



Confidence in a connected world.

High Availability for Databases

Protecting Oracle® Databases with Veritas™ Cluster Server

Eric Hennessey, Director
Technical Product Management

High Availability for Databases

Protecting Oracle Databases with Veritas™ Cluster Server

Contents

Introduction	4
The challenge of Oracle database availability	4
Veritas Cluster Server from Symantec	5
Key features	6
Components of Veritas Cluster Server with Oracle support	6
Clusters	6
Agents	7
How Veritas Cluster Server manages an Oracle database server	7
Anatomy of an Oracle failover with Veritas Cluster Server	8
How Veritas Cluster Server with the Oracle agent addresses customer challenges	9
The need to control costs	10
N+1 “roaming spare” clusters	10
N to N clusters	11
The need to reduce complexity	12
The need to manage dependent, multi-tier applications	12
The need to address disaster recovery requirements	13
Managing multiple clusters	14
Additional benefits of Veritas Storage Foundation™ HA for Oracle	15
Summary	16
About Symantec	18

High Availability for Databases: Protecting Oracle Databases with Veritas™ Cluster Server

Introduction

In years past, IT solutions were deployed to support back-end business functions such as inventory management, accounting, and human resources. In today's world, the IT solution increasingly *is* the business. These days the repositories for nearly all business application data are high-end database servers, and the majority of those are provided by Oracle.

This paper will discuss the various options available for providing high availability (HA) for single-instance Oracle databases, as well as dependent applications. It will demonstrate how Veritas Cluster Server from Symantec™ can provide cost-effective high availability for Oracle databases without the added cost and complexity of Oracle RAC.

Note that in those cases where Oracle RAC is chosen as the ultimate architecture, Veritas Cluster Server is still key because it is an integral component of the Veritas Storage Foundation for Oracle® RAC product, one of the most robust and “bulletproof” solutions in the market. Storage Foundation for RAC adds significant value to RAC deployments by keeping the Oracle private interconnect highly available and load-balanced and, at the same time, providing split-brain prevention using SCSI-III PRG I/O fencing. In addition, it provides raw-like performance through ODM and the Clustered File System (CFS), making management of all RAC environment components much easier and more reliable than any other alternative. The bottom line is that, whether the customer is deploying single-instance Oracle or RAC—or both—Veritas Cluster Server is a vital part of the solution.

The challenge of Oracle database availability

Customers who employ Oracle databases are, more often than not, using them for mission-critical applications. Their primary concern, therefore, is to keep those databases as available as possible. However, the concern is complicated by the need to:

- Control costs
- Reduce complexity
- Address dependent, multi-tier applications
- Address disaster recovery requirements

When asked about database availability, Oracle's answer will invariably be RAC, as this is the flagship Oracle product. In many cases, however, deploying RAC as a simple availability solution significantly increases the complexity and cost to a data center just to provide an increase in the level of availability that could easily be achieved with a simpler solution. Oracle RAC

High Availability for Databases: Protecting Oracle Databases with Veritas™ Cluster Server

provides value in environments where parallel access to multiple instances of the same database is required for application performance or very specific zero recovery-time failure-handling scenarios. Traditional high-availability failover solutions do not provide parallel access, but can provide a very short recovery time objective (RTO) at significantly reduced cost and complexity. In the case of Veritas Cluster Server, the RTO can be driven even lower by using the Veritas Cluster File System in place of the standard Veritas File System. The value-add of Veritas Cluster File System is that the file systems on the “standby” server can already be mounted so that, in case of a failover, the steps and time associated with mounting file systems can be avoided. The result: quicker application failover times.

Operating system (OS) vendors and several third-party software vendors provide classic high-availability failover products. Most of these products were designed for—or operate best in—a traditional active-passive failover configuration. In this configuration the “active” node or server runs a database instance, and the “passive” failover node stands idle, waiting for failure. This configuration does help sell additional servers, but it is not a cost-effective solution from the customer standpoint. Shifting to Oracle RAC utilizes the second server, eliminating the problem of idle hardware; however, this configuration drastically increases the costs and complexity of the Oracle software.

Because most companies do not operate just a single database but rather a large number of databases that support mission-critical applications, the cost-effective solution for providing solid high availability without the increased cost and complexity of Oracle RAC is Veritas Cluster Server. The Veritas Cluster Server high-availability solution from Symantec™ is designed to provide cost-effective “roaming spare” or “no spare” failover configurations for mission-critical Oracle environments.

Veritas Cluster Server from Symantec

Veritas Cluster Server from Symantec is an industry-leading cross-platform high-availability clustering solution, according to an IDC report. With its application-centric approach to availability, Veritas Cluster Server provides customers with a high degree of flexibility in determining the application failover behavior best suited to their business requirements.

Veritas Cluster Server provides high availability to mission-critical applications by managing each application within a construct known as a *service group*. Each service group contains the specific Oracle components needed to manage an Oracle instance, along with all of the network

High Availability for Databases: Protecting Oracle Databases with Veritas™ Cluster Server

and storage resources that the database requires. The service group is the single unit of failover within a Veritas Cluster Server cluster.

Key features

Veritas Cluster Server is a feature-rich, best-in-class high-availability clustering solution. Key features include:

- Intelligent failover logic with service group workload management
- Out-of-the-box support for most major applications
- Automated disaster recovery (DR)
- Disaster recovery and local high availability failover test suite
- Management console for multi-cluster and multi-site management

Components of Veritas Cluster Server with Oracle support

Clusters

A single Veritas Cluster Server cluster consists of multiple systems connected in various combinations to shared storage devices. Veritas Cluster Server monitors and controls applications running in the cluster, and can restart applications in response to a variety of hardware or software faults.

A cluster is defined as all systems that share a common cluster configuration and utilize a common interconnect network. The Veritas Cluster Server cluster interconnect consists of redundant physical Ethernet interconnects, generally over two or more dedicated private networks. The communications layer carries heartbeats between systems within the cluster, as well as membership and state change information.

Clusters can have from one to 32 member systems, or “nodes.” Within a single Veritas Cluster Server cluster, all member nodes must run the same operating system family. For example, a Solaris™ cluster would consist entirely of Solaris nodes—likewise with Linux®, AIX®, HP-UX®, and Windows® clusters. Within a given cluster, multiple versions of the operating system can be deployed to support specific upgrade scenarios.

Applications can be configured to run on specific nodes in the cluster. Storage is configured to provide access to shared application data for the systems that are hosting the application. In

High Availability for Databases: Protecting Oracle Databases with Veritas™ Cluster Server

that respect, the actual storage connectivity will determine where applications can be run. Nodes sharing access to storage are “eligible” to run an application.

Agents

Veritas Cluster Server agents handle the start, stop, and monitoring of all resources contained within a service group. Agents receive instructions on when to start, stop, or monitor a resource from the Veritas Cluster Server engine, and the agents return the results of those actions to the engine. Bundled with the Veritas Cluster Server package is a collection of agents to manage the storage and network resources that an application managed by Veritas Cluster Server requires.

Veritas Cluster Server also ships with agents to control all common system functions, such as file systems and network addresses. Additional agents are provided for out-of-the-box support for most enterprise applications, such as databases, application servers, and Web servers. This includes complete out-of-the-box (no customization required) support for Oracle database and network components.

How Veritas Cluster Server manages an Oracle database server

Veritas Cluster Server support for Oracle consists of two agents: one for the listener, and one for the database instance itself. The listener agent starts and stops the appropriate listener by issuing the standard *lsnrctl* start or *lsnrctl* stop command. The state of the listener is monitored by checking the process table for the *tnslsnr* process, and as an optional level of more detailed monitoring, by issuing the *lsnrctl* status command.

To control the Oracle instance, the agent starts and stops the instance using standard Oracle startup and shutdown commands (*sqlplus*). Preferred startup and shutdown options can be set by the customer.

For first-level (basic) monitoring, the agent verifies that required Oracle processes are running. To accomplish this, the agent checks the process table for the *ora_dbw*, *ora_smon*, *ora_pmon*, and *ora_lgwr* processes corresponding to a specific instance. For more detailed monitoring, the agent can also be set to connect periodically to the database instance and update a table, thereby verifying the instance’s ability to accept transactions.

High Availability for Databases: Protecting Oracle Databases with Veritas™ Cluster Server

Anatomy of an Oracle failover with Veritas Cluster Server

To understand the sequence of events that occurs when Veritas Cluster Server fails-over an Oracle database, it is helpful to understand how a Veritas Cluster Server service group is configured for Oracle. Figure 1 illustrates the Veritas Cluster Server console. It shows a fairly typical Oracle service group configuration in Veritas Cluster Server.

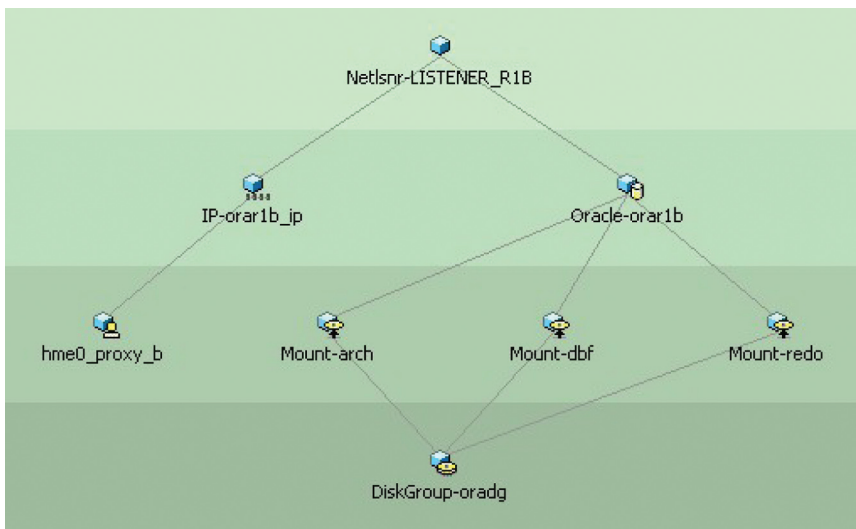


Figure 1. Oracle service group, as viewed from the Veritas Cluster Server console.

Each blue box in the figure represents a resource managed by Veritas Cluster Server within this service group. At the top is the Oracle Listener service, and beneath that on the right is the Oracle database instance itself. The remaining resources represent the storage and network resources; the virtual IP and NIC resources on the left-hand side, and the file systems and Veritas Volume Manager disk group on the right. The virtual IP address managed by Veritas Cluster Server here is the IP address configured in the listener's LISTENER.ORA file.

The lines connecting the resources depict the various dependencies between the resources, which dictate the correct start and stop sequence for all the resources in the service group. For example, one wouldn't want to attempt to start the Oracle database instance before the file systems containing its archive, database files and redo logs are mounted.

High Availability for Databases: Protecting Oracle Databases with Veritas™ Cluster Server

In a planned failover event—i.e., one in which a conscious decision has been made to move the Oracle instance from one system to another in the cluster—the following sequence of events will take place once the command is issued:

1. The Oracle listener agent stops the listener service.
2. The virtual IP address to which the listener had been bound is removed from its network interface, and at the same time, the Oracle instance is shut down.
3. All three file systems (archive, DB files, and redo logs) are unmounted.
4. The disk group containing the file systems is deported.

For the target node—the node chosen by the administrator to host the failover database—the above sequence is reversed.

Of course, not all failover events are planned. Depending on the nature of the problem that caused the failover, Veritas Cluster Server may or may not be able to gracefully shut down all the resources within the service group. In the case of a server panic or power outage, none of the resources will be brought down cleanly, and the events described above will take place in reverse sequence. What is different in this situation is that when the Oracle agent starts the Oracle database, Oracle will perform its database recovery routine to ensure transactional consistency; also, any needed file system checks on the file systems will be performed before mounting. It is vital to understand that the database recovery on a failover is identical to the normal Oracle instance recovery following any local server issue.

How Veritas Cluster Server with the Oracle agent addresses customer challenges

Beyond the simple ability to move an Oracle database instance from one system to another in response to a failure, Veritas Cluster Server also provides significant additional benefits to the customer. These include addressing the problem areas of controlling costs, reducing complexity, and providing high availability to the front-end tiers of the application stack that depend on the Oracle database (e.g. app tier, web tier, etc.).

In this section, we will explore the question of how Veritas Cluster Server addresses these customer challenges—and how it provides automated disaster recovery.

High Availability for Databases: Protecting Oracle Databases with Veritas™ Cluster Server

The need to control costs

Traditional approaches to high availability have entailed the clustering of mission-critical applications with pairs of servers in an active-passive (standby) configuration, which works well enough in a smaller environment with only one or two mission-critical applications. But in a larger environment with numerous mission-critical applications, the cost of the required hardware spares and additional software licenses is prohibitive. Veritas Cluster Server allows the deployment of large clusters with advanced failover logic that can make “intelligent” failover target selections based on server capacity and application load. With Veritas Cluster Server, all of a company’s servers that are running mission-critical applications can be grouped together in a single cluster. Then, one or two spare systems can easily be added to that cluster to handle failovers. This is called N+1 clustering, where *N* represents the number of servers with active applications and *1* indicates a single spare. In a larger cluster, this might actually be N+2 or N+3.

A variant of this cluster topology is N to N. In this topology, there is no spare server per se, but each server has some amount of spare capacity to host additional applications. The paragraphs that follow will discuss these topologies in more detail.

N+1 “roaming spare” clusters

The N+1 cluster topology has a number of active servers, plus one idle server as a spare. For this example, we will consider a three-node cluster with two Oracle instances, A and B, which are active on two separate servers, Server 1 and Server 2. Server 3 is, initially at least, the spare server.

At the initial cluster startup, Oracle A is started on Server 1 and Oracle B is started on Server 2. Server 3 is idle. If a fault occurs on Server 1, Oracle A fails-over to Server 3. When Server 1 is returned to service, there is no automatic fail-back of Oracle A, as a fail-back would introduce an additional unnecessary outage. Instead, Server 1 simply becomes the new spare. This is the meaning of the term “roaming spare.”

In larger N+1 clusters, the intelligent failover logic that Veritas Cluster Server offers with its *service group workload management* (SGWM) becomes evident. This management capability associates a numerical load value with each service group (application) and a numerical capacity value with each server. The cluster considers and evaluates these values when it selects a failover target. In the absence of an intelligent method for selecting a target server for failovers—such as

High Availability for Databases: Protecting Oracle Databases with Veritas™ Cluster Server

SGWM—in the event of multiple failures, applications tend to become stacked onto one or two servers.

N to N clusters

The N to N cluster topology does not include a spare server. All systems that host a mission-critical application are placed into the cluster, with the assumption that each server has some extra capacity to host additional applications in the event of a failure.

As with N+1 clusters, each service group (application) is assigned a numerical load value and each server in the cluster a numerical capacity value. Because the cluster has no idle servers, the cluster must be able to select failover target nodes based on load and capacity; conventional ordered list selection simply will not work.

This topology also lends itself well to server consolidation efforts, because the cluster will automatically attempt to “stack” applications on all of the cluster’s servers in a balanced fashion.

Veritas Cluster Server provides an additional construct known as *limits* (for servers) and *prerequisites* (for service groups). This permits a user to define the upper limits of a server’s resources, and the minimum amount of those resources that must be reserved for the application it hosts. For example, if a server in the cluster has 16 GB of memory installed, it can be defined with a *limit* value of “Mem=16.” Each service group in the cluster can then have a *prerequisite* value called *Mem* set to the amount of memory, in gigabytes, required by that service group. Each service group with a *Mem* value defined has that value deducted from the server’s *Mem* value. For a server to host a service group, it must have a remaining *Mem* value equal to or greater than the *Mem* value specified for that service group.

It is important to note that in both of these cluster architectures, the need to purchase spare servers is decreased: It is minimized in the case of N+1 clusters, and eliminated altogether with N to N clusters.

In contrast to this capability, other cluster products select a failover target from a simple ordered list of servers. In general, this list is set so that an application runs on one server in the cluster under normal circumstances, but will move to a designated spare upon failure. This practice is fairly adequate until subsequent failures cause the designated spare to become overloaded as multiple applications fail-over to it.

High Availability for Databases: Protecting Oracle Databases with Veritas™ Cluster Server

The need to reduce complexity

Veritas Cluster Server reduces complexity in a number of ways. Because of its consistency across OS platforms, Veritas Cluster Server reduces the administrative burden by providing a consistent clustering solution across the enterprise—so system administrators and operations staff members have only a single user interface to learn, regardless of the OS platform in use.

Because the agent for Oracle is pre-built specifically to manage Oracle databases, no custom scripting is needed to get an Oracle instance under Veritas Cluster Server control. The user must simply specify four parameters in order to cluster a given instance of Oracle. To cluster an Oracle instance, the user simply supplies to the Oracle agent the Oracle instance name (SID), the name of the Oracle user account, the Oracle home directory path, and the path to the Oracle instance's "pfile." Support for the use of Oracle's "spfile"—introduced with Oracle 9i—is also provided.

Veritas Cluster Server also ships with a fully functional Cluster Simulator for testing and training. Customers can use the Cluster Simulator to test new cluster configurations and changes prior to rolling them out in production. The Cluster Simulator also provides a perfect environment for new operator and administrator training. The Veritas Cluster Server Cluster Simulator can be installed on any Windows desktop system and is available on the Symantec corporate Web site.

The need to manage dependent, multi-tier applications

An Oracle database is rarely an island unto itself. Usually, the database is the back-end tier of a multi-tier application stack that also encompasses application and Web tiers on the front end. All too often, a customer will cluster the back-end database of a multi-tiered application stack, but neglect the upper portions of the tier. Then, when a failure occurs on a server that hosts the Oracle database, the cluster dutifully handles the failover of the database, but the upper-tier components of the application must be manually restarted because they have lost their connection with the database server. To sequence the startup and shutdown of the separate tiers correctly, these upper tiers—for example, a WebLogic application tier and an Apache web server tier—can be placed in the same cluster as the Oracle database instance with their service group (application) dependencies established when all tiers are running on the same OS platform.

In cases where the various tiers are running on separate OS platforms, which is true more often than not, Veritas Cluster Server provides a remote group agent to establish cross-cluster service group dependencies. This capability is key to multi-tier application control, and it is unique to Veritas Cluster Server.

High Availability for Databases: Protecting Oracle Databases with Veritas™ Cluster Server

The remote group agent is configured to point to a service group running on another cluster in the data center. When a failover occurs, the remote group resource will wait for the service group in the remote cluster to come online before bringing the local service group online. This enables a properly sequenced startup of applications even when the applications are running on separate OS platforms.

The need to address disaster recovery requirements

Oracle databases are often the cornerstone of an organization's most critical applications. In the event of a disaster that causes a complete data center outage, it is imperative that these applications be brought online as quickly as possible at an alternate site.

Veritas Cluster Server HA/DR enables global clustering, which extends the concepts of local high availability to wide-area disaster recovery failover. In a Veritas Cluster Server global cluster configuration, a Veritas Cluster Server cluster is deployed at each of two to four data centers. Communications are established between the clusters over an IP network, and each cluster is configured with an identical service group. Keep in mind that the primary Veritas Cluster Server cluster can have multiple service groups. In combination with the global clustering feature of Veritas Cluster Server, this capability provides flexibility in replicating the critical applications of the business—which can significantly reduce the cost of the total disaster recovery solution. The Oracle database files are replicated between storage devices at each site and the service group is capable of failing-over not only within its own local cluster in response to a localized failure, but also to another cluster in the event of complete loss of the entire data center. Veritas Cluster Server HA/DR includes agents for most enterprise storage replication technologies (e.g., EMC SRDF and IBM GlobalMirror) and Veritas™ Volume Replicator, as well as Oracle DataGuard. The idea behind the disaster recovery capabilities of Veritas Cluster Server is to protect the application, not just the data.

In order to provide regular validation that a wide-area failover will work when the time comes, Veritas Cluster Server provides a testing feature called Fire Drill that allows for non-disruptive testing of the disaster recovery configuration. When a test is run, Veritas Cluster Server creates a storage snapshot of the application data at the secondary data center and then brings the application up at the secondary data center without its network identity. In the case of an Oracle database, Fire Drill will start the database instance without its corresponding listener service, IP address, or any DNS changes. (The network components are omitted in order to prevent

High Availability for Databases: Protecting Oracle Databases with Veritas™ Cluster Server

interference with the production application running at the primary data center.) Because most disaster recovery failures stem from changes to the primary site's storage configuration that are not propagated to the secondary site, the completion of a Fire Drill test provides assurance that all required storage is being replicated and is valid.

Once testing is complete, the application is brought back down and the snapshot is released. The production application running at the primary data center is completely unaffected by the Fire Drill test at the secondary site.

In a real disaster scenario or planned application movement, the process is only slightly different. The actual database replica is mounted, and all related network components such as the Listener, IP address, and any DNS updates are started so that database users can continue to operate.

In the case of a planned application movement from one data center to another, the direction of replication is reversed so that database consistency is maintained at both sites.

Managing multiple clusters

In larger enterprises, it is common for multiple clusters that are managing Oracle databases to be deployed not just across multiple OS platforms, but across multiple data centers as well. This practice can pose an administrative challenge, especially when dependencies exist between applications running in different clusters or when applications are capable of failing-over across data centers. In response to this challenge, Symantec offers a consolidated Veritas Cluster Server Management Console, which provides full visibility into and management of all Veritas Cluster Server clusters throughout the enterprise, across multiple OS platforms and across multiple sites.

The console permits management and operational control over all Veritas Cluster Server clusters from a single "pane of glass." In addition, it further extends wide-area disaster recovery management capabilities by providing a simplified site migration function that allows all applications under Veritas Cluster Server control to be relocated from one data center to another.

A rich set of pre-defined reports, such as application uptime reports for SLA compliance, can be produced from the Management Console, and it can also provide complex metropolitan- and wide-area cluster visualizations.

The Management Console was introduced starting with the Veritas Cluster Server 5.0 release and is included as part of the base product. It can be downloaded from the Veritas Cluster Server product page at <http://go.symantec.com/vcs>.

Additional benefits of Veritas Storage Foundation™ HA for Oracle

When Veritas Cluster Server is bundled with Veritas Storage Foundation for Oracle, databases derive additional benefits in the form of increased performance, flexibility and reliability.

Storage Foundation for Oracle consists of the following components:

- **Veritas Volume Manager**—Veritas Volume Manager provides simplified disk management with striping and mirroring for improved performance and availability. This includes the ability to mirror across sites within metropolitan area cluster (MAC) distances of about 50 miles. Dynamic Multipathing (DMP) provides availability and load balancing across storage connections (host bus adapters [HBAs]).
- **Veritas File System**—Veritas File System is a high-performance, fast-recovery file system that is optimized for business-critical database applications and data-intensive workloads. Veritas File System offers online administration, letting you perform most frequently scheduled maintenance tasks (including online backup, resizing, and file system changes) without interrupting data or system availability. Veritas File System also provides support for large file systems (of more than eight exabytes in a 64-bit environment) and large files (in the exabyte range in a 64-bit environment). Veritas File System offers the following performance-enhancing features that are of particular interest in a database environment:
 - Veritas Quick I/O—Improves the throughput for Oracle databases built on the Veritas File System. It delivers raw device performance to databases, providing the administrative advantages of using file systems without the performance penalties.
 - Veritas Cached Quick I/O—Further enhances database performance by leveraging large system memory to selectively buffer frequently accessed data.
 - Veritas Concurrent I/O—Improves the performance of regular file systems without the need to extend namespaces and present the files as devices. This simplifies administrative tasks and allows relational databases that have no sequential read/write requirement (such as Oracle) to access files concurrently.
- **Veritas Enterprise Administrator**—Veritas Enterprise Administrator is the infrastructure that allows Veritas Storage Foundation for Oracle, Veritas Volume Manager, and Veritas File System information and features to be accessed through the graphical user interface (GUI).

High Availability for Databases: Protecting Oracle Databases with Veritas™ Cluster Server

Summary

Veritas Cluster Server from Symantec offers the highest degree of flexibility and cost-effectiveness in protecting Oracle databases from both localized and data center-wide failures. When deployed in an enterprise-wide environment, Veritas Cluster Server protects Oracle databases, along with their dependent applications, across multiple OS platforms and multiple locations. When Veritas Cluster Server is deployed in conjunction with Veritas Storage Foundation for Oracle, unparalleled levels of reliability, performance, and manageability are possible.

About Symantec

Symantec is a global leader in providing security, storage, and systems management solutions to help businesses and consumers secure and manage their information. Headquartered in Cupertino, Calif., Symantec has operations in more than 40 countries. More information is available at www.symantec.com.

For specific country offices and contact numbers, please visit our Web site. For product information in the U.S., call toll-free 1 (800) 745 6054.

Symantec Corporation
World Headquarters
20330 Stevens Creek Boulevard
Cupertino, CA 95014 USA
+1 (408) 517 8000
1 (800) 721 3934
www.symantec.com

Copyright © 2008 Symantec Corporation. All rights reserved. Symantec, the Symantec logo, Veritas, and Veritas Storage Foundation are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. HP and HP-UX are trademarks of Hewlett-Packard Company. IBM and AIX are trademarks of International Business Machines Corporation in the United States, other countries, or both. Linux is the registered trademark of Linus Torvalds. Microsoft and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Sun and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc., in the U.S. or other countries. Other names may be trademarks of their respective owners.
08/08 14216725