

Veritas[™] Disaster Recovery Advisor Deployment Requirements

AIX, ESX, HP-UX, Linux, Solaris, Windows
Server

6.1.1

Veritas Disaster Recovery Advisor Deployment Requirements

The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Product version: 6.1.1

Document version: 6.1.1 Rev 1

Legal Notice

Copyright © 2013 Symantec Corporation. All rights reserved.

Symantec, the Symantec Logo, Veritas and Veritas Storage Foundation are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

j-Interop: Pure Java - COM Bridge

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Symantec Corporation and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. SYMANTEC CORPORATION SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Symantec Corporation

350 Ellis Street

Mountain View, CA 94043

<http://www.symantec.com>

Technical Support

Symantec Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within Symantec to answer your questions in a timely fashion. For example, the Technical Support group works with Product Engineering and Symantec Security Response to provide alerting services and virus definition updates.

Symantec's support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

For information about Symantec's support offerings, you can visit our website at the following URL:

<http://www.symantec.com/business/support/index.jsp>

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

Contacting Technical Support

Customers with a current support agreement may access Technical Support information at the following URL:

http://www.symantec.com/business/support/contact_techsupp_static.jsp

Before contacting Technical Support, make sure you have satisfied the system requirements that are listed in your product documentation. Also, you should be at the computer on which the problem occurred, in case it is necessary to replicate the problem.

When you contact Technical Support, please have the following information available:

- Product release level
- Hardware information
- Available memory, disk space, and NIC information
- Operating system
- Version and patch level

- Network topology
- Router, gateway, and IP address information
- Problem description:
 - Error messages and log files
 - Troubleshooting that was performed before contacting Symantec
 - Recent software configuration changes and network changes

Licensing and registration

If your Symantec product requires registration or a license key, access our technical support Web page at the following URL:

<http://www.symantec.com/business/support/>

Customer service

Customer service information is available at the following URL:

<http://www.symantec.com/business/support/>

Customer Service is available to assist with non-technical questions, such as the following types of issues:

- Questions regarding product licensing or serialization
- Product registration updates, such as address or name changes
- General product information (features, language availability, local dealers)
- Latest information about product updates and upgrades
- Information about upgrade assurance and support contracts
- Information about the Symantec Buying Programs
- Advice about Symantec's technical support options
- Nontechnical presales questions
- Issues that are related to CD-ROMs or manuals

Support agreement resources

If you want to contact Symantec regarding an existing support agreement, please contact the support agreement administration team for your region as follows:

Asia-Pacific and Japan	customercare_apac@symantec.com
Europe, Middle-East, and Africa	semea@symantec.com
North America and Latin America	supportsolutions@symantec.com

Documentation

Product guides are available on the media in PDF format. Make sure that you are

using the current version of the documentation. The document version appears on page 2 of each guide. The latest product documentation is available on the Symantec website.

<https://sort.symantec.com/documents>

Your feedback on product documentation is important to us. Send suggestions for improvements and reports on errors or omissions. Include the title and document version (located on the second page), and chapter and section titles of the text on which you are reporting. Send feedback to:

doc_feedback@symantec.com

Contents

	About this document	
	Intended audience	9
Chapter 1	DRA deployment architecture	
	Deployment environment	12
	DRA server	13
	Oracle environment and licensing	14
	Web client requirements	15
	Credentials and collection methods used	15
	About setting up DRA user profiles	16
	About privilege control software	16
	Data collection from SYMCLI through a UNIX proxy	17
	Data collection from NaviCLI through a UNIX proxy	18
	Data collection from DSCLI through a UNIX proxy	21
	Data collection from XCLI through a UNIX proxy	22
	Data collection from SVC	24
	Data collection from HDS HiCommand/HP CommandView	25
	Data collection from NetApp	25
	Data collection from VMware vCenter	27
	Data collection from HMC	28
	Data collection from VIO	29
	Data collection from UNIX hosts	30
	Privileged commands on AIX	36
	Data collection from Windows	38
	Data collection from databases	40
	Data collection from Veritas Cluster Server	46
	Data collection from EMC Control Center	47
	Mail server configuration	47
	Network and other environmental recommendations	48
	Summary of ports and protocols used by DRA	49
Appendix A	Methods for secure privilege provisioning	
	sudo on UNIX hosts	51
	SymCLI proxy	51
	All UNIX flavors	52

Servers that use EMC storage	52
Servers that use HDS/HP XP storage	52
Servers that use NetApp storage	52
Servers that use IBM storage	52
Servers with Veritas SF and Cluster	52
Additional for Solaris servers	52
Additional for Linux servers	53
Additional for AIX servers	53
UNIX Privilege Manager	54
Suggested Oracle grant provisioning	56
Suggested MS SQL Server grant provisioning	57
Queries used to scan EMC ECC	66
Configuring DRA for ECC scanning over JDBC SSL (Oracle)	66
Configuring an SSL policy	67
Copying the wallet file from the ECC Repository server	69

About this document

This document summarizes Veritas Disaster Recovery Advisor (DRA) deployment requirements. It contains the following chapter and appendix:

- [Chapter 1, “DRA deployment architecture”](#) describes the deployment architecture for DRA.
- [Appendix A, “Methods for secure privilege provisioning”](#) describes suggested methods for secure privilege provisioning for the various entities supported by DRA.

Intended audience

This document is intended for the following:

- Project managers, who must understand DRA deployment requirements
- Security personnel, who need to know how DRA interacts with their environment and how it should adapt to their existing security standards
- Storage, system, and database administrators, who need to know the user account and credential settings required to support DRA


DRA deployment architecture

This chapter includes the following topics:

- [“Deployment environment”](#) on page 12
- [“DRA server”](#) on page 13
- [“Oracle environment and licensing”](#) on page 14
- [“Web client requirements”](#) on page 15
- [“Credentials and collection methods used”](#) on page 15

Deployment environment

As shown in the illustration below, you install DRA on a dedicated server:



The DRA deployment environment consists of the following:

- A Windows Server 2008 R2 64-bit application server and a dedicated Oracle 11g repository to store the collected and analyzed data (items 1 and 2 in the illustration).
- Various IT sources (3-9) that DRA collects data from for daily risk analysis.
- A web interface to use and manage DRA (10).

DRA server

DRA must run on a dedicated host, also referred to as the DRA application server. The host size depends on several parameters, including the size of your scanned environment and how much data you need to retain. The use of a virtual machine for the DRA application server should be limited to small to mid-size environments (typically, up to 100 scanned hosts). If you use VMware, you should take particular care to reserve the required CPU and memory. The recommended server configuration is as follows:

Table 1-1 DRA server requirements

# of scanned hosts	CPUs/cores	RAM	Free disk space	Operating system
Up to 100	2 Intel / AMD (4 recommended)	8 GB	80 GB	Windows Server 2008 R2 ¹ Standard or Enterprise Edition 64-bit ²
100-500	2 Intel/ AMD (4 recommended)	16 GB	100 GB	
500-1,000	4 Intel / AMD	32 GB	120 GB	
Above 1,000	Specific sizing required			

- 1:** If your organization does not plan to use Windows Remote Management (WinRM) to collect data from Windows servers, you can also use Windows Server 2003/2008 64-bit.
- 2:** Although a 32-bit operating system is not recommended, you can use one if you have a small to midsize environment. If you install a 32-bit operating system, you need at least 4 GB of RAM.

The following server requirements also apply:

- Database: Oracle 11g Standard or Enterprise installed with full database administrator rights.
- Server: Apache Tomcat 6.0 (the only supported version). If it is not on your system, it is automatically installed.
- Web client access:
 - Internet Explorer 6 or later with Java client 1.6 or later.
 - HTTP access from clients to the DRA server through port 8080 (configurable).

DRA requires administrator rights on the DRA application server.

Note: If you use VMware, you should take particular care to reserve the required CPU and memory.

The recommended collector configuration is as follows:

Table 1-2 Collector requirements

# of scanned hosts	CPUs/cores	RAM	Free disk space	Operating system
Up to 100	2 Intel / AMD	4 GB	40 GB	Windows Server 2008 R2 ¹ Standard or Enterprise Edition 64-bit
100-500	2 Intel/ AMD	8 GB	50 GB	
500-1,000	4 Intel / AMD	16 GB	60 GB	
Above 1,000	Specific sizing required			

1: If your organization does not plan to use Windows Remote Management (WinRM) to collect data from Windows servers, you can also use Windows Server 2003/2008 64-bit. DRA collectors require administrator rights on the server.

Note: If you use VMware, you should take particular care to reserve the required CPU and memory.

Oracle environment and licensing

DRA uses an Oracle 11g database to store and analyze the data collected from the scanned environment, also known as the DRA repository. Before you install DRA, you must install and configure the Oracle database.

To obtain an Oracle 11g 30-day trial license, go to:

<http://edelivery.oracle.com>

For a longer trial period or to install DRA permanently, you need an Oracle 11g standard edition license. To obtain an Oracle 11g standard edition license, contact Oracle.

DRA requires full DBA rights on the DRA repository.

No further maintenance is required for the Oracle database. DRA sets up and manages the Oracle database. It creates the required schema, handles routine database housekeeping, tuning, and so on.

Web client requirements

DRA has a web-based user interface. You need Internet Explorer 6 and later, with Java client 1.6 or later. HTTP IP access from the client to the DRA application server should be available. The default connection port is 8080, which you may change if needed.

In addition, the following must be configured:

- The address `http://dra_server:8080/DRA` must be defined as a trusted site in the Internet Explorer configuration.
- The client Internet Explorer must be configured to permit the running of JavaScripts, show applets, iFrames, and play animations.

Credentials and collection methods used

DRA mainly collects data from storage arrays, servers, and databases. For this reason, you need to set up certain dedicated user account profiles, or specify certain existing account profiles for the application's use.

Additional data that DRA collects from other logical IT elements, such as clustering software, logical volume management (LVM) software, network services, and so on, does not require further account provisioning. It can be retrieved through the operating system account profiles.

All query methods used by DRA using these account profiles have the following design principles:

- DRA collects data in read-only mode; it does not change your configuration.
- DRA only retrieves configuration data (metadata) – never actual production content. For example, DRA may read database startup files to learn how a database instance is configured. It may also connect to the database and issue system configuration queries to determine which files store database content. However, it does not query any production information, tablespace content, and so on.
- All queries use standard, well-known interfaces and commands. Nothing is hacked or retrieved in a non-standard way. In fact, all queries and commands used are well-known to the IT staff, who often use the same queries and commands during routine maintenance.
- None of the queries or commands put a noticeable load on servers, storage arrays, databases, or the network. The only significant computation is performed on the DRA application server and the DRA Oracle repository, which are dedicated computing resources.

You must enter the credential information into the DRA application, where it is kept strongly-encrypted using AES with a unique, per-customer encryption key.

DRA's flexible architecture lets it adapt to your specific customer security needs, and it complies with a wide array of security policies and doctrines. DRA has successfully adapted to the strictest security standards of many financial, government, and commercial organizations.

The following sections describe specific credential requirements and rights for each environment, and outline possible security adaptations. Note that your environment may not use all the components mentioned here. You may ignore those requirements.

About setting up DRA user profiles

When you set up a DRA user profile, keep in mind the following:

- All user account profiles provisioned for the use of DRA must have a password that does not need resetting after the first use.
- It is strongly recommended that you provision user profiles with non-expiring passwords. If that is not possible, allow the longest possible password expiration period. Symantec recommends at least six months. DRA uses the provisioned account profiles noninteractively, and the default connection method does not involve any plain-text password exchange. Therefore, these account profiles pose significantly lower risk than standard ones. Replacing the passwords on all hosts on the environment presents an administrative overhead that should be balanced against this low risk.
Finally, as long as expired passwords are not reset, DRA cannot detect risks in the environment. This should also be considered in favor of using non-expiring passwords, or ones with a long expiration period.
- It is strongly recommended to use the same user ID on many of the hosts and databases. The best practice is that you use the user ID drauser for all operating system and database account profiles.
- The user default shell should be sh.

About privilege control software

DRA mainly uses non-privileged queries and commands that do not require any administrative rights. There is a small number of read-only queries and commands that do require root privileges on UNIX. For these, Symantec recommends using privilege control software, such as sudo, PowerBroker (pbrun), UPM (pmrun), CA Access Control (seSUDO), super, and others.

For sudo, the suggested syntax for each UNIX platform is described in [Appendix A, “Methods for secure privilege provisioning”](#) on page 51. You can adapt this syntax to any other similar privilege control software.

Important: Configure sudo, PowerBroker (pbrun), UPM (pmrun), and similar privilege control software so a password is not required when executing privileged commands.

Data collection from SYMCLI through a UNIX proxy

DRA uses the standard SYMAPI interface and read-only SYMCLI commands to collect additional data from EMC Symmetrix arrays. The commands are run on one or more UNIX servers in the IT environment. Collectively, these servers can query all Symmetrix arrays in the scope. These servers are also known as SYMCLI proxies.

When you select SYMCLI proxies, use the following best practices:

- Use IT administrative servers rather than production servers.
- Use fewer proxies. For example, if one host can query all arrays on all sites, use it as a proxy rather than using two or more hosts.
- Use proxies that are consistently up and available, rather than ones that are sometimes down or unreachable.

As a standard, DRA opens a secure shell (SSH) session to the proxy, as it does to collect data from any UNIX host. Similarly, it requires sudo, PowerBroker (pbrun), UPM (pmrun), CA Access Control (seSUDO), super, or similar software to run privileged commands.

For more information, see [“Data collection from UNIX hosts”](#) on page 30.

DRA uses the following privileged, read-only SYMCLI commands:

- `/usr/symcli/bin/symcfg list`
- `/usr/symcli/bin/symdev list`
- `/usr/symcli/bin/symdisk list`
- `/usr/symcli/bin/symaudit list`
- `/usr/symcli/bin/symcli -def`

When you work with proxies, you should do the following:

- Provide each SYMCLI proxy name or IP address.
- Provide a user account profile on each SYMCLI proxy (existing or specifically created for DRA). This is the same as any UNIX host from the same vendor. For more information, see [“Data collection from UNIX hosts”](#) on page 30.
- If you prefer, provide sudo, PowerBroker (pbrun), UPM (pmrun), or similar definitions on each SYMCLI proxy. This is the same as any UNIX host from the same vendor. For more information, see [“Data collection from UNIX hosts”](#) on page 30.
- Make sure that IP connectivity through SSH is available between the DRA application server and each SYMCLI proxy. The default port is 22.

Important: Do not configure the same Symmetrix array to be scanned by more than one probe, because it may cause unpredictable results.

Note: By default, DRA connects to the proxies using SSH with user/password authentication. SSH with public key authentication is also supported. If you prefer, you can use Telnet; however, it is considered less secure than SSH. In terms of security provisioning, the only difference in using Telnet is that the default port is port 23, instead of the SSH port.

For suggestions on appropriate sudo definitions, see [Appendix A, “Methods for secure privilege provisioning”](#) on page 51.

You can adapt these suggestions to any other similar privilege control mechanism, such as PowerBroker.

Data collection from NaviCLI through a UNIX proxy

DRA uses read-only NaviSECCLI commands to collect additional data from EMC CLARiiON arrays. These commands run on one or more UNIX servers in the IT environment. Collectively, these servers can query all CLARiiON arrays in the scope. These servers are also known as NaviCLI proxies.

When you select NaviCLI proxies, use the following best practices:

- Use IT administrative servers rather than production servers.
- Use fewer proxies. For example, if one host can query all arrays on all sites, use it as a proxy rather than using two or more hosts.

- Use proxies that are consistently up and available, rather than ones that are sometimes down or unreachable.
- Use proxies that can access both Storage Processors (SPs) on CLARiiON arrays.

As a standard, DRA opens an SSH session to the proxy in the same way that it does to collect data from any UNIX host.

The connection from the proxy to the CLARiiON array requires authorization details (user/password/scope) for the array or for the host and user, which have already declared automatic (read-only) authorization for the array.

You can use sudo (or similar) software to achieve already declared authorization for the array. For more information, see “[Data collection from UNIX hosts](#)” on page 30.

DRA uses the following privileged, read-only NaviCLI commands. For AIX, the path is /usr/lpp/NAVICLI/.

- `/opt/Navisphere/bin/naviseccli -np -port port -v authorization -h array IP/name getall`
- `/opt/Navisphere/bin/naviseccli -np -port port -v authorization -h array IP/name getlun`
- `/opt/Navisphere/bin/naviseccli -np -port port authorization -h array IP/name metalun -list`
- `/opt/Navisphere/bin/naviseccli -np -port port authorization -h array IP/name mirror -async -listgroups`
- `/opt/Navisphere/bin/naviseccli -np -port port authorization -h array IP/name mirror -async -list`
- `/opt/Navisphere/bin/naviseccli -np -port port authorization -h array IP/name mirror -sync -listgroups`
- `/opt/Navisphere/bin/naviseccli -np -port port authorization -h array IP/name mirror -sync -list`
- `/opt/Navisphere/bin/naviseccli -np -port port -v authorization -h array IP/name snapview -listclonefeature`
- `/opt/Navisphere/bin/naviseccli -np -port port -v authorization -h array IP/name snapview -listclonegroup`
- `/opt/Navisphere/bin/naviseccli -np -port port -v authorization -h array IP/name getlog`
- `/opt/Navisphere/bin/naviseccli -np -port port -v authorization -h array IP/name sancopy -settings -list`

Where:

port	The port used for CLARiiON access. The default is 443.
authorization	Empty for already declared automatic authorization. Otherwise, use the following format: -User <i>user</i> -Password <i>password</i> -Scope <i>scope</i> Where: ■ User is the user name to be used for CLARiiON authorization. ■ Password is the password to be used for CLARiiON authorization. Avoid using the '<' and '>' characters in the password. ■ Scope is the scope to be used for CLARiiON authorization, represented as a numeric value (0: Global, 1: Local, and 2: LDAP).
array IP/name	The array DNS name or IP address of one of the CLARiiON storage processors (SPs).

You should provide the following information for each NaviCLI proxy:

- Name or IP address
- A user account profile (existing or specifically created for DRA)

You should provide the following information for each CLARiiON array:

- A user account profile (existing or specifically created for DRA). A profile with an empty password indicates that already declared automatic authorization is in use.

You should also verify that:

- IP connectivity through SSH (the default is port 22) is available between the DRA application server and each NaviCLI proxy.
- IP connectivity is available between the NaviCLI proxy and each CLARiiON array that it scans.

Important: Do not configure the same CLARiiON array to be scanned by more than one probe (even if the storage processors are different). This configuration may cause unpredictable results.

Note: By default, DRA connects to the proxies using SSH with user/password authentication. SSH with public key authentication is also supported. If you prefer, you can use Telnet; however, it is considered less secure than SSH. In terms of security provisioning, the only difference in using Telnet is that the default port is port 23, instead of the SSH port.

For suggestions regarding appropriate sudo definitions, see [Appendix A, “Methods for secure privilege provisioning”](#) on page 51.

You can adapt these suggestions to any other similar privilege control mechanism, such as PowerBroker.

Data collection from DSCLI through a UNIX proxy

DRA uses read-only DSCLI commands to collect additional data from IBM DS arrays. These commands run on one or more UNIX servers in the IT environment. Collectively, these servers can query all DS arrays in the scope. These servers are also known as DSCLI proxies.

When you select DSCLI proxies, use the following best practices:

- Use IT administrative servers rather than production servers.
- Use fewer proxies. For example, if one host can query all arrays on all sites, use it as a proxy rather than using two or more hosts.
- Use proxies that are consistently up and available, rather than ones that are sometimes down or unreachable.

As a standard, DRA opens an SSH session to the proxy in the same way that it does to collect data from any UNIX host.

The connection from the proxy to the IBM DS array requires authorization details (user/password) for the array.

DRA uses the following read-only DSCLI commands.

- `dscli -hmc1 array IP/name -user user -passwd passwd lssu`
- `dscli -hmc1 array IP/name -user user -passwd passwd lssi`
- `dscli -hmc1 array IP/name -user user -passwd passwd lsarraysite -dev SI`
- `dscli -hmc1 array IP/name -user user -passwd passwd lsarray -dev SI`
- `dscli -hmc1 array IP/name -user user -passwd passwd lsrank -dev SI`
- `dscli -hmc1 array IP/name -user user -passwd passwd lsextpool -dev SI`
- `dscli -hmc1 array IP/name -user user -passwd passwd lsfbvol -dev SI`
- `dscli -hmc1 array IP/name -user user -passwd passwd lssestg -dev SI`

- `dscli -hmc1 array IP/name -user user -passwd passwd> lslls -dev SI`
- `dscli -hmc1 array IP/name -user user -passwd passwd lsflash -dev SI`
- `dscli -hmc1 array IP/name -user user -passwd passwd ls session -dev SI`
- `dscli -hmc1 array IP/name -user user -passwd passwd lsprrcpath -dev SI`
- `dscli -hmc1 array IP/name -user user -passwd passwd lsprrc -dev SI`
- `dscli -hmc1 array IP/name -user user -passwd passwd showgmir -dev SI`

You should provide the following information for each DSCLI proxy:

- Name or IP address
- A user account profile (existing or specifically created for DRA)

You should provide the following information for each IBM DS array:

- Name or IP address
- A user account profile (existing or specifically created for DRA) with a *monitor* privilege (read-only). Avoid using the '<' and '>' characters in the password field.

You should also verify that:

- IP connectivity through SSH (default is port 22) is available between the DRA application server and each DSCLI proxy.
- IP connectivity is available between the DSCLI proxy and each IBM DS array that it scans.

Important: Do not configure the same IBM DS array to be scanned by more than one probe. This configuration may cause unpredictable results.

Note: By default, DRA connects to the proxies using SSH with user/password authentication. SSH with public key authentication is also supported. If you prefer, you can use Telnet; however, it is considered less secure than SSH. In terms of security provisioning, the only difference in using Telnet is that the default port is port 23, instead of the SSH port.

Data collection from XCLI through a UNIX proxy

DRA uses read-only XCLI commands to collect additional data from IBM XIV arrays. These commands run on one or more UNIX servers in the IT environment. Collectively, these servers can query all XIV arrays in the scope. These servers are also known as XCLI proxies.

When you select XCLI proxies, use the following best practices:

- Use IT administrative servers rather than production servers.
- Use fewer proxies. For example, if one host can query all arrays on all sites, use it as a proxy rather than using two or more hosts.
- Use proxies that are consistently up and available, rather than ones that are sometimes down or unreachable.

As a standard, DRA opens an SSH session to the proxy in the same way that it does to collect data from any UNIX host.

The connection from the proxy to the XIV array requires authorization details (user/password) for the array.

DRA uses the following read-only XCLI commands.

- `xcli -m array IP/name -u user -p passwd version_get`
- `xcli -m array IP/name -u user -p passwd config_get`
- `xcli -m array IP/name -u user -p passwd state_list`
- `xcli -m array IP/name -u user -p passwd fc_port_list`
- `xcli -m array IP/name -u user -p passwd pool_list`
- `xcli -m array IP/name -u user -p passwd vol_list`
- `xcli -m array IP/name -u user -p passwd mirror_list`
- `xcli -m array IP/name -u user -p passwd cg_list`
- `xcli -m array IP/name -u user -p passwd snap_group_list`
- `xcli -m array IP/name -u user -p passwd system_capacity_list`
- `xcli -m array IP/name -u user -p passwd time_list`

You should provide the following information for each XCLI proxy:

- Name or IP address
- A user account profile (existing or specifically created for DRA)

You should provide the following information for each XIV array:

- Name or IP address
- A user account profile (existing or specifically created for DRA) with a *read-only* privilege (Avoid using the '<' and '>' characters in the password field.)

Note: Due to an XIV bug, a read-only user cannot run the `mirror_list` command. Therefore, a *storage administrator* privilege may be required.

You should also verify that:

- IP connectivity through SSH (default is port 22) is available between the DRA application server and each XCLI proxy.
- IP connectivity is available between the XCLI proxy and each XIV array that it scans.

Important: Do not configure the same XIV array to be scanned by more than one probe. This configuration may cause unpredictable results.

Note: By default, DRA connects to the proxies using SSH with user/password authentication. SSH with public key authentication is also supported. If you prefer, you can use Telnet; however, it is considered less secure than SSH. In terms of security provisioning, the only difference in using Telnet is that the default port is port 23, instead of the SSH port.

Data collection from SVC

DRA uses read-only SVC CLI commands to collect additional data from SVC arrays.

As a standard, DRA opens an SSH session to the SVC in the same way that it does to collect data from any UNIX host.

DRA uses the following read-only SVC commands:

- `svcinfo lscluster`
- `svcinfo lsnode`
- `svcinfo lslicense`
- `svcinfo lsclusterip`
- `svcinfo lsportip`
- `svcinfo lsvdisk`
- `svcinfo lsiogrp`
- `svcinfo lsmdiskgrp`
- `svcinfo lsmdisk`
- `svcinfo lscontroller`
- `svcinfo lsfcmap`
- `svcinfo lsfeconsistgrp`
- `svcinfo lsrrrelationship`

- `svcinfo lsrrconsistgrp`
- `svcinfo lsquorum`
- `svcinfo lsvdiskcopy`
- `svcinfo lssevdiskcopy`
- `svcinfo lshost`
- `svcinfo lshostvdiskmap`
- `svcinfo lsfabric`
- `svcinfo lsmdiskextent`

You should provide the following information for each SVC array:

- Name or IP address
- A user account profile (existing or specifically created for DRA)

You should also verify that IP connectivity through SSH (default is port 22) is available between the DRA application server and each SVC array.

Data collection from HDS HiCommand/HP CommandView

DRA collects data from Hitachi Data Systems (HDS) and HP XP arrays by opening an HTTP connection to the HiCommand/CommandView server, or servers, if more than one is used.

DRA collects data using the following read-only requests:

- `GetServerInfo`
- `GetStorageArray`
- `GetHost`

To make sure that data collection goes smoothly, do the following:

- Provide each HiCommand server name or IP address.
- Provide the HiCommand Web application user name and password of a user with View rights that is assigned to a group. A default group is acceptable.
- Make sure that IP connectivity through HTTP (port 2001) is available between the DRA application server and each HiCommand server.

Data collection from NetApp

DRA collects data from NetApp storage arrays (also known as Filers) by connecting them using HTTP or HTTPS and issuing read-only commands using the NetApp ZAPI API.

You should provide the following information for each NetApp:

- Name or IP address
- A user account profile (existing or specifically created for DRA)
The user should be an administrative user assigned with the *Admin* role or a custom restricted user with the following capabilities:
 - login-http-admin
 - api-system-get-info
 - api-system-get-version
 - api-license-list-info
 - api-net-ifconfig-get
 - api-clock-get-clock
 - api-options-list-info
 - api-vfiler-list-info
 - api-aggr-list-info
 - api-volume-list-info
 - api-qtree-list
 - api-snapshot-list-info
 - api-lun-list-info
 - api-igroup-list-info
 - api-iscsi-node-get-name
 - api-lun-map-list-info
 - api-nfs-exportfs-list-rules
 - api-cifs-share-list-iter-start
 - api-cifs-share-list-iter-next
 - api-snapmirror-get-status
 - api-snapvault-primary-relationship-status-list-iter-start
 - api-snapvault-primary-relationship-status-list-iter-next
 - api-disk-list-info
 - api-net-resolve
 - api-volume-get-root-name
 - api-storage-adapter-get-adapter-info
 - api-storage-adapter-get-adapter-list

You should also verify that IP connectivity through HTTP (default is port 80) or HTTPS (default is port 443) is available between the DRA application server and each NetApp Filer.

The following example describes how to create a user with appropriate privileges:

- Create a new role using one of the following options. The latter option is more restricted (read-only):
 - `useradmin role add role_name -a login-http-admin,api-*`
 - `useradmin role add role_name -a login-http-admin,api-system-get-info,api-system-get-version,api-license-list-info,api-net-ifconfig-get,api-clock-get-clock,api-options-list-info,api-vfiler-list-info,api-aggr-list-info,api-volume-list-info,api-qtree-list,api-snapshot-list-info,api-lun-list-info,api-igroup-list-info,api-iscsi-node-get-name,api-lun-map-list-info,api-nfs-exportfs-list-rules,api-cifs-share-list-iter-start,api-cifs-share-list-iter-next,api-snapmirror-get-status,api-snapvault-primary-relationship-status-list-iter-start,api-snapvault-primary-relationship-status-list-iter-next,api-disk-list-info,api-net-resolve,api-snapvault-primary-relationship-status-list-iter-start,api-snapvault-primary-relationship-status-list-iter-next,api-disk-list-info`
- Create a new group:
 - `useradmin group add group_name -r role_name`
- Create a new user:
 - `useradmin user add user_name -g group_name`

Data collection from VMware vCenter

DRA collects data from vCenter using VMware's Virtual Infrastructure (VI) API by connecting to the vCenter server, or servers, if more than one is used. If SRM is used, DRA also collects SRM data using the SRM API.

DRA collects data by running read-only inquiries on the following entities:

- Data centers
- Data stores
- Host systems
- Virtual machines
- Clusters

To make sure that data collection goes smoothly, do the following:

- Provide each vCenter server name or IP address.
- Provide a read-only user name and password for each vCenter server.
- Make sure that IP connectivity through HTTPS (port 443) is available between the DRA application server and each vCenter server.

Data collection from HMC

DRA uses read-only HMC CLI commands to collect additional data from HMC.

As a standard, DRA opens an SSH session to the HMC in the same way that it does to collect data from any UNIX host.

DRA uses the following read-only HMC commands:

- `uname -a`
- `lshmc -V`
- `lssyscfg -r sys`
- `lssyscfg -r lpar -m hmc_system`
- `lshwres -r virtualio --subtype slot -m hmc_system --level slot`
- `lshwres -r virtualio --subtype scsi -m hmc_system`

You should provide the following information for each HMC:

- Name or IP address
- A user account profile (existing or specifically created for DRA) with *hmcviewer* role (read-only)

You should also verify that IP connectivity through SSH (default is port 22) is available between the DRA application server and each HMC.

Data collection from VIO

DRA uses read-only commands to collect additional data from VIO.

As a standard, DRA opens an SSH session to the VIO sever in the same way that it does to collect data from any UNIX host.

You should provide the following information for each VIO:

- Name or IP address.
- A user account profile (existing or specifically created for DRA)
The user should be either a regular user with *ksh* and permissions, as described in “[Privileged commands on AIX](#)” on page 36 or a restricted user with *rksh* (see below).
- If a restricted user is used, make sure that *PermitUserEnvironment* is set to *yes* in the */etc/ssh/sshd_config* file.

You should also verify that IP connectivity through SSH (default is port 22) is available between the DRA application server and each VIO.

The following example describes how to create a restricted user with appropriate privileges:

- 1 Log in to the VIO server using *padmin*.
- 2 Ensure that Enhanced RBAC is enabled:
 - `lsattr -El sys0 -a enhanced_RBAC`
 - If not, run `chdev -l sys0 -a enhanced_RBAC=true` and reboot.
- 3 Create the role:
 - `mkauth dfltmsg='Symantec' symantec`
 - `mkauth dfltmsg='Symantec DRA' symantec.dra`
 - `mkrole roletlist=ViewOnly authorizations=symantec.dra dfltmsg="Symantec DRA" dra`
 - `setkst`
- 4 Create the user:
 - `mkuser -attr pgrp=view drauser`
 - `chuser -attr roles=dra default_roles=dra drauser`

- 5 Create permission to run privileged commands. See “[Privileged commands on AIX](#)” on page 36 for the list of commands.
Do the following for each privileged command:

Note: In this example, the command is /usr/sbin/pcmpath.

- setsecattr -c eid=0 accessauths=symantec.dra innatprivs=PV_SU_
secflags=FSF_EPS authroles= /usr/sbin/pcmpath
- setkst
- oem_setup_env
- ln -s /usr/sbin/pcmpath /usr/ios/oem

Data collection from UNIX hosts

By default, DRA collects data from UNIX hosts by opening an SSH connection to the scanned hosts and issuing read-only commands. This is similar to SYMCLI data collection (described above). The commands DRA uses vary slightly, depending on the version of UNIX, as described below.

Note: If you use sudo, PowerBroker (pbrun), UPM (pmlrun), CA Access Control (seSUDO), super, or similar software, you should allow all privileged commands per platform on all hosts of that platform, even if some of the commands are not installed on certain hosts. This makes provisioning much simpler, and, if you install one of the commands in the future, ensures seamless compatibility. There is no harm in allowing non-existing commands on any of the specified command paths. These paths are read-only to all but root users.

To make sure that data collection goes smoothly, do the following:

- Provide the name or IP address of each UNIX host that is not auto-discovered through ECC, HiCommand/CommandView, or vCenter, as appropriate.
- Provide the user account profile on each UNIX host (existing or specifically created for DRA). You should use the same ID on all hosts, although you may use different IDs per platform or per individual host.
- If you prefer, use sudo, PowerBroker (pbrun), UPM (pmlrun), CA Access Control (seSUDO), super, or similar definitions on each UNIX host, as discussed above.
- Make sure that IP connectivity through SSH is available between the DRA application server and each UNIX host. The default SSH port is 22.

- On hosts using Symmetrix or CLARiiON, make sure you meet the following requirements:
 - DRA requires that at least one of the utilities – PowerPath, SymCLI, or inq (V7.3-487 and later) – is installed on each host.

Note: For AIX, these utilities are unnecessary when CLARiiON is used.

- If none of these utilities is available on a certain host, install the free EMC inq utility at /usr/local/bin.
- On hosts using HDS/HP XP, make sure you meet the following requirements:
 - DRA requires that at least one of the utilities (HDLM or inqraid) is installed on each UNIX host, and that inqraid is installed on Windows hosts.
 - If neither of these utilities is available on a certain host, install the free HDS inqraid utility at /HORCM/usr/bin/.
- On hosts using NetApp, make sure that at least one of the utilities (SnapDrive or sanlun) is installed on each host.

Note: By default, DRA connects to UNIX hosts using SSH with user/password authentication. SSH with public key authentication is also supported. If you prefer, you can use Telnet; however, it is considered less secure than SSH. In terms of security provisioning, the only difference in using Telnet is that the default port is port 23, instead of the SSH port.

Privileged commands on Solaris

Table 1-3 lists the privileged commands on Solaris.

Table 1-3 Privileged commands on Solaris

Command	Required when this is installed ...
/usr/sbin/fcinfo	
/etc/powermt display	EMC PowerPath
/usr/symcli/bin/symdg list	EMC Solution Enabler (SYMCLI)
/usr/symcli/bin/symcg list	EMC Solution Enabler (SYMCLI)
/usr/symcli/bin/sympd list	EMC Solution Enabler (SYMCLI)
/usr/symcli/bin/syminq	EMC Solution Enabler (SYMCLI)
/usr/local/bin/inq	EMC inq
/usr/sbin/vxdisk list	Veritas LVM/DMP
/usr/sbin/vxdisk path	Veritas LVM/DMP
/HORCM/usr/bin/inqraid	HDS horcm
/usr/local/bin/lunstat -t	HDS lunstat
/usr/DynamicLinkManager/bin/dlnmgr view	HDS HDLM
/opt/NetApp/snapdrive/bin/snapdrive storage show	NetApp SnapDrive
/opt/netapp/santools/sanlun lun show all	NetApp sanlun
/bin/cat */tnsnames.ora	Oracle
/HORCM/usr/bin/raidqry	HDS horcm
/HORCM/usr/bin/raidscan	HDS horcm
/HORCM/usr/bin/pairdisplay	HDS horcm
/sbin/xpinfo	XPINFO
/sbin/spmgr display	SecurePath
/sbin/autopath display	SecurePath
/usr/sbin/vxdmpadm list	Veritas LVM/DMP
/usr/sbin/datapath query device	IBM SDD
/opt/xiv/host_attach/bin/xiv_devlist	IBM XIV HAK

Table 1-3 Privileged commands on Solaris (Continued)

Command	Required when this is installed ...
/bin/cat	
/bin/ls	

Privileged commands on HP-UX

[Table 1-4](#) lists the privileged commands on HP-UX.

Table 1-4 Privileged commands on HP-UX

Command	Required when this is installed ...
/sbin/powermt display	EMC PowerPath
/usr/symcli/bin/symdg list	EMC Solution Enabler (SYMCLI)
/usr/symcli/bin/symcg list	EMC Solution Enabler (SYMCLI)
/usr/symcli/bin/sympd list	EMC Solution Enabler (SYMCLI)
/usr/symcli/bin/syminq	EMC Solution Enabler (SYMCLI)
/usr/local/bin/inq	EMC inq
/usr/sbin/vxdisk list	Veritas LVM/DMP
/usr/sbin/vxdisk path	Veritas LVM/DMP
/HORCM/usr/bin/inqraid	HDS inqraid
/usr/local/bin/lunstat -t	HDS lunstat
/usr/DynamicLinkManager/bin/dlnkmgr view	HDS HDLM
/opt/NetApp/snapdrive/bin/snapdrive storage show	NetApp SnapDrive
/opt/netapp/santools/sanlun lun show all	NetApp sanlun
/HORCM/usr/bin/raidqry	HDS horcm
/HORCM/usr/bin/raidscan	HDS horcm
/HORCM/usr/bin/pairdisplay	HDS horcm

Table 1-4 Privileged commands on HP-UX (Continued)

Command	Required when this is installed ...
/sbin/xpinfo	XPINFO
/sbin/spmgr display	SecurePath
/sbin/autopath display	SecurePath
/usr/sbin/vxdmpadm list	Veritas LVM/DMP
/usr/sbin/datapath query device	IBM SDD
/opt/xiv/host_attach/bin/xiv_devlist	IBM XIV HAK
/bin/cat	
/bin/ls	

Privileged commands on Linux

Table 1-5 lists the privileged commands on Linux.

Table 1-5 Privileged commands on Linux

Command	Required when ...
/sbin/powermt display	EMC PowerPath is installed.
/usr/symcli/bin/symdg list	EMC Solution Enabler (SYMCLI) is installed.
/usr/symcli/bin/symcg list	EMC Solution Enabler (SYMCLI) is installed.
/usr/symcli/bin/sympd list	EMC Solution Enabler (SYMCLI) is installed.
/usr/symcli/bin/syminq	EMC Solution Enabler (SYMCLI) is installed.
/usr/local/bin/inq	EMC inq is installed.
/usr/sbin/vxdisk list	Veritas LVM/DMP is installed.
/usr/sbin/vxdisk path	Veritas LVM/DMP is installed.
/usr/sbin/lvdisplay	LVM2 is used.
/usr/sbin/vgdisplay	LVM2 is used.
/usr/sbin/pvdisplay	LVM2 is used.

Table 1-5 Privileged commands on Linux (Continued)

Command	Required when ...
/sbin/multipath -l	MPIO is used.
/HORCM/usr/bin/inqraid	HDS inqraid is installed.
/usr/local/bin/lunstat -t	HDS lunstat is installed.
/usr/DynamicLinkManager/bin/dlnkmgr view	HDS HDLM is installed.
/opt/NetApp/snapdrive/bin/snapdrive storage show	NetApp SnapDrive is installed.
/opt/netapp/santools/sanlun lun show all	NetApp sanlun is installed.
/sbin/scsi_id	
/bin/raw -qa	A Linux raw character device is used.
/usr/bin/raw -qa	A Linux raw character device is used.
/HORCM/usr/bin/raidqry	HDS horcm is installed.
/HORCM/usr/bin/raidscan	HDS horcm is installed.
/HORCM/usr/bin/pairdisplay	HDS horcm is installed.
/sbin/xpinfo	XPINFO is installed.
/sbin/spmgr display	SecurePath is installed.
/sbin/autopath display	SecurePath is installed.
/usr/sbin/vxdmpadm list	Veritas LVM/DMP is installed.
/usr/sbin/datapath query device	IBM SDD
/opt/xiv/host_attach/bin/xiv_devlist	IBM XIV HAK
/bin/cat	
/bin/ls	

Privileged commands on AIX

[Table 1-6](#) lists the privileged commands on AIX.

Table 1-6 Privileged commands on AIX

Command	Required when this is installed ...
/usr/sbin/powermt display	EMC PowerPath
/usr/symcli/bin/symdg list	EMC Solution Enabler (SYMCLI)
/usr/symcli/bin/symcg list	EMC Solution Enabler (SYMCLI)
/usr/symcli/bin/sympd list	EMC Solution Enabler (SYMCLI)
/usr/symcli/bin/syminq	EMC Solution Enabler (SYMCLI)
/usr/local/bin/inq	EMC inq
/usr/es/sbin/cluster/utilities/cldisp	PowerHA
/usr/es/sbin/cluster/diag/clver	PowerHA
/usr/sbin/vxdisk list	Veritas LVM/DMP
/usr/sbin/vxdisk path	Veritas LVM/DMP
/HORCM/usr/bin/inqraid	HDS inqraid
/usr/local/bin/lunstat -t	HDS lunstat
/usr/DynamicLinkManager/bin/dlnkmgr view	HDS HDLM
/opt/NetApp/snapdrive/bin/snapdrive storage show	NetApp SnapDrive
/opt/netapp/santools/sanlun lun show all	NetApp sanlun
/HORCM/usr/bin/raidqry	HDS horcm
/HORCM/usr/bin/raidscan	HDS horcm
/HORCM/usr/bin/pairdisplay	HDS horcm
/sbin/xpinfo	XPINFO
/sbin/spmgr display	SecurePath
/sbin/autopath display	SecurePath
/usr/sbin/vxdmpadm list	Veritas LVM/DMP

Table 1-6 Privileged commands on AIX (Continued)

Command	Required when this is installed ...
/usr/sbin/datapath query device	IBM SDD
/usr/sbin/pcmpath query device	IBM PCMSDD
/opt/xiv/host_attach/bin/xiv_devlist	IBM XIV HAK
/bin/cat	
/bin/ls	

Data collection from Windows

DRA collects data from Windows hosts using Windows Management Instrumentation (WMI) queries and WMI remote command invocation. Windows Remote Management (WinRM) can also be used on hosts on which it is installed and configured.

The following WMI namespaces are queried:

- root/CIMV2
- root/WMI

The following WMI classes are queried:

- Under the CIMV2 namespace:
 - Win32_OperatingSystem
 - Win32_Processor
 - Win32_ComputerSystem
 - Win32_NetworkAdapterConfiguration
 - Win32_Service
 - Win32_PageFile
 - Win32_LogicalDisk
 - Win32_MappedLogicalDisk
 - Win32_Share
 - Win32_DiskDrive
- Under the WMI namespace:
 - MSiSCSIInitiator_SessionClass

If the corresponding binaries are installed on the server, the following commands are executed through WMI remote command invocation; otherwise, the commands are ignored:

- powermt display, syminq, symdg list, symcg list, inq (EMC)
- inqraid \$Phys -CLIWP, dlncmgr view -lu -item all, pairedisplay -CLI -l (HDS)
- sdcli disk list, iscsicli SessionList, dsmcli lun attributes (NetApp)
- datapath query device | version (IBM DS)
- xiv_devlist -t xml -o all (IBM XIV)
- haclus -state, haclus -display -localclus, hasys -display -localclus, hasys -nodeid, hagr -display -localclus, hares -display -localclus, hares -dep, llstat -c, gabconfig -v, gabconfig -l, vxfenadm -d, hagr -dep, hahb -display (Veritas cluster)

- vxprint, vxdisk list | path, vxdmpadm pathinfo, vxvol volinfo, vxlicrep (Veritas storage foundation)
- cluster /properties | /privproperties | node | node /properties | node /privproperties | group | group /properties | group /privproperties | resource | resource /properties | resource /privproperties | network | network /properties | network /privproperties | netinterface | netinterface /properties | netinterface /privproperties | /quorum | /listshares | /listnetpriority (Microsoft cluster)
- HbaCmd HbaAttributes, HbaCmd PortAttributes, HbaCmd GetDriverParams, HbaCmd TargetMapping, scli -Z, fcinfo /details, fcinfo /ports /details, fcinfo /mapping (HBA)
- diskpart list disk, diskpart list volume, diskpart detail disk, diskpart detail volume (LVM)
- mx matrix status, mx config list, mx server listsoftware, mx application status, mx mfs status, mx mni listinstances, mx vfs status, mx vfs_share status, sandiskinfo -ial, sandiskinfo -val, sandiskinfo -fal, sandiskinfo --dynvol_properties -al (HP PolyServe)

To make sure that data collection goes smoothly, you should do the following:

- Provide the name or IP address of each Windows host that is not auto-discovered through ECC, HiCommand/CommandView, or vCenter, as appropriate.
- Provide the user account profile on each Windows host (existing or specifically created for DRA). You should use the same ID on all hosts, although you may use different IDs per domain or per individual host. To simplify provisioning, it is also preferred to use non-privileged domain users rather than local users.
- Make sure each host has local administrative rights. Local administrative rights are required on each host for the designated user profile of that host. This is a Microsoft requirement that enables WMI access and remote command invocation.
- Make sure that IP connectivity is permitted through WMI on all TCP ports and UDP ports 135, 137, 138, and 139.
WMI is based on DCOM. You can further limit the TCP ports allowed for WMI/DCOM communication, but this requires a significantly more complex provisioning process.

For more information, see “Using Distributed COM with Firewalls” at:

<http://msdn.microsoft.com/en-us/library/ms809327.aspx>

Symantec recommends that all TCP port connections originating from the DRA server be allowed. Typically, the DRA server is placed on a secure management subnet.

An alternative to opening all the TCP ports is to use Windows Remote Management (WinRM). WinRM uses only one port (80/5985, by default). To use it, it should be installed and activated on the scanned hosts. When WinRM is used, avoid using the '<' and '>' characters in the password.

For more information, see [“Network and other environmental recommendations”](#) on page 48.

- Make sure CIFS connectivity is permitted (port 445). CIFS is used to run vendor-native commands.
- On hosts using Symmetrix or CLARiiion, make sure you meet the following requirements:
 - DRA requires that at least one of the utilities – PowerPath, SymCLI, or inq (V7.3-487 and later) – is installed on each host.
 - If none of these utilities is available on a certain host, install the free EMC inq utility.
- On hosts using HDS, make sure you meet the following requirements:
 - DRA requires that at least one of the utilities (HDLM or inqraid) is installed on each host.
 - If neither of these utilities is available on a certain host, install the free HDS inqraid utility.
- On hosts using NetApp, make sure that at least one of the utilities (sdcli or dsmcli) is installed on each host.
- On hosts using IBM DS, make sure that IBM SDD is installed on each host.
- On hosts using XIV, make sure that the IBM XIV Host Attachment Kit (HAK) is installed on each host.

Data collection from databases

DRA collects data from databases by opening a JDBC connection to each database and running read-only select queries on certain system tables. The specific queries vary from one database platform to another, as described below.

DRA automatically discovers databases and database instances. However, if you use virtual IP addresses or non-default vendor ports on certain instances, you should explicitly specify them for each instance.

To make sure that data collection goes smoothly, do the following:

- Provide the virtual IP (name or address) for each instance that does not use its primary host IP address.
- Provide a connection port for each instance that does not (also) listen on the vendor default port. The default ports are as follows:

MS SQL Server	1433
Oracle	1521
Sybase	5000
UDB	50000

- Provide the user account profile on each database (existing or specifically created for DRA). Specific information for each database vendor is provided below. You should use the same ID on all databases, although you may use different IDs per platform or per individual database.
- Make sure that the required specific rights have been granted to each user account. For more details, refer to the subsections below by vendor.
- Make sure that IP connectivity through JDBC is available between the DRA application server and each database server. You can use the default port number or specify another port.

Data collection from Oracle

As shown in [Table 1-7](#), data is collected from Oracle by connecting to each instance and querying the following V\$ views:

Table 1-7 Data collection from Oracle

View	Comments
v\$instance	
v\$database	
v\$datafile	
v\$controlfile	
v\$logfile	
v\$archive_dest	
v\$parameter	
v\$tablespace	
v\$backup	
v\$archive_dest_status	
v\$asm_diskgroup	Relevant only if ASM is used
v\$asm_disk	Relevant only if ASM is used
v\$archive_gap	Relevant only in a primary-standby database configuration
v\$log	
v\$log_history	
v\$archived_log	
v\$database_incarnation	
v\$diag_info	Required for Oracle version 11
v\$dataguard_config	Relevant only in a primary-standby database configuration
v\$dataguard_status	Relevant only in a primary-standby database configuration
v\$logstdby	Relevant only in a primary-standby database configuration

Table 1-7 Data collection from Oracle

View	Comments
v\$logstdby_stats	Relevant only in a primary-standby database configuration
v\$managed_standby	Relevant only in a primary-standby database configuration
v\$standby_log	Relevant only in a primary-standby database configuration
v\$logstdby_process	Relevant only in a primary-standby database configuration
v\$logstdby_progress	Relevant only in a primary-standby database configuration
v\$logstdby_state	Relevant only in a primary-standby database configuration
v\$logstdby_transaction	Relevant only in a primary-standby database configuration
v\$resource_limit	
v\$sga	
v\$tempfile	
v\$version	
v\$option	
v\$spparameter	
v\$active_instances	
v\$archived_log	
dba_logstdby_events	Relevant only in a primary-standby database configuration
dba_logstdby_log	Relevant only in a primary-standby database configuration
dba_logstdby_not_unique	Relevant only in a primary-standby database configuration
dba_logstdby_parameters	Relevant only in a primary-standby database configuration

Table 1-7 Data collection from Oracle

View	Comments
dba_logstdby_progress	Relevant only in a primary-standby database configuration
dba_logstdby_skip	Relevant only in a primary-standby database configuration
dba_logstdby_skip_transaction	Relevant only in a primary-standby database configuration
dba_logstdby_unsupported	Relevant only in a primary-standby database configuration
dba_logstdby_history	Relevant only in a primary-standby database configuration
dba_temp_files	
dba_free_space	
dba_tablespace	
dba_data_files	
dba_libraries	

Note: For suggestions about secure ways to grant these read-only privileges, see [Appendix A, “Methods for secure privilege provisioning”](#) on page 51. You should grant all the rights specified above for each instance. This makes provisioning much simpler, and, if you need these views in the future, ensures seamless compatibility.

Data collection from Sybase

DRA collects data from Sybase by connecting to each Sybase server master database and querying the following system tables:

- sysdatabases
- sysdevices
- sysusage
- @@ queries (including @@servername, @@version, @@boottime, @@pagesize, @@nodeid and @@version_as_integer)

Data collection from Microsoft MS SQL Server

DRA collects data from Microsoft MS SQL Server by connecting to each server instance master database and querying the following system tables:

- master.dbo.sysdatabases
- master.dbo.sysaltfiles
- master.sys.databases
- master.sys.configurations
- @@ queries (including @@servicename, @@servername and @@version)
- msdb.dbo.backupfile
- msdb.dbo.backupmediafamily
- msdb.dbo.backupmediaset
- msdb.dbo.backupset
- master.sysdevices

If you are using MS SQL make sure you specify whether the default is to use Windows authentication, MS SQL Server authentication, or both. If you use different settings for different databases, explicitly specify the method used for each database.

Note: When using Windows Authentication, make sure that the DRA server is a member of the same domain as the scanned MS SQL Server. In addition, make sure that the user used for scanning the database has login privileges on the DRA server.

Also, configure the permission *connect* for each database inside the server instance (map the user to all databases).

The following tables/procedures are executed/queried:

- sysfiles
- sysfilegroups
- sp_spaceused

Data collection from IBM UDB (DB2)

DRA collects data from UDB by connecting to each UDB database through JDBC and querying the sysibmadm.dbpaths (version 9 only) system table.

DRA also collects data by connecting to each UDB database through JDBC and running the following procedures:

- sysproc.env_get_inst_info
- sysproc.snapshot_container
- sysproc.snapshot_database
- sysproc.snap_get_sto_paths

Make sure that you meet the following UDB-specific requirements:

- On version 8.1 fix-packs 7 and later, the user profile assigned to DRA should be a member of the UDB SYSMON group.
- On version 8.1 fix-packs 1-6, the user profile assigned to DRA should be a member of the SYSCTRL, SYSMAINT, or SYSADM group.

Data collection from Veritas Cluster Server

DRA collects data from Veritas Cluster Server (VCS) by connecting to each VCS node through a VCS API and querying the following privileged read-only VCS commands. A non-root user needs sudo, PowerBroker (pbrun), UPM (pmsrun), CA Access Control (seSUDO), super, or similar software to execute the following privileged commands:

- /opt/VRTSvcs/bin/haclus -state
- /opt/VRTSvcs/bin/haclus -display -localclus
- /opt/VRTSvcs/bin/hasys -display -localclus
- /opt/VRTSvcs/bin/hasys -nodeid
- /opt/VRTSvcs/bin/hagrp -display -localclus
- /opt/VRTSvcs/bin/hagrp -dep
- /opt/VRTSvcs/bin/hares -display -localclus
- /opt/VRTSvcs/bin/hares -dep
- /opt/VRTSvcs/bin/hahb -display
- /sbin/lltstat -c
- /sbin/gabconfig -v
- /sbin/gabconfig -l
- /sbin/vxfenadm -d

Data collection from EMC Control Center

Data collection is based on opening a Java Database Connectivity (JDBC) connection to the EMC ControlCenter (ECC) repository (StorageScope sts view). The ECC repository and Storage Scope must be installed on the same server. Read-only select queries are used to obtain data.

Note: For a list of read-only queries, see [Appendix A, “Methods for secure privilege provisioning”](#) on page 51.

Note: For configuring ECC 6 with JDBC Secure Sockets Layer (SSL), see [“Configuring DRA for ECC scanning over JDBC SSL \(Oracle\)”](#) on page 66.

When you set up data collection, do the following:

- Provide the name or IP address of each ECC Repository server used in the scanned environment.
- Provide the user/password for the ECC Repository RAMBDB. The default account is stsview/sts.
- Make sure that IP connectivity through JDBC is available between the DRA application server and each ECC repository server. The default port is 1521.

Mail server configuration

You should allow DRA to send email messages containing automated, scheduled reports. To support this configuration, details of an existing customer email server should be provided.

You provide the following details for your email server:

- The mail server name or address.
- The connection port. The default is 25.
- Whether authentication is required. The default is No.
- The required user name used in the From: field. The default is *DRA@customer_domain*.
- Whether encryption is required. The default is No.

Make sure that IP connectivity is available between the DRA application server and the mail server. You can use the default port (25) or specify another.

Network and other environmental recommendations

You should place the DRA server on the least-active site (such as a passive DR site), rather than on the most active one. That way, if the main production site fails, DRA is available for possible postmortem troubleshooting.

If your IT environment uses an IP filtering device, such as a firewall, to create different security zones, specific port access must be allowed from the DRA server to the scanned elements. The sections of this document that discuss platform-specific data collection describe this scenario in more detail.

In this scenario, if you use subnets, you should place the DRA server on the same subnet as the storage management servers, such as ECC or HiCommand. The main advantages of this placement are:

- In most cases, these subnets are already provisioned with the necessary pass-through configuration that DRA needs.
- These subnets are often the most secure, and access to them is more highly-regulated than others. It is simpler and safer to allow outgoing connections from the DRA server to the scanned components.

Because the DRA server is not mission-critical, you do not have to configure it for high-availability or DR, although such configurations are possible. It may be beneficial to back up the entire server each week or to use replicated disks to store server data.

You can always restart the DRA repository, even if an IT failure led to its loss, with very little impact on DRA's usability.

Summary of ports and protocols used by DRA

The following table describes the ports and protocols used by DRA.

Table 1-8 DRA ports and protocols

From	To	Port/protocol	Description
Desktop of DRA users	DRA server	8080/HTTP	Access to DRA web UI
DRA server	HDS HiCommand/ HP CommandView	2001/HTTP	Connecting to HDS/ HP management consoles
DSCLI proxy	IBM DS arrays	1720, 1722, 1750, 8451-8455	DSCLI
XCLI proxy	IBM XIV arrays	7778	XCLI
DRA server	NetApp filers	80/HTTP or 443/HTTPS	Connecting to filers
DRA server	SVC arrays	22/SSH	Connecting to SVC
DRA server	HMC	22/SSH	Connecting to HMC
DRA server	UNIX servers	22/SSH	Connecting to UNIX servers
DRA server	Windows servers	All TCP, UDP 135-9/ WMI	Connecting to Windows servers
DRA server	Windows servers	80/5985 (default) and 445/WinRM	Connecting to Windows servers
DRA server	Oracle instances ip/vip	1521 (default)/JDBC	Connecting to Oracle
DRA server	SQL Server instances ip/vip	1433 (default)/JDBC	Connecting to MS-SQL
DRA server	IBM DB2/UDB databases ip/vip	50000 (default)/JDBC	Connecting to DB2/ UDB
DRA server	Sybase instances ip/vip	5000 (default)/JDBC	Connecting to Sybase
DRA server	vCenter/vSphere	443/SOAP	
DRA server	Mail server	25 (default)/SMTP	Sending emails from DRA

From	To	Port/protocol	Description
DRA server	EMC Control Center server	1521/JDBC 1575/JDBC SSL	Connecting to ECC RAMBDB views
DRA server	ftp.veritas.com	80 (default)/HTTP 20-22 (default)/FTP	Automatic update
DRA server	Active directory LDAP host	389 (default)	Optional when using active directory for users

Methods for secure privilege provisioning

This appendix suggests methods for secure privilege provisioning for the various entities supported by DRA. It includes the following topics:

- [“sudo on UNIX hosts”](#) on page 51
- [“UNIX Privilege Manager”](#) on page 54
- [“Suggested Oracle grant provisioning”](#) on page 56
- [“Suggested MS SQL Server grant provisioning”](#) on page 57
- [“Queries used to scan EMC ECC”](#) on page 66
- [“Configuring DRA for ECC scanning over JDBC SSL \(Oracle\)”](#) on page 66

sudo on UNIX hosts

This section provides suggestions for a sudo definition on UNIX hosts, per platform. You should verify that the specified path for vendor-specific utilities, such as SYMCLI, inq, or inqraid, matches your environment, and, if necessary, adjust the suggested definitions.

When you use sudo version 1.6.9 and later, DRA must run `sudo -E` to preserve environment variables.

Requiretty must be set to off, which is the default setting.

SymCLI proxy

```
username ALL= NOPASSWD: /usr/symcli/bin/* list*, /usr/symcli/bin/symcli  
-def
```

All UNIX flavors

```
username ALL= NOPASSWD: /bin/cat *, /bin/ls *
```

Servers that use EMC storage

```
username ALL= NOPASSWD: /usr/symcli/bin/symdg list *,  
/usr/symcli/bin/symcg list *, /usr/symcli/bin/sympd list *,  
/usr/symcli/bin/syminq *, /usr/local/bin/inq *, /usr/sbin/powermt display*,  
/sbin/powermt display*, /etc/powermt display*
```

Servers that use HDS/HP XP storage

```
username ALL= NOPASSWD: /HORCM/usr/bin/inqraid *,  
/HORCM/usr/bin/raidqry *, /HORCM/usr/bin/raidscan *,  
/HORCM/usr/bin/pairdisplay *, /usr/local/bin/lunstat *,  
/usr/DynamicLinkManager/bin/dlnkmgr view *, /sbin/xpinfo*, /sbin/spmgr  
display*, /sbin/autopath display*
```

Servers that use NetApp storage

```
username ALL= NOPASSWD: /opt/netapp/santools/sanlun lun show all,  
/opt/NetApp/snapdrive/bin/snapdrive storage show*
```

Servers that use IBM storage

```
username ALL= NOPASSWD: /usr/sbin/datapath query device,  
/usr/sbin/pcmpath query device, /opt/xiv/host_attach/bin/xiv_devlist
```

Servers with Veritas SF and Cluster

```
username ALL= NOPASSWD: /usr/sbin/vxdisk path, /usr/sbin/vxdisk list*,  
/usr/sbin/vxdmpadm list*, /sbin/lltstat -c, /sbin/gabconfig -v, /sbin/gabconfig -l,  
/sbin/vxfenadm -d, /opt/VRTSvcs/bin/haclus -state, /opt/VRTSvcs/bin/haclus  
-display*, /opt/VRTSvcs/bin/hagrp -display*, /opt/VRTSvcs/bin/hagrp -dep,  
/opt/VRTSvcs/bin/hares -display*, /opt/VRTSvcs/bin/hares -dep,  
/opt/VRTSvcs/bin/hasys -display*, /opt/VRTSvcs/bin/hasys -nodeid,  
/opt/VRTSvcs/bin/hahb -display*, /bin/cat */listener.log, /bin/cat *alert_*.log,  
/bin/cat */listener.ora
```

Additional for Solaris servers

```
username ALL= NOPASSWD: /usr/sbin/fcinfo
```

Additional for Linux servers

```
username ALL= NOPASSWD: /usr/sbin/vgdisplay, /usr/sbin/lvdisplay*,  
/usr/sbin/pvdisplay, /sbin/multipath -l, /sbin/scsi_id *, /bin/raw -qa,  
/usr/bin/raw -qa
```

Additional for AIX servers

```
username ALL= NOPASSWD:  
/usr/es/sbin/cluster/utilities/cldisp,/usr/es/sbin/cluster/diag/clver
```

UNIX Privilege Manager

This section provides suggestions for a UPM definition on UNIX hosts.

You should verify that the specified path for vendor-specific utilities, such as SYMCLI, inq, or inqraid, matches your environment, and, if necessary, adjust the suggested definitions.

Below is an example of a UPM profile definition for DRA:

```
#####
# Privilege Manager Profile
#
# This profile permits the drauser user to run read only commands as the root user.
#####

### Data

enableprofile      = true;           # set to false to disable the profile
profile            = "dra";          # Profile Name
enableKeystrokeLogging = false;      # Enable Keystroke Logging?
enableAuthentication = false;      # User Authentication Required for all commands?
enableTimeRestrictions = false;     # Apply time restriction to execution of commands
restrictionHours    = {"7:00","22:00"}; # Define using the 24 hour format without a leading
zero.
enableRemoteCmds    = false;         # Should remote cmds be allowed for privilege cmds
                                     # (ie submithost != runhost)?

authUser            = "root";         # runuser to use when running the authCommands
                                     # Set to empty string to run the command as the
                                     # submitting user - ie set runuser = user (ie the default)
authGroup            = "root";        # rungroup to use when running the authCommands
                                     # Set to empty string to run the command as the
                                     # submitting group - ie set rungroup = group (ie the
                                     # default)

shellProfile         = "restricted";  # If you want to allow users matching this profile to
run                                                          # privilege manager shells, then this is the name of
                                                           # the shell profile to include. The shell profiles
                                                           # are copied to <poicydir>/pro

files/shellprofiles,                                       # and defines shell-specific configuration.

### List of profile members ###

# Groups - Users can be assigned to this profile by their group membership
authGroups={
};                                                         # Description
                                                         # No groups assigned to this profile

# Users - Alternatively, users can be assigned to this profile individually
authUsers={
  "drauser"                                               # Description
                                                         # Allow all users to run a command as self
};

# Hosts - Hosts can be assigned individually by adding their FQDN
authHosts={
  ALL                                                       # Description
                                                         # Allow all hosts when running a command as self
};

### List of profile commands ###

# Authorized commands - these commands are executed as the authUser defined above
authCmds={
  "/usr/sbin/vgdisplay",                                   # Description
  "/usr/sbin/lvdisplay *",                                # Linux commands
  "/usr/sbin/pvdisplay",
  "/sbin/multipath -l",
  "/sbin/scsi_id *",
  "/bin/raw -qa",
  "/usr/bin/raw -qa",

```

```

"/usr/bin/fcinfo",
"/usr/es/sbin/cluster/utilities/cldisp",
"/usr/es/sbin/cluster/diag/clver",
"/usr/sbin/vxdisk list *",
"/usr/sbin/vxdisk path",
"/usr/sbin/vxdmpadm list *",
"/sbin/lltstat -c",
"/sbin/gabconfig -v",
"/sbin/gabconfig -l",
"/sbin/vxfsadm -d",
"/opt/VRTSvcs/bin/haclus -state",
"/opt/VRTSvcs/bin/haclus -display -localclus",
"/opt/VRTSvcs/bin/hasys -display -localclus",
"/opt/VRTSvcs/bin/hasys -nodeid",
"/opt/VRTSvcs/bin/hagrp -display -localclus",
"/opt/VRTSvcs/bin/hagrp -dep",
"/opt/VRTSvcs/bin/hares -display -localclus",
"/opt/VRTSvcs/bin/hares -dep",
"/opt/VRTSvcs/bin/hahp -display",
"/usr/symcli/bin/syminq *",
"/usr/symcli/bin/sympd list *",
"/usr/symcli/bin/symdg list *",
"/usr/symcli/bin/symcg list *",
"/usr/local/bin/inq *",
"/etc/powermt display *",
"/usr/sbin/powermt display *",
"/sbin/powermt display *",
"/HORCM/usr/bin/inraid *",
"/HORCM/usr/bin/raidqry *",
"/HORCM/usr/bin/raidscan *",
"/HORCM/usr/bin/pairdisplay *",
"/usr/local/bin/lunstat *",
"/usr/DynamicLinkManager/bin/dlnkmgr view *",
"/sbin/xpinfo *",
"/sbin/spmgr display *",
"/sbin/autopath display *",
"/opt/NetApp/snapdrive/bin/snapdrive storage show *",
"/opt/netapp/santools/sanlun lun show *",
"/usr/sbin/datapath query device",
"/usr/sbin/pcmpath query device",
"/opt/xiv/host_attach/bin/xiv_devlist",
"/bin/cat *",
"/bin/ls *",
"/usr/symcli/bin/symcfg list *",
"/usr/symcli/bin/symdev list *",
"/usr/symcli/bin/symaudit list *",
"/usr/symcli/bin/symdisk list *",
"/usr/symcli/bin/symcli -def"
};

processProfile();

```

Solaris commands
AIX commands
VXDMP on All UNIX Flavors
VCS on All UNIX Flavors
EMC Symmetrix/CLARiion on All UNIX Flavors
PowerPath on Solaris
PowerPath on AIX
PowerPath on HP-UX/Linux
HDS on All UNIX Flavors
HP on All UNIX Flavors
NetApp on All UNIX Flavors
IBM DS on All UNIX Flavors
IBM DS on AIX
IBM XIV on All UNIX Flavors
All UNIX Flavors
SymCLI Proxy on All UNIX Flavors

Note: This is a joined profile that contains all the commands possibly required for all supported UNIX options. Separated profiles may be created per platform.

Note: The fields shown in **bold red** text in the preceding example may be changed from those specified in order to match customer-specific security requirements.

Suggested Oracle grant provisioning

```
CREATE USER drauser IDENTIFIED BY [drauserpassword];  
grant create session to drauser;  
grant select any dictionary to drauser;  
grant select on dba_logstdby_events to drauser;  
grant select on dba_logstdby_log to drauser;  
grant select on dba_logstdby_not_unique to drauser;  
grant select on dba_logstdby_parameters to drauser;  
grant select on dba_logstdby_progress to drauser;  
grant select on dba_logstdby_skip to drauser;  
grant select on dba_logstdby_skip_transaction to drauser;  
grant select on dba_logstdby_unsupported to drauser;  
grant select on dba_logstdby_history to drauser;  
grant select on dba_temp_files to drauser;  
grant select on dba_free_space to drauser;  
grant select on dba_temp_files to drauser;  
grant select on dba_free_space to drauser;  
grant select on dba_tablespace to drauser;  
grant select on dba_data_files to drauser;  
grant select on dba_libraries to drauser;
```


Suggested MS SQL Server grant provisioning

The suggestions described in this section apply for MS SQL Server grant provisioning.

Important: When you create a user account for scanning, connect using an admin user account.


To create an account for scanning MS SQL Server 2008 using Microsoft SQL Server Studio

- 1 Connect to the relevant server using the admin user account.
- 2 Select **Security > Logins**.



- 3 Right-click on **Logins** and select **New Login**. The Login - New window opens.
- 4 Provide the login name and select the appropriate connection method (Windows Authentication or SQL Server Authentication).
- 5 Select the User Mapping page.
- 6 Select the checkboxes for the **master** and **msdb** databases under the Map column.
- 7 In the left pane, select **Databases > System Databases > Master**.
- 8 Right-click on **Master**, and select **Properties**.


9 Select the Permissions page.



- 10 Click the **Add** button, type in the user just created, and click **OK**.
- 11 In the Explicit permissions for user area, select the **Grant** checkbox for the Select, Create Database, and Connect rows.
- 12 To save your changes, click **OK**.


To create an account for scanning MS SQL Server 2005 using Microsoft SQL Server Studio

- 1 Connect to the relevant server using the admin user account.
- 2 Select **Security > Logins**.




- 3 Right-click on **Logins** and select **New Login**. The Login - New window opens.

- 4 Name the login and select the appropriate connection method (Windows authentication or SQL Server authentication).
- 5 Select the User Mapping page.
- 6 Under the Map column, select the check boxes for the **master** and **msdb** databases.
- 7 In the left pane, select **Databases > System Databases > Master > Views > System Views**.
- 8 In the right pane, right-click on **sysaltfiles**, and select **Properties**.
- 9 Select the Permissions page.



- 10 Click the **Add** button, type in the user just created, and click **OK**.
- 11 In the Users or roles area, select the relevant user.
- 12 In the Explicit permissions for user area, select the **Grant** check box for the Select permission row.
- 13 To save your changes, click **OK**.


- 14
- In the left pane, right-click on **SQL Server** and select **Permissions**. The Permissions page opens.



- 15
- In the Logins or roles area, select the relevant user.
- 16
- In the Explicit permissions for user area, select the **Grant** check box for the View any definition row, and click **OK** to save your changes.
Alternatively, you can execute **grant view ANY definition to user** from the command line.


To create an account for scanning MS SQL Server 2000 using SQL Server Enterprise Manager

1 Under the relevant SQL Server Group, select **Security > Logins**.




- 2 Right-click **Logins** and select **New Login**.
- 3 Name the login and select the appropriate connection method (Windows authentication or SQL Server authentication).
- 4 Select the Database Access tab.

- 5
- Under the Permit column, select the check boxes for the **master** and **msdb** databases.



- 6
- Under the relevant SQL Server group, select **Databases > master**.
- 7
- In the left pane, select **Tables**.
- 8
- In the right pane, select **sysaltfiles**.
- 9
- Right-click on **sysaltfiles**, and select **All Tasks > Manage Permissions**.


10 Check the **SELECT** check box for the user you just created.



11 To save your changes, click **OK**.


To create an account for scanning MS SQL Server 2000 using Microsoft SQL Server Studio

- 1 Connect to the relevant server using the admin user account.
- 2 Select **Security > Logins**.




3 Right-click on **Logins** and select **New Login**. The Login - New window opens.

- 4
- Name the login and select the appropriate connection method (Windows authentication or SQL Server authentication).
- 5
- Select the User Mapping page.



- 6
- Under the Map column, select the check boxes for the **master** and **msdb** databases.
- 7
- Select **Databases > System Databases > Master > Tables > System Tables**.
- 8
- In the right pane, select **sysaltfiles**.
- 9
- Right-click on **sysaltfiles**, and select **Properties**.

10 Select the Permissions page.



- 11 Click the **Add** button, enter the user you just created, and click **OK**.
- 12 In the Users or roles area, select the relevant user.
- 13 In the Explicit permissions for user area, select the **Grant** check box for the Select permission row.
- 14 To save your changes, click **OK**.

Queries used to scan EMC ECC

DRA uses read-only select queries to collect data from the following tables:

- STS_ARRAY
- STS_ARRAY_DEVICE
- STS_ARRAY_META_DEVICE
- STS_ARRAY_PORT_TO_DEV
- STS_ARRAY_REPLICA
- STS_HOST
- STS_HOST_FS
- STS_HOST_LOGICALVOLUME
- STS_HOST_VOLUMEGROUP
- STS_HOST_DEVICE
- STS_HOST_FS_DEVICE

Configuring DRA for ECC scanning over JDBC SSL (Oracle)

Scanning ECC with Oracle Advanced Security, where the ECC repository is configured to accept SSL-encrypted and authenticated connections only, requires the following steps:


- 1 Configure the policy to use the JDBC-SSL protocol.
See [“Configuring an SSL policy”](#) on page 67.
- 2 Copy the wallet file that holds the authentication information from the ECC Repository server to the server running DRA.
See [“Copying the wallet file from the ECC Repository server”](#) on page 69.
- 3 Configure credentials in the standard manner.

Configuring an SSL policy

This section describes how to configure an ECC probe over JDBC SSL.


To configure an SSL policy

- 1 Access the Policy window to create a new policy. The Policy window opens:




- 2 In the **Name** field, enter a policy name.
- 3 Check the **Use Protocol** check box and then select **JDBC-SSL** in the adjacent drop-down list.

- 4
- Add the ECC default credentials to the policy, as shown below, and then click **OK**.



- 5
- Access the Probe Properties window, select the newly created ECC over SSL policy, and click **OK**.



Copying the wallet file from the ECC Repository server

Because SSL authentication is used, an Oracle wallet file that holds the server authorization keys is required.

To copy the wallet file from the ECC Repository server

- 1 Copy the cwallet.sso file from the ECC Repository server using the following path on the ECC Repository server:

ECC_install_root_path\dbSafe\cwallet.sso


For example, if ECC is installed under E:\ECC, the wallet file location is E:\ECC\dbSafe\cwallet.sso.

- 2 Create a wallet directory in the DRA Home folder using the following path:

C:\Program Files\Symantec\Disaster Recovery Advisor\Data\security\wallets\ECC_Repository_IP\cwallet.sso

Let's look at an example, in which the ECC IP address is 192.168.68.57. Do the following in the order presented:

- Create the necessary directory.



- Copy the previously downloaded cwallet.sso file.

