

DOCUMENTATION

BDNA Discover™ 7.7.2

Installation Guide

November 11, 2016

LEGAL NOTICES

Use of the BDNA software and products are subject to the terms and conditions of a license agreement found in either a separately executed master license agreement or the click-through master license agreement that is accepted prior to delivery of the BDNA software and/or products.

Copyright © 2001-2016. BDNA Corporation

Information in this manual and all BDNA technical support policies are subject to change without notice. Check with your BDNA authorized representative to ensure that you have the most recent information.

BDNA®, the BDNA logo, Technopedia®, BDNA Discover™, BDNA Normalize®, and BDNA Analyze™ are trademarks or registered trademarks of BDNA Corporation in the United States and internationally.

The products described herein may be technically combined with third party products or other products not supplied by BDNA, including third party or customer software, hardware, and materials. Any combinations or potential combinations described herein are advisory only. BDNA expressly disclaims any liability, and any expressed or implied representation and warranty, resulting from any combinations of the BDNA products with any products not supplied by BDNA.

This document is provided “as is” and without warranty of any kind. BDNA and its licensors (hereinafter collectively referred to as “BDNA”) expressly disclaim all warranties, expressed or implied, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose and against infringement.

Oracle is a trademark or registered trademark of Oracle Corporation in the United States and in other countries. Red Hat Enterprise Linux (RHEL) is a trademark or registered trademark of Red Hat Incorporated in the United States and in other countries. Microsoft Internet Explorer is a trademark or registered trademark of Microsoft Corporation in the United States or in other countries. All other trademarks appearing herein are the property of their respective owners.

Confidential and Proprietary to BDNA.

BDNA Corporation
339 North Bernardo Avenue, Suite 206
Mountain View, CA 94043
USA
Phone +1 650 625 9530
Fax +1 650 625 9533
<http://www.bdna.com>
02500010101

Contents

Chapter 1. Installation Overview	11
About this Book	11
Intended Audience	11
Supported Browsers	11
Understanding the BDNA Discover Installation	11
Before You Begin the BDNA Discover Installation	14
Time Synchronization for BDNA Discover Servers	14
Recommended Subnet Location	14
Chapter 2. Hardware Requirements	15
About this Chapter	15
Oracle Database Server	15
Linux Collection Store and Factbase Servers Requirements	16
Windows Collection Server Requirements	17
Client Requirements	17
Sizing a Collection Store Cluster	17
Chapter 3. Installing BDNA Discover Database Servers	19
About this Chapter	19
Installation Prerequisites for the Oracle Database Server	20
Oracle Database Server Installation Requirements	20
Oracle Access Grants Requirements	21
Double-Byte Character Set Installation for East Asian Languages	21
Known Problems and Workarounds for Oracle11g	22
Ensuring Sufficient Oracle Tablespace Free Space	22
Impact of Oracle Temp Tablespace Bug on BDNA Discover Software Performance	24
Installing BDNA Discover Using an Existing Oracle Database Server	24
Verifying and Creating Tablespaces for BDNA Discover	25
Verifying Oracle Database Settings	25
Setting cursor_sharing Option to EXACT	26
Turning Off recyclebin Option	26
Chapter 4. Installing Linux for Discover Component Servers	27
About this Chapter	27
Installation Prerequisites	28

Installing the Oracle Client and Verifying Oracle Server Settings	28
Chapter 5. Installing Discover on the Linux Component Server	29
About this Chapter	29
Discover Linux Installer	29
Before You Begin	31
Copying the BDNA Discover Linux Installer to the Linux Component Server	31
Running the Linux Installer	32
Command-Line Options	32
Overview of Installer Operations	33
Database Compatibility Check	33
Hardware Compatibility Check	33
Software Compatibility Check	34
Shell	34
Kernel Parameters	35
Response File	35
Configuring BDNA Discover for the First Time	35
Before You Begin	35
Configuring the Discover Collection Store and FactBase Server(s)	36
Distributed Collection Store Installation	41
Distributed Component Layouts	42
Configuring the FactBase Server	44
Chapter 6. Installing Windows Collection Servers	49
About this Chapter	49
Installing and Configuring Windows Collection Server	49
Configuring Maximum Concurrent Tasks	54
Appendix A. Troubleshooting	55
Setting Up BDNA Discover for Access from a Browser	55
Time Synchronization for BDNA Discover Servers	55
Appendix B. Required RPM Packages	57
RHEL 6.8 RPM Packages	57
PERL RPM Packages	58
RPM Packages for Nmap 4.11	58

About this Book

Welcome to the BDNA Discover™ 7.7.2 Installation Guide. This guide provides instructions for installing and configuring BDNA Discover 7.7.2 components, including the following:

- [Chapter 1, “Installation Overview”](#)
- [Chapter 2, “Hardware Requirements”](#)
- [Chapter 3, “Installing BDNA Discover Database Servers”](#)
- [Chapter 4, “Installing Linux for Discover Component Servers”](#)
- [Chapter 5, “Installing Discover on the Linux Component Server”](#)
- [Chapter 6, “Installing Windows Collection Servers”](#)

This guide also provides instructions on how to prepare a client system to run the BDNA Discover client applications.

Intended Audience

You should be familiar with the following components and software:

- Oracle 11g R2 (64-bit) Standard or Enterprise Edition. or Oracle 12c Enterprise Edition
- Red Hat Enterprise Linux (RHEL) 6.5-6.8 (64-bit).
- Java Runtime Environment (JRE) 8 or higher.
- Windows Server 2008 R2 (64-bit) or Windows Server 2012 R2 (64-bit).

This chapter discusses the general architecture of BDNA Discover, and covers some key issues that you should consider prior to installation.

Supported Browsers

Browsers supported in this release of BDNA Discover are:

- Internet Explorer 11 or higher.
- Firefox 31 or higher.
- Google Chrome 36 or higher.

Understanding the BDNA Discover Installation

BDNA Discover is a Web Services application that is easily distributed across multiple servers. It consists of:

- **Collection Store:** A database containing the latest scan data. The scanned data is extracted to a file and is referred to as a “Collection.” The Collection is then imported into Factbase (see definition below).

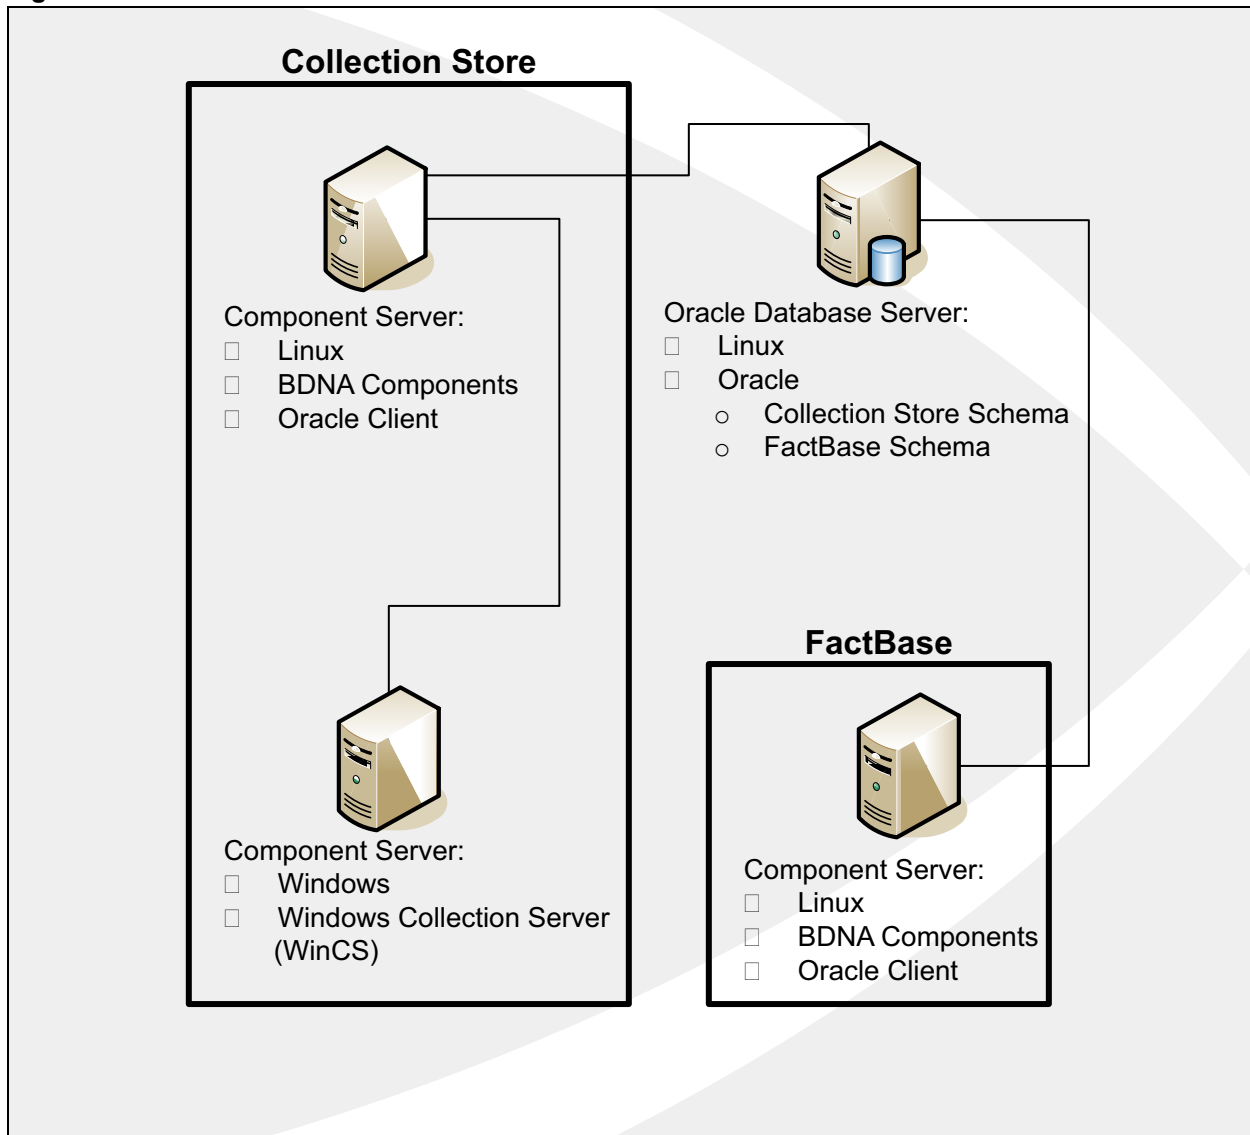
At a minimum, the Collection Store application requires two Linux machines; one for running the application that performs the scans, and one for the Collection Store database. To improve performance, the Collection Store application can be distributed across multiple servers. Additionally, a Windows collection server may be added to collect detailed Windows system and application information.

- **Factbase:** A persistent database containing data (Collections) from previous scans of all locations. Collections, created from scans performed using the Collection Store, can be added to inventories, and used for reporting.

Factbase requires two Linux machines; one for running the application and one for the Factbase database. (Although the application and database can run on the same machine, you may experience a slowdown in performance with this configuration.)

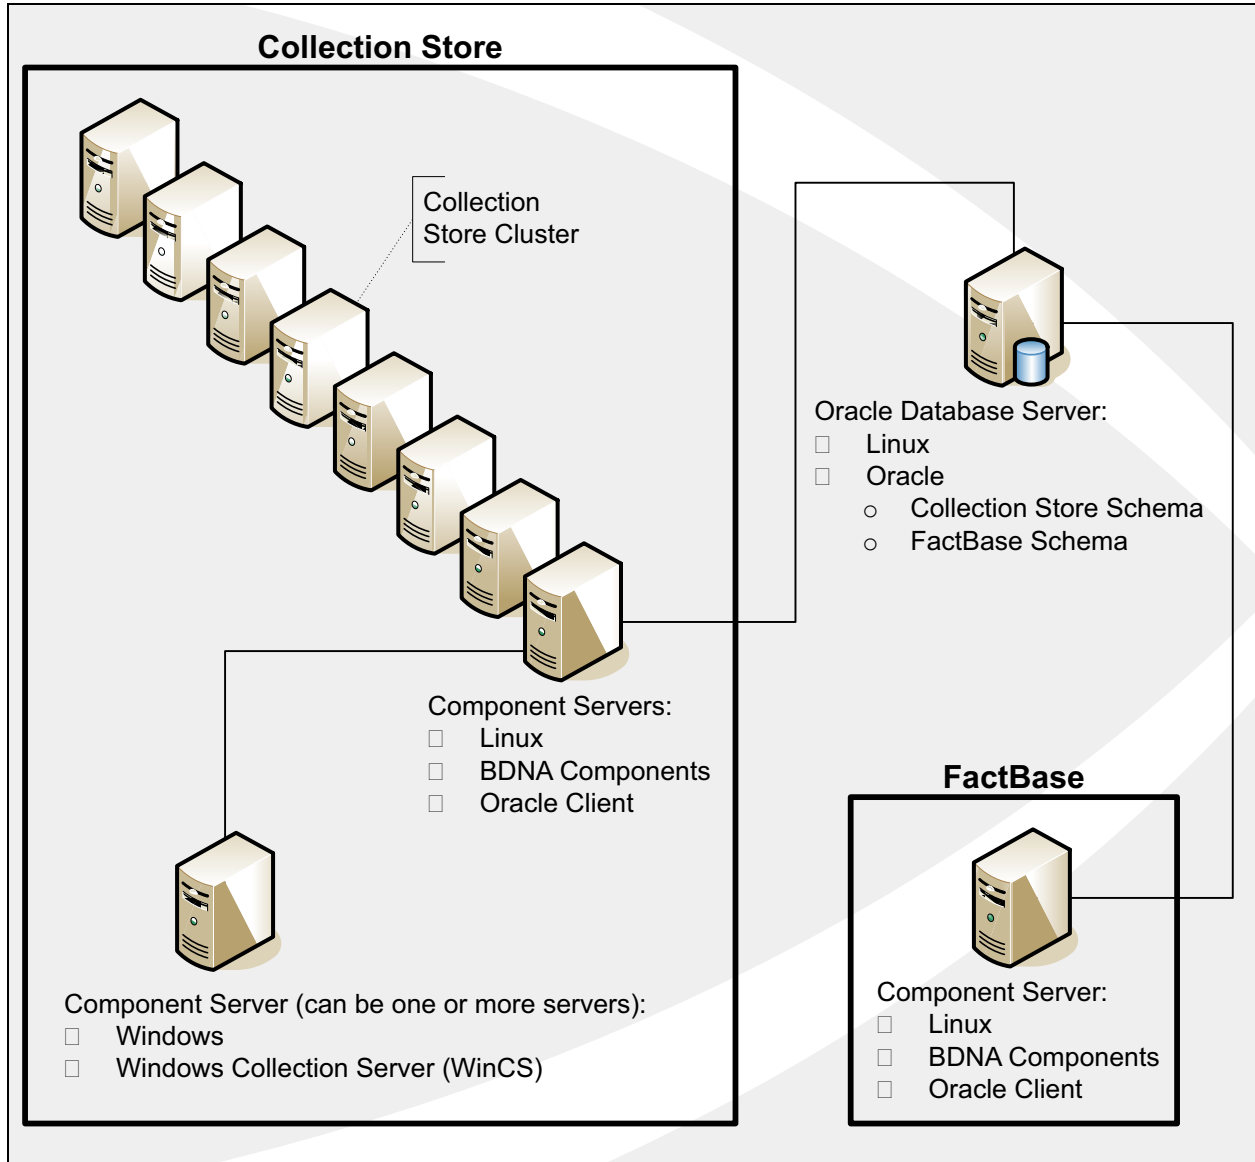
Note: The database server for Collection Store and Factbase can reside on the same machine.

Figure 1: Basic BDNA Discover Installation



- **Collection Store Cluster:** The BDNA Discover Collection Store component servers, together with the database server, are referred to as the Collection Store cluster. The Collection Store cluster is an installation that uses multiple servers. A Collection Store cluster is comprised of three types of servers:
 - Oracle Database Server
 - Linux Component Servers
 - Windows Collection Servers

Figure 2: Cluster-based BDNA Discover Installation



Before You Begin the BDNA Discover Installation

Before installing BDNA Discover, the BDNA Discover Administrator must consider the overall configuration of the BDNA Discover cluster. Questions to consider include:

Table 1: Configuring the BDNA Discover Cluster

Question	Consider the following...
Will the BDNA Discover cluster be used predominantly to collect Windows Level 2 inventory data from Windows systems?	If so, it may need a larger number of Windows Collection Servers.
Will the BDNA Discover cluster be used predominantly to collect SNMP Level 2 inventory network infrastructure discovery data (using SNMP community strings)?	If so, it may need more Linux Component Servers than Windows Collection Servers.
Will the BDNA Discover cluster be used predominantly to collect UNIX Level 2 inventory network infrastructure discovery data (using SNMP community strings)?	If so, it may need more Linux Component Servers than Windows Collection Servers.
How much inventory needs to be collected, and what is the time frame?	The database server for any BDNA Discover cluster will need to be appropriately sized to support the amount of inventory data being collected and the amount of time permitted for data collection (throughput requirements)

Note: See [Table 5](#) for specific information about BDNA Discover cluster setup.

Time Synchronization for BDNA Discover Servers

BDNA Discover Linux server clocks must be in sync with each other, and they also must be in sync with the Oracle database server clock. Linux servers must be configured to interact with a valid Network Time Protocol (NTP) server that is either part of your intranet or, if your firewall allows outbound NTP connections, with an Internet NTP server. Refer to the Linux documentation for more detailed information.

Similarly, if your BDNA Discover Windows servers are not configured to join a Windows domain, you will need to configure them to interact with a valid Windows time server that is either part of your intranet or, if your firewall allows outbound Windows time service connections, with the Internet Windows time servers. Refer to the Microsoft Windows documentation for more detailed information.

Recommended Subnet Location

BDNA recommends that all BDNA Discover servers be on the same subnet, and have 100 Mbits Fast Ethernet (or faster) network connections between them. The BDNA Discover servers should be located on a high-speed network that can access the machines to be scanned with as few network obstructions (firewalls, slow links, or packet filters) as possible.

About this Chapter

This chapter provides an overview of the hardware requirements for each of the BDNA Discover component servers:

- [“Oracle Database Server”](#)
- [“Linux Collection Store and Factbase Servers Requirements”](#)
- [“Windows Collection Server Requirements”](#)

Oracle Database Server

The Oracle databases (Collection Store and Factbase), act as a back-end repository for storing configuration, transactional, and analytical information. The hardware required to support the databases is dependent on the following factors:

- Number of devices scanned at Level 1
- Number of devices scanned at Level 2
- Time frame for scanning all of the systems
- Operating system running Oracle

An Oracle database server must have the following minimum configuration:

- Oracle 11g R2 (11.2.0.1 or higher) 64-bit Standard or Enterprise Edition
 - ...**or** Oracle 12c (12.1.0.2 or higher) 64-bit Enterprise Edition
- Required RPM Packages (see [“Required RPM Packages”](#) for information)
- Fast, dual processors
- 8GB of RAM
- At least 500GB of free disk space

This document does not provide detailed instructions on how to size an Oracle database. As a starting point, a typical BDNA Discover database server requires a minimum of 500GB of disk space, depending on the amount of data to be collected over time.

[Table 2, “Examples of Oracle Database Hardware Configurations for BDNA Discover Repositories,”](#) provides a few real-world examples of database hardware configurations for BDNA Discover repositories. These examples show the number of devices discovered at various scan levels over the scan times given. Use these examples as guidelines for database server sizing, in consultation with an Oracle DBA.

Table 2: Examples of Oracle Database Hardware Configurations for BDNA Discover Repositories

Level 1 Devices	Level 2 Devices	Scan Time	Hardware	USERS Free Space Required	BDNATEMP Tablespace Recommended Size
25,000	5,000	12 hours	2 CPU cores, 8GB RAM, 500GB Disk	10GB	1GB
100,000	25,000	24 hours	2 CPU cores, 8GB RAM, 1TB Disk	50GB	5GB
350,000	100,000	5 days	4 CPU cores, 8GB RAM, 2TB Disk	100GB	10GB

Note: These values are meant for rough planning purposes only. During initial operations, monitor the free space in both USERS and BDNATEMP carefully to ensure that sufficient free space is available.

Linux Collection Store and Factbase Servers Requirements

The Linux Collection Store component servers and the Factbase server are the workhorses of the BDNA Discover system. They run all major BDNA Discover Collection Store software components, including:

- Rule Engine
- Collection Engine
- Agenda Manager
- Remote Services
- Factbase Application

Similar to Windows collection servers, Linux-based servers are typically CPU- and memory-bound, rather than disk- or network I/O-bound. As such, Linux-based component servers should be configured to accommodate the high peak workloads they will encounter.

A BDNA Discover component server must have the following minimum configuration:

- Red Hat Enterprise Linux (RHEL) 6.5 or 6.8 (64-bit)
- Java SE 8 (64-bit) with Java Runtime Environment (JRE) 8 or higher
- Fast, dual processors
- 8GB of RAM
- At least 500GB of free disk space

Windows Collection Server Requirements

Windows Collection Servers run the Windows Collection Service (WinCS), which enables Windows Level 2 collection capabilities. Each WinCS machine must have the following minimum configuration:

- Windows Server 2008 R2 (64-bit) or Windows Server 2012 R2 (64-bit)
- Fast processor
- 8GB of RAM
- At least 250GB of free disk space

Client Requirements

To access the BDNA Discover User Interface (UI), you will need the following:

- Java Runtime Environment (JRE) 8 or higher

For instructions on installing JRE 8, go to the Oracle website:

http://docs.oracle.com/javase/8/docs/technotes/guides/install/install_overview.html

For access to the downloadable JRE 8, go to the Oracle website:

<http://www.oracle.com/technetwork/java/javase/downloads/jre8-downloads-2133155.html>

- Browser
 - Internet Explorer 11 or higher
 - Firefox 31 or higher
 - Google Chrome 36 or higher

Note: Before you can access the BDNA Discover UI for the first time, you must download and run the `AppConsole.jnlp` file.

A message will display when you try to access the BDNA Discover UI, asking if you want to open or save the `AppConsole.jnlp` file. BDNA recommends that you save the `,jnlp` file to your desktop so that in the future you only need to double-click the `AppConsole` icon to access BDNA Discover UI.

Sizing a Collection Store Cluster

Sizing a Collection Store cluster requires an evaluation of the database server, Linux component servers, and Windows collection servers as a whole. Sizing decisions require the administrator to review organizational goals for scan time, as well as many other factors.

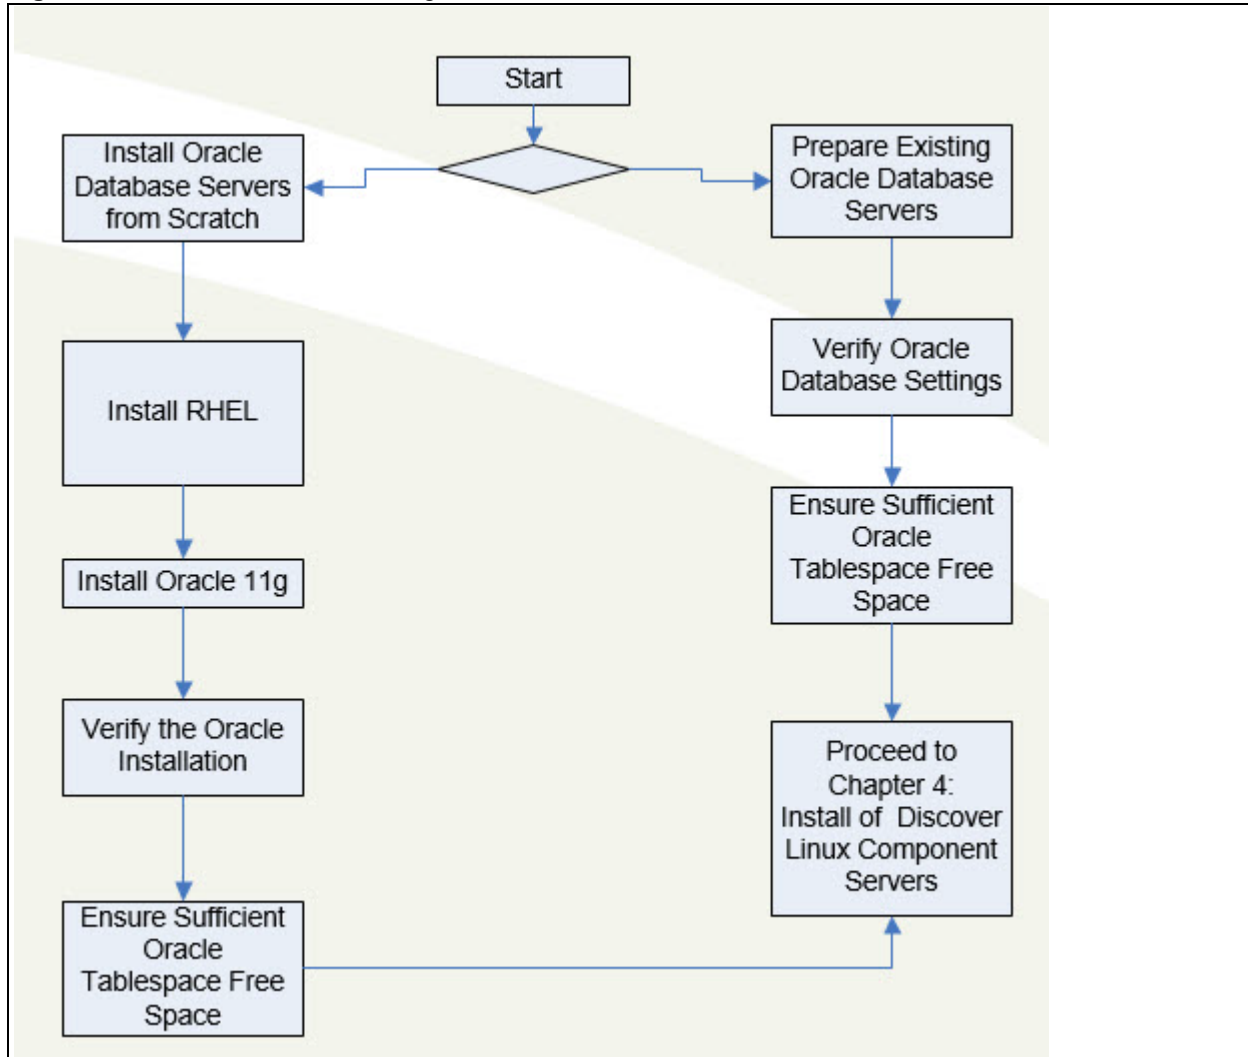
About this Chapter

This chapter provides information for installing the Oracle database server, specifically for a Collection Store and FactBase database.

Note: BDNA Discover is tested and certified with the Oracle database server running on Red Hat Linux. BDNA Discover supports users who have an Oracle instance on a platform supported by Oracle, as long as that version complies with the BDNA Discover requirements stated in the Release Notes for the version of BDNA Discover in use.

This chapter provides information about the following:

- [“Installation Prerequisites for the Oracle Database Server”](#)
- [“Ensuring Sufficient Oracle Tablespace Free Space”](#)
- [“Impact of Oracle Temp Tablespace Bug on BDNA Discover Software Performance”](#)
- [“Installing BDNA Discover Using an Existing Oracle Database Server”](#)

Figure 3: Process Flow for Installing BDNA Discover Database Servers

Note: To install Oracle Database Server and Client from scratch, refer to Oracle documentation.

Installation Prerequisites for the Oracle Database Server

Note: For double-byte installations, refer to [“Double-Byte Character Set Installation for East Asian Languages.”](#)

Oracle Database Server Installation Requirements

- Basic knowledge of Linux administration
- Basic knowledge of Oracle administration
- Server(s) for the Oracle and BDNA Discover installations

- Host machine to run the Oracle Client software (may be the same machine as the server on which Oracle is installed)
- Oracle 11g R2 (11.2.0.1 or higher) 64-bit Standard or Enterprise Edition
 - ...**or** Oracle 12c (12.1.0.2.0) 64-bit Enterprise Edition
- X-Windows capable system (for example, Cygwin on Windows).
- Putty (or other remote access app)
- WinSCP (or other file transfer app)
- Required RPM Packages (see “[Required RPM Packages](#)” for information about required RPM packages)

Note: Both the Oracle Server and the Oracle Client must match versions.

If the RHEL ISO image is from a download, use CD-writing software to burn it onto CDs.

The Oracle software can be installed from a CD or from a download. We recommend placing the Oracle installation files in a staging directory on the machine where Oracle is to be installed.

In addition, the installation requires the following items:

- Linux server built with certified RedHat compatible hardware that meets the BDNA Discover minimum server hardware requirements

Oracle Access Grants Requirements

The following Oracle access grants are required for the BDNA Discover user:

- Alter Session
- Create Session
- Create Synonym
- Create View
- Query Rewrite
- Create Materialized View
- Resource
- Create Catalog Role

Double-Byte Character Set Installation for East Asian Languages

BDNA Discover supports data collection with double-byte character sets (DBCS) such as Chinese and Japanese.

Refer to the *BDNA Discover 7.7.2 Release Notes* for updated information regarding prerequisites.

Be sure to configure the Oracle database with the UTF-8 character set.

Note: If your Oracle instance is already configured with a character set other than UTF8, you need to recreate the database by using the Oracle DBCA tool.

Known Problems and Workarounds for Oracle11g.

The following are known problems and workarounds for Oracle 11g.

SELinux:

After a successful Oracle 11g DB server installation, running sqlplus from \$ORACLE_HOME may produce the following error:

```
./sqlplus: error while loading shared libraries:
/u01/app/oracle/product/11.1.0/lib/libnzn11.so: cannot restore segment prot
after reloc: Permission denied
```

This is a security enhancement for Oracle 11g. To work around the issue, in `/etc/selinux/config` change:

```
SELINUX=enforcing
to
SELINUX=permissive
```

Case-Sensitive Passwords:

By default, the schema passwords in Oracle 11g are case-sensitive. If you would like to disable the case sensitive login, run the following SQL command.

```
sql> alter system set sec_case_sensitive_logon=False
```

BDNA Discover Scan Administration Reports:

When running BDNA Discover on Oracle 11g R2, you might experience an Oracle ORA-00600 internal error when accessing one of the Scan Administration reports such as Level 1/2 Summary Report. To work around this problem, update `init.ora` with the following setting:

```
_optimizer_push_pred_cost_based=false;
```

Building Reports:

Building reports fails due to Oracle bug #10408668 when running BDNA Discover 7.6.0 on Oracle 11g R2. To work around this problem, the user must log in as the Oracle system user and execute the following:

```
alter system set "_optimizer_distinct_agg_transform"=false;
```

Ensuring Sufficient Oracle Tablespace Free Space

In many cases, the default Oracle installation is pre-configured with a minimal tablespace sizing, which is not sufficient to store the BDNA Discover schema user. To make an Oracle database server ready for BDNA Discover, it is necessary to increase the tablespace sizes on the Oracle database server. You can use the following BDNA Discover SQL scripts to pre-allocate tablespace sizes:

To access the BDNA Discover SQL scripts:

1. Locate `Oracle_Tablespace_Update.zip`
2. Use WinSCP or any SFTP file transfer utility to transfer this zip file onto the Oracle Database server, under the `/home/oracle` directory.
3. Use Putty to log on to the Oracle Database server as "oracle" OS user.
4. Go to the directory where `Oracle_Tablespace_Update.zip` is located.
5. Run the following command to extract the file:

```
$ unzip Oracle_Tablespace_Update.zip
```

6. After unzipping the file, you should see two SQL script files:

```
postinstall_MEDIUM_67G.sql
```

```
postinstall_LARGE_152G.sql
```

To determine available disk space:

Use the `df -h` command to check the available disk space on the `/oracle` partition.

To run the SQL script to allocate tablespace:

After determining available disk space, complete the following steps.

1. Log in to the system as the Oracle OS user (`oracle` is the default name and password for the Oracle user):
2. Use one of the following SQL scripts to allocate tablespace:

postinstall_MEDIUM_67G.sql

The `postinstall_MEDIUM_67G.sql` script requires that the `/oracle` partition have a minimum of 67 GB of free disk space. This is the recommended tablespace sizing for scanning a network containing 10,000 active IPs.

postinstall_LARGE_152G.sql

The `postinstall_LARGE_152G.sql` script requires that the `/oracle` partition have a minimum of 152 GB of free disk space. This is the recommended tablespace sizing for scanning a network containing 20,000 active IPs.

For example:

```
$ sqlplus system/bdna
SQL> @postinstall_MEDIUM_67G.sql
SQL> exit
```

Note: The existing Oracle database should be ready to support BDNA Discover. BDNA recommends changing the default passwords of the root OS user, Oracle OS user, and the BDNA Discover OS user. The default password is "bdna" for the root and BDNA Discover OS users, and "oracle" for the oracle OS user.

Note: If problems with the Oracle installation or connectivity are encountered, consult the Oracle documentation. The Oracle documentation can be accessed online at the following URL:
www.oracle.com/technology/documentation/database11g.html

Impact of Oracle Temp Tablespace Bug on BDNA Discover Software Performance

A bug in the Oracle database software may affect the performance of report refresh in BDNA Discover. The workaround for this issue follows.

This issue applies to Oracle Server, Enterprise Edition, Versions 11gR1. The problem, which can occur on any OS platform, arises when running `refreshanalytics`, especially for large repositories. When running the respective reports-refresh command, the Oracle database server can run out of TEMP tablespace and generate the following Oracle error:

```
ORA-00600: internal error code, arguments: [32695], [hash aggregation can't be done]
```

Additionally, you might notice the following error message in `$BDNA_HOME/logs/dw.log`:

```
ORA-01652: unable to extend temp segment by 64 in tablespace TEMP
```

Moreover, because this problem manifests itself in different ways, even if you do not encounter these error messages, you could still be affected by this issue. If you suspect a performance problem when refreshing reports, BDNA recommends that you apply the workaround documented in this section.

To work around the problem, perform the following steps *before refreshing reports*:

1. Run SQL*Plus and connect as the Oracle system user or sysdba user.
2. Run the following SQL:

```
sql> alter system set "_gby_hash_aggregation_enabled"=false;
sql> exit
```

Note: Running the SQL commands in step 2 on ia a **temporary** fix that must be repeated each time the Oracle server is restarted. To preserve this setting after the Oracle server is restarted, edit the file `/u01/app/oracle/admin/ora11g/pfile/init.ora`.

Hash aggregation is new to Oracle 11gR1. The Oracle Optimizer in 11gR1 automatically uses hash aggregation for SQL `GROUP BY` clauses. This bug causes the database to consume large amounts of tablespace and results in a long-running operation. Oracle documented this issue in a technical article that is available on the My Oracle Support web portal; the document ID is 418823.1. You need login credentials in order to access this portal.

Installing BDNA Discover Using an Existing Oracle Database Server

Performing a BDNA Discover installation with an existing instance of Oracle involves two primary procedures:

- “[Verifying Oracle Database Settings](#)” on page 25
- “[Verifying and Creating Tablespaces for BDNA Discover](#)” on page 25

- [“Ensuring Sufficient Oracle Tablespace Free Space”](#) on page 22

To use an existing Oracle database with BDNA Discover, the database must have certain features enabled. BDNA Discover uses a number of standard features of the Oracle database that Oracle installs by default. However, some organizations may have removed these features. Verify that the following internal Oracle supplied packages are installed, operating, and executable by the BDNA Discover user:

- DBMS_UTILITY
- DBMS_LOB

Note: Although not a recommended configuration, if you have installed BDNA Discover and the Oracle instance on the same server, the XML DB feature of Oracle can cause conflicts with BDNA Discover. Specifically, if the database is running on a host that also runs the BDNA Discover Remote Service component, both XML DB and BDNA Discover attempt to attach to and listen on port 8080. If BDNA Discover Remote Services is not running on the same host, this is not an issue.

Verifying and Creating Tablespaces for BDNA Discover

Prior to creating the BDNA Discover user, the BDNA Discover Administrator must verify that there is sufficient storage space available in the database. The BDNA Discover user account must be created in the USERS tablespace. The USERS tablespace is created by Oracle during a default installation. The size of the USERS tablespace depends on the number of assets being discovered, the amount of data discovered for each asset, and the number of inventories retained in the FactBase. (Thus, it varies for every organization.)

Some organizations may have chosen to remove the USERS tablespace, or the existing USERS tablespace may not have the parameters recommended for BDNA Discover. If a USERS tablespace does not exist, the Oracle DBA will need to create one. BDNA Discover also requires a second tablespace named BDNATEMP to be created. The BDNATEMP tablespace is used to store temporary data during discoveries.

Verifying Oracle Database Settings

To ensure compatibility and good performance with the Oracle database, BDNA strongly recommends setting the following parameters to the values specified. These values should be sufficient for larger scan discovery. However, the Oracle DBA may choose to adjust these values, if only smaller scans will be performed. Be sure to verify these settings with the Oracle DBA:

- UNDO_RETENTION—recommended undo retention: 10800
- TRANSACTIONS—recommended transactions: 610
- CURSOR_SHARING—must be set to: EXACT
- RECYCLEBIN—must be set to: OFF
- `_gby_hash_aggregation_enabled`—must be set to: FALSE
- `_session_cached_cursors`—must be set to: 400
- `_optimizer_push_pred_cost_based`—must be set to: false

Setting cursor_sharing Option to EXACT

The `cursor_sharing` parameter affects the number of hard parses and soft parses in an Oracle database. It can be set to `EXACT`, `SIMILAR`, or `FORCE`. For a database setup used exclusively for BDNA Discover software, it must be set to `cursor_sharing=EXACT` in the database initialization parameters file, `init.ora`.

To verify that the `cursor_sharing` option is set to `EXACT`, run this query:

```
select value from v$parameter where name='cursor_sharing';
```

Turning Off recyclebin Option

The `DROP TABLE` command moves the table to the recycle bin, rather than actually dropping it. This could significantly impact the performance of the BDNA Discover system. BDNA strongly recommends that the `recyclebin` parameter be set to `OFF` in the Oracle initialization parameters file:

```
recyclebin=OFF
```

To verify that the `recyclebin` option is turned off, run this query:

```
select value from v$parameter where name = 'recyclebin';
```

The existing Oracle database is now ready to support BDNA Discover.

Installing Linux for Discover Component Servers 4

About this Chapter

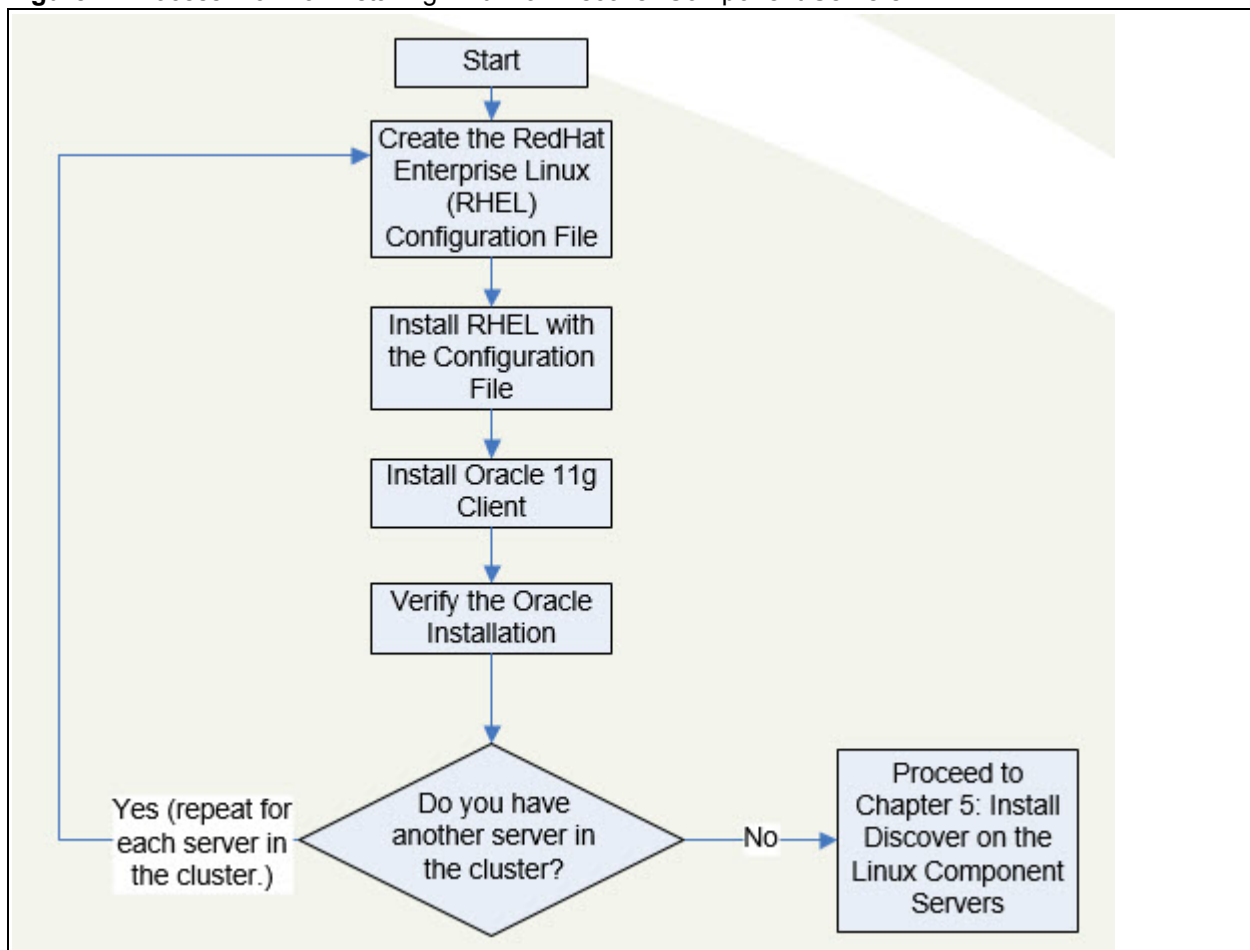
This chapter provides instructions for installing a single Linux server with the Oracle client, as well as the Discover components. Before beginning the installation, review the installation prerequisites.

Note: If Oracle and Discover are installed on the same physical box, there is no need to install the Oracle Client separately, because it is part of the Oracle Server installation.

This chapter provides the following information:

- “Installation Prerequisites”
- “Installing the Oracle Client and Verifying Oracle Server Settings”

Figure 4: Process Flow for Installing Linux for Discover Component Servers



Installation Prerequisites

First, obtain the following required files:

- RHEL 6.5 or 6.8 (64-bit)
- Oracle 11g R2 (11.2.0.1 or higher) 64-bit Standard or Enterprise Edition
 - ...**or** Oracle 12c (12.1.0.2.0) 64-bit Enterprise

If the Oracle and RHEL ISO images are from a download, use CD-writing software to burn them onto CDs.

In addition, the installation requires the following items:

- A Linux server built with certified Red Hat compatible hardware that meets the Discover minimum server hardware requirements
- Discover and Discover Linux Application Server Installer
- X-Windows capable system (for example, Cygwin on Windows).
- Putty (or other remote access app)
- WinSCP (or other file transfer app)

Installing the Oracle Client and Verifying Oracle Server Settings

All of the Discover component servers require installation of the Oracle Client libraries. Be sure these client libraries match the version of the database that Discover will be accessing. With the exception of configuring the network connection in the `tnsnames.ora` file, there are no special configuration requirements.

Installing Discover on the Linux Component Server 5

About this Chapter

The majority of Discover runs on Linux Servers. Once a Linux Server is properly configured, the Discover Administrator is ready to install the Discover application. This chapter tells you how to install and configure the Discover application on a Linux Server.

Discover Linux Installer

The Discover Linux Installer ensures that the base Linux operating system is properly configured to host a Discover installation. The installer helps automate the process so that Discover administrators can get up and running quickly, and begin scanning with the Discover application.

The installation process requires the following procedures:


- [“Copying the BDNA Discover Linux Installer to the Linux Component Server”](#)
- [“Running the Linux Installer”](#)

The basic operation of the installer is to:

- Check and set configuration of the Linux OS
- Check the configuration of the database
- Communicate with the Oracle databases

The installer is a command-line executable run on **each** Linux machine that is part of a Collection Store cluster server and the FactBase server. Command-line arguments are available to streamline installation for different circumstances. For example, the installer supports modes such as re-running, running on secondary machines in a BDNA Discover cluster, and running silently as part of a scripted installation process. To aid troubleshooting, the installer records all installation actions in the `/home/bdna/bdna_installation.log` file. The installer also returns exit codes so that it can be integrated into a scripting environment.

Note: Because the installer is a reusable file, be sure to retain it after executing it.

Figure 5: Process Flow for Installing Discover on the Linux Component Server

Before You Begin

The installer must be run by a BDNA Discover Linux user. Thus, the one pre-installation requirement is that the system administrator create a Linux user named “bdna”.

To ensure a smooth installation, review the installation requirements outlined in “[Hardware Requirements](#)” on page 15.

Note: To perform the installation, you must be logged in as the BDNA Discover Linux user, “bdna”, with the appropriate access privileges. For more information, refer to “[Oracle Access Grants Requirements](#).”.

Even though the installer needs the access privileges to accomplish some of its tasks, you cannot run the program as a “root” user. If you attempt to run the installer as “root,” the program displays an error message and exits.

If you do not have sudo privileges, you are prompted for the “root” password. If you provide the “root” password, you can continue with the installation. Refer to the description of the -r argument provided in Table 6, “Linux Installer Command-Line Arguments,” on page 51 for more detailed information.

During the installation process, the BDNA Discover Linux Installer prompts you for information about your environment. You should have the following:

- Hostname of the system that runs the Collection Store and FactBase Oracle databases. This is usually a single Linux machine
- SID of each database
- Port on which the Oracle listener is running; for example, 1521
- Oracle database user and password to run the check; you can use a DBA or non-DBA user; using a DBA user results in a more comprehensive check; i.e., a “sys” or “system” user
- At least 2GB available Oracle tablespace [SYSTEM]
- At least 20GB available Oracle tablespace [USERS]
- At least 4GB available Oracle tablespace [BDNATEMP]
- At least 4GB available Oracle tablespace [UNDOTBS]

To view a list of the Discover Linux Installer command-line options, refer to [Table 3, “Linux Installer Command-Line Arguments,”](#) on page 32. To begin the installation, log in to the Discover server as the Discover user.

Copying the BDNA Discover Linux Installer to the Linux Component Server

1. Login to your Linux system as the Linux user, “bdna”.

The BDNA Discover Linux Installer is a single executable named `bdna_x.x.x_####_yyyymmdd.installer`, where x.x.x is the product version number, #### is the product build number, yyyy is the year, mm is the month, and dd is the day.

2. Copy the *.installer to your Linux system. Remember the location.
3. See “[Running the Linux Installer](#),” for further information.

Note: \$BDNA_HOME is the base installation location for Discover. Typically, this is in the Discover user's home directory (e.g. `./`, `/home/bdna`), but it can be anywhere.

Running the Linux Installer

Use the following command to run the installer:

```
$ sh ./bdna_x.x.x_####_yyyymmdd.installer
```

Command-Line Options

To view a help message with the available command-line arguments, use the `-h` option of the installer command. For example:

```
$ sh ./bdna_x.x.x_####_yyyymmdd.installer -h
```

[Table 3, "Linux Installer Command-Line Arguments,"](#) lists the command-line arguments that are available with the installer. All arguments are optional.

Table 3: Linux Installer Command-Line Arguments

Argument	Description	Can Combine with Other Options?
-a	Automatically responds to prompts if prior answers are available from a previous installation, thereby requiring no action from the user. The question and answer are displayed on screen.	Y
-d	Skips the check for the Oracle database configuration.	Y
-f	Forces the installer to run and ignore error messages. During a typical installation, if the installer encounters an error (insufficient disk space, for example), it prompts you to continue. When this option is specified, the installer still prints any applicable warning messages. However, rather than prompting you, it ignores the warning and automatically proceeds with the installation. Upon completion, the installer returns a non-zero exit code to indicate that it completed with errors. Note that the <code>-f</code> option bypasses only the errors with a y/n prompt; it does not bypass other user inputs such as prompting for a root password, ORACLE_HOME path, etc. It also does not bypass impassable errors such as running the installer as a user who does not have root privileges.	Y
-h	Displays the help contents.	N
-i	Ignores all prompts (equivalent to <code>-frsd</code>).	N
-o	Checks only SQL*Net and database compatibility. Use this option to run a standalone check for proper Oracle database configuration and compatibility. When specified, all installation steps are skipped and the installer goes immediately to the SQL*Net dependency check, followed by the associated prompts.	N

Table 3: Linux Installer Command-Line Arguments (Continued)

-p	Uses a directory other than \$HOME as the base directory for \$BDNA_HOME. If this option is not specified, the base directory will be \$HOME.	Y
-r	Does not prompt for a root password. If you have the correct sudo privileges, the root actions of installation are performed. However, if you do not have sudo privileges, specifying this option causes the installer to skip both prompting as well as those actions that require root privileges.	Y
-s	Skips the Oracle SQL*Net client check.	Y

Note: When the installer encounters a warning during installation, the program interactively prompts you to continue (Ignore and continue (y/n)?). If you type y, the warning message is effectively ignored and installation continues; if you type n, you will need to confirm that you want to quit the installation. Both the warning message and your response are written to the log file.

Overview of Installer Operations

The following sections describe the sequence of operations that occur when you run the installer.

Database Compatibility Check

The installer performs the following database compatibility checks:

- Oracle database version is 11.x.x.x.
 - BDNA Discover is certified for 64-bit versions of Oracle 11g R2 (11.2.0.1 or higher) 64-bit Standard or Enterprise Edition or Oracle 12c (12.1.0.2.0) 64-bit Enterprise Edition

If the installer detects an incorrect version, the program displays a warning message and prompts you to continue.

Hardware Compatibility Check

The installer checks for the following compatibilities on the Linux machine. If a compatibility requirement is below the minimum, a warning message appears and you are prompted to continue.

- A minimum of 8GB of total system memory
- A minimum of 20GB of free disk space in your home directory (\$HOME).
- A minimum of 300MB of temporary disk space
- A minimum of 10GB of swap space

If the installer encounters insufficient disk space during its operation, the installer exits and a fatal message is logged.

Note: The default RHEL 6.5 creates a 4GB swap space. RHEL 6 recommends using a swap space according to this formula.

If: $M = \text{Amount of RAM in GB}$, and $S = \text{Amount of swap in GB}$, then $\text{If } M < 2 \text{ } S = M * 2 \text{ Else } S = M + 2$

Using that formula, change the swap space by following the instructions in the Linux documentation:
http://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/6/html/Storage_Administration_Guide/s1-swap-adding.html

Software Compatibility Check

The BDNA Discover installer performs compatibility checks for the operating system, client version, and database. If the requirements are not met, the installer displays a warning and prompts you to continue.

Shell

BDNA Discover requires the bash shell. If the installer encounters any shell other than bash, the program displays a warning message and prompts you to continue. The installer creates an environment variable file and sets the Discover environment variables for subsequent log-in sessions.

Environment variables must be exported using the `export` command. Whenever possible, directory paths should include variables instead of using absolute paths. For example, use `$HOME/installxxx` (where `xxx` is the product version number) instead of `/home/bdna/installxxx`. The following environment variables are needed by BDNA Discover and should be set in the BDNA Discover environment file.

BDNA_HOME

Points to the symbolic link directory specified. For example:

```
export BDNA_HOME=$HOME/installxxx (where xxx is the product version
number)
```

ORACLE_HOME

Points to the location determined by the Oracle SQL*Net client. For example:

```
export ORACLE_HOME=/u01/app/oracle/product/11.2.0
```

If the client check failed and `ORACLE_HOME` is unknown, a comment is used instead:

```
export ORACLE_HOME=/insert/oracle/home/path/here
```

LD_LIBRARY_PATH

Points to `$ORACLE_HOME/lib` and should build on existing `LD_LIBRARY_PATH`. For example: `export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:$ORACLE_HOME/lib`

If the client check failed and the `ORACLE_HOME` is unknown, a comment is used instead:

```
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:$ORACLE_HOME/lib
```

umask

The `umask` variable determines the default permissions for files that get created; this variable should be set to 0002; for example:

```
umask 0002
```

PATH

The executable search path of the OS user:

```
export PATH=$BDNA_HOME/bin:$ORACLE_HOME/bin:$PATH
```

If the `ORACLE_HOME` variable is unknown, the `PATH` line is added to the comments and an alternative `PATH` is set. For example:

```
#export PATH=$BDNA_HOME/bin:$ORACLE_HOME/bin:$PATH
export PATH=$BDNA_HOME/bin:$PATH # Warning: ORACLE_HOME missing
```

Kernel Parameters

The Discover installer can update the Linux kernel parameters to the Discover recommended values. This update is accomplished by editing the appropriate configuration files, which requires the availability of root privileges in order to use the Linux `sudo` command.

In order to operate correctly, Discover requires the following kernel parameters:

```
kernel.shmall = 2097152
kernel.shmmax = 2147483648
kernel.shmmni = 4096
kernel.sem = 250 32000 100 128
fs.file-max = 65536
net.ipv4.ip_local_port_range = 1024 65000
net.core.rmem_default = 1048576
net.core.rmem_max = 1048576
net.core.wmem_default = 262144
net.core.wmem_max = 262144
net.ipv4.neigh.default.gc_thresh3 = 65536
```

Response File

The response file stores responses to questions asked during installation. However, passwords are not saved in the file. Every time you respond to a question from the installer, the program stores the answer in the response file. All the responses get saved in `$HOME/.bdna_response_file`.

Note: There is a dot prefix to the filename (`.bdna_response_file`).

Configuring BDNA Discover for the First Time

The configuration process has the following steps:

- Copy the Enterprise Sequence
- Run `configure.sh` for the Collection Store
- Run `configure.sh` for the FactBase

Before You Begin

Before you begin configuring BDNA Discover, it is highly recommended that you become familiar with the following terminology and concepts used in the Discover product.

- **Collection Store:** A Collection Store is the schema into which a scan is performed. The results of a scan are packaged and exported from a Collection Store and are referred to as a collection. For more information, see the Working with Collection Stores in the Discover Discovery Guide.
- **Collection Store Distributed Environment (Collection Store cluster):** You can configure the Discover Collection Store on a distributed environment. Because Discover components can be off-loaded onto different servers in a Discover cluster, this can greatly improve the overall product performance in terms of network discovery. The `configure.sh` script supports a maximum of 8 machines, but custom configurations can support n machines.

During the configuration you are asked if the Collection Store is a distributed environment. If you are using several servers for the Collection Store, you will need the hostnames or IP addresses of all the Discover component servers.

For more information, see “[Distributed Collection Store Installation](#)” on page 41

- **FactBase:** The FactBase is a persistent store schema into which collections are placed. These collections are grouped into inventories that are then available for basic reporting, as well as time-based reporting. For more information, see the Discover FactBase User Guide for more information.
- **Sequence:** The data Discover uses to perform a scan, and understand the results of a scan. It contains the Technopedia, market data, fingerprints and other Discover information that lets the tool identify components, software and services when scanning. The Enterprise Sequence file is required for a complete installation.
- **Message Bus:** The message bus provides a communication mechanism for all Discover components and is the main coordinating component. For more information, see Working with Collection Stores in the Discover Discovery Guide.
- **Remote Services (RS):** Remote Services is the Discover component responsible for all interaction with interface components such as the Scan Administration and Analytics applications. For more information, see Remote Services in the Discover Administrators Manual and Reference.
- **Oracle Database Server:** Server running an Oracle database.

Configuring the Discover Collection Store and FactBase Server(s)

To configure the BDNA Discover component servers, you need to execute `configure.sh` on the Collection Store server and the FactBase server. If this is a distributed Collection Store installation, run `configure.sh` on the primary Collection Store server.

Note: Because of global variables set by the `*.installer`, you must log out of the session where you executed `*.installer` and start a new session to run the `configure.sh` script. The installer will prompt you to log out and log back in.

All of the inputs you enter in the `configure.sh` script are saved to the `$HOME/.bdna_setup_response` file. The response file lets you run `configure.sh` without needing to re-enter the same set of inputs. .

Note: Passwords are not saved in the response file.

To configure the BDNA Discover component servers:

1. In a new window, log in as the BDNA Discover user “bdna”.
2. Enter the configure.sh command (shell script for configuring the Discover installation). For example:

```
$ sh $BDNA_HOME/bin/configure.sh
```

```
Usage: sh ./configure.sh [-htwoi]
```

where:

[-t] skips the Oracle TNS configuration. (optional)

[-w] skips the BDNA Discover Web Administration pages deployment. (optional)

[-o] overrides JDBC connection string. (optional)

[-i] skips Ping test. (optional)

[-e interface] uses the specified interface. (optional)

[-h] shows help.

3. Answer the system prompts:

The following is a list of the prompts from `configure.sh` when configuring a system for Collection Store, and an explanation of each prompt. For the prompts to configure FactBase, see [“Configuring the FactBase Server”](#) on page 44.

1. Installing Discover Enterprise Sequence

Prompt	Explanation
Do you want to load the Enterprise Sequence as part of the configuration (y/n)?	<p>Enter ‘y’.</p> <p>By default, the *.installer includes and loads the Baseline Sequence, which is level 1 only. The Baseline Sequence contains the fingerprint modules that identify hardware and OS types only; it does not identify any applications.</p> <p>If you enter ‘n’, the Enterprise Sequence file will not be loaded during configuration. You can load the Enterprise Sequence after the server is configured, by re-running the configure.sh script.</p> <hr/> <p>Caution: If you rerun configure.sh, in order to preserve the credentials you set up previously, you must use the same encryption seed string when you rerun this script.</p> <hr/> <p>Note: If the Enterprise Sequence file previously has been successfully loaded, the configure.sh script skips this question and loads the Enterprise Sequence the next time the schema is initialized.</p> <hr/>
Enter the full path and file name of the Enterprise Sequence file.	The full path and file name of the Enterprise Sequence file.

2. Configuring Discover Schema

Prompt	Explanation
Will this server be running Collection Store (y/n)?	Indicates to the script if you are configuring a Collection Store server or a FactBase server. Enter 'y', to continue configuring the Collection Store server. If you enter 'n', the script configures a FactBase server. See " Configuring the FactBase Server " on page 44, for more information.
Will Collection Store be running on a distributed environment (y/n)?	Discover can run the Collection Store on a distributed environment (Discover cluster) consisting of up to 8 Linux servers. If you are using more than one server for the Collection Store, enter 'y'. For information about the prompts that appear, see " Distributed Collection Store Installation " on page 41 Otherwise enter 'n'. Note: You need to run *.installer on every Discover component server and configure.sh on the primary Collection Store server and the FactBase server.
Enter the Collection Store schema name:	The Collection Store schema name.
Enter the password for the Oracle user who will own the Collection Store schema:	The password of the Oracle user who owns the Collection Store schema.
Confirm the schema password:	Ensures the password was entered correctly.
Enter the IP address or hostname of the Collection Store database server instance	The IP address or hostname of the Collection Store database server instance.
Enter the TNS Listener port for the Oracle instance [default=1521]:	The TNS Listener port for the Oracle instance.
Enter the Oracle SID of the Collection Store database server:	The Oracle system identifier of the Collection Store database server.
Do you want to drop the schema '<schema name>' if it exists (y/n)?	If you enter 'y' and the schema name exists in the Collection Store database server, the script removes the schema, and then re-creates a new schema using the name shown.
Enter the encryption seed string for the Discover schema:	A seed string needed to create a private key, which is used for data encryption. This can be any string.
Confirm the seed string:	Ensures the seed string was entered correctly.
Enter the FactBase schema name:	The FactBase schema name.
Enter the password for the Oracle user who will own the FactBase schema:	The password for the Oracle user who will own the FactBase schema.

2. Configuring Discover Schema (Continued)

Confirm the schema password:	Ensures the password was entered correctly.
Enter the IP address or hostname for the FactBase database server instance:	The IP address or hostname of the FactBase database server.
Enter the TNS Listener port for the Oracle instance [default=1521]:	The TNS Listener port for the Oracle instance.
Enter the Oracle SID of the FactBase database server:	The Oracle system identifier of the FactBase database server.
Enter the hostname or IP address of the Discover server running FactBase:	The IP address or hostname of the Discover FactBase server. This is not the FactBase database server.

3. Configuring Discover Schema

Prompt	Explanation
Will this server be running Collection Store (y/n)?	Indicates to the script if you are configuring a Collection Store server or a FactBase server. Enter 'y', to continue configuring the Collection Store server. If you enter 'n', the script configures a FactBase server. See “Configuring the FactBase Server” on page 44, for more information.
Will Collection Store be running on a distributed environment (y/n)?	Discover can run the Collection Store on a distributed environment (Discover cluster) consisting of up to 8 Linux servers. If you are using more than one server for the Collection Store, enter 'y'. For information about the prompts that appear, see “Distributed Collection Store Installation” on page 41. Otherwise enter 'n'. Note: You need to run *.installer on every Discover component server and configure.sh on the primary Collection Store server and the FactBase server.
Enter the Collection Store schema name:	The Collection Store schema name.
Enter the password for the Oracle user who will own the Collection Store schema:	The password of the Oracle user who owns the Collection Store schema.
Confirm the schema password:	Ensures the password was entered correctly.
Is Collection Store database server stand-alone (enter N if it is clustered)?	Enter 'y' for stand-alone; 'n' for clustered.
Enter the IP address or hostname of the Collection Store database server instance	The IP address or hostname of the Collection Store database server instance.

3. Configuring Discover Schema (Continued)

Enter the TNS Listener port for the Oracle instance [default=1521]:	The TNS Listener port for the Oracle instance.
Does your listener use a SID (enter N if it is a SERVICE_NAME)?	Enter 'y' if it uses a SID; 'n' if it uses a service name.
Enter the Oracle SID of the Collection Store database server:	The Oracle system identifier of the Collection Store database server.
Would you like Setup to create a TNS entry in tnsnames.ora (y/n)?	Enter 'n' here?
Enter the TNS name to connect to the Oracle instance []:	(Enter TNS name).
Do you want to drop the schema '<schema name>' if it exists (y/n)?	If you enter 'y' and the schema name exists in the Collection Store database server, the script removes the schema, and then re-creates a new schema using the name shown.
Enter the encryption seed string for the Discover schema:	A seed string needed to create a private key, which is used for data encryption. This can be any string. Note: You should record this seed string PERMANENTLY, as you will need it for future upgrades.
Confirm the seed string:	Ensures the seed string was entered correctly.
Enter the FactBase schema name:	The FactBase schema name.
Enter the password for the Oracle user who will own the FactBase schema:	The password for the Oracle user who will own the FactBase schema.
Confirm the schema password:	Ensures the password was entered correctly.
Enter the IP address or hostname for the FactBase database server instance:	The IP address or hostname of the FactBase database server.
Enter the TNS Listener port for the Oracle instance [default=1521]:	The TNS Listener port for the Oracle instance.
Enter the Oracle SID of the FactBase database server:	The Oracle system identifier of the FactBase database server.
Enter the hostname or IP address of the Discover server running FactBase:	The IP address or hostname of the Discover FactBase server. This is not the FactBase database server.

4. Creating Discover Schema

Prompt	Explanation
Create <schema_name> schema (y/n)?	If you enter 'y', creates a schema by the name you specified. Note: Note: If you enter 'y', you will need to enter the SYSTEM user or equivalent and password. If you do not have this user, enter 'n'.
Initialize schema <schema_name> (y/n)?:	Lets you choose if whether you want to initialize the schema.
Enter the SYSTEM user for the Oracle instance server [default=system]:	The SYSTEM user for the Oracle instance server.
Enter the SYSTEM user password for the Oracle instance server [default=system]:	The password for the SYSTEM user for the Oracle instance server.

5. Setting Discover Administrator Passwords

Enter the Discover Administrator user password:	The password for the Discover user root.
Confirm the Discover Administrator user password:	Ensures the password was entered correctly.
Enter the Discover Analyst user password:	The password for the Discover user "Analyst".
Confirm the Discover Analyst user password:	Ensures the password was entered correctly.

Distributed Collection Store Installation

You can configure the Discover Collection Store on a distributed environment. Because Discover components can be off-loaded onto different servers in a Discover cluster, this can greatly improve the overall product performance. The Collection Store cluster can be distributed on a maximum of 8 machines when using the `config.sh` script. The cluster can be distributed on n number of machines when using a custom configuration.

During configuration of Collection Store, if you answer 'y' to the following question, you are asked to identify the additional machines.

```
Will Collection Store be running on a distributed environment (y/n)?
```

You will also need to run `*.installer` on each of the machines you specify as being in the Discover cluster and `configure.sh` on the primary Collection Store server and the FactBase server.

The following questions appear when configuring a distributed Collection Store.

Table 4: Distributed Collection Store Configuration Prompts

Prompt	Explanation
Enter the total number of Collection Store servers in the cluster (2-8):	The total number of Linux systems in the Collection Store cluster.
Enter the IP address or hostname of Collection Store Server 2:	The IP address or hostname of a secondary Linux server in the cluster. In this example, this is the 2nd Linux server in the cluster.
Enter the OS username managing the BDNA Discover installation of the above server [bdna]:	The OS user name used to manage the Linux servers across the Discover cluster. Typically this is the same OS account that is being used to install BDNA Discover on the primary BDNA Discover server in the Discover cluster.
Enter the password for the above OS username:	The password of the OS user.
Confirm the password of the OS user:	Ensures the password was entered correctly.
Install the SSH key on all machines in the cluster (y/n)?	In order for BDNA Discover to manage the Linux servers in a BDNA Discover component servers, the SSH key must be installed on each Linux server, so that the secondary BDNA Discover servers can be accessed by the primary BDNA Discover server without requiring a OS password. If you enter 'n', it is assumed that you have already installed the SSH key on all BDNA Discover component servers prior running the <code>configure.sh</code> script. If SSH key has not been installed, the <code>configure.sh</code> will fail.
Add component layout to each machine in the cluster (y/n)?	If you enter 'y', the BDNA Discover configuration script adds the component layouts to each machine in the BDNA Discover cluster based on the number of servers being used. See “ Distributed Component Layouts ” on page 42 for more detailed information. If you want to use your own component layout, enter 'n'. This means you will need to use the BDNA Discover shell commands such as <code>mkcomp</code> to add the components.

Distributed Component Layouts

The following tables show the pre-defined component layouts that are added to the Discover cluster by the `.installer` script based on the number of machines in the cluster. The tables show information for:

- [Table 5, “Component Layouts,”](#) on page 43
- [Table 6, “Property Settings,”](#) on page 44
- [Table 7, “WinCS Distribution,”](#) on page 44

Table 5: Component Layouts

Machine	Number of Servers in Collection Store Cluster						
	2	3	4	5	6	7	8
BDNA Discover Server 1	Agenda Manager (2.0G) RS (1.0 G) RULE0 (3.0G)	Agenda Manager (2.0G) RS (1.0 G) RULE0 (3.0G)	Agenda Manager (2.0G) RS (1.0G) RULE0 (3.0G)	Agenda Manager (2.0G) RS (1.0G) RULE0 3.0G)	Agenda Manager (2.0G) RS (1.0G)	Agenda Manager (2.0G) RS (1.0G)	Agenda Manager (2.0G) RS (1.0G)
BDNA Discover Server 2	RULE1 (4.0G) RULE2 (4.0G)	PerlCS1 (2.0G)	PerlCS1 (2.0G)	RULE0 (4.0G) bdnaPerlCS bdnaCLE CLE2 (2.0G)	RULE0 (4.0G) bdnaPerlCS bdnaCLE CLE2 (2.0G)	RULE0 (4.0G) bdnaPerlCS bdnaCLE CLE2 (2.0G)	RULE0 (4.0G) bdnaPerlCS bdnaCLE CLE2 (2.0G)
BDNA Discover Server 3	N/A	RULE1 (4.0G) RULE2 (4.0G)	PerlCS2 (2.0G)	RULE1 (4.0G) RULE2 (4.0G)	RULE1 (4.0G) RULE2 (4.0G)	RULE1 (4.0G) RULE2 (4.0G)	RULE1 (4.0G) RULE2 (4.0G)
BDNA Discover Server 4	N/A	N/A	RULE1 (4.0G) RULE2 (4.0G)	PerlCS1 (2.0G)	RULE3 (4.0G) RULE4 (4.0G)	RULE3 (4.0G) RULE4 (4.0G)	RULE3 (4.0G) RULE4 (4.0G)
BDNA Discover Server 5	N/A	N/A	N/A	PerlCS2 (2.0G)	PerlCS1 (2.0G)	RULE5 (4.0G) RULE6 (4.0G)	RULE5 (4.0G) RULE6 (4.0G)
BDNA Discover Server 6	N/A	N/A	N/A	N/A	PerlCS2 (2.0G)	PerlCS1 (2.0G)	RULE7 (4.0G) RULE8 (4.0G)
BDNA Discover Server 7	N/A	N/A	N/A	N/A	N/A	PerlCS2 (2.0G)	PerlCS1 (2.0G)
BDNA Discover Server 8	N/A	N/A	N/A	N/A	N/A	N/A	PerlCS2 (2.0G)

Table 6: Property Settings

Number of Servers in Discover Cluster							
Properties	2	3	4	5	6	7	8
CLE1:bdna.cle.claTaskBlockSize	1	1	1	1	1	1	1
CLE1:maxConcurrentTasks	650	650	650	300	300	300	300
CLE1:bdna.cle.taskBlockSize	150	150	150	150	150	150	150
PerlCS1:bdna.perlcs.maxTasks	100	250	250	250	250	250	250
PerlCS1:bdna.perlcs.maxConnections	100	250	250	250	250	250	250
PerlCS1:bdna.perlcs.cleName	CLE1	CLE1	CLE1	CLE2	CLE2	CLE2	CLE2
bdna.agenda.localHostMaxRsrc	400	400	400	500	500	500	500
CLE2:bdna.cle.claTaskBlockSize	N/A	N/A	N/A	500	500	500	500
CLE2:maxConcurrentTasks	N/A	N/A	N/A	250	250	250	250
CLE2:bdna.cle.taskBlockSize	N/A	N/A	N/A	1	1	1	1
PerlCS2:bdna.perlcs.maxTasks	N/A	N/A	250	250	250	250	250
PerlCS2:bdna.perlcs.maxConnections	N/A	N/A	250	250	250	250	250
PerlCS2:bdna.perlcs.cleName	N/A	N/A	CLE1	CLE2	CLE2	CLE2	CLE2

Table 7: WinCS Distribution

Number of Servers in Discover Cluster							
WinCS Instance	2	3	4	5	6	7	8
WinCS Instance 1	Connect to CLE1	Connect to CLE1	Connect to CLE1	Connect to CLE1	Connect to CLE1	Connect to CLE1	Connect to CLE1
WinCS Instance 2	N/A	N/A	Connect to CLE1	Connect to CLE1	Connect to CLE1	Connect to CLE1	Connect to CLE1
WinCS Instance 3	N/A	N/A	N/A	Connect to CLE1	Connect to CLE1	Connect to CLE1	Connect to CLE1
WinCS Instance 4	N/A	N/A	N/A	Connect to CLE1	Connect to CLE1	Connect to CLE1	Connect to CLE1

Configuring the FactBase Server

To configure the server to be used for running FactBase, execute the following command.

In a new window, enter the following command and answer the prompts:

```
% sh $BDNA_HOME/bin/configure.sh
```

The following is a list of the prompts from `configure.sh` when configuring a system for FactBase, and an explanation of each prompt.

1. Installing Discover Enterprise Sequence

Prompt	Explanation
Do you want to load the Enterprise Sequence file (y/n)?	<p>Enter 'y'.</p> <p>By default, the <code>*.installer</code> includes and loads a Baseline Sequence, which contains the fingerprint modules that identify hardware and OS types only; it does not identify any applications).</p> <p>If you enter 'n', the Enterprise Sequence file will not be loaded during configuration. You can load the Enterprise Sequence after the server is configured, by re-running the <code>configure.sh</code> script.</p> <hr/> <p>Caution: If you rerun <code>configure.sh</code>, in order to preserve the credentials you set up previously, you must use the same encryption seed string when you rerun this script.</p> <hr/> <p>Note: If the Enterprise Sequence file previously has been successfully loaded, the <code>configure.sh</code> script skips this question and loads the Enterprise Sequence the next time the schema gets initialized.</p> <hr/>
Enter the full path and filename of the Enterprise Sequence:	The full path and file name of the Enterprise Sequence file.

2. Configuring Discover Schema

Prompt	Explanation
Will this server be running Collection Store (y/n)?	Enter 'n' to continue the FactBase configuration.
Enter the FactBase schema name:	The FactBase schema name.
Enter the password for the Oracle user who will own the FactBase schema:	The password for the Oracle user who will own the FactBase schema.
Confirm the schema password:	Ensures the password was entered correctly.
Enter the IP address or hostname for the FactBase database server instance:	The IP address or hostname of the FactBase database server.
Enter the TNS Listener port for the Oracle instance [default=1521]:	The TNS Listener port for the Oracle instance.
Enter the Oracle SID of the FactBase database server:	The Oracle system identifier of the FactBase database server.
Do you want to drop the schema <schema name> if it exists (y/n)?	If you enter 'y' and the schema name exists in the Collection Store database server, the script removes the schema, and then re-creates a new schema using the name shown.

2. Configuring Discover Schema (Continued)

Enter the encryption seed string for the Discover schema:	A seed string needed to create a private key, which is used for data encryption. This can be any string.
Confirm the seed string:	Ensures the seed string was entered correctly.
Enter the Collection Store schema name:	The name of the schema to use for the Collection Store database server.
Enter the password for the Oracle user who will own the Collection Store schema:	The schema password for the Collection Store database server.
Confirm the schema password:	Ensures the password is entered correctly.
Enter the IP address or hostname of the Collection Store database server instance:	The IP address or hostname of the Collection Store database server instance.
Enter the TNS Listener port for the Oracle instance [default=1521]:	The TNS Listener port for the Oracle instance.
Enter the Oracle SID of the Collection Store database server [default=ora12c]:	The Oracle system identifier of the Collection Store database server.
Enter the message bus port for FactBase instance [default=8004]:	The message bus port of the FactBase component server.
Enter the RS port of the FactBase instance: [default=8081]:	The Remote Service port of the FactBase component server.
Enter the hostname or IP address of the primary Discover Collection Store server:	The IP address or hostname of the Collection Store server. Note: This is not the Collection Store database server. If you have a Discovery cluster, this is the IP address or hostname of the primary Collection Store server.

3. Creating Discover Schema

Prompt	Explanation
Create <schema_name> schema (y/n)?	If you enter 'y', creates a schema by the name you specified.
Initialize schema <schema_name> (y/n)?	Lets you decide whether you want to initialize the schema.
Enter the SYSTEM user for the Oracle instance server [default=system]:	The SYSTEM user for the Oracle instance server.
Enter the SYSTEM user password for the Oracle instance server [default=system]:	The password for the SYSTEM user for the Oracle instance server.

4. Setting Discover Administrator Passwords

Prompt	Explanation
Enter the Discover Administrator user password:	The password for the Discover user “Administrator”.
Confirm the Discover Administrator user password:	Ensures the password was entered correctly.
Enter the Discover Analyst user password:	The password for the Discover user “Analyst”.
Confirm the Discover Analyst user password:	Ensures the password was entered correctly.

About this Chapter

This chapter reviews the proper configuration of Windows Collection Server hardware, and the installation of the Windows Collection Service software.

Installing and Configuring Windows Collection Server

We recommend running only the BDNA Discover Collection Software and the required corporate security products, such as anti-virus scanners, on the Collection server.

Note: Refer to the *BDNA Discover 7.7.2 Release Notes* for updated information regarding supported versions. BDNA Discover Windows Collection Servers require Windows Server 2008 R2 or Windows Server 2012 R2. BDNA recommends a server with a minimum of 8GB RAM for use as a Windows Collection Server.

To install and configure WinCS:

1. Copy the BDNA Discover Windows Collection Service .exe file from the BDNA Support portal to the Windows Collection Server.

The Windows Server Components zip file is typically named:

`BDNA_Windows_Collection_Service_xxx_yyyy.exe`

where `xxx` is the product version and `yyyy` is the build number.

2. Run `setup.exe` in the `C:\tmp\winCS` directory.
3. Follow the on-screen Wizard instructions for accepting the License, choosing components to install, choosing the installation location, and specifying the CLE IP address.

Note: If the CLE IP address is known, enter it now. If not, provide any address; the value can be changed later.

4. Click Install.

An Installing progress window opens.

5. Click Finish when the installation is complete.

6. Stop the Discover Windows Collection Service.

After completion of the installation, the Discover Windows Collection Service will be running. Use Start > Programs > Administrative Tools > Services to locate and stop the Discover Windows Collection Service. You can restart it in the same manner.

The Discover Windows Collection Service is now installed.

Note: You need to run the `$BDNA_HOME/bin/mountWinCSLog.sh` script on the primary BDNA Discover server in the Collection Store cluster to enable archiving WinCS logs for Host-Based Diagnostics feature.

Table 8: WinCS Configuration Parameters

Configuration Item	Description
<code><add key="instanceName" value="WinCS1"/></code>	Database instance name.
<code><add key="cleHostName" value="PLACE_YOUR_CLE_NAME_HERE"/></code>	Host name or IP address (v4 or v6) of CLE to connect to. Only used when running standalone without the Message Bus. This is a Dynamic property. *Current Session Discarded*
<code><add key="clePort" value="8027"/></code>	Port CLE is listening on. Only used when running standalone without the Message Bus. This is a Dynamic property. *Current Session Discarded*
<code><add key="cleInstanceName" value="CLE1"/></code>	Instance name of CLE to connect to. This is a Dynamic property. *Current Session Discarded*
<code><add key="ttlInterval" value="600"/></code>	Time-To-Live interval in seconds. This value is reported to the CLE to indicate how often we expect the CLE to send us keep alive messages. This is a Static property.
<code><add key="clePollInterval" value="30"/></code>	CLE connection retry interval in seconds. This is a Dynamic property.
<code><add key="maxConcurrentTasks" value="75"/></code>	Maximum number of concurrent tasks to process. This is a Dynamic property *Current Session Discarded*
<code><add key="minWorkerThreads" value="0"/></code>	Minimum number thread pool worker threads. Set to 0 to use system default. This is a Dynamic property.
<code><add key="maxWorkerThreads" value="400"/></code>	Maximum number of thread pool worker threads. Set to 0 to use system default. This is a Dynamic property.
<code><add key="minIOCompletionThreads" value="0"/></code>	Minimum number of thread pool asynchronous IO threads. Set to 0 to use system default. This is a Dynamic property.
<code><add key="maxIOCompletionThreads" value="0"/></code>	Maximum number of thread pool asynchronous IO threads. Set to 0 to use system default. This is a Dynamic property.

Table 8: WinCS Configuration Parameters

Configuration Item	Description
<add key="scriptIsolation" value="2"/>	<p>Script isolation level. Controls the type of sandbox used to execute collection scripts in. efined values:</p> <p>0 - No isolation. Not currently supported and likely to result in unpredictable behavior.</p> <p>1 - AppDomain isolation. Scripts will execute in a separate AppDomain loaded into the same process as WinCS. Requires one AppDomain for each concurrent task.</p> <p>NOTE: Because of a known memory leak in WMI running long term in this mode will likely fail when the WinCS process reaches about 1.8Gb on a 32 bit machine.</p> <p>2 - Process isolation (default). Scripts will execute in a separate process from WinCS.Requires one process for each concurrent task.</p> <p>This is a Static property.</p>
<add key="logClaMessages" value="false"/>	<p>Trace all messages exchanged between WinCS and the CLE.</p> <p>This is a Dynamic property.</p>
<add key="logClaBufferProcessing" value="false"/>	<p>Trace buffer management for socket connected to the CLE.</p> <p>This is a Dynamic property.</p>
<add key="logTaskExecutorEvents" value="false"/>	<p>Trace event dispatching for the collection task execution FSM.</p> <p>This is a Dynamic property.</p>
<add key="disableTtl" value="false"/>	<p>Disable TTL enforcement. Prevents dropping the connection to the CLE when keep alive messages are not received in a timely fashion.</p> <p>This is a Dynamic property.</p>
<add key="ttlTimerInterval" value="300"/>	<p>How often (in seconds) WinCS checks to insure expected keep alive message have actually been received.</p> <p>This is a Dynamic property.</p>
<add key="messageDispatchTimeout" value="240"/>	<p>Amount of time in seconds a message from the CLE will wait to be dispatched to the appropriate message handler. f the timeout expires the message will be returned to the CLE with an RC_PROCESSING_EXCEPTION result code.</p> <p>This is a Dynamic property.</p>
<add key="processStartupTimeout" value="120"/>	<p>Amount of time in seconds WinCS will wait for a separate process sandbox to load and initialize. The value is only used in conjunction with scriptIsolation 2.</p> <p>This is a Dynamic property.</p>

Table 8: WinCS Configuration Parameters

Configuration Item	Description
<add key="taskCancelTimeout" value="120"/>	<p>The maximum time in seconds to wait for a cancelled task to exit. If the timeout expires the sandbox hosting the task being canceled will be discarded.</p> <p>This is a Dynamic property.</p>
<add key="scriptReferences" value="system.dll mscorlib.dll system.management.dll system.data.dll system.xml.dll System.Web.Services.dll"/>	<p>List of static references needed to compile connection/collection scripts. In general this list is only expected to contain references to Assemblies stored in the GAC. This list is combined with all Assemblies found in the WinCS installation Shared directory.</p> <p>This is a Dynamic property.</p>
<add key="tftpMessageReceiveTimeout" value="300"/>	<p>The maximum amount of time to wait (in seconds) for an expected TFTP message to be received. If this timeout expires the TFTP transfer is canceled.</p> <p>This is a Dynamic property.</p>
<add key="tftpReceiveTimerInterval" value="10"/>	<p>How often to check (in seconds) to see if the TFTP message received timer has expired. This value should always be less than tftpMessageReceiveTimeout.</p> <p>This is a Dynamic property.</p>
<add key="tftpUNCPath" value=""/> <add key="tftpPath_login" value=""/> <add key="tftpPath_password" value=""/>	<p>The flag of tftpUNCPath is only used if UNC Path is not empty.</p> <p>If tftpUNCPath is provided, any fingerprint that requires tftp will use the tftp executable from this UNC path.</p> <p>Access privilege is controlled by tftpPath_login and password.</p> <p>If tftpUNCPath option is used, a valid username and password must be provide.</p> <p>e.g.:</p> <pre><add key="tftpUNCPath" value="\\10.10.13.61\unc"/> <add key="tftpPath_login" value="normalizevm\administrator"/> <add key="tftpPath_password" value="Simple.0"/></pre> <p>This is a Static property.</p>
<add key="wmiMethodTimeout" value="30"/>	<p>Timeout in seconds for WMI queries.</p> <p>This is a Dynamic property.</p>
<add key="wmiConnectTimeout" value="120"/>	<p>Timeout in seconds for WMI connect operations.</p> <p>This is a Dynamic property.</p>
<add key="wmiBlockSize" value="100"/>	<p>Number of records per block fetched by WMI queries.</p> <p>This is a Dynamic property.</p>

Table 8: WinCS Configuration Parameters

Configuration Item	Description
<code><add key="leaveStdioTmpFilesOnRemoteHost" value="false"/></code>	<p>Debug switch. When set to true temporary files generated on remote hosts by the Remote Process library will not be deleted. Also files transferred to WinCS via TFTP will be written to the tmp directory.</p> <p>This is a Dynamic property.</p>
<code><add key="appDomainCleanupInterval" value="30"/></code>	<p>Number of seconds between checks to see there are AppDomains that need to be unloaded. When singleUseTaskContexts is true this value should be very small (1-2 seconds). Otherwise the default value of 30 should be adequate.</p> <p>This is a Dynamic property.</p>
<code><add key="logRegistryCalls" value="false"/></code>	<p>Debug switch. Set to true to enable registry call journaling. A log entry will be generated indicating what registry call was made and how many values were returned. Be careful, this can result in huge log entries of tens of megabytes!</p> <p>This is a Dynamic property.</p>
<code><add key="enableScriptDebugging" value="false"/></code>	<p>Debug switch. When set to true debug information will be compiled into each downloaded script and all of the compiler generated temporary files will be saved in the tmp directory.</p> <p>This is a Dynamic property.</p>
<code><add key="singleUseTaskContexts" value="false"/></code>	<p>When set to true each task context (script sandbox) will be used once and then discarded. Normally task contexts are re-used unless a script hosted in the context fails to exit properly.</p> <p>This is a Dynamic property.</p>
<code><add key="maximumProcessMemory" value="0"/></code>	<p>Maximum amount of memory in megabytes that a sandbox process is allowed to use. This value is checked against current Private Bytes value of a sandbox after each task execution. When the value is exceeded the sandbox will be discarded. Set this value to 0 (reccomended) to let WinCS compute this value based on the amount of physical memory installed and the configured maximum concurrent tasks (maxConcurrentTasks).</p> <p>This is a Dynamic property.</p>
<code><add key="sandboxBreakOnStartup" value="false"/></code>	<p>Causes the separate process sandbox to break into the debugger immediately on startup.</p> <p>This is a Dynamic property.</p>

Table 8: WinCS Configuration Parameters

Configuration Item	Description
<code><add key="tftpRetryCount" value="2"/></code>	Number of times to retry failed tftp file transfers. This retry will only apply to transfers that actually startup and then fail. It will not cause multiple attempts to spawn the tftp client on the remote host. This is a Dynamic property.
<code><add key="enableProfileCleanup" value="false"/></code>	Flag to enable profile cleanup in scripts that support this feature. This is a Dynamic property.
<code><add key="profileDirectoryDeleteRetryCount" value="30"/></code>	Number of retry attempts for profile directory deletion. Setting this value to zero will almost certainly guarantee failure as there is a race condition between when the profile becomes unlocked at the end of the collection script and when the script begins attempting to delete the directory. The first couple of delete attempts always yield a "sharing" violation. This is a Dynamic property.
<code><add key="profileDirectoryDeleteRetryInterval" value="1"/></code>	Amount of time in seconds to wait between profile directory delete retry attempts. This is a Dynamic property.

Configuring Maximum Concurrent Tasks

By default, the maximum concurrent tasks property is set to 75. If Windows Collection Service runs on a server with less than 8GB RAM, the maximum concurrent tasks property should be reduced to 25 using the following procedure.

To reconfigure the maximum concurrent tasks property:

1. Edit `C:\Program Files\BDNA\WinCS\WinCs.exe.config`.

Change the following property:

```
<add key="maxConcurrentTasks" value="75" />
```

2. Open the Services Panel and restart the Discover Collection Service.

This appendix provides troubleshooting for the following:

- “[Setting Up BDNA Discover for Access from a Browser](#)” on page 55
- “[Time Synchronization for BDNA Discover Servers](#)” on page 55

Setting Up BDNA Discover for Access from a Browser

The browsers supported in this release of BDNA Discover include the following:

- Internet Explorer 11 or higher
- Firefox 31.0 or higher
- Google Chrome 36 or higher

You can access BDNA Discover from a web browser at the URL:

```
http://myserver.domain.com:8080/bdna/AppConsole.jnlp
```

If you want to let users access the BDNA Discover applications using the short URL (<http://myserver.domain.com/>), then the template HTML files provided with the BDNA Discover installation must be transferred into the document root directory on the web server. The template web pages for setting up short URL redirection are located in the following directory under a BDNA Discover installation:

```
$ BDNA_HOME/bdna_web/templates/
```

The following command can be used to enable short URL redirection on a BDNA Discover server running httpd with a document root set to `/var/www/html/`. Before running this command, log in as the BDNA Discover OS user “bdna”, and make sure that the BDNA Discover user has sudo access to cp:

```
$ sudo cp -r $BDNA_HOME/bdna_web/templates/* /var/www/html
```

Note: The web server should be configured to recognize `index.html` as the default document to return on an unspecified request.

Time Synchronization for BDNA Discover Servers

Linux servers must be configured to interact with a valid NTP server that is either part of your intranet or, if your firewall allows outbound NTP connections, with the Internet NTP servers. Refer to the Red Hat Linux documentation for more detailed information.

Similarly, if your BDNA Discover Windows servers are not configured to join a Windows domain, you will need to configure them to interact with a valid Windows time server that is either part of your intranet or, if your firewall allows outbound Windows time service connections, with the Internet Windows time servers. Refer to the Microsoft Windows documentation for more detailed information.

The following RPM packages are required prerequisites for the installation of the Oracle Database Server, MBUS, and Component Servers.

All packages can be found on the RHEL DVD, unless otherwise specified.

RHEL 6.8 RPM Packages

- apr-1.3.9-5.el6_2.x86_64.rpm
- apr-util-1.3.9-3.el6_0.1.x86_64.rpm
- apr-util-ldap-1.3.9-3.el6_0.1.x86_64.rpm
- aspell-0.60.6-12.el6.x86_64.rpm
- aspell-en-6.0-11.el6.x86_64.rpm (located at http://rpm.pbone.net/index.php3/stat/4/idpl/15284672/dir/redhat_el_6/com/aspell-en-6.0-11.el6.x86_64.rpm.html)
- compat-libstdc++-33-3.2.3-69.el6.x86_64.rpm
- httpd-tools-2.2.15-53.el6.x86_64.rpm
- httpd-2.2.15-53.el6.x86_64.rpm
- lm_sensors-libs-3.1.1-17.el6.x86_64.rpm
- lm_sensors-3.1.1-17.el6.x86_64.rpm
- mod_ssl-2.2.15-53.el6.x86_64.rpm
- openssh-askpass-5.3p1-117.el6.x86_64.rpm
- rsh-0.17-64.el6.x86_64.rpm
- tcl-8.5.7-6.el6.x86_64.rpm
- telnet-0.17-48.el6.x86_64.rpm
- expect-5.44.1.15-5.el6_4.x86_64.rpm
- net-snmp-libs-5.5-57.el6.x86_64.rpm
- net-snmp-utils-5.5-57.el6.x86_64.rpm
- net-snmp-5.5-57.el6.x86_64.rpm
- sysstat-9.0.4-22.el6.x86_64
- glibc-devel-2.12-1.192.el6.x86_64.rpm
- libaio-0.3.107-10.el6.x86_64.rpm
- kernel-2.6.32-642.el6.x86_64.rpm
- openssh-clients-5.3p1-117.el6.x86_64.rpm
- openssh-server-5.3p1-117.el6.x86_64.rpm

- openssh-5.3p1-117.el6.x86_64.rpm
- sysstat-9.0.4-31.el6.x86_64.rpm

PERL RPM Packages

- perl-libs-5.10.1-141.el6_7.1.x86_64.rpm
- perl-5.10.1-141.el6_7.1.x86_64.rpm
- perl-URI-1.40-2.el6.noarch.rpm
- perl-XML-Parser-2.36-7.el6.x86_64.rpm
- perl-HTML-Tagset-3.20-4.el6.noarch.rpm
- perl-HTML-Parser-3.64-2.el6.x86_64.rpm
- perl-Compress-Raw-Zlib-2.021-141.el6_7.1.x86_64.rpm
- perl-IO-Compress-Zlib-2.021-141.el6_7.1.x86_64.rpm
- perl-XML-RegExp-0.03-7.el6.noarch.rpm
- perl-YAML-Syck-1.07-4.el6.x86_64.rpm
- perl-XML-DOM-1.44-7.el6.noarch.rpm
- perl-Date-Manip-6.24-1.el6.noarch.rpm

RPM Packages for Nmap 4.11

- gamin-0.1.10-9.el6.i686.rpm
- glib2-2.28.8-5.el6.i686.rpm
- glibc-2.12-1.192.el6.i686.rpm
- libgcc-4.4.7-17.el6.i686.rpm
- libselinux-2.0.94-7.el6.i686.rpm
- libstdc++-4.4.7-17.el6.i686.rpm
- nss-softokn-freebl-3.14.3-23.el6_7.i686.rpm
- zlib-1.2.3-29.el6.i686.rpm

To install the above RPM packages, run the following commands:

```
$ rpm -Uvh glib2-2.28.8-5.el6.i686.rpm glibc-2.12-1.192.el6.i686.rpm gamin-0.1.10-9.el6.i686.rpm  
libselinux-2.0.94-7.el6.i686.rpm nss-softokn-freebl-3.14.3-23.el6_7.i686.rpm zlib-1.2.3-29.el6.i686.rpm
```

Run the following commands to install the remaining RPM packages:

```
$ rpm -ivh libgcc-4.4.7-17.el6.i686.rpm
```

```
$ rpm -ivh libstdc++-4.4.7-17.el6.i686.rpm
```

Note: One RPM package has a potentially low-impact vulnerability. The affected RPM package is “nss-softokn-freebl-3.14.3-23.el6_7.i686.” See the full report at <https://access.redhat.com/security/cve/cve-2016-1938>.
