Customers:

Steps 1 through 3 below do not apply. See Special Installation Procedures for Oracle Big Data Appliance 4.10.

1. Start the Hadoop-Side Installation

Review the installation parameter options described in Chapter 2. The installation on the Hadoop side is where you make all of the decisions about how to configure Oracle Big Data SQL, including those that affect the Oracle Database side of the installation.

2. Edit the `bds-config.json` file provided with the bundle in order to configure the Jaguar installer as appropriate for your environment. You could also create your own configuration file using the same parameters.
3. Run the installer to perform the Hadoop-side installation as described in [Installing or Upgrading the Hadoop Side of Oracle Big Data SQL](#).

If the Database Authentication feature is enabled, then Jaguar must also output a “request key” (`.reqkey`) file for each database that will connect to the Hadoop cluster. You generate this file by including the `--requestdb` parameter in the Jaguar `install` command (the recommended way). You can also generate the file later with other Jaguar operations that support the `--requestdb`.

This file contains one half of a GUID-key pair that is used in Database Authentication. The steps to create and install the key are explained in more detail in the installation steps.

4. Copy the database-side installation bundle to any temporary directory on each Oracle Database compute node.
5. If a request key file was generated, copy over that file to the same directory.

6. Start the Database-Side Installation

Log on to the database server as the database owner. Unzip bundle and execute the run file it contained. The run file does not install the software. It sets up an installation directory under `$ORACLE_HOME`.

7. As the database owner, perform the Oracle Database server-side installation. (See [Installing or Upgrading the Oracle Database Side of Oracle Big Data SQL](#).)

In this phase of the installation, you copy the database-side installation bundle to a temporary location on each compute node. If a `.reqkey` file was generated for the

database, then copy the file into the installation directory before proceeding. Then run the `bds-database-install.sh` installation program.

The database-side installer does the following:

- Copies the Oracle Big Data SQL binaries to the database node.
- Creates all database metadata and MTA extprocs (external processes) required for access to the Hadoop cluster, and configures the communication settings.

Important:

Be sure to install the bundle on each database compute node. The Hadoop-side installation automatically propagates the software to each node of the Hadoop cluster. However, the database-side installation does not work this way. You must copy the software to each database compute node and install it directly.

In Oracle Grid environments, if cell settings need to be updated, then a Grid restart may be needed. Be sure that you know the Grid password. If a Grid restart is required, then you will need the Grid credentials to complete the installation.

8. If Applicable, Perform the “Database Acknowledge” Step

If Database Authentication or Hadoop Secure Impersonation were enabled, the database-side installation generates a ZIP file that you must copy back to Hadoop cluster management server. The file is generated in the installation directory under `$ORACLE_HOME` and has the following filename format.

```
<Hadoop cluster name>-<Number nodes in the cluster>-<FQDN of the  
cluster management server node>-<FQDN of this database node>.zip
```

Copy this file back to `/opt/oracle/DM/databases/conf` on the Hadoop cluster management server and then as `root` run the Database Acknowledge command from the `BDSJaguar` directory:

```
# cd <Big Data SQL install directory>/BDSJaguar  
# ./jaguar databaseack <bds-config.json>
```

Workflow Diagrams

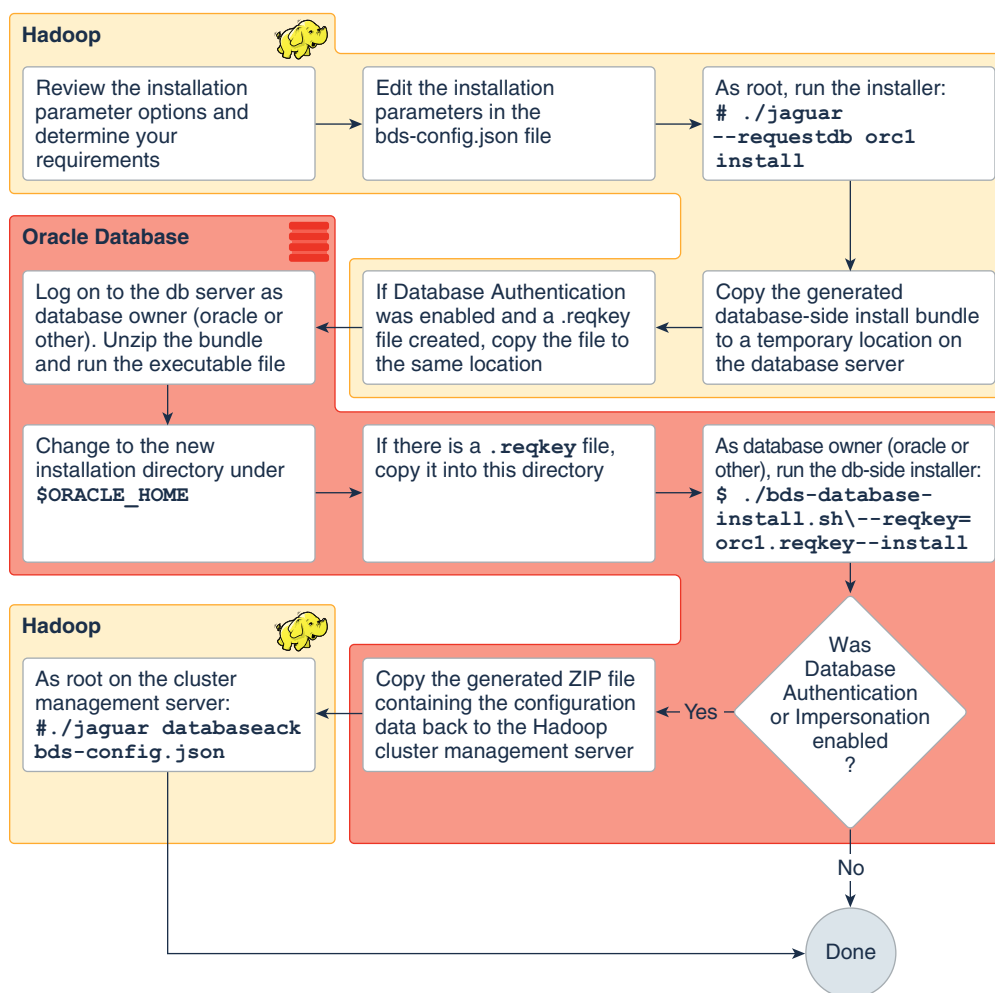
Complete Installation Workflow

The figure below illustrates the complete set of installation steps as described in this overview.

Note:

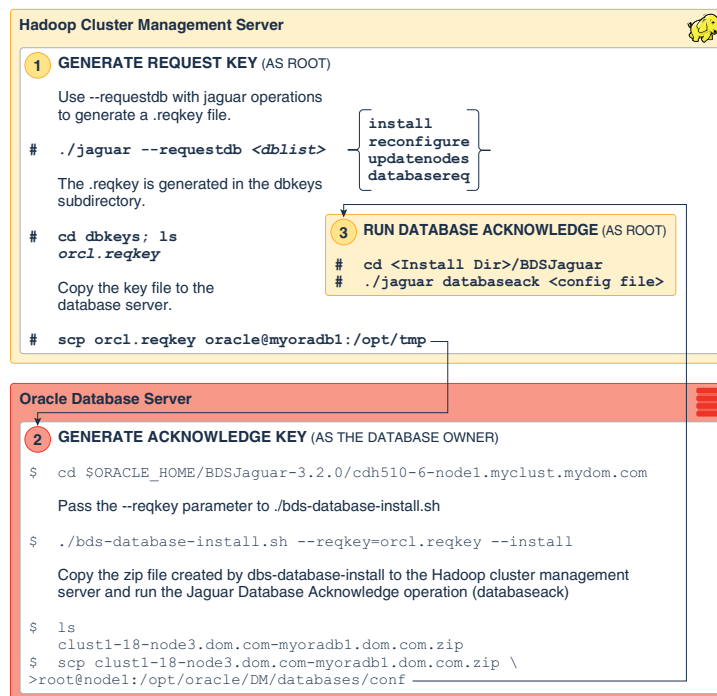
Before you start the steps shown in the workflow, be sure that both systems meet the installation prerequisites.

Figure 1-1 Installation Workflow



Key Generation and Installation

The figure below focuses on the three steps required to create and installing the GUID-key pair used in Database Authentication. The braces around parameters of the Jaguar command indicate that one of the operations in the list is required. Each of these operations supports use of the `--requestdb` parameter.

Figure 1-2 Generating and Installing the GUID-Key Pair for Database Authentication

1.9 Post-Installation Checks

Validating the Installation With `bdschecksw` and Other Tests

- The script `bdschecksw` now runs automatically as part of the installation. This script gathers and analyzes diagnostic information about the Oracle Big Data SQL installation from both the Oracle Database and the Hadoop cluster sides of the installation. You can also run this script as a troubleshooting check at any time after the installation. The script is in `$ORACLE_HOME/bin` on the Oracle Database server.

```
$ bdschecksw --help
```

See *Running Diagnostics With `bdachecksw`* in the *Oracle Big Data SQL User's Guide* for a complete description.

- Also see *How to do a Quick Test* in the user's guide for some additional functionality tests.

Checking the Installation Log Files

You can examine these log files after the installation.

On the Hadoop cluster side:

```
/var/log/bigdatasql  
/var/log/oracle
```

On the Oracle Database side:

```
$ORACLE_HOME/install/bds* (This is a set of files, not a directory)  
$ORACLE_HOME/bigdatasql/logs  
/var/log/bigdatasql
```



Tip:

If you make a support request, create a zip archive that includes all of these logs and include it in your email to Oracle Support.

Other Post-Installation Steps to Consider

- Read about measures you can take to secure the installation. (See [Securing Big Data SQL](#).)
- Learn how to modify the Oracle Big Data SQL configuration when changes occur on the Hadoop cluster and in the Oracle Database installation. (See [Expanding or Shrinking an Installation](#).)
- If you have used Copy to Hadoop in earlier Oracle Big Data SQL releases, learn how Oracle Shell for Hadoop Loaders can simplify Copy to Hadoop tasks. (See [Additional Tools Installed](#).)

1.10 Using the Installation Quick Reference

Once you are familiar with the functionality of the Jaguar utility on the Hadoop side and `bds-database-install.sh` on the Oracle Database side, you may find it useful to work from the [Installation Quick Reference](#) for subsequent installations. This reference provides an abbreviated description of the installation steps. It does not fully explain each step, so users should already have a working knowledge of the process. Links to relevant details in this and other documents are included.

2

Installing or Upgrading the Hadoop Side of Oracle Big Data SQL

After downloading the Oracle Big Data SQL deployment bundle and extracting the files, the next step is to configure the installer and then run the installation.



Note to Oracle Big Data Appliance 4.10 Customers:

The Hadoop-side installation instructions provided in this section does not apply to Oracle Big Data Appliance 4.10. See *Special Installation Procedures for Oracle Big Data Appliance 4.10*.

The installation of Oracle Big Data SQL on the Hadoop cluster is deployed using the services provides by the cluster management server (Cloudera Manager or Ambari). The Jaguar `install` command uses the management server API to register the BDS service and start the deployment task. From there, the management server controls the process and deploys the software to the nodes of the cluster and installs it.

The installation also generates the deployment bundle for the database side of the installation.

If a previous version of Oracle Big Data SQL is already installed, Oracle Big Data SQL upgrades the installation to Release 3.2. Users of previous Oracle Big Data SQL releases, please note that `setup-bds`, the previous Hadoop-side installer, had been replaced by the Jaguar installer and that the configuration parameters for the installation have changed significantly in this release.

2.1 Before You Start

Check to Ensure that all DataNodes of the Cluster Meet the Prerequisites

Prerequisites for both the Hadoop and Oracle Database sides of the installation are described earlier in this chapter.

If you miss a prerequisite, the installer will alert you. You can see this in the output examples in *Appendix B*.

Plan the Configuration and Edit the Jaguar Configuration File

Review the [Jaguar Configuration Parameter and Command Reference](#) in this chapter. You should know the answers to these questions before editing the configuration file:

- *Do you want to enable Database Authentication in order to validate connection requests from Oracle Database to Big Data SQL server processes on the Hadoop DataNodes?*

Database Authentication in the network connection between the Oracle Database and Hadoop is set to "true" in the configuration by default. You have the option to disable it by setting the `database_auth_enabled` to "false":

```
"database_auth_enabled" : "false",
```

- *Do you want to use the Multi-User Authorization feature?*

Multi-User Authorization enables you to grant users other than `oracle` permissions to run SQL queries against the Hadoop cluster.

The first step in setting up Multi-User Authorization is to set these parameters in the security section of the configuration file:

```
"impersonation_enabled" : "true",
"impersonation_usehosts" : "true",
"impersonation_blacklist" : "hdfs,hive"
```

Note that you can add any account to the blacklist.

- *Are the Hadoop cluster and the Oracle Database system going to communicate over Ethernet or over InfiniBand? Also, do the Hadoop nodes have more than one network interface?*

See the `use_infiniband` and `selection_subnet` parameters. (The `selection_subnet` does not apply to Oracle Big Data Appliance.)

```
"use_infiniband" : "false",
"selection_subnet" : "5.32.128.10/21"
```

By default `use_infiniband` is set to false. Ethernet is the default protocol.

- *Are you going pre-download the Hadoop and Hive client tarballs and set up a local repository or directory where the installer can acquire them, or, do you want the installer to download them directly from the Cloudera or Hortonworks repository on the Internet?*

On Oracle Big Data Appliance the tarballs are stored on and automatically retrieved from a repository within the cluster. Therefore, the related configuration parameters do not apply. However, for other CDH systems or Hortonworks, you can choose to specify an URL or directory path or both for the tarballs. If instead you choose to accept the default and let the installer download the files from the Internet, then ensure that there is Internet access from the node where you are running the installation.

```
"url" : [
"http://repo1.domain.com/loc/hadoop" ,
"http://repo2.domain.com/loc" ,
"http://alternate.domain.com/backup/repo"
],
"dir" : [ "/root/hadooprepo" ]
```


If you choose to let the installer download the tarballs from the Internet, you can configure the installer to work with HTTP/HTTPS proxies within your environment:

```
"http_proxy" : "http://my.proxy.server.com:80",  
"https_proxy" : "http://mysecure.proxy.server.com:80"
```

You can either set the proxy values in the configuration file or use preexisting proxy environment variables.

- *If the network is Kerberos-secured, do you want the installer to set up automatic Kerberos ticket renewal for the oracle account used by Oracle Big Data SQL?*

See the parameters in the `kerberos` section:

```
"principal" : "oracle/mycluster@MY.DOMAIN.COM",  
"keytab" : "/home/oracle/security/oracle.keytab",  
"hdfs-principal" : "hdfs/mycluster@MY.DOMAIN.COM",  
"hdfs-keytab" : "/home/hdfs/security/hdfs.keytab"
```

- *Do you want the Oracle Big Data SQL install process to automatically restart services that are in a stale state?*

By default, stale services are restarted automatically. If you want to suppress this, you can set the `restart_stale` parameter in the configuration file to `false`.

- *Is the Hadoop cluster using the default REST API port for CDH or Ambari?*

If not, set the `ports` parameter.

Note:

Setting these parameters in the configuration file does not complete the set up for some features. For example, to enable Database Authentication, you must also pass a special `--requestdb` parameter to the Jaguar utility in order to identify the target database or databases. There are also steps required to generate and install the security key used by this feature. To enable Multi-User Authorization, you start by setting the Hadoop Impersonation parameters in the configuration file, but also need to set up the authorization rules. The steps to complete these setups are provided where needed as you work through the instructions in this guide.

2.2 About the Jaguar Utility

Jaguar is a multifunction command line utility that you use to perform all Oracle Big Data SQL operations on the Hadoop cluster.

Jaguar currently supports these operations:

- `install`
 - Deploys Oracle Big Data SQL binaries to each cluster node that is provisioned with the DataNode service.

- Configures Linux and network settings for `bd_cell` (the Oracle Big Data SQL service) on each of these nodes.
- Generates the bundle that installs Oracle Big Data SQL on the Oracle Database side. It uses the parameter values that you set in the configuration file in order to configure the Oracle Database connection to the cluster.
- `updatenodes`
Checks and updates the installation on the DataNodes of the Hadoop cluster.
- `reconfigure`
Modifies the current configuration of the installation (according to the settings in the configuration file provided).
- `databasereq`
Generates a request key file that contains one segment of the GUID-key pair used in Database Authentication. (The `databasereq` operation performs this function only. For `install`, `updatenodes`, and `reconfigure`, request key generation is an option that can be included as part of the larger operation.)
- `databaseack`
Perform the last step in Database Authentication setup – install the GUID-key pair on all Hadoop DataNodes in a cluster in order to allow queries from the Oracle Database that provided it.
- `uninstall`
Uninstalls Oracle Big Data SQL from all DataNodes of the Hadoop cluster.



See Also:

[Jaguar Operations](#) in the next section provides details and examples.

2.2.1 Jaguar Configuration Parameter and Command Reference

Configuration Parameters

The table below describes all parameters available for use in `bds-config.json` or your own configuration file. Only the `cluster_name` parameter is always required. Others are required under certain conditions stated in the description.



Note:

When editing the configuration file, be sure to maintain the JSON format. Square brackets are required around lists, even in the case of a list with a single item.

Table 2-1 Configuration Parameters in bds-config.json (or in Customer-Created Configuration Files)


Parameter	Type	Description
name	String	<p>The name of the cluster.</p> <p>For CDH clusters (Oracle Big Data Appliance or other), this name can be either the physical cluster name or the display name. The installer searches first by physical name and then by display name.</p> <p>Required.</p>
port	Integer	<p>Cloudera Manager or Ambari REST API port.</p> <p>By default, on CDH clusters this port is 7183 for secured and 7180 for unsecured access. For Ambari, is 8443 for secured and 8080 for unsecured.</p> <p>Optional.</p>
restart_stale	Boolean	<p>If "true", then services with stale configurations are restarted at the end of install process. These services are HDFS NodeName, YARN NodeManager and/or Hive (depending upon the settings selected).</p> <p>If "false", the installation will finish but those services will remain on stale state. This is useful for avoid service interruption in order to restart when it will be more convenient.</p> <p>Optional. The default is "true".</p>
dir	List	<p>List of directories where the Hadoop clients for deployment on the database side are located. These directories can be on the local file system or on NFS. Directories are searched in the order listed. By default, the list is empty. If the dir list has any entries, these are searched before the URL list is searched, since this option should provide the fastest access to the clients. To give the installer the quickest access to the tarballs, you could set up a local repository, download the tarballs separately through a direct Internet connection, copy them into a directory on the same node where the Oracle Big Data SQL installer will run, and list that directory in the dir parameter.</p> <p>Optional.</p> <p>Not applicable to Oracle Big Data Appliance.</p>

Table 2-1 (Cont.) Configuration Parameters in `bds-config.json` (or in Customer-Created Configuration Files)

Parameter	Type	Description
<code>url</code>	List	<p>This is the list of URLs where the Hadoop client tarballs for deployment on the database side are located. If you data center already has repositories set up for access via HTTP, then you may prefer to maintain the Hadoop tarballs in that repository and use URL parameter for Oracle Big Data SQL installations. The URLs can be to the localhost, an internal network, or a site on the Internet (if the node has Internet access). The URLs are tried in the order listed. If Hadoop clients cannot be successfully located via the <code>directory</code> list or <code>url</code> list or if both lists are empty, then the Cloudera or Hortonworks public repo is used by default. In that case, the node must have Internet access at install time. Note that internal proxy values and/or OS environment proxy settings must be set to allow this access if needed.</p> <p>If access to all listed repositories fails and/or Internet access is blocked, the database installation bundle is not created and a warning message is displayed. After correcting any problems and providing access to a repository, you can re-run the installer using the <code>reconfigure</code> and the installer should successfully generated the database-side installation bundle. Note that <code>reconfigure</code> detects and implements changes according the current directives in the configuration file. It does not uninstall and reinstall Oracle Big Data SQL on the cluster.</p> <p>Not applicable to Oracle Big Data Appliance. The tarballs are stored in a local repository in the cluster and the location is automatically added to the configuration file.</p>
<code>use_infinib</code> and	Boolean	<p>If <code>"true"</code>, the communication will be set through private network interface, if <code>"false"</code> by client network interface.</p> <p>Used for Oracle Big Data Appliance clusters only.</p> <p>Default value: <code>"false"</code>.</p>

Table 2-1 (Cont.) Configuration Parameters in bds-config.json (or in Customer-Created Configuration Files)

Parameter	Type	Description
selection_subnet	String	<p>If Hadoop cluster nodes have several network interfaces, you can use <code>selection_subnet</code> to select one. The selected IP address will be the nearest to the selection subnetwork.</p> <p>If the Hadoop cluster nodes have only one network interface, this parameter is ignored.</p> <p>The default value depends upon these conditions:</p> <ul style="list-style-type: none"> On non-Oracle commodity Hadoop clusters (CDH or HDP) the default selection is 0.0.0.0/0. (If a cluster node has several IP addresses, the lowest address is selected.) On Oracle Big Data Appliance, the default is either the private or client IP address, depending upon the setting of the <code>use_infiniband</code> parameter.
http_proxy	String	<p>Specify the proxy settings to enable download the Hadoop client tarballs and cluster settings.</p> <p>If both of these strings are empty, the OS environment proxy settings are used.</p> <p>By default, both strings are empty.</p> <p>Using these two parameters in the configuration file is optional. If they are needed, you could instead set them externally as in <code>export http_proxy=<proxy value></code></p> <p>Not applicable to Oracle Big Data Appliance</p>
https_proxy	String	
database_authentication_enabled	Boolean	<p>If "true", the database authentication through the GUID-key mechanism is enabled. This requires an extra step in the installation process in order to set up the database GUID-key pair on the cluster side.</p> <p>If "false", the feature will not be enabled.</p> <p>Default value: "true".</p>
impersonation_enabled	Boolean	<p>If "true", Hadoop impersonation for Multi-user Authorization support is enabled. This sets up the <code>oracle</code> OS user as the Hadoop proxy user and propagates the proxy user's black list to the database nodes. If "false", this feature is not enabled.</p> <p>Default value: "false" .</p>


Note for Oracle Big Data Appliance Users:

It's possible to configure several networks on an Oracle Big Data Appliance. If multiple networks exist, then this parameter must be set in order to select a specific network.

Table 2-1 (Cont.) Configuration Parameters in bds-config.json (or in Customer-Created Configuration Files)

Parameter	Type	Description
<code>impersonation_blacklist</code>	String	The Hadoop proxy users blacklisted for impersonation. This parameter is used only if Hadoop impersonation is enabled. Since this is a required setting on the Oracle Database side, it is provided with a default value of "dummy" in order to avoid extproc errors that can occur if Hadoop Impersonation is not enabled.
<code>impersonation_usehosts</code>	Boolean	If "true", the proxy hosts variable is set to the IP address of the database node. If "false", the proxy hosts variable is set to the wildcard: "*". Default value: "true".
<code>principal</code>	String	The fully-qualified Kerberos principal name for the oracle user. The principal has three parts: <ul style="list-style-type: none"> • The User Name: (must always be oracle) • Qualifier: "/<qualifier>." This is optional information to help you organize and identify principals. • Domain: "@MY.DOMAIN.COM." This is required information managed by the KDC. The Oracle Big Data SQL installation uses the Kerberos <code>principal</code> field (and <code>keytab</code> field below) to set up automated Kerberos ticket renewal for the user represented by the principal. It does this on both the Hadoop and Oracle Database sides of the installation. The installer does not create the principal or the keytab. These must already exist. Required for secured clusters.
<code>keytab</code>	String	Fully-qualified location for the principal's keytab file name. Copy the keytab file to a location accessible to the Jaguar installer and set the path as the value of this parameter.
<code>hdfs-principal</code>	String	Fully-qualified Kerberos principal name for the "hdfs" user. It has three parts: User name, Qualifier, and Domain. The User name is the fully-qualified principal name for the hdfs user. Qualifier is the cluster name prefixed by a forward slash, as in /mycluster. Domain is specified in the form @MY.DOMAIN.COM. All three are required. The principal name is defined on the KDC. Required for secured clusters.
<code>hdfs-keytab</code>	String	Fully-qualified path to the principal keytab file. A keytab file is created for each principal on the KDC server. It must exist in a location accessible to the Jaguar installer. Required for secured clusters.
<code>min_hard_limit</code>	Integer	The minimum amount of memory reserved for Big Data SQL, in megabytes. This parameter is used on CDH clusters (Oracle Big Data Appliance and others). It is not used on HDP clusters. By default, the value is 8192 (8 GB).

Table 2-1 (Cont.) Configuration Parameters in bds-config.json (or in Customer-Created Configuration Files)

Parameter	Type	Description
max_percent age	Integer	<p>On CDH clusters (Oracle Big Data Appliance and others) this parameter specifies the percentage of memory on each node to reserve for Big Data SQL. This percentage is considered from a total amount of: NodeManager if YARN ResourceManager is enabled for that node. Physical memory if not. .</p> <p>If the YARN Resource Manager is enabled for the node, then percentage should be based on the total amount of memory used by the NodeManager. Otherwise it should be a percentage of physical memory.</p> <p>This parameter is ignored on HDP clusters.</p>

 **Note:**

After Oracle Big Data SQL is installed on the Hadoop cluster management server, you can find configuration file examples that demonstrate various parameter combinations in the *<Big Data SQL Install directory>/BDSjaguar* directory:

```
example-bda-config.json
example-cdh-config.json
example-kerberos-config.json
example-localrepos-config.json
example-subnetwork-config.json
example-unsecure-config.json
```

You can see all possible parameter options in use in `example-cdh-config.json` .

 **See Also:**

See the Appendix [Determining the Correct Software Version and Composing the Download Paths for Hadoop Clients](#) for suggestions that can help with the setup of client tarball downloads.

Jaguar Operations

The table below lists the full set of operations performed by the Jaguar utility on the Hadoop side of the Oracle Big Data SQL installation.

The general syntax for Jaguar commands is as follows. The `--requestdb` parameter does not apply to all Jaguar commands.

```
# ./jaguar {--requestdb <comma-separated database names> | NULL }
<action> { bds-config.json | <myfilename>.json | NULL }
```

Examples:

```
# ./jaguar install
# ./jaguar install bds-config.json
# ./jaguar install mycustomconfig.json
# ./jaguar --requestdb orcl,testdb,proddb install
# ./jaguar --requestdb orcl install
```

You can use the default `bds-config.json` or your own configuration file, or omit the configuration file argument (which defaults to `bds-config.json`).

 **About --requestdb:**

The `--requestdb` parameter is required for the `databasereq` command, optional for `install`, `updatenodes`, and `reconfigure`, and non-applicable for other Jaguar commands. The parameter must be passed in to one of these operations in order to enable Database Authentication in the connection between a Hadoop cluster and a database. Unless you prefer to disable Database Authentication, it is recommended that you include `--requestdb` with the initial `install` operation. Otherwise, you will need perform an additional step later in order to generate the request key.

This parameter is functional only when Database Authentication (`database_auth_enabled`) is set to `"true"` in the configuration. (This setting is a configuration default and does not need to be explicitly set in the configuration file.)

Jaguar needs the database names in order to generate a unique `.reqkey` (request key) file for each database. When `database_auth_enabled` is set `"true"` at installation time, the `--requestdb` parameter is still optional. Post-installation you have the same option to send the request key in the `updatenodes`, `reconfigure`, and `databasereq` commands. Database Authentication is not implemented until you do all of the following:

1. Ensure that `database_auth_enabled` is either absent from the configuration file or is set to `"true"`. (It is `"true"` by default.)
2. Include `--requestdb` in a Jaguar command:
 - a. Run the Jaguar `install`, `updatenodes`, or `reconfigure` and install the updated database-side installation bundle, or
 - b. Run Jaguar `databasereq` to generate an acknowledge key from the existing database side installation.
3. Copy the generated ZIP file that contains the `.ackkey` file from the database-side installation directory to `/opt/oracle/DM/databases/conf` on the Hadoop cluster management server.
4. Run the Jaguar `databaseack` command as described in the table below.

The table below shows the available Jaguar commands.

Table 2-2 Jaguar Operations

Jaguar Operation	Supports --requestdb?	Usage and Examples
install The --requestdb parameter is not strictly required by the install operation, but you cannot enable Database Authentication if you do not generate a request key for each database. --requestdb <comma-separated database list>	Y	<p>Installs Oracle Big Data SQL on the Hadoop cluster identified in the configuration file and creates an installation bundle for the database side based on the parameters included in the configuration file (or default values for parameters not explicitly assigned value in the configuration file). Examples:</p> <pre># ./jaguar --requestdb orcl,testdb,proddb install</pre> <p>No configuration file parameter is included in the above example. <code>bds-config.json</code> is the implicit default. You can specify a different configuration file as in <code>./jaguar --requestdb mydb install myconfig.json</code></p>
updatenodes	Y	Expand or shrink the cluster. Oracle Big Data SQL to any new DataNodes and update the cells inventory if the cluster has grown since the last Oracle Big data SQL installation.

 **Note:**

You may need to use the `scl` utility to ensure that the correct Python version is invoked:

```
scl enable python27 "./
jaguar install"
```

Table 2-2 (Cont.) Jaguar Operations

Jaguar Operation	Supports --requestdb?	Usage and Examples
reconfigure	Y	<p>Modify the current installation by applying changes you have made to the configuration file (<code>bds-config.json</code> or other).</p> <pre># ./jaguar reconfigure myconfigfile.json</pre> <p>Note that if you run <code>./jaguar reconfigure <config file></code> to reconfigure Oracle Big Data SQL on the Hadoop cluster, a corresponding reconfiguration is required on the Oracle Database side. The two sides cannot communicate if the configurations do not match. The Jaguar utility regenerates the database-side bundle files to incorporate the changes. You must redeploy the bundle on all database servers where it was previously installed.</p> <p>The <code>--requestdb</code> argument is required if <code>database_auth_enabled</code> is set to "true" in the updated configuration file. This is so that Jaguar will generate <code>.reqkey</code> files that are included in the database-side installation bundle. Note that we let the configuration file parameter default to <code>bds-config.json</code>.</p> <pre># ./jaguar --requestdb demodb,testdb,proddb1 reconfigure</pre>
databasereq	Y	<p>Use this command to create the <code>.reqkey</code> file without repeating the Hadoop-side installation, or doing an <code>updatenodes</code> or <code>reconfigure</code> operation. (For example, if you forgot to include the <code>--requestdb</code> argument with the Jaguar <code>install</code> command), you can create a request key later with <code>databasereq</code>. This operation requires that <code>database_auth_enabled</code> is set to "true" (the default value) in the configuration.</p> <pre># ./jaguar --requestdb demodb,testdb,proddb1 databasereq</pre>

Table 2-2 (Cont.) Jaguar Operations

Jaguar Operation	Supports --requestdb?	Usage and Examples
databaseack	N	<p>The “Database Acknowledge” process provides confirmation to the Oracle Big Data SQL installation on the Hadoop cluster that security features you enabled in the configuration file have been successfully implemented in the database-side installation. It then completes implementation of the selected security features on the Hadoop cluster side.</p> <pre>./jaguar databaseack bds-config.json</pre> <p>Only run databaseack if you chose to enable security features by setting either of these parameters in the configuration file to “true”:</p> <ul style="list-style-type: none"> “impersonation_enabled” : “true” “database_auth_enabled” : “true” <p>If a database-side installation bundle is built with any of these features set to “true”, then the database-side installation from that bundle generates a ZIP file in the installation directory under \$ORACLE_HOME on the database server. The format of the ZIP file name is <i><Hadoop cluster name>-<Number nodes in the cluster>-<FQDN of the cluster management server node>-<FQDN of this database node>.zip</i>. For example:</p> <pre>\$ ls \$ORACLE_HOME/BDSJaguar-3.2.0/cdh510-6- node1.my.domain.com/*.zip \$ cdh510-6-node1.my.domain.com- myoradb1.mydomain.com.zip</pre> <p>Copy this ZIP archive back to <code>/opt/oracle/DM/databases/conf</code> on the Hadoop cluster management server after the database-side installation is complete. Then, to fully enable the security features, run databaseack.</p>
uninstall		<p>Uninstall Oracle Big Data SQL from the Hadoop cluster.</p> <p>The uninstall process stops the <code>bd_cell</code> process (the Oracle Big Data SQL process) on all Hadoop cluster nodes, removes all instances from Hadoop cluster, and release all related resources.</p>

 **Note:**

When Oracle Big Data SQL is uninstalled on the Hadoop side, any queries against Hadoop data that are in process on the database side will fail. It is strongly recommended that you uninstall Oracle Big Data SQL from all databases systems shortly after uninstalling the Hadoop component of the software.

**See Also:**

[Uninstalling Oracle Big Data SQL.](#)

2.3 Steps for Installing on the Hadoop Cluster

After you have set up the Jaguar configuration file according to your requirements, follow these steps to run the Jaguar installer, which will install Oracle Big Data SQL on the Hadoop cluster and will also generate a database-side installation bundle that you deploy to the Oracle Database system. In these steps, `bds-config.json` is the configuration filename passed to Jaguar. This is the default. Any file name is accepted, therefore you can create separate configuration files for installation on different clusters and save them in different files.

**Note:**

Jaguar requires Python 2.7 to 3.0. If necessary, you can add a Jaguar-compatible version of Python as a secondary installation. Revisit the prerequisites section in the Introduction for details. If you are using Oracle Big Data Appliance (and possibly other Hadoop platforms), do not overwrite the installed Python release.

1. Log on to the cluster management server node as `root` and `cd` to the directory where you extracted the downloaded Oracle Big Data SQL installation bundle.
2. `Cd` to the `BDSJaguar` subdirectory under the path where you unzipped the bundle.

```
# cd <Big Data SQL Install Directory>/BDSJaguar
```

3. Edit the file `bds-config.json`. At a minimum, provide the required cluster name for the cluster where you want to install Oracle Big Data SQL:

```
{  
  "cluster": {  
    "name": "<Your cluster name>"  
  }  
}
```

Add the parameters that you want to use in this installation.

**See Also:**

The cluster name is the only required parameter in this version of Oracle Big Data SQL. See the [Jaguar Configuration Parameter and Command Reference](#) for a description of all available parameters. You can see an example of a `bds-config.json` file populated with all available parameters in [bds-config.json Configuration Example](#).

In the BDSJaguar directory, run the `Jaguarinstall` operation. Pass the `install` parameter and the configuration file name. (`bds-config.json` is the implicit default) as arguments to the Jaguar command. You may or may not need to include the `--requestdb` option.

```
[root@myclusteradminserver:BDSjaguar] # ./jaguar install <config file name>
```

 **Note:**

By default, Database Authentication is set to true unless you set `database_auth_enabled` to “false” in the configuration file. If you enable Database Authentication, then either as part of the install operation or later, generate a “request key.” This is half of a GUID/key pair used in the authentication process. To generate this key, include the `--requestdb` parameter in the Jaguar `install` command line:

```
[root@myclusteradminserver:BDSjaguar] # ./jaguar --requestdb mydb install
```

If the install was run with `database_auth_enabled` is “true”, you can use the Jaguar `databasereq` command to generate the key after the database-side installation. Several other Jaguar commands can also generate the request key if you pass them the `--requestdb` parameter.

Jaguar prompts for the cluster management service administrator credentials and then installs Oracle Big Data SQL throughout the Hadoop cluster. It also generates the database-side installation bundle in the `db-bundles` subdirectory. The following message is returned if the installation completed without error.

```
BigDataSQL: INSTALL workflow completed.
```

4. Check for the existence of the database side installation bundle:

```
# ls <Big Data SQL Install Directory>/BDSJaguar/db-bundles  
bds-3.2.0-db-<cluster>-<yymmdd.hhmi>.zip
```

This bundle is for setting up Oracle Big Data SQL connectivity Oracle database and the specific cluster defined in the `bds-config.json` (or other) configuration file. It contains all packages and settings files required except for an optional database request key file.

If you included `--requestdb` in the install command, then the installation also generates one or more database request key files under the `dbkeys` subdirectory. You should check to see that this key exists.

```
# ls <Big Data SQL Install Directory>/BDSJaguar/dbkeys  
clusterldb.reqkey
```

This completes the Oracle Big Data SQL installation on the Hadoop cluster.

What Next?

After Jaguar has successfully installed Oracle Big Data SQL on the Hadoop cluster, you are done with the first half of the installation. The next step is to install Oracle Big Data SQL on the Oracle Database system that will run queries against the data on the Hadoop cluster.

To do this, copy the database-side installation bundle to any location on the Oracle Database system. Unless you set `database_auth_enabled` to "false" in the configuration file, then also copy over the `.reqkey` file generated by Jaguar.

Tip:

You only need to send a request key to a database once. A single request key is valid for all Hadoop cluster connections to the same database. If you have already complete the installation to connect one Hadoop cluster to a specific database, then the database has the key permanently and you do not need to generate it again or copy it over to the database again in subsequent cluster installations.

Go to [Installing or Upgrading the Oracle Database Side of Oracle Big Data SQL](#) for instructions on unpacking the bundle and installing the database-side components of the software.

See Also:

An example of the complete standard output from a successful installation is provided in [Oracle Big Data SQL Installation Examples](#).

2.4 Special Installation Procedures for Oracle Big Data Appliance 4.10

If you have enabled or are planning to enable Big Data SQL 3.2, then download and install patch 27077322.

Oracle Big Data Appliance 4.10 includes an early, unofficial release of Oracle Big Data SQL 3.2 in its deployment bundle. [Patch 27077322](#) includes important changes which update the Oracle Big Data SQL package to the official 3.2 level. You can load the patch either before or after Oracle Big Data SQL 3.2 is enabled on Oracle Big Data Appliance 4.10.

For the patch installation steps, see [Installation of Mandatory One-Off Patch 27077322 On Big Data Appliance 4.10 OL5/OL6 with Big Data SQL 3.2 \(Doc ID 2340322.1\)](#) at [My Oracle Support](#).

This patch is not needed if you installed Oracle Big Data SQL 3.2 from the download available at the Oracle Software Delivery Cloud (eDelivery). However, if Release 3.2 is not already installed, our recommendation is to install the patch first and then use `bdaccli` to install the Oracle Big Data SQL package that is included with Oracle Big

Data Appliance 4.10. This installation method is better-integrated with the Oracle Big Data Appliance cluster and also includes checks before and after the installation.

! Important:

As described in the Oracle Big Data Appliance Owner's Guide, after the default installation of Oracle Big Data SQL, you can add customizations. If you have already customized your installation then before installing the patch, make a back up of all installation directories under `/opt/oracle/BDSJaguar` in order to preserve them. The `/opt/oracle/BDSJaguar` directory is deleted and then recreated empty when the patch is installed. Copy the directories back to `/opt/oracle/BDSJaguar` after the patch installation is complete.

3

Installing or Upgrading the Oracle Database Side of Oracle Big Data SQL

Oracle Big Data SQL must be installed on both the Hadoop cluster management server and the Oracle Database server. This section describes the installation of Oracle Big Data SQL 3.2 on Oracle Database systems, including single-node and RAC systems.

3.1 Before You Start the Database-Side Installation

Review the points in this section before starting the installation.

If the current version of Oracle Big Data SQL is already installed on the Oracle Database system and you are making changes to the existing configuration, you do not need repeat the entire installation process. See [Reconfiguring an Existing Oracle Big Data SQL Installation](#).

! Important:

- For multi-node databases (such as Oracle RAC systems), **you must repeat this installation on every compute node of the database.** If this is not done, you will see RPC connection errors when the Oracle Big Data SQL service is started.

You can reuse the same installation bundle on each database node by including the optional `--ip-cell` command line parameter to pass in the correct network interface address for the node when you run the installer. If this parameter is not used, the installer searches for and selects one.

- It is recommended that you set up passwordless SSH for root on the database nodes where Grid is running. Otherwise, you will need to supply the credentials during the installation on each node.
- If the `diskmon` process is not already running prior to the installation, a Grid infrastructure restart will be required in order to complete the installation.
- If you set up Oracle Big Data SQL connections from more than one Hadoop cluster to the same database, be sure that the configurations for each connection are the same. Do not set up one connection to use Infiniband and another to use Ethernet. Likewise, if you enable database authentication in one configuration, then this feature must be enabled in all Oracle Big Data SQL connections between different Hadoop clusters and the same database.
- If the database system is Kerberos-secured, then it is important to note that authentication for the database owner principal can be performed by only one KDC. Oracle Big Data SQL currently does not support multiple Kerberos tickets. If two or more Hadoop cluster connections are installed by the same database owner, all must use the same KDC.

The `KRB5_CONF` environment variable must point to only one configuration file. The configuration file must exist, and the variable itself must be set for the database owner account at system startup.

- For database running under Oracle Grid Infrastructure, if the system has more than one network interfaces of the same type (two or more Ethernet, two or more InfiniBand, or two or more Ethernet-over-InfiniBand interfaces), then the installation always selects the one whose name is first in alphabetical order.

 **See Also:**

Details are provide in [Special Considerations When a System Under Grid Infrastructure has Multiple Network Interfaces of the Same Type](#).

3.1.1 Potential Requirement to Restart Grid Infrastructure

In certain database environments `bds-database-install.sh` needs to create `cellinit.ora` and/or `celliniteth.ora`. In these cases, the script will attempt to propagate similar changes across all nodes in the Grid Infrastructure. To do this, the script expects that passwordless SSH is set up between `oracle` and `root` or it will prompt for the root password for each node during execution. If the nature of the changes requires a restart of the Grid Infrastructure, the script will also display messages indicating that grid infrastructure needs to be restarted manually. Because the installation cannot complete without Grid credentials if a restart is necessary, be sure that you have the Grid password at hand.

3.1.1.1 Understanding When Grid or Database Restart is Required

On the database side of Oracle Big Data SQL, the `diskmon` process is the agent in charge of communications with the Hadoop cluster. This is similar to its function on the Oracle Exadata Database Machine, where it manages communications between compute nodes and storage nodes.

In Grid environments, `diskmon` is owned by the Grid user. In a non-Grid environment, it is owned by the Oracle Database owner.

In Oracle Big Data SQL, `diskmon` settings are stored on `cellinit.ora` and `celliniteth.ora` files in the `/etc/oracle/cell/network-config/` directory. The installer updates these files in accordance with the cluster connection requirements.

This is how the installer determines when Grid or Oracle Database needs to be restarted:

- If the installer detects that no previous `cellinit.ora` or `celliniteth.ora` file exists, this means that no `diskmon` process is running. In this case, if the environment includes Oracle Grid then you must restart the Grid. If the environment does not include Grid, then you must restart Oracle Database.
- If previous `cellinit.ora` and/or `celliniteth.ora` file exist, this indicates that `diskmon` process is running. In this case, if the installer needs to make a change to these files, then only the database needs to be restarted.
- In multi-node Grid environments, `diskmon` works on all nodes as a single component and `cellinit.ora` and `celliniteth.ora` must be synchronized on all nodes. This task is done through SSH. If passwordless SSH is setup on the cluster, no user interaction is required. If not passwordless SSH is not set up, then the script will pause for you to input the root credentials for all nodes. When the `cellinit.ora` and `celliniteth.ora` files across all nodes are synchronized, then the script will continue. Then the script finishes and in this case, you must restart the Grid infrastructure.

3.1.2 Special Considerations When a System Under Grid Infrastructure has Multiple Network Interfaces of the Same Type

The Oracle Big Data SQL installation or SmartScan operation may sometimes fail in these environments because the wrong IP address is selected for communication with the cells.

When installing Oracle Big Data SQL within Oracle Grid Infrastructure, you cannot provide the installer with a specific IP address to select for communication with the Oracle Big Data SQL cells on the nodes. Network interface selection is automatically

determined in this environment. This determination is not always correct and there are instances where the database-side installation may fail, or, you may later discover that SmartScan is not working.

You can manually correct this problem, but first it is helpful to understand how the installation decides which network interfaces to select.

How the Installation Selects From Among Multiple Interfaces of the Same Type on a System Under Grid Infrastructure

- The `diskmon` process is controlled by Oracle Grid Infrastructure and not by the database. Grid manages communications with the Oracle Big Data SQL cells.
- The Oracle Big Data SQL installer in these cases does not create `cellinit.ora` and `celliniteth.ora`, nor does it update the cell settings stored in these files. In these environments, the task is handled by Grid, because it is a cluster-wide task that must be synchronized across all nodes.
- If there are multiple network interfaces, the Grid-managed update to the cells automatically selects the first network interface of the appropriate type on each node. It selects the interface whose name is first in alphabetical order. In other environments (with the exception of Exadata Database Machine systems) you can use the `--ip-cell` command line parameter with the installer to select a specific network interface. In a Grid environment, this parameter is ignored.

For example, here is a system that is under Grid Infrastructure. It has multiple InfiniBand, Ethernet, and Ethernet over InfiniBand (`bondeth*`) network interfaces. This is the list of interfaces:

```
[root@mynode ~]# ip -o -f inet addr show
1: lo      inet 127.0.0.1/8
2: eth0    inet 12.17.207.156/21
3: eth1    inet 16.10.145.12/21
19: bondeth0  inet 12.17.128.15/20
20: bondeth2  inet 16.10.230.160/20
21: bondeth4  inet 192.168.1.45/20
30: bondib0   inet 192.168.31.178/21
31: bondib1   inet 192.168.129.2/21
32: bondib2   inet 192.168.199.205/21
33: bondib3   inet 192.168.216.31/21
34: bondib4   inet 192.168.249.129/21
```

When the Oracle Big Data SQL installer runs on this system, the following interfaces would be selected.

- 192.168.31.178/21 is selected for the InfiniBand connection configured in `cellinit.ora`.
Among the InfiniBand interfaces on this system, the interface name `bondib0` is first in an ascending alphabetical sort.
- 12.17.128.15/20 is selected for an Ethernet-over-InfiniBand connection configured in `celliniteth.ora`.

 **Note:**

This example demonstrates an additional selection factor – Ethernet over InfiniBand takes precedence over standard Ethernet. The `bondeth0` interface name is first in this case.

How the Installation (or SmartScan) may Fail Under These Conditions

It is possible that `diskmon` cannot connect to the Oracle Big Data SQL cells via the network interface selected according to the logic described above. The correct subnet (one that `diskmon` can reach) may not appear first in an alphabetical sort.

How to Fix the Problem

You can manually change the IP addresses in the `cellinit.ora` and `celliniteth.ora` files. These files are at `/etc/oracle/cell/network-config` on each node.

1. Stop the CRS process. (Be sure to do this before the cell edit. If you do not, `diskmon` may hang.)
2. Edit `cellinit.ora` and/or `celliniteth.ora`. Change the IP addresses as needed.
3. Restart CRS.

3.2 About the Database-Side Installation Directory

You start the database side of the installation by unpacking the database-side installation bundle and executing the run file it contains. The run file creates an installation directory under `$ORACLE_HOME/BDSJaguar-3.2.0`. For example:

```
$ORACLE_HOME/BDSJaguar-3.2.0/cdh510-6-nodel.mycluster.mydomain.com
```

The installation of Oracle Big Data SQL is not finished when this directory is created. The directory is a staging area that contains all of the files needed to complete the installation on the database node.

There can be Oracle Big Data SQL connections between the database and multiple Hadoop clusters. Each connection is established through a separate database-side installation and therefore creates a separate installation directory. The segments in the name of the installation directory enable you to identify the Hadoop cluster in this specific connection:

```
<Hadoop cluster name>-<Number nodes in the cluster>-<FQDN of the  
cluster management server node>
```

You should keep this directory. It captures the latest state of the installation and you can use it to regenerate the installation if necessary. Furthermore, if in the future you need to adjust the database-side of the installation to Hadoop cluster changes, then the updates generated by the `Jaguar reconfigure` and/or `updatenodes` commands are applied to this directory.

Consider applying permissions that would prevent the installation directory from being modified or deleted by any user other than `oracle` (or other database owner).

3.3 Steps for Installing on Oracle Database Nodes

To install the database side of Oracle Big Data SQL, copy over the ZIP file containing the database-side installation bundle that was created on the Hadoop cluster management server, unzip it, execute the run file it contains, then run the installer.

Perform the procedure in this section as the database owner (`oracle` or other). You stage the bundle in a temporary directory, but after you unpack the bundle and execute the run file it contains, then the installation package is installed in a subdirectory under `$ORACLE_HOME`. For example: `$ORACLE_HOME/BDSJaguar-3.2.0/cdh510-6-node1.mycluster.mydomain.com`.

Be starting, be sure that `ORACLE_HOME` and `ORACLE_SID` are set correctly.

Copy Over the Components and Prepare for the Installation

1. If you have not already done so, use your preferred method to copy over the database-side installation bundle from the Hadoop cluster management server to the Oracle Database server. If there are multiple bundles, be sure to select the correct bundle for cluster that you want to connect to Oracle Database. Copy it to any location that you would like to use as a temporary staging area. To perform the copy operation, it may be easiest for you to log on to the Hadoop cluster management server as `root`, navigate down to the bundle location within the Oracle Big Data SQL installation directory, and push the bundle over to the database server. Use the database owner account (usually `oracle`) for the remote logon on the database side. For example:

```
# cd <Big Data SQL Install Directory>/BDSJaguar/db-bundles
# scp bds-3.2.0-db-<cluster>-<yymmdd.hhmi>.zip
oracle@<database_node>:/opt/tmp
```

2. If you generated a database request key, then also copy that key over to the Oracle Database server.

```
# cd <Big Data SQL Install Directory>/BDSJaguar/dbkeys
# scp <database name or other name>.reqkey
oracle@<database_node>:/opt/tmp
```

3. Log on to the database server host as `oracle` (or whichever user is the database owner) and `cd` to the directory where you copied the file (or files).
4. Unzip the bundle. You will see that it contains a single, compressed run file.
5. Check to ensure that the prerequisite environment variables are set – `$ORACLE_HOME` and `$ORACLE_SID`.
6. Run the file in order to unpack the bundle into `$ORACLE_HOME`. For example:

```
$ ./bds-3.2.0-db-cdh510-170309.1918.run
```

Because you can set up independent Oracle Big Data SQL connections between an Oracle Database instance and multiple Hadoop clusters, the `run` command unpacks the bundle to a cluster-specific directory under `$ORACLE_HOME/BDSJaguar-3.2.0`. For example:

```
$ ls $ORACLE_HOME/BDSJaguar-3.2.0
cdh510-6-node1.mycluster.mydomain.com
test1-3-node1.myothercluster.mydomain.com
```

If the `BDSJaguar-3.2.0` directory does not already exist, it is created as well.

7. If you generated a database request key, then copy it into the newly created installation directory. For example:

```
$ cp /opt/tmp/mydb.reqkey $ORACLE_HOME/BDSJaguar-3.2.0/cdh510-6-
node1.mycluster.mydomain.com
```

 **Tip:**

You also have the option to leave the key file in the temporary location and use the `--reqkey` parameter in the installation command in order to tell the script where the key file is located. This parameter lets you specify a non-default request key filename and/or path.

However, the install script only detects the key in the installation directory when the key filename is the same as the database name. Otherwise, even if the key is local, if you gave it a different name then you must still use `--reqkey` to identify to the install script.

Note that the installation directory for the cluster is in place, you are ready to install the database side of Oracle Big Data SQL.

Install Oracle Big Data SQL on the Oracle Database Node

 **Important:**

The last part of the installation may require a single restart of Oracle Database under either or both of these conditions:

- If Oracle Database does not include the Oracle Grid Infrastructure. In this case, the installation script makes a change to the `pfile` or `spfile` configuration file in order to support standalone operation of `diskmon`.
- If there are changes to the IP address and the communication protocol recorded in `cellinit.ora`. These parameters define the connection to the cells on the Hadoop cluster. For example, if this is a re-installation and the IP address for the Hadoop/Oracle Database connection changes from an Ethernet address to an InfiniBand address and/or the protocol changes (between TCP and UDP), then a database restart is required.

1. Start the database if it is not running.

2. Log on to the Oracle Database node as `oracle` and change directories to the cluster-specific installation directory under `$ORACLE_HOME/BDSJaguar-3.2.0`. For example:

```
$ cd $ORACLE_HOME/BDSJaguar-3.2.0/cdh510-6-node1.my.domain.com
```

3. Run `bds-database-install.sh`, the database-side Oracle Big Data SQL installer. You may need to include some optional parameters.

```
$ ./bds-database-install.sh [options]
```

4. Restart the database (optional).

See Also :

The `bda-database-install.sh` installer command supports parameters that are ordinarily optional, but may be required for some configurations. See the [Command Line Parameter Reference for bds-database-install.sh](#)

Extra Step If You Enabled Database Authentication or Hadoop Secure Impersonation

- If `database_auth_enabled` or `impersonation_enabled` was set to “true” in the configuration file used to create this installation bundle, copy the ZIP file generated by the database-side installer back to the Hadoop cluster management server and run the Jaguar “Database Acknowledge” operation. This completes the set up of login authentication between Hadoop cluster and Oracle Database.

Find the ZIP file that contains the GUID-key pair in the installation directory. The file is named according to the format below.

```
<name of the Hadoop cluster>-<number nodes in the cluster>-<FQDN of the node where the cluster management server is running>-<FQDN of this database node>.zip
```

For example:

```
$ ls $ORACLE_HOME/BDSJaguar-3.2.0/cdh510-6-node1.my.domain.com/
*.zip
$ mycluster1-18-mycluster1node03.mydomain.com-
myoradb1.mydomain.com.zip
```

1. Copy the ZIP file back to `/opt/oracle/DM/databases/conf` on the Hadoop cluster management server.
2. Log on to the cluster management server as `root`, `cd` to `/BDSJaguar` under the directory where Oracle Big Data SQL was installed, , and run `databaseack`

(the Jaguar “database acknowledge” routine). Pass in the configuration file that was used to generate the installation bundle (bds-config.json or other).

```
# cd <Big Data SQL Install Directory>/BDSJaguar
# ./jaguar databaseack bds-config.json
```

3.3.1 Command Line Parameter Reference for bds-database-install.sh

Table 3-1 Optional Parameters for bds-database-install.sh


Parameter	Function
--install	Install the Oracle Big Data SQL connection to this cluster.
<div style="border: 1px solid #0070C0; padding: 10px; margin: 10px 0;"> <p> Note:</p> <p>The --install parameter is now required in order to do an installation with bds-database-install.sh. This is different from releases of Oracle Big Data SQL earlier than 3.2, in which the parameter is an implicit default.</p> </div>	
--version	Show the bds-database-install.sh script version.
--info	Show information about the cluster, such as the cluster name, cluster management server host, and the web server.
--grid-home	Specifies the Oracle Grid home directory.
--crs	Use crsctl to set up Oracle Big Data SQL MTA extprocs. Ignored in non-Grid environments.
	--crs={true false}
	The installer always checks to verify that Grid is running. If Grid is not running, then the installer assumes that Grid is not installed and that the database is single-instance. It then automatically sets the crs flag to false.
	If --crs=true is explicitly set and Grid cannot be found, the installer terminates with an error message stating that GI_HOME must be set.
--cdb	Create database objects on all PDBs for CDB databases.
	--cdb={true false}

Table 3-1 (Cont.) Optional Parameters for bds-database-install.sh

Parameter	Function
<code>--ip-cell</code>	<p>IP address used by the diskmon process in single-instance database environments where there are several network interfaces.</p> <p>Ignored on non-commodity hardware (Oracle Exadata Database Machine).</p> <p><code>--ip-cell=<IP address></code></p>
<code>--db-resource</code>	<p>Deprecated.</p> <p>Oracle Big Data SQL scripts will run on the database specified by <code>ORACLE_SID</code> or <code>ORACLE_HOME</code> environment variables.</p> <p><code>--db-resource=\$ORACLE_SID</code></p>
<code>--db-name</code>	<p>Deprecated.</p> <p>Oracle Big Data SQL scripts will run on the database specified by <code>ORACLE_SID</code> or <code>ORACLE_HOME</code> environment variables.</p> <p><code>--db-name=\$ORACLE_HOME</code></p>
<code>--reqkey</code>	<p>This parameter tells the installer the name and location of the request key file. Database Authentication is enabled by default in the configuration. Unless you do not intend to enable this feature, then as one of the steps to complete the configuration, you must use <code>--reqkey</code>.</p> <ul style="list-style-type: none"> To provide the key file name and the path if the key file is not in the local directory: <pre>\$./bds-database-install.sh --reqkey=/opt/tmp/some_name.reqkey --install</pre> To provide the key file name if the file is local, but the filename is not the same as the database name: <pre>\$./bds-database-install.sh --reqkey=some_name.reqkey --install</pre> <p>If the request key filename is provided without a path, then the key file is presumed to be in the installation directory. The installer will find the key if the filename is the same as the database name.</p> <p>For example, in this case the request key file is in the install directory and the name is the same as the database name.</p> <pre>\$./bds-database-install.sh --install</pre> <p>This file is consumed only once on the database side in order to connect the first Hadoop cluster to the database. Subsequent cluster installations to connect to the same database use the configured key. You do not need to resubmit the <code>.reqkey</code> file.</p>
<code>--uninstall</code>	Uninstall the Oracle Big Data SQL connection to this cluster.

Table 3-1 (Cont.) Optional Parameters for bds-database-install.sh

Parameter	Function
<code>--reconfigure</code>	Reconfigures <code>bd_cell</code> network parameters, Hadoop configuration files, and Big Data SQL parameters. The Oracle Big Data SQL installation to connect to this cluster must already exist. Run <code>reconfigure</code> when a change as occurred on the Hadoop cluster side, such as a change in the DataNode inventory.
<code>--databaseack-only</code>	Create the Database Acknowledge zip file. (You must then copy the zip file back to <code>/opt/oracle/DM/databases/conf</code> on the Hadoop cluster management server and run <code>./jaguar databaseack <configuration file></code> .)
<code>--mta-restart</code>	MTA extproc restart. (Oracle Big Data SQL must already be installed on the database server.)
<code>--mta-setup</code>	Set up MTA extproc with no other changes. (Oracle Big Data SQL must already be installed on the database server.)
<code>--mta-destroy</code>	Destroy MTA extproc and make no other changes. (Oracle Big Data SQL must already be installed on the database server.)
<code>--aux-run-mode</code>	Because Oracle Big Data SQL is installed on the database side as a regular user (not a superuser), tasks that must be done as <code>root</code> and/or the Grid user require the installer to spawn shells to run other scripts under those accounts while <code>bds-database-install.sh</code> is paused. The <code>--aux-run-mode</code> parameter specifies a mode for running these auxiliary scripts. <code>--aux-run-mode=<mode></code> Mode options are: <ul style="list-style-type: none"> • <code>session</code> — through a spawned session. • <code>su</code> — as a substitute user. • <code>sudo</code> — through <code>sudo</code>. • <code>ssh</code> — through secure shell.
<code>--root-script</code>	Enables or disables the startup root script execution. <code>--root-script={true false}</code>
<code>--no-root-script</code>	Disables startup root script execution.

The `--root-script-only` parameter from previous releases is obsolete.

3.4 Granting User Access

In Oracle Big Data SQL releases prior to 3.1, access is granted to the `PUBLIC` user group. In the current release, you must do the following for each user who needs access to Oracle Big Data SQL:

- Grant the `BDSQL_USER` role.
- Grant read privileges on the `BigDataSQL` configuration directory object.

For example, to grant access to user1:

```
SQL> grant BDSQL_USER to user1;  
SQL> grant read on directory ORACLE_BIGDATA_CONFIG to user1;
```

 **See Also:**

Use the DBMS_BDSQL PL/SQL Package to indirectly provide users access to Hadoop data. The Multi-User Authorization feature that this package implements uses Hadoop Secure Impersonation to enable the `oracle` account to execute tasks on behalf of other designated users.

4

Expanding or Shrinking an Installation

Changes that occur to the Hadoop cluster over time, such as network reconfiguration or DataNode services added, moved, or removed, as well as changes on the Oracle Database server, can require corresponding changes to the Oracle Big Data SQL installation.

4.1 Adding or Removing Oracle Big Data SQL on Hadoop Cluster Nodes

On the Hadoop side, Oracle Big Data SQL must be installed on every cluster node where the DataNode role is running and should not be installed on other nodes. Note that when you run `jaguar updatenodes` as described below, you must also reconfigure the Oracle Database side of the installation.

The `Jaguar updatenodes` procedure uses the cluster management server to determine where the DataNode service is running and then installs or uninstalls Oracle Big Data SQL on each cluster node as needed.

Run `updatenodes` when the distribution of DataNode services has changed in the Hadoop cluster.

```
# cd <Big Data SQL install directory>/BDSJaguar
# ./jaguar updatenodes <bds-config.json>
```

When the update is done, this message is returned to the console:

```
BigDataSQL: UPDATENODES workflow completed.
```

The following are situations where you should run `updatenodes`:

- New nodes are added to the Hadoop cluster and some or all of them are provisioned with the DataNode role.
- Services and roles are redistributed in the cluster and some DataNode instances are moved.
- One or more nodes provisioned with the DataNode role, are removed from the cluster or deprovisioned.

! Important:

Be sure to run `updatenodes` after a cluster expansion. This will enable local processing for data stored on the updated nodes. Also, the new DataNode instances in the cluster provide additional resource for query processing that is not utilized until Oracle Big Data SQL is installed on those nodes. You must also run `bds-database-install.sh --reconfigure` on the database side after you run `updatenodes` on the Hadoop side.

✎ See Also:

The [Jaguar Configuration Parameter and Command Reference](#) provides a full description of the Jaguar `updatenodes` and `reconfigure` commands.

Communicating DataNode Changes to Oracle Database

The Oracle Database side of the installation maintains an inventory of the DataNodes where Oracle Big Data SQL is installed in the Hadoop cluster. This is a map that tells it which nodes to query. If you run `updatenodes`, this updated DataNode inventory must be provided to the database compute nodes. The `updatenodes` procedure generates a new database installation bundle that includes the new inventory. Install this updated bundle on the database compute nodes using the same method that you use in a full installation, with one difference – on the Oracle Database side, install the update using `bds-database-install.sh --reconfigure`

1. As the database owner, copy the bundle to temporary location on the database node.
2. Extract the contents of the ZIP file.
3. Execute the run file extracted from the ZIP file.
4. Go to the installation directory under `$ORACLE_HOME/BDSJaguar-3.2.0`.
5. Run the installer with the `--reconfigure` command line parameter. For example:

```
$ cd $ORACLE_HOME/BDSJaguar-3.2.0/cdh510-6-node1.mydomain.com
$ ./bds-database-install.sh --reconfigure
```

4.2 Adding or Removing Oracle Big Data SQL on Oracle Database Nodes

Add or remove Oracle Big Data SQL from Oracle Database nodes by manually running the install or uninstall program on individual database nodes as needed.

On the Oracle Database side, there is at this time no equivalent to the Hadoop-side Jaguar `updatenodes` command, which detects changes in the Hadoop cluster and installs or uninstalls Oracle Big Data SQL on cluster nodes as required.

To Install Oracle Big Data SQL on Additional Nodes of Oracle Database RAC System

If you want to extend Oracle Big Data SQL from a Hadoop cluster to additional Oracle Database nodes at any time after the initial installation of Oracle Big Data SQL, then install the database-side installation bundle for that cluster on these database nodes using the same configuration applied to other nodes of the database.

Note:

The configuration for the connection between a Hadoop cluster and Oracle Database must be the same for all nodes of the database.

- You may find the bundle at `<Big Data SQL Install Directory>/BDSJaguar/db-bundles` or at the temporary location on other database nodes where you copied the bundle in order to stage the installation. The bundle ZIP file should include the Hadoop cluster name and a time stamp: `bds-3.2.0-db-<cluster>-<yymmdd.hhmi>.zip`
- If you can no longer locate the original database-side installation bundle, but do have the configuration file that was used for other nodes of the database, then you can recreate the same bundle by running the Jaguar `reconfigure` command.

```
[root@myclusteradminserver:BDSjaguar]# ./jaguar reconfigure  
my_original_configfile.json
```

Deploy and install the bundle on the new database nodes as described in [Steps for Installing on Oracle Database Nodes](#).

To Uninstall Oracle Big Data SQL from Oracle Database Nodes

See [Uninstalling From an Oracle Database Server](#).

5

Reconfiguring an Installation

After the initial Oracle Big Data SQL installation, there may be changes on the Hadoop cluster or the Oracle Database system that require corresponding updates to the Oracle Big Data SQL configuration, including:

- A change in the status of Kerberos — if it was previously disabled and is now enabled, or vice versa.
- Implementation of TLS (Transport Layer Security) or other security measures, such as Database Authentication option now supported by Oracle Big Data SQL.
- Changes to the location of Hive service instances within the Hadoop cluster.
- A switch from Ethernet to InfiniBand (or InfiniBand to Ethernet) for networking between the Hadoop cluster and the Oracle Database system.
- A change to the version of Oracle Big Data SQL running on the Hadoop cluster.

Important:

If you are storing Oracle tablespaces in HDFS, then after you run `reconfigure`, take the tablespaces offline (or stop the database) and then restart the FUSE-DFS service running on all database compute nodes that are connected to the Hadoop cluster. See *Restarting FUSE-DFS After a Reconfigure Operation* at the end of this section for details.

The basic procedure for reconfiguring an installation is:

- Edit the `bds-config.json` file as needed.
- Run the Jaguar `reconfigure` command to push the configuration changes to the Hadoop cluster nodes and to generate an updated database-side installation bundle.
- Reinstall Oracle Big Data SQL on Oracle Database compute nodes using the new installation bundle. You can just overwrite the current installation. There is no prerequisite to uninstall the current Oracle Big Data SQL installation on either side.

Details are provided below.

See Also:

The addition or removal of Oracle Big Data SQL from Hadoop cluster nodes or nodes in a multinode Oracle Database system is not a considered configuration change. See [Expanding or Shrinking an Installation](#).

Run “jaguar reconfigure” to Reconfigure the Hadoop Side of the Installation and to Generate an Updated Database Installation Bundle

1. Log on to the cluster management server node as `root` and `cd` to the directory where you extracted the downloaded Oracle Big Data SQL installation bundle.
2. Go to the `BDSJaguar` subdirectory under the path where you unzipped the Oracle Big Data SQL bundle .

```
# cd <Big Data SQL Install Directory>/BDSJaguar
```

3. In the `BDSJaguar` directory, run the Jaguar utility. Pass it the `reconfigure` parameter and the configuration file name (`bds-config.json` or other) as arguments. You may or may not need to include the `--requestdb` option.

```
[root@myclusteradminserver:BDSjaguar] # ./jaguar reconfigure bds-config.json
```

If `bds-config.json` sets `database_auth_enabled` to “true”, then include the `--requestdb` option in the install command:

```
[root@myclusteradminserver:BDSjaguar] # ./jaguar --requestdb testdb reconfigure
```

The script prompts for the cluster management service administrator credentials and then reconfigures Oracle Big Data SQL on the management server and the cluster nodes according to the new settings in the configuration file. The script output terminates with the following message if the installation completed without error.

```
BigDataSQL: RECONFIGURE workflow completed.
```

Copy the Updated Database Installation Bundle to Each Oracle Database Node and Install It

1. Log on as either `root` or `oracle` on the cluster management server and `cd` to the directory where the new database bundle was generated.

```
# cd <Big Data SQL Install Directory>/BDSJaguar/db-bundles
```

Check that the new bundle is there.

```
[oracle@myclusteradminserver:db-bundles] # ls
[oracle@myclusteradminserver:db-bundles] # bds-3.2.0-db-<cluster>-<yymmdd.hhmi>.zip
```

2. Copy the database installation bundle over to any directory on the Oracle Database node, using the `oracle` account on the Oracle system. If there are

multiple bundles, be sure to select the correct bundle for cluster that you want to connect to Oracle Database.

```
[oracle@myclusteradminserver:db-bundles] # scp bds-3.2.0-db-
<cluster>-<yymmdd.hhmi>.zip oracle@<database_node>:/opt/tmp
```

3. Log on to the Oracle Database node as the database owner (`oracle` or other) in order perform the installation.
4. Cd to the directory where you copied the bundle, unzip the bundle, and then run the compressed executable that was extracted from the zip file. For example:

```
$ unzip bds-3.2.0-db-cdh510-170309.1918.zip
$ ./bds-3.2.0-db-cdh510-170309.1918.run
```

5. Change directories to the new directory under `ORACLE_HOME` and run `bds-database-install.sh`, the database-side installer.

```
$ cd $ORACLE_HOME/BDSJaguar-3.2.0/cdh510-6-node1.my.domain.com
$ ./bds-database-install.sh
```

Restarting FUSE_DFS After a Reconfigure Operation

The Oracle Tablespace in HDFS feature uses a FUSE-DFS connection for file transfers between Oracle Database and the Hadoop cluster. There is a FUSE-DFS service instance for each connection to cluster. The service is named as follows:

```
fuse_<cluster name>_hdfs
```

If this service exists, then after a reconfiguration of an Oracle Big Data SQL installation, then do the following:

1. Either set the tablespaces in HDFS to offline or stop the database.
The datafiles for these tablespaces are open and locked by database processes.
2. Restart FUSE-DFS. Do this on each compute node.

```
# service fuse_mycluster_hdfs restart
```

3. After FUSE-DFS has restarted, you can set the tablespaces to online (or you if you stopped the database, you can restart it).

The `status`, `stop`, and `start` operations are also available for the FUSE-DFS service, as in:

```
# service fuse_mycluster_hdfs status
```

5.1 Reconfiguring an Existing Oracle Big Data SQL Installation

When Oracle Big Data SQL is installed on the cluster management server, the database-side installation bundle that is generated contains configuration information

that is based on the current state of the cluster. When you install the bundle on the Oracle Database side, this configuration information is used to set up the connection between the Hadoop cluster and the database server. After the installation, changes in the Hadoop cluster may require an update to the Oracle Big Data SQL configuration both sides (Hadoop cluster and database server) in order to maintain connectivity.

Here are some Hadoop-side changes which require you to reconfigure the installation:

- Network security changes, such as a switch from HTTP to HTTPS or vice versa.
- Port reassignments in the cluster management service (CM or Ambari).
- Hive changes, such as migration of the service to different node or nodes or a change to the authentication method.
- Changes to other related Hadoop services, such as HDFS, MapReduce2, and YARN, may also affect connectivity.

 **Note:**

The `reconfigure` procedure cannot be used to update the installation of an older version of Oracle Big Data SQL. You should instead overwrite the old installation with the current Oracle Big Data SQL version.

Do not use this procedure if the only changes are to the DataNode inventory in the cluster (if DataNodes, are moved, added, or removed). In that case, use the `updatenodes` procedures instead.

1. As root on the Hadoop cluster management server, make any necessary updates to the configuration file for the connection between the Hadoop cluster and Oracle Database (`bds-config.json` or other).
2. Go to the `BDSJaguar` directory and run `Jaguar` to invoke the `reconfigure` operation. Pass in the configuration filename.

```
# cd <Big Data SQL install directory>/BDSJaguar
# ./jaguar reconfigure bds-config.json
```

This updates the Oracle Big Data SQL installation on the Hadoop cluster and generates a new database-side installation bundle, which you can find at `<Oracle Big Data SQL Install Directory>/BDSJaguar/dbbundles`.

 **Note:**

If you are using Database Authentication and already provided a request key to the database-side installer when you did the original installation, then you do not need to do again. That is, you do not need to include `--requestdb` in a `reconfigure` operation unless you want to set up a connection to another database.

3. As in the original installation, copy the new bundle file to a temporary directory each database node. Connect using the `oracle` account on the database node. For example:

```
# scp bds-3.2.0-db-<cluster>-<yymmdd.hhmi>.zip
oracle@dbnode:/opt/tmp
```

4. Log on to the database node as the `oracle` user.
5. Locate the bundle and unzip it.
6. Run the executable file extracted from the bundle.

```
[oracle@mydbserver: $ bds-3.2.0-db-<cluster>-<yymmdd.hhmi>.run
```

This extracts the installation files into a new directory under `ORACLE_HOME`.

7. Change directories to `ORACLE_HOME` and list the subdirectories. You should see the old installation directory along with the new one. The directory names will include the same cluster name, but will have different timestamps. Delete the old directory.

```
$ cd $ORACLE_HOME
$ ls
  bds-3.2.0-db-<cluster>-<new timestamp>
  bds-3.2.0-db-<cluster>-<old timestamp>
$ rm -rf bds-3.2.0-db-<cluster>-<old timestamp>
```

8. Change directories to the new directory.

Run `bds-database-install.sh` with the `--reconfigure` parameter.

```
$ ./bds-database-install.sh --reconfigure
```

6

Uninstalling Oracle Big Data SQL

If you want to completely uninstall Oracle Big Data SQL, the software must be removed from the Hadoop cluster management server and the Oracle Database nodes where it is currently installed.

If the Hadoop DataNode service is removed from a subset of the nodes, then Oracle Big Data SQL must be removed from those nodes only.

6.1 General Guidelines for Removing the Software

Oracle Big Data SQL can be uninstalled from the Hadoop cluster management server or from any Oracle Database servers connected to the cluster management server. The procedure is the same for all Hadoop platforms.

Guidelines for uninstalling Oracle Big Data SQL are as follows:

- To perform a complete uninstall of Oracle Big Data SQL, remove the software from the cluster management server and from each Oracle Database server connected to the BDS service.

This is a single script execution on each server. Not other manual steps are needed.

- You can uninstall from the cluster management server first or from the database servers first.

Tip:

If you are removing the software from both sides of the installation, uninstall it from the database side first. If you uninstall from the cluster management server first, queries in process will fail.

- On the database side, the first Oracle Big Data SQL connection installed is known as the “default cluster.”

Note:

In a Oracle RAC database, be sure to run the uninstall script on all nodes of the database.

6.2 Uninstalling From an Oracle Database Server

On any database server where you want to uninstall the current version of Oracle Big Data SQL, run the commands below as the database owner (usually the `oracle` user).

In the case of a RAC database, be sure to perform the uninstall on each instance of the RAC.

Note that the `--crs` command line parameter may or may not be required, depending upon conditions related to Grid support. See the *Command Line Parameter Reference for bds-database-install.sh* in this guide.

If FUSE is enabled for TableSpaces in HDFS, Uninstall it First

If you had previously set up FUSE to store Oracle tablespaces in HDFS, then you must first run the following command before uninstalling Oracle Big Data SQL:

```
$ ./bds-copy-tbs-to-hdfs.sh --uninstall
```

This command removes the FUSE mount point. You can then proceed with the Oracle Big Data SQL uninstall.

Uninstalling Oracle Big Data SQL

Run `bds-database-install.sh` with the `uninstall` command.

```
$ /bds-database-install.sh --uninstall --crs=false
```

Note:

If Grid is not running on this database node, or, if the database does not use Grid (CRS/ASM), then include the optional `--crs` parameter and set it to `false`.

Also note that the `--uninstall-as-primary` and `--uninstall-as-secondary` parameters from previous releases are deprecated in this release. It is no longer necessary to differentiate between primary and secondary clusters in the uninstall process. In addition, `--db-resource` has been deprecated and is therefore not used in the example provided here.

6.3 Uninstalling From a Hadoop Cluster

Oracle Big Data SQL works with the Hadoop DataNode service. When you run the Oracle Big Data SQL installer on the configuration management server, it automatically propagates the Oracle Big Data SQL components to all Hadoop cluster nodes where the DataNode service is running. If the DataNode service is removed from one or more nodes, then you must also remove the Oracle Big Data SQL components from those same nodes.

1. On the cluster management server (where CM or Ambari is running), log on as `root`.
2. Change directories to the `BDSjaguar` directory under the location where the Oracle Big Data SQL Installation bundle `BigDataSQL-<Hadoop distribution><version>.zip` was extracted.

3. Run Jaguar with the `uninstall` command and provide the configuration file previously used to install on the cluster (`bds-config.json` or other):

```
[root@myclusteradminserver:BDSjaguar] # ./jaguar uninstall bds-config.json
```

When the uninstall is complete, the following message is returned:

```
BigDataSQL: UNINSTALL workflow completed.
```

7

Securing Big Data SQL

This section describes security features provided by Oracle Big Data SQL, measures you can take to secure Big Data SQL, and to pointers to the information you need in order to configure Oracle Big Data SQL within secured environments.

7.1 Security Overview

Network Security

In Oracle Big Data SQL, network traffic between the database and the Hadoop cluster, is no longer guaranteed to be over a private InfiniBand network, but can occur over a client network. This network traffic is not currently secured. Therefore when operating a secured Hadoop cluster (e.g., Kerberos-enabled, RPC encryption), Oracle Big Data SQL requires either that all members of the client network be trusted, or that private network connectivity is used exclusively for communication between the Hadoop nodes and Oracle database instances. This private network is commonly referred to as the *Big Data SQL interconnect network*. The interconnect network must be a private network with only trusted users, use at least one switch, and 10 Gigabit Ethernet adapters. Ensure that only the nodes in the Hadoop cluster and Oracle RAC cluster can access the interconnect network. Do not use the interconnect network for other user communication.

Installer File Security

The new Jaguar installer incorporates the following best practices for secure Linux installers and applications:

- No persistent or temporary world-writable files are created.
- No setuid or setgid files are used.

In addition, the installer works with hardened Oracle Database environments as well as hardened CDH and HDP clusters as described in the Cloudera CDH and Hortonworks HDP security documentation

Password Safety

The Jaguar installer provides these password security measures:

- Passwords for the Ambari and Cloudera Manager servers (management servers) are not be passed in on the command line and are not be saved on any persistent file during and/or after the installation is complete.
- Passwords are not be logged to any log or trace files.

Security of Related Software

Oracle Big Data SQL relies on other installed software, including third party projects. It is the customer's responsibility to ensure that such software is kept up-to-date with the latest security patches. This software includes is (but not limited to):

- Java JDK
- Python
- Perl

7.2 Big Data SQL Communications and Secure Hadoop Clusters

It is generally a good security practice to ensure that HDFS file access permissions are minimized in order to prevent unauthorized write/read access. This is true regardless of whether or not the Hadoop cluster is secured by Kerberos.

Please refer to MOS Document 2123125.1 at [My Oracle Support](#) for detailed guidelines on securing Hadoop clusters for use with Oracle Big Data SQL.

7.3 Configuring Oracle Big Data SQL in a Kerberos-Secured Environment

If Kerberos is enabled on the Hadoop system, you must configure Oracle Big Data SQL on the database server to work with Kerberos. This requires a Kerberos client on each database node where Oracle Big Data SQL is installed. Also, the OS account that owns the database (`oracle` or another account) must be provisioned as a user principal.

When operating a secured Hadoop cluster (e.g., Kerberos enabled, RPC encryption), Oracle Big Data SQL requires either that all members of the client network be trusted, or that private network connectivity is used exclusively for communication between the Hadoop nodes and Oracle database instances. This private network is commonly referred to as the Big Data SQL interconnect network. The interconnect network must be a private network with only trusted users, use at least one switch, and 10 Gigabit Ethernet adapters. Ensure that only the nodes in the Hadoop cluster and Oracle RAC cluster can access the interconnect network. Do not use the interconnect network for other user communication.

7.3.1 Enabling Oracle Big Data SQL Access to a Kerberized Cluster

You must configure Oracle Big Data SQL to use Kerberos in environments where user access is Kerberos-controlled.

There are two situations when this is required:

- When enabling Oracle Big Data SQL on a Kerberos-enabled cluster.
- When enabling Kerberos on a cluster where Oracle Big Data SQL is already installed.

Oracle Big Data SQL processes run on the nodes of the Hadoop cluster as the `oracle` Linux user. On the Oracle Database server, the owner of the Oracle Database process is also (usually) the `oracle` Linux user. When Kerberos is enabled on the Hadoop system, the following is required in order to give the user access to HDFS.

- The `oracle` Linux user needs to be able to authenticate as a principal in the Kerberos database on the Kerberos Key Distribution Center (KDC) server. The

principal name in Kerberos does not have to be 'oracle'. However, the principal must have access to the underlying Hadoop data being requested by Oracle Big Data SQL.

- The following are required on all Oracle Database nodes and all Hadoop cluster nodes running Oracle Big Data SQL:
 - Kerberos client software installed.
 - A copy of the Kerberos configuration file from the KDC.
 - A copy of the Kerberos keytab file generated on the KDC for the `oracle` user.
 - A valid Kerberos ticket for the `oracle` Linux user.

Installing the Kerberos Client

If the Kerberos client is not installed, see [Installing a Kerberos Client on the Oracle Database Nodes](#) for instructions on installing the Kerberos client.

Creating a Kerberos Principal for the oracle User

On the Kerberos Key Distribution Center (KDC) server, become `root` and use `kadmin.local` to add a principal for the `oracle` user.

1. # `kadmin.local`

2. Within `kadmin.local`, type:

```
add_principal <user>@<realm>
quit
```

You have the option to include the password, as in:

```
add_principal <user>@<realm> -pw <password>
quit
```

Creating a Kerberos Keytab for the oracle User

1. On the KDC, become `root` and run the following:

```
# kadmin.local
```

2. Within `kadmin.local`, type:

```
xst -norandkey -k /home/oracle/oracle.keytab oracle
quit
```

This creates the `oracle.keytab` file for the Kerberos `oracle` user in the `/home/oracle` directory.

3. Ensure that `oracle.keytab` is owned by the `oracle` Linux user and is readable by that user only.

```
$ chown oracle oracle.keytab
$ chmod 400 oracle.keytab
```

Kerberos Tasks Automated by Oracle Big Data SQL

The following Kerberos tasks are now automated:

- Distributing Keytab and Kerberos configuration file distribution files.
The Oracle Big Data SQL installation can now be configured to automatically distribute the keytab and Kerberos configuration files for the `oracle` user or other database owner to the Hadoop DataNodes (and Oracle Database compute nodes). This is done if the principal name and keytab file location parameters are set in the Jaguar configuration file. This automation is performed on both the Hadoop and Oracle Database side.
On Oracle Big Data Appliance, the keytab file distribution is done default for the `oracle` account and you do not need to add the principal and keytab file path for this account to the configuration file.
- Acquiring a Kerberos Ticket for designated principals.
For `oracle` and other principals that were listed in the Jaguar configuration file, the installation acquires a Kerberos ticket on each Hadoop DataNode and Oracle DB compute node
- Ticket renewal
The installation automatically sets up cron jobs in the Hadoop cluster and on Oracle Database to kinit for new ticket for each principal in the configuration four times daily.

Cleaning up After Ticket Expirations

When the `bd_cell` process is running on the nodes of a secured Hadoop cluster but the Kerberos ticket is not valid, then the cell goes to quarantine status. You should drop all such quarantines.

1. Check that the `oracle` user has a valid Kerberos ticket on all Hadoop cluster nodes.
2. On each cluster node, become `oracle` and run the following:

```
$ /opt/oracle/bd_cell/cellsrv/bin/bdscli
```

3. In the `bdscli` shell, type:

```
list quarantine
```

4. While still in `bdscli`, drop each quarantine on the list:

```
drop quarantine <id>
```

5. Type `exit` to exit `bdscli`.

7.3.2 Installing a Kerberos Client on the Oracle Database Nodes

If the Oracle Database system is Kerberos secured, then Oracle Big Data SQL requires a Kerberos client. The client must be installed on each compute node of the database.

For commodity servers, download the Kerberos client software from a repository of your choice. If the database server is an Oracle Exadata Database Machine, download and install the software from the Oracle repository as shown below. The process should be similar for downloads from non-Oracle repositories.

Log on to the database server as root and use yum to install the `krb5-libs` and `krb5-workstation` packages. Download from the Oracle Linux 6 or Oracle Linux 5 repository as appropriate.

1. Check that the Oracle `public-yum-ol6` or `public-yum-ol5` repository ID is installed.

```
# yum repolist
```

2. Temporarily disable all repository IDs and then enable the Oracle repository only (Oracle Linux 6 in this example).

```
# yum --disablerepo="*" --enablerepo="public-yum-ol6" list
available
```

3. Install the Kerberos packages.

```
# yum install krb5-libs krb5-workstation
```

4. Copy the `/etc/krb5.conf` file from the Key Distribution Center (KDC) to the same path on the database server.

These steps must be performed for each Oracle Database node.

You must also register the `oracle` Linux user (or other Linux user) and password in the KDC for the cluster as described in [Enabling Oracle Big Data SQL Access to a Kerberized Cluster](#)

7.4 Using Oracle Secure External Password Store to Manage Database access for Oracle Big Data SQL

On the Oracle Database server, you can use the Oracle Secure External Password Store (SEPS) to manage database access credentials for Oracle Big Data SQL.

This is done by creating an Oracle wallet for the `oracle` Linux user (or other database owner). An Oracle wallet is a password-protected container used to store authentication and signing credentials, including private keys, certificates, and trusted certificates needed by SSL.

See [MOS Document 2126903.1](#) at [My Oracle Support](#) for information on using SEPS with Oracle Big Data SQL.

7.5 About Data Security on Oracle Big Data Appliance

If your Hadoop system is an Oracle Big Data Appliance, the following tools to strengthen security are already available.

- **Kerberos authentication:** Requires users and client software to provide credentials before accessing the cluster.

- **Apache Sentry authorization:** Provides fine-grained, role-based authorization to data and metadata.
- **HDFS Transparent Encryption:** Protects the data on disk and at rest. Data encryption and decryption is transparent to applications using the data.
- **HTTPS/ Network Encryption:** Provides HTTPS for Cloudera Manager, Hue, Oozie, and Hadoop Web UIs. Also Enables network encryption for other internal Hadoop data transfers, such as those made through YARN shuffle and RPC.

 **See Also:**

The Oracle Big Data Appliance Software User's Guide provides details on available security features. You can find this guide and other documentation for your release of the Oracle Big Data Appliance software in the "Big Data" section of the [Oracle Help Center](#).

7.6 Authentication Between Oracle Database and Oracle Big Data SQL Offload Cell Server Processes

The Database Authentication feature described in this guide prevents unauthorized and potentially malicious processes (which can originate from anywhere) from connecting to Oracle Big Data SQL cell server processes in the DataNodes of the Hadoop cluster.

When Ethernet is selected for the connection between Oracle Databases and Oracle Big Data SQL, then by default this secured authentication framework is set up automatically during the installation. Database Authentication is also available as a configuration option for InfiniBand connections.

 **See Also:**

[Jaguar Configuration Parameter and Command Reference](#)

7.7 The Multi-User Authorization Model

Multi-User Authorization gives you the ability to use Hadoop Secure Impersonation to direct the `oracle` account to execute tasks on behalf of other designated users. This enables HDFS data access based on the user that is currently executing the query, rather than the singular `oracle` user.

Administrators set up the rules for identifying the query user (the currently connected user) and for mapping this user to the user that is impersonated. Because there are numerous ways in which users can connect to Oracle Database, this user may be a database user, a user sourced from LDAP, from Kerberos, or other sources. Authorization rules on the files apply to the query user and audits will identify the user as the query user.

 **See Also:**

The DBMS_BDSQL PL/SQL Package in the *Oracle Big Data SQL User's Guide* describes the Multi-User Authorization security table and the procedures for adding user maps to the table and removing them from the table.

8

Additional Tools Installed

8.1 Copy to Hadoop and OSHH

Copy to Hadoop

Copy to Hadoop makes it easy to identify and copy data from an Oracle Database table to the Apache Hadoop Distributed File System (HDFS) and create an Apache Hive external table over that data. In HDFS, the data can be saved in Oracle Data Pump files (which store data as Oracle data types). This provides an exact copy of the data in the database, without the potential for data type conversion errors. You can then query the data via Hive (and Hadoop technologies that access data through Hive) and by Oracle Big Data SQL. Oracle Data Pump data can be converted to Parquet or ORC using Hive.

Oracle Shell for Hadoop Loaders (OSHH)

OSHH is an intuitive command line tool to move data between Hadoop and Oracle Database. It provides a set of declarative commands you can use to copy contents from an Oracle Database table to a Hive table.

OSHH is bundled with both Oracle Big Data SQL and Oracle Big Data Connectors. With Oracle Big Data SQL, OSHH works with Copy to Hadoop. With Oracle Big Data Connectors, OSHH provides an interface to copy contents from Hadoop and Hive to Oracle tables using Oracle Loader for Hadoop (OLH) and Oracle SQL Connector for Hadoop Distributed File System (OSCH).

Although OSHH can serve as a front end for OLH and OSCH as well as for Copy to Hadoop, OLH and OSCH are part of the Oracle Big Data Connectors product and are not installed by Oracle Big Data SQL.

Installation and Configuration of Copy to Hadoop and OSHH

You can use Copy to Hadoop and OSHH from any of these environments:

- A Hadoop cluster node
- An Oracle Database node
- An edge node.

Oracle Big Data SQL installs the kits for Copy to Hadoop and Oracle Shell for Hadoop Loaders on the Hadoop cluster management server and on any Oracle Database server node where you install Oracle Big Data SQL.

You can also manually copy the kits for both tools from the Hadoop cluster management server to an edge node and configure them for use from the edge node. The kits are stored at `/opt/oracle/orahivedp-<version>..` See `/opt/oracle/ohsh-<version>/doc/OSHH_EDGE_NODE_README.txt` on the Hadoop cluster management server for instructions.

 **Note:**

In the Oracle Big Data SQL installation, the Copy to Hadoop is fully automated except for the setting of the `CP2HADOOP_HOME` listed in the table below. You must set this manually. You must complete configuration of Oracle Shell for Hadoop Loaders manually, as described in this section.

Environment Variables for OSHH

On the Hadoop cluster management server, the OSHH configuration file is `/opt/oracle/ohsh-<version>/bin/ohsh_config.sh`. On an Oracle Database node, it is `$ORACLE_HOME/ohsh/bin/ohsh_config.sh`. This file contains environment required by OSHH. The set of variables needed the installation on a Hadoop cluster, Oracle Database Node or an edge node differ. The table below shows which variables you need to set for each environment. Some of the required variables are set for you automatically by the Oracle Big Data SQL installer.

8.1.1 Completing the OSHH Configuration on Oracle Database Nodes

The Oracle Big Data SQL installation automatically copies the OSHH installation kit to each node where Oracle Big Data SQL is installed. There are some manual steps you must perform to complete the installation.

 **Note:**

The `README` file at `$ORACLE_HOME/ohsh/doc/OSSH_DB_NODE_README.txt` on the Hadoop cluster management server where you installed Oracle Big Data SQL describes the full set of steps ordinarily needed to enable OSHH on an Oracle Database node. However, when OSHH is installed by Oracle Big Data SQL, some of the setup is done for you by the installer. To complete the configuration in this context, use this shortened procedure instead of the full procedure documented in the `README`.

Also notice below that patch 23237037 is required for 12.1.0.2 (or earlier Oracle Database releases) only. Do not install this patch if you are running a later release of Oracle Database.

1. Log on to the Oracle Database server as the `oracle` user.
2. Set up OSHH to work with Copy to Hadoop.
 - a. Create a staging directory in the file system on the Oracle Database node:

```
$ mkdir <OS_STAGE_DIR>
```

- b. Create an Oracle directory on this staging directory and grant read/write privileges to the Oracle Database user (`oracle` or other). In this example,

`ORACLE_STAGE_DIR` refers to the name of the directory object in Oracle Database. `OS_STAGE_DIR` refers to the path on the server.

```
SQL> CREATE DIRECTORY <ORACLE_STAGE_DIR> AS '<OS_STAGE_DIR>';
SQL> GRANT READ,WRITE ON DIRECTORY <ORACLE_STAGE_DIR> TO
<ORACLE_USER>;
```

 **Note:**

Later, if you employ the “stage” method to copy Oracle tables, use the OHSB `set locationdirectory` command to point to the Oracle directory object name (the name of the directory you created above in this step).

- c. If you are connecting to a CDH-based Hadoop system (Oracle Big Data Appliance or other), perform the substeps below to enable the OHSB “FUSE” method, in which you copy files to Hive tables over an HDFS Mountable FUSE directory. Note that HDP does not support Fuse, therefore you cannot create a mount point to HDFS on HDP using these steps.
 - i. If you are using **Oracle Database 12.1.0.2 or earlier**, install patch 23237037 (available on [My Oracle Support](#)) on all database nodes.
 - ii. On the Oracle Database host, create a FUSE mount to the root of the HDFS system on the Hadoop cluster. Refer to Cloudera’s documentation on installing FUSE. You may find current instructions at [Configuring Mountable HDFS](#) on Cloudera’s website.
 - iii. Create the Oracle directory on the FUSE mount and grant privileges. In this example, `ORACLE_FUSE_DIR` refers to the name of the directory object in Oracle Database. `OS_FUSE_MOUNT` refers to the path to the mount point on the server.

```
SQL> CREATE DIRECTORY <ORACLE_FUSE_DIR> AS
'<OS_FUSE_MOUNT>';
SQL> GRANT READ,WRITE ON DIRECTORY <ORACLE_FUSE_DIR>
TO <ORACLE_USER>;
```

Note that when you use OHSB later to copy Oracle tables to Hive using the “FUSE” method, use the OHSB `set locationdirectory` command to point to the same `<ORACLE_FUSE_DIR>`

3. Edit `$ORACLE_HOME/ohsb/bin/ohsb_config.sh` in order to configure the home directories of dependencies. When OHSB is installed by Oracle Big Data SQL, most of the required variables are already set for you. Set `HS2_HOST_PORT` and (if needed) `HIVE_SESS_VAR_LIST` as follows.

```
export HS2_HOST_PORT=<HS2_HOST>:<HS2_PORT>
```

Also set `HIVE_SESS_VAR_LIST` to any required Hive session variables. For example:

- To connect to HiveServer2 with Kerberos authentication:

```
export HIVE_SESS_VAR_LIST="principal=<The server principal of HiveServer2>"
```

The Hive principal is specified by the configuration property `hive.server2.authentication.kerberos.principal` in `hive-site.xml`.
 - To connect to HiveServer2 running in HTTP mode:

```
export HIVE_SESS_VAR_LIST="transportMode=http;httpPath=<The HTTP endpoint>"
```
 - To connect to HiveServer2 when SSL is enabled:

```
export HIVE_SESS_VAR_LIST="ssl=true;sslTrustStore=<Trust Store path>;trustStorePassword=<Trust Store password>"
```
4. If you want to configure OSH to use Oracle Wallet, set environment the variables `TNS_ADMIN` and `WALLET_LOCATION` to the directory path where TNS and Oracle Wallet are stored.

```
export WALLET_LOCATION="<wallet_location>"
export TNS_ADMIN="<tns_admin_location>"
```

 **Note:**

Later, when you start an OSH session, if the `TNS_ADMIN` and `WALLET_LOCATION` are different on the Hadoop cluster then set `hadooptnsadmin` and `hadoopwalletlocation`:

```
set hadooptnsadmin <tns_admin_hadoop_node>
set hadoopwalletlocation <wallet_location_hadoop_node>
```

5. Start an OSH session and check for resources. This check indicates whether or not the configuration is correct.

```
$ ohsh
ohsh> show resources
```

You should see the resources `hadoop0`, `hive0`, and `bash0`.

8.1.2 Completing the OSH Configuration on the Hadoop Cluster

On the Hadoop cluster management server, Oracle Big Data SQL installs OSH from an RPM to `/opt/oracle/ohsh`. Perform the steps below as `root` in order to complete the Hadoop side of the configuration.

 **Note:**

The installation instructions at `$ORACLE_HOME/ohsh/doc/OHSB_HADOOP_NODE_README.txt` on the Hadoop cluster management server contain extra information that is not applicable to the installation of Copy to Hadoop and OHSB in the context of the Oracle Big Data SQL installation. Use the instructions below instead.

1. If the Hadoop system is not an Oracle Big Data Appliance, then download the SQL*Plus Instant Client and Oracle Instant Client for Linux from the [Oracle Technology Network](#). Select the client versions that matches the version of the Oracle Database.

For example, you can find the client downloads for Oracle Database 12.2.0.1.0 at this address: <http://www.oracle.com/technetwork/topics/linuxx86-64soft-092277.html>

Extract both packages into the same directory.

Add the directory to `PATH` and `LD_LIBRARY_PATH`.

2. Create a local directory for TNS and from `TNS_ADMIN` on the Oracle Database host, copy the `tnsnames.ora` and `sqlnet.ora` file into this directory.

If you create an Oracle Wallet for the Oracle Database host, then also copy over the Oracle Wallet file from the database server. Also check that `sqlnet.ora` points to the correct Oracle Wallet location.

```
WALLET_LOCATION=
  (SOURCE=(METHOD=FILE)(METHOD_DATA=
    (DIRECTORY=<your wallet directory>)))
sqlnet.wallet_override=true
```

3. As root, edit `/opt/oracle/ohsh/bin/ohsh_config.sh`. Set the following environment variables.

```
export CP2HADOOP_HOME=/opt/oracle/orahivedp-<version>
export HS2_HOST_PORT=<HiveServer2 hostname>:<Port number>
```

Also set `HIVE_SESS_VAR_LIST` to any required Hive session variables. For example:

- To connect to HiveServer2 with Kerberos authentication:


```
export HIVE_SESS_VAR_LIST="principal=<The server principal of HiveServer2>"
```
- To connect to HiveServer2 running in HTTP mode:


```
export HIVE_SESS_VAR_LIST="transportMode=http;httpPath=<The HTTP endpoint>"
```
- To connect to HiveServer2 when SSL is enabled:


```
export HIVE_SESS_VAR_LIST="ssl=true;sslTrustStore=<Trust Store path>;trustStorePassword=<Trust Store password>"
```

4. If TNS and Oracle Wallet are enabled, also add the following variables to the file and export them:

```
export WALLET_LOCATION="<your wallet directory>"
export TNS_ADMIN="<your TNS directory>"
```

5. Add `/opt/oracle/ohsh` and `$HADOOP_HOME/bin` to the PATH variable

```
PATH=/opt/oracle/ohsh-<version>/bin:$PATH
PATH=$HADOOP_HOME/bin:$PATH
```

6. Start an OSH session and check for resources. This check indicates whether or not the configuration is correct.

```
$ ohsh
ohsh> show resources
```

You should see the resources `hadoop0`, `hive0`, and `bash0`.

8.1.3 Getting Started Using Copy to Hadoop and OSH

See `$ORACLE_HOME/ohsh/doc` and `$ORACLE_HOME/ohsh/examples` for examples that show how to run the OSH load methods.

See the Oracle Big Data SQL User's Guide more information on Copy to Hadoop and Oracle Shell for Hadoop Loaders.

A

Installation Quick Reference

This condensed description of the installation steps may be useful if you have installed the release before, or, feel that you do not need a detailed explanation of the steps. You can refer back to the full instructions in the main body of the guide for more information.

About Restarts

Hadoop Side:

- Cloudera Configuration Manager (or Ambari) may be restarted by the installation. This in itself does not interrupt any services.
- Hive, YARN, and any other services that have a dependency on Hive or YARN (such as Impala) are restarted.

The Hive libraries parameter is updated in order to include Oracle Big Data SQL JARs. On Cloudera installations, if the YARN Resource Manager is enabled, then it is restarted in order to set cgroup memory limit for Oracle Big Data SQL and the other Hadoop services. On Oracle Big Data Appliance, the YARN Resource Manager is always enabled and therefore always restarted.

Database Side:

- In some cases where Grid infrastructure is present, it must be restarted. If the system uses Grid then you should have the Grid user credentials on hand in case a restart is required
- The Linux users grid and oracle (or other database owner) must both be in the same group (usually oinstall).

Table A-1 Checking Prerequisites and Patches

Step	For More Information
1. Check for the prerequisites on both the Hadoop or Oracle Database systems. (The installation will alert you if any are missing.)	<ul style="list-style-type: none">• Prerequisites for Installation on the Hadoop Cluster• Prerequisites for Installation on Oracle Database Nodes
2. Ensure that the latest Database Proactive Bundle Patch is installed on database as well as any one-off patches required for this product. NOTE: The Database Proactive Bundle Patch is different than the Patch Set Update. Refer to the Master Compatibility Matrix for complete details on required patches.	Oracle Big Data SQL Master Compatibility Matrix (Doc ID 2119369.1 in My Oracle Support)
3. On the database server, check that opatch is included in PATH. Then run opatch to see what patches are installed on the database: \$./opatch lspatches	
4. Copy patch 26170659 to the database server if not already installed.	

Table A-1 (Cont.) Checking Prerequisites and Patches

Step	For More Information
5. Make sure database and listener are down.	Managing Oracle Software and Applying Patches (<i>Database 2 Day + Real Application Clusters Guide</i>)
6. Unzip the patch archive on the database server. Navigate to the patch and apply it: <code>\$./opatch apply</code> If the database server has both a database and Grid home, then you must apply the one-off patch to BOTH the grid home and database home.	
7. Navigate to <code>\$ORACLE_HOME/Opatch</code> and run the <code>datapatch</code> command: <code>\$./datapatch</code>	

Table A-2 Installing the Hadoop Side of Oracle Big Data SQL

Step	For More Information
8. Download the Oracle Big Data SQL installation bundle from OTN or eDelivery (Oracle Software Delivery Cloud).	<ul style="list-style-type: none"> • OTN • Oracle Software Delivery Cloud
9. Copy this file on to master node of the Hadoop cluster (where Ambari or CM is running). Unzip it and navigate to extracted directory.	
10. Create a configuration file:	
<pre>{ "cluster": { "name": <Your cluster name> } }</pre>	
Database Authentication is enabled by default. You can create a configuration file where Database Authentication is disabled:	
<pre>{ "cluster": { "name": <Your cluster name> }, "security": { "database_auth_enabled": "false" } }</pre>	<ul style="list-style-type: none"> • Jaguar Configuration Parameter and Command Reference • Steps for Installing on the Hadoop Cluster
Other configuration options are available.	

Table A-2 (Cont.) Installing the Hadoop Side of Oracle Big Data SQL

Step	For More Information
<p>11. As root, run: <code>./jaguar install <config file></code></p> <p>If you want to generate a key for Database Authentication, then include the <code>--requestdb</code> parameter and list the databases that should accept the key:</p> <pre>./jaguar --requestdb orcl,testdb,proddb install <config file></pre>	
<p>12. Copy the database-side installation bundle generated in <code><Big Data SQL Install directory>/BDSjaguar/dbundles</code> to the database server. If you used <code>--requestdb</code> to generate a request key, find the <code>.reqkey</code> file under <code>/BDSjaguar/dbkeys</code> and copy it to the database server as well.</p>	

Table A-3 Installing the Database Side of Oracle Big Data SQL

Step	For More Information
<p>13. Log on as the database owner, locate the database-side installation bundle, and unzip it.</p>	Steps for Installing on Oracle Database Nodes
<p>14. Execute the <code>.run</code> file extracted from the ZIP file.</p>	
<p>15. Navigate to <code>\$ORACLE_HOME/BDSJaguar-3.2.0</code> and find the installation directory created by the <code>.run</code> file. The directory name will indicate the cluster, node, and domain where the installation bundle originated. Cd into the directory and run the installer:</p> <pre>\$ cd cdh510-6-node1.my.domain.com \$./bds-database-install.sh</pre>	Command Line Parameter Reference for bds-database-install.sh
<p>You may need to include some parameters with this command.</p>	

Perform the steps in the table below if `database_auth_enabled` and/or `impersonation_enabled` was set to "true" in the configuration.

Table A-4 Extra Steps to Complete the Configuration for Database Authentication and/or Multi-User Authorization

Step	For More Information
<p>16. If Database Authentication and/or Multi-User Authorization are enabled, the database-side installation generates a ZIP file. Find this file in the installation directory and copy it back to the path indicated below on the Hadoop cluster management server:</p>	<p>Steps for Installing on Oracle Database Nodes</p>
<pre>\$ ls \$ORACLE_HOME/BDSJaguar-3.2.0/cdh510-6- node1.my.domain.com/*.zip mycluster1-18- mycluster1node03.mydomain.com- myoradb1.mydomain.com.zip \$ scp mycluster1-18- mycluster1node03.mydomain.com- myoradb1.mydomain.com.zip root@node1:/opt/ oracle/DM/databases/conf</pre>	
<p>17. Log on to the Hadoop cluster management server as root and run the Database Acknowledge operation:</p>	
<pre># cd <Big Data SQL Install Directory>/ BDSJaguar # ./jaguar databaseack <config file></pre>	

Table A-5 Creating Users and Setting Up Mappings for Multi-User Authorization

Action	For More Information
<p>18. To grant a user access:</p> <ul style="list-style-type: none"> • Grant the BDSQL_USER role. • Grant read privileges on the BigDataSQL configuration directory object. 	
<pre>SQL> grant BDSQL_USER to user1; SQL> grant read on directory ORACLE_BIGDATA_CONFIG to user1;</pre>	

Table A-5 (Cont.) Creating Users and Setting Up Mappings for Multi-User Authorization

Action	For More Information
<p>19. To grant users Hadoop access through the Multi-User Authorization feature, use the procedures of the DBMS_BDSQL PL/SQL Package:</p> <ul style="list-style-type: none"> • Grant the BDSQL_ADMIN role to designated administrators in order to allow them to invoke these procedures. • Use the ADD_USER_MAP procedure to define the rules for a user. This creates a row in the BDSQL_USER_MAP table. • Use REMOVE_USER_MAP to remove a row from the BDSQL_USER_MAP table. 	DBMS_BDSQL PL/SQL Package

See Also

These workflow diagrams provide a visual overview of the installation process as a whole and a close-up view of the steps needed to complete the setup of Database Authentication:

- [Installation Workflow](#)
- [Generating and Installing the GUID-Key Pair for Database Authentication](#)

The next two sections of this guide provide useful code examples:

- [bds-config.json Configuration Example](#)
- [Oracle Big Data SQL Installation Examples](#)

B

bds-config.json Configuration Example

The following is an example of a fully-populated bds-config.json file, which includes all available configuration parameters.

```
{
  "cluster": {
    "name": "cluster"
  },
  "api" : {
    "port" : 5050,
    "restart_stale" : "false"
  },
  "repositories" : {
    "url" : [
      "http://repo1.domain.com/loc/hadoop",
      "http://repo2.domain.com/loc",
      "http://alternate.domain.com/backup/repo"
    ],
    "dir" : [ "/root/hadooprepo" ]
  },
  "network" : {
    "use_infiniband" : "false",
    "selection_subnet" : "5.32.128.10/21",
    "http_proxy" : "http://my.proxy.server.com:80",
    "https_proxy" : "http://mysecure.proxy.server.com:80"
  },
  "security" : {
    "database_auth_enabled" : "true",
    "impersonation_enabled" : "true",
    "impersonation_usehosts" : "true",
    "impersonation_blacklist" : "hdfs,hive"
  },
  "kerberos" : {
    "principal" : "oracle/mycluster@MY.DOMAIN.COM",
    "keytab" : "/home/oracle/security/oracle.keytab",
    "hdfs-principal" : "hdfs/mycluster@MY.DOMAIN.COM",
    "hdfs-keytab" : "/home/hdfs/security/hdfs.keytab"
  },
  "memory" : {
    "min_hard_limit" : 16384,
    "max_percentage" : 30
  }
}
```

C

Oracle Big Data SQL Installation Examples

The following are samples of the console output for the Oracle Big Data SQL installation.

- Example 1 shows the output of an installation on Oracle Big Data Appliance.
- Example 2 shows the parallel installation on Hortonworks HDP.

The database side of the installation is the same for Exadata Database Machine and all commodity Oracle Database platforms.

The database side of these scripts shows the installation on a system running Oracle Database 12.1. The output is the same for an installation on an Oracle Database 12.2 system, except for minor path changes.

Example 1: Installation on Oracle Big Data Appliance

This example shows the Hadoop-side installation where the Hadoop system is Oracle Big Data Appliance. These are the factors to note in this scenario:

! Important:

Notice that we are using `scl` to ensure that the correct version of Python is invoked.

```
[root@clust0601bda14vm03 BDSJaguar]# scl enable python27 "./jaguar
install"
BigDataSQL:
=====
BigDataSQL: Oracle Big Data SQL 3.2.0.
BigDataSQL: Starting execution on: Mon Nov 13 13:47:13 2017.
BigDataSQL:           in directory: /root/BDSJaguar.
BigDataSQL:           and hostname: clust0601bda14vm03.vm.oracle.com.
BigDataSQL: Verifying configuration parameters...
BigDataSQL: Starting server...
BigDataSQL: API protocol: http.
BigDataSQL: API port: 7180.
BigDataSQL: Waiting for management server to respond...
BigDataSQL: WARNING: If there are services in the stale state, they
will be restarted !!!

Cloudera Manager user: admin
Password:

BigDataSQL: Authentication succeeded.
BigDataSQL: API version: 17.
BigDataSQL: Cluster type: BDA.
BigDataSQL: Cluster name: clust0601bda14vm.
```

```

BigDataSQL: Cluster version: CDH5.
BigDataSQL: Service: BIGDATASQL.
BigDataSQL: Verifying service configuration..
BigDataSQL: Connection type: Ethernet.
BigDataSQL: Initialization completed..
BigDataSQL: Copying CSD to Cloudera Manager..
BigDataSQL: Restarting management server...
BigDataSQL: Waiting for management server to respond...
BigDataSQL: Restarting Cloudera Management service..
BigDataSQL: Starting deployment server..
BigDataSQL: Deployment server listening
BigDataSQL: Analyzing cluster ....
BigDataSQL: Checking for Big Data SQL requirements..
BigDataSQL: Preparing installation..
BigDataSQL: Installing Big Data SQL agents..
BigDataSQL: No http_proxy variable defined.
BigDataSQL: No https_proxy variable defined.
BigDataSQL: Installing Big Data SQL servers..
BigDataSQL: Starting bigdatasql service..
BigDataSQL: Restarting stale services..
BigDataSQL: Restarting Cloudera Management service..
BigDataSQL: Big Data SQL task completed..
BigDataSQL: Creating database side bundle..
BigDataSQL: Creating bundle file.
BigDataSQL: Bundle file created: /root/BDSJaguar/db-bundles/bds-3.2.0-
db-clust0601bda14vm-171113.1347.zip.
BigDataSQL: WARNING: Database authentication is enabled, you will need
a reqkey file for set it up..
BigDataSQL:         Ignore this warning if you already have one or if
your database already has its key,
BigDataSQL:         otherwise, use './jaguar --requestdb <db name>
databasereq' to create one.
BigDataSQL: INSTALL workflow completed.

```

Requesting a .reqkey file for database "mashadb." Database Authentication is enabled but the key was not requested in the install operation.

```

[root@clust0601bda14vm03 BDSJaguar]# scl enable python27 "./jaguar --
requestdb mashadb databasereq"
BigDataSQL:
=====
BigDataSQL: Oracle Big Data SQL 3.2.0.
BigDataSQL: Starting execution on: Mon Nov 13 14:07:21 2017.
BigDataSQL:         in directory: /root/BDSJaguar.
BigDataSQL:         and hostname: clust0601bda14vm03.vm.oracle.com.
BigDataSQL: Verifying configuration parameters...
BigDataSQL: Starting server...
BigDataSQL: API protocol: http.
BigDataSQL: API port: 7180.
BigDataSQL: Waiting for management server to respond...
BigDataSQL: WARNING: If there are services in the stale state, they
will be restarted !!!

Cloudera Manager user: admin

```

Password:

```

BigDataSQL: Authentication succeeded.
BigDataSQL: API version: 17.
BigDataSQL: Cluster type: BDA.
BigDataSQL: Cluster name: clust0601bda14vm.
BigDataSQL: Cluster version: CDH5.
BigDataSQL: Service: BIGDATASQL.
BigDataSQL: Verifying service configuration..
BigDataSQL: Connection type: Ethernet.
BigDataSQL: Initialization completed..
BigDataSQL: Starting deployment server..
BigDataSQL: Deployment server listening
BigDataSQL: Requiring database keys only..
BigDataSQL: reqkey file created: /root/BDSJaguar/dbkeys/mashadb.reqkey.
BigDataSQL: DATABASEREQ workflow completed.

```

Copy the database-side installation bundle and the request key file to the database node.

```

[root@clust0601bda14vm03 BDSJaguar]# scp /root/BDSJaguar/dbkeys/
mashadb.reqkey masha@clust0601bda14vm04:/home/masha
masha@clust0601bda14vm04's password:
mashadb.reqkey
                                100% 252      0.3KB/s   00:00
[root@clust0601bda14vm03 BDSJaguar]# scp /root/BDSJaguar/db-bundles/
bds-3.2.0-db-clust0601bda14vm-171113.1347.zip masha@clust0601bda14vm04:/
home/masha
masha@clust0601bda14vm04's password:
bds-3.2.0-db-
clust0601bda14vm-171113.1347.zip
                                100% 1020MB  51.0MB/s   00:20

```

Unzip the bundle, execute the run file, and copy the request key file into the install directory created by the run file.

```

[masha@clust0601bda14vm04 ~]$ unzip bds-3.2.0-db-
clust0601bda14vm-171113.1347.zip
Archive:  bds-3.2.0-db-clust0601bda14vm-171113.1347.zip
  inflating: bds-3.2.0-db-clust0601bda14vm-171113.1347.run
[masha@clust0601bda14vm04 ~]$ ./bds-3.2.0-db-
clust0601bda14vm-171113.1347.run
bds-3.2.0-db-clust0601bda14vm-171113.1347.run: platform
is                : Linux
bds-3.2.0-db-clust0601bda14vm-171113.1347.run: script started
at                : Mon Nov 13 14:19:56 CST 2017
bds-3.2.0-db-clust0601bda14vm-171113.1347.run: running
location         : /home/masha
bds-3.2.0-db-clust0601bda14vm-171113.1347.run:
ORACLE_SID      : mashadb
bds-3.2.0-db-clust0601bda14vm-171113.1347.run:
ORACLE_HOME     : /u03/app/masha/12.1.0/dbhome_mashadb
bds-3.2.0-db-clust0601bda14vm-171113.1347.run:
install directory : /u03/app/masha/12.1.0/dbhome_mashadb/

```

```
BDSJaguar-3.2.0/clust0601bda14vm-3-clust0601bda14vm03.vm.oracle.com
bds-3.2.0-db-clust0601bda14vm-171113.1347.run: extracting...
bds-3.2.0-db-clust0601bda14vm-171113.1347.run: installer extracted
successfully
```

```
cd $ORACLE_HOME/BDSJaguar-3.2.0/clust0601bda14vm-3-
clust0601bda14vm03.vm.oracle.com
```

```
bds-3.2.0-db-clust0601bda14vm-171113.1347.run: for setup Big Data SQL
on this database instance
```

```
[masha@clust0601bda14vm04 ~]$ mv mashadb.reqkey $ORACLE_HOME/
BDSJaguar-3.2.0/clust0601bda14vm-3-clust0601bda14vm03.vm.oracle.com
```

Change to the install directory and run the install operation of bds-database-install.sh.

```
[masha@clust0601bda14vm04 ~]$ cd $ORACLE_HOME/BDSJaguar-3.2.0/
clust0601bda14vm-3-clust0601bda14vm03.vm.oracle.com
[masha@clust0601bda14vm04 clust0601bda14vm-3-
clust0601bda14vm03.vm.oracle.com]$ ./bds-database-install.sh --install
bds-database-install: platform is           : Linux
bds-database-install: script started at     : Mon Nov 13 14:20:59
CST 2017
bds-database-install:
running location           : /u03/app/masha/12.1.0/dbhome_mashadb/
BDSJaguar-3.2.0/clust0601bda14vm-3-clust0601bda14vm03.vm.oracle.com
bds-database-install: ORACLE_SID           : mashadb
bds-database-install: ORACLE_HOME         : /u03/app/masha/12.1.0/
dbhome_mashadb
bds-database-install: configuration spfile  : spfilemashadb.ora
bds-database-install: Grid home            : not detected
bds-database-install: db resource         : mashadb
bds-database-install: mta extproc sid     :
bds_mashadb_clust0601bda14vm
bds-database-install: use of crsctl       : false
bds-database-install: root
shell script /u03/app/masha/12.1.0/dbhome_mashadb/install/bds-database-
install-10657-root-script-clust0601bda14vm.sh
please run as root:
```

```
<enter> to continue checking..
```

```
q<enter> to quit
```

```
bds-database-install: root script seem to have succeeded, continuing...
bds-database-install: cluster type         : bda
bds-database-install: connection type     : Ethernet
bds-database-install: cluster name        : clust0601bda14vm
bds-database-install: dictionary cluster name : clust0601bda14vm
bds-database-install: Hadoop version      : hadoop-2.6.0-cdh5.12.1
bds-database-install: Hive version       : hive-1.1.0-cdh5.12.1
bds-database-install: HBase version      : hbase-1.2.0-cdh5.12.1
bds-database-install: BDSQL version      : 3.2.0
bds-database-install: BDSQL install date  : Mon Nov 13 13:59:09
```

```

2017
bds-database-install: bd_cell version          :
bd_cell-12.1.2.0.100_LINUX.X64_171109.1-1.x86_64
bds-database-install: cp2hadoop kit version    : orahivedp-3.2.0
bds-database-install: ohsh kit version         : ohsh-1.3.1
bds-database-install: cell config dir         : /etc/oracle/cell/
network-config
bds-database-install: hybrid cell model        : false
bds-database-install: cell network(s)         : 192.168.254.5/27
bds-database-install: allow multiple subnets  : _skgxp_ant_options=1
bds-database-install: use UDP protocol        :
_skgxp_dynamic_protocol=2
bds-database-install: cellaffinity.ora file   : missing
bds-database-install: listener end point      : EXTPROC1521
bds-database-install: installing JDK          : jdk1.8.0_151
bds-database-install: action                  : setup
bds-database-setup: validating 192.168.254.4/27 against network
interfaces
bds-database-setup: trying to ping cluster nodes from 10.0.2.15
bds-database-setup: response fail from 192.168.254.4
bds-database-setup: response fail from 192.168.254.2
bds-database-setup: response fail from 192.168.254.3
bds-database-setup: trying to ping cluster nodes from 192.168.254.5
bds-database-setup: response ok from 192.168.254.4
bds-database-setup: response ok from 192.168.254.2
bds-database-setup: response ok from 192.168.254.3
bds-database-setup: ping succeeded on all nodes, using 192.168.254.5/27
at eth1 instead
bds-database-setup: installing cluster configuration files
bds-database-setup: unpacking hadoop-2.6.0-cdh5.12.1.tar.gz
bds-database-setup: unpacking hadoop-2.6.0-nativelib-ol6.tar.gz
bds-database-setup: unpacking hive-1.1.0-cdh5.12.1.tar.gz
bds-database-setup: unpacking hbase-1.2.0-cdh5.12.1.tar.gz
bds-database-setup: creating database objects
bds-database-setup: mta will be setup
bds-database-setup: creating /u03/app/masha/12.1.0/dbhome_mashadb/hs/
admin/initbds_mashadb_clust0601bda14vm.ora
bds-database-setup: mta shutdown bds_mashadb_clust0601bda14vm
bds-database-setup: removing existing mta parameter settings
bds-database-setup: setting performance parameters for mta
bds-database-setup: setting listener_address to EXTPROC1521
bds-database-setup: starting extproc bds_mashadb_clust0601bda14vm
bds-database-setup: waiting upto 100 seconds for extproc to register
with listener
.....ok
bds-database-setup: patching view LOADER_DIR_OBJS
bds-database-setup: creating mta dblinks
bds-database-setup: installing cp2hadoop
bds-database-setup: installing ohsh
bds-database-setup: reconfiguring database instance mashadb
bds-database-setup: adding parameter _bigdata_external_table=true
bds-database-setup: setup cell contiguration
bds-database-setup: cellinit.ora already has the configured ip address
bds-database-setup: cellinit.ora already enabled with UDP protocol
bds-database-setup: celliniteth.ora already has the configured ip

```

```

address
bds-database-setup: cellniteteth.ora already enabled with UDP protocol
bds-database-setup: cellaffinity.ora already disabled
bds-database-setup: no database restart needed

bds-database-setup: activation zip file created: clust0601bda14vm-3-
clust0601bda14vm03.vm.oracle.com-clust0601bda14vm04.vm.oracle.com.zip

# /u03/app/masha/12.1.0/dbhome_mashadb/BDSJaguar-3.2.0/
clust0601bda14vm-3-clust0601bda14vm03.vm.oracle.com/clust0601bda14vm-3-
clust0601bda14vm03.vm.oracle.com-clust0601bda14vm04.vm.oracle.com.zip

bds-database-setup: copy this file
to root@clust0601bda14vm03.vm.oracle.com:/opt/oracle/DM/databases/conf
bds-database-setup: and run:

./jaguar databaseack

bds-database-setup: configuration successful
for cluster clust0601bda14vm

bds-database-setup: log files in /u03/app/masha/
12.1.0/dbhome_mashadb/install
bds-database-setup: bds-database-install-10657-* files
for more detailed info
bds-database-setup: full log in bds-
database-install-10657-main-script.log

```

This is the output from the root script execution:

```

[root@clust0601bda14vm04 ~]# /u03/app/masha/12.1.0/dbhome_mashadb/
install/bds-database-install-10657-root-script-clust0601bda14vm.sh
bds-database-install-10657-root-script-clust0601bda14vm.sh:
clust0601bda14vm: removing existing entries for 192.168.254.4 from /etc/
hosts
bds-database-install-10657-root-script-clust0601bda14vm.sh:
clust0601bda14vm: removing existing entries for 192.168.254.2 from /etc/
hosts
bds-database-install-10657-root-script-clust0601bda14vm.sh:
clust0601bda14vm: removing existing entries for 192.168.254.3 from /etc/
hosts
bds-database-install-10657-root-script-clust0601bda14vm.sh:
clust0601bda14vm: removing existing entries for
clust0601bda14vm03.vm.oracle.com from /etc/hosts
bds-database-install-10657-root-script-clust0601bda14vm.sh:
clust0601bda14vm: removing existing entries for
clust0601bda14vm01.vm.oracle.com from /etc/hosts
bds-database-install-10657-root-script-clust0601bda14vm.sh:
clust0601bda14vm: removing existing entries for
clust0601bda14vm02.vm.oracle.com from /etc/hosts
bds-database-install-10657-root-script-clust0601bda14vm.sh:
clust0601bda14vm: removing existing entries for clust0601bda14vm03
from /etc/hosts
bds-database-install-10657-root-script-clust0601bda14vm.sh:

```

```

clust0601bda14vm: removing existing entries for clust0601bda14vm01
from /etc/hosts
bds-database-install-10657-root-script-clust0601bda14vm.sh:
clust0601bda14vm: removing existing entries for clust0601bda14vm02
from /etc/hosts
bds-database-install-10657-root-script-clust0601bda14vm.sh:
clust0601bda14vm: added entry 192.168.254.4
clust0601bda14vm03.vm.oracle.com clust0601bda14vm03 to /etc/hosts
bds-database-install-10657-root-script-clust0601bda14vm.sh:
clust0601bda14vm: added entry 192.168.254.2
clust0601bda14vm01.vm.oracle.com clust0601bda14vm01 to /etc/hosts
bds-database-install-10657-root-script-clust0601bda14vm.sh:
clust0601bda14vm: added entry 192.168.254.3
clust0601bda14vm02.vm.oracle.com clust0601bda14vm02 to /etc/hosts
bds-database-install-10657-root-script-clust0601bda14vm.sh:
clust0601bda14vm: root script finished successfully

```

Copy the zip file generated by bds-database-install.sh back to the Hadoop cluster management server. This zip file contains the .ackkey file needed to complete the Database Authentication setup.

```

[masha@clust0601bda14vm04 clust0601bda14vm-3-
clust0601bda14vm03.vm.oracle.com]$ scp clust0601bda14vm-3-
clust0601bda14vm03.vm.oracle.com-clust0601bda14vm04.vm.oracle.com.zip
root@clust0601bda14vm03.vm.oracle.com:/opt/oracle/DM/databases/conf
root@clust0601bda14vm03.vm.oracle.com's password:
clust0601bda14vm-3-clust0601bda14vm03.vm.oracle.com-
clust0601bda14vm04.vm.oracle.com.zip          100%
789      0.8KB/s   00:00

```

As root on the Hadoop cluster management server, run Jaguar's Database Acknowledge operation. This completes the installation.

```

BigDataSQL: Oracle Big Data SQL 3.2.0.
BigDataSQL: Starting execution on: Mon Nov 13 14:26:37 2017.
BigDataSQL:           in directory: /root/BDSJaguar.
BigDataSQL:           and hostname: clust0601bda14vm03.vm.oracle.com.
BigDataSQL: Verifying configuration parameters...
BigDataSQL: Starting server...
BigDataSQL: API protocol: http.
BigDataSQL: API port: 7180.
BigDataSQL: Waiting for management server to respond...
BigDataSQL: WARNING: If there are services in the stale state, they
will be restarted !!!

```

```

Cloudera Manager user: admin
Password:

```

```

BigDataSQL: Authentication succeeded.
BigDataSQL: API version: 17.
BigDataSQL: Cluster type: BDA.
BigDataSQL: Cluster name: clust0601bda14vm.
BigDataSQL: Cluster version: CDH5.
BigDataSQL: Service: BIGDATASQL.

```



```

BigDataSQL: Verifying service configuration..
BigDataSQL: Connection type: Ethernet.
BigDataSQL: Initialization completed..
BigDataSQL: Starting deployment server..
BigDataSQL: Database acknowledge..
BigDataSQL: Deployment server listening
BigDataSQL: Refreshing Big Data SQL database connections..
BigDataSQL: Starting bigdatasql service..
BigDataSQL: Setting database key for Big Data SQL..
BigDataSQL: Restarting stale services..
BigDataSQL: Restarting Cloudera Management service..

Big Data SQL setup finished !!!
run bdschecksw tool on database node for sanity verification

BigDataSQL: DATABASEACK workflow completed.

```

Example 2: Installation on an HDP Cluster

In this example, the Hadoop-side installation is on an HDP system. To show what happens when Database Authentication is enabled but no request key is generated, we deliberately leave the `--requestdb` parameter out of the install operation. This is corrected later in the example.

```

[root@clust42bda07-adm BDSJaguar]# ./jaguar install
BigDataSQL:
=====
BigDataSQL: Oracle Big Data SQL 3.2.0.
BigDataSQL: Starting execution on: Mon Nov 13 12:41:05 2017.
BigDataSQL:           in directory: /root/BDSJaguar.
BigDataSQL:           and hostname: clust42bda07-adm.us.oracle.com.
BigDataSQL: Verifying configuration parameters...
BigDataSQL: Starting server...
BigDataSQL: API protocol: http.
BigDataSQL: API port: 8080.
BigDataSQL: Waiting for management server to respond...
BigDataSQL: WARNING: If there are services in the stale state, they
will be restarted !!!

Ambari user: admin
Password:

BigDataSQL: Authentication succeeded.
BigDataSQL: API version: 1.
BigDataSQL: Cluster type: HDP.
BigDataSQL: Cluster name: afghhdp.
BigDataSQL: Cluster version: HDP-2.6.
BigDataSQL: Service: BIGDATASQL.
BigDataSQL: Verifying service configuration...
BigDataSQL: Connection type: Ethernet.
BigDataSQL: Initialization completed..
BigDataSQL: Copying Stack to Ambari..
BigDataSQL: Restarting management server...
BigDataSQL: Waiting for management server to respond...

```

```

BigDataSQL: Starting deployment server..
BigDataSQL: Deployment server listening
BigDataSQL: Analyzing cluster ....
BigDataSQL: Checking for Big Data SQL requirements..
BigDataSQL: Preparing installation..
BigDataSQL: Installing Big Data SQL agents..
BigDataSQL: Installing Big Data SQL servers..
BigDataSQL: Starting bigdatasql service.
BigDataSQL: Restarting stale services..
BigDataSQL: Starting hdfs service.
BigDataSQL: Starting yarn service.
BigDataSQL: Starting mapreduce2 service.
BigDataSQL: Starting hbase service.
BigDataSQL: Big Data SQL task completed..
BigDataSQL: Creating database side bundle..
BigDataSQL: Creating bundle file.
BigDataSQL: Bundle file created: /root/BDSJaguar/db-bundles/bds-3.2.0-
db-afghhdp-171113.1241.zip.
BigDataSQL: WARNING: Database authentication is enabled, you will need
a reqkey file for set it up..
BigDataSQL:           Ignore this warning if you already have one or if
your database already has its key,.
BigDataSQL:           otherwise, use './jaguar --requestdb <db name>
databasereq' for create one..
BigDataSQL: INSTALL workflow completed.

```

Database Authentication was enabled but a key was not requested in the install operation and has not been request previously for the "laicadb" database where we want to perform the installation. To correct this, we invoke the Jaguar databasereq operation to create a request key for "laicadb":

```

[root@clust42bda07-adm BDSJaguar]# ./jaguar --requestdb laicadb
databasereq
BigDataSQL:
=====
BigDataSQL: Oracle Big Data SQL 3.2.0.
BigDataSQL: Starting execution on: Mon Nov 13 12:56:39 2017.
BigDataSQL:           in directory: /root/BDSJaguar.
BigDataSQL:           and hostname: clust42bda07-adm.us.oracle.com.
BigDataSQL: Verifying configuration parameters...
BigDataSQL: Starting server...
BigDataSQL: API protocol: http.
BigDataSQL: API port: 8080.
BigDataSQL: Waiting for management server to respond...
BigDataSQL: WARNING: If there are services in the stale state, they
will be restarted !!!

Ambari user: admin
Password:

BigDataSQL: Authentication succeeded.
BigDataSQL: API version: 1.
BigDataSQL: Cluster type: HDP.
BigDataSQL: Cluster name: afghhdp.

```

```

BigDataSQL: Cluster version: HDP-2.6.
BigDataSQL: Service: BIGDATASQL.
BigDataSQL: Verifying service configuration...
BigDataSQL: Connection type: Ethernet.
BigDataSQL: Initialization completed..
BigDataSQL: Starting deployment server..
BigDataSQL: Deployment server listening
BigDataSQL: Requiring database keys only..
BigDataSQL: reqkey file created: /root/BDSJaguar/dbkeys/laicadb.reqkey.
BigDataSQL: DATASERREQ workflow completed.

```

Now that we have a database-side installation bundle and a .reqkey file, we copy both to the database node.

```

[root@clust42bda07-adm BDSJaguar]# scp /root/BDSJaguar/db-bundles/
bds-3.2.0-db-afghhdp-171113.1241.zip laica@clust42bda05-adm:/home/laica
laica@clust42bda05-adm's password:
bds-3.2.0-db-
afghhdp-171113.1241.zip
          100% 604MB 112.0MB/s  00:05
[root@clust42bda07-adm BDSJaguar]# scp /root/BDSJaguar/dbkeys/
laicadb.reqkey laica@clust42bda05-adm:/home/laica
laica@clust42bda05-adm's password:
laicadb.reqkey
                          100% 253  453.7KB/s  00:00

```

Next, unzip the bundle, execute run file, and copy the request key to the install directory created by the run file.

```

[laica@clust42bda05-adm ~]$ unzip bds-3.2.0-db-afghhdp-171113.1241.zip
Archive:  bds-3.2.0-db-afghhdp-171113.1241.zip
  inflating: bds-3.2.0-db-afghhdp-171113.1241.run
[laica@clust42bda05-adm ~]$ ./bds-3.2.0-db-afghhdp-171113.1241.run
bds-3.2.0-db-afghhdp-171113.1241.run: platform is           : Linux
bds-3.2.0-db-afghhdp-171113.1241.run: script started at   : Mon
Nov 13 15:03:19 CST 2017
bds-3.2.0-db-afghhdp-171113.1241.run: running location    : /home/
laica
bds-3.2.0-db-afghhdp-171113.1241.run: ORACLE_SID         :
laicadb
bds-3.2.0-db-afghhdp-171113.1241.run:
ORACLE_HOME           : /u03/app/laica/12.1.0/dbhome_laicadb
bds-3.2.0-db-afghhdp-171113.1241.run:
install directory     : /u03/app/laica/12.1.0/dbhome_laicadb/
BDSJaguar-3.2.0/afghhdp-2-clust42bda07-adm.us.oracle.com
bds-3.2.0-db-afghhdp-171113.1241.run: extracting...
bds-3.2.0-db-afghhdp-171113.1241.run: installer extracted successfully

cd $ORACLE_HOME/BDSJaguar-3.2.0/afghhdp-2-clust42bda07-adm.us.oracle.com

bds-3.2.0-db-afghhdp-171113.1241.run: for setup Big Data SQL on this
database instance

```

```
[laica@clust42bda05-adm ~]$ mv laicadb.reqkey $ORACLE_HOME/
BDSJaguar-3.2.0/afghhdp-2-clust42bda07-adm.us.oracle.com
```

Change to the install directory and run the database-side install utility (bds-database-install.sh).

```
[laica@clust42bda05-adm ~]$ cd $ORACLE_HOME/BDSJaguar-3.2.0/afghhdp-2-
clust42bda07-adm.us.oracle.com
[laica@clust42bda05-adm afghhdp-2-clust42bda07-adm.us.oracle.com]$ ./
bds-database-install.sh --install
bds-database-install: platform is           : Linux
bds-database-install: script started at    : Mon Nov 13 15:04:54
CST 2017
bds-database-install: running location     : /u03/app/laica/12.1.0/
dbhome_laicadb/BDSJaguar-3.2.0/afghhdp-2-clust42bda07-adm.us.oracle.com
bds-database-install: ORACLE_SID          : laicadb
bds-database-install: ORACLE_HOME        : /u03/app/laica/12.1.0/
dbhome_laicadb
bds-database-install: configuration spfile : spfilelaicadb.ora
bds-database-install: Grid home           : /u03/app/12.1.0/grid
bds-database-install: db resource         : laicadb
bds-database-install: crs resource        : ora.laicadb.db
bds-database-install: database type       : SINGLE
bds-database-install: cardinality         : 1
bds-database-install: mta extproc sid     : bds_laicadb_afghhdp
bds-database-install: use of crsctl       : true
bds-database-install: cluster type        : hdp
bds-database-install: connection type    : Ethernet
bds-database-install: cluster name        : afghhdp
bds-database-install: dictionary cluster name : afghhdp
bds-database-install: Hadoop version      :
hadoop-2.7.3.2.6.3.0-235
bds-database-install: Hive version        : apache-
hive-1.2.1000.2.6.3.0-235-bin
bds-database-install: HBase version       :
hbase-1.1.2.2.6.3.0-235
bds-database-install: BDSQL version       : 3.2.0
bds-database-install: BDSQL install date  : Mon Nov 13 12:53:29
2017
bds-database-install: bd_cell version     :
bd_cell-12.1.2.0.100_LINUX.X64_171109.1-1.x86_64
bds-database-install: cp2hadoop kit version : orahivedp-3.2.0
bds-database-install: ohsh kit version    : ohsh-1.3.1
bds-database-install: cell config dir     : /etc/oracle/cell/
network-config
bds-database-install: hybrid cell model   : false
bds-database-install: cell network(s)     : 10.128.18.35/21
bds-database-install: allow multiple subnets : _skgxp_ant_options=1
bds-database-install: use UDP protocol    :
_skgxp_dynamic_protocol=2
bds-database-install: cellaffinity.ora file : missing
bds-database-install: listener end point  : EXTPROC1521
bds-database-install: installing JDK      : jdk1.8.0_151
bds-database-install: action              : setup
```

```
bds-database-setup: only one network interface detected,
10.128.18.35/21 at eth0
bds-database-setup: ping to cluster nodes
bds-database-setup: response ok from 10.128.18.36
bds-database-setup: response ok from 10.128.20.189
bds-database-setup: installing cluster configuration files
bds-database-setup: unpacking hadoop-2.7.3.2.6.3.0-235.tar.gz
bds-database-setup: unpacking hadoop-2.7.0-nativelib-ol6.tar.gz
bds-database-setup: unpacking apache-hive-1.2.1000.2.6.3.0-235-
bin.tar.gz
bds-database-setup: unpacking hbase-1.1.2.2.6.3.0-235.tar.gz
bds-database-setup: creating database objects
bds-database-setup: mta will be setup
bds-database-setup: creating /u03/app/laica/12.1.0/dbhome_laicadb/hs/
admin/initbds_laicadb_afghhdp.ora
bds-database-setup: mta shutdown bds_laicadb_afghhdp
bds-database-setup: removing existing mta parameter settings
bds-database-setup: stopping crs resource bds_laicadb_afghhdp
bds-database-setup: deleting crs resource bds_laicadb_afghhdp
bds-database-setup: registering crs resource bds_laicadb_afghhdp
bds-database-setup: using dependency db resource of laicadb
bds-database-setup: starting crs resource bds_laicadb_afghhdp
bds-database-setup: patching view LOADER_DIR_OBJ$
bds-database-setup: creating mta dblink$
bds-database-setup: installing cp2hadoop
bds-database-setup: installing ohsh
bds-database-setup: setup cell contiguration
bds-database-setup: cellinit.ora already has the configured ip address
bds-database-setup: cellinit.ora already enabled with UDP protocol
bds-database-setup: celliniteth.ora already has the configured ip
address
bds-database-setup: celliniteth.ora already enabled with UDP protocol
bds-database-setup: cellaffinity.ora already disabled
bds-database-setup: no database restart needed

bds-database-setup: activation zip file created: afghhdp-2-clust42bda07-
adm.us.oracle.com-clust42bda05-adm.us.oracle.com.zip

# /u03/app/laica/12.1.0/dbhome_laicadb/BDSJaguar-3.2.0/afghhdp-2-
clust42bda07-adm.us.oracle.com/afghhdp-2-clust42bda07-adm.us.oracle.com-
clust42bda05-adm.us.oracle.com.zip

bds-database-setup: copy this file to root@clust42bda07-
adm.us.oracle.com:/opt/oracle/DM/databases/conf
bds-database-setup: and run:

./jaguar databaseack

bds-database-setup: configuration successful for cluster afghhdp

bds-database-setup: log files in /u03/app/laica/12.1.0/dbhome_laicadb/
install
bds-database-setup: bds-database-install-29088-* files for
more detailed info
```

```
bds-database-setup: full log in bds-database-install-29088-
main-script.log
```

The zip file created by the install operation contains the .ackkey file needed to complete Database Authentication setup. Copy this zip file back to the Hadoop cluster management server.

```
[laica@clust42bda05-adm afghhdp-2-clust42bda07-adm.us.oracle.com]$
scp afghhdp-2-clust42bda07-adm.us.oracle.com-clust42bda05-
adm.us.oracle.com.zip root@clust42bda07-adm.us.oracle.com:/opt/
oracle/DM/databases/conf
root@clust42bda07-adm.us.oracle.com's password:
afghhdp-2-clust42bda07-adm.us.oracle.com-clust42bda05-
adm.us.oracle.com.zip                100% 739
0.7KB/s  00:00
```

As root on the Hadoop side run Jaguar's Database Acknowledge operation.

```
[root@clust42bda07-adm BDSJaguar]# ./jaguar databaseack
BigDataSQL:
=====
BigDataSQL: Oracle Big Data SQL 3.2.0.
BigDataSQL: Starting execution on: Mon Nov 13 13:08:43 2017.
BigDataSQL:           in directory: /root/BDSJaguar.
BigDataSQL:           and hostname: clust42bda07-adm.us.oracle.com.
BigDataSQL: Verifying configuration parameters...
BigDataSQL: Starting server...
BigDataSQL: API protocol: http.
BigDataSQL: API port: 8080.
BigDataSQL: Waiting for management server to respond...
BigDataSQL: WARNING: If there are services in the stale state, they
will be restarted !!!

Ambari user: admin
Password:

BigDataSQL: Authentication succeeded.
BigDataSQL: API version: 1.
BigDataSQL: Cluster type: HDP.
BigDataSQL: Cluster name: afghhdp.
BigDataSQL: Cluster version: HDP-2.6.
BigDataSQL: Service: BIGDATASQL.
BigDataSQL: Verifying service configuration...
BigDataSQL: Connection type: Ethernet.
BigDataSQL: Initialization completed..
BigDataSQL: Starting deployment server..
BigDataSQL: Database acknowledge..
BigDataSQL: Deployment server listening
BigDataSQL: Refreshing Big Data SQL database connections..
BigDataSQL: Starting bigdatasql service.
BigDataSQL: Setting database key for Big Data SQL..
BigDataSQL: Restarting stale services..
BigDataSQL: Starting hdfs service.
BigDataSQL: Starting yarn service.
```

```
BigDataSQL: Starting mapreduce2 service.  
BigDataSQL: Starting hbase service.
```

```
Big Data SQL setup finished !!!  
run bdschecksw tool on database node for sanity verification
```

```
BigDataSQL: DATABASEACK workflow completed.
```

D

Determining the Correct Software Version and Composing the Download Paths for Hadoop Clients

To configure `bds-database-create-bundle.sh` to download the Hadoop, Hive, and HBase tarballs, you must supply an URL to each these parameters:

```
--hive-client-ws  
--hadoop-client-ws  
--hbase-client-ws
```

To get the information needed to provide the correct URL, first check the content management service (CM or Ambari) and find the version of the Hadoop, Hive, and HBase services running on the Hadoop cluster. The compatible clients are of the same versions. In each case, the client tarball filename includes a version string segment that matches the version of the service installed on the cluster. In the case of CDH, you can then browse the public repository and find the URL to the client that matches the service version. For the HDP repository this would require a tool that can browse Amazon S3 storage. However you can also compose the correct URL using the known URL pattern along with information that you can acquire from Ambari, as described in this section.

For CDH (Both Oracle Big Data Appliance and Commodity CDH Systems):

1. Log on to Cloudera Manager and go to the **Hosts** menu. Select **All Hosts** , then **Inspect All Hosts**.
2. When the inspection is finished, select either **Show Inspector Results** (on the screen) or **Download Result Data** (to a JSON file).
3. In either case, scan the result set and find the service versions.

In JSON version of the inspector results, there is a `componentInfo` section for each cluster that shows the versions of software installed on that cluster. For example:

```
"componentInfo": [  
  ...  
  {  
    "cdhVersion": "CDH5",  
    "componentRelease": "1.cdh5.11.1.p0.6",  
    "componentVersion": "2.6.0+cdh5.11.1+2400",  
    "name": "hadoop"  
  },  
  ...  
]
```

4. Go to <http://archive.cloudera.com/cdh5/cdh/5>.

Look in the "hadoop," "hive," and "hbase" subdirectories of the CDH5 section of the archive. In the listings, you should find the client packages for the versions of the services installed on the cluster.

5. Copy the URLs and use them as the parameter values supplied to `bds-database-create-bundle.sh`. For example:

```
http://archive.cloudera.com/cdh5/cdh/5/hadoop-2.6.0-cdh5.11.1.tar.gz
http://archive.cloudera.com/cdh5/cdh/5/hive-1.1.0-cdh5.11.1.tar.gz
http://archive.cloudera.com/cdh5/cdh/5/hbase-1.2.0-cdh5.11.1.tar.gz
```

See Also:

Search for "Host Inspector" on [Cloudera](#) website if you need more help using this tool to determine installed software versions.

For HDP:

1. Log on to Ambari. Go to **Admin**, then **Stack and Versions**. On the **Stack** tab, locate the entries for the HDFS, Hive, and HBase services and note down the version number of each as the "service version."
2. Click the **Versions** tab. Note down the version of HDP that is running on the cluster as the "HDP version base."
3. Click **Show Details** to display a pop-up window that shows the full version string for the installed HDP release. Note this down as the "HDP full version"
4. The last piece of information needed is the Linux version ("centos5," "centos6," or "centos7"). Note this down as "OS version."

To search through the HDP repository in Amazon S3 storage to find the correct client URLs using this information acquired in this steps, you would need an S3 browser, browser extension, or command line tool. As alternative, you can piece together the correct URLs, using these strings.

For HDP 2.5 and earlier, the URLs pattern is as follows.

```
http://public-repo-1.hortonworks.com/HDP/<OS version>/2.x/updates/<HDP
version base>/tars/{hadoop|apache-hive|hbase}-<service version>.<HDP
full version>.tar.gz
```

Here are some examples. Note that the pattern of the gzip filename is slightly different for Hive. There is an extra "-bin" segment in the name.

```
http://public-repo-1.hortonworks.com/HDP/centos6/2.x/updates/2.3.2.0/
tars/hadoop-2.7.1.2.3.2.0-2950.tar.gz
http://public-repo-1.hortonworks.com/HDP/centos6/2.x/updates/2.3.2.0/
tars/apache-hive-1.2.1.2.3.2.0-2950-bin.tar.gz
http://public-repo-1.hortonworks.com/HDP/centos6/2.x/updates/2.3.2.0/
tars/hbase-1.1.2.2.3.2.0-2950.tar.gz
```

For HDP 2.5 and later releases, the pattern is almost the same except that there is an additional `hadoop`, `hive`, or `hbase` directory under the `tar` directory:

```
http://public-repo-1.hortonworks.com/HDP/centos6/2.x/updates/2.5.6.0/
tars/hadoop/hadoop-2.7.3.2.5.6.0-40.tar.gz
http://public-repo-1.hortonworks.com/HDP/centos6/2.x/updates/2.5.6.0/
tars/hive/apache-hive-1.2.1000.2.5.6.0-40-bin.tar.gz
http://public-repo-1.hortonworks.com/HDP/centos6/2.x/updates/2.5.6.0/
tars/hbase/hbase-1.1.2.2.5.6.0-40.tar.gz
```

Alternative Method for HDP:

You can get the required software versions from the command line instead of using Ambari.

- `# hdp-select versions`

Copy and save the numbers to the left of the dash as the “HDP version base”.

- `# hadoop version`
`# beeline --version`
`# hbase version`

Use the output from these commands to formulate the `<service version>.<HDP full version>` segment for each URL.

E

Oracle Big Data SQL Licensing

The licensing for Oracle Big Data SQL is separate from the licensing for other Oracle products.

When you purchase licensing for Oracle Big Data SQL, note the following:

- A separate license must be procured for each Hadoop cluster.
- All data nodes within the Hadoop cluster must be licensed. Partial licensing within a node is not available.
- All disks accessed by Oracle Big Data SQL (either within or outside of the Hadoop cluster) must be licensed. This includes disks used by Kafka or NoSQL stores when Oracle Big Data SQL accesses those sources.
- There is no additional license required for the Oracle Database server side.
- Oracle Copy to Hadoop licensing is included.

Third Party Licensing for Oracle Shell for Hadoop Loaders

Oracle Shell for Hadoop Loaders is included in the Oracle Big Data SQL installation bundle. The following are third-party projects contained in Oracle Shell for Hadoop Loaders.

- ANTLR 4.7
- Apache Commons Exec 1.3

Unless otherwise specifically noted, or as required under the terms of the third party license (e.g., LGPL), the licenses and statements herein, including all statements regarding Apache-licensed code, are intended as notices only.

E.1 ANTLR 4.7

Copyright (c) 2015 Terence Parr, Sam Harwell

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. The name of the author may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,

THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

E.2 Apache Commons Exec 1.3

Include the following License ONLY ONCE in the documentation even if there are multiple products licensed under the license.

The following applies to all products licensed under the Apache 2.0 License:

You may not use the identified files except in compliance with the Apache License, Version 2.0 (the "License").

You may obtain a copy of the License at <http://www.apache.org/licenses/LICENSE-2.0>. A copy of the license is also reproduced below.

E.3 Apache Licensed Code

The following is included as a notice in compliance with the terms of the Apache 2.0 License, and applies to all programs licensed under the Apache 2.0 license:

You may not use the identified files except in compliance with the Apache License, Version 2.0 (the "License.")

You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

A copy of the license is also reproduced below.

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.

E.4 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:
 - a. You must give any other recipients of the Work or Derivative Works a copy of this License; and
 - b. You must cause any modified files to carry prominent notices stating that You changed the files; and
 - c. 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
 - d. 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.

This product includes software developed by The Apache Software Foundation (<http://www.apache.org/>) (listed below):

F

Change History for Previous Releases

The following are changes in previous versions of the product.

F.1 Changes in Oracle Big Data SQL 3.1

The following are changes in Oracle Big Data SQL Release 3.1.

Expanded Deployments

Release 3.1 broadens support for Oracle Big Data SQL connectivity between Oracle Engineered Systems and commodity servers.

In earlier Oracle Big Data SQL releases, the following Oracle Database/Hadoop connections are possible:

- Oracle Exadata Database Machine to Oracle Big Data Appliance.
- Oracle Database on commodity servers to commodity Hadoop systems.

As of Release 3.1, Oracle Big Data SQL supports all of the following Oracle Database/Hadoop system connections:

- Oracle Database on commodity servers to Oracle Big Data Appliance.
- Oracle Database on commodity servers to commodity Hadoop systems.
- Oracle Exadata Database Machine to Oracle Big Data Appliance.
- Oracle Exadata Database Machine to commodity Hadoop systems.

The phrase “Oracle Database on commodity servers” refers to non-Exadata Linux systems that are officially-supported as Oracle Database platforms. “Commodity Hadoop systems” refers to Hortonworks HDP systems or to Cloudera CDH-based systems other than Oracle Big Data Appliance. In all cases, Oracle Database servers and Hadoop systems must meet the prerequisites identified in the [Oracle Big Data SQL Master Compatibility Matrix](#) (Doc ID 2119369.1 in [My Oracle Support](#)).

Oracle SPARC SuperCluster Support

Release 3.1 provides support for Oracle SPARC SuperCluster, with certain limitations:

- Ethernet connections between BDA (and commodity Hadoop systems) and SPARC SuperCluster are not supported.
- Oracle Database Tablepaces in HDFS (with Smart Scan technology) is not supported on this platform.

Ethernet Option for Connections to the Exadata Database Machine

The preferred method of connecting Oracle Big Data Appliance and Oracle Exadata Database for any purpose is through InfiniBand. Previous releases of Oracle Big Data SQL have required InfiniBand for these connections. In Release 3.1, Ethernet networking between the Exadata Database Machine and Oracle Big Data Appliance

is now supported. This enables you to use Oracle Big Data SQL with these two Engineered Systems in environments where InfiniBand is not feasible, such as when the two systems are geographically distant from each other.

Release 3.1 also enables Ethernet connections between commodity Hadoop systems and the Oracle Exadata Database Machine.

Oracle Big Data SQL connectivity between commodity Hadoop systems and commodity Oracle Database servers has been Ethernet-based throughout previous releases.

Note that Ethernet connections between Oracle Big Data Appliance (or commodity Hadoop systems) and Oracle SPARC SuperCluster are not supported at this time.

Simplified Deployment on Oracle Database Servers – Oracle Grid Infrastructure is now Optional

In previous releases of Oracle Big Data SQL, Oracle Grid Infrastructure is a prerequisite of the installation for all Oracle Database servers, including standalone servers that are not part of an Oracle RAC system. In Release 3.1, you have the option to install Oracle Big Data SQL on servers where Oracle Grid Infrastructure is not present. Note that in these cases, the installer makes some configuration file changes that require a restart of Oracle Database.

Unified Platform Support in the Oracle Big Data SQL Installer

Previous Oracle Big Data SQL releases included two separate installation procedures – one for Oracle Engineered Systems and another for commodity servers. In Release 3.1, you use the same installation process for both Oracle and non-Oracle platforms. This is also true for maintenance. For all of the supported Hadoop/Oracle Database combinations there is a uniform set of steps to update the Oracle Big Data SQL configuration when there are changes to the Hadoop cluster or Oracle Database server.

New Features to Simplify ILM – Oracle Database Tablespaces in HDFS (With Smart Scan Technology)

Oracle Database ILM (Information Lifecycle Management) can now be extended to use Hadoop to store read-only Oracle Database tablespaces. When you move tablespaces from Oracle Database to HDFS, the tables, partitions, and data retain their original Oracle Database internal format, remain accessible to queries, and support the full range of Oracle Database performance optimizations and security features, including the following:

- Smart Scan for HDFS, which enables off-load of query processing to Oracle Big Data SQL on the Hadoop cluster. Smart Scan also provides filtering of query results in Hadoop prior to the return of the data to Oracle Database. In most circumstances, this can be a significant performance optimization. Indexing, Hybrid Columnar Compression, Partition Pruning, and Oracle Database In-Memory are also supported.
- Oracle Advanced Security Option (ASO) Transparent Encryption and Data Redaction.

Tablespaces stored in HDFS are read-only, therefore this storage is best suited to data archiving.

See Section 3.2 in the *Oracle Big Data SQL User's Guide* for details.

Enhancements in Oracle Shell for Hadoop Loaders 1.2

Oracle Shell for Hadoop Loaders (OHS) is an intuitive command line tool for data migration. You can set up resources to connect to Hive, HDFS, or Oracle Database, and then access each of these data sources through the uniform OHS interface. OHS is one of the ways to use Copy to Hadoop. Copy to Hadoop users can download OHS from [OTN](#).

OHS 1.2 includes the following changes:

- On-disk logging of load operations in the `$HOME/.ohsh` shadow directory.
- The ability to minimize output when doing load commands. (See the help command for `set outputlevel`.)
- Loading Hive tables from Oracle tables not living in the `oracle` user's schema.
- Wallet and TNS usage by OHS relies on the setting of environmental variables `WALLET_LOCATION` and `TNS_ADMIN`. The `set tnsadmin` and `set walletlocation` commands are no longer supported.

In addition, you no longer set `HIVE0_URL` to the fully-qualified URL of remote HiveServer2 in order to create a `%hive0` resource. In OHS 1.2, set the environmental variable `HS2_HOST_PORT` in `bin/ohsh`, which is the `<hostname>:<port>` pair of HiveServer2.

Enhancements to Copy To Hadoop

The new method `directcopy` is added to Copy to Hadoop.

This is a direct, single-step method of copying data from Oracle Database to HDFS. See *Using Copy to Hadoop to do Direct Copies* in the *Oracle Big Data SQL User's Guide* for more information.

Granting Access – Users Now Require the BDSQL_USER Role

In prior Oracle Big Data SQL releases, all users are granted Big Data SQL access implicitly. Release 3.1 introduces the `BDSQL_USER` role. Users requiring Oracle Big Data SQL access must be granted this role explicitly.

You must also now grant read privileges on the `BigDataSQL` configuration directory object.

For example, to grant access to `user1`:

```
SQL> grant BDSQL_USER to user1;
SQL> grant read on directory ORACLE_BIGDATA_CONFIG to user1;
```

Installation Instructions Moved to Oracle Big Data SQL Installation Guide

The *Oracle Big Data SQL Installation Guide* provides instructions how to install and uninstall the software. In releases prior to Oracle Big Data Appliance 3.1, installation instructions are in the user's guide.

F.2 Changes in Oracle Big Data SQL 3.0.1

The following are changes in Oracle Big Data SQL Release 3.0.1.

In addition to continued support for Oracle Big Data Appliance and the Exadata Database Machine, Release 3.0.1 includes further improvements to Big Data SQL support for platforms other than Oracle Engineered Systems. These currently include Hortonworks Data Platform (HDP) and Cloudera's Distribution Including Apache Hadoop (CDH) running on non-Big Data Appliance commodity hardware where these Hadoop systems are used in conjunction with Oracle Database running on non-Exadata servers.

Support for Newer CDH and HDP Platforms

Oracle Big Data SQL 3.0.1 runs on Hortonworks Hadoop Data Platform (HDP) 2.3 – 2.4.x and Cloudera's Distribution Including Apache Hadoop (CDH) 5.4.7 – 5.7. CDH 5.4.7 support includes both commodity CDH systems and Oracle Big Data Appliance 4.3.

Improved Installation Automation

The installation on HDP and CDH now automatically saves the YARN resource management configuration for later restoration.

The `oracle` user and `oinstall` group setup are now managed by the installation program.

Uniform Uninstall Procedure for HDP and CDH

The uninstall procedure is now the same for both HDP and CDH systems.

Other Changes for HDP and CDH

- The default JDK is upgraded to `jdk-8u66-linux-x64`
- `bds-database-create-bundle.sh` includes a new `--jdk-tar-path` parameter for specifying a JDK other than the default.
- `bds-database-install.sh` includes two new parameters:
 - `--cell-config-dir`
 - `--root-script-only`
- The installation process now creates and populates `cellinit.ora` if the file does not exist or is empty.

Oracle Shell for Hadoop Loaders

Oracle Shell for Hadoop Loaders is a helper shell that provides a simple-to-use command line interface to Copy to Hadoop, Oracle Loader for Hadoop, and Oracle SQL Connector for HDFS.

Support for Oracle Big Data Appliance 4.5

Oracle Big Data SQL Release 3.0.1 is a mammoth installation option in Oracle Big Data Appliance 4.5. You can also enable Oracle Big Data SQL Release 3.0.1 at any time after the mammoth installation by using the `bdacli` utility.

 **Note:**

In Release 3.0.1, Oracle Big Data SQL does not yet support connections between Oracle Big Data Appliance and non-Exadata Oracle Database servers, or between non-Big Data Appliance Hadoop systems and the Exadata Database Machine.

Index

Symbols

--aux-run-mode, [3-11](#)
--crs, [3-11](#), [6-2](#)
--db-name, [3-11](#)
--db-resource, [6-2](#)
--grid-home, [3-11](#)
--ip-cell, [3-1](#), [3-11](#)
--no-root-script, [3-11](#)
--reconfigure, [3-11](#), [4-2](#)
--reqkey, [3-6](#), [3-11](#)
--root-script, [3-11](#)
.reqkey file, [2-14](#), [3-6](#)

A

Acknowledge key, [1-10](#)

B

bdacli, [2-1](#)
bds-config.json, [2-4](#), [2-14](#), [3-6](#), [4-2](#), [5-1](#), [5-4](#)
bds-database-install.sh, [6-2](#)
bdschecksw, [1-12](#)
BDSQL_USER, [3-6](#), [3-12](#), [F-3](#)
BigDataSQL configuration directory object, [3-12](#),
[F-3](#)

C

cellint_ora, [3-6](#)
Commodity systems, [1-1](#)
configuration file, [2-4](#), [2-14](#), [4-2](#), [5-1](#)
Copy to Hadoop, [8-1](#)
CRS, [3-11](#), [6-2](#)

D

Database Acknowledge, [1-12](#), [3-11](#)
Database Authentication, [1-10](#), [2-1](#), [2-4](#), [5-1](#)
database_auth_enabled, [2-1](#), [2-4](#)
databaseack, [1-10](#), [2-3](#), [3-6](#)
databasereq, [1-10](#), [2-3](#)
DataNode, [4-2](#), [5-4](#)

DataNodes, [2-1](#)
dbkeys directory, [2-14](#), [3-6](#)
Default cluster, [1-10](#), [6-1](#)
dir, [2-4](#)
diskmon process, [3-3](#)

E

Ethernet, [1-1](#), [2-1](#), [5-1](#)

F

FUSE-DFS, [5-1](#), [6-2](#), [8-2](#)

G

Grid, [1-12](#), [3-1](#), [3-3](#), [3-11](#)
GUID-key pair, [2-3](#)

H

Hadoop Secure Impersonation, [2-4](#)
HDFS Transparent Encryption, [7-6](#)
HIVE_SESS_VAR_LIST, [8-2](#), [8-5](#)
HTTPS/ Network Encryption, [7-6](#)

I

impersonation_blacklist, [2-1](#)
impersonation_enabled, [2-1](#), [2-4](#)
InfiniBand, [2-4](#)
InfiniBand, [1-1](#), [2-1](#), [5-1](#)
install, [2-3](#)
Installation directory, [2-14](#), [3-5](#), [3-6](#), [4-2](#)

J

Jaguar, [1-12](#), [2-1](#), [2-3](#), [2-14](#), [3-5](#), [3-6](#), [4-2](#), [5-1](#),
[5-4](#)

K

KDC, [3-1](#)
Kerberos, [1-6](#), [2-1](#), [5-1](#), [8-2](#), [8-5](#)

KRB5_CONF, [3-1](#)

L

licenses, third-party, [E-1](#)

locationdirectory, [8-2](#)

M

MTA

--mta-destroy, [3-11](#)

--mta-restart, [3-11](#)

--mta-setup, [3-11](#)

N

name, [2-4](#)

O

OHS, [8-1](#), [8-2](#), [8-5](#)

Oracle Data Pump, [8-1](#)

Oracle Instant Client, [8-5](#)

Oracle Loader for Hadoop, [8-1](#)

Oracle Shell for Hadoop Loaders, [8-1](#)

Oracle Tablespaces in HDFS, [1-7](#), [5-1](#)

Oracle Wallet, [8-5](#)

ORACLE_HOME, [3-6](#), [5-4](#)

ORACLE_SID, [3-6](#)

ORC, [8-1](#)

P

Parquet, [8-1](#)

port, [2-4](#)

ports, [2-1](#)

Python, [1-2](#)

R

RAC database, [6-1](#)

RAC systems, [3-1](#)

reconfigure, [2-3](#), [3-5](#), [4-2](#), [5-1](#), [5-4](#)

Request key, [1-10](#), [2-14](#), [3-6](#), [3-11](#)

restart_stale, [2-4](#)

Restarts, [3-3](#)

S

selection_subnet, [2-1](#)

selection_subset, [2-4](#)

SQL*Plus Instant Client, [8-5](#)

SSH, [3-1](#)

SSL, [8-5](#)

T

third-party licenses, [E-1](#)

TNS, [8-5](#)

U

uninstall, [2-3](#), [6-1](#), [6-2](#)

updatenodes, [2-3](#), [3-5](#), [4-2](#), [5-4](#)

url, [2-4](#)

use_infinband, [2-1](#)

use_infiniBand, [2-4](#)

User access, [3-12](#)

users

BDSQL_USER, [F-3](#)