

Veritas[™] Disaster Recovery Advisor User's Guide

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

6.1.1

Veritas Disaster Recovery Advisor User's Guide

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 2

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

About this guide

The Veritas Disaster Recovery Advisor (DRA) User's Guide is for IT staff, system integrators, services providers, and resellers who want to assess and optimize their data protection and also detect and handle disaster recovery problems, also known as “gaps.”

This guide includes the following chapters:

- [Chapter 1, “Introducing Veritas Disaster Recovery Advisor”](#) on page 17 summarizes DRA's functionality and features.
- [Chapter 2, “Getting started with DRA”](#) on page 21 describes the typical workflow for using DRA.
- [Chapter 3, “Installing or upgrading DRA”](#) on page 31 describes how to deploy or upgrade DRA.
- [Chapter 4, “A quick tour of the DRA interface”](#) on page 49 describes the DRA user interface and its main modules.
- [Chapter 5, “Using the Dashboard”](#) on page 63 describes how to view details about the last full system scan using the Dashboard.
- [Chapter 6, “Configuring DRA”](#) on page 69 describes how to configure DRA for your environment.
- [Chapter 7, “Using the Tickets module”](#) on page 255 describes the Tickets module and how to use it to investigate and resolve gaps in your environment.
- [Chapter 8, “Using the Topology module”](#) on page 279 describes the Topology module and its operation.
- [Chapter 9, “DRA reporting”](#) on page 317 describes how to generate and view DRA reports.
- [Chapter 10, “Using the Comparison module”](#) on page 353 describes the Comparison module and its operation.

- [Chapter 11, “Viewing DRA SLA information”](#) on page 371 describes how to view SLA information for business entities.
- [Glossary](#) on page 377 defines terms and acronyms used in this guide.

Note: Screens/windows in this document are examples only.

Contents

Chapter 1	Introducing Veritas Disaster Recovery Advisor	
	What is DRA?	17
	DRA benefits	18
	DRA components	19
	IT Discovery and Scanning engine	19
	Data Dependency engine	19
	Gap Analysis engine	20
	Gap and risk detection	20
	Report generator	20
Chapter 2	Getting started with DRA	
	Installing DRA	22
	Configuring DRA	25
	Full scan cycle	26
	Data discovery and collection	27
	Gap analysis	27
	Viewing, investigating, and tuning	28
	Monitoring the Dashboard	28
	Handling a ticket	28
	Topology module	29
	Reports	29
	Ticket lifecycle	29
	Automatic opening and closing of tickets	29
	Handling an open ticket	30
Chapter 3	Installing or upgrading DRA	
	Installation procedure	31
	Pre-installation requirements	31
	Preparing the computer for installation or upgrade	32
	Installing or upgrading DRA	32
	Logging in	42
	Uninstalling	43
	Installing a Collector	43
	Security considerations	46
	DRA security	46

User management	46
Oracle database security	46
Server hardening	47
Data collection	47

Chapter 4 A quick tour of the DRA interface

Main window	50
Dashboard	51
Topology module	51
Tickets module	53
Reports module	54
SLA module	55
Comparison module	56
Configuration module	57
Session pane	58
Controlling the window display layout	60
Controlling table displays	61

Chapter 5 Using the Dashboard

Viewing Dashboard information	64
Scan Coverage Summary area	64
Scan Coverage Details area	65
Collector Status area	65
Ticket Summary area	66
Ticket Tabs area	66
Business Summary area	67

Chapter 6 Configuring DRA

DRA probes	70
Configuration Wizard	71
Management console definition – configuring EMC ECC, Hitachi Data Systems (HDS) HiCommand, IBM HMC, and VMware vCenter	72
Configuring EMC ECC consoles	72
Configuring HDS HiCommand consoles	72
Configuring IBM HMC consoles	73
Configuring VMware vCenter consoles	73
Deleting a management console	78
Adding hosts	78
Storage discovery – configuring EMC SYMCLI, EMC NaviCLI, IBM DSCLI, IBM XCLI, IBM SVC, Hitachi HiCommand, and NetApp proxies	81
Configuring EMC SYMCLI proxies	81

Configuring EMC NaviCLI proxies	82
Configuring IBM DSCLI proxies	85
Configuring IBM XCLI proxies	85
Configuring SVC proxies	86
Configuring Hitachi HiCommand proxies	86
Configuring NetApp proxies	86
Site definition – configuring sites	91
Adding a site from a CSV file	94
Hosts groups and policies – configuring hosts	95
Managing host policies	100
Managing host groups	101
Adding databases	102
Database groups and policies – configuring databases	105
Managing database policies	107
Managing database groups	108
Managing automatic association of databases to groups and policies	110
Viewing the status of an entity	113
Configuring distributed collection	115
Configuring standby definitions	119
Adding a standby definition from a CSV file	122
Configuring standby policies	123
Attaching a standby policy	125
Configuring high availability (HA) clusters	126
Scheduling	129
Scan Scheduling	129
Scheduling a scan task	132
Scheduling a data analysis task	137
Scheduling full cycle scans	139
Report scheduling	144
Adding a new email report schedule	145
Email settings	147
Configuring policies	148
Adding a policy	148
Editing a policy	151
Configuring credentials	151
Configuring a proxy	156
Configuring sudo	157
Configuring business entities	158
Defining business entity types	168
Defining business entity roles	170
Defining SLA policies	172
Configuring expansion packages	179

Defining data collection packages	180
Defining host comparison expansion packages	193
Defining risk detection expansion packages	198
Defining reporting packages	208
System log	211
Scan Status	212
Scan troubleshooting	213
Configuring advanced settings	217
Configuring user management settings	217
Configuring user roles	220
Configuring external ticketing	223
Configuring an external ticketing connector	224
Managing labels	228
Configuring gap management settings	229
Configuring licenses	231
Viewing license details	233
Configuring system properties	235
Active directory	236
Additional collection	237
Aggregation	237
Automatic import	237
Collection	239
Collection timeouts	241
Command line interface	243
Comparison	243
Configuration	244
Cyber-Ark configuration	244
Dashboard	245
Distributed collection	246
Email settings	246
External ticketing system integration	247
Full Cycle	247
Gap Rules	247
General	247
Housekeeping	248
Maintenance	249
Reports	249
Rule Engine	250
SLA	251
System API	251
Tickets module	252
Configuring system tools	252
Using the support tool	253

	System API views	253
	Using the analyzer	253
	Restarting DRA	253
Chapter 7	Using the Tickets module	
	Overview	255
	Searching for a ticket	258
	Tickets in the Topology module	266
	Performing operations on tickets	268
	Adding a note to a ticket	270
	Adding a rating to a ticket	271
	Generating a unique ID for a ticket	272
	Sending tickets directly from the application	272
	Exporting tickets directly from the Tickets module	273
	Performing operations on multiple tickets	274
	Assigning a label to a ticket	277
Chapter 8	Using the Topology module	
	Overview	280
	Topology graphical view	281
	Topology details tabs	283
	Topology toolbar	285
	Navigator	292
	Topology tree	293
	Accessing the Topology module	294
	Opening the Topology module from the Tickets module	294
	Opening the Topology module from the Dashboard	295
	DRA object model	295
	Item display methods in the graphical view	298
	Connection display methods	299
	Direct connections	299
	Indirect connections	299
	Searching the topology	300
	Performing Topology module operations	303
	Filter displayed content	304
	Hide Selected Entities	305
	Expand Entity	306
	Collapse Selected Entities	307
	Expand Selected Entities	308
	Display Connected Entities	309
	Display Storage	310
	Display Replications	311

Displaying item details in the Topology module	311
Properties tab	312
Related Tickets tab	313
Connections tab	313
History tab	314

Chapter 9

DRA reporting

Overview	317
Exporting reports	319
Report types	319
Scan Status report	321
System Event Log report	322
Ticket Details report	323
Optimization report	324
Structure of the report	325
Unused storage grouping criteria	325
Summary of storage capacity per criterion	325
Distribution analysis of unused storage by array	327
Distribution analysis of unused storage by host/cluster	328
Unreplicated Data on Replicated Hosts report	329
NetApp Filer Replication Summary report	330
Unsynchronized Remote Replication report	331
Old Replicas report	332
Standby Pairs report	333
New Databases Notification report	334
Host Configuration Differences report	335
Infrastructure Change Log report	336
Database Replica Synchronization report	337
Unscanned Servers/Arrays report	338
Hosts with a Single SAN I/O Path report	339
Storage Summary report	340
Scan Troubleshooting report	341
Database Storage Utilization report	342
VMware Summary report	343
Host HBA Comparison report	344
MS SQL Server Backup Summary report	345
Replicated Temporary Databases report	346
Business Entity Scan Coverage report	347
License Usage report	349
Data Collection Expansion Package Summary report	350
Data Collectors Summary report	350
Installed Software report	351
Custom report	351

Chapter 10	Using the Comparison module	
	Overview	353
	Configuring Worksheets	355
	Configuring Comparison Groups	356
	Comparing Host Configurations	366
Chapter 11	Viewing DRA SLA information	
	Overview	371
	Using the SLA window	372
Glossary		377
Index		385

Introducing Veritas Disaster Recovery Advisor

This chapter includes the following topics:

- [“What is DRA?”](#) on page 17 describes the DRA application, its technology, and key benefits.
- [“DRA benefits”](#) on page 18 describes the benefits of using DRA for data protection risk assessment.
- [“DRA components”](#) on page 19 describes DRA’s software components.

What is DRA?

DRA is a data protection risk assessment solution that lets you diagnose disaster recovery problems (also called gaps) and optimize data protection. DRA enables enterprises to effectively manage business continuity implementations to ensure that critical business data is protected. DRA automatically detects and alerts you to any potential gaps, best practice violations, or service level agreement (SLA) breaches.

DRA is an agentless enterprise discovery and monitoring tool that automatically scans your infrastructure and detects gaps and infrastructure vulnerabilities in your disaster recovery (DR) implementation.

DRA gathers information about your environment and does the following:

- Provides automated insight into your data replication environment to provide an online, detailed, and up-to-date disaster recovery topology.
- Automatically detects and analyzes gaps and unprotected production areas using a signature knowledge base of over 5,000 signatures.
- Discovers the current data protection status of your critical applications and compares it to the state needed to comply with disaster recovery SLAs.

DRA uses this information to provide the following:

- Detailed recommendations on how you can improve your environment, based on best practices and recovery objectives.
- Detailed lists and information about current data protection and disaster recovery risks and the prioritized actions for fixing them. DRA also provides a variety of tools that let you drill down and analyze your environment using detailed tables and topology maps. You can use this information to fix the problems that DRA detects.
- Auditing and compliance documentation, including a map of your production environment, disaster recovery configuration, and dependencies.

DRA benefits

DRA offers the following benefits:

More protection	DRA is an intelligent automated tool. It can detect a level of detail that far surpasses any manual process. DRA lets you exponentially increase your level of protection and recovery. DRA detects known, unknown, and suspected vulnerabilities in your environment.
Timely protection	Because a manual audit is so complex and time-consuming, it is only done periodically. Your systems remain exposed and risks can remain undetected for a long time. After the initial deployment, DRA constantly monitors your environment and business continuity configurations to automatically detect and alert you of any potential issues as they occur.
Efficiency	DRA generates risk reports much faster than a manual process and initially only involves a few hours of time by your staff. Later, the time required is negligible.
Budget justification	IT personnel may suspect that their configurations do not meet business recovery service levels, but they cannot support this suspicion with hard facts. DRA clearly articulates the current environment status and business service levels that are not met. DRA's reports let you identify the resources you need to achieve the required level of protection. They also give you all the information you need to prepare for a meeting with management to discuss your budget.

DRA components

The DRA technology consists of the following technical components.

IT Discovery and Scanning engine

Based on a schedule you define, DRA's IT Discovery and Scanning engine periodically scans all the hosts, storage arrays, and databases in your environment. DRA is an agentless technology; you do not have to install any software or hardware throughout your network.

For example, DRA collects information from host servers using protocols, such as Windows Management Instrumentation (WMI), secure shell (SSH), and Telnet. It logs into each host with the credentials that you supply and extracts information about the host (its configuration, storage asset details, and installed databases). It stores the information in the DRA database. DRA can also connect to storage resource management software, such as EMC Control Center (ECC), to gather information about storage elements and their replication environment.

Data Dependency engine

DRA's Data Dependency engine builds an object model of your data and its behavior. The information compiled in this process is subsequently used to analyze gaps.

The DRA Data Dependency engine automatically discovers which data is protected, its actual protection service levels, and which data areas are not protected. This proprietary engine identifies the dependencies between servers, storage devices, their databases, and their replications in the production environment.

DRA then automatically verifies the data recovery protection side to check if there is an identical relationship that ties back to the server from the production environment.

This technology also identifies the actual protection levels, which are then tied back to the business services protection to identify service-level violations.

Gap Analysis engine

Symantec has compiled a database of over 5,000 potential gaps. DRA's proprietary Gap Analysis engine uses this information to automatically analyze the data map compiled by the DRA Data Dependency engine. The DRA Gap Analysis engine compares its extensive database of potential gaps and best practice signatures with the object models compiled from your enterprise, identifies any violation patterns, and recommends solutions.

DRA Gap Analysis engine also verifies customer-specific SLAs. For example, it can identify areas that are either under-protected or too aggressively protected, or identify components of your disaster recovery environment that cannot provide the required level of performance or availability.

Gap and risk detection

The DRA Gap Analysis engine detects the following types of gaps and risks, and gives you suggestions about how to fix them:


Data completeness	Incomplete data comparison between source and standby systems
Data inconsistency	Inconsistent data comparison between source and standby systems
Data accessibility and data path problems	Incorrect mapping of data or applications, or zoning and masking problems, which prevent standby data access, and so on
Incorrect deployment or maintenance actions	Incorrect deployment or incorrectly performed maintenance actions
Data tampering	Evidence of inappropriate modification of data copies at the DR site
General faults	Specific or sporadic problems, including events out of their recommended sequence, inefficient practices, and so on
Saving opportunities	Inefficient resource usage, not called for by the required corporate Business Continuity SLA

Report generator

DRA's automatic report generator automatically produces detailed reports describing your configuration and the gaps that it detected. These reports are more detailed than any manual process.


Getting started with DRA

This chapter presents a quick tour of how to work with DRA.



Installing DRA

As shown in the following illustration, DRA is installed on a dedicated server:



- A DRA installation consists of 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).

The DRA server configuration has the following requirements:

Table 2-1 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 Tomcat 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, take particular care to reserve the required CPU and memory.

The recommended collector configuration is as follows:

Table 2-2 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, take particular care to reserve the required CPU and memory.

Configuring DRA

After you install DRA, you must configure the elements it uses and the schedule for running the discovery process. This is done in the Configuration module, which is accessed by selecting the Configuration tab, as shown in the following figure:



Note: Configuration operations can only be performed by Admin users.

Use the configuration wizard to configure the following elements:

- Management Console Definition
See [“Management console definition – configuring EMC ECC, Hitachi Data Systems \(HDS\) HiCommand, IBM HMC, and VMware vCenter”](#) on page 72.
- Storage Discovery
See [“Storage discovery – configuring EMC SYMCLI, EMC NaviCLI, IBM DSCli, IBM XCLI, IBM SVC, Hitachi HiCommand, and NetApp proxies”](#) on page 81.
- Site Definition
See [“Site definition – configuring sites”](#) on page 91.
- Host Groups and Policies
See [“Hosts groups and policies – configuring hosts”](#) on page 95.
- Database Groups and Policies
See [“Database groups and policies – configuring databases”](#) on page 105.

You may also configure the following elements, which are accessed in the tree structure in the Configuration window's left pane:

- Standby Definition
See [“Configuring distributed collection”](#) on page 115.
- HA Cluster Definition
See [“Configuring high availability \(HA\) clusters”](#) on page 126.
- Scheduling
See [“Scheduling”](#) on page 129.
- Policies
See [“Configuring policies”](#) on page 148.
- Business Entities
See [“Configuring business entities”](#) on page 158.
- SLA Definition
See [“Defining SLA policies”](#) on page 172.
- Expansion Packages
See [“Configuring expansion packages”](#) on page 179.
- System Log
See [“System log”](#) on page 211.
- Advanced Settings
See [“Configuring advanced settings”](#) on page 217.

For more information about the configuration process, see [Chapter 6, “Configuring DRA”](#) on page 69.

Full scan cycle

Based on a schedule you define, DRA periodically scans the hosts and storage arrays in your environment. The default is once a day.

DRA is an agentless technology; you do not have to install any hardware or software in your environment.

After DRA performs its first full cycle, it regularly repeats the scan, based on your scheduling criteria.

Data discovery and collection

DRA investigates your environment and data protection configuration and collects information about IT assets, business applications, configurations, current SLA, and additional customer-specific policies.

DRA collects the information that you configured from each of its storage and host probes using the credentials that you defined. The collection phase of the scanning process involves connecting to the devices to be scanned, collecting data from them, and then closing the connections. DRA scans probe-by-probe, accessing all the instances of a probe type. For example, it first handles all EMC SYMCLI probe instances, then all NetApp probe instances, and then all EMC ECC probe instances.

After DRA collects data during the scanning process, it reviews the collected data and applies a set of enhancement rules. These enhancement rules add logical connections between items and properties of items in your environment. DRA uses these logical connections, but they do not physically exist in your environment.

For example, if your environment contains entities A, B, and C, where C is a replica of B and B is a replica of A, DRA applies an enhancement rule that logically connects A and C, even though they are not physically connected in the environment.

After the actual assets of your configuration are detected, DRA maps and models the dependencies between these IT elements. DRA also discovers their data protection and disaster recovery configurations.

Gap analysis

After DRA collects and enhances the data obtained during discovery, it analyzes the gaps in your environment.

As part of this process, DRA issues tickets for each gap that it finds. A gap is a definition of a particular problem type. A ticket is a unique instance of that problem found in your environment. It specifies the problem and the environment entities associated with that problem.

Symantec has compiled a database of hundreds of potential gaps, which is used to automatically analyze the data map compiled by the DRA Data Dependency engine. The DRA Gap Analysis engine compares its extensive database of potential gaps and best-practice signatures with the data models compiled from your enterprise to identify violation patterns and to recommend solutions. It also verifies customer-specific, service-level requirements.

For example, DRA can identify critical data that is not protected properly, standby hosts that do not have access to the replicated data, SLAs that are too strict or too lenient, and so on.

Viewing, investigating, and tuning

The last phase in the DRA workflow, and the one in which you are most involved, is viewing the information that was collected and analyzed, and fixing the problems in your environment.

DRA has a variety of viewing and investigation tools to help you analyze and fix problems. These include many tables of linked information and a graphical topology environment map that lets you quickly see the items in your environment.

DRA has extensive drill-down capabilities for quick navigation, so you can easily examine how items relate to one another.

You can also use queries to search your environment and drill down to an item's children and subcomponents.

For more information, see the following:

- [Chapter 7, “Using the Tickets module”](#) on page 255.
 - [Chapter 8, “Using the Topology module”](#) on page 279.
- [“Monitoring the Dashboard”](#) describes one option for investigating a gap.

Monitoring the Dashboard

The Dashboard page summarizes the information collected during the most recent full scan of the environment and includes details about any problems that may have occurred. It also lists open tickets, beginning with the most severe.

For more information, see [“Dashboard”](#) on page 51.

Click a ticket in the Dashboard to open it in the Topology module, as described in [“Handling a ticket.”](#)

Handling a ticket

The Topology module lets you view details, including lists and a graphical object topology of the tickets that DRA detects. It shows the object model of your environment relevant to the detected gap, identifying the hosts and storage arrays involved in the detected data protection gap.

The Topology module describes the gap in detail, its impact on your environment, and suggests ways to fix it.

For more information, see [Chapter 7, “Using the Tickets module”](#) on page 255 and [Chapter 8, “Using the Topology module”](#) on page 279.

Topology module

The Topology module shows information about the scanned object model of your environment (hosts and storage arrays, their relationships, and so on) and provides extensive drill-down capabilities. For each item, it also shows all open gaps and relevant information about each object.

Reports

DRA's Reports module lets you generate and view system reports.

For more information, see [Chapter 9, “DRA reporting”](#) on page 317.

Ticket lifecycle

[Table 2-3](#) lists the states that a ticket goes through in its lifecycle.

Table 2-3 Ticket lifecycle

Current states	New states	Event description
None	Open	The system found a ticket.
Open	Closed	The system can no longer find the ticket.
Open	Suppressed	The user marked this ticket as suppressed. Note: Suppressing the gap does not change the ticket status.
Open	Solved	The user marked the ticket as solved.
Solved	Reopen	The system reopened the ticket.
Solved	Closed	The system can no longer find the ticket.
Closed	Reopen	The system reopened the ticket.

Automatic opening and closing of tickets

The simplest case of a ticket is as follows:

- During a full scan, DRA detects the ticket and opens it.
- If DRA does not detect the ticket in the next full scan (for example, because you resolved the issue), DRA closes it.
- During the next full scan, if DRA detects the problem again, it reopens the ticket that you marked resolved. The ticket returns to the open state and shows the time and date that it was first detected.

Handling an open ticket

You can handle an open ticket in the following ways:

Resolve the ticket	You can mark it as resolved, which closes it before the next scan. If the gap has actually been resolved, the ticket is marked as closed during the next full scan.
Suppress the ticket	You can suppress a ticket, because it does not impact your environment, and you do not want it to keep appearing. This suppresses a specific gap for a specific host or storage element in your organization. You can suppress a ticket permanently or from the current date to a specific date in the future.
Suppress the gap	You can suppress a gap, because it does not impact your environment, and you do not want it to keep appearing. This also suppresses any specific tickets opened for all the hosts or storage elements in your organization.

Installing or upgrading DRA

This chapter includes the following topics:

- [“Installation procedure”](#) on page 31 describes pre-installation requirements and provides a procedure for installing or upgrading DRA.
- [“Security considerations”](#) on page 46 provides complete security guidelines and instructions for DRA, including user management, Oracle database security, server hardening, and data collection.

Installation procedure

The following provides complete DRA installation instructions, including pre-installation requirements, followed by detailed guidelines for preparing your computer for DRA's installation and itemized instructions for installation.

Pre-installation requirements

For more information about pre-installation requirements, see *Veritas Disaster Recovery Advisor Support Requirements*.

Preparing the computer for installation or upgrade

Before you install or upgrade DRA, if Apache Tomcat is installed on your computer, stop the Tomcat server.

To stop the Apache Tomcat server

- 1 Do one of the following:
 - In the Tomcat Monitor, click **Stop**.
 - On the command line, enter the following:
net stop tomcat6

Note: If you need to start the Tomcat server, enter
`net start tomcat6` on the command line.

- 2 If you have DRA installed on your computer and want to perform an upgrade, close any Web browser logged to it.

Note: It is assumed that an Oracle 11g database is installed on a computer with an accessible path that must be specified during installation. If not, install the Oracle 11g database before installing DRA. This database should be dedicated to DRA.

Installing or upgrading DRA

This section describes the default procedure for installing DRA and the following required software: JRE 1.6u20 and Tomcat 6.0 (32- or 64-bit versions) on drive C. Click Next to progress through the installation wizard. This section only describes important windows or ones on which you enter data.


The procedure for upgrading DRA is almost identical to that for installing DRA. To upgrade DRA to the latest version, run the installer and then follow the instructions described in “[To install DRA](#)” on page 32.

To install DRA

- 1 Run the `DRA_6.1.1.exe` file to display the first DRA installation wizard window.


Note: When you launch the executable file, the wizard may display a message that your computer is running an unsupported operating system. If this message is displayed, read it carefully and decide whether you want to continue.

- 2 Follow the displayed instructions and click **Next** to progress through the steps. The following window opens:




- 3 Do one of the following:
- If you want to install Java Runtime Environment (JRE) 1.6, select **Yes, install JRE 1.6 and use it when running Veritas Disaster Recovery Advisor**. When you select this option, a wizard opens to guide you through the JRE 1.6 installation process.
 - If JRE 1.6 is already installed on the designated path, select **No, use the currently installed JRE when running Veritas Disaster Recovery Advisor**.

- 4 Click **Next**. The following window opens:



- 5 Do one of the following:
- If you want to install Apache Tomcat 6.0, select **Yes, install Tomcat 6.0 and use it when running Veritas Disaster Recovery Advisor**. When you select this option, a wizard opens to guide you through the Apache Tomcat installation process.
 - If Apache Tomcat 6.0 is already installed on the designated path, select **No, install Veritas Disaster Recovery Advisor under the current Tomcat installation**.

- 6 Click **Next**. The following window opens if Windows Remote Management (WinRM) is not installed:




- 7 If you intend to use WinRM to scan Windows hosts, select the **Yes** radio button. When you select this option, a wizard opens to guide you through the WinRM installation process.
- 8 Click **Next**.
If an earlier version of DRA is installed, an upgrade window opens. If you want to upgrade, select the **Yes** radio button and click **Next**.
If no previous DRA version is installed, the window in [step 9](#) on page 36 opens.


- 9 Click **Next**. The following window opens:



- 10 Select **I accept the agreement** and click **Next** to display the following window:




- 11 Select **I accept this agreement** and click **Next** to display the following window:



The screenshot shows a Windows-style window titled "Setup - Veritas Disaster Recovery Advisor". The main heading is "Database information". Below it, a message states: "Please enter the required database information. The installer will not create the database. Create it before proceeding." There is a yellow gear icon in the top right corner. The form contains three input fields: "Database SID:" with the value "qddb", "Port number:" with the value "1521", and "Server name: (Modify only if using remote database)" with the value "localhost". At the bottom left is a URL "http://www.symantec.com". At the bottom right are three buttons: "< Back", "Next >", and "Cancel".


- 12 Fill in the fields and click **Next**. The following window opens:



The screenshot shows the same window titled "Setup - Veritas Disaster Recovery Advisor", but the main heading is "Schema information". The message below it says: "Please enter the name of an existing schema or of a new schema to be created. You can create this new schema in the next step." There is a yellow gear icon in the top right corner. The form contains three input fields: "Schema name:" with the value "rguser", "Schema password:" with masked characters "*****", and "Confirm schema password:" with masked characters "*****". At the bottom left is a URL "http://www.symantec.com". At the bottom right are three buttons: "< Back", "Next >", and "Cancel".

Note: If you are running an upgrade, continue with [step 20](#) on page 41.

- 13 Specify the database schema name and password and click **Next**. The following window opens:




- 14 This window deals with creating the schema required by DRA. Do one of the following:

- To have DRA create the schema automatically, select **Yes, create the required schema on this computer**.
- To have the database owner create the schema, select **No, show me the SQL to execute** and execute the following text:

```
DROP USER drauser CASCADE;  
CREATE USER drauser PROFILE "DEFAULT" IDENTIFIED BY "<password>"  
ACCOUNT UNLOCK;  
GRANT UNLIMITED TABLESPACE TO drauser;  
GRANT "DBA" TO drauser;  
QUIT;
```

- 15 Click **Next**. The following window opens:



Setup - Veritas Disaster Recovery Advisor

DBA Login Details
Please enter the database's SYSDBA user name and password to create the required schema.

SYSDBA User Name:
sys


SYSDBA Password:
••••••••

Confirm Password:
••••••••

<http://www.symantec.com>

< Back Next > Cancel

- 16 Enter the username and password for the DRA database administrator and click **Next**. The following window opens:



Setup - Veritas Disaster Recovery Advisor

DRA User configuration
Please enter the required DRA administrator and user passwords

User password:
••••••••

Confirm user password:
••••••••


Admin password:
••••••••

Confirm admin password:
••••~•••

<http://www.symantec.com>


< Back Next > Cancel

- 17 Specify the password for the user and admin DRA users and click **Next**. The following window opens:




- 18 Specify where to store the DRA program shortcuts by clicking **Browse** and navigating to the required folder.

19 Click **Next**. The following window opens:




20 Specify whether or not to launch Tomcat during installation. Select or deselect the **Start Apache Tomcat 6.0** check box.

21 Click **Next**. If you are installing DRA, rather than upgrading, the following window opens:



- 22 After specifying your installation settings, verify your installation selections and click **Install** to install DRA. At the end of the installation process, the following window opens:



- 23 Click **Finish** to complete the installation.

Logging in

After you install DRA, you can log in to the application.

To log in to DRA

Enter the following string:

`http://server_name:Tomcat_port/DRA`

For example:

`http://localhost:8080/DRA`

Alternatively, click the DRA desktop icon or access the application from the Start menu.

Uninstalling

DRA can be easily uninstalled, when necessary.

To uninstall DRA

Select **Start > Programs > Symantec > Disaster Recovery Advisor > Uninstall Veritas Disaster Recovery Advisor**. This removes DRA from the machine. Note that the database and the data folder that were created during the installation are not deleted.

Installing a Collector

This procedure describes how to install a collector for use by DRA to scan an environment, as described in [“Configuring distributed collection”](#) on page 115.

Java is required to install a collector. If it is not already installed, use the `jre-6u22-windows-x64.exe` file provided by Symantec to install Java.

Defining a collector for DRA involves the following steps:


- 1 Download the collector installation file, as described in [“Configuring distributed collection”](#) on page 115.
- 2 Install the collector on the machine on which it will run, as described below.
- 3 Use the **Add New Collector** button to configure the collector in DRA, as described in [“Configuring distributed collection”](#) on page 115.

To install a collector


- 1 Extract the contents of the DRACollector-6.1.0.0.zip collector installation file.

Note: The name of the directory in which the zip file is extracted cannot contain spaces.


- 2 Run the setup.bat file. The following window opens:




- 3 Click **Next**. The following window opens:



- 4 Specify the Java home directory and click **Next**. The following window opens:



- 5 Specify the connectivity port number through which the collector will communicate with the DRA server, and the path where the collector will be installed. The default port is 27000.
- 6 Check the **Start DRA Collector after installation is finished** check box to launch DRA after collector installation.
- 7 Click **Next**. The following window opens:



- 8 Click **Finish** to complete collector installation. You must now configure sites for this collector, as described in “[Configuring distributed collection](#)” on page 115.

Security considerations

This section describes various aspects of DRA security relating to the various users that can access the system, the Oracle database that DRA uses, the server on which the product is installed, and the ways data is collected and stored in DRA.

DRA security

User management

DRA has two kinds of default users: user and admin. The admin username should only be used to administer the application. The user username should perform routine tasks or add other users, as needed.

The user username has user privileges and can do the following:

- Use the Dashboard, Reports, Tickets, and Topology modules
- Search and view tickets
- Change the ticket lifecycle to resolve or suppress types/gaps
- Change the topology browser items that are displayed
- Generate reports

The admin username has administrator privileges. It has all the privileges of the user account, and can also do the following:

- Add items and connections
- View and change the configuration

Both the admin and user usernames can access DRA through the web.

The password for both users is set during the installation. Keep the password in a safe place. If you forget it, contact technical support.

The system log records login operations that have been done for audit.

Oracle database security

DRA uses an Oracle database. The schema name, user name, and password are all set during installation.

The Oracle database stores all the configuration and collected data. All sensitive data (credentials) are stored encrypted using a 128-bit encryption key.

Server hardening

The DRA server is hardened. This includes the following:

- All unnecessary services are disabled.
- Only Web ports are opened for incoming communication.
- The following ports are opened for outgoing communication:
 - Oracle listener port
 - Java Database Connectivity (JDBC)
 - SSH
 - Telnet
 - Remote shell (rsh)
 - WMI

Note: If you enable Automatic Updates, you must use the required FTP port. For more information, see “[Automatic import](#)” on page 237.

Data collection

Which data is collected and how – general overview

DRA uses an agentless solution to collect data. You do not have to install anything on the scanned servers. There is no footprint on the servers themselves during the collection process.

DRA collects data by connecting to the external machine using a standard protocol and standard commands.

Executed commands during data collection

When you configure DRA, you determine the data collection commands it runs. Before each command is run, it is validated to make sure it has not been tampered with since DRA was installed.

Neither of the DRA authenticated users (admin and user) can use Web access to change commands, either intentionally or unintentionally. All commands are read-only, so no changes can be made that cause external damage.

Some of the commands require root credentials to the server. However, it is possible to log in to the server as a non-privileged user and execute a privileged command using a sudo mechanism. For more information, as well as a list of commands, contact technical support.

Credentials

DRA uses an agentless solution to collect data. You do not need to install anything on the scanned servers. Instead, DRA communicates with the server to retrieve the required information.

Connecting with the server requires the setting of the correct credentials. The credential level depends on the data that is retrieved. It can be credentials with user privileges or with administrator privileges. In both cases, all credentials are stored encrypted in the Oracle database, which is password-protected. No one (not even the DRA administrator) can view these credentials. Only DRA uses them in an encrypted manner to authenticate against the server.

Collected data

DRA only collects configuration information and not the actual data. For example, DRA does not collect the data in a storage device, but only the configuration data (its volume, size, and other properties). It also collects configuration information for the data that is collected and stored in the Oracle database, which is password-protected.

A quick tour of the DRA interface

This chapter includes the following topics:

- [“Main window”](#) on page 50 describes the DRA main window.
- [“Dashboard”](#) on page 51 introduces DRA's Dashboard, which summarizes the information collected during the most recent full scan of your environment. It also lists any problems that may have occurred during the last scan.
- [“Topology module”](#) on page 51 introduces the Topology module and describes its operation.
- [“Tickets module”](#) on page 53 introduces the Tickets module, which lets you search for specific tickets or ticket types.
- [“Reports module”](#) on page 54 introduces the Reports module, which lets you view and generate system reports.
- [“SLA module”](#) on page 55 introduces the SLA module, which lets you view SLA information about business entities.
- [“Comparison module”](#) on page 56 introduces the Comparison module, which lets you compare the configurations of various hosts.
- [“Configuration module”](#) on page 57 introduces the Configuration module, which lets you configure all aspects of DRA for an environment.
- [“Session pane”](#) on page 58 describes the Session pane, which you can display or hide in the Topology, Tickets, and Reports modules.
- [“Controlling the window display layout”](#) on page 60 describes how you can control the window display in the Topology, Tickets, and Reports modules.
- [“Controlling table displays”](#) on page 61 describes how you can control table displays in DRA modules.


Main window

The DRA main window is shown below. After you log in, the Dashboard page, described on page 51, opens by default. DRA has the following modules, each of which is designated by a tab at the top of the window:

- Dashboard
- Topology module
- Tickets module
- Reports module
- SLA module
- Comparison module
- Configuration module

Note: The Configuration module is for use by Admin users only.

Click a tab to open its respective module.



The right side of the main window includes the following links.

[About](#) | [Suggest Gap](#) | [Help](#) | [FAQ](#) | [Logout](#)

Suggest Gap lets you give input to Symantec about potential disaster recovery problems.

Your suggestions help Symantec address the specific needs of your enterprise, and other companies benefit from a more comprehensive gap database.

Clicking Suggest Gap takes you to Symantec Connect, a secure site for Symantec customers and partners. To enter your comments and suggestions, you must have a SymAccount and register on the Symantec Connect. If you do not have a SymAccount, you will be prompted to create one.

Dashboard

The Dashboard page summarizes the information collected during the most recent full scan of your environment, and describes any problems that may have occurred.


For more information, see [Chapter 5, “Using the Dashboard”](#) on page 63.

Topology module

The Topology module opens the topology browser. The browser provides a map of your environment that lets you quickly see the items in it. You can search the objects in your environment by clicking on the objects in the browser, clicking the links in the table, or using queries to supplement your data analysis and problem solving.

When you perform a search, data that was collected during the last scan and then analyzed is retrieved from the DRA database. The parameters you specify for a search may vary, depending on the item(s) you specify in your query.

To open the Topology module, click the Topology tab. When you first open this module, it displays your environment’s high-level topology.



The Topology page has the following areas:

Session pane	Displays or hides topologies and search results
Graphical view	Displays the topology map, including the following: ■ Navigator Provides an easy way of navigating to the area of interest in the topology browser ■ Topology tree Provides access to all the items shown in the topology browser in a hierarchical tree
Detail tabs	Displays tables of detailed topology information at the bottom of the window in a series of tabs

You can also perform searches in the Topology module. After you perform a search query, the browser shows a graphical map of your environment relevant to the searched item. You can navigate to display additional item details. Click an item in the graphical or tabular views to drill down to the item's children and/or subcomponents.

The Topology module also has a series of operation buttons with more viewing and navigation options.


For more information on using the Topology module and its browser, see [Chapter 8, “Using the Topology module”](#) on page 279.

Tickets module

The Tickets module lets you filter the list of DRA tickets in the system to search for a specific ticket or ticket type. After DRA displays the query search results, you can drill down to see the ticket’s details, as well as view the ticket’s topology in the Topology module.

The ticket describes the problem, the items it affects, its impact on the system, and suggests how to fix it.

To open the Ticket Search page, click the Tickets tab.




For more information about using the Tickets module, see [Chapter 7, “Using the Tickets module”](#) on page 255.

Reports module

The Reports module lets you generate and view system reports. To open the Reports page, click the Reports tab.

Reports tab




For more information, see [Chapter 9, “DRA reporting”](#) on page 317.

SLA module

The SLA module lets you to view service level agreement data for business entities to determine which business entity has a problem with one of its resources. To open the SLA page, click the SLA tab.


SLA tab



For more information, see [Chapter 11, “Viewing DRA SLA information”](#) on page 371.

Comparison module

The Comparison module helps you identify the host configuration drifts hiding in your IT. Such drifts often fail cluster failover processes, and reduce the availability of your organization’s applications. In this module, you can create *comparison groups* that include hosts, clusters, or business entities, and easily track configuration differences between them. To open the Comparison page, click the Comparison tab.



For more information, see [“Using the Comparison module”](#) on page 353.

Configuration module


The Configuration module lets you configure all aspects of your environment, including sites, business entities, probes, scheduling, and advanced settings. Only the admin user can access this module.

For more information about DRA users, see “Configuring user management settings” on page 217.

To open the Configuration page, click the Configuration tab.

For more information, see Chapter 6, “Configuring DRA” on page 69.

Configuration tab



Session pane


The Session pane is displayed in the left pane of the Topology, Tickets, and Reports modules. This list is a useful tool for tracking the operations performed since the current session began, and helps you analyze and fix problems in the system.

The Session pane has the following sections:

Topologies	Includes item and ticket topologies viewed in the Topology module
Search Results	Includes the search results of ticket, topology and report search queries
Reports	Includes reports viewed in the Reports module

DRA automatically adds report items to the pane. You must manually add

Topologies and Search Results items using the Capture button  in the toolbar. If you need to, you can also remove elements from the Session pane.



The item names in the Session pane match the item labels in the Topology module graphical view.

The Session pane uses icons to identify the types of activities performed in the system. They include the following:



Item topology

Shows the topology in the context of a specific item, which was probably selected in the Topology Search window.

Note: The icon that displays to indicate an item in the topology depends on the type of item. The sample icon shown here represents a site.

For a description of all the item (object) icons used in the system, see “[DRA object model](#)” on page 295.



Ticket topology

Shows the topology for a specific ticket.



Topology search

Shows the Topology Search window.



Tickets search

Shows the Tickets Search window.



Report

Shows a specific report generated by the system.

To add an element to the Session pane

- ◆ To add an element, click the  button. If the element is a topology, it is added to the Topologies section in the Session pane. If the element is a search query result, it is added to the Search Results area of the pane.


Note: DRA automatically adds reports to the Session pane. You must manually add Topology and Search Results items to this pane using the  button.

To delete an element in the Session pane

- ◆ Next to the element you want to delete, click the  button. The element is deleted from the Session pane.


Controlling the window display layout

You can manually adjust the window layout for the Topology, Tickets, and Reports modules. Between the panes on the right, drag the separator bar up or down to the position you want.



You can also minimize and restore the two panes on the right (above and below the separator bar) using the Minimize/Restore buttons for each area.

For example, in the window above, to maximize the Ticket search area so that only it appears in the right pane, click  at the top of the Ticket area. After you maximize the window, it looks similar to the following:






Note: To automatically minimize and restore the Tickets area, click Ctrl+Space in the Tickets module.


Controlling table displays

You can adjust how tables are displayed in DRA modules using the following toolbar:



It contains the following tools:

	First page Displays the first page of the table.
	Previous page Displays the previous page of the table.
	Pages Displays the current page number out of the total number of pages in the table.
	Next page Displays the next page of the table.
	Last page Displays the last page of the table.
	Refresh Refreshes the table information. During the refresh, the icon spins.
	Show Indicates the number of items to display.
	Search Searches for the text you enter in the box in the Summary column of the table. Click  to begin the search. Click  to end it.



Regular Expression

Select the check box to search for a regular expression. Then, enter the expression in the adjacent Search field. You can use wildcards.



Export to Excel

Exports ticket summaries to Excel.

Using the Dashboard


This chapter includes the following topic:

- [“Viewing Dashboard information”](#) on page 64

Viewing Dashboard information

The Dashboard page summarizes the information collected during the most recent full scan of the environment and describes any problems that may have occurred. Information is refreshed every five minutes and is read-only.


After you log on, the Dashboard page opens by default. The top-left area of the window, just above the Dashboard, displays the name of the currently logged on user.



The Dashboard has the following areas, each of which summarizes a different type of information:

- “Scan Coverage Summary area” on page 64
- “Scan Coverage Details area” on page 65
- “Collector Status area” on page 65
- “Ticket Summary area” on page 66
- “Ticket Tabs area” on page 66
- “Business Summary area” on page 67

Scan Coverage Summary area



You may decide to use DRA to protect all your hosts, storage devices, and databases or only part of them.

This area summarizes the hosts, storage devices, databases, and clusters that are protected and indicates how many of them are not protected.

The section title bar displays the number of sites and business entities that have been scanned.

Each component is represented by a color-coded pie chart: one for hosts, one for storage devices, one for databases, and one for clusters. Green indicates how many items were scanned. Red indicates how many items were not scanned. Orange indicates how many items were partially scanned. Yellow indicates how many items were authenticated. Purple indicates how many items were down.

The ratio of scanned items to the total is shown above each pie chart. For example, above the Storage pie chart, shown above, “23/23” indicates that DRA scanned 23 of 23 storage devices.

Scan Coverage Details area

The Scan Coverage Details area shows information about the scan status of Hosts, Storage, Databases, and Clusters.

Scan Coverage Details							
Dashboard last update		Aug 08, 2012 2:51:14 PM					
Hosts		Storage		Databases		Clusters	
Scanned	116	Scanned	2	Scanned	0	Scanned	20
Not Scanned	87	Not Scanned	2	Not Scanned	6	Not Scanned	13
Authentication Errors	412	Authentication Errors	0	Authentication Errors	0	Partially Scanned	0
				Down			0


Collector Status area

This area displays the status of the data collector.

Collector Status		
Collector Name	Location	Status
Boston Collector	Boston	Down
Local Scan Server	Chicago	Idle

Ticket Summary area

The Ticket Summary area shows how many tickets were opened today, in the past week, and the total that are currently open. DRA groups tickets by category.



Ticket Tabs area

The Ticket Tabs area includes the following:

- **Top 5 Tickets Tab**
As shown in the following figure, this tab lists the most severe tickets in descending order. Generally, these are the tickets that you should handle first.

Top 5 Tickets			New 5 Tickets		
ID	Detection Date	Status	Summary	Business Entities	Severity
150	Sep 02, 2009 6:11:12 PM		A passive node in VCS VACLUSTER at site Boston is missing a required mount point directory	Billing	
68	Sep 02, 2009 6:08:18 PM		Passive node of VCS VACLUSTER at site Boston is not connected to all required storage volumes	Billing	
71	Sep 02, 2009 6:08:41 PM		Different replication methods for data and log files were detected on Oracle database db29 on Windows host nev1 at site Chicago	DWH	
154	Sep 10, 2009 11:16:29 PM		VG udbyg on AIX host subhost at site Boston is stored on Symmetrix devices that have partial replica sets	ERP	
125	Sep 02, 2009 6:10:04 PM		Suboptimal storage volume layout of multiplexed redo log files for Oracle database db24 on Solaris host mia at site Chicago		

- **New 5 Tickets**
This tab lists the most severe tickets that were detected in the last scan in descending order.

Business Summary area

The Business Summary area uses color codes to indicate ticket severity. DRA groups the open tickets for each business entity by Data Risk, Availability Risk, and Optimization (by super-category). You can click on any of the colored balls to view those types of tickets for the given business entity.


To use this section, you must configure business entities.

For more information, see [“Configuring business entities”](#) on page 158.

Business Summary					
Scan	Business Entity	Data Risk	Availability R...	Optimization	SLA
 100%	 Customers	 Total: 0 A:- M:-	 Total: 0 A:- M:-	 Total: 0 A:- M:-	 Total: 0 A:- M:-
 91%	 ERP Scanning issues: 3	 Total: 0 A:- M:-	 Total: 0 A:- M:-	 Total: 0 A:- M:-	 Total: 0 A:- M:-
 11%	 CRM Scanning issues: 411	 Total: 7 A:-2  M:-5 	 Total: 46 A:-33  M:-3 	 Total: 49 A:-49  M:-7 	 Total: 4 A:-4  M:-4 
 100%	 NTP	 Total: 0 A:- M:-	 Total: 0 A:- M:-	 Total: 0 A:- M:-	 Total: 0 A:- M:-
 5%	 SOA Scanning issues: 18	 Total: 0 A:-8  M:-8 	 Total: 7 A:-2  M:-8 	 Total: 7 A:-7  M:-	 Total: 0 A:- M:-
 52%	 Energy Scanning issues: 62	 Total: 7 A:-2  M:-5 	 Total: 43 A:-32  M:-4 	 Total: 49 A:-49  M:-7 	 Total: 4 A:-4  M:-4 

The Scan column shows a pie chart that indicates how much of the business entity has been scanned. The percentage below the pie chart indicates the percentage of the business entity scanned.

When you hover over a pie chart, a tooltip displays the number of hosts, databases, and storage volumes scanned during the current scan. This information can help you determine the reliability of the information in the columns to the right. For example, if you see that only 10% of a business entity has been scanned, you may not have a true indication of all the scanning issues for that business entity.



If you click on a pie chart, the Business Entity Scan Coverage report automatically opens for that business entity. For more details about this report, see [“Business Entity Scan Coverage report”](#) on page 347.

The Data Risk, Availability Risk, Optimization, and SLA columns represent different types of ticket categories. Each column shows a total, which is the number of tickets opened on the host for that business entity of that ticket category. Note that a ticket can be related to more than one ticket category (Data Risk, Availability Risk, and so on).

Scan	Business Entity	Data Risk	Availability R...	Optimization	SLA
 11%	 CRM Scanning issues: 410	 Total: 12 A: +12  M: -	 Total: 57 A: +55  M: +28 	 Total: 49 A: +49  M: +12 	 Total: 0 A: - M: -

Each ticket category column also contains an A entry and an M entry. In this context, A represents the number of tickets since the last analysis and M represents the number of tickets in the last 30 days. The arrows adjacent to the A and M entries indicate whether the tickets situation is improving or worsening for that ticket category. An arrow pointing upwards () indicates that the value in the column is higher than the previous scan. This means that the tickets situation is worsening in that ticket category for the business entity.



Configuring DRA

This chapter includes the following topics:

- [“DRA probes”](#) on page 70 describes the probes installed on your DRA server.
- [“Configuration Wizard”](#) on page 71 describes how to configure management consoles, storage proxies, hosts, databases, and sites in DRA.
- [“Configuring distributed collection”](#) on page 115 describes how to configure collectors in DRA.
- [“Configuring standby definitions”](#) on page 119 describes how to configure primary/standby host/cluster pairs.
- [“Configuring high availability \(HA\) clusters”](#) on page 126 describes how to configure HA clusters in DRA.
- [“Scheduling”](#) on page 129 describes how to configure DRA scanning and data analysis tasks and email report deliveries.
- [“Configuring policies”](#) on page 148 describes how to configure system policies.
- [“Configuring business entities”](#) on page 158 describes how to configure business entity groups, types, and roles.
- [“Configuring expansion packages”](#) on page 179 describes how to configure various types of expansion packages in DRA.
- [“System log”](#) on page 211 describes how to view the DRA system event log.
- [“Configuring advanced settings”](#) on page 217 describes how to manage and configure gap information, user passwords, system properties, external ticketing, ticket labels, and system tools.

DRA probes

You install DRA probes on the DRA server to collect data from the configured hosts and storage devices in your enterprise. DRA uses an agentless system, so you do not have to install any agents on endpoint devices.

DRA uses a suite of probes to retrieve information from different types of entities. For example, EMC ECC and EMC SYMCLI probes collect information from EMC arrays. The probes are created automatically in the background as you configure your environment in DRA; for example, when you configure EMC ECC as a management console in DRA, an ECC probe is created in the background. Each probe can have multiple instances.

Note: DRA is provided with a fixed set of probes. If you need another type of probe, contact Technical Support.

You do not have to configure a probe type under the following conditions:

- Your environment does not have that type of device.
- You do not want to analyze that type of device.

The Configuration Wizard lets you configure the following types of probes:

EMC ECC, HDS HiCommand, and VMware vCenter	See “Management console definition – configuring EMC ECC, Hitachi Data Systems (HDS) HiCommand, IBM HMC, and VMware vCenter” on page 72.
EMC SYMCLI, EMC NavCLI, and NetApp	See “Storage discovery – configuring EMC SYMCLI, EMC NaviCLI, IBM DSCLI, IBM XCLI, IBM SVC, Hitachi HiCommand, and NetApp proxies” on page 81.

Configuration Wizard

The Configuration Wizard lets you configure various system components, including management consoles, storage proxies, sites, and hosts. You should run the wizard the first time you run DRA and again whenever you need to update the configuration.

Complete the Configuration Wizard step-by-step, in the following order:

- [“Management console definition – configuring EMC ECC, Hitachi Data Systems \(HDS\) HiCommand, IBM HMC, and VMware vCenter”](#) on page 72
- [“Storage discovery – configuring EMC SYMCLI, EMC NaviCLI, IBM DSCLI, IBM XCLI, IBM SVC, Hitachi HiCommand, and NetApp proxies”](#) on page 81
- [“Site definition – configuring sites”](#) on page 91
- [“Hosts groups and policies – configuring hosts”](#) on page 95
- [“Database groups and policies – configuring databases”](#) on page 105

Note: Many of the pages in the Configuration Wizard contain a Status column, indicating the current state of the selected entity (host, management console, storage proxy, database, and so on). For more information about determining an entity's status, see [“Viewing the status of an entity”](#) on page 113.

Click the View Scanning Issues  icon in each step of the wizard to open the Scanning Issues window in a new window of your Internet browser. Click the

View System Event Log  icon to open the System Event Log in a new window.

On the bottom right of each wizard page, you can use the Next>> and <<Back buttons to move forward or backward through the configuration pages.

After you complete the Configuration Wizard, configure the other features, according to the nodes in the tree. Follow the instructions beginning with [“Configuring high availability \(HA\) clusters”](#) on page 126.

Management console definition – configuring EMC ECC, Hitachi Data Systems (HDS) HiCommand, IBM HMC, and VMware vCenter

The first step in the Configuration Wizard is configuring the management consoles. You can configure the following types of consoles:

- EMC Control Center (ECC)
See [“Configuring EMC ECC consoles”](#) on page 72.
- HDS HiCommand
See [“Configuring HDS HiCommand consoles”](#) on page 72.
- IBM HMC
See [“Configuring IBM HMC consoles”](#) on page 73.
- VMware vCenter
See [“Configuring VMware vCenter consoles”](#) on page 73.

Configuring EMC ECC consoles

EMC ECC consoles configure, maintain, and monitor EMC arrays along with other entities in your environment. DRA uses the ECC to gather information about hosts in your environment.

To configure an EMC ECC console, you must specify the IP address of the machine on which the ECC Repository is installed and the policy required to access this device (user/password).

A popular policy username is stsvie and password is sts (these are ECC Repository RAMBDB credentials). Most ECC Repository systems are configured with this default login.

Typically, you only need one EMC ECC instance. However, sometimes you may have more than one copy of the EMC ECC software running to control your environment. In this case, you may need to configure more than one EMC ECC console.

Configuring HDS HiCommand consoles

HDS HiCommand consoles configure, maintain, and monitor HDS arrays. DRA uses HiCommand consoles to gather information about hosts and HDS Storage Arrays in your environment.

To configure a HiCommand console, you must specify the IP address of the machine on which HiCommand is installed and the policy required to access this device (user/password).

Typically, you only need one HiCommand instance. However, sometimes you may have more than one copy of the HiCommand software running to control your environment. In this case, you may need to configure more than one HDS HiCommand.

Configuring IBM HMC consoles

IBM HMC consoles configure and maintain IBM-managed systems. DRA uses HMC consoles to gather information about LPARs VIO configuration in your environment.

To configure an HMC console, you must specify the IP address of the HMC and the policy required to access it (user/password).

Configuring VMware vCenter consoles

VMware vCenter consoles configure, maintain, and monitor VMware host systems (ESX, ESXi).

DRA uses vCenter consoles to gather information about virtual machines and host systems in your environment.

To configure a VMware vCenter console, you must specify the IP address of the vCenter server and the policy required to access this server (user/password).

Some environments have more than one VMware vCenter running. In this case, you may need to configure more than one VMware vCenter console.

You can perform the following operations from the Management Consoles page:

Add Management Console		Adds a new management console definition.
Edit Management Console		Modifies a management console's properties.
Delete Management Console		Deletes a management console.
Verify Connectivity		Note: If the management console has discovered hosts that were inserted into the system, the Delete Management Console Wizard opens. For more information, see “Deleting a management console” on page 78. Verifies if the management console is reachable by trying to log in to it using the specified policy. If it is successful, the Verified field in the table on the Host Discovery page is checked. (The Verified field can also be checked if the Discover Hosts and Execute Scan operations succeed.)

Discover Hosts 

Uses the management console to discover all the hosts in your environment. To display the hosts, click View Hosts.

Execute Scan 

Connects to the selected management console and collects information. When you press this button, the system runs a full scan on the selected EMC ECC, HiCommand, or vCenter console:

- For ECC, it gathers information about entities in your environment (hosts, arrays) that will be used for the Unscanned Servers/Arrays report.
- For HiCommand, it gathers a variety of information about HDS arrays (disks and replications), and about hosts in your environment.
- For vCenter, it gathers a variety of information about entities in your virtual environment (ESXs, virtual machines, and so on).

Usually, you should not run a full scan in the Configuration Wizard. Because a full scan may take some time to complete, you should only run one after the initialization process ends or during a full cycle.

View Hosts 

Lists the hosts that were discovered when the Discover Hosts button was clicked. View Hosts lets you manually add, edit, or delete a host.

For more information, see [“Adding hosts”](#) on page 78.

Add Host Manually 


Opens the Host window so you can manually add a new host.

For more information, see [“Adding hosts”](#) on page 78.


To configure an EMC ECC, HDS HiCommand, or VMware vCenter console

- 1 To access the Configuration page, click the Configuration tab.
- 2 In the tree hierarchy, click **Configuration Wizard**. The Host Discovery page opens.

Configuration Wizard



- 3 At the bottom of the page, click **Add Management Console** (). The following window opens:



For an EMC ECC console, the **Type** field value is always ECC.

For an HDS HiCommand console, the **Type** field value is always HiCommand.


For a VMware vCenter console, the **Type** field value is always vCenter.

- 4 In the **Description** field, enter a name or description for this management console.
- 5 In the **IP** field, enter the IP address for this management console.
- 6 In the **Policy** field, select a policy from the drop-down list. Initially, one default policy is defined, named ECC default policy, which enables you connect to an ECC.
To add a new policy, click the **Add Policy** button, which is to the right of the drop-down list.
For instructions on how to add a policy, see [“Adding a policy”](#) on page 148.

Note: Typically, the user name and password are the same in all installations. To change the user name or password, select the Policies branch in the tree on the Configuration page as described in [“Configuring policies”](#) on page 148.

- 7 In the **Site** field, select a site in the drop-down list.
- 8 To enable the scanning of the ECC/HiCommand/vCenter servers during the full system scan, select the **Enable** check box. If you do not select this check box, information from the management console is not collected during the next full system scan. To return to previous menu, click **OK**.
- 9 To save your settings and exit the window, click **OK**. The management console is added as a new row in the Management Consoles list.
- 10 To verify connectivity to the management console, click **Verify Connectivity** (). If it can connect, the **Verified** field in the table on the Host Discovery page is checked.
- 11 To use this management console to discover all the hosts that can be accessed from this management console, click **Discover Hosts** (.

To display the discovered hosts, click **View Hosts** (). The following window opens:




You can also manually add a host.

For more information, see [“Adding hosts”](#) on page 78.

- 12 Click **Next**. The Storage Discovery page is displayed, letting you configure EMC SYMCLI, EMC NaviCLI, Hitachi HiCommand, and NetApp proxies, as described in [“Storage discovery – configuring EMC SYMCLI, EMC NaviCLI, IBM DSCLI, IBM XCLI, IBM SVC, Hitachi HiCommand, and NetApp proxies”](#) on page 81.

Deleting a management console

To delete a management console, select the management console in the Management Consoles list and click Delete Management Console. If this management console was never scanned, it is deleted immediately. Otherwise, a deletion wizard opens to guide you through the safe removal of the management console.



Adding hosts

From the Management Console Definition page, you can select Discover Hosts to use the selected console to find all the hosts in your environment. Then, you can click View Hosts to list the discovered hosts. To connect to a specific host that does not appear in the list, you must add it in one of the following ways:


- Using the Add Host window
- Importing host information from a comma-separated values (CSV) file

Adding a host using the Add Host dialog box

The Add Host dialog box lets you add a host quickly and accurately by providing drop-down lists with predefined values for the Type, Policy, and Group values.

To add a host using the Add Hosts window

- 1 On the Management Console Definition page, click **Add Host** (). The following window opens:



Note: You can also access the window by clicking View Hosts () on the Management Console Definition page, and then in the Host window clicking Add Host.

- 2 Enter the following host details:

Name	The host name. Note: If you enter the FQDN of a host, DRA truncates it to the simple host name.
Type	The type of operating system on the host, such as Windows, HP-UX, Solaris, AIX, or Linux.
Primary IP	The host primary IP address.
Site	The host site.
Policy	The type of policy to be used to connect to the host.
Group	The group to which the host belongs.

- 3
- Click **OK**. The window closes and the Hosts window lists the new host.

Notes: To edit a host, select the host in the View Hosts list and click Edit Host. Note that you cannot change the host name, only the operating system type, and the primary IP.

If you want to change the host name, you must delete the existing host and recreate it with a new name.

To delete a host, select the host in the View Hosts list and click Delete Host.

Adding a host from a CSV file

You can add hosts to DRA by creating a comma-separated values (CSV) file and importing it. When you create a CSV it must exactly match the format shown in the CSV sample figure.

To add a host from a CSV file

- 1
- On the Host Discovery page, click **Import Hosts from CSV File** (.
- 2
- On the Import Hosts from CSV file screen, review the file format. Your file must match this format exactly.

Import Hosts from CSV file

CSV file format:

- **Name ***: The host name
- **IP ***: The primary IP address of the host
- **OS ***: The Operating System, must be one of the following:
 - Solaris
 - HP-UX
 - Linux
 - AIX
 - Windows

	A	B	C
1	Name	IP	OS
2	server1	192.168.2.45	Windows
3	server2	192.168.2.46	HP-UX
4	server3	192.168.2.47	Linux
5	server4	192.168.2.48	Solaris
6	server5	192.168.2.49	AIX

[Download sample](#)



Importing a host list might take several moments. Use the "Refresh" button to refresh the host list. If the hosts are not added after more than one minute, please verify that the file format is according to the instruction above

Import from file...

Close 

- 3 Create a CSV file using a text editor or by saving an Excel spreadsheet to CSV format. Your file must include the following:

Heading	The first line of your file must be the following: Name, IP, OS
Name	The name of the host.
IP	The primary IP address of the host.
OS	The operating system must be one of the following: ■ AIX ■ HPUX ■ Linux ■ Solaris ■ Windows

- 4 Save your file with the extension .csv.
- 5 On the Import Hosts from CSV file screen, click **Import from file**.
- 6 Browse to your file, select it, and click **OK** to import it.

Storage discovery – configuring EMC SYMCLI, EMC NaviCLI, IBM DSCLI, IBM XCLI, IBM SVC, Hitachi HiCommand, and NetApp proxies

Configuring EMC SYMCLI proxies

To run Symmetrix Command Line Interface (SYMCLI) commands on EMC Symmetrix storage arrays, you must run the commands from a machine that is directly connected to the storage array. Because DRA is not directly connected to the storage array, it cannot issue the SYMCLI commands required to retrieve data from the storage array.

Therefore, for scanning EMC SYMCLI proxies, DRA must make an intervening connection using either SSH or Telnet to a machine that serves as a gateway. The gateway is physically connected to the storage array from which data will be obtained.

DRA can then connect to the gateway machine, from which it executes the necessary SYMCLI commands to retrieve data from the storage array. You can use either Linux or UNIX machines for these gateway connections.

For EMC SYMCLI proxies, there are one or more gateways that connect physically to all the Symmetrix storage arrays in the environment.

Note: If no single (gateway) machine is physically connected to all the devices from which you want to collect data, you must define multiple gateways.

Configuring an EMC SYMCLI gateway is similar to configuring an EMC ECC management console. However, when you configure an EMC SYMCLI gateway, you must specify the name or IP address of the gateway machine, as described below.

The designated credentials must enable access to the gateway machine, and they must have root privileges to enable DRA root access to run SYMCLI commands.

You can choose to scan some of the discovered Symmetrix arrays using a specific proxy. To do so, select the proxy you added and click Edit Proxy Scope

(). Select the arrays to scan and click Scan By This Proxy (). To disable scanning of an array by the proxy, click Do Not Scan By This Proxy (.

Configuring EMC NaviCLI proxies

To run NaviSphere Secure Command Line Interface (NaviSecCLI) commands on EMC CLARiiON storage arrays, you must run the commands from a machine that has a connection to the storage array.

Because DRA does not support a Windows NaviCLI gateway, it cannot issue the NaviSecCLI commands required to retrieve data from the storage array.

Therefore, for EMC NaviCLI scanning, DRA must make an intervening connection using either SSH or Telnet to a machine that serves as a gateway, which is connected to the storage array from which data will be obtained.

DRA can then connect to the gateway machine, from which it executes the necessary NaviCLI commands to retrieve data from the storage array. You can use either Linux or UNIX machines for these gateway connections.


For EMC NaviCLI proxies, there are one or more gateways that connect to all the CLARiiON storage arrays in the environment.

Note: If more than one gateway machine is connected to all the arrays from which you want to collect data, you must define multiple gateways.

Configuring a NaviCLI proxy is similar to configuring an EMC SYMCLI proxy.

To configure the scope of an EMC NaviCLI proxy and to discover the CLARiiON arrays


- 1 Select the newly created proxy and then click . The following window opens:



You can perform the following operations from this window:

- Add CLARiiON Array  Adds a new CLARiiON array
- Edit CLARiiON Array  Modifies the properties of the selected CLARiiON array
- Delete CLARiiON Array  Deletes a CLARiiON array
- Scan by This Proxy  Enables scanning of the selected storage arrays using the currently selected proxy
- Do Not Scan by This Proxy  Disables the scanning of the selected storage arrays using the currently selected proxy
- Change Policy  Changes the policy assigned to the selected CLARiiON array

- 2 Click **Add CLARiiON Array** (). The following window opens:



- 3 In the **Target IP** field, enter the IP address of one of the storage processors for the CLARiiON array.
- 4 In the **Target Port** field, enter the port number. The default port is 443.
- 5 In the **Policy** field, select the policy for the CLARiiON probe.
To add a new policy, click the **Add Policy** button, which is to the right of the drop-down list.
For instructions on how to add a policy, see [“Adding a policy”](#) on page 148.
Make sure that the credential for the policy is a CLARiiON type.

Note: Typically, the user name, password, and scope are the same in all installations. To change the user name, password, or scope, display the Configuration page and select the Policies branch in the tree, as described in [“Configuring policies”](#) on page 148.

- 6 To save your settings and exit the window, click **OK** and wait until array discovery completes. The array is added as a new row in the Storage Arrays Discovered by Proxy window.
- 7 DRA retrieves information from all the CLARiiON arrays it discovers using the NaviCLI proxy. If you do not want DRA to use this proxy to retrieve information on all arrays, select the arrays to be scanned and disable all others. To return to the previous menu, click **Close**.

Note: To delete an array, select the array in the Storage Arrays Discovered by

Proxy window and click **Delete CLARiiON Array** (.

Configuring IBM DSCLI proxies

To run DS Command Line Interface (DSCLI) commands on IBM DS storage arrays, you must run the commands from a machine that has a connection to the storage array.

For IBM DSCLI scanning, DRA must make an intervening connection using either SSH or Telnet to a machine that serves as a gateway, which is connected to the storage array from which data is to be obtained.

DRA can then connect to the gateway machine, from which it executes the necessary DSCLI commands to retrieve data from the storage array. You can use either Linux or UNIX machines for these gateway connections.

For IBM DSCLI proxies, there are one or more gateways that connect to all the DS storage arrays in the environment.

DSCLI proxy configuration and array discovery is similar to EMC CLARiiON proxy.

Note: If the DS arrays are partitioned, multiple DS arrays may be discovered for the same IP address provided in the Add IBM DS Array page.

Configuring IBM XCLI proxies

To run XSCLI commands on IBM XIV storage arrays, you must run the commands from a machine that has a connection to the storage array.

For IBM XCLI scanning, DRA must make an intervening connection using either SSH or Telnet to a machine that serves as a gateway, which is connected to the storage array from which data is to be obtained.

DRA can then connect to the gateway machine, from which it executes the necessary XCLI commands to retrieve data from the storage array. You can use either Linux or UNIX machines for these gateway connections.

For IBM XCLI proxies, there are one or more gateways that connect to all the XIV storage arrays in the environment.

XCLI proxy configuration and array discovery is similar to EMC CLARiiON proxy.

Configuring SVC proxies

The scanning of SVC storage arrays is performed by connecting directly to the SVC. Therefore, you must configure a proxy for each SVC storage array you want to scan.

For example, if you have five SVC storage arrays, you must configure five SVC proxies.

Configuring an SVC is similar to configuring an EMC SYMCLI proxy, as described below.

Configuring Hitachi HiCommand proxies

For more information about configuring Hitachi HiCommand proxies, see [“Management console definition – configuring EMC ECC, Hitachi Data Systems \(HDS\) HiCommand, IBM HMC, and VMware vCenter”](#) on page 72.


Configuring NetApp proxies

The scanning of NetApp filers is performed by connecting directly to the NetApp operating system. Therefore, to configure NetApp filers for scanning, you must configure a proxy for each NetApp storage filer. For example, if you have five NetApp storage arrays, you must configure five NetApp proxies. Configuring a NetApp proxy is similar to configuring an EMC SYMCLI proxy, as described below.

Note: As there is no actual need for NetApp filers in a proxy, the term proxy in this context refers to the definition of the IP address of the filer and the policy used for scanning.

To configure a storage proxy

- 1 From the Management Console Definition page (described in “[Management console definition – configuring EMC ECC, Hitachi Data Systems \(HDS\) HiCommand, IBM HMC, and VMware vCenter](#)” on page 72), click **Next**. The Storage Discovery page opens.




You can perform the following operations in this window:

- Add Storage Proxy**  Adds a new proxy.
- Edit Storage Proxy Definitions**  Modifies the properties of a proxy.
- Edit Proxy Scope**  Enables you to select the exact storage arrays to scan for each storage proxy. In a large environment, this can reduce the full system scan duration.
- Delete Storage Proxy**  Deletes a proxy.
- Verify Connectivity**  Verifies the connectivity to the proxy by trying to log in to it. If it is successful, the Verified field in the table on the Storage Discovery page is checked. (The Verified field can also be checked if the Discover Storage Arrays and Execute Scan operations succeed.)
- Discover Storage Arrays**  Uses the proxy to discover all the storage units. You can display the discovered storage units by clicking View Storage Arrays.
- Execute Scan**  Connects to the selected proxy and collects information. This may take several minutes.

View Storage Arrays 

Lists the storage arrays discovered by the proxies. You cannot manually edit or add to this list.

- 2 At the bottom of the page, click **Add Storage Proxy** (). The following window opens:



The image shows a 'Storage Proxy' dialog box with a 'Properties' tab. It contains the following fields and controls:


- Type:** A dropdown menu with 'SymCLI' selected.
- Description:** An empty text input field.
- IP address:** An empty text input field.
- Policy:** A dropdown menu with 'CLARiiON Policy' selected. To the right of the dropdown is a small icon with a plus sign and a gear.
- Site:** A dropdown menu with 'Unknown Site' selected.
- Enable:** A checked checkbox.
- Buttons:** 'OK' (with a green checkmark icon) and 'Cancel' (with a red X icon) buttons at the bottom.

- 3 Enter the proxy details. In the **Type** field, select SymCLI, NetApp, or NaviCLI according to your requirements. If HiCommand was already defined on the Management Console Definition page, it is automatically available on this page.
- 4 In the **Description** field, enter the name or a description for this proxy.
- 5 In the **IP** field, enter the proxy's IP address.
- 6 In the **Policy** field, select a policy from the drop-down list.
To add a new policy, click the **Add Policy** button, which is to the right of the drop-down list.
For instructions on how to add a policy, see ["Adding a policy"](#) on page 148.

Note: Typically, the user name and password are the same in all installations. To change the user name or password, see ["Configuring policies"](#) on page 148.

- 7 In the **Site** field, select a site from the drop-down list.
- 8 Select the **Enable** check box to enable scanning using this proxy during the full system scan. If you do not select this check box, the information from this proxy is not collected during the next full system scan.
- 9 Click **OK** to save your settings and exit the window. The proxy is added as a new row in the Discovery Probes list on the Storage Discovery page.

- 10 Click **Verify Connectivity** () to verify if the probe is connected and if the policy is correct by trying to log in to it. If it is successful, the Verified field in the table on the Storage Discovery page is checked.
- 11 Click **Discover Storage Arrays** () to use this proxy to discover the storage arrays that can be accessed from this proxy.
The storage discovered can be displayed by clicking **View Storage Arrays** ()



Type	Name	Source	Scan Status	Connectivity Status
Symmetrix	000187870349	Symmetrix proxy - storagesrv (192.168.68.151)		
Symmetrix	000187910448	Symmetrix proxy - storagesrv (192.168.68.151)		
Symmetrix	000187880222	Symmetrix proxy - storagesrv (192.168.68.151)		


Page 1 of 1 Search: Reg Exp No data to display Show: 50 items

Close

Note: The Discover Storage Arrays option finds only storage arrays in your environment. The Execute Scan option collects information about all other data related to those arrays: disks, replications, and so on. It is not recommended to execute a scan during the initial configuration of the system, because the scan may take some time. You should only run the Discover Storage Arrays option. The Execute Scan itself is a part of the Full System Scan process.

- 12 You can select the Symmetrix/CLARiiON storage arrays that you want to scan using a specific proxy.

Click **Edit Proxy Scope** () to open the following window:



Uncheck the **Scan all storage arrays discovered by this proxy** check box and then select the storage arrays you want to scan.

- 13 Click **Scan By This Proxy** () , and then **Close**.
- 14 Click **Next**. The Sites page is displayed, letting you configure the sites.
See “[Site definition – configuring sites](#)” on page 91.

To delete a storage array

- 1 On the Management Console Definition page (described in “[Management console definition – configuring EMC ECC, Hitachi Data Systems \(HDS\) HiCommand, IBM HMC, and VMware vCenter](#)” on page 72), click **Next**. The Storage Discovery page opens.
- 2 Click **View Storage Arrays** () to view the storage arrays that exist in the system.
- 3 Select the storage arrays you want to delete and click **Delete Storage Array** () .

Site definition – configuring sites

This step lets you configure the physical sites of your environment in DRA. A site is an abstract group of hosts and storage arrays, such as the Boston site, the Finance department, the 2nd floor, and so on, along with the criteria for connecting to these objects. The objects you can map to a site include hosts and storage arrays. An item can only belong to one site.

You can add new sites and edit or delete existing ones.

DRA includes a default site called Unknown Site. It contains all the entities that are not connected to any other site.

Note: If you are using LVM site tagging in your environment, you should use the same site names as the names in the LVM. Configuring the same names automatically activates site-tagging-related gap signatures.

You can perform the following operations from the Sites page:

Add Site 

Adds a new site

Edit Site Definitions 


Modifies a site's properties and changes the host/storage association

Delete Site 

Deletes a site


To configure a new site

- 1
- From the Storage Discovery page (described in “Storage discovery – configuring EMC SYMCLI, EMC NaviCLI, IBM DSCSI, IBM XCLI, IBM SVC, Hitachi HiCommand, and NetApp proxies” on page 81), click **Next**. The Sites page opens.



The right pane lists the sites defined in DRA, including the site name, description, and the number of connected items.

- 2
- To define a new site, click **Add Site** (). The following window opens:



- 3
- In the **Name** field, enter a name for the site. This name identifies the site throughout the application.
- 4
- In the **Description** field, enter a description for the site. You can now define related item information for the site.

- 5 In the Site Entries Configuration area, in the **Available Entities** list, highlight the items you want to connect to the site, and click the  right arrow. The selected item(s) is moved into the **Entities Associated to Site** list. To remove an item from the **Entities Associated to Site** list, select the item and click the  left arrow.

Note: You can only assign a host/storage array to a single site. After you assign a host/storage array and it appears in the Entities Associated to Site list, it is no longer available in the Available Entities list to be assigned to other sites.

Note: You can find specific hosts in the lists by entering all or part of their name in the Filter field, selecting the Match Case check box (if relevant), checking the Regular Expression check box (if the value you entered in the Filter field is a regular expression), and clicking Go.

- 6 To save your settings and exit the window, click **OK**. The site is added to the Sites list.
- 7 To proceed to the next step, click **Next**. The Host Groups & Policies page is displayed, letting you configure the hosts groups and policies.
See "[Hosts groups and policies – configuring hosts](#)" on page 95.

Note: After an item mapping is listed in the Entities Associated to Site list on the Add Site or Edit Site pages, you can remove it by selecting it and clicking the  left arrow.

To edit a site

- 1 On the Sites page of the Configuration Wizard, click **Edit Site Definitions** ().
- 2 Modify the site's properties as required and click **OK** to save your changes.
- 3 Click **Next**.

Adding a site from a CSV file


You can add sites to DRA by creating a comma-separated values (CSV) file and importing it. When you create a CSV file, it must exactly match the format shown in the CSV sample figure.

To add a site from a CSV file

- 1 On the Site Definition page, click **Import Site Definitions From CSV File**




The following window opens:



- 2 Review the file format. Your file must match this format exactly.
- 3 Create a CSV file using a text editor or by saving an Excel spreadsheet to CSV format. Your file must include the following:

Heading	The first line of the file must be the following: Site, Host
Site	The site name.
Host	The server name.
	Note: Only physical server associations are imported.

- 4 Save the file with the extension *.csv.

- In the Import Site definition from CSV file window, click **Import from file ...** ().
- Browse to your file, select it, and click **OK** to import it.

Hosts groups and policies – configuring hosts

You can organize hosts into logical groups so you can manage which hosts are to be scanned now and which should not be scanned. You can assign each host to a policy which defines the connection protocols and credentials DRA uses to communicate with it.

The following information is shown for each host:

Name	The host machine name.
IP	The host machine's IP address.
Type	The type of operating system.
Policy	The name of the policy used to connect to the host. The policy is a virtual item that defines the way to communicate with the host, management consoles, and storage proxies. The policy contains the credentials used to log in and other definitions, such as proxy, and sudo. For more information, see “Configuring policies” on page 148.
Site	The site where the host is located.
Associated Groups	The scanning groups to which the host is assigned.
Physical Host	The name of the physical host. For a VM, the name of the physical host running that VM is shown.
Note	A note attached to the host (if provided by the user).
Scan Status	The last scan status of the host. This status changes only after a full scan of the host.
Connectivity Status	The status of the last connection attempt to the host.

You can perform the following operations from the Host Groups & Policies page:

- Manage Groups  Lets you define logical groups of hosts.
For more information, see [“Managing host groups”](#) on page 101.
An empty default group is defined during installation.

Manage Host Connectivity Policies		Lets you select the policy associated with the selected host. For more information, see “Managing host policies” on page 100.
Rules for Automatic Association of Databases to Groups and Connectivity Policies		You can configure automatic association of databases to groups and policies. The application automatically assigns newly discovered databases to the scanning group and policy defined in the rule, and the database is scanned in the first full system scan after it is discovered. For more information, see “Managing automatic association of databases to groups and policies” on page 110.
Ping Hosts		Tries to ping the specified host. The status of this operation is shown in the Connectivity Status column of the Host Groups & Policies page.
Verify Connectivity		Verifies whether the selected hosts are connectable by trying to connect to them. If it is successful, the Verified field in the table on the Host Groups & Policies page is checked.
Discover Databases		Performs a light scan to discover the databases installed on the selected hosts. To display the discovered databases, click View Databases.
Execute Scan		Scans the selected hosts.
Verify and Correct IP Address		Resolves the host name using the DNS server. If the resolution is successful, the IP address of the host changes accordingly.
View Hosts		Lists the hosts discovered by ECC, HiCommand, or vCenter, or those that were added manually.
View Databases		Lists the databases found when Discover Databases is clicked. Lets you manually add, edit, or delete a database. For more information, see “Adding databases” on page 102.
Add Host		Manually adds a host to the system.
Edit Policy		Changes the policy assigned to a host probe.
Add Note		Attaches a note to the host for easier maintenance of the scanned environment.

View Custom Collection



Lists the data collected using data collection expansion packages.

Note: You can select more than one host and click Verify/Discover/Execute/Edit Policy.


To define host groups and policies

- 1 From the Sites page, click **Next**. The Host Groups & Policies page is displayed, listing the hosts to be scanned when you perform a full scan for all the enabled groups:

Name	IP	Type	Policy	Site	Associated Groups	Physical Host	Note	Scan St...	Connect...
star1	star1	Solaris	-	Boston	Default	moon3	-		
cshram1	cshram1	HP/UX	-	Chicago	Default	Physical host	-		
oraprod1	oraprod1	Linux	-	Boston	Default	Ajpn2	-		
nev1	10.2.0.218	Windows	Windows Policy	Chicago	Default	Physical host	-		
mia	10.2.0.212	Solaris	Unix Policy	Chicago	Default	Physical host	-		
neviah	10.1.0.212	Windows	Windows Policy	Boston	Default	Physical host	-		
qa4	qa4	Linux	-	Boston	Default	Ajpn5	-		
sqbrd2	172.16.0.2	Windows	-	Boston	Default	Physical host	-		
dwhp	10.5.180.240	Linux	Unix Policy	Boston	Default	Physical host	-		
aw	10.1.0.213	Solaris	Unix Policy	Boston	Default	Physical host	-		
qa5	qa5	Linux	-	Boston	Default	Ajpn2	-		
blared		Windows	Windows Policy	Chicago	Default	Physical host	-		
star2	star2	Solaris	-	Boston	Default	moon2	-		
avad	10.1.0.223	Solaris	Unix Policy	Chicago	Default	Physical host	-		
prodsm1	prodsm1	Linux	-	Boston	Default	Physical host	-		

For more information on the Sites page, see “[Site definition – configuring sites](#)” on page 91.


- 2
- To assign a policy to a specific host, click **Manage Host Connectivity Policies** (). The Policy page opens.



For more information on managing policies, see [“Managing host policies”](#) on page 100.

- 3
- To save your changes and close the Policy window, click **OK**.
- 4
- To define a logical group of hosts, click **Manage Groups** (). The Group Definitions window is displayed, letting you control which hosts are scanned according to the group(s) you select.
For more information managing groups, see [“Managing host groups”](#) on page 101.
- 5
- To save your changes and close the Group Definitions window, click **OK**.
- 6
- To discover the databases that are installed on a specific host, go to the **Hosts to be Scanned in the Next Scheduled Full Scan** table, select the host that you want to scan, and click **Discover Databases** (). You can select more than one host.

- 7 To display the discovered databases, select **View Databases** (). The following window opens:



The screenshot shows a window titled "Databases" with a table of discovered databases. The table has columns: Host Name, Host IP, Type, Name, Virtual IP, Port, Policy, Scan S..., and Connec... (Connection). The table lists various databases including Oracle, MS SQL server, and Sybase. At the bottom of the window, there are navigation buttons (back, forward, search, etc.) and a "Close" button.

Host Name	Host IP	Type	Name	Virtual IP	Port	Policy	Scan S...	Connec...
crmd	10.2.0.217	Oracle	vanxml			Oracle policy		
crmd	10.2.0.217	Oracle	vandesk			Oracle policy		
tristan	10.1.0.216	Oracle	stat	tristan	1529	Oracle policy		
blaked	blaked	MS SQL server	sgbdr04	10.1.3.18, 10.1.3.3...	3279	MSSQLServer Policy		
blakep	blakep	MS SQL server	sgbdr03	10.1.3.18, 10.1.3.3...	2630	MSSQLServer Policy		
blaked	blaked	MS SQL server	sgbdr03	10.1.3.18, 10.1.3.3...	2630	MSSQLServer Policy		
blakep	blakep	MS SQL server	sgbdr02	10.1.3.18, 10.1.3.3...	1859	MSSQLServer Policy		
blaked	blaked	MS SQL server	sgbdr02	10.1.3.18, 10.1.3.3...	1859	MSSQLServer Policy		
blakep	blakep	MS SQL server	sgbdr01	10.1.3.18, 10.1.3.3...	1033	MSSQLServer Policy		
blaked	blaked	MS SQL server	sgbdr01	10.1.3.18, 10.1.3.3...	4904	MSSQLServer Policy		
lb1	10.1.0.208	Oracle	rbch10			Oracle policy		
udbhost	udbhost	DB2	roy		N/A	DB2		
rgdb-xp	192.168.68.59	Sybase	rgdbxp	rgdb-xp	5000	Sybase policy		
lnwx01	192.168.68.41	Oracle	rgdb	lnwx01	1521	Oracle policy		
rgdb-xp	192.168.68.59	Oracle	ora9			Oracle policy		
udbhost	udbhost	DB2	moss		N/A	DB2		
crmd	10.2.0.217	Oracle	vanxml			Oracle policy		

To manually add more databases in this window, see [“Adding databases”](#) on page 102.

- 8 Optionally, you can click **Execute Scan** () to scan the defined hosts, and then select **View Databases** to display the discovered databases.
- 9 Click **Next** to proceed to the next step. The Database Groups & Policies page is displayed, letting you configure the database Groups and Policies, as described in [“Database groups and policies – configuring databases”](#) on page 105.


Managing host policies

This section describes how to assign and unassign policies to and from hosts.

To manage host policies

- 1
- To assign policies to one or more specific hosts, click **Manage Host**

Connectivity Policies (). The Policy window opens.



- 2
- Select the policy to associate with the hosts.
- 3
- On the left side of the page, in the Hosts that are not assigned to current Policy list, select the hosts to assign the policy.
- 4
- To assign hosts to the selected policy, click the  button.

Note: To define a new policy in the Policy list, click Add Policy and follow the procedure defined in “[Adding a policy](#)” on page 148. After you assign a policy to a host, the host is removed from the Hosts that are not assigned to current Policy list.

Note: To remove the association between hosts and policies, go to the right side of the page, in the Host assigned to current policy list, and select the

policy and mark the hosts that you want to unassign, and click the  button.

- 5
- To save your changes and return to the Hosts page, click **OK**.

Managing host groups


This section describes how to define groups and manage the association of hosts with those groups.

You can perform the following operations from the group definitions page:

- Add Group  Defines a new group
- Edit Group  Modifies a group's properties
- Delete Group  Deletes a group

To manage a group


- To define a logical group of hosts, click **Manage Groups** () in the Configuration Wizard Host Groups & Policies page. The Group Definitions window is displayed, letting you control which hosts are scanned according to the group(s) you select.



Note: You can assign a host to multiple groups.

Selecting a group in the Groups list at the top of the page displays the hosts in and not in that group in the Hosts in Group and Hosts not in Group tables at the bottom of the page.

- 2 At the top right, click **Add** () to define a new group. The following window opens:



- 3 Enter the group name and enable the group by selecting the **Enable** check box.

Note: If you do not enable the group, it is not scanned in the next scan cycle. When a host is assigned to multiple groups, it is scanned only once if at least one of its groups is enabled.

- 4 Click **OK**. The window closes and the Group Definitions page is displayed.


Adding databases

To discover all databases installed on the selected hosts, go to the Host Groups & Policies page and click Discover Databases. After this, you list the discovered databases discovered by clicking View Databases. To connect to a specific database that is not in the list, you must manually define that database, as described below.


Note: To discover Oracle databases, the Oracle process should be started or the `/etc/oratab` or `/var/opt/oracle/oratab` file should be present.

To define a database

- 1 On the Host Groups & Policies page, click **View Databases** (). The Databases window opens.



- 2 Click **Add Database** (). The following window opens:



3 Enter the following database details:

- Name


The database name.
- Port

The port from which to connect to the database.
- Virtual IP

The IP address from which to connect to the database.
- Type

The database type; for example, Oracle, DB2, Sybase, MSSQL Server, and so on.
- Hostname

The host to which the database is connected. To specify a host, click the Select Host button. The following window opens:



- Policy

The policy to use to connect to the database.
- Group

The group to which the database belongs.

4 Select a host and click **OK**. The window closes and the new database appears in the list in the Databases window.

Note: To edit a database's properties, select the database from the list and click **Edit Database**.

Note: To delete a database, select it from the list and click **Delete Database**.

Database groups and policies – configuring databases

Like hosts, you can organize databases in logical groups so you can manage the scan schedule of specific databases.

You can assign a policy to each database which defines the connection protocols and credentials DRA uses to communicate with the database.

The following information is shown for each database:

Host Name	The host machine name.
Database Type	The database type.
Database Name	The database name.
IP	The IP address used to connect to the database. It can be the host machine's primary IP address or a virtual address on it.
Port	The port on which the database accepts client connections.
Policy	The name of the policy used by the database. The policy is a virtual item that defines the way to communicate with the host, database, management console, and proxies. The policy contains the credentials used for login and other definitions, such as proxy, sudo, and so on. For more information about policies, see “Configuring policies” on page 148.
Scan Status	The last scan status of the database. This status changes only after a full scan of the database.
Connectivity Status	The status of the last connection attempt to the database.
Associated Groups	The group associated with the database.
Business Entities	The business entity associated with the database.
Roles	The logical role associated with the database. For example, Backup, DR, Production, and so on.

You can perform the following operations from the Database Groups & Policies page.

Manage Groups	 Lets you define a logical group of databases. An empty default group is defined during installation.
Manage Database Connectivity	Lets you assign policies to databases, as described below.
Policies	

Rules for Automatic Association of Databases to Groups and Connectivity Policies 	You can configure automatic association of databases to groups and policies. The application automatically assigns newly discovered databases to the scanning group and policy defined in the rule, and the database is scanned in the first full system scan after it is discovered. For more information, see “Managing automatic association of databases to groups and policies” on page 110.
Verify Connectivity 	Verifies whether the selected databases are connected by trying to connect to them. If it is successful, the Verified field in the table on the Databases page is checked.
Execute Scan 	Scans the selected databases.
View Databases 	Lists discovered databases.
Add Database 	Manually adds a database.
Edit Policy 	Changes the policy assigned to the selected databases.
Edit Port and Virtual IP 	Changes the port or virtual IP for the selected databases.

Note: You can select more than one database and click Verify/Execute.

To configure database groups and policies

- 1 From the Configuration Wizard Host Groups & Policies page, described in [“Hosts groups and policies – configuring hosts”](#) on page 95, click **Next**. The Database Groups & Policies page is displayed, listing all the databases defined in the system.
- 2 To assign a policy to a specific database, click **Manage Database Connectivity Policies** (). The Policy page opens.

For more information about managing policies, see [“Managing host policies”](#) on page 100.
- 3 To save your changes and close the Policy window, click **OK**.


- 4 To define a logical group of hosts and/or databases, click **Manage Groups** (). The Group Definitions window is displayed, letting you control which databases are scanned according to the group(s) you select. For more information on managing groups, see [“Managing database groups”](#) on page 108.
- 5 To save your changes and close the Group Definitions window, click **OK**.

Managing database policies

This section describes how to assign and unassign policies to and from databases.

To manage database policies

- 1 To assign a policy to one or more databases, click **Manage Data base Connectivity Policies** () in the Configuration Wizard Database Groups & Policies page. The Policy window opens.



- 2 Select the policy to associate with the databases.
- 3 On the left side of the page, in the **Databases not assigned to current Policy** list, select the databases to assign the policy.

- 4
- To assign the databases to the selected policy, click the  button.

Note: To define a new policy that is not in the Policy list, click Add Policy and follow the procedure defined in “Adding a policy” on page 148. After a policy is assigned to a database, the database is removed from the Databases not assigned to current Policy list.

Note: To remove the association between databases and policies, go to the right side of the page, in the Databases not assigned to current Policy list, and select the policy and mark the databases that you want to unassign, and click the  button.

- 5
- To save your changes and return to the Database Groups & Policies page, click **OK**.

Managing database groups


This section describes how to define groups and to assign and unassign databases to and from those groups.

You can perform the following operations from the Group Definitions page:

Add Group		Defines a new group
Edit Group		Modifies the properties of a group
Delete Group		Deletes a group

To manage a database group


- 1 To define a logical group of databases, click **Manage Groups** () in the Configuration Wizard Database Groups & Policies page. The Group Definitions window is displayed, letting you control which databases are scanned according to the group(s) you select.



Note: You can assign a database to multiple groups.

Selecting a group in the Groups list at the top of the page displays the databases in and not in that group in the Databases in Group and Databases not in Group tables at the bottom of the page, respectively.

- 2 To define a new group, go to the top right and click **Add** () . The following window opens:



- 3
- Enter the group name and enable the group by selecting the **Enable** check box.

Note: If you do not enable a group, it is not scanned in the next scan cycle. When a host is assigned to multiple groups, it is scanned only once if at least one of its groups is enabled.

- 4
- Click **OK**. The window closes and the Group Definitions page is displayed.

Managing automatic association of databases to groups and policies

This section describes how to automatically assign groups and policies to a newly discovered database.

A policy and scan group rule specifies a criteria definition that is based on the host and the database type. Each newly discovered database that matches a rule's criteria is automatically assigned to the policy and scanning group configured in that rule.

You can perform the following operations from the Automatic Policy and Scanning Group Assigning page:

- Add Rule 


Defines a new rule. If the new rule is based on the same host and database type, the old rule is overridden.
- Delete Rule 

Deletes the rule.

To manage an automatic policy and scanning group assignment rule

- 1 To define a new rule, click **Rules for Automatic Association of Databases to**

Groups and Connectivity Policies () in the Configuration Wizard Host Groups & Policies page or in the Database Groups & Policies page. The Automatic Policy and Scanning Group Assigning window is displayed, which enables you to control the rules defined in the system.



- 2
- To define a new rule, click **Add Rule** (). The following window opens:

-- Webpage Dialog

Policy and Scan Group Rule

Host

☒ All hosts

[select host]

Select the host that you would like to apply this rule on. You can select **All hosts**, or select a specific host from a list.
Note: when selecting a specific host - this definition will override rule that was defined for "All hosts".

Database Type

Oracle

The discovered database type to be automatically assigned. For example: when selecting "Oracle", this rule will take affect on all of the newly discovered Oracle instances installed on the selected host (or all hosts, if checked).

Default Policy

Ora-Policy

The default policy that will be assigned to the new databases.
In MS SQL servers, make sure that the type of the credential of the policy is of **Windows Authentication** type.

Default Group Name

New Oracle databases

The default scan group that the new databases discovered will be assigned to.

OK

Cancel

3

Enter the following rule details and then click **OK**:

Host	The host to which to apply this rule. You can select a specific host or check the All Hosts check box if you want this rule to be the default rule for the selected database type.
Database Type	The database type to which to apply this rule.
Default Policy	The policy to which newly discovered databases that match the rule criteria are automatically assigned.
Default Group Name	The scanning group to which newly discovered databases that match the rule criteria are automatically assigned.

In the example window shown above, every new Oracle database that is discovered by DRA is automatically assigned to the policy called Ora-Policy and to the scanning group called New Oracle databases.

Viewing the status of an entity

You may need to view the status of entities such as hosts and databases after they are configured in the system. DRA provides an icon status indicator for each entity. A tool tip is available for each icon.

[Table 6-1](#) lists the DRA status indicators.

Table 6-1 DRA status indicators

Status	Icon	Description
Not Verified		The initial status assigned to an entity, before connectivity to it has been attempted. This is the default status.
Connection Verified		The system successfully connected to the entity, but the entity has not yet been scanned.
Connection Failed		The system failed to connect to the entity.
Authentication Failed		The system established a connection to the entity, but authentication for the requested user failed. This status typically indicates that an incorrect password was provided.
Database Down		The system tried to connect to a database and was informed that the database was down.
Host Not Scanned		The requested database was not scanned because the host was not scanned during the last 24 hours.
Connection Partially Scanned		The entity was scanned successfully, but some required information could not be retrieved. Contact vendor support.
Missing Credentials		The user did not provide proper credentials.
Multiple Credentials		The user provided more than one set of credentials for the entity.
Fully Scanned		The entity was successfully scanned.
Ping Failed		The ping operation failed.
Ping Succeeded		The ping operation was successful.

Table 6-1 DRA status indicators

Status	Icon	Description
Data Collector Not Available		The data collector is not currently available.
Scan Timed Out		The scan timed out.

Configuring distributed collection

DRA uses a distributed collection mechanism for scanning environments and gathering data. Distributed collection enables multiple collectors to be defined, with each one collecting the data of a single site and then sending it back to the server. For distributed collection, different collectors can be installed on various hosts, which improves network traffic as well as enables scanning of secure, demilitarized environments.

Communication between the DRA server and a collector uses a single port (default port is 27000), with communication flowing from the server to the collector. A local collector is installed for DRA by default, named Local Scan Server. If no additional collectors are defined in DRA, then DRA runs with only a single collector.

When you define a collector, you define its primary sites and its failover sites. Because it is possible to have many collectors per site, DRA must first determine which collector to use for scanning a host. To do so, DRA searches for a collector whose primary site is the same as the site of the host. Collectors must be online (meaning, either idle or scanning) when making this determination, in order to be considered for scanning. Any collectors that are offline (down) are not considered.

If a single collector is identified that meets these criteria, that collector is used to scan that host. If more than one collector is identified, then all applicable collectors are used to scan the host, using load balancing between these collectors. For example, let's say that you want to scan 100 hosts, and that there are two collectors that meet the required criteria. In this case, each collector will scan 50 hosts.

If a collector is not identified that has the same primary site as the site of the host, then the DRA attempts to identify a collector that has the host's site defined as its failover site, and whose status is online. If such a collector is identified, it is used to scan the host.

To summarize, each site can have one or more collectors that are the primary for that site and one or more collectors that are failovers for that site. All of them are considered when determining the collector to use to scan a host.

If no primary sites or failover sites are defined for a collector, then the host is scanned using the Local Scan Server (default collector).

Defining a collector to run in DRA involves the following steps:

- 1 Download the collector installation file, as described on page 118.
- 2 Install the collector on the machine on which it will run, as described in [“Installing a Collector”](#) on page 43.
- 3 Use the **Add New Collector** button, as described on page 117, to configure the collector in DRA. Communication with the collector begins after you define the collector in DRA.

The following information is shown for each collector:


Name	The name of the collector.
Hostname	The name or IP address of the machine.
Port	The port of the machine on which the collector runs.
Location	The site on which the collector runs.
Primary Sites	The name of the primary site on which the collector runs.
Failover Sites	The name of the failover site on which the collector runs.
Status	The current status of the collection, which can be Down, Idle, or Scanning.
Connection Status	An icon indicating whether the collector is connected. The icon is green when the collector is connected and red when it is not.

To configure a collector

- 1 Select the Configuration tab.

Note: You must be an Admin user to perform this operation.

- 2 In the tree hierarchy, click **Distributed Collection**. The Distributed Collection Management page opens.



- 3 Click **Add New Collector** () to add a collector. The following window opens:

The 'Collector Properties' dialog box contains the following fields and controls:

- Name:** Text input field.
- Hostname:** Text input field.
- Port:** Text input field with the value '27000'.
- Location:** Dropdown menu with 'Unknown Site' selected.
- Primary Sites:** Button with a list icon and '[none]' text.
- Fallover Sites:** Button with a list icon and '[none]' text.
- Enable:** A checked checkbox.
- Buttons:** 'OK' and 'Cancel' buttons at the bottom.

4 Enter the following details and then click **OK** to save the collector definition:

Name	The name of the collector.
Hostname	The name or IP address of the machine.
Port	The port of the machine on which the collector runs.
Location	The site on which the collector runs.
Primary Sites	The name of the primary site(s) on which the collector runs.
Failover Sites	The name of the failover site(s) on which the collector runs.
Enable	Indicates whether the collector is enabled or disabled. When a collector is disabled, its Status field is yellow  Disabled .

You can perform the following operations from the Distributed Collection Management page:

Add New Collector		Adds a new collector
Edit Collector Properties		Edits the properties of a collector
Delete Collector		Deletes a collector
Receive Logs		Retrieves the collector logs and saves them in the Logs directory on the DRA server
Restart Collector		Restarts the collector
Reconnect		Disconnects the communication of the server from the collector and establishes a new connection
Search Connections		Attempts to connect immediately to all offline collectors; by default, DRA attempts to connect to offline collectors every 30 seconds
Download Collector Installer		Retrieves the installation file for a collector and stores it on the machine on which the DRA GUI is being viewed.

Configuring standby definitions

DRA lets you define which host/cluster is the standby host/cluster for a specific primary host/cluster. You do this by defining a pair of primary and standby hosts/clusters. Pairs can be defined by pairing:

Primary	Standby
Host	Host
Cluster	Host
Host	Cluster
Cluster	Cluster

Symantec strongly recommends that you define primary/standby host/cluster pairs so the system is more precise in its gap reporting. Defining primary/standby host/cluster pairs lets the system verify both primary and standby host/cluster issues.


When a primary/standby host/cluster pair is defined, it is assigned a default standby policy. You can create other standby policies, as described in [“Adding a standby definition from a CSV file”](#) on page 122. You can edit and delete standby definitions as required.

To configure a primary/standby host pair

- 1 Select the Configuration tab.


Note: You must be an Admin user to perform this operation.

- 2 In the tree hierarchy, click **Standby Definition**. The Standby Definition page opens.



- 3 In the Available Primary Hosts/Clusters pane, select a primary host or cluster. The selected host/cluster is highlighted.
- 4 In the Available Standby Hosts/Clusters pane, select a standby host or cluster. The selected host/cluster is highlighted.
- 5 Click **Configure Standby Relationship Between the Selected Hosts/Clusters** . The primary/standby host /cluster pair is added to the Standby Pairs list at the bottom of the page.


- 6
- After the first full system scan, the storage paths are available. As a result, DRA can suggest the standby pair according to the storage paths between hosts/clusters. Click **Automatically Suggest Standby Pairs** () to open the following window:



- 7
- Select the standby pair that you want to use and click **OK**.
- 8
- To assign a standby policy to a primary/standby host/cluster pair, see [“Attaching a standby policy”](#) on page 125.

To delete a primary/standby host /cluster pair

- ◆
- In the Standby Pairs list, select a primary/standby host/cluster pair and click **Delete Standby Definition** ().




Adding a standby definition from a CSV file

You can add standby definitions to DRA by creating a CSV file and importing it. When you create a CSV file, it must exactly match the format shown in the CSV sample figure.

To add a standby definition from a CSV file

- 1 On the Standby Definition page, click **Import Standby Definitions From CSV**

 **File** (). The following window opens:



- 2 Review the file format. Your file must match this format exactly.
- 3 Create a CSV file using a text editor or by saving an Excel spreadsheet to CSV format. Your file must include the following:

Heading	The first line of the file must be the following: Primary, Standby
Primary	The primary host name
Standby	The secondary host name

- 4 Save the file with the extension *.csv.

- 5 In the Import Standby definition from CSV file window, click **Import from file ...** (**Import from file...**).
- 6 Browse to your file, select it, and click **OK** to import it.

Configuring standby policies

When a pair of primary and standby hosts/clusters is initially defined, it is assigned a default standby policy. A standby policy specifies the type of relationship between the primary host/cluster and the standby host/cluster. This relationship determines how the members of the pair behave with respect to one another, such as the type of validation and replication to be used between pair members.


When the DRA system performs scanning, it looks for primary/standby host/cluster pairs. After a pair is found, the system performs the checks that are specified in the standby policy assigned to that pair.

Standby policies can be added and modified as required.

You can change the standby policy assigned to a primary/standby host/cluster pair, as described in [“Attaching a standby policy”](#) on page 125.


To configure a standby policy

- 1 Select the Configuration tab.
- 2 In the tree hierarchy, click **Standby Policies**. The Standby Policies page opens.



The Standby Policies area lists only those standby policies that you create. It does not include the default policy in this list.

- 3
- To add a new standby policy, click **Add Standby Policy** (). The following window opens:



Webpage Dialog

Standby Policy

Standby Policy properties

Name:

Description:

Standby Policy gap detection configuration

☒ **Replication Visibility**
Verify that all remotely replicated file systems, volume groups and database files of the primary host/cluster are mapped to the standby host/cluster.

☒ **Compare Local File Systems**
Verify that all local (unreplicated) file systems of primary hosts/cluster exist in each standby host/cluster.

☒ **DNS Search Order**
Verify that the DNS best practices are met.

OK  Cancel 

- 4
- Enter the following details and then click **OK** to save the policy definition:

Name	The name of the standby policy.
Description	A description of the standby policy. This field is optional.
Replication Visibility	When this check box is selected, verification is performed to ensure that all remotely replicated devices on the primary host/cluster are mapped to the standby host/cluster.
Compare Local File System	When this check box is selected, verification is performed to ensure that all local file systems of the primary host/cluster exist on each standby host/cluster. This means that tickets are opened if either of the following conditions are not met: ■ The primary host/cluster has a file system that the standby host/cluster does not have. ■ The standby side has the file system that the primary side has, but its size is less than 90% of that on the primary.

DNS Search Order

When this check box is selected, verification is performed to ensure that DNS best practices are used, and that the DNS server configuration of the primary and standby is reversed.

Note: In the default policy, all the checks are disabled.

To delete a standby policy

- ◆ In the Standby Policies list, select a policy and click **Delete Standby Policy**



Attaching a standby policy

By default, the default standby policy is initially assigned to a primary/standby host/cluster pair. You can attach a different policy, as described below.

To attach a standby policy

- 1 Select the Configuration tab.
- 2 In the tree hierarchy, click **Standby Definition**. The Standby Definition page opens.
- 3 In the Standby Pairs pane, select a primary/standby host/cluster pair. The selected host/cluster is highlighted.

Standby Policy	Primary Name	Primary IP	Primary Type	Standby Name	Standby IP	Standby Type
(Default)	HAC1		PowerHA	ha2		PowerHA
(Default)	SC-Prod		ESX cluster	SC-TEST		ESX cluster

- 4 Click **Attach Standby Policy** () to open the following window:

Attach Standby Policy

Standby Policy


Standby Policy:

- NO CHANGE -

OK

Cancel

- 5 In the **Standby Policy** drop-down list, select the standby policy to attach. The drop-down lists all the currently defined standby policies. For example:



Note: Select **NO CHANGE** to leave the currently assigned standby policy.

- 6 Click **OK**.

Configuring high availability (HA) clusters

DRA can detect some cluster types automatically. Unsupported cluster configuration can be manually defined as an HA cluster.

Symantec strongly recommends that you define the unsupported HA clusters in your environment so that the system is more precise in its gap reporting. Defining the unsupported clusters activates advanced gap signatures and reporting.

You can edit and delete manually added cluster definitions, as required.


Note: It is recommended to configure the cluster after the first full system scan. This is for two key reasons. Firstly, some of the clusters are automatically detected by DRA. Secondly, DRA can automatically suggest the cluster configuration once it has storage path information.

To configure a cluster


- 1 Select the Configuration tab.

Note: Only an Admin user can perform this operation.

- 2 In the tree hierarchy, click **HA Cluster Definition**. The HA Cluster Definition page opens.



- 3 Click **Add HA Cluster** () to add a cluster. The following window opens:




- 4 Enter the following details and then click **OK** to save the cluster definition:


- Cluster Name
- The name of the cluster
- Cluster Nodes
- Opens a window in which you select the cluster nodes for this cluster

- 5 To edit an existing cluster definition, click **Edit HA Cluster** ().

- 6 After the first full system scan, the storage paths are available. As a result, DRA can suggest the HA clusters according to the storage paths between hosts. Click **Suggest HA Clusters** () to open the following window:



- 7 In the HA Cluster Suggestion window, double-click a cluster to open the following window, in which you can provide the name of the cluster and change the cluster nodes (if needed).



- 8 Click **OK** to accept the suggested cluster.

To delete a cluster definition


- ◆ You can delete an existing cluster definition by selecting it and clicking **Delete HA Cluster** ()

Scheduling

DRA lets you schedule several types of scanning and data analysis tasks to implement scan windows using scan tasks as well as email report deliveries.

Note: Only Administrators can access these options.

Scan Scheduling



The Scan Scheduling page has several areas:


- **Scan Scheduling Configurations area:** Displays a separate row for each scheduled task. The following information is shown for each task row:

Task Type	Specifies the type of task, which can be a scan task, data analysis task, or full cycle task.
Task Identifier	Specifies the name of the scheduled task. Each type of task must have a unique name. You can assign the same name to different types of tasks. For example, a scan task and a data analysis task can have the same name.
Schedule	Indicates when the task is to run, which can be Monthly, Weekly, or One Time.
Scan Scope	Specifies the scope to be scanned, including which management consoles, storage proxy servers, hosts, and databases to scan.

- Execute Data Analysis
- Indicates whether or not data analysis is to be performed. For a data analysis task, this value is always Yes, and for a scan task it is always No. Typically, this value is Yes for a full cycle task.
- Last Execution
- A link that when clicked, displays the status of the last execution of the task. If the task was not yet executed, the following message is displayed:

Task status was not found. Perhaps the task was never executed?

If the task is currently executing, then a window opens showing the task's status.



This window shows who created the task and when it began running. It also shows a separate bullet for each component to be scanned, and its scan status.

- Status
- Displays an icon that indicates the status of the task, as follows:
- : Task executed successfully.
 - : Task is currently executing. The icon rotates clockwise while the task is running.

- **Weekly Scan Plan area:** Displays a weekly calendar that shows an at-a-glance view of the tasks with a weekly schedule. Scheduled tasks for all task types (scan task, data analysis task, and full cycle task) are shown.

Weekly Scan Plan							
	Sunday	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday
12am	12:15 AM (Local Time) Weekly_00:15 Scan Task	12:15 AM (Local Time) Weekly_00:15 Scan Task	12:15 AM (Local Time) Weekly_00:15 Scan Task	12:15 AM (Local Time) Weekly_00:15 Scan Task	12:15 AM (Local Time) Weekly_00:15 Scan Task	12:15 AM (Local Time) Weekly_00:15 Scan Task	12:15 AM (Local Time) Weekly_00:15 Scan Task
1am							
2am							
3am							
4am							
5am							

The length of the blue rectangle for a task designates the duration of that task. For example, in the window above, the task is scheduled to begin at 12:15 AM every day. This task has a duration of three hours. Therefore, the blue rectangle goes from 12:15 AM to 3:15 AM every day on the calendar.

You can perform the following operations from the Scan Scheduling page:

Add Scan Task Scheduling 

Creates a scheduled scan task that includes data collection. The maximum duration of the scan is specified, which once reached, stops the scan.

Add Data Analysis Task 

Creates a scheduled task that includes mapping and modeling dependencies between scanned IT elements. This data analysis task discovers high-availability and data-protection configurations and detects risk signatures.

Add Full Cycle Task 

Creates a full cycle scheduled task.

Edit Task Scheduling 

Edits the selected task's scheduling or scan scope configuration.

Suspend/Revive Task Scheduling



Suspends/revives a scheduled task. Suspend a task to deactivate it so that the system ignores its configuration. Revive the task to reactivate it.

A suspended task appears grayed out in the Scan Scheduling Configurations area.

Delete Task Scheduling 

Deletes a scheduled task.


Refresh Data		Refreshes the data on the page.
Execute Selected Task		Executes the selected task immediately.
Configuration		

Scheduling a scan task

A scan task only includes data collection. This type of task is performed for the initial scan in an organization or when a new component is added in the system. No data analysis is performed as part of this task. For more details about adding a data analysis task or a full cycle task, see [“Scheduling a data analysis task”](#) on page 137 and [“Scheduling full cycle scans”](#) on page 139, respectively.

To add a scan task

- 1 Select the Configuration tab.
- 2 In the tree hierarchy, click **Scan Scheduling**. The Scan Scheduling page opens.
- 3 Click **Add Scan Task Scheduling** () to add a scan task. The following window opens:



4 Enter the following task details:

- Scan Task Identifier** A free-text name for the task.
- Time Zone** Select the time zone with which the task is associated. For execution, the task starts at the designated time, based on the location you specify here. For example, if you select America/New York in the drop-down list and you specify that the task should start at 10:00 AM, then the task will run when it is 10:00 AM in New York.
- Max Task Duration** The limit for the scan task. If the task reaches this limit, it is automatically terminated.

5 Select one of the following radio buttons to specify when the task should run:

- Weekly schedule** Schedules the task to run on a weekly schedule. Specify the hour and minutes when the task is to run. Specify the days of the week by selecting the corresponding check box(es).

☒ Weekly schedule

12 : 15 : AM

☒ Sunday ☒ Monday ☒ Tuesday ☒ Wednesday ☒ Thursday ☒ Friday ☒ Saturday

- One time** Schedules the task to run only once. Specify the date and time when the task is to run.

☒ One time

2012-6-22 12 : 15 : AM


- Monthly schedule** Schedules the task to run on a specific day during the month. Specify the day and time when the task is to run.

☒ Monthly schedule

1st 12 : 15 : AM

You can specify **Last** in the day drop-down list to designate the last day of the current month, since the last day of the month can vary.

- 6 Click **Next**. The following window is displayed:




- 7 Select one of the following radio buttons to specify the scope for the scan task:

Scan default scope	Select to scan all active management consoles, storage proxies, and host/database scanning groups. This is the default setting.
--------------------	---

Scan specific groups


Select to specify specific management consoles, storage proxies, and host/database scanning groups for scanning, and then click **Next** to open the following window (for example):




- Select the **Scan all management consoles and storage proxies** check box to scan all enabled management consoles and storage proxy servers.
Or, specify the specific management consoles and storage proxy servers to be scanned. Leave those that you do not want to scan in the left pane and use the  arrow to move those to be scanned to the right pane.
- Select the **Scan all hosts in scope** check box to scan all hosts that are members of enabled scanning groups.
Or, specify the specific hosts to be scanned. Leave those that you do not want to scan in the left pane and use the  arrow to move those to be scanned to the right pane.
- Select the **Scan all databases in scope** check box to scan all databases that are members of enabled scanning groups.
Or, specify the specific databases to be scanned. Leave those that you do not want to scan in the left pane and use the  arrow to move those to be scanned to the right pane.

Scan specific business entities

Select to specify the business entities for scanning and then click **Next** to open the following window (for example):




- Select the **Scan all management consoles and storage proxy servers** check box to scan all enabled management consoles and storage proxy servers. Or, specify the specific management consoles and storage proxy servers to be scanned. Leave those that you do not want to scan in the left pane and use the  arrow to move those to be scanned to the right pane.
- In the Scan the business entities you want to scan area, select the business entity(ies) to be scanned by selecting the respective check box(es).



Use the  box for filtering to easily find a business entity of interest.

- 8 Click **Next**. The following window is displayed, which summarizes the configuration for the scan task (for example):



- 9 Click **Finish** to complete the scan task definition.

Scheduling a data analysis task


A data analysis task is a task that includes mapping and modeling dependencies between scanned IT elements. It discovers high-availability and data-protection configurations, as well as detects risk signatures. It does not include data collection.

For more details about adding a scan task or a full cycle task, see "[Scheduling a scan task](#)" on page 132 and "[Scheduling full cycle scans](#)" on page 139, respectively.

To add a data analysis task

- 1 Select the Configuration tab.
- 2 In the tree hierarchy, click **Scan Scheduling**. The Scan Scheduling page opens.

- 3
- Click **Add Data Analysis Task**  to add a data analysis task. The following window is displayed:



- 4
- Enter the following task details:

Task Identifier	A free-text name for the task.
Time Zone	Select the time zone with which the task is associated. For execution, the task starts at the designated time, based on the location you specify here. For example, if you select America/New York in this drop-down list, and you specify that the task should start at 10:00 AM, then the task will run when it is 10:00 AM in New York.

- 5
- Select the scheduling configuration type by selecting one of the following radio buttons:
- Weekly schedule
- One time
- Monthly schedule
- 6
- Click **OK** to save the data analysis task definition.

Scheduling full cycle scans


The full cycle scheduling configuration specifies when the product executes a full system scan and how often. A full system scan includes a comprehensive scan of your entire environment using all the enabled configured proxies. It also applies enhancement rules and identifies gaps. This process is repeated for each full system scan.

Note: You can define multiple full cycle scheduling configurations; however, you can only run one full system scan at a time.

To configure a full cycle scheduled task

- 1 In the Configuration tab, in the tree hierarchy, click **Scan Scheduling**.


- 2 Click **Add Full Cycle Scheduling** () to open the following wizard:



The image shows a 'Full Cycle Schedule' configuration window. It has a title bar with a standard Windows icon and a close button. The main area is titled 'Full Cycle Scheduling Configuration'. It contains a text field for 'Full Cycle Identifier:' and a dropdown menu for 'Default: DRA Master Server Time Zone'. Below these are three radio buttons: 'Weekly schedule' (selected), 'One time', and 'Monthly schedule'. Under 'Weekly schedule', there are two dropdowns for time (12 and 15) and a unit dropdown (AM). Below the time settings are seven checkboxes for days of the week: Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, and Saturday, all of which are checked. At the bottom, there are four buttons: 'Back' (with a left arrow), 'Next' (with a right arrow), 'Finish' (with a checkmark), and 'Cancel' (with an X).

- 3 In the **Full Cycle Identifier** field, enter an identifier for the scheduling configuration.
- 4 Select the time zone with which the task is associated. For execution, the task starts at the designated time, based on the location you specify here. For example, if you select America/New York in this drop-down list, and you specify that the task should start at 10:00 AM, then the task will run when it is 10:00 AM in New York.
- 5 Select the scheduling configuration type by selecting one of the following radio buttons:
 - Weekly schedule
 - One time
 - Monthly schedule


6 Click **Next**. The following window opens:



7 Select the scanning scope for this full cycle scheduling configuration by selecting one of the following radio buttons:


- **Scan default full cycle scope**
The default scope includes all enabled management consoles and storage proxies, and all hosts and databases that are assigned to enabled scanning groups.
- **Scan specific groups**
This option enables you to select specific scanning groups for scanning. In addition, you can select whether to scan the storage and the management consoles.
- **Scan specific business entities**
This option enables you to scan only servers and databases assigned to specific business entities. Note that when you select this option, the storage arrays storing these hosts and databases are also scanned.
- **Don't scan anything**
This option skips the scanning stage and proceeds directly to data analysis. This option provides seamless backwards compatibility.

- 8 Click **Next**. The following window is displayed, in which you specify the scope for the full cycle task (for example).



- 9 In the Scan all management consoles and storage proxy servers area, select the check box to scan all enabled management consoles and storage proxy servers.
Or, specify the specific management consoles and storage proxy servers to be scanned. Leave those that you do not want to scan in the left pane and use the  arrow to move those to be scanned to the right pane.
- 10 In the Scan all hosts in scope area, select the check box to scan all hosts that are members of enabled scanning groups.
Or, specify the specific hosts to be scanned. Leave those that you do not want to scan in the left pane and use the  arrow to move those to be scanned to the right pane.
- 11 In the Scan all databases in scope area, select the check box to scan all databases that are members of enabled scanning groups.
Or, specify the specific databases to be scanned. Leave those that you do not want to scan in the left pane and use the  arrow to move those to be scanned to the right pane.

- 12 Click **Next**. The following window is displayed, which summarizes the full cycle scheduled task configuration (for example).




- 13 To execute data analysis at the end of scanning, select the **Execute data analysis** check box.
- 14 Click **Finish** to save the full cycle scheduling configuration.

Report scheduling

DRA lets you schedule when reports are automatically generated and sent to one or more email destinations that you configure. You may also choose to save the generated reports in the file system and access them later.

To configure the schedule by which reports are emailed

- 1
- To access the Configuration page, click the Configuration tab.
- 2
- In the tree hierarchy, click **Report Scheduling**. The Report Scheduling window opens.



This window shows a row for each report that you schedule to be emailed. It provides the following options:

- Add

Lets you create a new schedule for emailing a report, as described in [“Adding a new email report schedule”](#) on page 145.
- Edit

Lets you edit the selected email report schedule. The Report Task Properties window is displayed, as described in [“Email settings”](#) on page 147.
- Delete

Deletes the selected report schedule definition.
- Execute Now

Generates and sends the selected report immediately even if its scheduled time has not yet arrived and even if the selected report's schedule is not enabled. You may use this option to test your email report schedule definition or when an emergency report must be sent out immediately to the defined email recipients.

Enable All Scheduled Reports  Enables all scheduled reports.

Disable All Scheduled Reports  Disables all scheduled reports. When using this option, DRA stops sending reports.

Default Distribution List  Defines the default distribution list to be used when creating new scheduled reports. The separator between multiple emails is a semicolon (;).


Default distribution list:
manager@company.com

Adding a new email report schedule

You can create new email report schedules at anytime. A schedule defines which report is sent, where the email is sent, and how often it is sent – daily, weekly or monthly. After you define and enable an email report schedule, the report is automatically generated and sent at the scheduled time. To stop transmissions you can uncheck the Enabled check box or delete the email report schedule.

To define a new email report schedule

- 1 In the Report Scheduling window, click **Add** () to display the Report Task Properties window:



Report Task Properties

Report Details

Name: Ticket Details

Use Filter: Recently opened | Opened in the last: day

Status: Open | Reopen | Severity: All | Category: General

☒ Show Summary ☐ Show Categories And Severity Distribution

☐ Show Closure Date

☐ Show First Detection Date

☐ Show Suppression Date

☒ Show Details ☒ Show Impact ☒ Show Resolution

☒ Show Notes ☐ Show Topology

☐ Show Rating ☐ Show Unique ID

☒ Show Ticket Summary

Report output: Send by Email | Report format: PDF

Email Details

Recipients: myname@mycompany.com

Subject: Tickets opened in the last day

Body:

Scheduling Details

Daily | Send report every day at 0 : 0

☒ Enable ☐ Send if empty

OK Cancel

- 2 In the **Report** field, select the report to send. Most of the DRA reports described in [Chapter 9, “DRA reporting”](#) on page 317 are provided for your selection.
- 3 After you select a report, criteria are displayed for you to define the data to be included in the report, such as host names or date ranges. You should fill in these fields, as described for each report in [Chapter 9, “DRA reporting”](#) on page 317.

A system property controls the email gateway.
For more information, see [“Configuring system properties”](#) on page 235.
- 4 In the **Report output** field, select whether to send the generated report by email, or save it in the file system. If you choose to save the report to the file system, you need to provide the directory path where the report should be saved.

-- Webpage Dialog

Report Task Properties

Report Details

Name: Ticket Details

Use Filter: Recently opened Opened in the last day

Status: Open | Reopen Severity: All Category: General

☒ Show Summary ☐ Show Categories And Severity Distribution

☐ Show Closure Date

☐ Show First Detection Date

☐ Show Suppression Date

☒ Show Details ☒ Show Impact ☒ Show Resolution

☒ Show Notes ☐ Show Topology

☐ Show Rating ☐ Show Unique ID

☒ Show Ticket Summary

Report output: Save to file system Report format: PDF

Output Folder

Save to folder: C:\MyReports

Scheduling Details

Daily Send report every day at 0 : 0

☒ Enable ☐ Send if empty

OK Cancel

- 3 To edit the fields, click **Edit** ().
- 4 To verify the mail server and authentication configuration, enter your email address and click **Send Test Email** ().

Configuring policies

A policy is a virtual item that defines how to communicate with a proxy, host, or database, such as proxy, sudo, and so on.

When you configure policies, you can also configure the following:

- Credentials
See [“Configuring credentials”](#) on page 151.
- Proxy
See [“Configuring a proxy”](#) on page 156.
- Sudo
See [“Configuring sudo”](#) on page 157.

DRA includes a default policy for each of these types.


Adding a policy

If you do not want to use the default policy, you can create a new policy for a particular host probe by doing one of the following:


- Performing Step 1 of the Configuration Wizard.
See [“Management console definition – configuring EMC ECC, Hitachi Data Systems \(HDS\) HiCommand, IBM HMC, and VMware vCenter”](#) on page 72.
- Selecting the Policies option in the tree in the Configuration window. The procedure for adding a policy is described below.

To add a policy

- 1 To access the Configuration page, click the Configuration tab. You can also access this page from the Configuration Wizard.
- 2 In the tree hierarchy, click **Policies** and **Add Policy** (). The following window opens:



- 3 In the **Name** field, enter a policy name.
- 4 Select the connection protocol type with which to connect to the host. When you select the connection protocol type, the  button appears next to the field, letting you add a specific protocol as required.



- 5
- (Optional) If the protocol you need is not in the list or your connection requires the use of a proxy machine or a sudo mechanism, click the  button corresponding to the protocol, sudo, or proxy to add it. Depending on your selection, one of the following windows opens:

protocol

Protocol Properties

Name *

My Custom SSH

Type *

SSH

Port *

1134

OK

Cancel

sudo

Sudo Properties

Name *

Sesudo

Path *

/usr/bin/sesudo

Credential

ECC default credential

PTY

System Default

OK

Cancel

proxy

Proxy Properties

Name *

IP *

Protocol *

Credential *

OK

Cancel

- 6 Enter the details and click **OK**. The window closes, and the new protocol, sudo, or proxy is added to the drop-down list.
- 7 To select the credential to attach to the selected policy, highlight it in the **Available Credentials** list and click the  right arrow. The credential appears in the **Credentials in Policy** list.
If the credential you need does not appear in the list, click **Add Credential** () to define a new one.
For information, see “[Configuring credentials](#)” on page 151.

Note: To delete a credential from the policy, select the credential in the

Credentials in Policy list and click the  left arrow. The credential then appears in the **Available Credentials** list.

- 8 In the Policy window, click **OK**. The window closes.

Editing a policy

This section describes how you can edit a policy using the Policies branch in the tree in the Configuration window.

To edit a policy

- 1 To access the Configuration page, click the Configuration tab.
- 2 In the tree hierarchy, click **Policies**. The Policies window is displayed.
- 3 Select a policy and click **Edit Policy** ()

Configuring credentials


The Credentials page lists all the credentials defined in the system for all management consoles, proxies, hosts, and databases. You can use this page to add, change, or delete credentials.

You can add credentials in the following ways:

- Using the Credentials page
- Importing credential information from a comma-separated values (CSV) file

To add a new credential using the Credentials page

- 1
- To access the Configuration page, click the Configuration tab.
- 2
- In the tree hierarchy, under Policies, click **Credentials**. The Credentials page opens.



- 3
- Click **Add Credential** (). The following window opens:

A screenshot of a 'Credential Properties' dialog box. The dialog has a title bar that says '-- Webpage Dialog'. Inside, there's a section titled 'Credential Properties:'. It contains four fields: 'Credential Type:' with a dropdown menu showing 'Username/Password', 'Description' (empty text box), 'Username' (empty text box), and 'Password' (empty text box). Below the Password field is a 'Retype password' field. At the bottom of the dialog are 'OK' and 'Cancel' buttons with green and red checkmarks respectively.

4 Select one of the following credential types:

Username/Password	Typically used for entities such as hosts and databases
Windows Authentication	A specific add-on for MS SQL server scanning
CLARiiON	Used to define CLARiiON credentials
SSH Public Key Authentication	Used when SSH public key authentication is needed
Cyber-Ark	Used when using Cyber-Ark to manage credentials
Rotating Password	Used to define a mechanism by which the password is automatically changed.

5 In the **Description** field, enter a description for the credential.

6 In their respective fields, enter the credential's user name and password.

7 If a CLARiiON credential was selected, select a scope for the credential based on the following:

Select Global for scope 0	Authorization is checked against the CLARiiON domain.
Select Global for scope 1	Authorization is checked against the scanned CLARiiON array.
Select LDAP for scope 2	A user account exists in the LDAP directory and applies to any storage system that is configured to authenticate against the LDAP server (for example, Active Directory).

8 To save your settings and exit the window, click **OK**. The credential is added to the DRA master list of credentials.


To add a new credential by importing a CSV file

1 To access the Configuration page, click the Configuration tab.

2 In the tree hierarchy, under Policies, click **Credentials**.

3 Click **Import Credentials from CSV file** (.

- 4
- On the Import Credentials from CSV file screen, review the file format. Your file must match this format exactly.



- 5
- Create a CSV using a text editor or by saving an Excel spreadsheet to CSV format. Your file must include the following:

Heading	The first line of your file must be the following: Type,Description,Username,Password,Scope Even if your credential does not include all these attributes, you must include them in the heading row.
Type	The credentials type must be one of the following, capitalized and spaced exactly as shown. ■ Username/Password ■ WindowsAuthentication ■ CLARiION
Description	The description field can contain any text you want. If you omit the description, include two commas between the Type and Username fields.
Username	The username of the credential.

Password	The password of the credential. If Password is the last field in your credential, you must add a comma after it. See the sample CSV file in step 4 on page 154.
Scope	The Scope field only applies if you are creating CLARiiON credentials. This field must be one of the following, capitalized exactly as shown: ■ Global ■ Local ■ LDAP

- 6 Save your file with the extension .csv.
- 7 On the Import Credentials from CSV file screen, click **Import from file**.
- 8 Browse to your file, select it, and click **OK** to import it.

To delete a credential

- 1 To access the Configuration page, click the Configuration tab.
- 2 In the tree hierarchy, under Policies, click **Credentials**. The Credentials page is displayed.
- 3 Select a credential in the list and click **Delete Credential** (). The credential is deleted from the list.

To edit credential properties


- 1 To access the Configuration page, click the Configuration tab.
- 2 In the tree hierarchy, under Policies, click **Credentials**. The Credentials page is displayed.
- 3 Select a credential in the list and click **Edit Credential** (). The Credential Properties window is displayed.
- 4 Modify the credential's properties as needed.
- 5 To save your changes and exit the window, click **OK**.

Configuring a proxy

The Proxy page lists all the proxies defined in the system. You can use this page to add, change, or delete proxies.

To add a new proxy

- 1 To access the Configuration page, click the Configuration tab.
- 2 In the tree hierarchy, under Policies, click **Proxy**. The Proxy page is displayed.
- 3 Click **Add Storage Proxy** (). The following window opens:



The screenshot shows a window titled "proxy" with a tab labeled "Proxy Properties". Inside the window, there are four labeled text input fields: "Name *", "IP *", "Protocol *", and "Credential *". Each field has a corresponding input box. The "Credential *" field has a dropdown arrow on its right side. At the bottom of the window, there are two buttons: "OK" with a green checkmark icon and "Cancel" with a red X icon.

- 4 Enter the proxy name, IP address, protocol, and credential.

Note: This protocol is used by the proxy machine to connect to the desired host, and not the one that is used to connect to itself.

- 5 To save your settings and exit the window, click **OK**. The proxy is added to the Proxy list.

To delete a proxy

- 1 To access the Configuration page, click the Configuration tab.
- 2 In the tree hierarchy, under Policies, click **Proxy**. The Proxy page is displayed, as shown above.
- 3 Select a proxy in the list and click **Delete Proxy** (). The proxy is deleted from the list.

To edit proxy properties


- 1 To access the Configuration page, click the Configuration tab.
- 2 In the tree hierarchy, under Policies, click **Proxy**. The Proxy page is displayed, as shown above.
- 3 Select a proxy in the list and click **Edit Storage Proxy Definitions** (). The Proxy Properties window opens.
- 4 Modify the proxy's properties as needed.
- 5 To save your changes and exit the window, click **OK**.

Configuring sudo

The Sudo page lists all sudos defined in the system. You can use this page to add, change, or delete sudos. DRA provides a default sudo.

To add a new sudo

- 1 To access the Configuration page, click the Configuration tab.
- 2 In the tree hierarchy, under Policies, click **Sudo**. The Sudo page is displayed.
- 3 Click **Add Sudo** (). The following window opens:



Sudo Properties	
Name *	Sesudo
Path *	/usr/bin/sesudo
Credential	ECC default credential
PTY	System Default

OK  Cancel 

- 4 Enter the sudo name, path, credential, and PTY.
- 5 To save your settings and exit the window, click **OK**. The sudo is added to the Sudo list.

To delete a sudo

- 1 To access the Configuration page, click the Configuration tab.
- 2 In the tree hierarchy, under Policies, click **Sudo**. The Sudo page is displayed.
- 3 Select a sudo in the list and click **Delete Sudo** (). The sudo is deleted from the list.

To edit sudo properties

- 1 To access the Configuration page, click the Configuration tab.
- 2 In the tree hierarchy, under Policies, click **Sudo**. The Sudo page is displayed.
- 3 Select a sudo in the list and click **Edit Sudo** (). The Sudo Properties window is displayed.
- 4 Modify the properties of the sudo as needed.
- 5 To save your changes and exit the window, click **OK**.

Configuring business entities

DRA lets you organize items into logical groups called business entities. Such groups are useful because they let you classify various items in your environment for a particular purpose. For example, you could group all the billing applications in your environment into a business entity called Billing. This group can contain many hosts and databases.

You can also create hierarchical relationships between business entities, where one business entity may contain one or more child business entities.

You can add, change, or delete business entities.

Note: An item can belong to only one site. However, an item may belong to multiple business entity groups.

Configuring a business entity is similar to configuring a site, the one exception being that you can assign multiple business entities to a single host. A business entity may include hosts, databases, and/or file systems. During configuration, you define the criteria for mapping items to the business entity. The Business Entity Configuration wizard guides you through the configuration process.

You can perform the following operations from the Business Entity Management page:

Add Business Entity		Adds a new business entity
Edit Business Entity		Modifies the properties of a business entity and changes the hosts/databases association
Delete Business Entity		Deletes a business entity
Add Business Entity Using Wizard		Adds a new business entity using the Business Entity Wizard described below
Edit Business Entity Using Wizard		Modifies the properties of a business entity and changes the hosts/databases association using the Business Entity Wizard described below
Import Business Entity from External Database		Imports business information from an external data source
Import Business Entities		Imports business entities from a *.CSV file


Note: Configuring a business entity external data source involves view/table creation in the organization’s business management systems. To learn more about external data source configuration, follow the instructions provided in the popup link.

There are two ways to define business entities. The first one provides a simple configuration that enables you to define the business entity and its contained hosts and databases. The second method uses the advanced Business Entity Wizard, which enables you to define advanced properties, assign hosts, databases, VGs, file systems, NetApp filers, volumes, and Q-trees to the business service, and to select the role and SLA policy for each contained resource.

To configure a business entity


- 1
- To access the Configuration page, click the Configuration tab.
- 2
- In the tree hierarchy, click **Business Entities**. The Business Entity Management window opens.

Note: Only Administrators can access this page.




The right pane lists each business entity defined in DRA, including its name, description, type, and the number of connected items.

- 3
- To define a new simple business entity, click **Add Business Entity** (). The Add Business Entity window opens.



- 4 In the **Name** field, enter a business entity name. This name identifies the business entity throughout DRA.
- 5 In the **Description** field, enter a description for the business entity. You can now define related-item information for the business entity.
- 6 In the **Type** field, select a business entity type. The options listed in the drop-down list are defined in the Business Entity Types branch.
- 7 In the tabs below the Type field, you can assign hosts and databases to the business entity. Select the hosts from **Available Hosts** list and click  to assign them to the business entity. In the same way, you can assign databases in the Databases tab.
- 8 To remove a host from the business entity, select the host from the **Hosts assigned to Business Entity** list and click . You can remove a database from the business service in the same way, in the Databases tab.
- 9 To define a new business entity using the wizard, click **Add Business Entity**

Using Wizard (). The Business Entity Configuration wizard opens displaying the General Details page:



Business Entity Configuration wizard, General Details page.

Webpage Dialog

Business Entity Details

1 General Details | 2 Contained Business Entities | 3 Contained Resources | 4 Roles & SLA Configuration

Details

Name: DWH

Description: Marketing analysis subsystem


Type: Business Service

☒ Top Level

Back Next Finish Cancel

- 10 In the **Name** field, enter a business entity name. This name identifies the business entity throughout DRA.


- 11 In the **Description** field, enter a description for the business entity. You can now define related-item information for the business entity.
- 12 In the **Type** field, select a business entity type. The options listed in the drop-down list are defined in the Business Entity Types branch.
For more information, see [“Defining business entity types”](#) on page 168.
- 13 Optionally, if you want this business entity to be connected directly to the business entity tree root node, select the **Top Level** check box. If you want child business entities to automatically be displayed on the Dashboard, you should also select this check box. Click **Next**. The Contained Business Entities page opens.



Use this window to select another business entity that is contained in this business entity. A tree hierarchy displays all previously-defined business entities.

- 14 In the tree hierarchy, select a business entity. If you select a parent business entity, its children are automatically assigned to this business entity. If you select a child business entity, only it is assigned to this business entity.


15 Click **Next**. The Contained Resources page opens:



Use this window to specify the resources, including databases, hosts, file systems, and NetApp filers that belong to this business entity. Assigning resources to a business entity lets you easily identify where a problem occurred in the organization. A tree hierarchy lists the resources.

16 Select the required resource(s) in the tree.

17 Click **Next**. The Roles & SLA Configuration page opens.




Use this page to assign roles and SLA policies to the resources of this business entity.

For more information, see [“Defining business entity roles”](#) on page 170 and [“Defining SLA policies”](#) on page 172.

18 From the **Role** drop-down list, select the role of the resources inside the business entity (for example, Production, DR, Standby, Reporting, and so on). You can add custom roles and determine the importance of each role, as described in [“Defining business entity roles”](#) on page 170.

If you want to assign an SLA policy to a business entity resource, you can select the policy from the drop-down list. After an SLA policy is assigned to the resource, DRA ensures that the resource is enforced in the policy requirements. Any deviation from the SLA policy results in a ticket. You can create custom SLA policies according to your environment requirements in the SLA Definition page. For more information, see [“Defining SLA policies”](#) on page 172.

- 19 To complete configuring the business entity, click **Finish**. The business entity appears in the Business Entities tree.




Note: All Top Level business entities are connected directly to the root node. All parent business entities are automatically Top Level. Child business entities that are defined as Top Level entities are shown in two places in the tree hierarchy: directly under the root and as a child entity under its parent. Indicating Top Level business entities in this manner lets you quickly see them in the tree, even when the tree is collapsed.


To edit business entity resource roles and SLA policies

- 1 To access the Configuration page, click the Configuration tab.
- 2 In the tree hierarchy, click **Business Entities**. The Business Entity Management window opens.

- 3
- Select the business entity you want to change, and then select the Roles & SLA Definition tab. The following page opens:



- 4
- In the Roles & SLA Definition tab, you can select the resources you want to change. To do so, click **Edit Role and SLA Definition** (). The following window opens:




- 5
- Select the required role and/or SLA policy, and click **OK**. The role and/or SLA policy is changed for all the selected resources.

To add a new business entity by importing a CSV file

- 1
- To access the configuration page, click the Configuration tab.
- 2
- In the tree hierarchy, click **Business Entities**. The Business Entity Management window opens.

- 3 Click **Import Business Entities from CSV file** (). The Import Business Entities from CSV file page opens.



CSV file format:

- Business Entity** *: Business Entity name.
Note: If the business entity does not exist, the system will create it.
- Business Entity Type**: Not mandatory, unless the business entity does not exist - this field determines the type of the newly created entity.
Note:
 - If the business entity exists, this field is ignored.
 - If the business entity does not exist, and the type is empty - the business entity will not be created.
- Host**: The host name to assign to the business entity (see sample figure, lines 2,3).
- Site**: If the host name is unique across all sites - this field is not mandatory. Otherwise (e.g. virtual machines with the same name on different sites), this field helps the system to determine the correct host to assign to the business entity, the field must contain the name of the site exactly as it appears in the Configuration Wizard, step 3. (see sample figure, lines 2,3).
- Database**: The database name to assign to the business entity. If you want to assign host only - this field value should stay empty.
Note: When assigning database to business entity, the host will not be assigned to the business entity. To assign the host you must add a specific row for the host, without any database. (see sample figure, lines 4,5).
- Role**: Not mandatory. If not provided, the system will use the default role. If the role does not exist, the system will create it with the rating given in "Role Rating". (see sample figure, lines 2-8).
- Role Rating**: Not mandatory. If the role does not exist, the system will create it with the rating provided in this field. (see sample figure, line 8). If this field is empty, the default value is 4.
Note: If the role exists, this field will be ignored.
- SLA Policy**: Not mandatory. This field represents the name of the SLA policy to assign to the host/database for this business entity (see sample figure, lines 2-8). If the SLA policy does not exist, the system will not create it, and the value will be ignored.
- Sub Business Entity**: Not mandatory. To create hierarchical business structure, you can add sub-business entities under other business entities (see sample figure, lines 9,10). If the sub-business entity does not exist, the system will not create it. Therefore, if the sub business entity is a new entity, it must appear previously in the file.
Note: When you assign sub-business entity, the rest of the fields (host, database, role, SLA policy, etc.) are ignored.

	A	B	C	D	E	F	G	H	I
	Business Entity	Business Entity Type	Host	Site	Database	Role	Role Rating	SLA Policy	Sub Business Entity
1	SAP	Application	erpapp1	New York		Production		Tier 0 PRD	
2	SAP	Application	erpapp2	New York		Production		Tier 0 PRD	
3	SAP	Application	erpback	New York		Backup		Tier 0 BCK	
4	SAP	Application	erpback	New York	ORCL	Backup		Tier 0 BCK	
5	SAP	Application	erpdr	Chicago		DR		Tier 0 DR	
6	ERP DEV	Division	erpdevenv	Chicago		Development		Tier 6 Dev	
7	ERP DEV	Division	erpstenv	Chicago		CustomRole	4	Tier 6 Test	
8	ERP	Business Service							SAP
9	ERP	Business Service							ERP DEV
10	ERP	Business Service							

[Download sample](#)

 If the business entities are not added, please verify that the file format is according to the instruction above

- 4 On the Import Business Entities from CSV file screen, review the file format. Your file must match this format exactly.
- 5 Create a CSV file using a text editor or by saving an Excel spreadsheet to CSV format.
- 6 Save your file with the extension .csv.
- 7 On the Import Business Entities from CSV file screen, click **Import from file**.
- 8 Browse to your file, select it, and click **OK** to import it.

Defining business entity types

Business entity types categorize business entities into groups, based on your criteria. When you define these groups, you may use any criteria you like. You can perform the following operations from the Business Entity Type Management page:

- Add Business Entity Type



Adds a new business entity type
- Edit Business Entity Type



Modifies the properties of a business entity type
- Delete Business Entity Type




Deletes a business entity type

To define a business entity type


- 1 To access the Configuration page, click the Configuration tab.
- 2 In the tree hierarchy, click **Business Entity Types**. The Business Entity Type Management window opens.

Note: Only Administrators can access this page.



In the right pane, this page lists each business entity type defined in DRA, including its name and description.

- 3 To define a new business entity type, click **Add Business Entity Type** (). The Business Entity Type window opens.



- 4 In the **Name** field, enter business entity type name. This name identifies the business entity type throughout DRA.
- 5 In the **Description** field, enter a description for the business entity type.
- 6 Select an icon to represent the business entity type and click **OK**. The business entity type appears in the Business Entity Type Management window.

Defining business entity roles

Business entity roles rate a business entity to indicate its importance in the system. DRA uses a 5-star system. One star (★) indicates the lowest level of importance and five stars (★★★★★) the highest.

You can perform the following operations from the Roles Management page:

- Add Business Entity Role



Adds a new business entity role
- Edit Business Entity Role




Modifies the properties of a business entity role
- Delete Business Entity Role



Deletes a business entity role


To define a business entity role

- 1 To access the Configuration page, click the Configuration tab.
- 2 In the tree hierarchy, click **Business Entity Roles**. The Roles Management window opens.



In the right pane, this page lists each business entity role defined in DRA, including its role name and its rating. The default role has a green checkmark  next to it.

- 3 To define a new business entity role, click **Add Role** (). The Role window opens.



Webpage Dialog

Role

Role Properties

Name

☐ Set as Default Role

☒ Calculate for SLA
Items assigned to business entities with role that is not marked with **Calculate for SLA** are ignored when calculating the high-level replication SLA data. Setting this property correctly is the key for usable and informative high-level SLA view.

Please rate the role importance ★★★★★

OK  Cancel 

- 4 In the **Name** field, enter a business entity role name. This name identifies the business entity role throughout the application.
- 5 Optionally, to automatically assign this business entity the default role, select the **Set as Default Role** check box.
- 6 To indicate the importance of this role, click the relevant number of stars. The more stars, the more important the role. Click **OK**. The business entity role appears in the Role Management window.

Defining SLA policies

SLA policies define a set of data protection and availability properties. These properties can be applied on business entity resources. Using the Rule engine, DRA verifies an SLA policy to ensure that its components are enforced on business entity resources. Any deviation from the defined SLA results in a ticket.

You can perform the following operations from the SLA Policies page:

- Add SLA Policy



Adds a new SLA policy
- Clone SLA Policy



Clones the selected SLA policy as a new editable policy
- Edit SLA Policy



Modifies the properties of an SLA policy
- Delete SLA Policy




Deletes an SLA policy


To define an SLA policy

- 1 To access the Configuration page, click the Configuration tab.
- 2 In the tree hierarchy, click **SLA Definition**. The SLA policies window opens.

Note: Only Administrators can access this page.



- 3 To define a new SLA policy, click **Add SLA Policy** (). The SLA Policy window opens.




The screenshot shows the 'SLA Policy' window with the following details:


- SLA Policy properties**
 - Name: Tier 0
 - Description: Production Tier 0 SLA
- Tabs:** Data Protection (selected), Availability, Standby/DR Capacity Plan, Database Replication
- RPO (double-click any item to edit)**
 - Buttons: Add RPO

Enforce	RPO	Site	Replication Path	Remove
☒	Up to date	Local	SRDF	
- Retention (double-click any item to edit)**
 - Buttons: Add Retention

Enfo...	Frequency	Number of Copi...	Copy Type	Site	Replication Path	Seen By Hosts	Remove
☒	Hourly	24	Logically protec...	Local	Snapshot	-	
☒	Daily	6	Physically prote...	Local	BCV	-	
- Buttons:** OK, Cancel

- 4 In the **Name** field, enter the SLA policy name. This name identifies the SLA policy throughout DRA.
- 5 In the **Description** field, enter a description for the SLA policy (for example, Remote RPO 0, 3 local daily copies).


- 6 In the Data Protection tab, you can add Recovery Point Objective (RPO) and retention definitions. To do so, click . A new RPO definition is added to the table. Double-click an RPO definition to edit it. The following window opens:




The image shows a 'Webpage Dialog' window titled 'RPO'. It contains the following fields and options:

- RPO:** A dropdown menu set to 'Up to date' with a green '+' icon to its right.
- Description:** Text explaining RPO (Recovery Point Objective) as the acceptable amount of data loss in the event of recovery (measured as time-delta). It instructs to select the desired RPO from the list (to add your own RPO, click the "+" button). A note specifies: 'Note: for synchronous replication choose **Up to date**.'
- Site:** A dropdown menu set to 'Local'.
- Instructions:** 'Select a site to which this RPO apply:' followed by a bulleted list:
 - You can select any specific site defined in the application; OR;
 - Local:** The same site of the component (host, filer, and so on) this policy is assigned to; OR;
 - Any remote:** Any site other than **Local**
- Replication path:** A dropdown menu set to 'Don't care' with green '+' and red '-' icons to its right.
- Instructions:** 'Specify the replication path for this RPO configuration by clicking "+".'
- Enforce:** A checked checkbox labeled 'Enforce'.
- Description:** 'Check this option to generate a ticket whenever this RPO policy is violated'
- Buttons:** 'OK' with a green checkmark icon and 'Cancel' with a red X icon.

- 7 To save the RPO changes, click **OK**.

- 8 To create a new retention definition, click . The new retention definition is added to the table. To edit a retention definition, double-click it. The following window opens:



Retention

Frequency: 

The frequency of Point in Time copies.
Select the desired frequency from the list, or add your own to the list by clicking the "+" button.

Number of copies:

The minimum number of copies to retain, in the frequency specified above.
For example: if recovery to any of the last 24 hours is required, you should specify 24 x Hourly.

Copy type:

The type of copy can be one of the following:

- **Physically protected:** indicates a full, high-performance copy is required (in other words, not a snapshot, or partial clone). Even if the source is physically damaged, the copy will remain usable.
- **Logically protected:** indicates the copy can be linked, or dependent on source. This is usually faster, but does not ensure the same data protection, or performance as **Physically protected**.

Site:

Select a site to which this Retention apply:

- You can select any specific site defined in the application; OR;
- **Local:** The same site of the component (host, filer, and so on) this policy is assigned to; OR;
- **Any remote:** Any site other than **Local**

Replication path:  

Specify the replication path of the point-in-time copy you want to monitor by clicking "+".

Seen by hosts: 

Specify the hosts you would like make sure that the point-in-time copies are mapped to.

☒ Enforce

Check this option to generate a ticket whenever this retention policy is violated

- 9 To save the retention definition, click **OK**.

- 10
- In the Availability tab, you can define the minimum number of SAN I/O paths, the minimum number of LUN maps, and the RAID level.

-- Webpage Dialog

SLA Policy

SLA Policy properties

Name: Tier 0

Description: Production Tier 0 SLA

Data Protection

Availability

Standby/DR Capacity Plan

Database Replication

Availability properties

☒ Minimum number of I/O paths: 4

Enforces the minimum required number of SAN I/O paths from the server to the storage volumes. It is highly recommended to set this property to minimum of 2 I/O paths to avoid a single-point-of-failure and to gain better I/O performance.

☒ Standby number of I/O paths ratio: 80 %

Enforces the number of SAN I/O paths ratio between primary and standby servers (production vs. DR).
For example: if your production tier SLA requires minimum 8 SAN I/O paths, and you want to verify that the standby servers of this tier have a minimum of 4 I/O paths - set the value of this field to "50%".
Select the detection method for the standby servers:

☒ Enforce the ratio only on replicated storage volumes of the primary server, which are mapped to the standby server.
This configuration is suitable for **storage-based replication** environments.

☐ Enforce the ratio on all storage volumes **used** by the standby server.
This configuration is suitable for **host-based replication** environments or for active-active environments.

☐ Enforce the ratio on all storage volumes **mapped** to the standby server (both used and unused storage).
This configuration is suitable for **backup servers** that do not mount any volume group or file system on the storage. It provides a complete I/O redundancy assurance (to the ratio defined) for all storage volumes mapped to the standby server.

Note: this configuration will be enforced only on servers that have a **Standby Server**. To define the standby relationships, go to *Configuration Module - Standby Definition*.

☒ Minimum number of LUN maps: 2

Enforces the minimum required number LUN maps (array port mappings) of the SAN storage volumes used by the server. It is highly recommended to set this property to minimum of 2 LUN maps to avoid a single-point-of-failure and to gain better I/O performance.

☒ RAID level: RAID5

Enforces a specific RAID level of the SAN storage volumes used by the server.

OK

Cancel

- 11 In the Standby/DR Capacity Plan tab, you can define the minimum capacity ratio for standby hosts.

The screenshot shows a web-based configuration window titled "SLA Policy". It has a tabbed interface with four tabs: "Data Protection", "Availability", "Standby/DR Capacity Plan" (which is selected), and "Database Replication".

SLA Policy properties

Name: Tier 0
Description: Production Tier 0 SLA

Standby/DR Capacity Plan

- The *Standby/DR Capacity Plan* allows you to define the desired capacity ratio between the production and the standby servers.
- Each ratio configured in this section will be enforced only on servers that have a standby server, according to the **Standby** configuration. To define the standby relationships, go to *Configuration Module - Standby Definition*.
- You should fill this section only if this policy is intended to primary servers. Assigning a policy with *Standby/DR Capacity Plan* to standby servers will result in false tickets.

☒ Inform me (with a ticket) when hosts assigned to this policy have no standby server

Production vs. DR/Standby Server Capacity

☒ **Number of HBA ports ratio:** 50 %
The number of HBA ports installed on the standby server must be at least XX% of the number of HBA ports installed on the primary server.

☒ **Memory size ratio:** 70 %
Standby server must have at least XX% of the memory size (RAM) that the primary server has.

Production vs. DR/Standby Computing Power


☐ **Number of cores ratio:** %
The number of cores (CPUs) installed on the standby server must be at least XX% of the number of cores installed on the primary server.

☐ **CPU speed ratio:** %
CPU speed of the standby server must be at least XX% of the primary server CPU speed.

☒ **Total computing power ratio:** 80 %
Standby server total computing power must be at least XX% of the primary server.
Total computing power = (Number of CPUs) X (CPU speed).

At the bottom, there are "OK" and "Cancel" buttons.

12 In the Database Replication tab, you can define the preferred RPO and RTO for the Oracle Data Guard application.



13 To save the SLA policy, click **OK**.

The following is an example of a common SLA policy:

SLA Policy

SLA Policy properties

Name: Platinum

Description: RPO 0 at the DR site, with one copy every 30 minutes. Local retention using snapshots

Data Protection | Availability | Standby/DR Capacity Plan | Database Replication

RPO (double-click any item to edit)

+ Add RPO

Enforce	RPO	Site	Replication Path	Remove
☒	Up to date	Any remote	Don't care	

Retention (double-click any item to edit)

+ Add Retention

Enfo...	Frequency	Number of Copi...	Copy Type	Site	Replication Path	Seen By Hosts	Remove
☒	Daily	7	Logically protec...	Local	Don't care	-	
☒	Weekly	4	Logically protec...	Local	Don't care	-	
☒	Every 30 minutes	1	Physically prote...	Any remote	Don't care	-	
☒	Monthly	12	Logically protec...	Local	Don't care	-	
☒	Hourly	24	Logically protec...	Local	Don't care	-	

OK Cancel

Configuring expansion packages

DRA expansion packages enable you to add new functionality to the product. For example, you can run your own scripts or commands, or gather data specific to your environment.

Defining data collection packages

Data Collection Packages define what additional information should be collected from your servers. You can specify new configuration items such as files and directories (content and meta data), registry entries as well as output of scripts or commands.


Data Collection Packages can tightly integrate with Host Comparison Packages (see [“Defining host comparison expansion packages”](#) on page 193), Risk Detection Packages (see [“Defining risk detection expansion packages”](#) on page 198), and Report Packages (see [“Defining reporting packages”](#) on page 208). The new configuration items included in the package are automatically added to the System API (SAPI).

A Data Collection Package is a logical container of configuration items that you want to add and the methods for obtaining the required data. There are two basic types of Data Collection Expansion Packages: one for UNIX platforms and one for Windows. The following table lists the types of available configuration item types:

Table 6-2 Data collection package structure

Type	Supported Platform	Description
File Meta Data	UNIX, Windows	Contains the file privileges, ownership, create/modify date, checksum, and size.
File Content	UNIX, Windows	Text file content.
Directory Structure	UNIX, Windows	Directory structure and meta data.
Command	UNIX, Windows	Executes the specified command and command arguments, and captures the output.
Registry Entry	Windows	Registry entry values for the entries that match a specified pattern.
WMI Query	Windows	Collects data from WMI using a query statement.
VB Script	Windows	Executes a VB script on the Windows host.
Shell Script	UNIX	Executes a script on the UNIX host.

Multiple data collection expansion packages can be defined.




Each row in the Data Collection Packages page represents a separate data collection package. The following information is shown for each package:



Name	The name of the data collection package
Author	The author of the data collection package
Supported Platforms	The platforms on which the package can run
Description	A description of the package

You can perform the following operations from the Data Collection Packages page:

Add UNIX Platform Package		Adds a new UNIX collection package.
Add Windows Platform Package		Adds a new Windows collection package.
Edit Package		Modifies the properties of the data collection package.
Delete Package		Deletes the selected package.
Import Package		Imports the selected package.
Export Package		Exports the selected package.



Exports the Data Collection Expansion Package Summary report in either PDF or RTF format. Use this option to back up a package, or to share it with other parties. This report can be accessed [here](#), or from the Reports tab.



Specifies the permissions that the user must have in order to run the package.

To add a UNIX data collection package

- 1 Click **Add UNIX Platform Package** (). The following window opens:

A dialog box titled "Data Collection Package" with a "Webpage Dialog" title bar. It contains a "Properties" section with fields for "Package Name", "Author", and a "Description" text area. Below is an "Assigned Items" section with a list of items: "Sites", "Business Entities", "Clusters", and "Hosts", each with a "none" button and an icon. At the bottom is an "Add..." button and a table with columns "Configuration Item" and "Remove". The "OK" and "Cancel" buttons are at the very bottom.


- 2 In the **Package Name** field, enter a name for the package.
- 3 In the **Author** field, enter a name for the package author.

- 4 In the **Supported Platforms** area, select the check box(es) corresponding to the operating systems on which the package is supported.
- 5 In the **Description** field, enter a description for the package.
- 6 In the Assigned Items area, specify the hosts on which to run the package. This area acts like a filter to define the specific scope of the package, by selecting specific entities on which to run the package, including:
 - **Sites:** The package applies to all hosts located at the selected site.
 - **Business Entities:** The package applies to hosts that are members of the selected business entities (and their sub business entities).
 - **Clusters:** The package applies to nodes of the selected clusters.
 - **Hosts:** The package applies to the selected hosts.

Hosts from different item categories (if more than one is used) are joined to define the final scope for the package. For example, if you select *Boston* under Sites, and *ERP* under Business Entities, all hosts in Boston (even if not part of the ERP business entity), as well as all ERP hosts (even if not in Boston) are part of the package scope.

To specify a particular item, click the adjacent **None** link, to open a window in which you make your selection for that item (sites, business entities, clusters and hosts).

- 7 Click the  **Add...** button to define what the package does. This action opens the following drop-down list, in which you select the type of functionality the package is to perform.




You can select the following functionality types:

Command

Command functionality executes the specified command and command arguments and captures the output. The command must be installed on your servers in the specified path. You can also configure an Awk parser to parse the output and input the results of the parsing to DRA. To enable parsing, check the **Awk Script** check box, and then enter the script in the area below it. If the **Awk Script** check box is not checked, then no parser is used.

The screenshot shows a web browser window titled "-- Webpage Dialog" displaying a "Configuration" dialog box. The dialog has a blue header bar with the title "Configuration". Below the header, there is a section titled "Command" with a small icon. The text in this section reads: "Executes the specified command and command arguments, and captures the output. Note that the specified command must be installed on your servers, in the specified path. You can also configure an 'awk' parser that will be applied to the output before saving it in the RecoverGuard database." Below this text is a "Properties" section with three input fields: "Identifier", "Command", and "Command arguments". At the bottom of the dialog is a checkbox labeled "Awk Script". Below the checkbox is a large text area for entering the script. At the very bottom of the dialog are two buttons: "OK" with a green checkmark icon and "Cancel" with a red X icon.

Directory Structure Directory Structure functionality enables you to compare the structure of a directory, meaning its folders, subfolders, signatures, size, permissions and so on. Enter the name of the directory and indicate whether the structure is recursive (includes subfolders) or non-recursive (does not include subfolders).



The screenshot shows a 'Webpage Dialog' window with a 'Configuration' tab. Under the 'Directory Structure' section, there is a description: 'Directory structure and meta-data.' Below this, a 'Properties' section contains two fields: 'Directory name' with an empty text input box, and 'Recursive' with a dropdown menu currently set to 'Not Recursive'. At the bottom of the dialog are 'OK' and 'Cancel' buttons with green and red checkmarks respectively.

File Content

File Content functionality retrieves the contents of a file, and uses it for comparison. You can also configure an Awk parser to parse the output and input the results of the parsing to DRA. To enable parsing, check the **Awk Script** check box, and then enter the script in the area below it. If the **Awk Script** check box is not checked, then no parser is used.

-- Webpage Dialog

Configuration

File Content

Text file content. You can configure an 'awk' (after the famous Unix utility) parser that will be applied to the content before storing it in the RecoverGuard database.

Properties

Identifier

File name


☐ Awk Script

OK

Cancel


File Meta Data

File Meta Data functionality checks the meta data for a file, such as its privileges, ownership, creation date and so on.



Shell Script


Shell Script functionality executes the specified script on the UNIX host. You can also configure an Awk parser to parse the output and input the results of the parsing to DRA. To enable parsing, check the **Awk Script** check box, and then enter the script in the area below it. If the **Awk Script** check box is not checked, then no parser is used.



- 8 Click **OK** to save the package definition.


To add a Windows package


- 1 Click **Add Windows Platform Package** (). The following window opens:



The screenshot shows a 'Webpage Dialog' window titled 'Data Collection Package'. It features a 'Properties' section with input fields for 'Package Name', 'Author', and a multi-line 'Description'. Below this is an 'Assigned Items' section with a list of categories: 'Sites', 'Business Entities', 'Clusters', and 'Hosts', each accompanied by a '[none]' button. At the bottom of the dialog is an 'Add...' section containing a table with one header 'Configuration Item' and a 'Remove' button. The dialog concludes with 'OK' and 'Cancel' buttons at the very bottom.

- 2 Follow steps **2** through **6** in the *To add a UNIX data collection package* procedure on page 182.

- 3
- Click the  button to define what the package does. This action opens the following drop-down list, in which you select the type of functionality the package is to perform.




You can select the following functionality types:

Command	Command functionality for a Windows data collection package is the same as that for a UNIX data collection package. See page 184 for details.
Directory Structure	Directory Structure functionality for a Windows data collection package is the same as that for a UNIX data collection package. See page 185 for details.
File Content	File Content functionality for a Windows data collection package is the same as that for a UNIX data collection package. See page 186 for details.
File Meta Data	File Meta Data functionality for a Windows data collection package is the same as that for a UNIX data collection package. See page 187 for details.

Registry Entry

Registry Entry functionality compares registry values for the entries that match the pattern specified by the **Registry Key** field.



The image shows a web-based configuration dialog box titled "Registry Entry". The dialog has a title bar with a close button. Inside, there's a section header "Registry Entry" with a green checkmark icon. Below it, a text description reads: "Registry entry values for the entries that match the pattern provided in the 'Registry Key' parameter." Under the heading "Properties", there is a label "Registry key" followed by an empty text input field. At the bottom of the dialog, there are two buttons: "OK" with a green checkmark icon and "Cancel" with a red X icon.

VB Script

VB Script functionality executes the specified script on the Windows host. You can also configure an Awk parser to parse the output and input the results of the parsing to DRA. To enable parsing, check the **Awk Script** check box, and then enter the script in the area below it. If the **Awk Script** check box is not checked, then no parser is used.

-- Webpage Dialog

Configuration

☒ VB Script

Provide VB script that will be executed on the Windows host. You can configure an 'awk' (after the famous Unix utility) parser that will be applied to the content before storing it in the RecoverGuard database.

Properties

Identifier

Script

☐ Awk Script

OK


✓

Cancel

✗

WMI Query

WMI Query functionality collects data from WMI using a query statement. Enter the required parameters and the query statement to be used.



The screenshot shows a dialog box titled "Webpage Dialog" with a "Configuration" tab. Inside the tab, there is a section for "WMI Query" with the instruction "Collect data from WMI using a query statement." Below this, there are three input fields: "Identifier", "Namespace", and "Query". The "Query" field is a large text area with a vertical scrollbar. At the bottom of the dialog, there are "OK" and "Cancel" buttons with green and red checkmarks respectively.


- 4 Click **OK** to save the package definition.

Defining host comparison expansion packages

Host Comparison Packages extend the functionality of the Host Comparison Module and the Host Configuration Differences Report, by introducing means to compare configuration items collected using Data Collection Packages. In effect, they define the rules to be used to compare the data you collected.

Within a Host Comparison Package, you can choose which configuration items to compare, and define how to compare them (you can even add more than one comparison method for a given item).

For each package, DRA automatically adds a new comparison category (activated through a check box) to the Host Comparison Module and the Host Configuration Differences Report. Differences found based on the comparison methods included in the package are highlighted.



Each row in the Data Collection Packages page represents a separate data collection package. The following information is shown for each package:

Name	The name of the host comparison package
Status	Indicates whether the host comparison package is valid or invalid.
Author	The author of the host comparison package
Description	A description of the package

You can perform the following operations from the Host Comparison Packages page:


Add Package		Adds a new host comparison package
Edit Package		Modifies the properties of the host comparison package
Delete Package		Deletes the selected package
Import Package		Imports the selected package

Export Package 

Exports the selected package

To add a host comparison package


- 1 Click **Add Package** (). The following window opens:




The screenshot shows a web-based dialog box titled "Host Comparison Package". It is divided into two main sections. The top section, labeled "Properties", contains three input fields: "Package name", "Author", and "Description". The "Description" field is a larger text area with scrollbars. The bottom section, labeled "Add", contains a table with three columns: "Comparison", "Comparator", and "Remove". The "Comparison" column is currently selected. At the bottom of the dialog are "OK" and "Cancel" buttons.

- 2 In the **Package Name** field, enter a name for the package.
- 3 In the **Author** field, enter a name for the package author.
- 4 In the **Description** field, enter a description for the package.

- 5
- Click  to define the rules for comparing data. This action opens the following window:



- 6
- In the **Identifier** field, enter a name for the comparison. This identifier is used in the Host Configuration Differences report and in the Host Comparison Module.
- 7
- In the **Data Collection Expansion Package** drop-down list, select the applicable data collection package. This value, along with the value you specify in step 8, define what to compare.
- 8
- In the **Configuration Item** drop-down list, select the configuration item you want to compare. The items presented in the list vary, depending on the items you defined for the data collection package you selected in step 7.



9 In the **Comparison Method** drop-down list, select the comparison method to be used. The options that display vary, depending on the items selected in steps 7 and 8 above. The following methods are available:

- **Textual:** Advanced text comparison that lets you define various text manipulations before comparison is executed.
- **XML:** Comparison method designed for XML files and outputs, which ignores indentation, spaces, and attribute order in tags.
- **Registry:** Comparison method designed for registry entries (recursive and non-recursive).
- **Directory structure:** Comparison method designed for directory structure (recursive and non-recursive).
- **File meta data:** Comparison method designed for file meta data. This method collectively compares size, checksum, owner, and permissions.
- **Custom Java comparator:** Enables you to implement your own comparison method using Java. To do so, create a class that follows the template below, compile it (with Java 6 Compiler) into a jar file and copy this file to `${TOMCAT_HOME}/webapps/DRA/WEB-INF/lib/` on the DRA server.

```
public class MyCustomComparator {
    /*
     * param: inputMap: java.util.Map - Unmodifiable map that contains the values for each host (host -> value).
     * param: outputMap: java.util.Map - An output map that you need to fill with the final result for each host (host -> value).
     * By default, the output map contains the same values as the input map.
     *
     * returns: boolean value which indicates if the values are different.
     */
    public boolean areValuesEqual(java.util.Map<Integer, String> inputMap, java.util.Map<Integer, String> outputMap) {
        // TODO - check if values are different...
    }
}
```

After the jar deployment is completed, provide the full class name (including the package) in the **Comparison Method** field.

10 In the Comparison Parameters area, define the parameters for the comparison.

Comparison Parameters	
Show results in diff mode	Yes ▼
Sort lines before comparing	Yes ▼
Trim spaces in lines	Yes ▼
Ignore empty lines	Yes ▼
Case sensitive	Yes ▼
Ignore lines that matches regular expression	
Accept only lines that matches regular expression	


11 In the Difference Impact area, provide a textual description of the impact expected when an inconsistency between hosts is detected in the specified configuration item.

12 Click **OK**.

Defining risk detection expansion packages

Risk Detection Packages enable you to define two types of custom gap signatures on collected data:

- Custom Collection Based, which are based on the configuration items defined in Data Collection Packages.
- Generic Java, which provides full custom functionality using Java code. This custom signature type gives you maximum flexibility by providing your own Java code, which can access the entire DRA System API (SAPI).



Each row in the Risk Detection Packages page represents a separate package. The following information is shown for each package:

Name	The name of the risk detection package
Status	Indicates whether the risk detection package is valid or invalid
Author	The author of the risk detection package
Description	A description of the package

The operations that you can perform from the Risk Detection Packages page are the same as those described on page 194.

To create a gap signature based on a custom collection package

- 1 Click **Add Package** (+). The following window opens:

The screenshot shows a 'Risk Signature Package' dialog box. It has a title bar with a 'Webpage Dialog' icon and a close button. The main area is divided into two sections. The top section, titled 'Properties', contains three text input fields: 'Package name', 'Author', and 'Description'. The 'Description' field has a vertical scrollbar. Below the 'Properties' section is a section with an 'Add...' button (green plus icon) and a table. The table has three columns: 'Signature', 'Type', and 'Remove'. The table is currently empty. At the bottom of the dialog are 'OK' and 'Cancel' buttons.

- 2 In the **Package Name** field, enter a name for the package.
- 3 In the **Author** field, enter a name for the package author.
- 4 In the **Description** field, enter a description for the package.
- 5 Click **Add** (+) to define the gap signature. This action opens the following drop-down list:

The screenshot shows a drop-down list that appears when the 'Add' button is clicked. The list has a header 'Add...' with a green plus icon and a dropdown arrow. Below the header are three items, each with a small icon to its left: 'Custom Collection Based Signature', 'SAPI Query Based Signature', and 'Generic Java Signature'.

- 6 Select the **Custom Collection Based Signature** option. The following window opens:

The screenshot shows the 'Gap Configuration' dialog box. The 'Properties' section includes a 'Signature Name' text field, an 'Area of Responsibility' section with checkboxes for 'Storage', 'Servers', and 'Databases', a 'Data Collection Expansion Package' dropdown menu, and a 'Configuration Item' dropdown menu. The 'Detection Method' section features a 'Value analysis' dropdown menu set to 'Equals to', a 'Value' text field, and three checkboxes: 'Case sensitive' (Yes), 'Trim spaces' (Yes), and 'Ignore line breaks' (Yes). The 'Summary' section has a text field. The 'Description' section has a large text area. The 'Impact' section has a large text area. The 'Severity' section has a dropdown menu set to 'Error'. The 'Categories' section has a link to 'none'. The 'Resolution' section has a large text area. At the bottom are 'OK' and 'Cancel' buttons.

- 7 In the **Signature Name** field, enter a name for the gap signature.
- 8 In the **Area of Responsibility** field, select the check box for the team responsible for the related open ticket.
- 9 In the **Data Collection Expansion Package** drop-down list, select the package to which the signature applies.
- 10 In the **Configuration Item** drop-down list, select the configuration item you want to compare. The items presented in the list vary, depending on the items you defined for the data collection package you selected in step 8.

- 11** In the **Value analysis** field, select the value analysis method for the configuration item from among the following:

Empty	The value of the configuration item is empty.
Not empty	The value of the configuration item is not empty.
Equals to	The value of the configuration item equals the provided parameter.
Different from	The value of the configuration item is different than the value provided as a parameter.
Contains	The value of the configuration item contains the string provided as a parameter.
Does not contain	The value of the configuration item does not contain the string provided as a parameter.
Matches	The value of the configuration item matches the regular expression provided as a parameter.
Doesn't match	The value of the configuration item does not match the regular expression provided as a parameter.
Greater than	The value of the configuration item is greater than the value provided as a parameter (numeric only).
Less than	The value of the configuration item is less than the value provided as a parameter (numeric only).
Between	The value of the configuration item is within the range provided by the minimum/maximum parameters (numeric only).
Not between	The value of the configuration item is not within the range provided by the minimum/maximum parameters (numeric only).
Registry key existence	The registry key provided as a parameter does not exist.
Registry value	The value of the registry key provided matches the defined criteria.

Custom Java Analyzer You can use your own Java class. To do so, create and compile a Java class that complies with the following format, using Java 6. Export a jar file and copy it to `${TOMCAT_HOME}/webapps/DRA/WEB-INF/lib/` on the DRA server. The required format is illustrated by the following sample:

```
public class MyCustomGapAnalysis {  
    public boolean shouldOpenTicket(String hostname, String value) {  
        // TODO - check if value is not OK...  
    }  
}
```

To use your code, provide the full class name (including the package name) in the Java Class Name section.

- 12 In the **Value** field, select the value for the comparison.
- 13 Define the method of comparison using the following fields:
 - **Case Sensitive**
 - **Trim spaces**
 - **Ignore line breaks**
- 14 [Optional] Click  to view information about ticket patterns.
- 15 Enter the applicable details in the Summary, Description, and Impact areas.
- 16 In the **Severity** drop-down list, select the severity of the gap.
- 17 Click the **Categories** link to select the impact categories of the risk signature.
- 18 In the **Resolution** field, specify the resolution patterns for the signature.
- 19 Click **OK** to save the gap signature.

To create a gap signature based on a SAPI query-based package

- 1 Click **Add Package** (+). The following window opens:

The screenshot shows a 'Risk Signature Package' dialog box. It has a title bar with a 'Webpage Dialog' icon and a close button. The main area is divided into a 'Properties' section and a table. The 'Properties' section has three fields: 'Package name', 'Author', and 'Description'. Below the 'Description' field is an 'Add...' button with a green plus icon. Underneath the 'Add...' button is a table with three columns: 'Signature', 'Type', and 'Remove'. The table is currently empty. At the bottom of the dialog are 'OK' and 'Cancel' buttons.

- 2 In the **Package Name** field, enter a name for the package.
- 3 In the **Author** field, enter a name for the package author.
- 4 In the **Description** field, enter a description for the package.
- 5 Click **Add** (+) to define the gap signature. This action opens the following drop-down list:

The screenshot shows a drop-down list that appears when the 'Add' button is clicked. The list has a title bar with a green plus icon and the text 'Add...'. It contains three items, each with a plus icon to its left: 'Custom Collection Based Signature', 'SAPI Query Based Signature', and 'Generic Java Signature'.

- 6 Select the **SAPI Query Based Signature** option. The following window opens:

The screenshot shows a 'Gap Configuration' dialog box with the following fields and options:

- Signature Name:** A text input field.
- Area of Responsibility:** Checkboxes for ☐ Storage, ☐ Servers, and ☐ Databases.
- Query:** A large text area for entering the SQL query.
- Query analysis:** A dropdown menu currently showing 'Empty result set'.
- Opens a ticket when the query is returned empty:** A checkbox.
- Summary:** A text input field.
- Description:** A large text area.
- Impact:** A large text area.
- Severity:** A dropdown menu currently showing 'Error'.
- Categories:** A link labeled 'none'.
- Resolution:** A large text area.

At the bottom of the dialog are 'OK' and 'Cancel' buttons.

- 7 In the **Signature Name** field, enter a name for the gap signature.
- 8 In the **Area of Responsibility** field, select the check box for the team responsible for the related open ticket.
- 9 In the **Query** field, enter the SQL query. The SQL query is based on an SAPI schema. The query result is the raw material for gap detection.
- 10 Use the **Query analysis** field to represent the manipulation of the query results, meaning the data returned by the query. Possible values are:
- Empty result set: Opens a ticket when the query is returned empty.
 - Not empty result set: Opens a ticket when the query returns results.
 - Ticket per row: Opens a ticket for each row returned by the query.

- Number of rows: Opens a ticket if the number of rows returned by the query matches the provided criteria.
 - Key column: Opens a ticket for each group of rows, according to the specified column.
- 11 In the **Value** field, select the value for the comparison.
 - 12 Define the method of comparison using the following fields:
 - **Case Sensitive**
 - **Trim spaces**
 - **Ignore line breaks**
 - 13 [Optional] Click  to view information about ticket patterns.
 - 14 Enter the applicable details in the Summary, Description, and Impact areas.
 - 15 In the **Severity** drop-down list, select the severity of the gap.
 - 16 Click the **Categories** link to select the impact categories of the risk signature.
 - 17 In the **Resolution** field, specify the resolution patterns for the signature.
 - 18 Click **OK** to save the gap signature.

To create a gap signature based on a generic Java package

- 1
- Click **Add Package** (+). The following window opens:

The screenshot shows a 'Risk Signature Package' dialog box. It has a title bar with a close button. The main area is divided into two sections. The top section, titled 'Properties', contains three text input fields: 'Package name', 'Author', and 'Description'. The 'Description' field has a vertical scrollbar. Below the 'Properties' section is a section with a green plus icon and the text 'Add...'. Underneath this is a table with three columns: 'Signature', 'Type', and 'Remove'. The table is currently empty. At the bottom of the dialog are 'OK' and 'Cancel' buttons.

- 2
- In the **Package Name** field, enter a name for the package.
- 3
- In the **Author** field, enter a name for the package author.
- 4
- In the **Description** field, enter a description for the package.
- 5
- Click **Add** (+) to define the gap signature. This action opens the following drop-down list:

The screenshot shows a drop-down list that is open. The list has a title bar with a green plus icon and the text 'Add...'. Below the title bar are three items, each with a plus icon to its left: 'Custom Collection Based Signature', 'SAPI Query Based Signature', and 'Generic Java Signature'.


- 6 Select the **Generic Java Signature** option. The following window opens:

- 7 In the **Signature Name** field, enter a name for the gap signature.
- 8 In the **Area of Responsibility** field, select the check box for the team responsible for the related open ticket.
- 9 In the **Java Class Name** field, enter the name of the class that implements the signature.
- 10 [Optional] Click  to display information about how to write your own Java code.
- 11 Enter the applicable details in the **Summary**, **Description**, and **Impact** areas.
- 12 In the **Severity** drop-down list, select the severity of the gap signature.
- 13 Click the **Categories** link to select the impact categories of the risk signature.

- 14 In the **Resolution** field, specify the resolution patterns for the signature.
- 15 Click **OK** to save the gap signature.

Defining reporting packages

Reporting Packages enable you to create custom reports, based on the DRA System API (SAPI), which are automatically added to the available report list (under Custom Reports). Custom reports have the same functionality as built-in reports (including exporting, scheduling, and so on).



Each row in the Reporting Packages page represents a separate package. The following information is shown for each package:

Name	The name of the reporting package
Author	The author of the reporting package
Report Name	The name of the new customer report
Description	A description of the package

The operations that you can perform from the Reporting Packages page are the same as those described on page 194.

To create a reporting package

- 1 Click **Add Package** (📁➕). The following window opens:

The screenshot shows a 'Webpage Dialog' window titled 'Reporting Package'. It contains two main sections. The top section, also titled 'Reporting Package', has a 'Properties' group with three fields: 'Package Name', 'Author', and a large 'Description' text area. The bottom section has a 'Properties' group with two fields: 'Report Name' and a large 'SQL Query' text area. At the bottom of the dialog are 'OK' and 'Cancel' buttons.

- 2 In the **Package Name** field, enter a name for the package.
- 3 In the **Author** field, enter a name for the package author.
- 4 In the **Description** field, enter a description for the package.
- 5 In the **Report Name** field, enter a name for the report. The name that you specify is used to identify the report in the Reports Module, under Custom Report.
- 6 In the **SQL Query** field, enter an SQL query, which extracts the data for the report.

Note: You can use parameters that are text variables (or macros) that are passed to the query when the report is generated. If parameters are specified, the user is presented with additional form fields to complete when configuring the report. To use parameters, simply include the expression in the format `${ParamName}` inside the query, where `ParamName` is the parameter name (may include spaces, and are case sensitive).

For example, when you specify the following query;

```
SELECT PRIMARY_IP  
FROM SAPI_HOST  
WHERE NAME LIKE '%${HostName Pattern}%'
```

A new text field labeled **Host Name Pattern** is added to the report generation form.

- 7 Click **OK** to save the package.

System log

DRA stores internal events in the system event log, including when a full scan starts or stops, or when the system fails to connect to a host. These events may be assigned a severity of Info, Warning, Error, or Critical. Most events are information-only; you do not have to do anything about them. However, you may have to take action to fix a Warning, Error, or Critical event.

Event ID	Severity	Created At	Description
Dashboard	Topology	Tickets	Reports
SLA	Compliance	Configuration	
Menu	System Event Log		
System	Shows events notified in the last 1993		
Configuration			
Configuration			
Basic Scan Configuration			
Configuration Wizard			
Distributed Collection			
Policies			
Advanced Scan Configuration			
Security Definition			
File Cluster Definition			
Business Entities			
SLA Definition			
Expansion Packages			
Data Collection			
Host Collection			
Report Scheduling			
Report Scheduling			
System Log			
Scan Status			
Scan Troubleshooting			
Advanced Configuration			

Each system event includes the following information:


Event ID	A unique identifier for the event
Severity	Defines the event severity, which can be one of the following: ■ Info ■ Warning ■ Error ■ Critical
Created At	When the event was created
Description	A clear, descriptive summary of the event

To view the system event log

- 1 To access the Configuration page, click the Configuration tab.
- 2 In the tree hierarchy, click **System Log**. The System Event Log page is displayed.


Scan Status

The Scan Status branch shows details for active and previously run scheduled tasks.



This branch contains two tabs:


- **Active Scan Tasks:** Displays a gray bar for each scheduled task that is currently running. This bar shows who created the task and when it started.
- **History:** Displays a gray bar for each scheduled task that has been run, including details about who created the task, when it started, and when it ended.



You can perform the following operations from the Scan Status page:

Refresh  Refreshes the data on the page

Data in the History tab can be filtered using the filter options at the top of the page. You can filter by time and by task identifier.




Scan troubleshooting

During the scanning process, DRA may encounter several types of errors. Scanning errors may occur in the following situations:

- The credential does not have sufficient privileges to execute commands.
- The sudo definitions are incorrect.
- A command fails.
- There is a connectivity problem.
- A password has expired.

DRA reports these errors in the Scan Troubleshooting page.



In this page, the errors are divided into the following categories, each of which corresponds to a tab in the Scan Troubleshooting area:

- Hosts Host scan-related errors
- Databases Database scan-related errors
- Storage Arrays Storage array scan-related errors
- Scope Issues Missing configuration information, such as the network file system target is not scanned, the replica Symmetrix is not scanned, and so on

The Scan Troubleshooting area also contains an Installed Applications tab that shows the relevant installed applications that DRA uses to scan the servers, and an Installed Applications (Proxy) tab that shows the relevant installed applications that DRA uses to scan the storage proxies.

Each scanning issue presented in the Scan Troubleshooting page includes the following information:

- Severity The impact severity on the data reliability.
- Type The type of scanning issue.
- Summary A summary of the issue.
- Host The name of the relevant host on which the scanning issue is opened.
- Application The related application.
- Proxy The proxy used to scan the component.

- Action

Indicates the ability to suppress/unsuppress the issues.

Note: To suppress/unsuppress an issue, first click Show/Hide Suppressed Issues.
- Issue Preview

When you select an issue in the Scan Troubleshooting area, its details are shown in the Issue Preview area at the bottom of the window.

To view scanning troubleshooting information

- 1

To access the Configuration page, click the Configuration tab.
- 2

In the tree hierarchy, click **Scan Troubleshooting**. The Scan Troubleshooting page opens.


To modify the path of the installed application

- 1

To access the Configuration page, click the Configuration tab.
- 2

In the tree hierarchy, click **Scan Troubleshooting**. The Scan Troubleshooting page opens.
- 3

Click the Installed Applications tab. The following page opens:




- 4

Click a host in the Scan Troubleshooting area to display its installed applications in the **Application Path and Binary Filers Used for Scanning** area. If an application is reported as not installed, it means that DRA could not find the application on the host.

- 5 Usually, DRA detects all the installed applications automatically. If an application was not found, you can modify the application path so that DRA looks for that application in the specified path during the next scan. To do so, double-click the application row, or select the row and click **Edit Paths**

(). The following window opens:



- 6 In the **Path** field, specify the correct path.
- 7 If you want this path to be the default path for the selected operating system, check the **Apply to all AIX hosts** check box. Note that the operating system changes according to the selected hosts.
- 8 [Optional] Use the **Activate/Deactivate** button to define whether or not to scan an application.
- 9 To save the changes, click **OK**. In the next scan, DRA verifies the path that was manually entered.

Configuring advanced settings

The DRA application lets you view and configure the following information that relates to the application's operation:

- User Management
See [“Configuring user management settings”](#) on page 217.
- External Ticketing
See [“Configuring external ticketing”](#) on page 223.
- Manage Labels
See [“Managing labels”](#) on page 228.
- Gaps
See [“Configuring gap management settings”](#) on page 229.
- Licenses
See [“Configuring licenses”](#) on page 231.
- System Properties
See [“Configuring system properties”](#) on page 235.
- System Tools
See [“Configuring system tools”](#) on page 252.

Configuring user management settings

The Users page lets you add, change, or delete DRA users. You can also define user roles, which specify what information this type of user may see.

For more information see [“Configuring user roles”](#) on page 220.

DRA includes two default user names: admin and user.

The DRA administrator uses the admin account. This user can modify both the admin and user passwords. admin users can access all DRA modules.

Important: Keep the admin password in a safe place. If you lose it, you have to reinstall DRA.

You can change the admin and user passwords, but you cannot create or change these user names or their assigned user roles.

Important: A role does not provide or configure privileges. It relates to the scope of the business entities that are exposed to this role. The only privileged user is admin.

You can perform the following operations from the Users page:

- Add User



Adds a new user
- Edit User



Modifies the properties of a user
- Delete User




Deletes a user


Note: The Delete User operation cannot be performed on the default user or the Admin user.

To add a new user

- 1 To access the Configuration page, click the Configuration tab.
- 2 In the tree hierarchy, under Advanced, click **Users**. The Users page opens.




- 3 Click **Add User** (). The User window opens.



- 4 In the **Login** field, enter the user's name.
- 5 Check the **Active Directory User** check box to disable the passwords and the secret question option for this user. When doing so, authentication is performed using the organization's Active Directory. To use this option, you must configure the Active Directory properties in DRA, as described in ["Active directory"](#) on page 236.
- 6 In the **Password** field, enter the user's password.
- 7 In the **Confirm** password field, re-enter the password.
- 8 In the **Role** drop-down list, select a role.
- 9 In the **Secret question** drop-down list, select a predefined question the user must answer when they enter password details. You can also create your own question.
- 10 To save the settings and exit the window, click **OK**.

Configuring user roles

A user role can be assigned to a user. The user role specifies the type of DRA information a user can view. For example, if an organization has multiple divisions, you can assign a different user role to each division, so that each division can only view the information that is relevant to it.



Important: A role does not provide or configure privileges. It relates to the scope of the business entities that are exposed to this role. The only privileged user is admin.


You can perform the following operations from the User Roles page:

- Add User Role  Adds a new user role
- Edit User Role  Modifies the properties of a user role
- Delete User Role  Deletes a user role

To add a new user role

- 1 To access the Configuration page, click the Configuration tab.
- 2 In the tree hierarchy, under Advanced, click **User Roles**. The User Roles page opens.

- 3 Click **Add User Role** (). The User Role window opens.



- 4 In the **Name** field, enter the user role name.
- 5 Optionally, if you want the user role to have access to all business entities, select the **All Business Entities Allowed** check box. When you select this option, this user role does not limit the type of information these users can see.
- 6 To limit the user role to specific business entities, make sure you do not select the **All Business Entities Allowed** check box. Then, in the left pane, select the business entities that you want the users of this role to see and click  to move the entities to the right pane. The list of business entities on the left includes all business entities in the system. The list on the right should include all business entities that the user role can see. Those business entities in the right pane do not appear in reports, tickets, the dashboard, and so on for this role.
- 7 To save the settings and exit the window, click **OK**.

To add users by importing a CSV file

- 1 To access the configuration page, click the Configuration tab.
- 2 In the tree hierarchy, click **Users**.

- 3
- Click **Import Users from CSV file** (). The Import Users from CSV file page opens.

Import Users from CSV file

CSV file format:

- **Username *** : The user login name.
- **Password** : The password (not required for Active Directory users).
- **Role**: The user role (default: system). If the user role does not exist, it will be automatically created.
- **Active Directory** : Indicates whether the user authentication is done against the Active Directory. Possible values: **Yes/No**.
- **Admin** : Indicates whether the user is administrator. Possible values: **Yes/No**.

	A	B	C	D	E
1	Username	Password	Role	Active Directory	Admin
2	user1	mypassword	system		
3	user2	mypassword	ERP Administrators		yes
4	user3		Management	yes	yes
5	user4	mypassword	Management		yes

[Download sample](#)

 If the users are not added, please verify that the file format is according to the instruction above

Import from file...


Close 

- 4
- On the Import Users from CSV file screen, review the file format. Your file must match this format exactly.
- 5
- Create a CSV using a text editor or by saving an Excel spreadsheet to CSV format.
- 6
- Save your file with the extension .csv.
- 7
- On the Import Users from CSV file screen, click **Import from file**.
- 8
- Browse to your file, select it, and click **OK** to import it.

Configuring external ticketing

DRA can easily integrate with your enterprise ticket / incident management system, through the use of either a generic email API or a customized Java plug-in.

Currently, the HP Service Manager connector is supported.



You can perform the following operations from the Enterprise Ticket Management Connector Configuration page:

Save Configuration 

Saves an external ticketing configuration

Delete Configuration 


Deletes an external ticketing configuration

Configuring an external ticketing connector

To configure a connector to your Enterprise Ticket Management system

- 1 To access the Configuration page, click the Configuration tab.
- 2 In the tree hierarchy, under Advanced, click **External Ticketing**. The Enterprise Ticket Management Connector Configuration page opens.

Note: Only Administrators can access this page.




- 3 Under **Connector Configuration**, perform the following:
 - Select the **Connection mode** (Manual, Automatic, or Inactive) in the drop-down list.
 - Select the **Connector type** in the drop-down list.A list of **Connector Attributes** loads, based on your selection. DRA provides out-of-the-box support for:
 - A generic email connector
 - A Java plug-in for HP Service Manager 9.2

Notes: Additional custom connectors can be loaded. Refer to the documentation that came with your custom connector, as required. See the instructions below for details about configuring various connectors.

- 4 Add one or more filtering criteria to determine which DRA tickets are synchronized with your Enterprise Ticket Management system. To do so:

- a Click  **Add Criteria**. The following window opens in which you can:



- Select the scope: To filter by site or business entity, click one (or both) of the respective links and then select which hosts or business entities to include.
- Further refine the selection by filtering based on ticket category, ticket severity, or ticket label.

- b

You can add more than one criterion. Tickets that match the filtering configuration of different criteria are appended to one another.
- c

After making your selections, click **OK**. The filtering criteria are added to the Filter area of the window.

Filter							
Add Criteria							
Category	Severity	Hosts	Business Entities	Sites	Clusters	Ticket Label	Delete
SLA Breach, Data Risk	WARNING, ERROR, INFO	All hosts	ERP, Billing	All sites	All clusters	All ticket label	X
Optimization	ERROR, WARNING, INFO	All hosts	All business entities	Chicago, Boston	All clusters	All ticket label	X
General	ERROR	win2k, xp10	All business entities	All sites	All clusters	All ticket label	X

Notes: To change an existing criterion, double-click it.
To delete a criterion, click the **X** under the **Delete** column.

- 5

Click **Save Configuration** () to save the settings.

To configure an e-mail connector

- 1

To access the Configuration page, click the Configuration tab.
- 2

In the tree hierarchy, under Advanced, click **External Ticketing**. The Enterprise Ticket Management Connector Configuration page opens.
- 3

Select Email as the **Connector Type**.
- 4

Fill in the following connector attributes in the Connector Attributes area of the window, all of which are mandatory:

General

Connector Configuration

Connector modeInactive

Connector typeE-mail

Connector Attributes

E-mail address

Report formatPDF

Show topology in reportNo

- **E-mail address:** Specify the email address to which to send tickets.
- **Report format:** Select the format (HTML, PDF, or RTF).
- **Show topology in report:** Select whether or not to include a topology diagram.

Note: You must have a mail server already configured in DRA to use an email connector. For details, see [“Email settings”](#) on page 147.

To configure an HP Service Manager connector

- 1

To access the Configuration page, click the Configuration tab.
- 2

In the tree hierarchy, under Advanced, click **External Ticketing**. The Enterprise Ticket Management Connector Configuration page opens.
- 3

Select HP Service Manager 9.2 as the **Connector Type**.

- 4 Fill in the following connector attributes in the Connector Attributes area of the window, all of which are mandatory:

The screenshot shows a software window titled 'General'. It is divided into two main sections. The left section, 'Connector Configuration', contains two dropdown menus: 'Connection mode' set to 'Inactive' and 'Connector type' set to 'HP Service Manager 9.2'. The right section, 'Connector Attributes', contains several input fields and dropdown menus: 'Host name /IP' (empty), 'Host tcp /IP port' set to '13080', 'Report format' set to 'PDF', 'Show topology in report' set to 'No', 'Ticket area' (empty), and 'Ticket assignment group' (empty).

- **Host name/IP:** Enter your Enterprise Ticket Management system host name or IP address.
 - **Host TCP/IP port:** Use the default (13080) port, or change it if necessary.
 - **Report format:** Select the format (HTML, PDF, or RTF).
 - **Show topology in report:** Select whether or not to include a topology diagram.
- 5 Fill in the following information, which is statically assigned to all synchronized tickets:
 - **Ticket area:** Enter the HP Service Manager area.
 - **Ticket assignment group:** Enter the HP Service Manager assignment group.
 - **Ticket category:** Enter the HP Service Manager category.
 - **Ticket contact:** Enter the HP Service Manager contact.
 - **Ticket impact:** Enter the HP Service Manager ticket impact (for example, 1=enterprise, 2=site, and so on.).
 - **Ticket service:** Enter the HP Service Manager service.
 - **Ticket site category:** Enter the HP Service Manager site category (or keep the default).
 - **Ticket sub area:** Enter the HP Service Manager sub area.
 - **Ticket username:** Enter the HP Service Manager user.

Note: If you are unsure about which values to use, consult you HP Service Manager administrator.


- 6 Enter the login credentials to use with the HP Service Manager (user name and password).

- 7 Click **Save Configuration** () to save the settings.

Managing labels

The Manage Labels page lets you create labels that can be assigned to tickets. There are default labels that are automatically assigned by the system when a ticket is detected. The default label can be changed, added, or deleted by assigning other labels to a ticket. This feature helps to manage tickets more efficiently.

Ticket labels are assigned to tickets in the Tickets module. See [“Assigning a label to a ticket”](#) on page 277 for more details.



You can perform the following operations from the Labels page:


- | | | |
|--------------|---|-------------------------|
| Add Label |  | Adds a new ticket label |
| Edit Label |  | Modifies a ticket label |
| Delete Label |  | Deletes a ticket label |

Note: The Edit Label and Delete Label operations cannot be performed on default labels. They can only be performed on user-created labels.

To add a label

- 1 To access the Configuration page, click the Configuration tab.
- 2 In the tree hierarchy, under Advanced, click **Manage Labels**. The Labels page opens.

- 3 Click **Add Label** (). The following window opens:



- 4 In the **Name** field, enter a name for the label and click **OK**. The label is added to the list of labels in the Labels page.

To edit a label


- 1 Select a label in the Labels page and click **Edit Label** (). The Edit Label window opens.
- 2 Modify the label name and click **OK**.

Configuring gap management settings

The Gap Management page lets you view all currently suppressed gaps and, as needed, to unsuppress them. When a gap is suppressed, problems of this type are not identified during the next full system scan, and DRA does not collect any information about it.

Symptoms are issues that DRA recognizes as important for gap detection. DRA detects symptoms and then analyzes them and opens tickets.

For more information about suppressing gaps and tickets, see [“Performing operations on tickets”](#) on page 268.



To unsuppress a gap

- 1 To access the Configuration page, click the Configuration tab.
- 2 In the tree hierarchy, under Advanced, click **Gap Management**. The Gap Management page opens.
- 3 In the list of gaps, select the gap you want to unsuppress and click **Unsuppress**.

To unsuppress a gap symptom

- ◆ In the Gap Management page, select the Suppressed Gap Symptoms tab and then click the  button.

Configuring licenses

Your DRA license controls the number of hosts that you can configure on your system. It also has an expiration date that determines how many days are left for the system to operate. After a license has expired, you can still log on to DRA, but scanning is disabled. You should submit a request to renew your license as early as possible.

License

Product Licenses

Status	Number of hosts	Architecture	Expiration date	Days till expiration	Licensed for	License key	Actions
	300	Intel (Linux/Windows)	Perpetual	Perpetual	My Company	8D48A4A447...	
	200	Non-Intel (Unix)	Perpetual	Perpetual	My Company	8299F7F060...	

Add license

Maintenance Licenses

Status	Number of hosts	Architecture	Expiration date	Days till expiration	Licensed for	License key	Actions
	300	Intel (Linux/Windows)	Dec 31, 2011	429	My Company	4F8A294D05...	
	200	Non-Intel (Unix)	Dec 31, 2011	429	My Company	14507AF181...	

Number of physical hosts in scope

285 Intel (Linux/Windows)
31 Non-Intel (Unix)

Linux

3

Windows

282

ADK

8

Solaris

23

Activate license key by voucher

Voucher:

Browse...

Licensed for:

Activate

You can perform the following operations from the License page:

- Add License

Adds a new license to the system
- Copy to Clipboard

Copies the license key to the clipboard
- Delete

Deletes a license from the system
- Browse and Activate

Browses for the Symantec voucher and activates the license

DRA comes with a 14-day trial period.

To renew your license, you can choose one of two procedures, depending on whether your DRA server has Internet connectivity.

Contact Symantec to receive a new license. If your DRA server has Internet connectivity, Symantec gives you a voucher file. Save this file in a location accessible by the DRA server. If your DRA server is not connected to the Internet, Symantec gives you a license key, which is a long line of encrypted characters.

To activate a new DRA license

- 1
- Click the Configuration tab to access the Configuration page.
- 2
- In the left pane, under Advanced, click **License**. The License page is displayed.
- 3
- If the DRA server is connected to the Internet, do the following:
- Click the **Browse** button and find the voucher file that you previously saved. Select the file and click **Open**.

■

Click the **Activate** button to activate the DRA license. DRA then connects to a Symantec system over the Internet, which verifies the voucher. After the voucher has been verified, your system is automatically activated.
- 4
- If the DRA server is not connected to the Internet, do the following:
- Click **Add License**. The License page displays the following:

License

Product Licenses

Status	Number of hosts	Architecture	Expiration date	Days till expiration	Licensed for	License key	Actions
	300	Intel (Linux/Windows)	Perpetual	Perpetual	My Company	8D48A44447...	 Delete
	200	Non-Intel (Unix)	Perpetual	Perpetual	My Company	8299F7F060...	 Delete

Please enter the license key:

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

Maintenance Licenses

Status	Number of hosts	Architecture	Expiration date	Days till expiration	Licensed for	License key	Actions
	300	Intel (Linux/Windows)	Dec 31, 2011	429	My Company	4F8A294D05...	 Delete
	200	Non-Intel (Unix)	Dec 31, 2011	429	My Company	14507AF181...	 Delete

Number of physical hosts in scope

285 Intel (Linux/Windows)

31 Non-Intel (Unix)

Linux3

Windows282

AIX8

Solaris23

Activate license key by voucher

Voucher:

[Browse...](#)

Licensed for:

[Activate](#)

-
- In the License Key text box, enter the license key you received from Symantec.
-
- To save the license, click **Save**. The license is immediately updated and activated.

Note: Click Cancel to cancel the license activation or Copy to Clipboard

 to copy the license key to the clipboard. You can use the Copy to Clipboard feature to copy the license key to the clipboard. From there, you can paste the license into an email to your vendor, asking to renew your license.

Viewing license details

You can view license details in both the License page and the About window, as shown below.

License Page

License

Product Licenses

Status	Number of hosts	Architecture	Expiration date	Days till expiration	Licensed for	License key	Actions
	300	Intel (Linux/Windows)	Perpetual	Perpetual	My Company	8D48A4A447...	 Delete
	200	Non-Intel (Unix)	Perpetual	Perpetual	My Company	8299F7F060...	 Delete

Add license

Maintenance Licenses

Status	Number of hosts	Architecture	Expiration date	Days till expiration	Licensed for	License key	Actions
	300	Intel (Linux/Windows)	Dec 31, 2011	429	My Company	4F8A294D05...	 Delete
	200	Non-Intel (Unix)	Dec 31, 2011	429	My Company	14507AF181...	 Delete

Number of physical hosts in scope

285 Intel (Linux/Windows)

31 Non-Intel (Unix)


Linux	3
Windows	282
AIX	8
Solaris	23

Activate license key by voucher

Voucher:

Licensed for:

About Window (example)




To access the About window, go to the quick-access bar to the right of the Configuration tab, and click **About**:



About button

Configuring system properties

The System Properties page lets you configure parameters that control various aspects of the DRA application. System properties are grouped into categories, such as Aggregation, Dashboard, Rule Engine, and so on. The list in the System Properties page is sorted by category.




The following information is displayed for each system property:

Category	The category to which the property belongs
Description	A description of the system property
Value	The value of the system property
Info	If the value was manually changed, an indication appears with the default value as a tooltip.

To edit a system property

- 1
- In the System Properties page, click the row corresponding to the property value you want to set and click **Edit System Property** (). A windows opens, similar to the following. The window varies, depending on the type of property.



- 2
- Enter your property value and click **OK**.
- 3
- To restore the default value that came with the installation, click .

The following topics describe DRA system properties.

Active directory

Active Directory base DN for search	Specifies the Active Directory base search. This property usually looks like the following: dc=yourDomain,dc=com or dc=yourDomain,dc=local
Active Directory LDAP host URL	Specifies the URL to connect to the Active Directory. The value for this property typically looks like the following: ldap://hostname:389
Domain	Specifies the Active Directory domain. DRA uses this domain to authenticate the Active Directory users.
Use Active Directory username on application header	Shows the real name of the logged-in user instead of the login name.

Additional collection

Once a week, additional data is collected.

Additional collection schedule	Determines the day of the week.
Enable additional collection	Enables/disables this feature.

Aggregation

Note: If you change any system properties in the Aggregation category, they only take effect after the next full cycle.

Data age threshold, in seconds	Specifies the data age calculation threshold, in seconds, when calculating age inconsistency. For example, a three-second difference would not be a gap.
Date threshold, in seconds	Specifies when to combine storage volumes, based on their date. If the difference between the date of two storage volumes equals this value, the two volumes are combined.
Minimum number of items to aggregate in topology	Specifies when items in the topology are aggregated. If the number of items in the topology is greater than the value set for this property, then they are aggregated. For example, if this property value is set to 3, when there are three items or more in the topology, they are aggregated.
Minimum rows in ticket table to perform aggregation	Specifies the minimum number of rows that a ticket table must have for the rows to be aggregated. Row aggregation in the table only occurs when the table has more rows than the value specified here.

Automatic import

DRA integrates with the one of the following DBMSs:

- DB2
- MS SQL Server
- Oracle
- Sybase

The database administrator should define the table (or view) to be used for integration.

Automatically import business entities during full cycle	For DRA to import business entities automatically during a full cycle, set this property value to Yes.
Automatically import hosts during full cycle	For DRA to import hosts automatically during a full cycle, set this property value to Yes.
Automatically import sites during full cycle	For DRA to import site configuration automatically during a full cycle, set this property value to Yes.
Database authentication - Password	Specifies the database password to be used for authentication.
Database authentication - Username	Specifies the database username to be used for authentication.
Database connection - Database name (for MS SQL Server and DB2)	Specifies the database name for the database connection, which is only required for MS SQL Server and DB2.
Database connection - Database type	Specifies the type of database (Oracle, DB2, MS SQL Server, or Sybase) for the database connection.
Database connection - Host name	Specifies the host name or IP address for the database connection.
Database connection - Instance name	Specifies the instance name for the database connection.
Database connection - port	Specifies the port for the database connection.
Integration table/view name	Specifies the name of the integration table (or view). The default name for this table is MAPPINGS_INTEGRATION.
Override (re-create) Business Entities manually deleted using the user interface	For imported data to override changes made manually using the interface, set this property value to No. The default value is Yes.

Override SLA policies, roles, and associations that were manually configured using user interface	For imported data to override changes made manually using the interface, set this property value to Yes. The default value is No.
---	---

Collection

Check if the database port is available before connecting	Validates that the database port is responding before trying to connect to the database.
Database port check timeout, in seconds	Specifies the timeout for trying to connect to a scanned database port.
Dispatcher flush mode	Specifies whether or not to flush items before the end of the response in order to reduce memory consumption.
ECC port	Specifies the port that DRA uses to connect to EMC ECC in order to collect information.
ECC SSL port	Specifies the port that DRA uses to connect to Secured EMC ECC in order to collect information.
HiCommand port	Specifies the port that DRA uses to connect to HDS HiCommand in order to collect information.
Maximum allowed host processes	To ensure the safety of your environment, before DRA executes a command on a server, it checks if the command is already running, and how many processes for that command are currently running. If there are more processes than the specified value, the command is not executed, and a proper message is displayed on the Scan Troubleshooting page.
Maximum number of items per iteration	Specifies the number of items handled in one iteration. This property affects the required CPU memory when loading responses.
Memory load percentage	Specifies the threshold, in percent, for the loading process. The lower the value, the less memory the system consumes, but the longer it takes for the loading process to complete.
NetApp ZAPI major version	Specifies the minimum NetApp ZAPI major version that is supported.

NetApp ZAPI minor version	Specifies the minimum NetApp ZAPI minor version that is supported.
Ports list for Oracle	Lists the ports that DRA uses to connect to an Oracle server database.
Ports list for Oracle with SSL	Lists the ports that DRA uses to connect to an Oracle with SSL server database.
Ports list for Sybase	Lists the ports that DRA uses to connect to a Sybase server database.
Ports list for SQL server database	Lists the ports that DRA uses to connect to an SQL server database.
Ports list for UDB	Lists the ports that DRA uses to connect to a UDB server database.
SRM port	Specifies the port used to connect to SRM in order to collect information.
Sudo regexp to open SSH pty	A regular expression that determines whether or not PTY should be used in the SSH connection.
Time (in seconds) between external collector lookup	Specifies the internal time to wait before attempting to reconnect to non-responsive remote collectors.
Timeout for collector connection (in collector lookup cycles)	Specifies the timeout to wait for a remote collector to respond.
Timeout on waiting for searching a binary on Windows hosts, in minutes	Specifies the timeout interval, in minutes, for DRA to search binaries in Windows platforms.
Try to re-find binaries on Windows hosts that failed on timeout in past cycles	Specifies whether or not DRA should attempt to find binaries there were not found due to a timeout in past full system scans.
Try to re-find binaries on Windows hosts that were not found in past cycles	Specifies whether or not DRA should attempt to find binaries there were not found in past full system scans.
Validate database VIP	Validates that the database virtual port is a configured IP address on the database host.
Wait time between two consecutive server connections, in seconds	Specifies the internal time to wait before connecting to the next scanned server.

Whether to check ping for new added/discovered hosts	Specifies whether or not the application tries to ping the newly discovered/added hosts. Setting this property to No provides better performance in the initial configuration of DRA.
--	---

Collection timeouts

Timeout for a single database scan, in minutes	Specifies the time to wait for a database scan to finish. If the timeout expires before the scan completes, the scan fails.
Timeout for a single DSCLI proxy scan, in minutes	Specifies the time to wait for a DSCLI proxy scan to finish. If the timeout expires before the scan completes, the scan fails.
Timeout for a single ECC scan, in minutes	Specifies the time to wait for an ECC scan to finish. If the timeout expires before the scan completes, the scan fails.
Timeout for a single HiCommand scan, in minutes	Specifies the time to wait for a HiCommand console scan to finish. If the timeout expires before the scan completes, the scan fails.
Timeout for a single HMC scan, in minutes	Specifies the time to wait for an HMC scan to finish. If the timeout expires before the scan completes, the scan fails.
Timeout for a single host scan, in minutes	Specifies the time to wait for a host scan to finish. If the timeout expires before the scan completes, the scan fails.
Timeout for a single NaviCLI proxy scan, in minutes	Specifies the time to wait for a NaviCLI proxy scan to finish. If the timeout expires before the scan completes, the scan fails.
Timeout for a single NetApp filer scan, in minutes	Specifies the time to wait for a NetApp filer scan to finish. If the timeout expires before the scan completes, the scan fails.
Timeout for a single SymCLI proxy scan, in minutes	Specifies the time to wait for a SymCLI proxy scan to finish. If the timeout expires before the scan completes, the scan fails.
Timeout for a single SVC proxy scan, in minutes	Specifies the time to wait for an SVC proxy scan to finish. If the timeout expires before the scan completes, the scan fails.

Timeout for a single vCenter scan, in minutes	Specifies the time to wait for a vCenter scan to finish. If the timeout expires before the scan completes, the scan fails.
Timeout for a single XCLI proxy scan, in minutes	Specifies the time to wait for an XCLI proxy scan to finish. If the timeout expires before the scan completes, the scan fails.
Timeout for dead Data Collector, in minutes	Specifies the timeout to wait for a remote collector to respond for cases when a request was sent and the collector was down.
Timeout for scanning all databases, in minutes	In a full cycle, specifies the time to wait for all database scans to finish. If the timeout expires, the remaining databases are not scanned and the full cycle continues to the next stage.
Timeout for scanning all DSCLI proxies, in minutes	In a full cycle, specifies the time to wait for all DSCLI proxy scans to finish. If the timeout expires, the remaining DSCLI proxies are not scanned and the full cycle continues to the next stage.
Timeout for scanning all ECCs, in minutes	In a full cycle, specifies the time to wait for all ECC scans to finish. If the timeout expires, the remaining ECCs are not scanned and the full cycle continues to the next stage.
Timeout for scanning all HiCommand consoles, in minutes	In a full cycle, specifies the time to wait for all HiCommand console scans to finish. If the timeout expires, the remaining HiCommand consoles are not scanned and the full cycle continues to the next stage.
Timeout for scanning all HMC proxies, in minutes	In a full cycle, specifies the time to wait for all HMC scans to finish. If the timeout expires, the remaining HMCs are not scanned and the full cycle continues to the next stage.
Timeout for scanning all hosts, in minutes	In a full cycle, specifies the time to wait for all host scans to finish. If the timeout expires, the remaining hosts are not scanned and the full cycle continues to the next stage.
Timeout for scanning all NaviCLI proxies, in minutes	In a full cycle, specifies the time to wait for all NaviCLI proxy scans to finish. If the timeout expires, the remaining NaviCLI proxies are not scanned and the full cycle continues to the next stage.

Timeout for scanning all NetApp filers, in minutes	In a full cycle, specifies the time to wait for all NetApp filer scans to finish. If the timeout expires, the remaining NetApp filers are not scanned and the full cycle continues to the next stage.
Timeout for scanning all SymCLI proxies, in minutes	In a full cycle, specifies the time to wait for all SymCLI proxy scans to finish. If the timeout expires, the remaining SymCLI proxies are not scanned and the full cycle continues to the next stage.
Timeout for scanning all SVC proxies, in minutes	In a full cycle, specifies the time to wait for all SVC proxy scans to finish. If the timeout expires, the remaining SVC proxies are not scanned and the full cycle continues to the next stage.
Timeout for scanning all vCenters, in minutes	In a full cycle, specifies the time to wait for all vCenter scans to finish. If the timeout expires, the remaining vCenters are not scanned and the full cycle continues to the next stage.
Timeout for scanning all XCLI proxies, in minutes	In a full cycle, specifies the time to wait for all XCLI proxy scans to finish. If the timeout expires, the remaining XCLI proxies are not scanned and the full cycle continues to the next stage.

Command line interface

CLI communication port	Specifies the CLI port.
------------------------	-------------------------

Comparison

Golden Copy comparison - Show real value for equal servers (otherwise, display: OK)	The golden copy comparison shows the difference between multiple hosts compared to one <i>golden</i> host. When this property is set to No (the default), OK is displayed when the comparison values are equal and the real value is displayed when the comparison values are not equal. When this property is set to Yes, the real value is displayed, even if it is the same as the golden copy value.
---	--

Configuration

Display scan status according to last scan	Specifies whether or not DRA displays the scan status of storage arrays, hosts, and databases according to the last full cycle results. For example, if a host was scanned successfully in the past, but was not scanned in the last full scan (at all), if the value of this property is Yes, the status of this host returns to “Not verified yet”.
Host and database scan status refresh interval, in seconds	Specifies the interval for refreshing the scan/connectivity statuses in the interface.

Cyber-Ark configuration

Default AppID	Specifies the default AppID.
Default Folder	Specifies the default Folder.
Default Port	Specifies the default Port.
Default Safe	Specifies the default Safe.

Dashboard

Note: If you change any system properties in the Dashboard category, they only take effect after the Dashboard is refreshed (every five minutes).

Drill down high level business services to sub business entities	When this property is set to Yes, when logging in with a user role that is limited to one or more business entities, the sub-business entities are shown (expand view). When set to No, only the high-level business entities are shown.
Number of tickets to show in Top/New Tickets section	Specifies the number of tickets shown in the bottom dashboard pane that shows the Top/New tickets.
Scan summary – minimum error tickets for red notification	Specifies the minimum number of Error tickets that cause the relevant service and category light in the Scan Summary area to turn red. Note: See the figure below for reference.



Scan summary – minimum error tickets for yellow notification	Specifies the minimum number of Error tickets that cause the relevant service and category light in the Scan Summary area to turn yellow.
Scan summary – minimum warning tickets for red notification	Specifies the minimum number of Warning tickets that cause the relevant service and category light in the Scan Summary area to turn red.
Scan summary – minimum warning tickets for yellow notification	Specifies the minimum number of Warning tickets that cause the relevant service and category light in the Scan Summary area to turn yellow.

Scan Summary – Scan percentage considered as insufficient information	When the business entity scan coverage percentage is above the number shown in this property, the circles are colored and not grayed out.
Show Running/Failed Scan Tasks	When this property is set to Yes, Running/Failed scan tasks are displayed in a separate pane on the dashboard.
Sort top tickets by rating	When set to Yes, the top tickets by rating are sorted in descending order, based on the ticket's rating. When set to No, the sort is based on the ticket's severity.

Distributed collection

Default collector port	Specifies the default port of the collector.
Remote collector heartbeat interval, in seconds	Specifies the interval for the remote collector heartbeat.

Email settings

Default distribution list for reports	Specifies the default recipients list that is shown each time a new scheduled report is created, or when the user sends a ticket from the Tickets module. This property represents only the default distribution list. You can always change the recipients when creating scheduled reports.
List of domains for sending non-secured emails	Provides a domain list to prevent the sending of non-secure email to domains other than those specified. This mechanism prevents human error.
Mail sender address	Specifies the email sender's email address to use when sending reports by email.
Mail server	Specifies the mail server to use when sending reports by email.
Mail server authentication password	Specifies the mail server authentication password to use when sending reports by email.
Mail server authentication username	Specifies the mail server authentication username to use when sending reports by email.
Mail server port	Specifies the mail server port to use when sending reports by email.
Use authentication	Specifies whether or not to use authentication when sending reports by email.

Use TLS (authentication type)	Specifies whether or not to use TLS authentication when sending reports by email.
Zip attached files	Specifies whether or not to send zipped files, in order to reduce the size of the emails sent by DRA.

External ticketing system integration

External Ticketing Working Mode	Indicates the status of the external ticketing integration. Values are: - Manual: Manually sends tickets to an external system. - Auto: Automatically sends tickets to an external system, based on user-defined criteria. - Inactive: Never sends tickets to an external system.
---------------------------------	---

Full Cycle

Enable Full Cycle fault tolerance	Specifies whether a full cycle should continue from the point it stopped upon failure.
Restart Tomcat before scheduled full cycle	Specifies whether or not the Tomcat service is restarted before a scheduled full cycle.

Gap Rules

Contact support before changing any of the properties in this section.

General

Allow login during full cycle	Indicates whether or not to permit the login during a full cycle scan.
Product environment	Shows the name of the product environment next to the name of the logged-in user. This property is useful when there are multiple DRA installations.
Suggest statistics on login timeout, in seconds	When the system recognizes that performance is poor, it enables the user to perform data analysis. This property specifies the threshold used by the system to identify poor performance. The criterion used to determine poor performance is how long it takes to query license data.

Housekeeping

Advanced - Always collect large files (ignore large file size threshold)	Collects large files on every execution.
Advanced - Enable database maintenance	Indicates whether database maintenance should be performed on DRA. Setting this field to enabled is highly recommended for backup purposes.
Advanced - Enable logs and responses packing	Indicates whether DRA packs the responses and logs to save disk space. Note: The support tool uses the email settings configured in System Settings for sending emails.
Advanced - Heavy-duty maintenance schedule, day of week	The day of the week on which the database maintenance tasks should run.
Advanced - Keep all daily large files to be sent on weekly runs	Keeps collected large files for weekly execution.
Advanced - Minimum free drive space requirement, in MB	Specifies the minimum free drive space that is required for the support tool to run. If there is less than the specified amount of free space, the support tool cannot run.
Advanced - Number of daily DB dumps to keep	Specifies the number of days to keep daily database backups. Dumps older than this setting are deleted.
Advanced - Number of days back to bring Windows eventlog history	Specifies the number of days back to retrieve Windows event log history.
Advanced - Number of days to keep history	Specifies the maximum number of days to retain packed history (logs, responses). Files older than this setting are deleted.
Advanced - Number of weekly DB dumps to keep	Specifies the number of weeks to keep weekly database backups. Database dumps older than this setting are deleted.
Advanced - Zip split size, in MB	Specifies the zip parts' maximum file size, in MB. Larger files are split.
Daily Schedule	Sets the time of day when the support tool runs. Note: The configured time should not conflict with the full cycle schedule.
Enable Housekeeping	Support tool status. To override all other support tool settings, set this value to false.

List of recipients to receive unprotected support logs	Specifies the emails that are to receive the support tool logs in an unencrypted zip file.
--	--

Maintenance

Contact support before changing any of the properties in this section.

Reports

Acceptable hardware difference percentage in Host Configuration Differences report	Specifies the maximum accepted memory differences between hosts, in percent. Any difference greater than this value is displayed in the report.
Display Reports module to admins only	Specifies whether or not to display the Reports module only for Admin users.
ESX options to ignore in Host Configuration Differences report	Specifies the ESX options to be ignored in the Host Configuration Differences report.
Maximum business entities to be shown in Standby report summary table	Specifies the maximum number of business entities to be shown in the Standby report summary table.
Maximum databases to show in Database Storage Utilization report	Specifies the maximum number of databases to show in the Database Storage Utilization report.
Maximum number of differences to show in Host Configuration Differences report	When the number of differences in a single comparison is greater than the number specified, the differences for this comparison are not shown.
Maximum report size to be shown in the Internet browser, in bytes	Specifies the maximum report size, in bytes, that can be displayed in the Internet browser. Reports larger than this value do not display in the browser, and you are automatically directed to export the report (to a PDF, RTF, or Excel file).

The report is too large to be displayed in the internet browser.
Please use either PDF, RTF or Excel export options on the toolbar to view the report.

Page size	Specifies the page size to use when exporting a report.
Product name filter regexp in Host Configuration Differences report	Product name to ignore when comparing in the Host Configuration Differences report.
Summarize long tables factor	Specifies the number of rows for long tables. For example, if a table has 200 rows and this value is set to 20, then the number of rows in the table is 20. In this case, the 200 rows in the table are summarized into 20 rows to comprise the table.

Summarize long tables in ticket details report	Specifies whether or not to summarize tables that have many rows in the Ticket Details report. When set to Yes, only the most important rows are shown in the report, to reduce the number of rows.
Ticket details report topology number of tickets limitation	Specifies the maximum number of tickets that the Ticket Details report should have in order to include the tickets topologies (when requested for inclusion).
Use decorated headers	When set to Yes, the report title and report section heading are shown with a blue bar, as shown below.

SCAN REPORT

The following report presents the scan status of storage arrays, hosts, clusters and databases from the last full cycle of RecoverGuard.

LAST FULL CYCLE STATUS

Information is available only after running the first full cycle

When set to No, the report title and report section headings do not have blue bars, as shown below. Setting this value to No makes the RTF output easier to edit.

SCAN REPORT

The following report presents the scan status of storage arrays, hosts, clusters and databases from the last full cycle of RecoverGuard.

LAST FULL CYCLE STATUS

Information is available only after running the first full cycle

Rule Engine

Contact support before changing any of the properties in this section.

SLA

Data age threshold, in seconds	Specifies SLA tolerance for data age differences. A replicated unit's subcomponents can have different data ages. SLA tolerates deviations only as defined by this threshold. Above this threshold, DRA sends an alert. Note: If you change this property, it only takes effect after the next full cycle.
DataGuard synchronization - critical sequence difference threshold	Determines the threshold where the color of the replica bullet in the SLA module becomes red.
DataGuard synchronization - significant sequence difference threshold	Determines the threshold where the color of the replica bullet in the SLA module becomes yellow.
Ignore volume groups, logical volumes and file systems that match regular expression for host level calculation	Specifies the FSs, VGs, and LVs to ignore when calculating the SLA. The default is usually non-replicated devices that cause false positives.
Load SLA on system startup	Indicates whether or not to load SLA data to the cache on startup.
Retention frequency acceptable deviation percentage factor	Specifies the acceptable deviation percentage factor used in SLA retention verification. When DRA verifies that a component exists in the retention SLA defined for it, it checks the existence of the copies and that the age of these copies matches the SLA. This property defines the acceptable age deviation from the SLA policy defined.
SLA RPO - continuous replication threshold, in minutes	When replication is asynchronous, specifies the threshold when to open an SLA ticket.

System API

Enable System API feature	Indicates whether or not to create the SAPI schema when data analysis completes.
SYSAPI password	Specifies the password for the SYSAPI schema. When setting the password, it is automatically changed in the next SAPI regeneration.

Tickets module

- Display business impact in ticket description

Specifies whether or not to show the related business entities within the ticket's Impact tab.
- Display host and database scan status in the filter tree

Specifies whether or not to display the scan status of hosts and databases in the filter tree.
- Display scan results in ticket description when ticket involves items not scanned in the last cycle

Specifies whether or not to display the last successful scan data of items that are involved in a ticket, but were not scanned in the last full cycle.
- Display warning icon when items involved in a ticket have scanning issues

Note: If you change this property, the change only takes effect after the next full cycle.

Specifies whether or not to display a warning icon when items involved in a ticket have scanning issues.


ID	Labels	Detection Date	Status	Summary	Rating	Severity
Configuration mismatch between hosts on the same cluster (2)						
203	Alt: Cluster	Jun 04, 2012 11:38:23 AM		Cluster ches at site Afula has nodes that differ in key operation system properties	★★★★★	
202	Cluster: VMware	Jun 04, 2012 11:38:23 AM		Cluster ches2 at site Tel Aviv has nodes that differ in key operation system properties	★★★★★	

- Update labels when ticket is updated

Specifies whether or not to overwrite the ticket labels for each ticket validation process.

Configuring system tools

The System Tools page lets you configure and use various tools used within DRA.



The available tools are as follows:

Support tool	See “Using the support tool” on page 253.
Analyzer	See “Using the analyzer” on page 253.
Restart Veritas Disaster Recovery Advisor	See “Restarting DRA” on page 253.
System API Views	See “System API views” on page 253.

Using the support tool

This tool performs various maintenance functions, such as backing up system information (database, log files and so on). Click the adjacent  button to activate the support tool.

Configurable support tool settings are defined in the Housekeeping section of the System Properties page (see [“Housekeeping”](#) on page 248).

System API views

DRA scans the IT environment and builds an internal CMDB which you can access using sysapi.

Note: In order to use this feature, you must enable it within the System API properties.

Using the analyzer

The analyzer is a process designed to improve overall performance. The analyzer automatically runs with every full cycle. You can also initiate the analyzer process on demand, whenever you want to improve system

performance. Click the adjacent  button to activate the process.

Restarting DRA

The DRA restart tool runs a DRA restart process that restarts the Apache Tomcat service and reloads the database. The restart process takes less than one minute to complete, and maintains current settings and data. Click the adjacent



button to activate the process.

Using the Tickets module

This chapter includes the following topics:

- [“Overview”](#) on page 255 introduces the Tickets module and describes its organization.
- [“Searching for a ticket”](#) on page 258 describes the ticket search operation that enables you to search for a specific ticket or ticket type.
- [“Performing operations on tickets”](#) on page 268 describes the operations you can perform during the lifecycle of a ticket.

Overview

The Tickets module is an investigative tool that lets you view information and resolve the gaps that DRA detects.

The Tickets module lets you filter the list of tickets in your system to find a specific ticket or ticket type. After DRA displays your query search results, you can do the following:


- Group the results by criteria that you specify.
- Drill down to see ticket details.
- View the topology related to the ticket in the Topology module.

A full description of the ticket is provided, including an explanation of the problem and the items in your environment that are involved.

The impact on your system is also shown with suggestions on how to resolve the ticket.


Note: To see new tickets after a full cycle, you must run a search.

To access the Tickets module, click the Tickets tab.




The Filter Tree is displayed by default.

The Ticket Search page contains a toolbar for convenient access to module operations:

-  Click to display a pane in which you can filter the tickets to be shown in the results.
-  For more information about Filter Trees, see page 264.
-  Click to capture the result of a ticket search query and add it to the Session pane.
-  For more information, see “[Session pane](#)” on page 58.

The table at the bottom of the Ticket Details area logically groups all the items relevant to the currently selected ticket. To view more information about a ticket, click the **+** on the row you want to expand, as shown below:




Plus signs

To expand all rows in the Ticket Details Area table click the Expand/Collapse

Table  icon.

After the table has been expanded entirely, an Excel  icon appears beneath the table. You can click this icon to open the ticket in Excel.



Excel icon

Searching for a ticket

DRA has a ticket search query that lets you enter search criteria and find specific tickets or ticket types in the DRA database. Searching for specific ticket information is a valuable tool for analyzing and resolving problems in your environment.


The following procedure explains how to:

- Specify the range of dates, status, severity, and category of the tickets in which you are interested.
- Filter the search results even further by using the Filter Tree to view tickets that only belong to specific hosts or databases.
- Group similar tickets in the display.

To perform a ticket query

- 1 To open the Ticket Search page, click the Tickets tab.

Tickets tab




- 2 In the **Use Filter** drop-down list, select the filter to use for the query. Each filter has default parameters that are predefined for the filter.


You can select from the following filter options:

Open Tickets	Includes tickets with a status of Open or Reopen. Includes tickets from all categories with no specific detection date. This is the default.
Solved Tickets	Includes tickets with a status of Solved. This filter does not use a specific detection date.
Closed Tickets	Includes tickets with a status of Closed. This filter does not use a specific detection date.
Suppressed Tickets	Includes tickets with a status of Suppressed. This filter does not use a specific detection date.
Recently Opened Tickets	Includes tickets with a status of Open or Reopen that were opened in the last 24 hours.
Recently Closed Tickets	Includes tickets closed in the last 24 hours.
Recently Suppressed	Includes tickets suppressed in the last 24 hours.
Find Ticket by ID	Locates a ticket with the specified ID.
Find Ticket by Rating	Locates a ticket with the specified importance rating. For more details about using ticket ratings, see “Adding a rating to a ticket” on page 271.

Find Ticket by Note	Locates a ticket whose note contains the specified text string. For more details about using ticket notes, see “Adding a note to a ticket” on page 270.
---------------------	--




Find Ticket by Unique ID	Locates a ticket based on its unique ID. For more details about using unique ticket IDs, see “Generating a unique ID for a ticket” on page 272.
--------------------------	--




Note: You can change the default parameters for any of these filtering options as needed. For example, you can specify a detection date range (from/to) for the Solved tickets or Closed tickets filters.

- 3
- In the **Detected between** (from/to) fields, enter the date range for the search.
- 4
- To open a window in which you specify the business entity for the query, click the Business Entities link. The following window opens:



- 5
- Select the relevant business entity check box(es).

- 6 To open a window in which you specify the ticket status for the query, click the **Status** field.




- 7 Select the **Filter by Status** check box and the relevant status check boxes. You can select from the following options:

Open	Includes all currently open tickets
Resolved	Includes all resolved tickets
Suppressed	Includes all currently suppressed tickets
Reopen	Includes all reopened tickets
Closed	Includes all closed tickets

- 8 To save your settings and exit the window, click **OK**.

- 9
- To open a window in which you specify the ticket severity for the query, click the **Severity** field.




- 10
- To specify that all tickets with this severity are included, select the **Filter by Severity** check box and the relevant severity check boxes. You can select from the following options:


Error	Refers to tickets with a high risk level
Warning	Refers to tickets with a medium risk level
Information	Refers to tickets with a low risk level

- 11
- To save your settings and exit the window, click **OK**.

- 12 To open a window in which you specify the ticket category for the query, click the **Category** field.




- 13 To select any of the listed categories for the query, mark the relevant check boxes.
- 14 To open a window in which you specify the ticket label for the query, click the **Ticket Label** field.




- 15 To select any of the listed labels, mark the relevant check boxes.

- 16 To open a window in which you indicate whether to display the query with or without the external ticket, click the **External Ticket** field.
- 17 Mark the relevant check boxes to either show the query with or without the external ticket.





- 18 To save your settings and exit the window, click **OK**.
- 19 To display tickets that match the specified criteria, at the top right of the window, click **Search**.
- 20 To filter your search results further, click the **Filter Tree** button. This displays the following pane on the left:



The Filter Tree's highest branch specifies item types, such as Business Entities, Sites, or Hosts. You can select one of these branches to specify that only tickets that apply to these types of items appear in the search results. The Filter Tree's subbranches list specific items of the parent branch. For example, all the specific host names, as shown above. You can select a specific subbranch so that only tickets that apply to these items, such as a specific Host named Oranus, appear in the search results.


21 To group your search results, click the **Group Tickets By**

 button and select the grouping criterion from the drop-down list.




After you select an option, the tickets in the search results are grouped by that criterion.

For example, the following shows the search result tickets grouped by business entity, one business entity per row. To the left of each business entity, you can click the + to display the business entity's tickets that were found in your search.



Note: A ticket may belong to more than one group; therefore, when you use the Group By option, some tickets may appear more than once. For example, a host's ticket may belong to two or three business entities and may appear each business entity's group.

- 22
- To view more details for a specific ticket, click a row in the table. The ticket is highlighted in orange and tabs at the bottom of the page show more information about it. The tabs include summary information about the ticket, the gap to which this ticket applies, when DRA detected the ticket, the ticket's severity, any supplemental notes, and symptoms related to the ticket (when applicable).



After you get your ticket query results, you can perform a variety of operations on a ticket.

For more information, see [“Performing operations on tickets”](#) on page 268.


Note: You can also find a specific ticket by going to the bottom of the Tickets Search area, entering text in the **Search** field, and clicking **Find**.

Tickets in the Topology module

You can also open a ticket from this view to display its topology. To do so, select a ticket and click **Open Topology**. This action displays the ticket's topology at the top of the page in the Topology module and detailed information in a series of tabs at the bottom of the page. Information in these tabs is linked, and gives you easy access to complete details for the items related to this ticket.

In the following example, the elements shown in the Topology module relate to the item that triggered the ticket.

Note: In some tickets, the items that triggered the ticket may be highlighted.



Both the Tickets module and Topology module include the following tabs for each ticket:

Description tab	Describes the problem reflected in this ticket and lists the items involved.
-----------------	--

Impact tab	Describes the impact of this gap on your environment, such as the expected malfunction of the associated items with this ticket. The information varies, depending on the ticket.
------------	---

Resolution tab	Describes various solutions for resolving the gap.
----------------	--

History tab	Describes the states through which this ticket has gone.
-------------	--

Notes tab	Displays a note that the operator added to the ticket. For more information, see “Adding a note to a ticket” on page 270.
-----------	--

Symptoms tab	Displays symptoms related to the ticket.
--------------	--

Note: The Symptoms tab is available only when related symptoms exist.

See “Performing operations on tickets” on page 268 for more details.

Note: If you click an item in the graphical view, the Topology module changes its focus to that item and displays details for that item in the browser.


Performing operations on tickets

You can perform a variety of operations on a ticket during its lifecycle. To execute a ticket operation, go to the ticket list at the top of the Ticket Search page, select the ticket, and perform one of the following operations on it:

- Mark as Resolved


Resolves the ticket. A resolved ticket is treated as closed during the next full system scan.
- Suppress Ticket

Suppresses the ticket. A suppressed ticket indicates that you are aware of this problem, but you do not want to be informed about it again during the next full system scan. As shown below, when you suppress a ticket, you can specify a date when the suppression expires:



You can either select a date in the calendar and click Suppress or click Suppress Forever to suppress the ticket indefinitely.

After suppressing the ticket, the following window opens, in which you can add a comment about the suppression reason:



- Unsuppress Ticket

Unsuppresses a ticket that is currently suppressed. You can only use this option with a suppressed ticket.
- Suppress Gap

Suppresses a gap. A suppressed gap indicates that you are aware of this type of problem, but you do not want to be informed about any instance of this gap again.

For details on how to unsuppress a gap type, see [“Configuring gap management settings”](#) on page 229.
- Open Topology

Opens the Topology module displaying the items relevant to the ticket.

For more information, see [“Displaying item details in the Topology module”](#) on page 311.
- Copy Unique ID to Clipboard

Copies the ticket unique id to the clipboard.

For more information, see [“Generating a unique ID for a ticket”](#) on page 272.
- Assign a Ticket Label

Assigns a ticket label to a ticket.

For more information, see [“Assigning a label to a ticket”](#) on page 277.

To perform the following operations on a ticket, click the Symptoms tab in a ticket:

- Suppress Symptom



Suppresses a symptom. A suppressed symptom indicates that you are aware of this type of symptom, but you do not want to be informed about it again.

Suppress Symptom Type	Suppresses a symptom type. A suppressed symptom type indicates that you are aware of this type of symptom type, but you do not want to be informed about it again.
	

Note: Not all tickets have the Symptom tab. This tab is present only when a symptom exists.


For more information on working with tickets, see “[Ticket lifecycle](#)” on page 29. You can also perform operations on multiple tickets simultaneously. For more information, see “[Sending tickets directly from the application](#)” on page 272.

Adding a note to a ticket

An operator can add a note to a ticket that describes it in more detail. You can filter by note text to make it easier to find a specific ticket(s).

To add a note to a ticket

- ◆ In the Ticket Search page, select a ticket. At the bottom of the window select its Notes tab. In the **Ticket Notes** field, enter your note and click **Save**.



Ticket Notes field

Adding a rating to a ticket

You can rate a ticket to indicate its level of importance. DRA uses a five star rating scale. One star (★) indicates the lowest level of importance and five stars (★★★★★) the highest.

Group Tickets By: (no grouping) - Export Tickets						Ratings	
ID	Labels	Detection Date	Status	Summary	Rating	Severity	
252		Feb 25, 2011 3:21:07 PM		The shared Symmetrix devices storing SRDF/copy of datastore prodrep of datacenter PROD Datacenter .	★★★★★	△	
213		Mar 02, 2011 3:12:39 PM		The LVM DisallowSnapshotLun option is disabled on ESXi host esx1 at site Boston	★★★★★	△	
191		Feb 25, 2011 3:17:32 PM		Inconsistent access to replica Symmetrix devices of VMFS datastore datastore1 on ESX host esx2 at si...	★★★★★	△	
150		Sep 28, 2010 5:13:00 PM		A passive node in VCS VIAcluster at site Boston is missing a required mount point directory	★★★★★	△	
72		Sep 28, 2010 5:10:31 PM		Essential paths of Apache resource ApacheDummy in VCS VIAcluster at site Boston were not found.	★★★★★	△	
229		Mar 03, 2011 7:18:15 AM		Windows VM vcprod at site Boston has swap files stored on RAID5 devices	★★★★★	△	
287		Feb 27, 2011 3:14:02 PM		Standby hosts esx2 at site Boston and esx2-dr at site Chicago use the same DNS server from distant...	★★★★★	△	

You can filter ticket ratings to make it easier to find a specific ticket(s).

To add an importance rating to a ticket

- 1 In the Ticket Search page, select a ticket.
- 2 In the selected ticket row, click the appropriate star to indicate the level of importance of the ticket. The more stars, the more important the ticket.

Stars indicate importance

DashboardTopologyTicketsReportsSLAComparisonConfiguration

Tickets Search - Found 82 distinct tickets

User Filter: Open Tickets

Detected between: [] and []

Search

Status Open / Reopen Security All Category General Ticket Label All External Ticket All Reset Filter

Group Tickets By: (no grouping) ... | Exclude Tickets: ...

ID	Labels	Detection Dates	Status	Summary	Rating	Severity
248		Feb 25, 2013 12:20:21 PM	[icon]	Inconsistent access to the Symmetrix devices storing database prod of datacenter PROD Datacenter at site Chicago.	★★★★★	▲
194		Feb 25, 2013 11:38:17 PM	[icon]	HPUX host cbsmrst at site Chicago is using replicas of storage volumes which are only used by non-	★★★★★	▲
78		Sep 25, 2010 5:11:29 PM	[icon]	Solaris cluster node cmdm of VCS cksr2 at site Boston is utilizing storage volumes located at site Chic.	★★★★★	▲
262		Feb 25, 2013 12:23:03 PM	[icon]	ESXi cluster node moonf of ESXi cluster room at site Boston has detected SAN I/O paths	★★★★★	▲
218		Mar 02, 2011 12:13:39 PM	[icon]	The LUN DisallowOnSaphanaLun option is disabled on ESXi host esx2 at site Boston	★★★★★	▲
166		Feb 25, 2011 11:08:00 PM	[icon]	Missing license or target NetApp file dgrlcr at site Chicago compared to its SnapImmaged source Net.	★★★★★	▼
162		Oct 08, 2010 06:25:25 PM	[icon]	Storage Solaris host angel at site Chicago does not have the same mounted network file systems as its.	★★★★★	▼
148		Feb 25, 2010 5:13:00 PM	[icon]	A mount response of VCS MAC200 at site Boston is configured for automatic mount on boot on a	★★★★★	▼
129		Sep 28, 2010 5:11:52 PM	[icon]	Suboptimal storage volume layout of multiplexed control files for Oracle database dbf1 on Linux host tyf.	★★★★★	▼
258		Feb 25, 2011 12:58:34 PM	[icon]	Storage array inconsistency was detected on Symmetrix devices storing VMS databases datacenter2 on	★★★★★	▲
1		Feb 25, 2011 12:28:34 PM	[icon]	Windows host blaked at site Chicago is using replicas of unified storage volumes	★★★★★	▲
223		Jun 01, 2011 17:18:15 AM	[icon]	Linux VM errorpr2 at site Boston has swap files located on RAID5 devices	★★★★★	!
64		Sep 28, 2010 5:10:06 PM	[icon]	Liv's support on the DataGroup resource log_inq_dg are not configured as resources in VCS cksr2 at sit.	★★★★★	▼
19		Feb 25, 2010 4:51:19 PM	[icon]	Veritas DQ sleep on Solaris host rileap at site Chicago contains a mixture of database files	★★★★★	▼
70		Feb 25, 2013 11:36:39 PM	[icon]	Datastore disk of datacenter PROD Datacenter at site Boston is stored on VGP local tapes that have.	★★★★★	▼
205		Sep 28, 2010 5:10:45 PM	[icon]	Veritas DQ sleep on Solaris cluster node cmdm of VCS cksr2 at site Boston is stored on Symmetrix device.	★★★★★	▼
211		Feb 25, 2013 11:18:39 PM	[icon]	Datastore prodreg of datacenter PROD Datacenter at site Boston is stored on Symmetrix device that h.	★★★★★	▲
263		Feb 25, 2013 12:23:03 PM	[icon]	Solaris host rileap at site Chicago has detected SAN I/O paths	★★★★★	▲
217		Feb 25, 2013 12:10:07 PM	[icon]	The shared Symmetrix devices storing SDRFS copy of datastore prodreg of datacenter PROD Datacent.	★★★★★	▲
191		Feb 25, 2013 11:17:32 PM	[icon]	Inconsistent access to replica Symmetrix devices of VMS databases datacenter1 on ESXi host esx2 at si.	★★★★★	▲

Page 1 of 2

Search

Log Exit

Displaying 1 - 50 of 82 Show 50 ▼

Generating a unique ID for a ticket

DRA lets you generate a unique ID for a ticket. You can use this ID across environments to find the ticket, regardless of the computer or environment on which it resides.

When DRA generates unique ID for a ticket, it automatically copies the ID to the clipboard. You can paste the ID into an email and send to the person handling the ticket. They can use the Find Ticket by Unique ID filter option to search for the ticket in their system.

For more information, see page 259.


To generate a unique ID for a ticket

- ◆ In the Ticket Search page, select a ticket and click the **Copy Unique ID to**

Clipboard  button to generate a unique ID for the ticket and copy it to the clipboard. After the ID is in the clipboard, you can select it and paste it where needed.


Sending tickets directly from the application

DRA lets you send a selected ticket directly from the application. You can access this option from the Details tab.



To send a ticket from the application

- 1 On the Tickets Search page, in the Details tab, click **Send Ticket** (). The Send Ticket Options window opens.



The image shows a 'Webpage Dialog' window titled 'Send Ticket Options'. It contains the following fields:

- Mail Recipients:** A text box containing 'john.smith@company.com'.
- Export Type:** A dropdown menu currently set to 'PDF'.
- Message Body:** A large text area containing the text:


```
Hi John.
Can you please take a look at this TrueCopy
replication?

Thanks,
Tommy.{'
```


An 'OK' button is located at the bottom right of the dialog.

- 2 Enter details in the **Mail Recipients**, **Export Type**, and **Message Body** fields and click **OK**. By default, the **Mail Recipients** field contains the default distribution list specified in the Default distribution list for the reports system property (see “[Email settings](#)” on page 246). You can change this default setting, when required.

Exporting tickets directly from the Tickets module


DRA lets you export tickets directly from the Tickets module. You can access this option from the Ticket Search area, as shown below.

Export Tickets



To export tickets from the Tickets module


- 1
- In the Ticket Search area, click the down arrow on the **Export Tickets** button to open a drop-down list.



- 2
- Select the format for the export. The tickets in the Ticket Search area are exported as a Ticket Details report in the specified format (PDF or Word).

Performing operations on multiple tickets

DRA lets you perform operations on multiple tickets simultaneously. To do so, use Ctrl + Shift to select multiple tickets in the Ticket Search page and then perform one of the operations listed in the Details area.



When you select multiple tickets, the Details area displays the following tabs:

- **Actions:** Lets you perform the following operations on the group of selected tickets:

Details

Tickets 194, 195

Actions Notes Labels

Action on selected tickets

☒ Take no action

☐ Mark selected tickets as resolved

☐ Suppress selected tickets with no time limit

☐ Suppress selected tickets for limited period of time

☐ Suppress gap type of selected tickets

☐ Change rating for selected tickets

★★★★★ (clear)

Submit

- Take no action.
- Mark selected tickets as resolved.
- Suppress selected tickets with no limit.
- Suppress selected tickets for limited period of time.
- Suppress gap type of selected tickets.
- Change rating for selected tickets.

To change the rating for the tickets, select the appropriate rating (number of stars) to be applied to the tickets and click **Submit**. The rating for all selected tickets is changed to the designated rating.

- **Notes:** Lets you add or change a note for the selected tickets.

Details

Tickets 194, 195

Actions Notes Labels

Ticket Notes

☐ Do not change any ticket note ☐ Add note for tickets without note only ☐ Override notes of all selected tickets ☒ Append text on all selected tickets


Clear

Submit

Available options are as follows:

- Do not change any ticket note.
- Add note for tickets without note only.
- Override notes of all selected tickets.
- Append text on non-empty ticket notes.

- Labels: Lets you add labels and remove labels for the selected tickets.



The screenshot shows a web interface for the 'Details' tab of a ticketing system. At the top, there's a header 'Details' and a sub-header 'Tickets 194, 195'. Below this, there are three tabs: 'Actions', 'Notes', and 'Labels'. The 'Labels' tab is currently selected. Under the 'Labels' tab, there are two radio buttons. The first one is labeled 'Add labels to selected tickets' and has a green plus icon next to it. The second one is labeled 'Remove labels from selected tickets' and has a green minus icon next to it. At the bottom of the form, there is a 'Submit' button.

To perform operations on multiple tickets

- 1 Select the tickets in the Ticket Search page. This action automatically lists the operations you can perform in the Details area.
- 2 Select the radio button for the applicable operation and click **Submit**. The selected operation is performed on all selected tickets.


Assigning a label to a ticket

You can add one or more labels to a ticket in order manage tickets more efficiently. Default labels are automatically assigned by the system when a ticket is detected. Default labels can be changed, added, or deleted.


To add a label to a ticket

- 1 Select the ticket to which you want to add a label and click **Manage Labels**

() in the Details pane in the Description tab. The following window opens:




- 2 Select the label(s) to add to the ticket in the left pane and then click .
- 3 Click **OK** to save the settings. The label is assigned to the ticket and displays in the Details pane.



Manage Labels button Ticket labels

Note: To edit the labels assigned to a ticket, click **Manage Labels** () and then modify the labels currently assigned to the ticket.

To remove a label from a ticket, click the **X** adjacent to the label name in the Details pane.



Using the Topology module

This chapter includes the following topics:

- [“Overview”](#) on page 280 introduces the Topology module and describes its organization.
- [“Accessing the Topology module”](#) on page 294 describes the multiple methods available for accessing the Topology module.
- [“DRA object model”](#) on page 295 describes the DRA object model and its representation of basic item types and connection types in the application.
- [“Item display methods in the graphical view”](#) on page 298 describes the item types that you can display in the topology.
- [“Connection display methods”](#) on page 299 describes the multiple connection types that you can display in the topology.
- [“Searching the topology”](#) on page 300 describes how to perform a search query for a specific item or item type in the topology.
- [“Performing Topology module operations”](#) on page 303 describes the Topology module toolbar and its operations.
- [“Displaying item details in the Topology module”](#) on page 311 describes general principles that apply when viewing information in the Topology module's browser.

Overview

The Topology module is an investigative tool that lets you view information and resolve the tickets that DRA detects.


The Topology module lets you view and browse through a variety of graphs and tables about any item in your environment.

To open the Topology module, click the Topology tab. The Topology module shows a high-level map of your environment and a series of tabs of information.

The Topology module maps the devices in your environment and their interrelationships. Extensive drill-down capabilities are available for quick navigation, and enable you to easily examine relationships between items, gaps and tickets. You can search in your environment using queries to locate a particular item or item type in the topology.

The Topology page consists of the following primary areas:

- Session pane
See [“Session pane”](#) on page 58.
- Topology graphical view
See [“Topology graphical view”](#) on page 281.
- Topology detail tabs
See [“Topology details tabs”](#) on page 283.
- Topology toolbar
See [“Topology toolbar”](#) on page 285.
- Navigator
See [“Navigator”](#) on page 292.
- Topology tree
See [“Topology tree”](#) on page 293.




Topology graphical view

The DRA Topology module graphical view shows an icon for each item. You can click on an icon to display more details in the detail tabs at the bottom of the page. You can also click a direct connection line in the graphical view to view the properties for that item, other items connected to this item, history of this item and the item's related open tickets. In general, the Topology module and its browser display the complete details for the currently selected item and make that item the focus in the browser.


Note: Graphical topology shows both direct connections and indirect connections. Direct connections, such as between a volume group (VG) and a physical volume, are blue and can be clicked to show details. Indirect connections, such as between two arrays, are light blue and do not have details.

When you select an item or a connection line in the graphical view, it is highlighted. The selected item in the topology is highlighted in orange. Correspondingly, the table at the bottom of the page displays details for this item.



You can move any item in the topology graphical view by dragging it with the mouse. The item's connection lines move with it.


When you click on an item in the graphical view, its details are shown in the detail tabs, which are described on page 283, and in the topology tree, which is described on page 293. The icon representing the element is also highlighted. Items in the graphical view are aggregated, as shown in the following screen:



The topology aggregation function organizes devices with the same type into a single group. The x# value in the item's label indicates the number of devices that are aggregated for this specific item in the topology. For example, x7 indicates that seven devices are aggregated for that item.

To display the aggregated devices in the graphical view, click the item to select it

and then click **Expand Selected Entities** .



Topology details tabs

The table at the bottom of the page displays detailed topology information. This area consists of the Ticket and Item tabs.

The Ticket tab shows detailed information for a ticket, when a specific ticket is the focus of the topology. It presents ticket details in a series of subtabs: the Description, Impact, Resolution, History, Notes, and Symptoms (when applicable) tabs.

For more information, see [“Tickets in the Topology module”](#) on page 266.

Similarly, the Item tab displays detailed information for an item when it is the focus of the topology. It presents ticket details in a series of subtabs: the Properties, Related Tickets, Connections, and History tabs.

These tabs display the following information for the item of focus:


- Properties tab

Presents descriptive information about this item
- Related Tickets tab

Identifies tickets affecting this item
- Connections tab

Shows relationships between this item and other items
- History tab

Lists all the changes applied to this item



Topology toolbar

The Topology module toolbar, shown below, provides easy access to various viewing and navigation options. In addition, other viewing and navigation options are also available:



Zoom In/Zoom Out

Each click on one of these buttons either zooms in or zooms out the graphical view. Zooming in lets you see more detail about the items. Zooming out lets you see a larger overview of the topology. You can also use the mouse wheel to zoom in or out. Roll it away from you to zoom in, and roll it towards you to zoom out.



Zoom Area


Select an area and the map zooms in to this area.




Fit Contents

Sizes the display of the contents of the graphical view to fit it within the current size of the graphical view pane. You can use this option for any zoom or layout that you use in the graphical view pane.

Before Fit Contents



After Fit Contents




If you enlarge the graphical view pane, for example by hiding the Session pane, the sizing of the topology does not change. In this case, you should click the Fit Content tool again to display it in the pane's new size.



Hierarchy Layout

Displays the items in the topology browser in a hierarchical manner, letting you see the containing items in one layer and the items that it contains in another layer.

For example, a site may contain a host, which contains a volume group (VG). This same site may contain an array also. In Hierarchical View, the site is shown on the top level, the host and array in a level below it and the VG in a level below that.




The first click of this button shows the topology in a hierarchical manner. Each subsequent click rotates the topology 90 degrees clockwise. You can view the topology in Hierarchical layout or Incremental layout as best suits your needs.



Incremental Layout

This option displays the items in the topology browser in an incremental manner enabling you to see the containing items and the items that it contains inside it.


For example, a site may contain a host, which contains a VG. This same site may contain an array also. Incremental View shows the site containing the host and the array. It shows the host containing the VG.



The first click of this button shows the topology in an incremental manner. Each subsequent click rotates the topology 90 degrees clockwise. You can view the topology in Hierarchical layout or Incremental layout as best suits your needs.




Show Connection Labels (On/Off)




Centralize on Select (On/Off)

When you select an entity from the ticket section or from the left tree, the item gets focus.



High Level

Shows a high level view of your IT structure that contains the sites, storage arrays, and hosts that were scanned, and the relationships between them.



Search

Accesses the Search function in the Topology module.

For more information, see [“Searching the topology”](#) on page 300.



Capture

Lets you capture the current topology and search related elements, and add it to the Session pane.

For more information, see [“Session pane”](#) on page 58.



Filter Displayed Content

Lets you manage the displayed items of the selected item's subtopology.

For more information, see [“Filter displayed content”](#) on page 304.



Expand Entity

Expands the selected item to show its children.

For more information, see [“Expand Entity”](#) on page 306.



Expand Selected Entities

Displays all entities under the selected entity.

For more information, see [“Expand Selected Entities”](#) on page 308.



Collapse Selected Entities

Hides all entities under the selected entity.

For more information, see [“Collapse Selected Entities”](#) on page 307.



Display All Connected Entities

Displays all entities connected to the selected entity.

For more information, see [“Display Connected Entities”](#) on page 309.



Display Storage.

Displays all storage arrays connected to the selected entity.



Display Replications.

Displays the replications/sources of the selected storage volumes.

Note: If the selected entity is a DB/VG/FS, it displays the storage volumes and their replications. If the selected entity is a host, it displays the hosts that are connected to the replica storage devices of the volumes storing the selected host.

For more information, see “[Display Replications](#)” on page 311.



Hide Selected Entities

Removes the selected entities from the topology.

For more information, see “[Hide Selected Entities](#)” on page 305.




Export Topology

This option exports the current topology to a JPEG file, which can be downloaded directly to your file system.




Print Preview

The first time this option is used in a session, the following window opens, in which you authorize the use of the printer by the DRA applet.




Click Yes to display the following window:



At the top of this window, buttons let you change the page format, print, zoom in, and zoom out.

You can also click the Options button to display the following window:



This window lets you tile the printout of this topology so that you can print it very large and cover more than one sheet of paper. The options in this


window let you specify the number of posters (that is, the number of sheets of paper or tiles) that comprise the printout.

Poster Rows	Specifies the number of horizontal pages across which the topology is printed.
Poster Columns	Specifies the number of vertical pages across which the topology is printed.
Add Poster Cords	Specifies whether the tiling coordinates are printed on each page, making it easier for you to organize the pages after they are printed.
Clip Area	Specifies which area to print. You can print the view that is currently displayed on the screen (which may be a portion of the graph or larger than the graph depending on the current zoom setting), or you can print the entire graph in which case the whole graph is printed, no matter the current zoom shown on the screen.


For more information about using these tools to perform operations in the Topology module, see [“Performing Topology module operations”](#) on page 303.

Navigator

The Navigator provides another option for controlling the graphical Topology View. It has a gray viewing rectangle that lets you specify the area shown in the center of the graphical view area.




To move the area shown in the center of the graphical topology view, click and drag the gray viewing rectangle.



Topology tree

The topology tree provides another option for controlling the graphical view. It provides a tree that lists the elements displayed in the graphical view and lets you double-click on any item to bring that item to the center of the graphical view, as follows:

Topology tree



The filtering and display options that apply to the items in the graphical view also affect the items in the topology tree.

When you click on an item in the topology tree, its details are shown in the detail tabs, which are described on page 283. The element's icon is also highlighted.

Accessing the Topology module

You can access the Topology module in the following ways:


- Select a ticket in the Tickets module and then click the Open Topology button.
See “Opening the Topology module from the Tickets module” on page 294.
- Click on a ticket in the Dashboard.
See “Opening the Topology module from the Dashboard” on page 295.
- Click the Topology tab.

Opening the Topology module from the Tickets module

You can open the Topology module with a specific ticket as its focus. To do so, select a ticket in the Tickets module and click Open Topology.

At the bottom of the page, the following subtabs present detailed information for the items relevant to the ticket: Description, Impact, Resolution, History, Notes, and Symptoms (when applicable).

Note: You can also access the Topology module by clicking the links in a ticket, as presented in the ticket’s Description, Impact and Resolution tabs in the Tickets module.



Opening the Topology module from the Dashboard

You can also open the Topology module by clicking a ticket ID in the Dashboard. In this case, the application switches to the Topology module, displaying the selected ticket as its focus.

DRA object model

The Topology module creates its representation of your environment using an object model. The object model defines the basic item types and connection types that are used in the browser to present relationships between them.

Note: Not all item types can use all relationship types.

[Table 8-1](#) lists some examples of the graphical representations of the objects that appear in the topology browser. Each object is described in the [Glossary](#) on page 377.

Table 8-1 Topology browser icon examples

Icon	Object
	Site
	Business entity
	Domain
	Host
	File system, local file system
	Network file system
	LV
	VG
	PV, SAN PV
	DAS PV

Table 8-1 Topology browser icon examples

Icon	Object
	Storage array
	Symmetrix DG, Symmetrix CG, Symmetrix RDF group, or CLARiiON Snapshot Session
	NetApp Aggregate
	NetApp Volume
	NetApp Q-Tree
	Network Share, NFS Export/NFS Share, or CIFS Share
	Database instance
	Database data file, database data files
	Database control file, database control files
	Database log file, database log files
	Database archive destination, database archive destinations
	SV, Symmetrix Device, Symmetrix Meta Device, Symmetrix Meta Head, Symmetrix Meta Member, NetApp LUN, unknown array device, BCV, SRDF, clone, or snapshot
	Any Type of Connection
	Veritas Cluster Server
	Veritas cluster service group
	Veritas cluster resource
	Veritas cluster node
	Veritas cluster local application
	Symmetrix array


Table 8-1 Topology browser icon examples

Icon	Object
	HDS
	NetApp Filer
	CLARiiON
	Cluster
	Global/campus cluster


Item display methods in the graphical view

DRA distinguishes between the following types of items in the Topology module graphical view:

- Compound items
Compound items contain other items that are currently displayed in the Topology module graphical view. The following figure shows one compound item (named as17), which contains multiple simple items:



- Simple items
Simple items do not contain other items that are currently displayed in the Topology module graphical view. These simple items most likely have child items, but the children are not currently displayed in the browser. For more information, see [“Connection display methods”](#) on page 299.
All the items in the following topology are simple items:



Connection display methods

In the Topology module's graphical view, DRA uses a line to display connections between items. Connections can be either direct or indirect.

Direct connections

Direct connections directly connect two simple items. A bold blue line indicates a direct connection between two items and always shows the type of the connection below the line. An arrow indicates the direction of the connection. Such connections are always one to one.


In addition, when the user clicks on a direct connection, the connection details are shown in the Item tab.

Indirect connections

Indirect connections indicate that one of the children of a simple item has a direct connection with another item or one of its children in the graphical view. A light blue line indicates an indirect connection between items.

For example, in the figure below, the local file systems in the rgdb-xp host have a direct connection to their physical volumes and indirect connections between the physical volumes and the NetApp Filer.

Unlike a direct connection, an indirect connection does not show the name or direction of the connection. When you click on the connection, DRA does not display any details.





Searching the topology

You can perform Topology module queries to supplement your data analysis and problem solving. These queries let you search for a specific item or item type in the topology. Data is retrieved from the DRA database when executing the queries. The parameters you specify for a query may vary, depending on the item(s) specified for the query.

To perform a search in the Topology module

- 1 To access the Topology module, click the Topology tab.

- 2 Click . The Topology Search page opens.



- 3 In the **Search For** drop-down list, select the item type for the query. Most search options also require that you specify one or more parameters. You can select from the following search options:
 - **Sites**
When you select this option, no additional parameter is required.
 - **Entities**
For this option, you must also specify the Entity Type parameter. In the Entity Type drop-down list, select the required element/item type. Entity types include elements used in the DRA application, such as hosts, storage arrays, file systems, domains, and so on.
 - **Hosts**
When you select this option, you must also specify the site(s) on which to search for the host.

- **Storage Arrays**
When you select this option, you must also specify the site(s) on which to search for the storage array.
- **Relationships**
For this option, you specify three additional parameters for the search function, including the Entity Type From (entity type), Entity Type To (entity type) and Relationship Type, meaning the type of connection between the from/to entity types.

Topology Search - Relationships

Search For: Relationships

Show -- All -- relationships

between -- All -- entities and -- All -- entities

Search

- 4 Click **Search**. Results that meet your query criteria are displayed in a table at the top of the page. To display additional details for a specific item in the table, click the relevant row to select it. The selected row is highlighted in orange. At the bottom of the page, a series of tabs displays more information about the selected item. Click a tab to view its details.

Topology Search - Entities (20)

Search For: Entities

Entity Type: Host

Search

Entity Parents	Entity	Entity type
Site Unknown Site	host1	host
Site Unknown Site -> Datacenter prd -> Host localhost	wn-pool3	host
Site site Q -> Cluster Q2	hostg2-3	host
Site Unknown Site	c1p2	host
Site site Q	hostb2	host
Site site R	hostr1	host

Page 1 of 1

Search

Reg Exp

Displaying 1 - 20 of 20

Show 20 items

Entity

hostb2

Open Topology

Properties

Related Tickets

Connections

History

General

Item Name:

hostb2

Item Type:

HPUX host

Detection Date:

Jul 9, 2012 5:45:23 PM

Properties

IR: [139, 141]


Name

hostb2

Primary IP

111.11.1.28

When you click Open Topology, a map of some of the items in your environment is displayed at the top of the page (graphical view).




To select an item in the graphical view, click on it. At the bottom of the page, a series of tabs displays more information about the item. Click a tab to view its details

For more information on each tab, see [“Topology details tabs”](#) on page 283.

Note: When you click an item in the graphical view, the Topology module changes its focus to that item. The page refreshes to display details for that item.


Performing Topology module operations

The Topology module contains a suite of tools that let you perform operations in the browser. The following operations can be performed:



High Level

Shows a high level view of your IT structure that contains the sites, storage arrays and hosts that were scanned, together with the relationships between them.



Topology Search

See [“Searching the topology”](#) on page 300.



Capture

Adds the current topology to the Session pane.

See [“Session pane”](#) on page 58.



Filter Displayed Content

Filters the items' subtopology display.

See [“Filter displayed content”](#) on page 304.



Expand Entity

See [“Expand Entity”](#) on page 306.



Expand Selected Entities

See [“Expand Selected Entities”](#) on page 308.



Collapse Selected Entities

See [“Collapse Selected Entities”](#) on page 307.



Display Connected Entities

See [“Display Connected Entities”](#) on page 309.



Display Storage

See [“Display Storage”](#) on page 310.



Display Replications

See [“Display Replications”](#) on page 311.



Hide Selected Entities

See [“Hide Selected Entities”](#) on page 305.



Print Preview

Provides a print preview of the printout before it is sent to the printer.



Export Topology

Exports the topology as a JPEG image and stores it on the client.

Filter displayed content



The Filter Displayed Content button lets you add items to or remove items from the topology. For example, as shown in the figure on page 306, /loglv00 was initially selected as the child item to add to the topology in the Expand window.

To add additional child items to the topology, click Filter Displayed Content




. A wizard opens that takes you through the process of specifying the types of child items that you would like to see. It then displays those items. The following options are provided:

- **Contained Items**
Lets you specify the descendent items to be shown. Select the descendents type to view and then the actual descendent items to be included in the view. For example, you could select the following:
 - Storage array contains storage volumes
 - Host contains file systems, volume groups, and so on

■ Related Items

Lets you specify the related items (connected to the select item or one of its descendents) to be shown. Select the related item type to view and then the actual items to be included in the view. For example, the volume group (VG) is stored on a storage volume. When you click the filter for the VG, storage volumes become a related item for the VG.




Hide Selected Entities

The Hide Selected Entities  button deletes all selected items. To hide child items in the graphical view, click a parent item to select it and then click Hide Selected Entities.

Note: Using this tool performs the same action as clicking the Expand Entity tool, deselecting the currently selected child check box(es) and then clicking OK.

Expand Entity


The Expand Entity  tool expands the current topology view to show the children of the currently selected item. You specify which child item(s) you want to include in the current view. When you click the Expand Entity tool, the Expand window opens, in which you select the children and subcomponents you want to display.



Note: This window only shows direct children of the selected item.

To filter the rows, use the search field at the bottom of the table. You can also perform searches in this window to easily locate an item, based on its name. To do so, enter an item name in the Find field. You can also specify the name using a regular expression by checking the Reg Exp check box and then entering your text in the Find field. To match a regular expression for an item name, select the Reg Exp check box. After you define your selection criteria, click OK.


The item of focus in the topology plus the children you selected to view (in the Expand window) is shaded with orange in the graphical view. In addition, the child item(s) you select are shown in the Expand window at the top of the list.



Collapse Selected Entities

The Collapse Selected Entities  tool lets you collapse the topology view so that an item's children are no longer displayed. For example, in the following figure, the bold, red item shows the topology after clicking Collapse Selected Entities.

To remove child items from the graphical view, click a parent item to select it and then click Collapse Selected Entities.




Expand Selected Entities

The  Expand Selected Entities tool lets you expand all items in the topology to show all their children and connections. The following example shows a before and after view when using this tool:

Before




After




Display Connected Entities

The Display Connected Entities  tool lets you show only the internal direct connections for a selected item. The following example shows a before and after view when using this tool:

Before




After




Display Storage

The Display Storage  tool displays all the storage devices connected to the selected entity. The following example shows a before and after view when using this tool:

Before



After




Display Replications

The Display Replications  tool displays all the replications of the selected storage device. The following example shows a before and after view when using this tool:

Before



After



Displaying item details in the Topology module

By clicking a link in the graphical view of the Topology module, the complete details for the selected item are displayed at the bottom of the window in a series of tabs. All information in the Topology module is linkable, which means that you can click an item to drill down to its contained elements and/or subcomponents.


In general, whatever item you select in the Topology module becomes the focus for display in the browser.

The focus for the browser varies depending on the way you access it.

- Accessing the Topology module by selecting an item in the module
The browser shows the topology of that item in the graphical view and relevant details in the tabs at the bottom of the page, including the properties of the item, related open tickets, connections and history.
- Accessing the Topology module from the Tickets module
Selecting Tickets > Open Topology opens the Ticket Topology view and displays only the most relevant elements for that specific ticket. The upper pane displays a graphical description of the ticket. The following tabs are shown:
 - The Ticket tab shows the ticket information.
 - The Item tab shows the item information, as detailed in the previous element.

Properties tab

The Properties tab displays information about the properties of the item as follows:



The display includes following information about the properties of the item:

Item Name	Name selected for the item
Item Type	Type of item
Detection Date	Date the item was first detected on the system
Item Properties	List of properties that define the item

Related Tickets tab

The Related Tickets tab shows the tickets relevant to the selected elements on the topography as follows:

ID	Detection Date	Status	Summary	Severity
154	Sep 11, 2009 11:29 AM		VG udwbg on Aix host udwbg at site Boston is stored on Symmetrix devices that have partial replica sets	
33	Sep 2, 2009 7:51:31 PM		Multiple databases are stored on VG udwbg on Aix host udwbg at site Boston	

Connections tab

The Connections tab shows a tabular list of the connections in the graphical view, as follows:

Relationship	Related Item	Related Item Type
Stored on	2407	Symmetrix device

Relationship Related Item Related Item Type

The connections described in the tabular list are defined as follows:

Relationship	Type of connection.
Connection Direction	The relationships are directional. For example, in a “stored on” relationship, the direction is important.
Related Item	Item the selected item connects to or is connected by.
Related Item Type	Type of connected item or connection.

History tab

The History tab shows a tabular list of all the changes made to the selected item, as follows:

Date	Type	Subject	Previous Value	Value
Sep 02, 2009 3:44:06 PM	Entity found	udbvg	-	-
Sep 02, 2009 3:44:06 PM	Entity verified	udbvg	-	-
Sep 13, 2009 9:04:30 AM	Property detected	maxSnapAge	-	NoReplicas
Sep 13, 2009 9:04:30 AM	Property detected	primaryItemID	-	109580
Sep 13, 2009 9:04:30 AM	Property detected	minSnapAge	-	NoReplicas
Sep 13, 2009 9:04:30 AM	Property detected	correctness	-	VALID
Sep 13, 2009 9:04:30 AM	Property detected	numOfSnaps	-	0
Sep 13, 2009 9:04:30 AM	Property detected	replicationSet	-	-
Jan 13, 2010 9:04:30 AM	Property detected	callName	-	106800

Labels below the table: Date, Type, Subject, Previous Value, Value

The list of changes includes the following:

- The creation of an item
- The expiration of an item
- Addition, expiration, and changes of value of a property
- Addition, change, and expiration of a connection for items

The following information is included about these changes:

Date	The date that the change occurred.
Type	The type of change.
Subject	The element that was changed. This is either the item/connection or one of its properties.
Previous Value	The value that the item had before it was changed. If the item is new, there is no previous value.
Value	The current value of the element in the list.

DRA reporting

This chapter includes the following topics:

- [“Overview”](#) on page 317 introduces the report generator and describes how to run and open a report.
- [“Exporting reports”](#) on page 319 describes how to export reports into other formats.
- [“Report types”](#) on page 319 describes the various report types provided in DRA and how they are organized.

Overview

DRA's report generator automatically produces detailed reports describing your configuration and the gaps that it detected from information extracted from the DRA database.

Some reports have parameters that let you specify the information that you are interested in before you run the report.


Some reports may take some time to run. You can run multiple reports at the same time.

Note: DRA also lets you schedule when reports are automatically generated and sent to one or more email destinations that you configure. New reports automatically generate each time before they are sent.

For more information, see [“Report scheduling”](#) on page 144.

To run a report

- 1 Click the Reports tab. The Reports page opens.



- 2 In the **Report** drop-down list, select the report that you want to run. The relevant reports page is displayed.
- 3 For a report that has parameters that you need to specify, enter the required parameters and click **Generate**. The report is displayed at the bottom of the Reports page.

Note: Some reports may take some time to execute. While waiting for a report to finish running, you can browse and do other work in the DRA application.


You can open a report directly from the Session pane. For more information, see [“Session pane”](#) on page 58.

To open a report from the Session pane

- ◆ To select a report, click it in the Session pane. The report automatically opens in the right pane of the window.

Exporting reports

Many DRA reports let you export their content into PDF or Word format. The reports that provide this feature display the relevant tools at the top left of the report window, as follows:



Exports to a PDF file with a portrait orientation



Exports to a PDF file with a landscape orientation



Exports to a Word file with a portrait orientation



Exports to a Word file with a landscape orientation



Exports to an Excel file

Report types

DRA has the following reports:

- Scan Status
See [“Scan Status report”](#) on page 321.
- System Event Log
See [“System Event Log report”](#) on page 322.
- Ticket Details
See [“Ticket Details report”](#) on page 323.
- Optimization
See [“Optimization report”](#) on page 324.

- Un-replicated Data on Replicated Hosts
See [“Unreplicated Data on Replicated Hosts report”](#) on page 329.
- NetApp Filer Replication Summary
See [“NetApp Filer Replication Summary report”](#) on page 330.
- Unsynchronized Remote Replication
See [“Unsynchronized Remote Replication report”](#) on page 331.
- Old Replicas
See [“Old Replicas report”](#) on page 332.
- Standby Pairs
See [“Standby Pairs report”](#) on page 333.
- New Databases Notification
See [“New Databases Notification report”](#) on page 334.
- Host Configuration Differences
See [“Host Configuration Differences report”](#) on page 335.
- Infrastructure Change Log
See [“Infrastructure Change Log report”](#) on page 336.
- Database Replica Synchronization
See [“Database Replica Synchronization report”](#) on page 337.
- Unscanned Servers/Arrays
See [“Unscanned Servers/Arrays report”](#) on page 338.
- Hosts with a Single SAN I/O Path
See [“Hosts with a Single SAN I/O Path report”](#) on page 339.
- Storage Summary
See [“Storage Summary report”](#) on page 340.
- Scan Troubleshooting
See [“Scan Troubleshooting report”](#) on page 341.
- Database Storage Utilization
See [“Database Storage Utilization report”](#) on page 342.
- VMware Summary
See [“VMware Summary report”](#) on page 343.
- Host HBA Comparison
See [“Host HBA Comparison report”](#) on page 344.
- MS SQL Server Backup Summary
See [“MS SQL Server Backup Summary report”](#) on page 345.

- Replicated Temporary Databases
See [“Replicated Temporary Databases report”](#) on page 346.
- Business Entity Scan Coverage
See [“Business Entity Scan Coverage report”](#) on page 347.
- License Usage
See [“License Usage report”](#) on page 349.
- Data Collection Expansion Package Summary
See [“Data Collection Expansion Package Summary report”](#) on page 350.
- Data Collectors Summary
See [“Data Collectors Summary report”](#) on page 350.
- Installed Software
See [“Installed Software report”](#) on page 351.
- Custom
See [“Custom report”](#) on page 351.

Scan Status report

The Scan Status report, as shown in the following figure, displays the scan status of storage arrays, hosts, clusters, and databases from the last DRA full cycle. A historical view is available for auditing purposes.

Reports

Report: Scan Status

Show scan history from the last No History-

☐ Display databases in one table

[Generate](#)

SCAN REPORT

The following report presents the scan status of storage arrays, hosts, clusters and databases from the last full cycle of DRA. History view is available for auditing purposes.

LAST FULL CYCLE STATUS

Last full cycle start time	Apr 26, 2009 2:23:39 AM
Last full cycle end time	Apr 26, 2009 9:14:43 AM
Collection stage duration	01:49:12
Enrichment stage duration	03:49:38
Gaps stage duration	01:12:14
Full scan status	Completed successfully

SCAN SUMMARY

STORAGE SCANNING STATUS

Scanned successfully: 6/22

Details:

System Event Log report

The System Event Log report summarizes all user actions and system events that are recorded in the system log. This information includes the suppression of tickets, when DRA was available, when users logged in (and with which user role), when tickets were closed, and so on.

You can specify the number of days to be included in the report, as follows:


Reports

Report: System Event Log

Show events notified in the last 2 days

Please click the **Generate** button to generate the report.


The following figure shows a sample System Event Log report:



Ticket Details report

This report provides complete details about each ticket in the report. The top of the Reports page displays filtering options you can use to control the information that is displayed for each ticket.

The following figure shows a sample Ticket Details report:



The report contains the following information for each ticket:


- Header information, including the ticket status, impact categories, and ticket severity. This information comprises the first four lines of the report.
- A full description of the ticket.

Except for the header information, the ticket details are identical to those displayed in the Tickets module.

For more information, see [“Tickets in the Topology module”](#) on page 266.

Optimization report

DRA Optimization reports identify areas of your environment where you may be wasting resources. Storage administrators need this information when they allocate new storage space. The storage volumes identified in DRA Optimization reports may be mistakenly classified as used, when actually they can be reclaimed and reused.



The following conventions apply to all Optimization reports:

- Connected to host

A host is connected to the storage volume, but is not using it. It does not store data on the storage volume. The storage volume appears as a physical volume (PV) on the host, but no volume group (VG) or logical volume (LV) was defined on the physical volume.
- Used by host

A host is connected to the storage volume and is using it. The storage volume appears as a physical volume on the host and a volume group or logical volume is defined on the physical volume. The VG/LV can be used as raw devices or a file system can be defined on the VG/LV.

Structure of the report

This section of the Optimization report summarizes the following report sections:

- Unused Storage Grouping Criteria
See [“Unused storage grouping criteria”](#) on page 325.
- Summary of Storage Capacity per Criterion
See [“Summary of storage capacity per criterion”](#) on page 325.
- Distribution Analysis of Unused Storage by Array
See [“Distribution analysis of unused storage by array”](#) on page 327.
- Distribution Analysis of Unused Storage by Host/Cluster
See [“Distribution analysis of unused storage by host/cluster”](#) on page 328.

Unused storage grouping criteria

This section of the Optimization report specifies the grouping criteria used to describe unused storage volumes. This classification uses the following groups:

- | | |
|---------|--|
| Group 1 | Storage volumes allocated to hosts or clusters that are not used by any host |
| Group2 | Configured replicas of storage volumes belonging to Group 1 |

Summary of storage capacity per criterion

This section of the Optimization report summarizes the Group 1 and Group 2 criteria described in [“Unused storage grouping criteria”](#) on page 325.

This section of the report displays the following:

- Storage volumes allocated to hosts or clusters that are not used by any host. This category provides the following information:
 - The unused storage volumes that are connected to hosts, but that are not used by any host for storage. These storage volumes are not copies/replicas. When a storage administrator needs to allocate new storage space, these storage volumes and their unused replicas might mistakenly be classified as used when they can be reclaimed and used.
 - The storage volumes which are not copies/replicas that are not connected to any host, but which have unused replicas. When a storage administrator needs to allocate new storage space, these storage volumes and their unused replicas might mistakenly be classified as used, when they can actually be reclaimed and used.

Usually, these devices are storage volumes that were removed, but whose replicas were not removed or reclaimed. If the replicas serve no purpose, you should remove the replication configuration and keep a clean, up-to-date configuration.

- The storage volumes which are not copies are not connected to any hosts and do not have copies, but do have LUN maps. When a storage administrator needs to allocate new storage space, these storage volumes might mistakenly be classified as used, when they can actually be reclaimed and used.


Typically LUN map configuration is performed to connect hosts to storage volumes; therefore, when a storage volume has a LUN map(s), you can assume that the storage volume is also connected to a host(s).

- Configured replicas of storage volumes that may be available for reuse. This category provides the following information:
 - Identifies BCV replicas with a Never Established state that are not used by any host. These may be obsolete pairs that were paired some time ago, but never used. When a storage administrator needs to allocate new storage space, these storage volumes might mistakenly be classified as used, when they can actually be reclaimed and used. Details such as the source host, VG/LV and array, as well as the target array and size are shown.

It is also possible that these BCV replicas are part of a new configuration. In this case, they should not be reclaimed.
 - Identifies the replication source target pair in which the target volume is larger than the source volume. This means that there may be unused storage space in the target volume. In such cases, the additional space in the replica volume will never be used. If the replica volume is replaced by a storage volume that is the same size as the source storage volume, the additional space can be reclaimed and used.


Distribution analysis of unused storage by array

This section of the Optimization report summarizes unused storage volumes. It includes the total size of unused storage volumes per array, which is presented as a table and a chart.



Distribution analysis of unused storage by host/cluster

This section of the Optimization report summarizes unused storage volumes. It includes the total size of unused storage volumes per host, which is presented as a table and a chart.



Unreplicated Data on Replicated Hosts report

This report provides details about inconsistencies in host-level replication. It includes information for the following:

- Hosts on which some (but not all) of the volume groups have local replicas
- Hosts on which some (but not all) of the volume groups have remote replicas
- Hosts on which some (but not all) of the volume groups have local replicas on top of remote replicas

The following shows a sample report:

Reports

Report: Un-replicated Data on Replicated Hosts

☐ Show database ☐ Filter file systems smaller than: GB

☐ Display report in one table ☐ Filter file systems named (regexp):

☒ Use checkboxes

Generate

UN-REPLICATED DATA ON REPLICATED HOSTS

The following report presents host level replication inconsistencies. Commonly, SAN volume groups on the same host should have the same protection level in terms of local/remote replicas.

The report searches for:

- Hosts on which some of the volume groups have local replicas, but not all
- Hosts on which some of the volume groups have remote replicas, but not all
- Hosts on which some of the volume groups have local replicas on top of remote replicas, but not all

Hosts with such inconsistencies will be listed in this report, along with their volume groups and replica types.

SITE: CHICAGO

HOST: BLAKED

The following table presents the volume groups of host *blaked* at site *Chicago* and their replica types:

VG/LV	Local copy	Remote copy	Local copy of remote copy	Comments
G:\	☒	☐	☐	
J:\	☐	☐	☐	

SITE: BOSTON

HOST: PRCRM

The following table presents the volume groups of host *prcrm* at site *Boston* and their replica types:

VG/LV	Local copy	Remote copy	Local copy of remote copy	Comments
Archive	☐	☒	☐	

VG/LV Local copy Remote copy Local copy of remote copy Comments

NetApp Filer Replication Summary report

This report provides summarizes the replication configuration for NetApp filers.

Each filer has a table that lists the replication configuration for volumes and Q-Trees, including the following information:

- Whether or not the volume or Q-Tree is remotely replicated
- The number of local point-in-time snapshots
- The number of remote point-in-time snapshots
- The number of configured SAN LUNs
- The hosts accessing the volume or Q-Tree
- The type of volume or Q-Tree (source or replica)
- The filers to which the volume or Q-Tree are replicated

Reports

Report: Filer Replication Summary

Generate

FILER REPLICATION SUMMARY

NETAPP FILER PRODFILER AT BOSTON

4/12 of the volumes are replicated,
All QTrees are not replicated.

The following table presents the volumes of the NetApp filer *prodfiler* at site *Boston* and whether they are replicated:

Volume Name	Has Remote Replica	# Local Snaps	# Remote Snaps	# LUNs	Accessing hosts	Is Replica?	Standby Filers
NFSVol	☐	0	0	1	N/A	No	N/A
db2_data	☒	1	3	1	- Inxw01(uses)	No	- drpfiler
db2_data2	☒	1	3	1	- Inxw01(uses)	No	- drpfiler
db2_data3	☒	1	3	1	- Inxw01(uses)	No	- drpfiler
db2_data4	☒	1	3	1	- Inxw01(uses)	No	- drpfiler
InxArch	☐	0	0	1	- Inxw01(uses)	No	N/A
InxCont	☐	1	0	1	- Inxw01(uses)	No	N/A
InxData	☐	0	0	1	- Inxw01(uses)	No	N/A
InxLog	☐	0	0	1	- Inxw01(uses)	No	N/A
rgdb_Arc_Cont	☐	1	0	2	- rgdb-xp(uses)	No	N/A
rgdb_Data_Logs	☐	1	0	2	- rgdb-xp(uses)	No	N/A
vol0	☐	1	0	0	N/A	No	N/A

The following table presents the Q-Trees of the NetApp filer *prodfiler* at site *Boston* and whether they are replicated:

Q-Tree Name	Has Remote Replica	# Local Snaps	# Remote Snaps	# LUNs	Accessing hosts	Is Replica?	Standby Filers
NFSVol/QNFS	☐	0	0	0	N/A	Yes	N/A

Volume Name/
Q-Tree Name

Has Remote
Replica

Local
Snaps

Remote
Snaps

LUNs

Accessing
hosts

Is Replica?


Standby
Filers

Unsynchronized Remote Replication report

This report lists the unsynchronized remote replications for the devices used by hosts or clusters. The report provides the following information:

- Devices that are used by hosts or clusters
- Devices that have remote replicas configured for synchronous or asynchronous mode
- Devices with unsynchronized remote replicas

The following shows a sample report:



Old Replicas report

The Old Replicas report identifies old replicas that have not been synchronized for a long period of time. Often, many of these replicas are obsolete and can be reclaimed.

Reports

Report: Old Replicas

from: 1 month to: 6 months

Generate

OLD REPLICATIONS

The following report presents replicas which their age matches the given user criteria.
The purpose of this report is to identify replicas which were not synchronized for a long period of time.
Typically, many of these replicas are obsolete and can be reclaimed. Thus, this report is another DRA-tool for storage optimization and identification of saving opportunities.

The report searches for replications between 1 month and 6 months.
Report generation time: Nov 06, 2008 15:49:56

SITE: SANITY PROD

HOST: RGDB-XP

The following table presents the replicas of host **rgdb-xp** at site **Sanity Prod** which their age matches the given user criteria:

Source Data	Source Device	Target Device	Target Host	Target Data	Layout Path	Target Age	Replication State	Size
Database: ORA9	prodfilevolrgdb_CONFFilesLun1	prodfilevolrgdb_CONFFilesLun1.dpfiler00999148951_rgdb_CONFFiles_DR.1316	-	-	NetAppSnapshot	2m 74 22:08:38	N/A	392 MB
Database: ORA9	prodfilevolrgdb_CONFFilesLun1	prodfilevolrgdb_CONFFilesLun1.weekly.0	-	-	NetAppSnapshot	2m 74 22:08:38	N/A	392 MB
Database: ORA9	prodfilevolrgdb_DataFilesLun1	prodfilevolrgdb_DataFilesLun1.dpfiler00999148951_rgdb_DataFiles_DR.77	-	-	NetAppSnapshot	2m 88 19:16:35	N/A	392 MB

Without host

The following table presents the replicas which their age matches the given user criteria:

Source Data	Source Device	Target Device	Target Host	Target Data	Layout Path	Target Age	Replication State	Size
-	prodfilevolrgdb_CONFFiles	dpfilevolrgdb_CONFFiles_DR	-	-	NetAppSnapshot	2m 88 00:42:34	Snapshotmored	600 MB
-	prodfilevolrgdb_DataFiles	dpfilevolrgdb_DataFiles_DR	-	-	NetAppSnapshot	2m 88 21:51:41	Snapshotmored	600 MB
-	prodfilevolrgdb_m1	dpfilevolrgdb_m1_DR	-	-	NetAppSnapshot	3m 236 07:29:07	Snapshotmored	600 MB
-	prodfilevolrgdb_m1	dpfilevolrgdb_m1	-	-	NetAppSnapshot	3m 224 12:08:55	Snapshotmored	100 MB

Source Data

Source Device

Target Device

Target Host

Target Date

Layout Path

Target Age

Replication State

Size

This report shows one table per site, which contains a table for each host. These tables show the storage devices for the given device.

At the top of the page, in the From and To fields, you can specify the date range for the report. If you select - (hyphen) in the To field, the report has no definite end date and includes all relevant replicas from the From date forward.

Report: Old Replicas

from: 1 month to: -

Please click the **Generate** button to generate the report.

3 months

6 months

1 year

-

Standby Pairs report

The Standby Pairs report presents information for standby host pairs. This report distinguishes between the following types of standby pairs:

- Explicit pairs

Standby pairs that are explicitly defined by the user and are monitored by DRA as standby pairs.
- Implied pairs

Standby pairs that are not monitored by DRA as standby pairs, but the replication paths imply that they are standby pairs. The criteria for listing such pairs in the report are as follows:
 - A host (source host) is using SAN storage volumes that are replicated to a remote SAN array.
 - A second host (target host) has access to these replicas.

The following shows a sample report:

Reports

Report Standby Pairs

Generate

STANDBY HOSTS

The following report presents the standby host pairs.
The report separates between two kinds of standby pairs:

1. Explicit pairs
These are standby pairs that were explicitly defined by the user and are monitored by DRA as standby pairs.

2. Implied pairs
These are standby pairs that are not monitored by DRA as standby pairs, but the replication paths imply that these are standby pairs.
The criteria for listing such pairs in the report are:

- 2.1. A host ("Source Host") is using SAN storage volumes which are replicated to a remote SAN array.
- 2.2. A second host ("Target Host") has access to these replicas.

It is highly recommended to review the implied pairs and consider defining them with proper standby relationship.

STANDBY PAIRS SUMMARY

Source Host (Site)	Target Host (Site)	Related Business Entities	Defined?	Details
Inxx01 (Sanity Prod)	usacosol10 (Sanity Prod)	Prod	☒	Press Here
rgtb-xp (Sanity Prod)	w2k3db2v0 (Sanity Prod)	Prod	☒	Press Here

STANDBY PAIRS DETAILS

SOURCE: Inxxv01 (Sanity Prod), TARGET: usacosol10 (Sanity Prod)
Source Host Business Entities: Prod
Target Host Business Entities: Prod

Source Host (Site)

Target Host (Site)

Related Business Entities

Defined?

Details

New Databases Notification report

The New Databases Notification report presents information about the databases discovered by DRA that are not assigned to a scanning scope. DRA does not monitor these databases.

You should configure credentials for the databases identified in this report, and you should assign them to a scanning scope.

Reports

Report: New Databases Notification

From date: ☒ From last full scan

☐ Display databases of replicated servers only

☐ Display report in one table

Generate

NEW DATABASES NOTIFICATION

This report presents all the databases that were discovered by RecoverGuard and that are not assigned to a scanning scope. These databases are not monitored by RecoverGuard, so it is recommended to configure credentials for them and assign them to a scanning scope.

Report generation time: Jul 29, 2012 9:49:40 AM

Number of databases found: 8

DATABASES NOT ASSIGNED TO ANY SCOPE

DATABASES OF HOST CWVQ-2K5E1D

Database Type	Database Service	Status	Roles	Discovery Date
MS SQL Server instance	sqlserv	Up	N/A	04/22/2012
MS SQL Server instance	sqlserv	Up	N/A	04/22/2012
MS SQL Server instance	sqlserv	Up	N/A	04/22/2012

DATABASES OF HOST IGOR

Database Type	Database Service	Status	Roles	Discovery Date
Oracle instance	rgdb	Up	N/A	06/14/2012

DATABASES OF HOST STORAGESIM

Database Type	Database Service	Status	Roles	Discovery Date
Oracle instance	orcl	Down	N/A	06/16/2012
Oracle instance	orcl1	Down	N/A	06/16/2012

DATABASES OF HOST WAMDMORAPRD1

Database Type	Database Service	Status	Roles	Discovery Date
Oracle instance	-sem	Down	GA	04/22/2012

Database Type

Database Service

Status

Roles

Discovery Date

Host Configuration Differences report

The Host Configuration Differences report analyzes potential vulnerabilities based on inconsistencies in the host configurations. For hosts protected by DRA, the report lists the configuration differences between standby hosts and clustered hosts. You can manually select the list of hosts for comparison. If you select a single host and the Show all host configuration check box, the report displays its entire configuration and installed applications. You can filter the compared hosts by business entities. The default is All.

Reports

Report: Host Configuration Differences

☒ Standby aware comparison

☒ Manual hosts selection

02401smesshp,localhost

☒ Show all host configuration (including non-different parameters)

Compare:

☒ Hardware

☒ Software

☒ Operating System

☒ Users and Groups

☒ OS/Kernel Parameters /Limits

Generate

HOST CONFIGURATION DIFFERENCES

This report analyzes potential vulnerabilities based on inconsistencies between host configurations.
The report lists the configuration differences between standby hosts and clustered hosts that are protected by RecoverGuard.

Report generation time: Jul 29, 2012 9:57:42 AM

Hardware

Component	localhost
CPU Speed	3400MHz
CPU Vendor	intel
Memory	2GB
Number of CPUs	1
Server Model	VMware Virtual Platform

Operating System

Component	localhost
Type	ESXi
Version	5.0.0
Version Edition	VMware ESXi 5.0.0 build-280247
Version Level	280247

Component	02401smesshp
Type	ESXi
Version	3.5.0
Version Edition	VMware ESXi 3.5.0 build-453859
Version Level	453859

Infrastructure Change Log report

The Infrastructure Change Log report presents a log of the changes made in the scanned environment for specific items. To specify the items to include in the report, go to the top of the page, and select the relevant Scope and Change Type check box(es).

Reports

Report: Infrastructure Change Log

From date: 2008-12-30 to: From last full scan

Scope: ☒ Host ☒ Database ☒ Storage ☒ Cluster

Change Type: ☒ New ☒ Changed ☒ Removed

Generate

INFRASTRUCTURE CHANGE LOG

Report generation time: Dec 31, 2008 12:29:23 PM

Date	Type	Context	Description
12/30/2008	Changed	Cluster	Attribute ClusterTime value of Veritas Cluster Server VACluster has changed from 1217985668 to 1217996492
12/30/2008	Changed	Cluster	Attribute GlobalCounter value of Veritas Cluster Server VACluster has changed from 3645 to 6351
12/30/2008	Changed	Host	Property used of swap device idevVolGroup00.LogVol01 on host hoxv001 has changed from 295860 to 291084
12/30/2008	Changed	Cluster	Attribute ClusterTime value of Veritas Cluster Server VACluster has changed from 1217996492 to 1217996524
12/30/2008	Changed	Cluster	Attribute GlobalCounter value of Veritas Cluster Server VACluster has changed from 6351 to 6359
12/30/2008	Changed	Cluster	Attribute ClusterTime value of Veritas Cluster Server VACluster has changed from 1218004216 to 1218004200
12/30/2008	Changed	Cluster	Attribute GlobalCounter value of Veritas Cluster Server VACluster has changed from 8282 to 8278
12/30/2008	Changed	Cluster	Attribute ClusterTime value of Veritas Cluster Server VACluster has changed from 1217996524 to 1218004216
12/30/2008	Changed	Cluster	Attribute GlobalCounter value of Veritas Cluster Server VACluster has changed from 6359 to 8282
12/30/2008	Changed	Host	Property used of swap device idevVolGroup00.LogVol01 on host hoxv001 has changed from 291084 to 289328
12/30/2008	Changed	Cluster	Attribute ClusterTime value of Veritas Cluster Server VACluster has changed from 1218004200 to 1218021872
12/30/2008	Changed	Cluster	Attribute GlobalCounter value of Veritas Cluster Server VACluster has changed from 8278 to 12696
12/30/2008	Changed	Host	Property used of swap device idevVolGroup00.LogVol01 on host system1 has changed from 2216 to 3352
12/30/2008	Changed	Cluster	Attribute ClusterTime value of Veritas Cluster Server VACluster has changed from 1218021872 to 1218021880

DateTypeContextDescription

Database Replica Synchronization report

The Database Replica Synchronization report shows the last refresh date for the replicas of each database. This refresh date is computed for replicas of database data files. Each replication category (local, remote, and local copies on remote replicas), shows the latest synchronization date. If there are multiple devices for each database data file, the most recent refresh date is shown.

This report is limited to those databases stored on EMC Symmetrix storage volumes.

Reports

Report: Database Replica Synchronization

Generate

DATABASE REPLICA SYNCHRONIZATION

The following report presents the last refresh date of the replicas for each database.
About the data presented in this report:

The last refresh date is computed for replicas of database data files. For each replication category (local, remote and local copies on remote replicas), the latest synchronization date is presented.
If multiple devices exist per database data files, the most recent refresh data is presented.

Note: The report is currently limited to databases stored on EMC Symmetrix storage volumes.

Report generation time: Dec 20, 2000 4:31:42 AM

Database	Last SRDF Establish Time	Last Snap Activate Time
Oracle BOISDEV	N/A	N/A
Oracle BOISRV	N/A	N/A
Oracle BOISTST	N/A	Dec 21, 2008 4:00:30 AM
Oracle CTM	N/A	N/A
Oracle EMREP	N/A	N/A
Oracle MINT	N/A	N/A
Oracle MINTDEV	N/A	N/A
Oracle MINTTST	N/A	N/A
Oracle MONDB	N/A	N/A
Oracle MONTSTA	N/A	N/A
Oracle MONTSTB	N/A	N/A

Database

Last SRDF Establish Time

Last Snap Activate Time

Unscanned Servers/Arrays report

The Unscanned Servers/Arrays report identifies servers and storage arrays that DRA is not scanning. The list of entities in the report includes the following:

- Storage arrays not scanned by DRA.
- Replicated servers not scanned by DRA. These include servers that access NAS/SAN storage volumes that are either locally or remotely replicated.
- Servers using SAN storage volumes.
- Servers that access replicas of storage volumes used by other servers. These servers are typically standby, DR, backup, or reporting servers.
- Servers that are automatically identified by DRA as clusters, but are not part of the scope.

Note: The report relies on data collected by the EMC Control Center Storage Scope. For up-to-date results, you should execute these probes before you run the report.

The following shows a sample report:

Reports

Report: Un-scanned Servers / Arrays

Generate

UN-SCANNED SERVERS / ARRAYS

The following report presents storage arrays and servers which are currently not scanned by DRA, But potentially should be.

The list includes:

- Storage arrays not scanned by DRA
- Replicated servers not scanned by DRA - Servers which access NAS/SAN storage volumes which are either locally or remotely replicated
- Servers utilizing SAN storage volumes
- Servers which access replicas of storage volumes used by other servers. Such servers are likely to be standby, DR, backup, reporting (etc) servers
- Clusters not scanned by DRA - Servers which share access to SAN storage volumes and are most likely members in a Cluster High Availability configuration

Note: Data collected for servers not scanned by DRA is based mostly on external data sources such as EMC Control Center Storage Scope.

Report generation time: Dec 31, 2009 12:42:28 PM

STORAGE ARRAYS WHICH ARE NOT MONITORED

Type	Name
Symmetrix	000187870349
Symmetrix	000187880222
Symmetrix	000187910448

Type

Name

Hosts with a Single SAN I/O Path report

The Hosts with a Single SAN I/O Path report lets you filter Windows-based hosts. It searches for SAN devices that meet one of the following criteria:

- The number of pseudo paths is one or zero.
- The number of LUN maps is zero.

The following is a sample report:

Reports

Report: Hosts with a Single SAN I/O Path

☐ Include Windows ☒ Include AIX ☒ Include Solaris ☒ Include Linux ☒ Include HP-UX ☒ Include ESX/ESXi ☒ Include VIO

Generate

HOSTS WITH A SINGLE SAN I/O PATH

This report presents devices which have either one LUN map or one SAN I/O path (in both cases, potential single point of failure).
The presented devices are grouped by host.

Following operating systems were excluded from the report:

- Windows

Report generation time: Jul 10, 2012 1:26:52 PM

Host	PV	Databases	VG/LV	SAN device	Number of LUN maps	Number of paths
rac1 (Unknown Site)	/dev/sdy		/dev/sda6	1/dev/sdy	N/A	1
	/dev/sdz		/dev/sda7	1/dev/sdz	N/A	1
	/dev/sdx		/dev/sda5	1/dev/sdx	N/A	1
osa2 (site R)	PV29		vgA	AA10/SV30	N/A	1
osa4 (site R)	PV133		vgA	AA10/SV45	N/A	1
	PV40		vgA	AA10/SV44	N/A	1

Host

PV

Databases

VG/LV

SAN device


No. of LUN maps

No. of paths

Storage Summary report

The Storage Summary report displays a summary of SAN storage per host. It displays the following for each host:

- The number of SAN devices connected and their size.
- The number of SAN devices used and their size.



Scan Troubleshooting report

The Scan Troubleshooting report displays the scanning and scope issues that occurred during the last full cycle.

Reports

Report: Scan Troubleshooting

☐ Long Format

Show scanning issues detected in the last -All-

Generate

SCAN TROUBLESHOOTING

This report presents the scanning issues of hosts, HA clusters, storage arrays and databases.

Report generation time: Jun 27, 2012 1:28:29 PM

HOSTS

Hosts	Issue	Error Message	Suggested Resolution	Severity
storagesim	SymCLI command error: /usr/symcli/bin/symmg-sym -nocache -nocapacity -v -output xml	Return code: 1	Login to the server and check the reason for the command error. Also make sure that the user used for scanning meets the deployment requirements and has all required privileges	ERROR
			Login to the server and check the	

Host

Issue


Error Message

Suggested Resolution

Severity

Database Storage Utilization report

The Database Storage Utilization report computes database space utilization opportunities, based on allocated versus used space.



VMware Summary report

The VMware Summary report gives you an overview of the VMware environment scanned by the product. It describes the storage utilization (datastores and RDM) and virtual machines (including detailed storage allocation and VMware tools status) for each ESX server/ESX cluster.

Reports

Report: VMware Summary

Generate

ESXi: lab-server-1

Storage capacity

- Datastores
 - Total: 165.5GB
 - Used: 138.16GB
 - Free: 27.34GB
- Raw Device Mappings (RDM)
 - Total: 0

Virtual machine summary

- 1 Solaris server
- 5 Windows servers (1 powered off)

Virtual machine details

Type	VM	Power State	Storage Utilization	VMware Tools Status
Solaris	sdl10-1	On	16GB	Not installed
Windows	w2003r1	On	16GB	Installed and running (Version is up-to-date)
Windows	w2003r2 - vcenter	On	16GB	Installed and running (Version is up-to-date)
Windows	w2003r3	On	20GB	Installed and running (Version needs upgrade)
Windows	spn1	On	16GB	Installed and running (Version is up-to-date)
Windows	spn2	On	16GB	Installed and running (Version is up-to-date)
Windows	w2008r1	Off	8GB	Installed (Version is up-to-date)

Type

VM

Power State

Storage Utilization

VMware Tools Status

Host HBA Comparison report

This report lists all hosts that have multiple HBA ports with different parameters. It shows the inconsistent parameters for each HBA installed on these hosts.

Reports

Report: Host HBA Comparison

Generate

HBA CONFIGURATION OF SOLARIS SERVERS											
Site	Cluster	Host	HBA port	Used?	Serial number	Vendor	Driver version	Firmware version	Model	Opt ROM version	Port speed
-	thor8	thor88	/dev/dgic2	No	BG53693687	Emulex	2.3.1p (2008.12.11.10.30)	3.92a2 (C2D3.92A2)	LP9002L	Boot:none Fcode:none	1
-	thor8	thor88	/dev/dgic3	No	BG53693709	Emulex	2.3.1p (2008.12.11.10.30)	3.92a2 (C2D3.92A2)	LP9002L	Boot:none Fcode:none	2
-	thor8	thor88	/dev/dgic5	Yes	-	QLogic	20081115-2.29	3.03.27	QLA2342	BIOS: 1.34;	2
-	thor8	thor88	/dev/dgic6	Yes	-	QLogic	20081115-2.29	3.04.27	QLA2342	BIOS: 1.34;	2

Site

Cluster

Host

HBA port

Used?

Serial number

Vendor

Driver version

Firmware version

Model

Opt ROM version

Port speed

MS SQL Server Backup Summary report

This report lists all MS SQL Server backups of the MS SQL Server databases scanned by the system.

Reports

Report: MS SQL Server Backup Summary

Show backup data for: [Full history](#)

Show backup history from the last: 2 weeks

Generate

Report generation time: Oct 27, 2010 2:29:41 PM

Host	Instance	Database	Type	Size	Date
Windows w2k3ds2v8 at MySite	sq0005mixed	model	database	1.33MB	Sep 19, 2010 9:56:27 AM
Windows w2k3ds2v8 at MySite	sq0005mixed	test123	database	1.33MB	Sep 19, 2010 9:35:53 AM
Windows w2k3ds2v8 at MySite	sq0005mixed	test123	database	1.33MB	Sep 19, 2010 9:58:45 AM
Windows w2k3ds2v8 at MySite	sq0005mixed	test123	database	1.33MB	Sep 19, 2010 10:01:11 AM
Windows w2k3ds2v8 at MySite	sq0005mixed	test123	database	1.33MB	Sep 19, 2010 10:02:20 AM
Windows w2k3ds2v8 at MySite	sq0005mixed	test123	database	1.48MB	Sep 19, 2010 10:04:38 AM
Windows w2k3ds2v8 at MySite	sq0005mixed	test123	database	1.48MB	Sep 19, 2010 11:03:16 AM
Windows w2k3ds2v8 at MySite	sq0005mixed	test123	database	1.48MB	Sep 19, 2010 11:05:41 AM
Windows w2k3ds2v8 at MySite	sq0005mixed	test123	differentialFile	657KB	Sep 19, 2010 9:40:12 AM
Windows w2k3ds2v8 at MySite	sq0005mixed	test123	differentialFile	404KB	Sep 19, 2010 10:06:37 AM
Windows w2k3ds2v8 at MySite	sq0005mixed	test123	fileOrFilegroup	1.39MB	Sep 19, 2010 10:04:52 AM
Windows w2k3ds2v8 at MySite	sq0005mixed	test123	fileOrFilegroup	148KB	Sep 19, 2010 10:05:30 AM
Windows w2k3ds2v8 at MySite	sq0005mixed	test123	fileOrFilegroup	148KB	Sep 19, 2010 10:05:55 AM
Windows w2k3ds2v8 at MySite	sq0005mixed	test123	log	79KB	Sep 19, 2010 9:40:37 AM
Windows w2k3ds2v8 at MySite	sq0005mixed	test123	log	278KB	Sep 19, 2010 10:06:58 AM

HostInstanceDatabaseTypeSizeDate

Replicated Temporary Databases report

This report shows all temporary databases (MS SQL Server and Sybase) that are remotely replicated.

As a temporary database does not contain data that should be protected, these replications present storage and performance optimization opportunities.

The report provides the total replicated storage size and detailed information for each replicated temporary database file, as well as details about whether other databases are stored on the same file system and whether un-replicated drives exist on the same server.

Reports

Report: Replicated Temporary Databases

Generate

QUANTITATIVE SUMMARY OF REPLICATED TEMPORARY DATABASES

Since a temporary database does not contain data that should be protected, these replications present a storage and performance optimization opportunities. Overall, 8.5MB of temporary database files are replicated.

The following table lists the replicated temporary database files as well as whether other databases are stored on the same file system and whether an un-replicated drives exist on the same server.

- If no other database is stored on the same file system, then potentially the replication of this drive can be stopped (considering scan coverage issues).
- If an un-replicated file system exists on the same server, then potentially the temporary database files can be relocated to this disk drive.

Report generation time: Oct 27, 2010 2:28:55 PM

Host	MS SQL server (Last scan)	Type	File	File size	File system	Storage Device	Other databases stored on the same file system	Un-replicated file systems/VGs on this host
w2k3dc2v6	sq005msed	data	C:\Program Files\Microsoft SQL Server\MSSQL_2\MSSQL\DATA\tempdb.mdf	8MB	C:\	Disk0 -> 0428011	sq005msedmodel, sq005msedmsdb, sq005msedmaster, sq005msedtest123	N/A
w2k3dc2v6	sq005msed	log	C:\Program Files\Microsoft SQL Server\MSSQL_2\MSSQL\DATA\templog.ldf	512KB	C:\	Disk0 -> 0428011	sq005msedmodel, sq005msedmsdb, sq005msedmaster, sq005msedtest123	N/A

Host

MS SQL server (Last scan)

Type

File

File size

File system

Storage Device

Other databases stored on the same file system

Un-replicated file systems/VGs on this host

Business Entity Scan Coverage report

This report shows scan-coverage status information for a business entity. It summarizes all information collected for the selected business entities.

Reports

Report: Business Entity Scan Coverage

Selected business entities: TEST

Generate

BUSINESS ENTITY SCAN COVERAGE
The following report provides scan coverage information by business entity.
Selected business entities:

- TEST

The report contains the following sections:

- Scan Summary:
 - Summary chart
 - Scan status distribution by entity type (host, database and storage array)
- Scan Details:
 - Host scan status
 - Database scan status
 - Storage array scan status
- Scan Issues:
 - Host issues
 - Database issues
 - Storage array issues

Notes:

- Some of the sections above may be omitted if no corresponding items have been configured in the scanning scope (for example, if no databases are configured for scanning, the corresponding Details and Issues sections will not appear in the report).
- The Scan Issues section does not present suppressed issues

Report generation time: Aug 08, 2012 3:13:34 PM

Division TEST

SCAN SUMMARY

Entity Type	Scanned	Partially Scanned	Unscanned	Authentication
Hosts	1	0	0	0
Databases	0	0	0	0
Storage Arrays	2	0	0	0

SCAN STATUS DISTRIBUTION BY ENTITY TYPE

HOST SCAN STATUS DISTRIBUTION

Scan Status	Number of hosts
Not verified yet	1

STORAGE SCAN STATUS DISTRIBUTION

Scan Status	Number of storage arrays
Not verified yet	2

SCAN DETAILS

HOST SCAN STATUS

0 out of 1 hosts were scanned successfully.

Type	Host Name	Site	Region	Last Successful Scan	Last Scan Attempt	Scan Status	# Scan Issues
ESX	host4	Unknown Site	N/A	Never	Never	Not verified yet	1

STORAGE SCAN STATUS

0 out of 2 storage arrays were scanned successfully.

Type	Array Name	Site	Region	Last Successful Scan	Last Scan Attempt	Scan Status	# Scan Issues
Symmetrix	SA1	Unknown Site	N/A	Never	Never	Not verified yet	0
Symmetrix	SA2	Unknown Site	N/A	Never	Never	Not verified yet	0

SCANNING ISSUES

HOST ISSUES

Issue Type	(#) Hosts
Site definition missing	(1) host4

You can generate this report for one or more business entities. For example, click the [\[all\]](#) link to open the following window in which you select the business entities of interest for the report.

Business Entities

Search:

Reg Exp

Name

Customers

ERP

CRM

Energy

SOA

Page

1 of 1

Displaying 1 - 5 of 5

Show

20

Items

OK

Cancel

License Usage report

This report shows license usage information, such as details about the number of licenses allocated and their usage, by host.

Reports

 Report: License Usage

Generate

LICENSE USAGE SUMMARY

The following report presents license usage information, in the following sections:

- List of licensed physical (non-virtualized) hosts and Unix LPARs (each host / LPAR = 1 license)
- List of licensed Virtual Infrastructure ("VI") hosts (ESX, ESXi and Solaris Global Zones) running VMs or Private Zones in scope (each VI host = 1 license)

Notes:

- VMware VMs and Solaris Private Zones ("Guests") are not counted towards the license usage.
- Guests listed in this report, appear only to designate which VI hosts are configured to run them. In this context, the term "configured" denotes both the actual VI hosts running the guests and those VI hosts to which the guest can fail-over to.

Report generation time: Aug 08, 2012 3:31:37 PM

LICENSE USAGE SUMMARY BY OPERATING SYSTEM TYPE

- Windows - 2
- AIX - 2

License usage: 4 out of 20 licenses.

PHYSICAL HOSTS AND UNIX LPARS - 4 LICENCES

Operating System	Host Name	Business Entities	Site
AIX	aix_host1	-	Chicago
AIX	aix_host2	-	Chicago
Windows	windows_host1	Gold	Boston
Windows	windows_host2	Gold	Boston

Operating System

Host Name

Business Entity

Site

Installed Software report

This report shows which hosts have a particular product name.

Reports

 Report: Installed Software

Filter product name by Regular Expression: yp.*

Generate

INSTALLED SOFTWARE

THE FOLLOWING TABLE LISTS THE HOSTS AND INSTALLED SOFTWARE VERSION

Host	Product	Version
Linux storagesim	yp-tools	2.8-7
Linux storagesim	ypbind	1.17.2-3

Host

Product

Version

Use the filter area to filter the report to show details for a specific software application.

 Report: Installed Software

Filter product name by Regular Expression: yp.*

Custom report

A custom report is created as an expansion package. Any custom reports that you define are shown in the Custom Report drop-down list. For more details, see [“Defining reporting packages”](#) on page 208.

Reports

 Report: Custom Report

Custom Report: Hosts not assigned to any Business Entity

Generate

Please click the **Generate** button to generate the report.

Using the Comparison module

This chapter includes the following topics:

- [“Overview”](#) on page 353 introduces the Comparison module and describes its purpose.
- [“Configuring Worksheets”](#) on page 355 describes how to define and manage worksheets.
- [“Configuring Comparison Groups”](#) on page 356 describes how to define and manage comparison groups.
- [“Comparing Host Configurations”](#) on page 366 describes how to compare the configurations of various hosts.


Overview

The Comparison module helps you identify the host configuration drifts hiding in your IT. Such drifts often fail cluster failover processes, and reduce the availability of your organization’s applications. In this module, you can create *comparison groups* that include hosts, clusters, or business entities, and easily track configuration differences between them.

The Comparison module uses a *worksheet*, which is a logical container of comparison groups. In addition, it contains all suppressions and difference monitoring information.

You begin by creating a worksheet, and then by creating comparison groups. Once defined, comparison groups can be assigned to a worksheet.

To access the Comparison module, click the Comparison tab.




You can perform the following operations from the Comparison Panel page:

Load Worksheet



Opens the Worksheet Selection window, in which you select the worksheet to load



Create New Worksheet



Adds a new worksheet

Save Worksheet



Saves a worksheet

Clone Worksheet



Clones the current worksheet by closing the current worksheet, creating a copy of it, and then opening it

Close Worksheet



Closes a worksheet

Delete Host Configuration

Differences Tickets



Executes the rule engine, which initiates the detection process. Tickets are opened, based on your selection. A message is displayed once the process finishes. Detection results are shown in the Tickets module.

Delete Worksheet




Deletes a worksheet

Configuring Worksheets

A worksheet is a logical container of comparison groups. It also contains all suppressions and difference monitoring information. Worksheets are defined and saved at the user level, which means that each user has his/her own worksheet.

To create a new worksheet

- 1
- Click **Create New Worksheet** (). The Comparison Panel page displays as shown below:



- 2
- In the **Name** field, enter a name for the worksheet.

- 3
- In the **Privacy** drop-down list, specify the privacy level for the worksheet. The privacy level determines which users can access the worksheet. You can select from the following options:
- Private

Only the worksheet owner can view or modify the worksheet.
- Protected


Other users can view the worksheet in read-only mode, and can copy it, but not use it.
- Public

Other users can view and modify the worksheet.
- 4
- Select the **Standby aware comparison** check box if you want to perform a smart comparison between the primary and the standby. When selecting this option, the primary and its standby are considered as comparable items.
- 5
- Click **Save Worksheet** ().

Configuring Comparison Groups

A comparison group is a dynamic group of hosts that you want to compare. Four types of comparison groups are available. Each group type behaves differently in terms of the group scope and comparison functionality. The following group types are available:

Hosts	A static list of hosts.
Clusters	A list of clusters. The hosts compared are cluster members. If you add a cluster node to the selected cluster, it is automatically added to the group's scope.
Business Entities	A list of business entities. The hosts compared are the hosts that are members of the selected business entity. When configuring a group of this type, you should specify the operating system you want to compare.
Golden Copy	This group enables you to select a host as a <i>Golden Copy</i> , and compare a list of hosts to its configuration.



The Groups pane lists each defined comparison group as a separate primary node. Each subnode under a primary node represents the members belonging to that comparison group, such as hosts, cluster nodes, business entities, and so on.


You can perform the following operations from the Groups pane:

Create New Comparison Group  Adds a new comparison group


Automatically Create New
Comparison Group Based on HA

Clusters 

Adds a new comparison group, based on specific clusters. You select the clusters for the group in the following window.



Click the [\[none\]](#) link next to **Clusters** to open the Clusters window in which you select the clusters for the group.




After making your selection, click **Create**, which automatically creates a new group for each cluster selected.

Automatically Create New
Comparison Groups Based on
Standby Pairs



Adds a new comparison group, based on standby pairs. You select the pairs for the group in the following window.



After making your selection, click **Create**, which automatically creates a new group for each standby pair selected.

Collapse All Groups



Collapses all group subnodes in the Groups pane tree

Expand All Groups




Expands all group subnodes in the Groups pane tree

Compare All Groups in Worksheet



Compares all comparison groups defined in the Groups pane simultaneously. A separate results tab is shown in the Comparison Panel for each group that was compared.

Group results tab



Edit Group



Modifies a comparison group


Delete Group 

Deletes a comparison group

You can click the name of a group in the Groups pane to open it for comparison, as described in [“Comparing Host Configurations”](#) on page 366.

To create a comparison group

- 1
- Click **Create New Comparison Group** (). The following window opens:




- 2
- In the **Name** field, enter a name for the comparison group.
- 3
- In the **Type** drop-down list, select the comparison group type. You can select from the following options:
- Hosts
 - Clusters
 - Business Entities
 - Golden Copy


Note: The rest of the windows in the wizard vary, depending on the group type that you select.

- 4
- Click **Next**.

- For the host group type, the following window opens:




Click the [\[none\]](#) link to open the following window, in which you select the hosts for the comparison group.




Select the hosts for the comparison group and then click **OK**.

- For the clusters group type, the following window opens:



Click the [\[none\]](#) link next to **Selected Clusters** to select the cluster nodes for the comparison group. Click the link next to **Selected Hosts** to select the hosts for the group. After making your selections, click **OK**.

- For the business entities group type, the following window opens:




Click the [\[none\]](#) link next to **Selected Business Entities** to select the business entities for the comparison group.

In the **Platform** drop-down list, select the operating system you want to compare.


After making your selections, click **OK**.

- For the Golden Copy group type, the following window opens:



In this context, a Golden Copy represents the preferred configuration. Use this option to compare a list of hosts to the Golden Copy configuration.

Click the [\[none\]](#) link next to **Golden Copy** to open the following window:




Select the host to use as the Golden Copy, and then click **OK**. Only one host can be designated as the Golden Copy.

Click the [\[none\]](#) link to select the hosts to be compared to the Golden Copy.

After making your selections, click **OK**.

- 5 Click **Finish**. The comparison group is added to the list of groups in the Groups pane.



- 6 It is recommended to save the worksheet at this point. To do so, click **Save**

Worksheet (.

Note: You can automatically create comparison groups based on discovered clusters by clicking **Automatically create Comparison Groups based on HA**


Clusters (). See page 358 for more details.

Comparing Host Configurations

After you have created the requisite worksheets and comparison groups, use the options described in this section to compare host configurations.


To compare host configurations

- 1
- In the Groups pane, click the name of the comparison group that you want to compare. The following window opens:



A new tab is opened for the selected group in the Comparison Panel on the right.

- 2
- Select the categories that you want to compare by checking the respective check boxes, and then click **Compare!**.



The following options are always available for comparisons:

- Hardware
- Software
- Operating System
- Users and Groups
- OS Kernel Parameters / Limits


In addition, the list of comparison options includes a check box for each host comparison package that is currently defined.


Results are displayed in the Comparison Panel. Each check box option that you selected for the comparison typically has one results tab. However, some options, such as Software, may have more than one. For example, there may be a results tab for each distinct software product and patch.

The tables in the results tab contain all the differences between the two comparators in the comparison group. The following information is displayed for each row in a results tab:


Component	The name of the item being compared
Difference Impact	The impact of the differences found for the comparison
<Host Name 1>	The name of the first host for the comparison
<Host Name 2>	The name of the second host for the comparison

You can perform the following operations in the Comparison Panel:


☐ Show suppressed values	When checked, hides rows containing ignored values so that they do not display. When not checked, rows with ignored values are grayed out.
PDF 	Exports the comparison results in PDF format.
Word 	Exports the comparison results in RTF format.
Excel 	Exports the comparison results in Excel format.
	Initiates the comparison process, and displays comparison results differences in tabular format in the Comparison Panel page.




Enables you to perform the following actions:




- **Ignore selected values:** Ignores the difference, meaning that the difference does not identify a gap.
- **Ignore selected values in the entire worksheet:** Ignores the difference for the detected rows in all other groups of the same worksheet.
- **Stop ignoring selected values:** Stops ignoring selected values and resets them. Such rows then show as gray.




Enables you to perform the following actions:




- **Show results in worksheet level:** Compares the selected component for all hosts in the worksheet from all groups.
- **Show statistics for all hosts:** Shows the value distribution for the selected component per operating system, meaning for all hosts in the system and not just those in groups.




Use the **Monitor** button to select components and then monitor their changes in the Tickets module. You can perform the following actions for this purpose:



- **Monitor changes in Tickets module:** Monitors changes for the selected components in the Tickets module. Components selected for monitoring display  next to their name. After you save the worksheet, these changes automatically open tickets, which can be viewed in the Tickets module.



- **Stop monitoring in Tickets module:** Stops the monitoring of selected components in the Tickets module.



Opens a window to define the impact. A default impact value is available, which is configurable.

Note: Some of the operations above are not available for the Golden Copy group type.

- 3
- It is recommended to save the worksheet after you change the group configuration. To do so, click **Save Worksheet** ().

Viewing DRA SLA information

This chapter includes the following topics:


- [“Overview”](#) on page 371 introduces the SLA module and describes its purpose.
- [“Using the SLA window”](#) on page 372 describes the organization and usage of the DRA SLA window.

Overview

The SLA module identifies business entities that have a problem with one of their resources. It lets you view data for a selected business entity, including its resources and subresources (volume groups, databases, file systems, and so on). You can view this information in the SLA tab of the main window.

Using the SLA window

To access the SLA page, click the SLA tab. The following screen is displayed:



The SLA page has several areas. The tree on the left lists all business entities defined in the system. When you select a business entity in this tree, all of its associated resources are shown in a tree structure in the right pane. A separate row is shown for each resource, which contains the following information about that resource:

SLA information					
Number of copies		RPO		Retention	
Local	Remote	Local	Remote	Local	Remote
0 - 6	0 - 2	29d 12:02:35 - 1y 5m 10d 10:05:38	N/A	29d 12:02:35 - 1y 5m 10d 10:05:38	N/A

- **Number of copies (range):** The number of local copies and the number of remote copies for the resource, including all of its subresources (if any). If the resource contains subresources, these values are expressed as a range. The first number indicates the fewest number of replicas for one of its subresources, and the second number indicates the greatest number of replicas for one of its subresources. For example, the 0-6 range for Local copies shown in the following sample window means that 0 is the fewest number of replicas for one subresource and 6 is the maximum number of replicas for another subresource that comprise this resource.

Number of copies	
Local	Remote
0 - 6	0 - 2

- RPO range: The timestamp of the replica’s Recovery Point Objective (RPO). This information is provided for both local and remote copies. The information is expressed as a range, indicating both the oldest and latest timestamps for the subresources comprising this resource.

RPO	
Local	Remote
29d 12:02:35 - 1y 5m 10d 10:05:38	N/A

- Retention range: The timestamp of both the oldest and latest complete copies available. This information is provided for both local and remote copies. The information is expressed as a range, indicating both the oldest and latest timestamps for the subresources comprising this resource.

Retention	
Local	Remote
29d 12:02:35 - 1y 5m 10d 10:05:38	N/A

Each resource can be comprised of subresources. When you expand the resource tree in the right pane of the SLA page to its lowest level, details for the subresources are displayed.



This level of detail shows a timeline for each subresource. Time on this timeline is not proportionally distributed. Distribution across the timeline is based solely on an equal distribution of the number of replicas for the subresource and not on time.

Each subresource contains one or more replicas. Each replica is represented by a colored icon.




The icon color indicates whether there is a problem on the replica, as follows:


- Green  The replica is OK.
- Yellow  The replica has a non-critical problem.
- Red  The replica has a critical problem.
- Gray  The item is only partially replicated.

An L or an R in the center of the icon indicates a local or remote replica, respectively. In addition, **L(0) R(13)** at the beginning of a subresource row indicates the number of local (L) and remote (R) replicas for the subresource. Thus, L(0) R(13) indicates that this subresource has 0 local replicas and 13 remote replicas, as shown in the figure above.

When you hover over one of the colored replica icons, information is displayed for that replica, such as its data age, layout path, state, RAID type, and which host sees the replica. This information is the same for all replica icons, except that a yellow icon indicates a less severe error than a red icon. A green replica icon indicates that there are no errors.



The Show status for field at the top of the right pane lets you select the information you want to display. When you select a day other than today, it displays information that was correct for that day.



Below each icon is a timestamp and the state of the replica. The timestamp indicates when the replica was copied.

L(1) R(1)		00:00:00 Copied		N/A Split
L(1) R(1)		8m 4d 20:47:21 Copied		N/A Split
L(1) R(1)		29d 12:02:35 Copied		N/A Split
L(1) R(1)		00:00:00 Copied		N/A Split
L(1) R(1)		1y 5m 10d 10:05:38 Copied		N/A Split
L(1) R(1)		30d 01:43:38 Copied		N/A Split
L(0) R(2)		N/A Split		N/A Split

Glossary

Any Type of Connection	Any relationship between two objects. For example, a physical volume is stored on a storage volume.
BC	Business Continuity.
BCV	Business Continuous Volumes, previously known as Business Suite Continuous Volumes. An EMC Symmetrix term for a special mirror volume managed by an array. BCVs are created and maintained by the TimeFinder software module of EMC.
Business entity	A collection of IT resources that provide application services to a defined business process. Business entities can comprise servers, storage resources, applications, and other business entities in a hierarchical way.
CG	See Symmetrix CG.
CIFS Share	A Network Share implemented using the Common Internet File System (CIFS) protocol. The protocol is usually implemented as a service by hosts (most commonly by Windows hosts or NAS devices, but it may also be implemented by UNIX hosts through a special service, such as samba).
CLARiiON	Historically, a Data General product brand acquired by EMC. The CLARiiON family includes mid-range storage arrays.
Clone	A type of storage volume replica representing a point-in-time state of its source storage volume. Several vendors, sometimes with slight interpretations, use the term.
DASPV	A DRA term for a PV that represents a disk device that is directly connected to a host, using a disk bus protocol, such as SCSI, IDE, FC-AL, SAS, and so on. In DRA, a DAS PV may represent either a physical disk device on a JBOD or a logical disk device on a directly-attached storage RAID.
DG	See Symmetrix DG.

Domain	■ A collection of servers, network devices, network services, users, user roles, and so on, that is administered collectively as a single entity. ■ A repository containing information regarding elements described above; for example, Microsoft AD.
DP	Data Protection.
DPRA	Data Protection Risk Assessment.
ECC	EMC Control Center.
EMC SYMCLI	EMC Symmetrix Command Line Interface.
FC-AL	Fiber Channel Arbitrated Loop.
File system	■ A method of organizing files on a disk, usually in a hierarchical directory. ■ A disk area managed by an operating system formatted according to the method in (a) above. See also local file system and network file system.
Gap	A definition of a particular problem type.
HDS HiCommand	Hitachi Storage Command Suite. Used to manage and retrieve information about Hitachi Data Systems (HDS) arrays.
Host	A computing device or server, physical, partitioned, or logically virtualized, running an instance of an operating system.
ISCSI	Internet SCSI. An IP-based standard for linking data storage devices over a network and transferring data by carrying SCSI commands over IP networks.
IT	Information Technology.
JBOD	Just a Bunch of Disks. A storage array with no intelligence, add-on software, and so on. A JBOD is a convenient way to connect multiple disks to hosts, especially when there is no physical space for disk installation in the server chassis itself. Typically, JBODs can be occupied by hot-pluggable disks using special disk casing, providing the disks with a common power source (possibly redundant), and a common connection interface (possibly redundant) to hosts or SAN networks.
Layout	A short representation of a replication path. For example, the layout of a storage volume that is an SRDF replica of a BCV replica is “BCV, SRDF.”

Local file system	A file system managed by a host, which is stored on storage resources directly attached to and managed by that host.
LUN	Logical Unit. Originally a SCSI term that was adapted for use in FC-AL and iSCSI. A LUN is the smallest addressable entity in a SCSI or Fiber-channel bus. In simple installations, each LUN usually stands for a whole disk, tape drive, CD-ROM drive, or any other single unit of a physical media. Devices such as RAIDs or SAN switches can map physical resources in various ways, so that a LUN may represent logical units smaller or larger than a single device (and many combinations thereof).
LUN map	A LUN map maps a storage volume to a specific port. Without a LUN map, hosts cannot see a storage volume.
LV	Logical volume. A UNIX term denoting a data storing portion of a volume or disk. An LV can be implemented as either a partition of a single disk or as a partition of a volume group (see also VG), if a logical volume manager is being used (see also LVM).
LVM	Logical volume manager. A software module that runs on a host, as a low-level service. An LVM can manage multiple storage volumes accessible by the host, using them to provide other entities in the host's operating system with logical volumes (LVs). Mapping between physical and logical volumes can be of various complexities; that is, an LV may be physically stored on a portion of one physical volume, on multiple portions of multiple physical volumes, on an entire volume, on a set of volumes, and so on.
NetApp	Network appliance. A product of Network Appliance, Inc.
NetApp Aggregate	A NetApp term denoting a manageable unit of RAID protected storage.
NetApp Filer	A NetApp term for one of their storage arrays.
NetApp LUN	A special file stored on a NetApp storage array that can be represented as a LUN or physical disk to hosts. A NetApp LUN can be made accessible to hosts either using SAN or using iSCSI.
NetApp Q-Tree	A NetApp Filer object comprising a special subdirectory of the root of a volume that acts as a virtual subvolume with special attributes.

NetApp Volume	A NetApp term denoting a file system that can store data. Volumes appear in two forms: Traditional Volumes and Flexible Volumes (the latter supported from ONTAP 7).
Network file system	A file system accessed by a host through a network protocol (such as CIFS or NFS), using an appropriate client component, which is implemented and managed by a remote host providing a corresponding server-side service. Typically, a server can serve more than one client.
Network Share	A mechanism for sharing data for remote accessing by a different host. Sharing is done by hosts (Windows, NAS and UNIX, which use the term export instead of share) through publishing a directory.
NFS Export/NFS Share	A Network Share implemented through the use of the Network File System (NFS) protocol. Usually, the protocol is implemented as a service by hosts (most commonly by UNIX hosts or NAS devices, but may also be implemented by Windows hosts through a special service).
ODBC	Open Database Connectivity. An open standard API for accessing a database.
Oracle Archive Destination	In the context of a specific Oracle Instance, one of the Instance's Archive Destinations.
Oracle Archive Destinations	In the context of a specific Oracle Instance, a container for all the Instance's Archive Destinations, such as the locations that store archive files.
Oracle Control File	In the context of a specific Oracle Instance, one of the Instance's Control Files.
Oracle Control Files	In the context of a specific Oracle Instance, a container for all the Instance's Control Files.
Oracle Data File	In the context of a specific Oracle Instance, one of the Instance's Data Files.
Oracle Data Files	In the context of a specific Oracle Instance, a container for all the Instance's Data Files.
Oracle Instance	An active executable entity, running on a host and implementing an Oracle Database service.
Oracle Log File	In the context of a specific Oracle Instance, one of the Instance's Log Files.
Oracle Log Files	In the context of a specific Oracle Instance, a container for all the Instance's Log Files.

OS	Operating system.
PE	Path Element.
PV	Physical volume. A UNIX term denoting an operating system representation of a physical disk as seen by a host.
R1	The source volume of an SRDF replication.
R2	The replicated volume of an SRDF replication.
RAID	Redundant Array of Independent Disks. ■ A term describing a mechanism of grouping several storage devices into a logical group which provides improvements in one or more of the following aspects: capacity, redundancy or performance, or ■ A term describing a hardware device that can store multiple disks and provide intelligent services of storage management and virtualization through the use of mechanisms such as described above.
RDF	See SRDF.
RPO	Recovery Point Objective.
RTO	Recovery Time Objective.
SAN	A collection of specific communication network and storage devices that provides storage services to hosts: ■ SAN communication devices Comprised of Fiber Channel (FC) switches or hubs connected in islands called Fabrics. More than one fabric may be used in an organization. All switches and hubs in a fabric form a network with physical FC ports to which storage devices and hosts can be connected. ■ SAN services A combination of management software and hardware tools enables configuring a fabric to form subnets called zones and to define accessibility rules between storage volumes and hosts (masking). ■ AN storage devices A collection of storage devices connected to a fabric, such as disks, tapes, storage arrays, JBODs, and so on.
SAN I/O Path	A DRA term for a physical volume that represents a logical disk device made addressable by a Storage Area Network. See also SAN.

SAS	Serial Attached SCSI. See also SCSI.
SCSI	Small Computer Systems Interconnect.
Site	An IT site.
SANPV	A DRA term for a PV that represents a logical disk device made addressable by a Storage Area Network. See also SAN.
SLA	Service Level Agreement.
Snapshot	A type of storage volume replica representing a point-in-time view of its source storage volume. Any changes that occur to the source storage volume after the snapshot begins do not affect the replica.
SRDF	Symmetrix Remote Data Facility. An EMC Symmetrix term for a special replication technology that can copy storage devices from one array to another in various modes.
SRM	Storage Resource Management.
SSH	Secure Shell (SSH) or Secure Socket Shell. A UNIX-based command interface and protocol for securely accessing a remote computer.
Storage Array	A method for storing information on multiple devices.
SV	Storage Volume. A DRA term denoting a physical or logical storage entity that can be connected to a host and recognized as a disk.
SYMCLI	Symmetrix Command Line Interface.
SymmCG	See Symmetrix CG.
SymmDG	See Symmetrix DG.
SymmRDFGroup	See Symmetrix RDF Group.
Symmetrix	An EMC product name used for high-end storage arrays. Newer models are called Symmetrix DMX.
Symmetrix CG	EMC Symmetrix Composite Group. A Symmetrix CG is very similar to a Symmetrix DG, except it can contain devices that span multiple Symmetrix arrays and multiple RDF groups.

Symmetrix Device	An EMC term denoting a logical segment managed by an EMC Symmetrix Storage Array that can be represented as a LUN or physical disk to hosts.
Symmetrix DG	EMC Symmetrix Device Group. Symmetrix device groups are user-defined objects that collectively manage groups of Symmetrix devices.
Symmetrix DMX	An EMC product name used for high-end storage arrays. Gradually replacing Symmetrix arrays.
Symmetrix Meta Device	An EMC term denoting a metadvice in an EMC Symmetrix Storage Array. Metadevices let individual devices be concatenated to create larger devices.
Symmetrix Meta Head	An EMC term denoting a device in an EMC Symmetrix Storage Array with a Meta device head.
Symmetrix Meta Member	An EMC term denoting a device in an EMC Symmetrix Storage Array with a Meta device member.
Symmetrix RDF Group	EMC Symmetrix Remote Data Facility Group. A logical representation that defines a relationship between two Symmetrix Storage Arrays that are logically connected with EMC Symmetrix SRDF.
Telnet	A user command and an underlying TCP/IP protocol for accessing remote computers.
Ticket	A specific instance of a particular gap problem found in your environment.
TimeFinder	The name of an EMC product suite that creates various types of copies of storage volumes in a single array. The suite includes: TimeFinder /BCV, TimeFinder /Clone and TimeFinder /Snap. See also Snapshot and Clone.
TR	Traversal Rule.
Unknown Array Device	A DRA term denoting a device of an unknown Storage Array.
Unknown Storage Array	A DRA term for representing a Storage Array of an unknown vendor.
VCS	Veritas Cluster Server

VG	Volume group. A UNIX term denoting a collection of one or more physical volumes (see also PV) from which space can be allocated to one or more logical volumes. See also LV.
WMI	Windows Management Instrumentation. A set of Microsoft specifications for consolidating the management of devices and applications in a network from a Windows computing system.

Index

A

- Accessing
 - Topology module 294
- Active Directory system properties 236
- Adding
 - databases 102
 - hosts 78
 - new email report schedule 145
 - note to ticket 270
 - policy 148
 - rating to ticket 271
 - site from CSV file 94
 - standby definitions from CSV file 122
- Additional collection system properties 237
- Advanced settings
 - configuring 217
- Aggregation system properties 237
- Analyzer 253
- Area
 - Business Summary 67
 - Scan Coverage Details 65
 - Scan Coverage Summary 64
 - Ticket Summary 66
 - Ticket Tabs 66
- Automatic association of databases to groups and policies
 - configuring 110
- Automatic import system properties 238
- Automatic opening/closing
 - tickets 29

B

- Benefits
 - DRA 18
- Business entities
 - configuring 158

- Business entity roles
 - defining 170
- Business Entity Scan Coverage
 - report 347
- Business entity types
 - defining 168
- Business Summary area 67
- C**
- CLARiiON arrays 83
- Closing
 - tickets 29
- Collapse
 - selected entity 307
- Collected data 48
- Collection system properties 239
- Collection timeouts system properties 241
- Collector
 - installing 43
- Collectors
 - configuring 115
- Command line interface system properties 243
- Commands
 - executed during data collection 47
- Comparing
 - host configurations 366
- Comparison
 - module 56
- Comparison group
 - configuring 356
- Comparison module 353
- Comparison system properties 243
- Components
 - DRA 19
- Computer preparations
 - installation 32

Configuration

- module 57

Configuration Wizard 71**Configuring**

- advanced settings 217
- automatic association of databases to groups and policies 110
- business entities 158
- collectors 115
- comparison groups 356
- credentials 151
- data collection packages 180
- distributed collection 115
- DRA 69
- email settings 147
- EMC ECC 72
- EMC ECC and HDS HiCommand, management consoles 72
- EMC NaviCLI proxies 82
- EMC SYMCLI proxies 81
- expansion packages 179
- external ticketing 223
- gap management settings 229
- HA clusters 126
- HDS HiCommand 72
- Hitachi HiCommand proxies 86
- host comparison packages 193
- IBM DSCli proxies 85
- IBM HMC consoles 73
- IBM HMC management console 72
- IBM XCLI proxies 85
- licenses 231
- NetApp 86
- policies 148
- proxy 156
- reporting packages 208
- risk detection packages 198
- SLA policies 172
- standby host definitions 119
- standby policies 123
- Sudo 157
- SVC proxies 86
- system properties 235

- ticket labels 228

- user management settings 217

- user roles 220

- VMware vCenter 73

- VMware vCenter management console 72

- worksheets 355

Configuring databases

- database groups and policies 105

- Configuring EMC SYMCLI, EMC NaviCLI, IBM DSCli, IBM XCLI, IBM SVC, Hitachi HiCommand, and NetApp proxies

- storage discovery 81

Configuring hosts

- hosts groups & policies 95

Configuring sites

- site definitions 91

Connection

- display methods 299

Connections

- tab 313

Connector

- configuring 224

Considerations

- security 46

Controlling

- table display layouts 61

Controlling window

- display layout 60

Credentials 48

- configuring 151

Custom report 351**Cyber-Ark configuration**

- system properties 244

D**Dashboard 51**

- monitoring 28

- opening Topology module 295

- system properties 245

- using 63

- viewing information 64

Data analysis task

- scheduling 137

- Data collection 47
 - executed commands 47
- Data Collection Expansion Package Summary
 - report 350
- Data collection packages 180
- Data Collectors Summary
 - report 350
- Data Dependency engine 19
- Database
 - groups, managing 108
 - policies, managing 107
- Database Replica Synchronization report 337
- Database Storage Utilization
 - report 342
- Databases
 - adding 102
 - automatic association to groups and policies 110
 - Oracle, security 46
- Defining
 - business entity roles 170
 - business entity types 168
- Definitions
 - configuring standby host 119
- Details tabs
 - topology 283
- Detection
 - gap and risk 20
- Disaster 1
- Discovery
 - storage, configuring EMC SYMCLI, EMC NaviCLI, IBM DSCSI, IBM XCLI, Hitachi HiCommand, and NetApp proxies 81
- Display
 - connected entities 309
 - replications 311
 - storage 310
- Display connected entities 309
- Display layout
 - window, controlling 60
- Display methods
 - connection 299
- Display replications 311
- Display storage 310
- Displayed content
 - filter 304
- Displaying item details
 - Topology module 311
- Distributed collection
 - configuring 115
 - system properties 246
- DRA
 - benefits 18
 - components 19
 - configuring 25, 69
 - getting started 21
 - installing 22, 31, 32
 - introducing 17
 - logging in 42
 - main window 50
 - object model 295
 - overview 17
 - probes 70
 - reporting 317
 - security 46
 - uninstalling 43
- DRA interface
 - quick tour 49
- DRA reporting
 - overview 317
- DRA SLA
 - viewing information 371
- E**
- Editing
 - policy 151
- Email report schedule
 - adding new 145
- Email settings
 - configuring 147
- Email system properties 246
- EMC ECC
 - configuring 72
 - configuring, management consoles 72

- EMC NaviCLI
 - configuring 82
- EMC NaviCLI proxies
 - configuring, storage discovery 81
- EMC SYMCLI proxies
 - configuring 81
 - configuring, storage discovery 81
- Engine
 - Data Dependency 19
 - Gap Analysis 20
 - IT Discovery and Scanning 19
- Entities
 - business, configuring 158
 - hide 305
- Entity
 - collapse 307
 - display connected 309
 - expand 306
- Entity roles
 - business, defining 170
- Entity types
 - business, defining 168
- Executed commands
 - data collection 47
- Expand
 - entity 306
 - selected entities 308
- Expansion packages
 - data collection 180
 - host comparison 193
 - reporting 208
 - risk detection 198
- Exporting
 - reports 319
 - tickets from the Tickets module 273
 - topology 289
- External ticketing
 - configuring 223
 - configuring a connector 224
- External ticketing system integration
 - system properties 247

F

- Filter

- displayed content 304
- Full cycle
 - system properties 247

G

- Gap Analysis engine 20
- Gap and risk detection 20
- Gap management settings
 - configuring 229
- Gap rules
 - system properties 247
- General system properties 247
- Generating
 - ticket unique ID 272
- Generator
 - report 20
- Getting started
 - DRA 21
- Glossary 377
- Graphical view
 - item display methods 298
 - topology 281
- Groups
 - database, configuring 105
 - database, managing 108
 - host, configuring 95
 - host, managing 101

H

- HA clusters
 - configuring 126
- Handling
 - open tickets 30
- Hardening
 - server 47
- HDS HiCommand
 - configuring 72
 - configuring, management consoles 72
- HiCommand proxies
 - configuring 86
 - configuring, storage discovery 81
- Hide
 - selected entities 305
- History
 - tab 314

- Hitachi HiCommand proxies
 - configuring, storage discovery 81
- Host
 - comparing configurations 366
 - groups, managing 101
- Host comparison packages 193
- Host Configuration Differences report 335
- Host HBA Comparison
 - report 344
- Host policies
 - managing 100
- Hosts
 - adding 78
- Hosts with a Single SAN I/O Path report 339
- Housekeeping system properties 248
- I**
- IBM DCLI proxies
 - configuring 85
 - configuring, storage discovery 81
- IBM HMC
 - configuring 73
- IBM SVC proxies
 - configuring, storage discovery 81
- IBM XCLI proxies
 - configuring 85
 - configuring, storage discovery 81
- ID
 - unique ticket, generating 272
- Information
 - Dashboard, viewing 64
- Infrastructure Change Log report 336
- Installation
 - computer preparations 32
 - procedure 31
- Installed Software
 - report 351
- Installing
 - collector 43
 - DRA 31, 32
- Interfaces
 - DRA, quick tour 49
- Introducing
 - DRA 17
- Investigation 28
- IT Discovery and Scanning engine 19
- Item display methods
 - graphical view 298
- L**
- License
 - configuring 231
 - viewing details 233
- License Usage
 - report 349
- Lifecycle
 - ticket 29
- Log
 - system 211
- Logging in 42
- M**
- Main window
 - DRA 50
- Maintenance system properties 249
- Management console
 - removing 78
- Management consoles
 - configuring EMC ECC and HDS HiCom-
mand 72
 - configuring HDS HiCommand 72
 - configuring IBM HMC 73
 - configuring VMware vCenter 73
 - EMC ECC, configuring 72
- Management settings
 - gap, configuring 229
 - user, configuring 217
- Managing
 - automatic association of databases to
groups and policies 110
 - database groups 108
 - database policies 107
 - host groups 101
 - host policies 100
- Managing users 46
- Module
 - Comparison 56, 353
 - Configuration 57
 - Reports 54

- SLA 55
- SLA, overview 353, 371
- Tickets 53
- Tickets, opening Topology module 294
- Tickets, overview 255
- Tickets, using 255
- Topology 51
- Topology, accessing 294
- Topology, displaying item details 311
- Topology, opening Dashboard 295
- Topology, opening Tickets module 294
- Topology, operations 303
- Topology, overview 280
- Topology, tickets 266
- Topology, using 279
- MS SQL Server Backup Summary
 - report 345
- N**
- Navigator 292
- NetApp Filer Replication Summary
 - report 330
- NetApp proxies
 - configuring 86
 - configuring, storage discovery 81
- New Databases Notification report 334
- Note to ticket
 - adding 270
- O**
- Object model
 - DRA 295
- Open tickets
 - handling 30
- Opening
 - tickets 29
 - Topology module from Dashboard 295
 - Topology module from Tickets module 294
- Operations
 - multiple tickets 274
 - tickets 268
 - Topology module 303
- Optimization
 - report 324

- Oracle database
 - security 46
- Overview
 - Comparison module 353
 - DRA 17
 - DRA reporting 317
 - security 46
 - SLA module 371
 - Tickets module 255
 - Topology module 280
- P**
- Policies
 - configuring 148
 - database, configuring 105
 - database, managing 107
 - host, configuring 95
 - host, managing 100
- Policy
 - adding 148
 - editing 151
- Pre-installation
 - requirements 31
- Probes
 - DRA 70
- Procedure
 - installation 31
- Properties
 - tab 312
- Proxies
 - configuring EMC NaviCLI 82
 - configuring EMC SYMCLI 81
 - configuring EMC SYMCLI, EMC NaviCLI, IBM DSCLI, IBM XCLI, Hitachi Hi-Command, and NetApp, storage discovery 81
 - configuring Hitachi HiCommand 86
 - configuring IBM DSCLI 85
 - configuring IBM XCLI 85
 - configuring SVC 86
- Proxy
 - configuring 156
- Q**
- Quick tour

- DRA interface 49
- R**
- Rating to ticket
 - adding 271
- Related Tickets
 - tab 313
- Removing
 - management console 78
- Replicated Temporary Databases
 - report 346
- Report
 - Business Entity Scan Coverage 347
 - Custom 351
 - Data Collection Expansion Package
 - Summary 350
 - Data Collectors Summary 350
 - Database Replica Synchronization 337
 - Database Storage Utilization 342
 - Host Configuration Differences 335
 - Host HBA Comparison 344
 - Hosts with a Single SAN I/O Path 339
 - Infrastructure Change Log 336
 - Installed Software 351
 - License Usage 349
 - MS SQL Server Backup Summary 345
 - NetApp Filer Replication Summary 330
 - New Databases Notification 334
 - Optimization 324
 - Replicated Temporary Databases 346
 - Scan Status 321
 - Scan Troubleshooting 341
 - Standby Pairs 333
 - Storage Summary 340
 - structure 325
 - System Event Log 322
 - system properties 249
 - Ticket Details 323
 - types 319
 - Unscanned Servers/Arrays 338
 - VMware Summary 343
- Report generator 20
- Report schedule
 - adding new email 145
- Reporting
 - DRA 317
 - DRA, overview 317
- Reporting packages 208
- Reports
 - exporting 319
 - module 54
 - scheduling 144
- Requirements
 - pre-installation 31
- Restarting
 - DRA 253
- Restarting DRA 253
- Risk detection packages 198
- Roles
 - business entity, defining 170
 - user, configuring 220
- Rule engine system properties 250
- S**
- Scan Coverage Details area 65
- Scan Coverage Summary area 64
- Scan status 212
- Scan Status report 321
- Scan task
 - scheduling 132
- Scan Troubleshooting
 - report 341
- Scan troubleshooting information 213
- Scanning group assignment rules 111
- Scans
 - full system, scheduling 139
 - scheduling 129
- Schedule
 - adding new email report 145
- Scheduling
 - data analysis task 137
 - full cycle scans 139
 - reports 144
 - scan task 132
 - scans 129
- Searching
 - ticket 258
- Security

- considerations 46
- DRA 46
- Oracle database 46
- Sending tickets directly from the application 272
- Server hardening 47
- Session pane 58
- Settings
 - advanced, configuring 217
 - gap management, configuring 229
 - user management, configuring 217
- Site
 - adding from CVS file 94
- SLA
 - DRA, viewing information 371
 - module 55
 - system properties 251
- SLA module
 - overview 353, 371
- SLA policies
 - configuring 172
- SLA window
 - using 372
- Standby definitions
 - adding from CSV file 122
 - configuring 119
- Standby Pairs report 333
- Standby policies
 - attaching 125
 - configuring 123
- Status
 - viewing 113
- Storage
 - unused, grouping criteria 325
- Storage capacity
 - summary, per criterion 325
- Storage Summary report 340
- Structure
 - report 325
- Sudo
 - configuring 157
- Summary
 - storage capacity per criterion 325
- Support tool 253
- SVC proxies
 - configuring 86
- System
 - log 211
- System API system properties 251
- System API views 253
- System Event Log
 - report 322
- System properties
 - active directory 236
 - additional collection 237
 - aggregation 237
 - automatic import 238
 - collection 239
 - collection timeouts 241
 - command line interface 243
 - comparison 243
 - configuring 235
 - Cyber-Ark configuration 244
 - Dashboard 245
 - distributed collection 246
 - email 246
 - external ticketing system integration 247
 - full cycle 247
 - gap rules 247
 - general 247
 - housekeeping 248
 - maintenance 249
 - report 249
 - rule engine 250
 - SLA 251
 - system API 251
 - tickets module 252
- System scans
 - full, scheduling 139
- System tools 252
 - analyzer 253
 - API views 253
 - DRA restart 253
 - support tool 253

T**Tab**

- Connections 313
- History 314
- Properties 312
- Related Tickets 313

Table display layouts

- controlling 61

Tabs

- details, topology 283

Ticket

- adding note 270
- exporting from the Tickets module 273
- handling 28
- lifecycle 29
- rating, adding 271
- searching 258
- sending directly from the application 272
- unique ID, generating 272

Ticket Details

- report 323

Ticket labels

- assigning to a ticket 277
- defining 228

Ticket Summary area 66**Ticket Tabs area 66****Tickets**

- assigning a label 277
- automatic opening/closing 29
- module 53
- open 30
- operations 268
- operations on multiple tickets 274
- Related, tab 313
- Topology module 266

Tickets module

- opening the Topology module 294
- overview 255
- using 255

Tickets module system properties 252**Toolbar**

- Topology 285

Topology

- details tabs 283
- exporting 289
- graphical view 281
- module 51
- search 300
- toolbar 285
- tree 293

Topology module

- accessing 294
- displaying item details 311
- opening the Tickets module 294
- opening, Dashboard 295
- operations 303
- overview 280
- tickets 266
- using 279

Tree

- Topology 293

Troubleshooting

- viewing scanning errors 213

Tuning 28**Types**

- business entity, defining 168
- report 319

U**Uninstalling DRA 43****Unique ID**

- ticket, generating 272

Unscanned Servers/Arrays report 338**Unused storage**

- grouping criteria 325

User management 46**User management settings**

- configuring 217

User roles

- configuring 220

Using

- Dashboard 63
- SLA window 372
- Tickets module 255
- Topology module 279

V

Veritas Disaster Recovery Advisor
see DRA 17

View

graphical, topology 281

Viewing

Dashboard information 64
entity status 113
license details 233
scan status 212
scan troubleshooting information 213

Viewing DRA SLA
information 371

VMware Summary
report 343

VMware vCenter
configuring 73

W

Window

SLA, using 372

Window display layout
controlling 60

Wizard

Configuration 71

Worksheet

configuring 355