Veritas™ Risk Advisor Deployment Requirements

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

7.3.3

Veritas Risk Advisor Deployment Requirements

Legal Notice

Document version: 7.3.3

Last updated: 2019-02-03

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. VERITAS TECHNOLOGIES LLC 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, et seq. "Commercial Computer Software and Commercial Computer Software Documentation," as applicable, and any successor regulations, whether delivered by Veritas as on premises or hosted services. 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.

Veritas Technologies LLC
500 East Middlefield Road
Mountain View, CA 94043

http://www.veritas.com

Technical Support

Veritas 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 Veritas to answer your questions in a timely fashion. For example, the Technical Support group works with Product Engineering and Veritas Security Response to provide alerting services and virus definition updates.

Veritas's maintenance 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 Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers automatic software upgrade protection
- Global support that is available 24 hours a day, 7 days a week
- Advanced features, including Account Management Services

For information about Veritas's Maintenance Programs, you can visit our Web site at the following URL:

www.veritas.com/techsupp

Contacting Technical Support

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

http://www.veritas.com/business/support/assistance_care.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 Veritas
- Recent software configuration changes and network changes

Licensing and registration

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

www.veritas.com/techsupp

Customer service

Customer service information is available at the following URL:

www.veritas.com/techsupp

Customer Service is available to assist with 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 maintenance contracts
- Information about the Veritas Buying Programs
- Advice about Veritas's technical support options
- Nontechnical presales questions
- Issues that are related to CD-ROMs or manuals

Documentation feedback

Your feedback on product documentation is important to us. Send suggestions for improvements and reports on errors or omissions to storage_management_docs@veritas.com.

Include the title and document version (located on the second page), and chapter and section titles of the text on which you are reporting.

Maintenance agreement resources

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

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

Additional enterprise services

Veritas offers a comprehensive set of services that allow you to maximize your investment in Veritas products and to develop your knowledge, expertise, and global insight, which enable you to manage your business risks proactively.

Enterprise services that are available include the following:

Veritas Early Warning Solutions — These solutions provide early warning of cyber attacks, comprehensive threat analysis, and countermeasures to prevent attacks before they occur.

Managed Security Services — These services remove the burden of managing and monitoring security devices and events, ensuring rapid response to real threats.

Consulting Services — Veritas Consulting Services provide on-site technical expertise from Veritas and its trusted partners. Veritas Consulting Services offer a variety of prepackaged and customizable options that include assessment, design, implementation, monitoring, and management capabilities. Each is focused on establishing and maintaining the integrity and availability of your IT resources.

Educational Services — Educational Services provide a full array of technical training, security education, security certification, and awareness communication programs.

To access more information about Enterprise services, please visit our Web site at the following URL:

www.veritas.com

Select your country or language from the site index.

Contents

About this document 9

Chapter 1 VRA deployment architecture

Appendix A Methods for secure privilege provisioning

About this document

This document summarizes Veritas Risk Advisor (VRA) deployment requirements. It contains the following chapter and appendix:

- Chapter 1, "VRA deployment architecture" describes the deployment architecture for VRA.
- Appendix A, "Methods for secure privilege provisioning" describes suggested methods for secure privilege provisioning for the various entities supported by VRA.

Intended audience

This document is intended for the following:

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

Chapter 1

VRA deployment architecture

This chapter includes the following topics:

Deployment environment

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

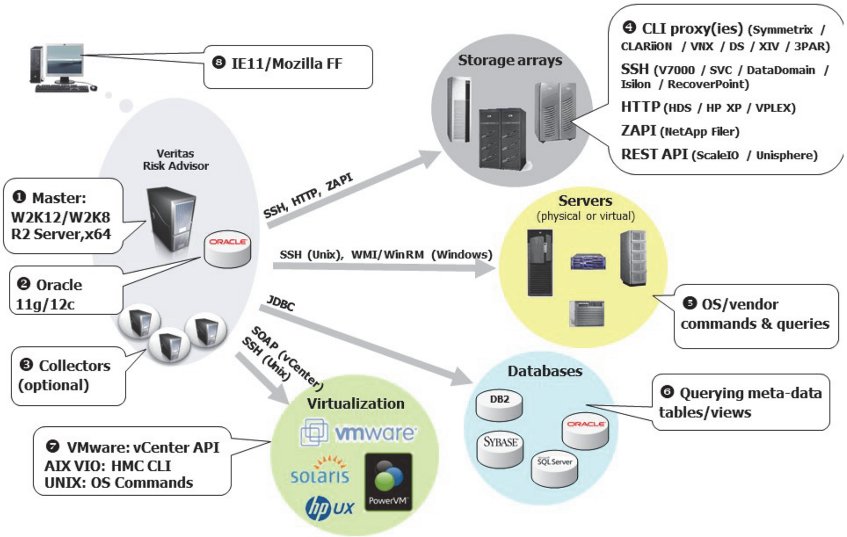

The VRA deployment environment consists of the following:

- A Windows Server 2008 R2, Windows Server 2012 or Windows Server 2016 64-bit application server, a dedicated Oracle 11g or 12c (12.1.0.2.0 or 12.2.0.1.0) repository to store the collected and analyzed data and optionally collectors to scan large environments (items 1-3 in the illustration).
- Various IT sources (4-7) that VRA collects data from for daily risk analysis.
- A Web interface to use and manage VRA (8).

VRA server

VRA must run on a dedicated host, also referred to as the VRA application server. The host size depends on several parameters, including the size of your scanned environment and how much data you need to retain. The use of a virtual machine for the VRA application server should be limited to small to mid-size environments (typically, up to 1000 scanned hosts). If VMware is used, it is advised to take particular care to reserve the required CPU and memory.

The recommended server configuration is as follows:

Table 1-1 VRA server requirements

# of scanned hosts	# of cores[a]	RAM	Free disk space[b]	Operating system
Up to 1,000	4	32 GB	250 GB	Windows Server 2008 R2 or Windows Server 2012 or Windows Server 2016, Standard or Enterprise Edition 64-bit
1,000-2,500	4	64 GB	350 GB	
2,500-5,000	4	96 GB	400 GB	
5,000-15,000	8	128 GB	500 GB	
15,000-25,000	16	192 GB	750 GB	
25,000-50,000	16	256 GB	1 TB	
Above 50,000	Specific sizing required			

[a] If Oracle database is installed locally, it is recommended to have at least two different sockets

[b] If Oracle database is installed remotely, the free disk space can be up to 40% lower

The following server requirements also apply:

- Database: Oracle 11g or 12c (12.1.0.2.0 or 12.2.0.1.0) Standard or Enterprise installed with full database administrator rights. Oracle client must be installed if using a remote Oracle database. In this case it is also highly important to maintain low latency (1ms) and high-speed connectivity.
- It is recommended to use Solid State Drive (SSD) for VRA server and especially for Oracle database server to achieve optimal performance
- Server: Java 8 and Apache Tomcat 8.5 (the only supported version). Standard Apache Software Foundation Tomcat package will be installed as a service during the installation of VRA if not already installed on your system. The setup wizard will create a local user account (member of the Local Administrators group) named "tomcatuser" and set it as the logon account for the Tomcat service. It is important to verify that the "tomcatuser" user account is not blocked or restricted by any security tool.

Furthermore, as part of VRA Tomcat architecture, a watchdog service named "Apache Tomcat Watchdog" will be deployed on the server during the VRA installation.

- Web client access:
  - Internet Explorer 11, Mozilla Firefox or Google Chrome.
  - HTTP/HTTPS access from clients to the VRA server through port 8080/8443 (configurable).

VRA web user interface is best displayed and operated with Full HD resolution (1080p) on minimum 21" screens with aspect ratio of 16:9.

VRA requires administrator rights on the VRA application server.

Note: If VMware is used, it is advised to take particular care to reserve the required CPU and memory.

The recommended collector configuration is as follows:

Table 1-2 Collector requirements

# of scanned hosts	# of cores	RAM	Free disk space	Operating system
Up to 1,000	4	12 GB	60 GB	Windows Server 2008 R2, Windows Server 2012 or Windows Server 2016, Standard or Enterprise Edition 64-bit Linux[a]
1,000-5,000	4	16 GB	120 GB	
5,000-15,000	4	32 GB	180 GB	
Above 15,000	Specific sizing required			

a.Collectors installed on Linux cannot scan Windows hosts

VRA collectors require administrator rights on the server.

Note: If VMware is used, it is advised to take particular care to reserve the required CPU and memory.

Oracle environment and licensing

VRA uses an Oracle 11g / 12c database to store and analyze the data collected from the scanned environment, also known as the VRA repository. Before you install VRA, you must install and configure the Oracle database.

To obtain an Oracle 11g / 12c 30-day trial license, go to:

http://edelivery.oracle.com

For a longer trial period or to install VRA permanently, you need an Oracle 11g / 12c standard edition license. To obtain an Oracle 11g /12c standard edition license, contact Oracle.

VRA requires full DBA rights on the VRA repository.

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

Web client requirements

VRA has a Web-based user interface. You need Internet Explorer 11, Mozilla Firefox or Google Chrome. HTTP/HTTPS IP access from the client to the VRA application server should be available. The default connection port is 8080 for HTTP and 8443 for HTTPS, which you may change if needed.

VRA web user interface is best displayed and operated with Full HD resolution (1080p) on minimum 21" screens with aspect ratio of 16:9.

In addition, the following must be configured:

- The address http://*vra_server*:8080/VRA or https://*vra_server*:8443/VRA must be defined as a trusted site in the Internet Explorer configuration.
- The client browser must be configured to permit the running of JavaScripts, iFrames, and play animations.

Credentials and collection methods used

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

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

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

- VRA collects data in read-only mode; it does not change your configuration.
- VRA only retrieves configuration data (metadata) – never actual production content. For example, VRA may read database startup files to learn how a

database instance is configured. It may also connect to the database and issue system configuration queries to determine which files store database content. However, it does not query any production information, tablespace content, and so on.

- All queries use standard, well-known interfaces and commands. Nothing is hacked or retrieved in a non-standard way. In fact, all queries and commands used are well-known to the IT staff, who often use the same queries and commands during routine maintenance.
- None of the queries or commands put a noticeable load on servers, storage arrays, databases, or the network. The only significant computation is performed on the VRA application server and the VRA Oracle repository, which are dedicated computing resources.
- By default VRA uses agent-less remote data collection through standard protocols such as SSH and WMI. VRA also supports agent-based scan for hosts and storage CLI proxies.

Important: The information provided in this document refer to the agent-less data collection. For information regarding data collection using an agent, See “VRA Agent” on page 80.

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

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

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

About setting up VRA user profiles

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

- All user account profiles provisioned for the use of VRA must have a password that does not need resetting after the first use.
- It is strongly recommended that you provision user profiles with non-expiring passwords. If that is not possible, allow the longest possible password expiration period. Veritas recommends at least six months.

  VRA uses the provisioned account profiles noninteractively, and the default connection method does not involve any plain-text password exchange. Therefore, these account profiles pose significantly lower risk than standard ones. Replacing the passwords on all hosts on the environment presents an administrative overhead that should be balanced against this low risk.

  Finally, as long as expired passwords are not reset, VRA cannot detect risks in the environment. This should also be considered in favor of using non-expiring passwords, or ones with a long expiration period.
- It is strongly recommended to use the same user ID on many of the hosts and databases. The best practice is that you use the user ID vrauser for all operating system and database account profiles.
- The user default shell should be sh.

About privilege control software

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

For sudo, the suggested syntax for each UNIX platform is described in Appendix A, “Methods for secure privilege provisioning” on page 83.

You can adapt this syntax with any other privilege control software.

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

Data collection from SYMCLI through a UNIX / Windows proxy

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

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

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

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

For more information, see “Data collection from UNIX hosts” on page 54.

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

- /usr/symcli/bin/symcfg list
- /usr/symcli/bin/symdev list
- /usr/symcli/bin/symdisk list
- /usr/symcli/bin/symaudit list
- /usr/symcli/bin/symevent list
- /usr/symcli/bin/symmaskdb list
- /usr/symcli/bin/symaccess list
- /usr/symcli/bin/symcli -def

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

- Provide each SYMCLI proxy name or IP address.
- Provide a user account profile on each SYMCLI proxy (existing or specifically created for VRA). This is the same as any UNIX host from the same vendor. For more information, see “Data collection from UNIX hosts” on page 54.

- If you prefer, provide sudo, PowerBroker (pbrun), UPM (pmrun), or similar definitions on each SYMCLI proxy. This is the same as any UNIX host from the same vendor.

  For more information, see "Data collection from UNIX hosts" on page 54.

- Make sure that IP connectivity through SSH is available between the VRA application server and each SYMCLI proxy. The default port is 22.

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

Note: By default, VRA connects to the proxies using SSH with user/password authentication. SSH with public key authentication is also supported (The key size is limited to 4096 characters). If you prefer, you can use Telnet; however, it is considered less secure than SSH. In terms of security provisioning, the only difference in using Telnet is that the default port is port 23, instead of the SSH port.

For suggestions on appropriate sudo definitions, see Appendix A, "Methods for secure privilege provisioning" on page 83.

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

Note: Windows server can also be used as SYMCLI proxies. VRA opens a WMI or WRM connection to the proxy and runs SYMCLI commands to collect data. For more information, see "Data collection from Windows" on page 63.

Note: There is an option to define a remote server that will query other SYMCLI proxies with the single set of credentials. Contact Support for further details.

Data collection from NaviCLI through a UNIX / Windows proxy

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

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

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

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

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

You can use sudo (or similar) software to achieve already declared authorization for the array. For more information, see “Data collection from UNIX hosts” on page 54.

VRA uses the following NaviCLI command syntax.

- /opt/Navisphere/bin/naviseccli -np -port *<port>* *<authorization>* -h *<array IP/name>* *<command>*

Where:

port — The port used for CLARiiON access. The default is 443.

authorization — Empty for already declared automatic authorization. Otherwise, use the following format:

-User *<user>* -Password *<password>* -Scope *scope*

Where:

- User is the user name to be used for CLARiiON authorization.
- Password is the password to be used for CLARiiON authorization. Avoid using the **'<'** and **'>'** characters in the password.
- Scope is the scope to be used for CLARiiON authorization, represented as a numeric value (0: Global, 1: Local, and 2: LDAP).

array IP/name — The array DNS name or IP address of one of the CLARiiON storage processors (SPs).

For AIX, the path is /usr/lpp/NAVICLI.

VRA uses the following privileged, read-only NaviCLI commands.

- getall
- getlun
- storagepool -list -all
- metalun -list
- mirror -async -listgroups
- mirror -async -list
- mirror -sync -listgroups
- mirror -sync -list
- snapview -listclonefeature
- snapview –listclonegroup
- getlog
- sancopy -settings -list

You should provide the following information for each NaviCLI proxy:

- Name or IP address

- A user account profile (existing or specifically created for VRA)

You should provide the following information for each CLARiiON array:

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

You should also verify that:

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

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

Note: By default, VRA connects to the proxies using SSH with user/password authentication. SSH with public key authentication is also supported (The key size is limited to 4096 characters). If you prefer, you can use Telnet; however, it is considered less secure than SSH. In terms of security provisioning, the only difference in using Telnet is that the default port is port 23, instead of the SSH port.

For suggestions regarding appropriate sudo definitions, see Appendix A, "Methods for secure privilege provisioning" on page 83.

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

Note: Windows server can also be used as NaviCLI proxies. VRA opens a WMI or WRM connection to the proxy and runs NaviCLI commands to collect data. For more information, see "Data collection from Windows" on page 63.

Data collection from DSCLI through a UNIX proxy

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

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

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

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

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

VRA uses the following DSCLI command syntax.

- dscli -hmc1 <array IP/name> -user <user> -passwd <passwd> <command>

VRA uses the following read-only DSCLI commands.

- lssu
- lssi
- lsarraysite -dev <SI>
- lsarray -dev <SI>
- lsrank -dev <SI>
- lsextpool -dev <SI>
- lsfbvol -dev <SI>
- lssestg -dev <SI>
- lslss -dev <SI>
- lsflash -dev <SI>
- lssession -dev <SI>
- lspprcpath -dev <SI>
- lspprc -dev <SI>
- showgmir -dev <SI>

You should provide the following information for each DSCLI proxy:

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

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

- Name or IP address

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

You should also verify that:

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

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

Note: By default, VRA connects to the proxies using SSH with user/password authentication. SSH with public key authentication is also supported (The key size is limited to 4096 characters). If you prefer, you can use Telnet; however, it is considered less secure than SSH. In terms of security provisioning, the only difference in using Telnet is that the default port is port 23, instead of the SSH port.

Data collection from SMCLI through a UNIX proxy

VRA uses read-only SMCLI commands to collect additional data from IBM DS 3000/4000/5000 family arrays. These commands run on one or more UNIX servers in the IT environment. Collectively, these servers can query all DS arrays in the scope. These servers are also known as SMCLI proxies.

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

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

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

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

VRA uses the following read-only SMCLI commands.

- smcli -d -i

- smcli <array IP/name> -n <storage subsystem> -c show storageSubsystem summary;
- smcli <array IP/name> -n <storage subsystem> -c show storageSubsystem;
- smcli <array IP/name> -n <storage subsystem> -c show volumeCopy allLogicalDrives;

You should provide the following information for each SMCLI proxy:

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

You should provide the following information for each IBM DS array that is not defined in the SMCLI configuration:

- Name or IP address
- A user account profile (existing or specifically created for VRA). Avoid using the '<' and '>' characters in the password field.

You should also verify that:

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

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

Note: By default, VRA connects to the proxies using SSH with user/password authentication. SSH with public key authentication is also supported (The key size is limited to 4096 characters). If you prefer, you can use Telnet; however, it is considered less secure than SSH. In terms of security provisioning, the only difference in using Telnet is that the default port is port 23, instead of the SSH port.

Data collection from XCLI through a UNIX proxy

VRA uses read-only XCLI commands to collect additional data from IBM XIV or A9000 arrays. These commands run on one or more UNIX servers in the IT environment. Collectively, these servers can query all XIV or A9000 arrays in the scope. These servers are also known as XCLI proxies.

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

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

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

The connection from the proxy to the XIV or A9000 array requires authorization details (user/password) for the array.

VRA uses the following XCLI command syntax.

- xcli -m <array IP/name> -u <user> -p <passwd> <command>

VRA uses the following read-only XCLI commands.

- version_get
- config_get
- state_list
- fc_port_list
- pool_list
- vol_list
- mirror_list
- cg_list
- snap_group_list
- system_capacity_list
- time_list

You should provide the following information for each XCLI proxy:

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

You should provide the following information for each XIV or A9000 array:

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

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

You should also verify that:

- IP connectivity through SSH (default is port 22) is available between the VRA application server and each XCLI proxy.
- IP connectivity is available between the XCLI proxy and each XIV or A9000 array that it scans.

Important: Do not configure the same XIV or A9000 array to be scanned by more than one probe. This configuration may cause unpredictable results.

Note: By default, VRA connects to the proxies using SSH with user/password authentication. SSH with public key authentication is also supported (The key size is limited to 4096 characters). If you prefer, you can use Telnet; however, it is considered less secure than SSH. In terms of security provisioning, the only difference in using Telnet is that the default port is port 23, instead of the SSH port.

Data collection from 3PAR

VRA can scan HP 3PAR arrays either through Unix/Windows proxy server or directly through SSH. VRA uses read-only CLI commands to collect additional data from HP 3PAR arrays.

Data collection from 3PAR using InForm CLI Proxy

InForm CLI commands run on one or more UNIX or Windows servers in the IT environment. Collectively, these servers can query all 3PAR arrays in the scope. These servers are also known as InForm CLI proxies.

When you select InForm CLI proxies, use the following best practices:

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

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

The connection from the proxy to the 3PAR array requires authorization details (user/password) for the array.

VRA uses the following InForm CLI command syntax.

- cli -sys <array IP/name> -user <user> -password <passwd> <command>

VRA uses the following read-only InForm CLI commands.

- showsys -d
- showversion -a
- showlicense
- showdate
- showinventory
- showeeprom
- showcage -d
- showpd
- showld -d
- showcpg -d
- showvv
- showvv -cpgalloc
- showvv -d
- showvv -pol
- showvvset -d
- showdomain -d
- showdomainset -d
- showhost -d
- showhostset -d
- shownode -d
- shownet
- showport
- showrctransport
- showvlun
- showrcopy -d

You should provide the following information for each InForm CLI proxy:

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

You should provide the following information for each 3PAR array:

- Name or IP address
- A user account profile (existing or specifically created for VRA) with a *read-only* privilege. In the password field, you should enter the encrypted password of the 3PAR storage array.

You should also verify that:

- IP connectivity through SSH (default is port 22) is available between the VRA application server and each InFrom CLI proxy.
- IP connectivity is available between the InForm CLI proxy and each 3PAR array that it scans.

Note: By default, VRA connects to the proxies using SSH with user/password authentication. SSH with public key authentication is also supported (The key size is limited to 4096 characters). If you prefer, you can use Telnet; however, it is considered less secure than SSH. In terms of security provisioning, the only difference in using Telnet is that the default port is port 23, instead of the SSH port.

Note: Windows server can also be used as Inform CLI proxies. VRA opens a WMI or WRM connection to the proxy and runs Inform CLI commands to collect data. For more information, see “Data collection from Windows” on page 63.

Data collection from 3PAR through SSH

Make sure IP connectivity through SSH (default is port 22) is available between the VRA application server and each 3PAR array. All other requirements are identical to those listed above.

Data collection from HDS HiCommand/HP CommandView

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

VRA collects data using the following read-only requests:

- GetServerInfo
- GetStorageArray
- GetHost
- GetVStorageArray

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

- Provide each HiCommand server name or IP address.
- Provide the HiCommand Web application user name and password of a user with View rights that is assigned to a group. A default group is acceptable.
- Make sure that IP connectivity through HTTP or HTTPS (port 2001) is available between the VRA application server and each HiCommand server.
- Scanning of Hitachi G series systems requires HiCommand API version 8 or higher. VRA automatically discovers the right version of HiCommand which can be overridden using the system property “HiCommand API major version” which can be found under Collection system properties.
- In case of HiCommand version lower than 5, the aforementioned system properties should be set explicitly to the right version. In case of HiCommand version 5 and higher, setting these properties to 0 (zero) will assure the automatic version detection.

Data Collection from HCP

VRA collects data from Hitachi Content Platform (HCP) object storage platform using REST API by opening an HTTPS connection to the HCP. VRA collects data using the following REST API requests on the HCP cluster and tenants:

- services/statistics
- services/replication/links
- services/replication/links/linkname
- nodes/statistics
- tenants/
- tenants/tenantname/
- tenants/tenantname/namespaces
- tenants/tenantname/namespaces/namespacename
- tenants/tenantname/namespaces/namespacename/permissions
- tenants/tenantname/namespaces/namespacename/complianceSettings
- tenants/tenantname/namespaces/namespacename/replicationCollisionSettins
- tenants/tenantname/namespaces/namespacename/protocols/http
- tenants/tenantname/namespaces/namespacename/protocols/cifs
- tenants/tenantname/namespaces/namespacename/protocols/nfs

- tenants/tenantname/namespaces/namespacename/protocols/smtp
- tenants/tenantname/userAccounts
- tenants/tenanatname/userAccounts/username
- tenants/tenanatname/userAccounts/username/dataAccessPermissions
- tenants/tenanatname/groupAccounts/groupname/dataAccessPermissions

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

- Provide HCP DNS name
- Provide a user name and password for each HCP, where user role is monitor
- Make sure that IP connectivity through HTTPS (port 9090) is available between the VRA application server and each HCP.

Data collection from V7000/SVC

VRA uses read-only V7000/SVC CLI commands to collect additional data from V7000/SVC arrays.

As a standard, VRA opens an SSH session to the V7000/SVC in the same way that it does to collect data from any UNIX host.

VRA uses the following read-only V7000/SVC commands:

- svcinfo lscluster
- svcinfo lsnode
- svcinfo lslicense
- svcinfo lsclusterip
- svcinfo lsportip
- svcinfo lsvdisk
- svcinfo lsiogrp
- svcinfo lsmdiskgrp
- svcinfo lsmdisk
- svcinfo lscontroller
- svcinfo lsfcmap
- svcinfo lsfcconsistgrp
- svcinfo lsrcrelationship
- svcinfo lsrcconsistgrp
- svcinfo lsquorum

- svcinfo lsvdiskcopy
- svcinfo lssevdiskcopy
- svcinfo lshost
- svcinfo lshostvdiskmap
- svcinfo lsfabric
- svcinfo lsmdiskextent
- svcinfo lsportfc
- svcinfo lsvdiskhostmap

You should provide the following information for each V7000/SVC array:

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

You should also verify that IP connectivity through SSH (default is port 22) is available between the VRA application server and each V7000/SVC array.

Data collection from VPLEX

VRA collects data from EMC VPLEX arrays by opening RESTful API by opening an HTTPS connection to the VPLEX Management Server, or servers, if more than one is used.

VRA collects data using the following read-only requests:

- GET https://<VPLEX Management Server IP>/vplex/**

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

- Provide each VPLEX Management Server name or IP address. In VPLEX Metro or Geo systems enter only one Management Server
- Provide the VPLEX Management Server user name and password.
- Make sure that IP connectivity through HTTPS (port 443) is available between the VRA application server and each VPLEX Management Server.

Data collection from InfiniBox

VRA collects data from Infinidat InfiniBox arrays by using REST API by opening an HTTPS connection to the InfiniBox array.

VRA collects data using the following REST API requests:

- system
- components

- pools
- datasets
- cgs
- hosts
- clusters
- links
- replicas

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

- Provide each InfiniBox name or IP address
- Provide the InfiniBox user name and password.
- Make sure that IP connectivity through HTTPS (port 443) is available between the VRA application server and each InfiniBox array.

Data collection from NetApp

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

You should provide the following information for each NetApp Filer (7-MODE):

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

  The user should be an administrative user assigned with the *Admin* role or a custom restricted user with the following capabilities:

  - login-http-admin
  - api-system-get-info
  - api-system-get-version
  - api-license-v2-list-info
  - api-net-ifconfig-get
  - api-clock-get-clock
  - api-options-list-info
  - api-vfiler-list-info
  - api-aggr-list-info
  - api-volume-list-info
  - api-qtree-list
  - api-snapshot-list-info

- api-lun-list-info
- api-igroup-list-info
- api-iscsi-node-get-name
- api-lun-map-list-info
- api-nfs-exportfs-list-rules
- api-nfs-exportfs-list-rules-2
- api-cifs-share-list-iter-start
- api-cifs-share-list-iter-next
- api-cifs-share-list-iter-end
- api-snapmirror-get-status
- api-snapvault-primary-relationship-status-list-iter-start
- api-snapvault-primary-relationship-status-list-iter-next
- api-snapvault-primary-relationship-status-list-iter-end
- api-disk-list-info
- api-net-resolve
- api-volume-get-root-name
- api-storage-adapter-get-adapter-info
- api-storage-adapter-get-adapter-list

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

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

- Create a new role using one of the following options. The latter option is more restricted (read only):
  - useradmin role add <role_name> -a login-http-admin,api-*
  - useradmin role add <role_name> -a login-http-admin,api-system-get-info,api-system-get-version, api-license-v2-list-info,api-net-ifconfig-get,api-clock-get-clock, api-options-list-info,api-vfiler-list-info,api-aggr-list-info, api-volume-list-info,api-qtree-list,api-snapshot-list-info, api-lun-list-info,api-igroup-list-info,api-iscsi-node-get-name, api-lun-map-list-info,api-nfs-exportfs-list-rules, api-nfs-exportfs-list-rules-2,api-cifs-share-list-iter-start, api-cifs-share-list-iter-next,api-cifs-share-list-iter-end, api-snapmirror-get-status, api-snapvault-primary-relationship-status-list-iter-start, api-snapvault-primary-relationship-status-list-iter-next, api-snapvault-primary-relationship-status-list-iter-end, api-disk-list-info,api-net-resolve,api-volume-get-root-name, api-storage-adapter-get-adapter-info, api-storage-adapter-get-adapter-list
- Create a new group:
  - useradmin group add <group_name> -r <role_name>
- Create a new user:
  - useradmin user add <user_name> -g <group_name>

You should provide the following information for each NetApp Cluster:

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

  The user should be with the predefined *readonly* role or with a custom role with the following readonly permissions:
  - cluster identity show
  - cluster peer show
  - network interface show
  - snapmirror show
  - storage aggregate show
  - storage disk show
  - system node show
  - vserver show

  - vserver cifs share show
  - vserver iscsi show
  - vserver options
  - vserver peer show
  - volume show
  - volume qtree show
  - volume snapshot show
  - lun show
  - lun igroup show
  - lun mapped show
  - version

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

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

- Use the predefined *readonly* role or create a new role as described below:
  - role create -role <role_name> -cmddirname "cluster identity show" -access readonly
  - role create -role <role_name> -cmddirname "cluster peer show" -access readonly
  - role create -role <role_name> -cmddirname "network interface show" -access readonly
  - role create -role <role_name> -cmddirname "snapmirror show" -access readonly
  - role create -role <role_name> -cmddirname "storage aggregate show" -access readonly
  - role create -role <role_name> -cmddirname "storage disk show" -access readonly
  - role create -role <role_name> -cmddirname "system node show" -access readonly
  - role create -role <role_name> -cmddirname "vserver show" -access readonly
  - role create -role <role_name> -cmddirname "vserver cifs share show" -access readonly
  - role create -role <role_name> -cmddirname "vserver iscsi show" -access readonly

  - role create -role <role_name> -cmddirname "vserver options" -access readonly
  - role create -role <role_name> -cmddirname "vserver peer show" -access readonly
  - role create -role <role_name> -cmddirname "volume show" -access readonlyrole create -role <role_name> -cmddirname "volume qtree show" -access readonly
  - role create -role <role_name> -cmddirname "volume snapshot show" -access readonly
  - role create -role <role_name> -cmddirname "lun show" -access readonly
  - role create -role <role_name> -cmddirname "lun igroup show" -access readonly
  - role create -role <role_name> -cmddirname "lun mapped show" -access readonly
  - role create -role <role_name> -cmddirname "version" -access readonly
- Create a new user:
  - security login create -username <user_name> -application ontapi -authmethod password -role <role_name>

Data collection from Cisco MDS Fabric Switch

VRA uses either read-only Cisco MDS CLI show commands or Cisco Data Center Network Manager to collect data from Cisco MDS fabric switch.

When using a direct switch scan, VRA opens an SSH session to the Cisco MDS switch in the same way that it does to collect data from any UNIX host.

VRA uses the following read-only Cisco MDS CLI show commands:

- show switchname
- show hardware
- show module
- show interface
- show port internal info
- show flogi database
- show topology
- show vsan
- show zoneset active

- show zone status
- show fcs database
- show npv status
- show npv flogi-table
- show fcns database npv detail
- show cfs application
- show aaa authentication
- show feature
- show hardware
- show http
- show ip interface mgmt0
- show logging server
- show ntp peer-status
- show role brief
- show role name default-role
- show role network admin
- show snmp user
- show sprom backplane 1
- show tacacs-server
- show tacacs-server groups
- show telnet server
- show user-account
- show version
- show zoneset

You should provide the following information for one switch in each fabric (assuming that all the switches in the fabric have the same user/password):

- Name or IP address
- A user account with permission to run show commands

You should also verify that IP connectivity through SSH (default is port 22) is available between the VRA application server and each Cisco MDS switch.

When using Cisco Data Center Network Manager (DCNM), VRA collects data from DCNM by running the following read only soap queries:

- SanWSService/SanWS:getFabrics
- SanWSService/SanWS:getSwitchesByFabric
- SanWSService/SanWS:getFcPorts
- SanWSService/SanWS:getEndports
- SanWSService/SanWS:getVsans
- SanWSService/SanWS:getIslsWithPCMembers
- SanWSService/SanWS:getNpvLinksWithPCMembers

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

- Provide each DCNM server name or IP address
- Provide a user name and password for each DCNM server
- Make sure that IP connectivity through Soap (port 443) is available between the VRA application server and each DCNM server

Data collection from Brocade Fabric Switch

VRA uses either read-only Brocade CLI commands or Brocade Network Advisor to collect data on Brocade fabric switch.

When using a direct switch scan, VRA opens an SSH session to the Brocade switch in the same way that it does to collect data from any UNIX host.

VRA uses the following read-only Brocade CLI commands:

- chassisshow
- configshow
- slotshow
- version
- lscfg --show
- fosconfig --show
- userconfig --show
- cfgshow (used with fosexec as well)
- fabricshow (used with fosexec as well)
- islshow (used with fosexec as well)
- portcfgshow (used with fosexec as well)

- portflagsshow (used with fosexec as well)
- portshow (used with fosexec as well)
- switchshow (used with fosexec as well)
- trunkShow (used with fosexec as well)
- zoneshow (used with fosexec as well)
- agshow (used with fosexec as well)
- fosexec (when Logical switches with the same IP address are used. Commands: agshow, cfgshow, fabricshow, islshow, portcfgshow, portflagsshow, portshow, switchshow, trunkshow, zoneshow)
- agautomapbalance --show
- ag --show
- ag --modeshow
- ag --mapshow
- ag --wwnmapshow
- ag --adsshow
- ag --failbackshow
- ag --failovershow
- ag --pgshow
- ag --policyshow

You should provide the following information for one switch in each fabric (assuming that all the switches in the fabric have the same user/password):

- Name or IP address
- A user account with permission to run read-only commands. In a Virtual Fabric environment, chassis permission is required.

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

When using Brocade Network Advisor (BNA), VRA opens a REST API session to the BNA and collects data by running read only requests on the following entities:

- fcfabrics
- fcswitches
- fcports

- zonedbs

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

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

Data collection from HP Virtual Connect

VRA uses read-only HP Virtual Connect CLI to collect data from HP Virtual Connect switches.

As a standard, VRA opens an SSH session to the HP Virtual Connect switch in the same way that it does to collect data from any UNIX host.

VRA uses the following read-only HP Virtual Connect commands:

- show all * -output=script1

You should provide the following information for each HP Virtual Connect domain:

- Name or IP address of the HP Virtual Connect domain
- A user account with user privilege level

You should also verify that IP connectivity through SSH (default is port 22) is available between the VRA application server and each HP Virtual Connect domain.

Data collection from EMC RecoverPoint

VRA uses read-only RecoverPoint CLI commands to collect data from EMC RecoverPoint.

As a standard, VRA opens an SSH session to the EMC RecoverPoint CLI in the same way that it does to collect data from any UNIX host.

VRA uses the following read-only RecoverPoint CLI commands:

- get_version
- get_system_report
- get_group_volumes
- get_group_state
- get_group_settings
- get_group_statistics

You should provide the following information for one RPA in each RecoverPoint installation:

- Name or IP address
- A user account with a view permission. You can use the predefined monitor user.

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

Data collection from EMC Isilon

VRA uses read-only CLI commands to collect data from EMC Isilon.

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

VRA uses the following read-only CLI commands:

- isi status
- isi version
- isi networks list interfaces
- isi networks list subnets
- isi license status
- isi snapshot snapshots list -v
- isi sync policy list -v
- isi sync target list -v
- isi nfs exports list -v
- isi smb shares list -v

You should provide the following information for each Isilon:

- Name or IP address
- A user account with permission to run read-only isi commands

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

Data collection from EMC DataDomain

VRA uses read-only CLI commands to collect data from EMC DataDomain.

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

VRA uses the following read-only CLI commands:

- system show all
- net show all
- filesys show space
- mtree list
- nfs show clients
- cifs share show
- replication show config all
- replication show stats all
- log host show
- system show modelno
- system show serialno
- system show version

The following non-mandatory commands require elevated privileges:

- adminaccess ssh option show
- adminaccess web option show
- user password aging show
- user show list

You should provide the following information for each DataDomain:

- Name or IP address
- A user account with a user role

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

Data Collection from EMC XtremIO

VRA collects data from EMC XtremIO arrays by opening an HTTP or HTTPS connection to the XtremIO Management server or servers (XMS), if more than one is used.

VRA collects data using REST API (api/json/v2/types) with the following read-only requests:

- clusters
- bricks
- targets
- volumes

- consistency-groups
- initiator-groups
- initiators

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

- Provide each XMS name or IP address.
- Provide the XMS user name and password of a user.
- Make sure that IP connectivity through HTTP or HTTPS (port 443) is available between the VRA application server and each XMS.

Data collection from EMC Unity

VRA collects data from EMC Unity arrays by opening an HTTP or HTTPS connection to the EMC Unisphere management server or servers of more than one is used.

VRA collects data using REST API (api/json/v2/types) with the following read-only requests:

- system
- storageProcessor
- fcPort
- lun
- snap
- filesystem
- volumesList
- storageResource
- pool
- host
- hostInitiator
- nfsShare
- cifsShare
- hostLUN
- remoteSystem
- replicationSession

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

- Provide each Unisphere server name or IP address.

- Provide the Unisphere server user name and password of a user.
  - Make sure that IP connectivity through HTTP or HTTPS (port 443) is available between the VRA application server and each Unity Server.

Data collection from EMC ScaleIO

VRA collects data from EMC ScaleIO Software Defined Storage by opening an HTTP or HTTPS connection to the EMC ScaleIO Gateway instance server or instances if more than one is used.

VRA collects data using REST API with the following read-only

requests:

- api/api/types/ProtectionDomain/instances
- api/types/FaultSet/instances
- api/instances/FaultSet::<fault-set>/relationships/Sds
- api/instances/Sds::<fault-set>/relationships/Device
- api/types/StoragePool/instances
- api/types/Sdc/instances
- api/types/Volume/instances
- api/types/VTree/instancestypes/System/instances

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

- Provide each ScaleIO Gateway server name or IP address.
- Provide the user name and password of a user with Monitoring role or higher.
- Make sure that IP connectivity through HTTP or HTTPS (port 443) is available between the VRA application server and each ScaleIO Gateway.
- All hosts running ScaleIO SDSs and MDMs should be explicitly scanned.
- For each Linux host running SDS, the user used for scanning the host should have permission to access the following file:
  - /opt/emc/scaleio/sds/cfg/rep_tgt.txt
- For each Linux host running MDM, the user used for scanning the host should have permission to run the following command:
  - /sbin/service

Data collection from VMware vCenter

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

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

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

If Zerto is used, VRA also collects Zerto data using the Zerto Virtual Replication REST API.

VRA collects data by running read-only requests on the following entities:

- /v1/localsite
- /v1/vpgs
- /v1/vras
- /v1/vms

If VMware NSX is used and it is accessible with the same credentials as the corresponding vCenter then VRA also collect the NSX data using its REST API. If VMware NSX is used but requires different credentials then VRA automatically creates a probe which should be configured with the corresponding credentials. See “Data Collection from VMware NSX” on page 46 for details.

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

- Provide each vCenter server name or IP address.
- Provide a read-only user name and password for each vCenter server.
  If SRM is used, make sure the user has a read-only privilege in each SRM.
  If Zerto is used, the user cannot be a member of vCenter internal security domain.
- Make sure that IP connectivity through HTTPS (port 443) is available between the VRA application server and each vCenter server.

Data Collection from VMware NSX

VRA collects data from VMware NSX by running the following read-only REST API queries on the VMware NSX Manager:

- api/1.0/appliance-management/backuprestore/backups
- api/1.0/appliance-management/backuprestore/backupsettings
- api/1.0/appliance-management/summary/system
- api/2.0/services/vcconfig
- api/2.0/services/vcconfig/status
- api/2.0/vdn/controller
- api/2.0/vdn/scopes
- api/2.0/vdn/virtualwires
- api/4.0/edges
- api/4.0/edges/<edge-id>/appliances
- api/4.0/edges/<edge-id>/highavailability/config
- api/4.0/edges/<edge-id>/interfaces
- api/4.0/edges/<edge-id>/mgmtinterface
- api/4.0/edges/<edge-id>/summary
- api/4.0/edges/<edge-id>/vnics

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

- Provide each VMware NSX Manager name or IP address.
- Provide a user with Auditor permission level.
- Make sure that IP connectivity through HTTPS (port 443) is available between the VRA application server and each NSX Manager.

Note: VRA supports VMware NSX for vSphere (NSX-V)

Data collection from Microsoft System Center Virtual Machine Manager

VRA collects data from System Center Virtual Machine Manager (SCVMM) by opening a WMI or WRM connection to the SCVMM server, as it does to collect data from any Windows host.

Scanning MS Windows Cluster 2008 R2 or higher requires PowerShell version 3.0 or higher.

For more information, see “Data collection from Windows” on page 63.

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

- SCVMMServer
- SCVMHost
- SCVMHostCluster
- SCVirtualMachine
- SCVMHostGroup
- SCVMCheckpoint
- SCVirtualHardDisk
- SCVMHostNetworkAdapter
- SCApplicableVMHostGroup
- SCDynamicOptimizationConfiguration

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

- Provide each SCVMM server name or IP address.
- Provide a user name and password for each SCVMM server.
- Make sure each user has local administrative rights
- Make sure that IP connectivity is permitted through WMI or WRM between the VRA application server and each SCVMM server.

Data collection from Oracle Enterprise Manager

VRA uses read-only JDBC queries to collect additional data from Oracle Enterprise Manager (OEM).

VRA uses the following OEM repository views:

- MGMT$DB_DBNINSTANCEINFO
- MGMT$HA_INFO
- MGMT$DB_INIT_PARAMS
- MGMT$DB_TABLESPACES
- MGMT$DB_DATAFILES
- MGMT$DB_REDOLOGS
- MGMT$DB_CONTROLFILES
- MGMT$METRIC_CURRENT (metrics DGPrimaryDBName, dataguard*, Response, Disk_Path, CRS_output, resource_status, Exadata CellDisk Metric)

- MGMT$TARGET
- MGMT$TARGET_ASSOCIATIONS
- MGMT$AVAILABILITY_CURRENT
- MGMT$APPLIED_PATCHES
- MGMT$AGENTS_MONITORING_TARGETS (OEM 12 only)
- MGMT$MANAGEABLE_ENTITIES (OEM 12 only)
- MGMT$DB_FEATUREUSAGE (OEM 12 only)
- CM$MGMT_ASM_DISKGROUP_ECM (OEM 12 only)
- CM$MGMT_ASM_DISK_ECM (OEM 12 only)

You should provide the following information for each OEM:

- Name or IP address
- Database name
- An OEM user (existing or specifically created for VRA) with a view any target privilege (EM_ALL_VIEWER in OEM 12)

You should also verify that IP connectivity through JDBC is available between the VRA application server and each OEM.

Note: Oracle Exadata machines are scanned using Oracle Enterprise Manager version 12 or higher.

Data collection from Oracle GoldenGate Monitor

VRA uses read-only JDBC queries to collect additional data from Oracle GoldenGate Monitor.

VRA uses the following Oracle GoldenGate Monitor repository views:

- MONITOR_INFO
- GGS_OBJECTS
- LINKS
- CONNECTIONS
- MPS
- MPS_COMPOSITE_VALUES

You should provide the following information for each GoldenGate Monitor:

- Name or IP address
- Database name
- A user (existing or specifically created for VRA) with read privileges on the GoldenGate Monitor repository views

You should also verify that IP connectivity through JDBC is available between the VRA application server and each GoldenGate Monitor.

Data collection from Cisco UCS Manager

VRA collects data from Cisco Unified Computing System (UCS) blade servers by opening an HTTP or HTTPS connection to the Cisco UCS Manager server or servers of more than one is used.

VRA collects data using XML API with the following read-only requests:

- method="configScope" dn="sys" inHierarchical="true"
- method="configScope" dn="org-root" inHierarchical="true"
- method="configResolveClass" classId="lsServer" inHierarchical="true"

You should provide the following information for each UCS Manager server:

- Name or IP address.
- A locally authenticated user account with read-only role

You should also make sure that IP connectivity through HTTP (port 80) or HTTPS (port 443) is available between the VRA application server and each UCS Manager server.

Data collection from NetApp OnCommand Unified Manager Core

VRA can collect data about NetApp Filers/Clusters using OnCommand Unified Manager Core (DFM). It uses the DFM as a proxy to run read-only commands on the NetApp Filers/Clusters with the ZAPI API.

Each NetApp Filer/Cluster which is scanned using DFM must have a valid login and password defined in DFM. This login should have capabilities as defined in "Data collection from NetApp" on page 33.

You should provide the following information for each DFM:

- Name or IP address
- A user account profile (existing or specifically created for VRA) with a Storage Administrator role

You should also verify that IP connectivity through HTTP (default is port 8088) or HTTPS (default is port 8488) is available between the VRA application server and each DFM.

Data Collection from F5 BIG-IP

VRA collects data from F5 Load Balancers by opening HTTPS connection to the BIG-IP server or servers if more than one is used. VRA collects data using REST API with the following read-only requests:

- mgmt/tm/ltm/pool
- mgmt/tm/ltm/pool/<pool_name> /members

You should provide the following information for each BIG-IP management server:

- Name or IP address.
- A user account with “guest” privileges

You should also make sure that IP connectivity through HTTPS (port 443) is available between the VRA application server and each BIG-IP server.

Data collection from HP OneView

VRA collects data from HP OneView by opening HTTPS connection to the OneView server or servers if more than one is used. VRA collects data using REST API with the following read-only requests:

- rest/uplink-sets
- rest/connections
- rest/ethernet-networks
- rest/interconnects
- rest/server-profiles
- rest/server-hardware
- rest/enclosures
- rest/domains
- rest/version
- rest/login-sessions

You should provide the following information for each OneView server:

- Name or IP address.

- A user account with read only privileges

You should also make sure that IP connectivity through HTTPS (port 443) is available between the VRA application server and each OneView server.

Data collection from HMC

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

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

VRA uses the following read-only HMC commands:

- uname -a
- lshmc -V
- lssyscfg -r sys
- lssyscfg -r lpar -m <hmc_system>
- lshwres -r virtualio --rsubtype slot -m <hmc_system> --level slot
- lshwres -r virtualio --rsubtype scsi -m <hmc_system>

You should provide the following information for each HMC:

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

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

Data collection from VIO

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

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

You should provide the following information for each VIO:

- Name or IP address.
- A user account profile (existing or specifically created for VRA)
  The user should be either a regular user with *ksh* and permissions, as described in "Privileged commands on AIX" on page 61 or a restricted user with *rksh* (see below).
- If a restricted user is used, make sure that *PermitUserEnvironment* is set to *yes* in the /etc/ssh/sshd_config file.

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

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

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

5 Create permission to run privileged commands. See "Privileged commands on AIX" on page 61 for the list of commands.
Do the following for each privileged command:

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

- setsecattr -c euid=0 accessauths=veritas.vra innateprivs=PV_SU_ secflags=FSF_EPS authroles= /usr/sbin/pcmpath
- setkst
- oem_setup_env
- ln -s /usr/sbin/pcmpath /usr/ios/oem

Data collection from UNIX hosts

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

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

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

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

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

    Note: For AIX, these utilities are unnecessary when CLARiiON is used.
  - If none of these utilities is available on a certain host, install the free EMC inq utility at /usr/local/bin.
- On hosts using HDS/HP XP, make sure you meet the following requirements:
  - VRA requires that at least one of the utilities (HDLM or inqraid) is installed on each UNIX host.
  - If neither of these utilities is available on a certain host, install the free HDS inqraid utility at /HORCM/usr/bin/.
- On hosts using NetApp, make sure that at least one of the utilities (SnapDrive or sanlun) is installed on each host.
- On hosts running Oracle RAC, make sure that the **oratab** file is accessible (read-only) from /etc/oratab or /var/opt/oracle/oratab
- On Solaris hosts running Oracle VM for SPARC (formerly LDOM), make sure that the user has the read privilege to LDOMs (solaris.ldoms.read). In this case VRA auto-discovers logical Solaris hosts when scanning the Control Domain server running Oracle VM for SPARC.

Note: By default, VRA connects to UNIX hosts using SSH with user/password authentication. SSH with public key authentication is also supported (The key size is limited to 4096 characters). If you prefer, you can use Telnet; however, it is considered less secure than SSH. In terms of security provisioning, the only difference in using Telnet is that the default port is port 23, instead of the SSH port.

Privileged commands on Solaris

Table 1-3 lists the privileged commands on Solaris.

Table 1-3 Privileged commands on Solaris

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

Table 1-3 Privileged commands on Solaris (Continued)

Command	Required when this is installed ...
${oracleHome}/cemutlo -w	Oracle RAC
${oracleHome}/olsnodes -l	Oracle RAC
${oracleHome}/crsctl check crs	Oracle RAC
${oracleHome}/crsctl check crsd	Oracle RAC
${oracleHome}/crsctl check cssd	Oracle RAC
${oracleHome}/crsctl check evmd	Oracle RAC
${oracleHome}/crsctl status resource -l	Oracle RAC
${oracleHome}/crsctl status resource -f	Oracle RAC
${oracleHome}/crsctl status resource -t	Oracle RAC
${oracleHome}/crsctl query css votedisk	Oracle RAC
/usr/sbin/ldm	Oracle VM for SPARC
/usr/sbin/ldmpower	Oracle VM for SPARC
/bin/cat	
/bin/ls	
/usr/sbin/ping	

Privileged commands on HP-UX

Table 1-4 lists the privileged commands on HP-UX.

Table 1-4 Privileged commands on HP-UX

Command	Required when this is installed ...
/sbin/powermt display	EMC PowerPath
/usr/symcli/bin/symdg list	EMC Solution Enabler (SYMCLI)
/usr/symcli/bin/symcg list	EMC Solution Enabler (SYMCLI)

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

Command	Required when this is installed ...
/usr/symcli/bin/sympd list	EMC Solution Enabler (SYMCLI)
/usr/symcli/bin/syminq	EMC Solution Enabler (SYMCLI)
/usr/local/bin/inq	EMC inq
/usr/sbin/vxdisk list	Veritas LVM/DMP
/usr/sbin/vxdisk path	Veritas LVM/DMP
/HORCM/usr/bin/inqraid	HDS inqraid
/usr/local/bin/lunstat -t	HDS lunstat
/usr/DynamicLinkManager/bin/dlnkmgr view	HDS HDLM
/opt/NetApp/snapdrive/bin/snapdrive storage show	NetApp SnapDrive
/opt/netapp/santools/sanlun lun show all	NetApp sanlun
/HORCM/usr/bin/raidqry	HDS horcm
/HORCM/usr/bin/raidscan	HDS horcm
/HORCM/usr/bin/pairdisplay	HDS horcm
/sbin/xpinfo	XPINFO
/sbin/spmgr display	SecurePath
/sbin/autopath display	SecurePath
/usr/sbin/vxdmpadm list	Veritas LVM/DMP
/usr/sbin/datapath query device	IBM SDD
/opt/xiv/host_attach/bin/xiv_devlist	IBM XIV HAK
/usr/sbin/cmviewcl	HP Serviceguard
/usr/sbin/cmquerycl	HP Serviceguard
${oracleHome}/cemutlo -w	Oracle RAC
${oracleHome}/olsnodes -l	Oracle RAC
${oracleHome}/crsctl check crs	Oracle RAC
${oracleHome}/crsctl check crsd	Oracle RAC

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

Command	Required when this is installed ...
${oracleHome}/crsctl check cssd	Oracle RAC
${oracleHome}/crsctl check evmd	Oracle RAC
${oracleHome}/crsctl status resource -l	Oracle RAC
${oracleHome}/crsctl status resource -f	Oracle RAC
${oracleHome}/crsctl status resource -t	Oracle RAC
${oracleHome}/crsctl query css votedisk	Oracle RAC
/bin/cat	
/bin/ls	

Privileged commands on Linux

Table 1-5 lists the privileged commands on Linux.

Table 1-5 Privileged commands on Linux

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

Table 1-5 Privileged commands on Linux (Continued)

Command	Required when this is installed ...
/usr/sbin/pvdisplay	LVM2 is used
/sbin/lvdisplay	LVM2 is used
/sbin/vgdisplay	LVM2 is used
/sbin/pvdisplay	LVM2 is used
/sbin/multipath -l	MPIO is used
/HORCM/usr/bin/inqraid	HDS inqraid
/usr/local/bin/lunstat -t	HDS lunstat
/usr/DynamicLinkManager/bin/dlnkmgr view	HDS HDLM
/opt/NetApp/snapdrive/bin/snapdrive storage show	NetApp SnapDrive
/opt/netapp/santools/sanlun lun show all	NetApp sanlun
/sbin/scsi_id	
/bin/raw –qa	A Linux raw character device is used
/usr/bin/raw –qa	A Linux raw character device is used
/HORCM/usr/bin/raidqry	HDS horcm
/HORCM/usr/bin/raidscan	HDS horcm
/HORCM/usr/bin/pairdisplay	HDS horcm
/sbin/xpinfo	XPINFO
/sbin/spmgr display	SecurePath
/sbin/autopath display	SecurePath
/usr/sbin/vxdmpadm list	Veritas LVM/DMP
/usr/sbin/datapath query device	IBM SDD
/opt/xiv/host_attach/bin/xiv_devlist	IBM XIV HAK
/usr/sbin/ccs --getconf	Linux cluster
/usr/sbin/ccs --checkconf	Linux cluster
/usr/sbin/clustat	Linux cluster

Table 1-5 Privileged commands on Linux (Continued)

Command	Required when this is installed ...
/sbin/crm_mon -s	Linux cluster
/sbin/pcs status corosync	Linux cluster
/sbin/pcs cluster cib	Linux cluster
/sbin/pcs status xml	Linux cluster
lsnrctl status	Oracle
${oracleHome}/cemutlo -w	Oracle RAC
${oracleHome}/olsnodes -l	Oracle RAC
${oracleHome}/crsctl check crs	Oracle RAC
${oracleHome}/crsctl check crsd	Oracle RAC
${oracleHome}/crsctl check cssd	Oracle RAC
${oracleHome}/crsctl check evmd	Oracle RAC
${oracleHome}/crsctl status resource -l	Oracle RAC
${oracleHome}/crsctl status resource -f	Oracle RAC
${oracleHome}/crsctl status resource -t	Oracle RAC
${oracleHome}/crsctl query css votedisk	Oracle RAC
/sbin/service	EMC ScaleIO MDM
/bin/cat	
/bin/ls	
/usr/sbin/ping	

Privileged commands on AIX

Table 1-6 lists the privileged commands on AIX.

Table 1-6 Privileged commands on AIX

Command	Required when this is installed ...
/usr/sbin/powermt display	EMC PowerPath

Table 1-6 Privileged commands on AIX (Continued)

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

Table 1-6 Privileged commands on AIX (Continued)

Command	Required when this is installed ...
${oracleHome}/cemutlo -w	Oracle RAC
${oracleHome}/olsnodes -l	Oracle RAC
${oracleHome}/crsctl check crs	Oracle RAC
${oracleHome}/crsctl check crsd	Oracle RAC
${oracleHome}/crsctl check cssd	Oracle RAC
${oracleHome}/crsctl check evmd	Oracle RAC
${oracleHome}/crsctl status resource -l	Oracle RAC
${oracleHome}/crsctl status resource -f	Oracle RAC
${oracleHome}/crsctl status resource -t	Oracle RAC
${oracleHome}/crsctl query css votedisk	Oracle RAC
/usr/bin/uname -L	
/bin/cat	
/bin/ls	

Data collection from Windows

VRA collects data from Windows hosts using Windows Management Instrumentation (WMI) queries and WMI remote command invocation. Alternatively, WMI can be activated using PowerShell - controlled by the system property “Login scanner WMI use PowerShell”, which can be found under “Collection - Admin” system properties. The latter method requires PowerShell version 5.1 or higher on the VRA server.

Windows Remote Management (WinRM) can also be used on hosts on which it is installed and configured.

The following WMI namespaces are queried:

- root/CIMV2
- root/WMI

The following WMI classes are queried:

- Under the CIMV2 namespace:
  - Win32_OperatingSystem

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

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

- powermt display, syminq, symdg list, symcg list, inq (EMC)
- inqraid $Phys -CLIWP, dlnkmgr view -lu -item all, pairdisplay -CLI -l (HDS)
- sdcli disk list, iscsicli SessionList, dsmcli lun attributes (NetApp)
- datapath query device | version (IBM DS)
- xiv_devlist -t xml -o all (IBM XIV)
- haclus -state, haclus -display -localclus, hasys -display -localclus, hasys -nodeid, hagrp -display -localclus, hares -display -localclus, hares -dep, lltstat -c, gabconfig -v, gabconfig -l, vxfenadm -d, hagrp -dep, hahb -display (Veritas cluster)
- vxprint, vxdisk list | path, vxdmpadm pathinfo, vxvol volinfo, vxlicrep (Veritas storage foundation)
- cluster /properties | /privproperties | node | node /properties | node /privproperties | group | group /properties | group /privproperties | resource | resource /properties | resource /privproperties | network | network /properties | network /privproperties | netinterface | netinterface /properties | netinterface /privproperties | /quorum | /listshares | /listnetpriority (Microsoft cluster 2008 R2 or earlier)
- import-module FailoverClusters -ErrorAction Stop, Get-Cluster* (Microsoft cluster 2012 or higher)
- HbaCmd HbaAttributes, HbaCmd PortAttributes, HbaCmd GetDriverParams, HbaCmd TargetMapping, scli -Z, fcinfo /details, fcinfo /ports /details, fcinfo /mapping (HBA)

- diskpart list disk, diskpart list volume, diskpart detail disk, diskpart detail volume (LVM)
- mx matrix status, mx config list, mx server listsoftware, mx application status, mx mfs status, mx mni listinstances, mx vfs status, mx vfs_share status, sandiskinfo -ial, sandiskinfo -val, sandiskinfo -fal, sandiskinfo --dynvol_properties -al (HP PolyServe)
- ping

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

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

  WMI is based on DCOM. You can further limit the TCP ports allowed for WMI/DCOM communication, but this requires a significantly more complex provisioning process.

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

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

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

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

  For more information, see “Network and other environmental recommendations” on page 75.
- Make sure CIFS connectivity is permitted (port 445). CIFS is used to run vendor-native commands. Note that there is a workaround to avoid the usage of CIFS. However, there are certain trade-offs for doing that. For more details, contact Veritas Technical Support.

- On hosts using Symmetrix or CLARiion, make sure you meet the following requirements:
  - VRA requires that at least one of the utilities — PowerPath, SymCLI, or inq (V7.3-487 and later) — is installed on each host.
  - If none of these utilities is available on a certain host, install the free EMC inq utility.
- On hosts using HDS, make sure you meet the following requirements:
  - VRA requires that at least one of the utilities (HDLM or inqraid) is installed on each host.
  - If neither of these utilities is available on a certain host, install the free HDS inqraid utility.
- On hosts using NetApp, make sure that at least one of the utilities (sdcli or dsmcli) is installed.
- On hosts using IBM DS, make sure that IBM SDD is installed.
- On hosts using XIV, make sure that IBM XIV Host Attachment Kit (HAK) is installed.

Data collection from databases

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

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

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

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

MS SQL Server	1433
Oracle	1521
Sybase	5000
DB2	50000

- Provide the user account profile on each database (existing or specifically created for VRA). Specific information for each database vendor is provided below. You should use the same ID on all databases, although you may use different IDs per platform or per individual database.

- Make sure that the required specific rights have been granted to each user account. For more details, refer to the subsections below by vendor.
- Make sure that IP connectivity through JDBC is available between the VRA application server and each database server. You can use the default port number or specify another port.

Data collection from Oracle

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

Table 1-7 Data collection from Oracle

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

Table 1-7 Data collection from Oracle

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

Table 1-7 Data collection from Oracle

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

Note: For suggestions about secure ways to grant these read-only privileges, see Appendix A, "Methods for secure privilege provisioning" on page 83.
You should grant all the rights specified above for each instance. This makes provisioning much simpler, and, if you need these views in the future, ensures seamless compatibility.

Note: You can collect data from Oracle databases using Oracle Enterprise Manager instead of connecting to each database. See "Data collection from Oracle Enterprise Manager" on page 48.

Data collection from Sybase

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

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

Data collection from Microsoft MS SQL Server

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

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

In addition, VRA collect data from MS SQL Server Always On

Availability Groups by querying the following tables:

- sys.availability_replicas
- sys.dm_hadr_availabilty_replica_states
- sys.dm_hadr_database_replica_states
- sys.availability_groups
- sys.availability_databases_cluster

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

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

In order to scan MS SQL Always On successfully, make sure that the user

configured to scan this SQL instance has 'View server state' permission.

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

The following tables/procedures are executed/queried:

- sysfiles
- sysfilegroups
- sp_spaceused

Data collection from IBM DB2

VRA collects data from DB2 by connecting to each DB2 database through JDBC and running the following queries:

- syscat.tablespaces
- sysibmadm.dbpaths
- sysibmadm.tbsp_utilization

VRA also collects data by connecting to each DB2 database through JDBC and running the following procedures:

- sysproc.env_get_inst_info
- sysproc.db_get_cfg
- sysproc.snap_get_db (or sysproc.snapshot_database in versions before 10)
- sysproc.snap_get_container (or sysproc.snapshot_container in version before 10)
- sysproc.snap_get_tbsp_part

Make sure that you meet the following DB2-specific requirements:

- On version 9 and above, the user profile assigned to VRA should be a member of the DB2 SYSMON group.
- On version 8.1 fix-packs 7 and above, the user profile assigned to VRA should be a member of the DB2 SYSMON group.

- On version 8.1 fix-packs 1-6, the user profile assigned to VRA should be a member of the SYSCTRL, SYSMAINT, or SYSADM group.

Data collection from Veritas Cluster Server

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

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

Data collection from Veritas Volume Replicator

VRA collects data from Veritas Volume Replicator (VVR) by connecting to each server and querying the following privileged read-only VVR commands. A non-root user needs sudo, PowerBroker (pbrun), UPM (pmrun), CA Access Control (seSUDO), super, or similar software to execute the following privileged commands:

- /usr/sbin/vradmin printvol
- /usr/sbin/vradmin printrvg
- /usr/sbin/vradmin -g <diskgroup> -l repstatus <rvg>

Data collection from Linux Cluster Server

VRA collects data from RedHat Cluster Suite version 7 and higher
by connecting to each host and querying the following read-only privileged commands:

- /sbin/crm_mon
- /sbin/pcs

VRA collects data from RedHat Cluster Suite version 6 by connecting to each host and querying the following read-only privileged commands:

- /usr/sbin/ccs
- /usr/sbin/clustat.

RedHat is using the ricci service to operate the cluster. You should manually run the command "ccs -h localhost --checkconf", and if it requests a password - please enter the ricci service password. This is a one-time action required for every cluster node. Alternatively, the authentication can be done by read-only access of the below file on each cluster node:

- /etc/cluster/cluster.conf

Data collection from EMC Control Center

Data collection is based on opening a Java Database Connectivity (JDBC) connection to the EMC ControlCenter (ECC) repository (StorageScope sts view). The ECC repository and Storage Scope must be installed on the same server.

Read-only select queries are used to obtain data.

Note: For a list of read-only queries, see Appendix A, "Methods for secure privilege provisioning" on page 83.

Note: For configuring ECC 6 with JDBC Secure Sockets Layer (SSL), see "Configuring VRA for ECC scanning over JDBC SSL (Oracle)" on page 90.

When you set up data collection, do the following:

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

Mail server configuration

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

You provide the following details for your email server:

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

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

Network and other environmental recommendations

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

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

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

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

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

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

Summary of ports and protocols used by VRA

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

Table 1-8 VRA ports and protocols

From	To	Port/Protocol	Description
Desktop of VRA users	VRA server	8080/HTTP or 8443/HTTPS	Access to VRA web UI
VRA server	HDS HiCommand/ HP CommandView	2001/HTTP	Connecting to HDS/ HP management consoles
DSCLI proxy	IBM DS 6000/8000 arrays	1720, 1722, 1750, 8451-8455	DSCLI
SMCLI proxy	IBM DS 3000/4000/5000 arrays	3260, 49152-65536	SMCLI
XCLI proxy	IBM XIV/A9000 arrays	7778	XCLI
VRA server	NetApp filers/clusters	80/HTTP or 443/HTTPS	Connecting to filers
VRA server	V7000/SVC arrays	22/SSH	Connecting to V7000/SVC
VRA server	InfiniBox arrays	443/HTTPS	Connecting to InfiniBox
VRA server	VPLEX Management Server	443/HTTPS	Connecting to VPLEX Management Server
VRA server	RecoverPoint appliance	22/SSH	Connecting to RecoverPoint
VRA server	Isilon arrays	22/SSH	Connecting to Isilon
VRA server	DataDomain arrays	22/SSH	Connecting to DataDomain
VRA server	XMS	80 / HTTP or 443 / HTTPS	Connecting to XtremIO
VRA server	Unisphere	80 / HTTP or 443 / HTTPS	Connecting to Unity

Table 1-8 VRA ports and protocols

From	To	Port/Protocol	Description
VRA server	ScaleIO Gateway	443 / HTTPS	Connecting to ScaleIO
VRA server	HMC	22/SSH	Connecting to HMC
VRA server	HCP	9090/HTTPS	Connecting to HCP
VRA server	UNIX servers	22/SSH	Connecting to UNIX servers
VRA server	Windows servers	All TCP, UDP 135-9/ WMI	Connecting to Windows servers
VRA server	Windows servers	80/5985 (default) and 445/WinRM	Connecting to Windows servers
VRA server	Oracle instances ip/vip	1521 (default)/JDBC	Connecting to Oracle
VRA server	SQL Server instances ip/vip	1433 (default)/JDBC	Connecting to MS-SQL
VRA server	IBM DB2 databases ip/vip	50000 (default)/JDBC	Connecting to DB2
VRA server	Sybase instances ip/vip	5000 (default)/JDBC	Connecting to Sybase
VRA server	Oracle Enterprise Manager (OEM)	1521 (default)/JDBC	Connecting to Oracle Enterprise Manager
VRA server	Oracle GoldenGate Monitor	1521 (default)/JDBC	Connecting to Oracle GoldenGate Monitor
VRA server	NetApp OnCommand Unified Manager Core (DFM)	8088/HTTP or 8488/HTTPS	Connecting to DFM
VRA server	vCenter/vSphere	443/SOAP	Connecting to VMware vCenter
VRA server	VMware SRM	9007/SOAP	Connecting to VMware SRM
VRA server	Zerto Virtual Manager	9669/HTTPS	Connecting to Zerto Virtual Manager

Table 1-8 VRA ports and protocols

From	To	Port/Protocol	Description
VRA server	VMware NSX Manager	443/HTTPS	Connecting to VMware NSX Manager
VRA server	System Center Virtual Machine Manager	All TCP, UDP 135-9/ WMI OR 80/5985 (default) and 445/WinRM	Connecting to Microsoft System Center Virtual Machine Manager
VRA server	Cisco UCS Manager	80 / HTTP or 443 / HTTPS	Connecting to Cisco UCS Blade Server
VRA server	HP OneView	443 / HTTPS	Connecting to HP Blade System
VRA server	F5 BIG-IP	443 / HTTPS	Connecting to F5 Load Balancer
VRA server	Cisco MDS SAN Fabric Switch	22/SSH	Connecting to Cisco MDS switches
VRA server	Cisco Data Center Network Manager	443/SOAP	Connecting to Cisco Data Center Network Manager
VRA server	Brocade SAN Fabric Switch	22/SSH	Connecting to Brocade switches
VRA server	Brocade Network Advisor	443/HTTPS	Connecting to Brocade Network Advisor
VRA server	HP Virtual Connect	22/SSH	Connecting to HP Virtual Connect
VRA server	Mail server	25 (default)/SMTP	Sending emails from VRA
VRA server	EMC Control Center server	1521/JDBC 1575/JDBC SSL	Connecting to ECC RAMBDB views
VRA server	Active directory LDAP host	389 (default)	Optional when using active directory for users

VRA Agent

VRA agents provide an alternative to the standard agent-less scan. By installing an agent on target systems, users may scan target hosts and storage CLI servers without entering host credentials into VRA and without maintaining a privileges policy through access control solutions such as sudo and PowerBroker. The two scan methods (agent-less and agent-based) co-exist within VRA and the user may select to use the most appropriate method for each scanned host and storage CLI server.

VRA agent has the following design principles:

- Data is collected in read-only mode; It does not change your configuration.
- Only standard well-known vendor and OS commands are used.
- Low footprint. None of the scripts or commands put a noticeable load on servers or the network. The only significant computation is performed on the VRA application server and the VRA Oracle repository which are dedicated computing resources.
- Single host port required (configurable per host).
- Master/Collector to agent communication is digitally signed using RSA with unique, per-customer keys.
- Master/Collector and agent communication is encrypted using AES.

Requirements

Before installing and agent, verify that the following pre-requisites are met on the target host:

- Java Runtime Environment (JRE) is installed (version 8 and above).
- The agent installation requires the root (id 0) account for Unix and Linux, and an Administrator for Windows.
- A port on the host is available and can be assigned to the agent for communication with the master/collector server. TCP communication originating from the master or collector servers to the target host should be allowed on the designated port.
- Ensure at least 500MB of free disk space is available.

Supported Platforms and Limitations

VRA agent can be installed on the following operating systems:

- AIX (4 and above)
- HP-UX (11 and above)
- Linux (RedHat Advanced Server / SUSE)
- Solaris (8 and above)
- Windows (XP / 2000 / 2003 / 2008 / 2012)

VRA agent can be used to accomplish the following scan tasks:

- Data collection for the host on which the agent is installed
- Agents can be installed on storage CLI hosts and can be used to collect data from EMC Symmetrix, EMC VNX/CLARiiON, HP 3PAR, IBM DS Series and IBM XIV

Important Notes:
Credentials for the storage systems are still required when data is collected using an agent on a storage CLI host (except for EMC Symmetrix).
Databases and management consoles cannot be scanned using an agent.

Methods for secure privilege provisioning

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

sudo on UNIX hosts

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

When you use sudo version 1.6.9 and above, VRA must run sudo -E to preserve environment variables.

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

Note: There is a possibility to configure the whole script to be executed with elevated privileges rather than per specific commands. Contact Support for details.

SymCLI proxy

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

All UNIX flavors

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

Note: In order to enable using sudo for these commands, please set the following system properties to “yes” (the default is “no”)

Use sudo for the cat command
Use sudo for the ls command

These properties can be found under Collection system properties.

Servers that use EMC storage

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

Servers that use HDS/HP XP storage

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

Servers that use NetApp storage

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

Servers that use IBM storage

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

Servers with Veritas SF and Cluster

username ALL= NOPASSWD: /usr/sbin/vxdisk path, /usr/sbin/vxdisk list*, /usr/sbin/vxdmpadm list*, /sbin/lltstat -c, /sbin/lltstat -nvv, /sbin/gabconfig -v, /sbin/gabconfig -l, /sbin/vxfenadm -d, /opt/VRTSvcs/bin/haclus -state, /opt/VRTSvcs/bin/haclus -display*, /opt/VRTSvcs/bin/hagrp -display*, /opt/VRTSvcs/bin/hagrp -dep, /opt/VRTSvcs/bin/hares -display*, /opt/VRTSvcs/bin/hares -dep, /opt/VRTSvcs/bin/hasys -display*, /opt/VRTSvcs/bin/hasys -nodeid, /opt/VRTSvcs/bin/hahb -display*, /usr/sbin/vradmin printvol, /usr/sbin/vradmin printrvg, /usr/sbin/vradmin * repstatus *

Servers with Oracle database

username ALL= NOPASSWD: /bin/cat */listener.log, /bin/cat *alert_*.log, /bin/cat */listener.ora

Additional for Solaris servers

username ALL= NOPASSWD: /usr/sbin/fcinfo

Additional for Linux servers

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

Additional for AIX servers

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

Additional for HP-UX servers

username ALL= NOPASSWD: /usr/sbin/cmviewcl *, /usr/sbin/cmquerycl *

UNIX Privilege Manager

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

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

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

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

### Data

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

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

shellProfile           = "restricted";      # If you want to allow users matching this profile to
run

                                             # privilege manager shells, then this is the name of
                                             # the shell profile to include. The shell profiles
                                             # are copied to <poicydir>/pro

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

### List of profile members ###

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


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


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

### List of profile commands ###

# Authorized commands - these commands are executed as the authUser defined above
authCmds={                                             # Description
"/usr/sbin/vgdisplay",                                 # Linux commands
"/usr/sbin/lvdisplay *",
"/usr/sbin/pvdisplay",
"/sbin/vgdisplay",
"/sbin/lvdisplay *",
"/sbin/pvdisplay",
"/sbin/multipath -l",
"/sbin/scsi_id *",
"/bin/raw -qa",
"/usr/bin/raw -qa",
“/usr/sbin/ccs --getconf”,
“/usr/sbin/ccs --checkconf”,
“/usr/sbin/clustat”,
“/sbin/crm_mon -s”,
“/sbin/pcs status corosync”,
“/sbin/pcs cluster cib”,
“/sbin/pcs status xml”,
"/usr/bin/fcinfo",                                     # Solaris commands
"/usr/es/sbin/cluster/utilities/cldisp",               # AIX commands
"/usr/es/sbin/cluster/diag/clver",
"/usr/sbin/cmviewcl *",                                # HPUX commands
"/usr/sbin/cmquerycl *",
"/usr/sbin/vxdisk list *",                             # VXDMP on All Unix Flavors
"/usr/sbin/vxdisk path",
"/usr/sbin/vxdmpadm list *",
"/sbin/lltstat -c",                                    # VCS on All Unix Flavors
"/sbin/lltstat -nvv",
"/sbin/gabconfig -v",
"/sbin/gabconfig -l",
"/sbin/vxfenadm -d",
"/opt/VRTSvcs/bin/haclus -state",
"/opt/VRTSvcs/bin/haclus -display -localclus",
"/opt/VRTSvcs/bin/hasys -display -localclus",
"/opt/VRTSvcs/bin/hasys -nodeid",
"/opt/VRTSvcs/bin/hagrp -display -localclus",
"/opt/VRTSvcs/bin/hagrp -dep",
"/opt/VRTSvcs/bin/hares -display -localclus",
"/opt/VRTSvcs/bin/hares -dep",
"/opt/VRTSvcs/bin/hahb -display",
"/usr/sbin/vradmin printvol",                         # VVR on All Unix Flavors
"/usr/sbin/vradmin printrvg",
"/usr/sbin/vradmin * repstatus *",
""/usr/symcli/bin/syminq *",                          # EMC Symmetrix/CLARiiON on All Unix Flavors
"/usr/symcli/bin/sympd list *",
"/usr/symcli/bin/symdg list *",
"/usr/symcli/bin/symcg list *",
"/usr/local/bin/inq *",
"/etc/powermt display *",                              # PowerPath on Solaris
"/usr/sbin/powermt display *",                         # PowerPath on AIX
"/sbin/powermt display *",                             # PowerPath on HPUX/Linux
"/HORCM/usr/bin/inqraid *",                            # HDS on All Unix Flavors
"/HORCM/usr/bin/raidqry *",
"/HORCM/usr/bin/raidscan *",
"/HORCM/usr/bin/pairdisplay *",
"/usr/local/bin/lunstat *",
"/usr/DynamicLinkManager/bin/dlnkmgr view *",
"/sbin/xpinfo *",                                      # HP on All Unix Flavors
"/sbin/spmgr display *",
"/sbin/autopath display *",
"/opt/NetApp/snapdrive/bin/snapdrive storage show *",  # NetApp on All Unix Flavors
"/opt/netapp/santools/sanlun lun show *",
"/usr/sbin/datapath query device", # IBM DS on All Unix Flavors
"/usr/sbin/pcmpath query device", # IBM DS on AIX
"/opt/xiv/host_attach/bin/xiv_devlist",# IBM XIV on All Unix Flavors
"/bin/cat *",          # All Unix Flavors
"/bin/ls *",
"/usr/symcli/bin/symcfg list *",                       # SymCLI Proxy on All Unix Flavors
"/usr/symcli/bin/symdev list *",
"/usr/symcli/bin/symaudit list *",
"/usr/symcli/bin/symevent list *",
"/usr/symcli/bin/symdisk list *",
"/usr/symcli/bin/symmaskdb list *",
"/usr/symcli/bin/symaccess list *",
"/usr/symcli/bin/symcli -def"
};

processProfile();
```

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

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

Suggested Oracle grant provisioning

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

Suggested MS SQL Server grant provisioning

The following SQL creates a login with the appropriate permissions:

```
USE [master];

CREATE LOGIN vrauser
    WITH PASSWORD = N'vrauserpassword',
    DEFAULT_DATABASE = [master],
    DEFAULT_LANGUAGE = [us_english],
    CHECK_POLICY = OFF,
    CHECK_EXPIRATION = OFF;

CREATE USER vrauser FOR LOGIN vrauser;

GRANT VIEW ANY DEFINITION TO vrauser
GRANT VIEW SERVER STATE TO vrauser;
GRANT SELECT ON sys.sysaltfiles TO vrauser;
GO

EXEC sp_MSforeachdb '
    USE ?
    CREATE USER vrauser
    GRANT CONNECT TO vrauser
    '
```

Note: You will need to GRANT CONNECT for each new database added.

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

Queries used to scan EMC ECC

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

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

Configuring VRA for ECC scanning over JDBC SSL (Oracle)

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

1. Configure the policy to use the JDBC-SSL protocol.
   See "Configuring an SSL policy" on page 91.
2. Copy the wallet file that holds the authentication information from the ECC Repository server to the server running VRA.
   See "Copying the wallet file from the ECC Repository server" on page 93.
3. Configure credentials in the standard manner.

Configuring an SSL policy

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

To configure an SSL policy

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

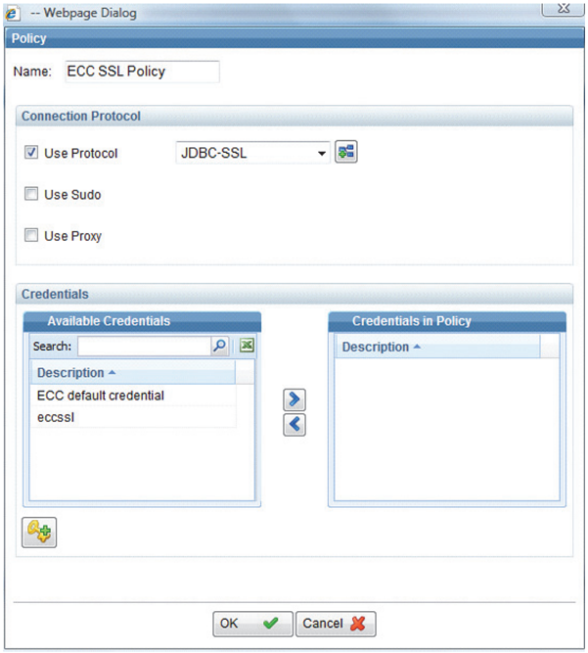

-- Webpage Dialog

Policy

Name: ECC SSL Policy

Connection Protocol

☑ Use Protocol JDBC-SSL

☐ Use Sudo

☐ Use Proxy

Credentials

Available Credentials

Description
ECC default credential
eccssl

Credentials in Policy

Description

2 In the **Name** field, enter a policy name.

3 Check the **Use Protocol** checkbox and then select **JDBC-SSL** in the adjacent drop-down list.

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

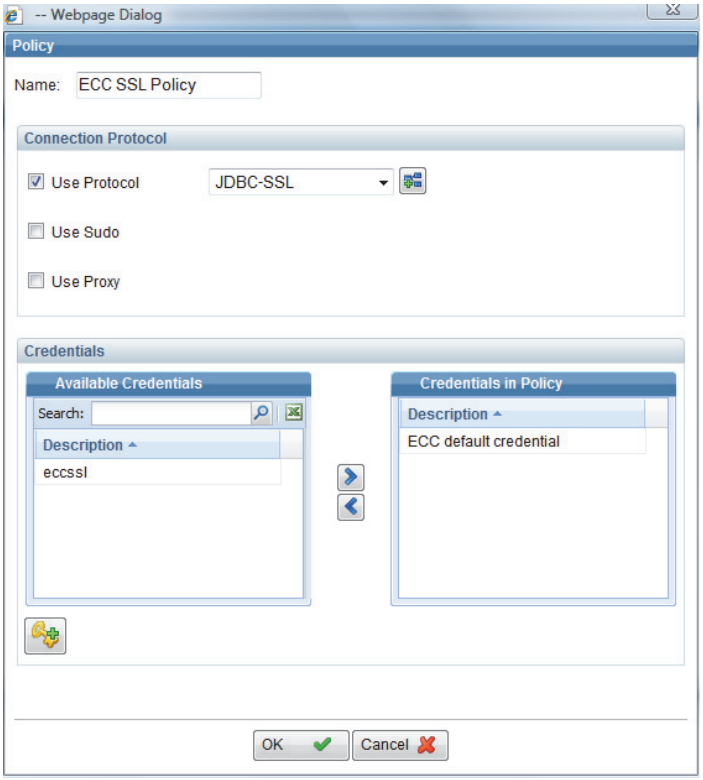

-- Webpage Dialog

Policy

Name: ECC SSL Policy

Connection Protocol

☑ Use Protocol JDBC-SSL

□ Use Sudo

□ Use Proxy

Credentials

Available Credentials

Search:

Description ▲
eccssl

Credentials in Policy

Description ▲
ECC default credential

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

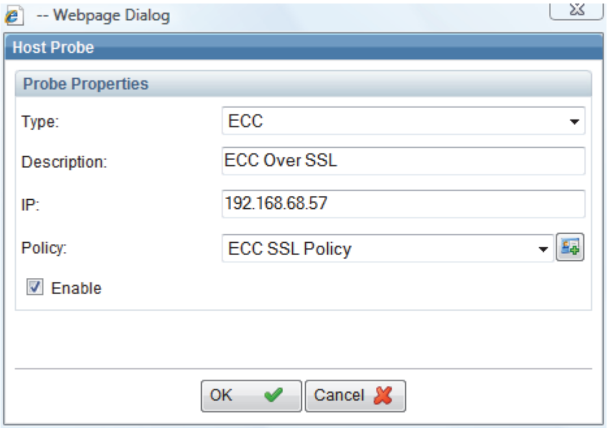

-- Webpage Dialog

Host Probe

Probe Properties

Type: ECC

Description: ECC Over SSL

IP: 192.168.68.57

Policy: ECC SSL Policy

☑ Enable

Copying the wallet file from the ECC Repository server

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

To copy the wallet file from the ECC Repository server

1. Copy the cwallet.sso file from the ECC Repository server using the following path on the ECC Repository server:
   ECC_install_root_path\dbSafe\cwallet.sso
   For example, if ECC is installed under E:\ECC, the wallet file location is E:\ECC\dbSafe\cwallet.sso.

2. Create a wallet directory in the VRA Home folder using the following path:
   C:\VRA\security\wallets\ECC_Repository_IP\cwallet.sso
   Let's look at an example, in which the ECC IP address is 192.168.68.57. Do the following in the order presented:

   - Create the necessary directory.

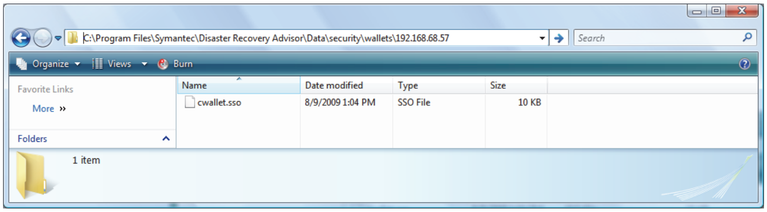

C:\Program Files\Symantec\Disaster Recovery Advisor\Data\security\wallets\192.168.68.57

Name	Date modified	Type	Size
cwallet.sso	8/9/2009 1:04 PM	SSO File	10 KB

1 item

   - Copy the previously downloaded cwallet.sso file.

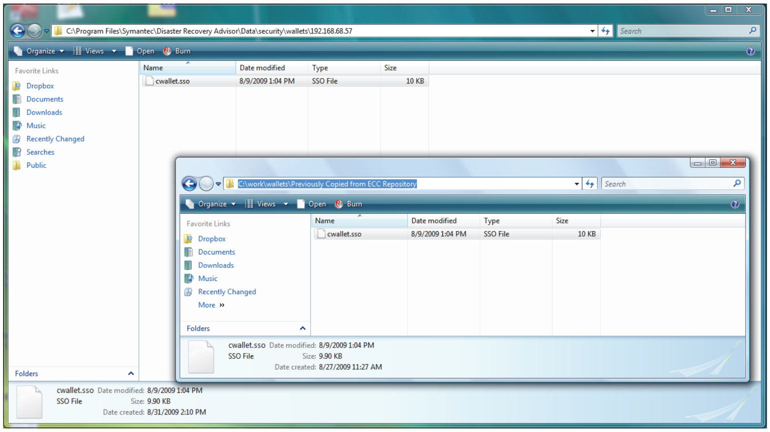

C:\Program Files\Symantec\Disaster Recovery Advisor\Data\security\wallets\192.168.68.57

Name	Date modified	Type	Size
cwallet.sso	8/9/2009 1:04 PM	SSO File	10 KB

cwallet.sso Date modified: 8/9/2009 1:04 PM
SSO File Size: 9.90 KB
Date created: 8/31/2009 2:10 PM

C:\work\wallets\Previously Copied from ECC Repository

Name	Date modified	Type	Size
cwallet.sso	8/9/2009 1:04 PM	SSO File	10 KB

cwallet.sso Date modified: 8/9/2009 1:04 PM
SSO File Size: 9.90 KB
Date created: 8/27/2009 11:27 AM