Storage Replicator for Volume Manager™

Successful Replication and Disaster Recovery

November 1999

VERITAS

BUSINESS WITHOUT INTERRUPTION™

Table of Contents

Introduction

Companies today rely to an unprecedented extent on online, frequently changing, and frequently accessed data to run their businesses. Therefore unplanned events that interfere with the availability of this data interferes with the successful running of the business. Additionally, permanent loss of parts of this data (from natural disaster or any other source) will likely have serious negative consequences for the continued viability of a business. Therefore, when data or its availability is threatened, a company must plan to recover quickly with usable data, and eliminate or minimize data loss.

To protect themselves against the fallout from data loss, many enterprises have created disaster preparedness plans that include detailed steps for the protection of data, and their recovery after a disastrous event. Successful data recovery assumes that there is an uncompromised data source with full integrity of replicated data. The data to be used in the recovery must be fully up to date, complete, and free from errors. It should allow the operation to continue to run without a significant break in continuity. To keep it removed from the impact of a disastrous event, it makes good sense for recovery data to be replicated and stored in a different geographical location than the primary data storage. Some businesses contract with an external data storage provider for this purpose, but many will want to fully implement and be in control of their own strategy. To allow for failures that concern less than an entire site, and to arrive at high availability of data on a local level, the disaster recovery plan should also include a reliable clustering and data mirroring strategy.

There is no doubt that replication is becoming a standard industry practice. While this paper focuses on disaster recovery, replication can also enable access to geographically distributed file systems and off-host data processing. Alternatives to generalized policy based software replication solutions are usually (a) labor intensive, ad hoc, and therefore error prone, (b) limited to specific storage hardware, or specific data formats, or (c) limited in delivering data currency or integrity on recovery. The labor intensive alternatives to full replication typically do not deliver the same completeness and data integrity, because they are limited as to data copied or hardware that can be used in the process.

VERITAS Software is in the business of making it possible for enterprises to do business without interruption. As the world's leading provider of application software storage management solutions, VERITAS' array of products includes Storage Replicator, an optimal data replication tool that can make an important contribution in successful data protection and recovery. In what follows, we briefly discuss some of the key concepts of data replication, explain some of the features and strengths of Storage Replicator, and show how it might be used as part of a disaster recovery plan.

Replication: Some Basic Issues

Replication, such as the term is used in the context of disaster recovery, is an automated and rules-based method for the geographical distribution of identical data to a safe location. In an optimal environment, replication will be automated and policy-based. This reduces the opportunity for human error, and minimizes the need for administrator intervention. Replication should make efficient use of resources and, after an initial replication of all data, keep WAN network traffic down by only replicating the blocks of data that actually change. The replication solution needs to be flexible enough to be easy to configure and modify for the network manager and help to optimize administrative resources.

There are two main types of replication, synchronous and asynchronous. Both have their advantages, and should be available options for the network manager. They use a different process to arrive at the same goal, and deal somewhat differently with network conditions. The performance and effectiveness of both depend ultimately on business requirements such as how soon updates must be reflected in replication, but are also strongly determined by the available bandwidth, network latency, the number of participating servers, the amount of data to be replicated, and even the geographical distance. Unlike some competing replication products, VERITAS replication solutions maintain full data integrity in both synchronous and asynchronous modes.

Synchronous replication ensures that a write update has been posted to the secondary node before it is acknowledged back to the primary application. This way, in the event of a disaster at the primary data location, data can be recovered from any surviving secondary server, because all servers share the same data state. Synchronous replication produces full data consistency and integrity, but a series of rapid data writes may confront a level of network latency or limited bandwidth that can reduce application performance. Synchronous replication is most effective in application environments with low update rates that need all sites to always reflect a common data state. The possible degradation in update response time experienced by the application can be mitigated by adding network bandwidth and reducing network round-trip time between each primary-secondary pair. If a delay in updates between the primary and any secondary location is not acceptable, synchronous replication would be the preferred strategy.

Under asynchronous replication, while application updates are immediately reflected at the primary, they are persistently queued to be forwarded to each secondary. When the writing application experiences temporary surges in update rate, this queue may grow. Unlike with synchronous replication, the writing application does not suffer from the response time degradation caused by each update incurring the cost of a network round-trip. At all times, the queue is being drained as fast as the available network bandwidth allows. During periods when the update rate is less than the available network bandwidth, this queue drains faster than it grows, allowing the secondary data state to catch up rapidly with that of the primary. The improvement in response time comes at the low cost of the data state at the secondary being slightly behind the data state at the primary during peak update times.

Let's imagine a disaster strikes during a period of peak update activity while asynchronous replication is in effect. Following the disaster, the application, for example a database, is successfully started using the data available at the secondary. It is now possible that a group of updates most recently committed to the primary is not reflected in the data state at the secondary. This is because of the possible lag, already mentioned, between the primary and secondary data states. The network manager should have the option of configuring how much of a lag is tolerable in the event of a disaster. Asynchronous replication would then never let the lag exceed this configured maximum.[1]

It is worth noting that in both asynchronous and synchronous replication the data states at each secondary faithfully track data states at the primary. This is called 'write-order fidelity.' Without it, in the event of a disaster the data state at a secondary might be inconsistent with itself and could make it impossible for an application like a database to use this data successfully. Write order fidelity as applied to any replication process means that the order of updates within a group of replicated data volumes is consistently preserved in all secondary copies. Updates are made to both the log and data spaces of a database management system in a fixed sequence. The log and data space are usually in different volumes, and the data itself can be spread over several additional volumes. This is a vital concern: if write order fidelity is not maintained, a database application at the primary site, managing several volumes of data, may not be successfully recoverable with the secondary data. A well designed replication solution needs to consistently safeguard write order fidelity. This may be accomplished by a logical grouping of data volumes so that the order of updates within that group is preserved across all secondary copies of these volumes.

Its ability to queue updates persistently and hold them at the primary for later transmission makes asynchronous replication able to better deal with temporary outages of the network or the secondary site. Synchronous replication, however, due to its requirement that each update be reflected at the secondary before it can be acknowledged, handles outages by making the update in question fail at the primary. Network managers may want to configure replication in a "soft synchronous" mode, where replication is synchronous under normal circumstances, but converts to asynchronous during a temporary outage. In this mode, after the outage passes and the secondary catches up, replication reverts to synchronous. Depending on specific needs, most network managers desiring synchronous behavior will probably use the "soft synchronous" variant for maximum continuity.

Let's take a look at how VERITAS has successfully addressed all of these concerns in its Storage Replicator for Volume Manager.

[1]When considering replication solutions, it is important to know that not all vendors use this "store and forward" model to maintain write order fidelity for an asynchronous replication link. That omission could compromise data integrity at the secondary location. Again, *VERITAS solutions always maintain complete data integrity in both synchronous and asynchronous modes*.

Storage Replicator for Volume Manager: An Optimal Disaster Recovery Tool

VERITAS Storage Replicator for Volume Manager (SRVM) belongs to the **online** product family of integrated solutions that enable operations to thrive in business without interruption and enjoy continuous high availability of their mission critical data. SRVM works as a fully integrated module within VERITAS Volume Manager, the industry leading, highly popular online data storage management solution. Volume Manager, used in over 100,000 enterprises worldwide, is highly regarded for the reliability of its operation. Storage Replicator benefits from the robustness, ease of use and high performance consequent to its deep integration with Volume Manager. At the same time, it maintains the integrity and manageability of Volume Manager data. SRVM will be effective in any volume management based application environment using raw partitions or a file system. It was built to be the optimal replication tool for deployment in an effective disaster recovery strategy. Any application, even with existing data, can be configured to use SRVM transparently.

Unlike most competing products, SRVM is not dependent on any specific storage hardware platform. It also has complete database management system support including Oracle, Sybase and any other commercial database system. Businesses are therefore free from any restrictions exercised by hardware or software vendors. Depending on setup and configuration, SRVM will perform comprehensive server or application specific volume replication in both synchronous and asynchronous replication modes, described above.

Storage Replicator for Volume Manager is the only product of its kind that can scale to support up to 32 secondary data storage sites. SRVM is able to work with symmetric configurations on a network, where one single machine can function as a primary site for some data sets and as a secondary data site for other data sets. Storage Replicator for Volume Manager elegantly handles one to one, one to many, many to many, and many to one data replication configurations, all of which need to be available for today's global enterprise.

As is characteristic for the entire VERITAS product family, Storage Replicator for Volume Manager gives the network manager a very high level of control and transparency. All administrative tasks, such as installation or modification of primary or secondary data sets, can be performed online. SRVM uses a simple command line interface to create replicated data sets. The continuity of business operations and availability of data will not be negatively impacted while adjustments are made to replication policies or secondary servers added. Recovery after a failure event can be configured to be mostly automated, but gives the administrator complete control depending on the operation's changing needs.

While operations may have their own definitions for these events, a temporary outage or transient failure is very different from a full failure or disastrous event. A wide-area failover under the wrong conditions may seriously impact enterprise productivity. With SRVM, the network administrator can define the parameters for when to implement an actual failover to a hot site, or when to treat an event as a transient failure defined by time and configuration variables.

SRVM can replicate over any IP network, LAN or WAN. It does not specifically require a network dedicated to itself. It is resilient to temporary network outages and failures. Its error handling capabilities will be alerted to a network outage through a failed heartbeat communication attempt. If it becomes necessary to migrate the application from the primary to the secondary site, either because of disaster at the primary or for any other reason, SRVM allows the application to be successfully restarted using the data available at the secondary.

While replication goes on, the copy of the data at the secondary is constantly undergoing change outside the control of any application at the secondary that may be attempting to read it. In order for an application at the secondary to be able to safely read data volumes that are currently undergoing replication, SRVM includes a synchronization capacity called In-Band Control (IBC). It lets a writing application at the primary tell reading applications at a secondary when they can safely read the replicated data. IBC helps to increase the consistency and availability of geographically distributed data, and allows an application to apply its own consistency protocols.

In order to have an application read and write data at the secondary, SRVM provides the ability to break off a point-in-time copy of the replicated data, as long as the secondary data volumes were configured with an adequate number of mirrors. The broken-off copy can then be used by an application at the secondary to do backups, data mining etc., while replication continues on the intact mirrors at the secondary.

Depending on available bandwidth, excessive lag between primary and secondary as well as primary log overflow may be a liability in asynchronous replication. Storage Replicator for Volume Manager can easily recover from primary log overflow and allows automatic control of excessive lag between primary and secondary sites. A feature called latency protection guards the network from having a secondary site fall too far behind in updating replicated data. Latency protection gives the network administrator the option to set a high watermark to keep the secondary from an excessive write lag. When the threshold is reached, all update activity is held until the update backlog has reached a preset level.

SRVM will always maintain consistent write order fidelity at the secondary sites. The replicated volume groups in Storage Replicator bundle related data volumes together, and make it possible to perform a successful recovery from a secondary site in both asynchronous and synchronous modes.

Storage Replicator for Volume Manager extends true high availability into disaster recovery practices. VERITAS takes continuous improvement of its products and practices very seriously. Recent efforts have very much focused on the possibilities of the integration of solutions with each other to leverage their strengths. SRVM can be cluster aware to maintain local high availability, as well as integrated with wide-area failover/migration solutions. SRVM can be a cornerstone to a solid high availability and disaster recovery strategy.

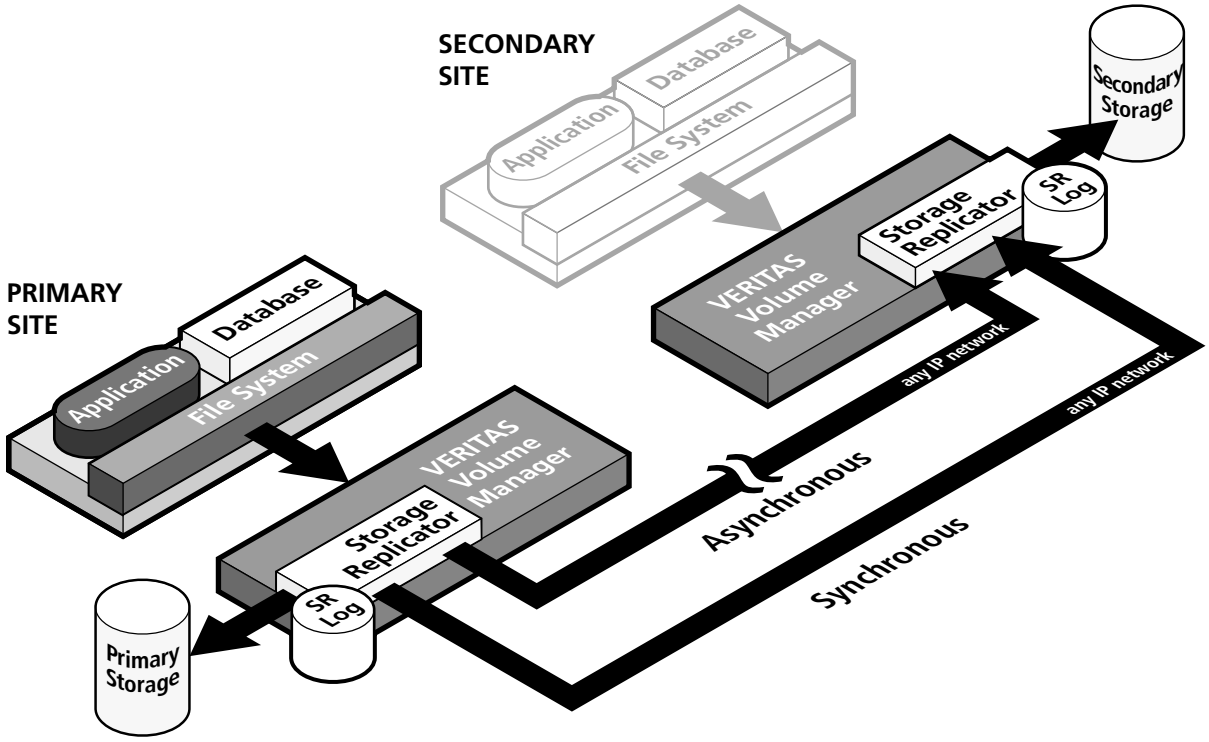

Figure 1. SRVM Architecture

Two Possible Application Scenarios

To illustrate how Storage Replicator for Volume Manager can be used in a real-life situation, here are two possible scenarios.

Consider the case of **Company A**, a Wall Street business with enterprise hosts that need to replicate to a secondary location in Philadelphia. Company A's main business critical application is its Oracle database. At the Wall Street location they have accomplished high availability by setting up a clustered environment with VERITAS Cluster Server (VCS). They then arrange for another VCS pair at the designated Philadelphia hot site for data replication. Company A uses SRVM for replication, but also runs an SRVM agent on their VCS Clusters to enjoy highly available replication services. Ordinarily, Company A does not want to experience any lag at all in their critical financial data updates, and for that reason they use the synchronous application mode. However, they also configured SRVM to enable it to switch from synchronous into asynchronous mode (or soft synchronous) when a temporary or transient failure occurs. This way, they can benefit from the buffering of updates in asynchronous replication, and still maintain data integrity, though there may be a small lag in the update process. Once the temporary failure has been resolved, the replication link will dynamically return to full synchronous mode as soon as the secondary catches up. Using this configuration, Company A can seamlessly migrate their critical Oracle application or failover to the hot site, if a true disaster occurs. Business productivity then continues after successful quick migration to the secondary site.

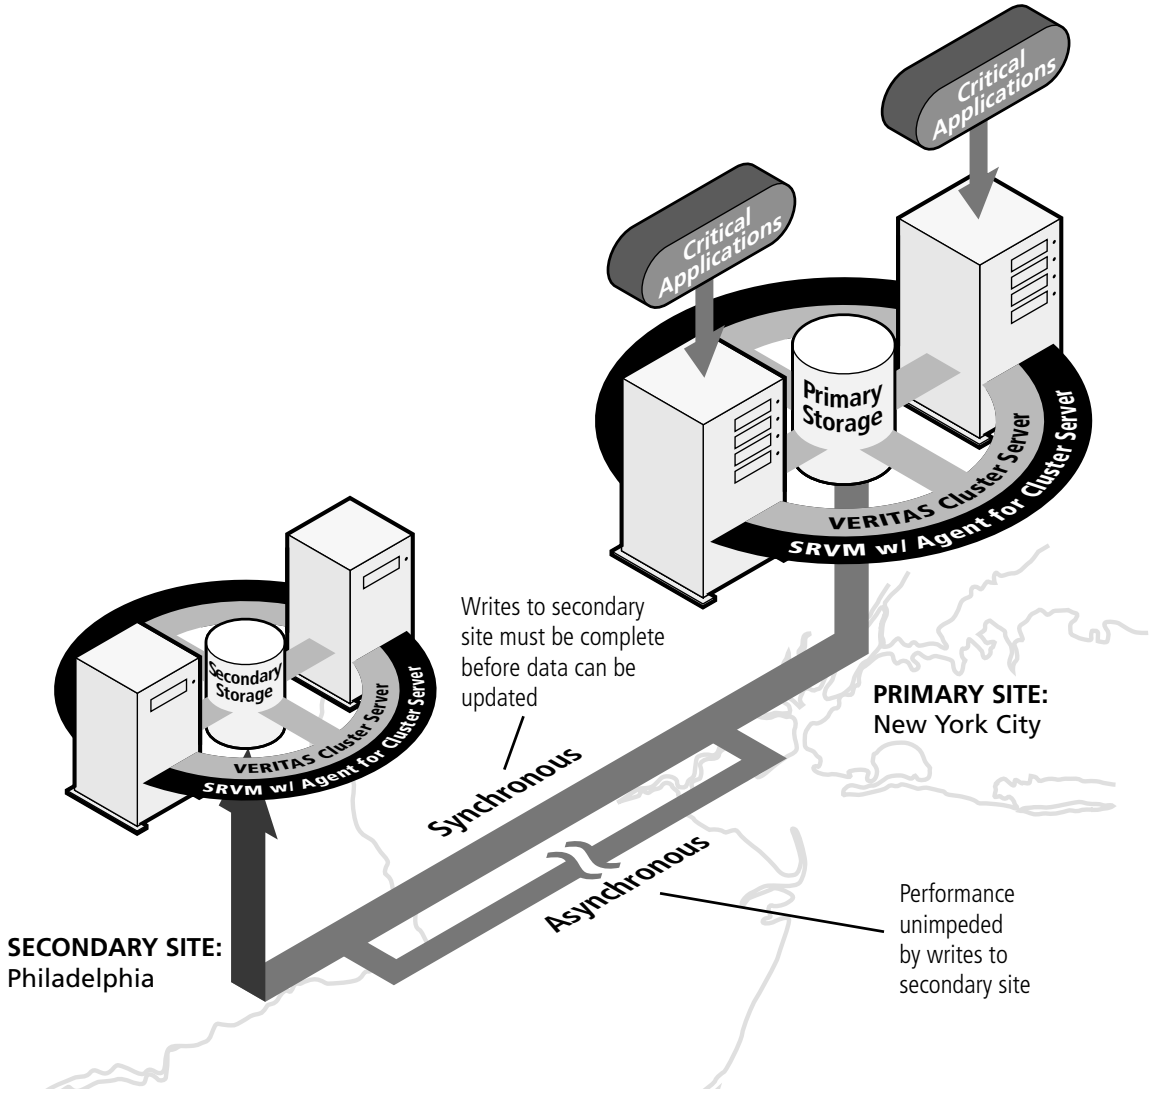

Figure 2. Company A – Synchronous Replication with Clustered Secondary

Company B is a manufacturing enterprise. Their primary location is in Detroit, and their secondary site is in Nashville. They have several mission critical applications, including financial data and enterprise resource planning (ERP) on an SAP/Oracle database. They also use VERITAS Cluster Server locally, together with the SRVM agent to assure highly available replication. They have a single enterprise server at the secondary site. They use SRVM in asynchronous replication mode, because they want to maximize the performance at the primary location, and are willing to allow a temporary lag in update writing for that. The slight update lag (usually measured in milliseconds) is deemed inconsequential, and preferable to a possible performance degradation because of network latency or limited bandwidth. Data integrity will still be entirely undiminished. Company B uses latency protection to set the write thresholds to avoid undue discrepancies between primary and secondary sites.

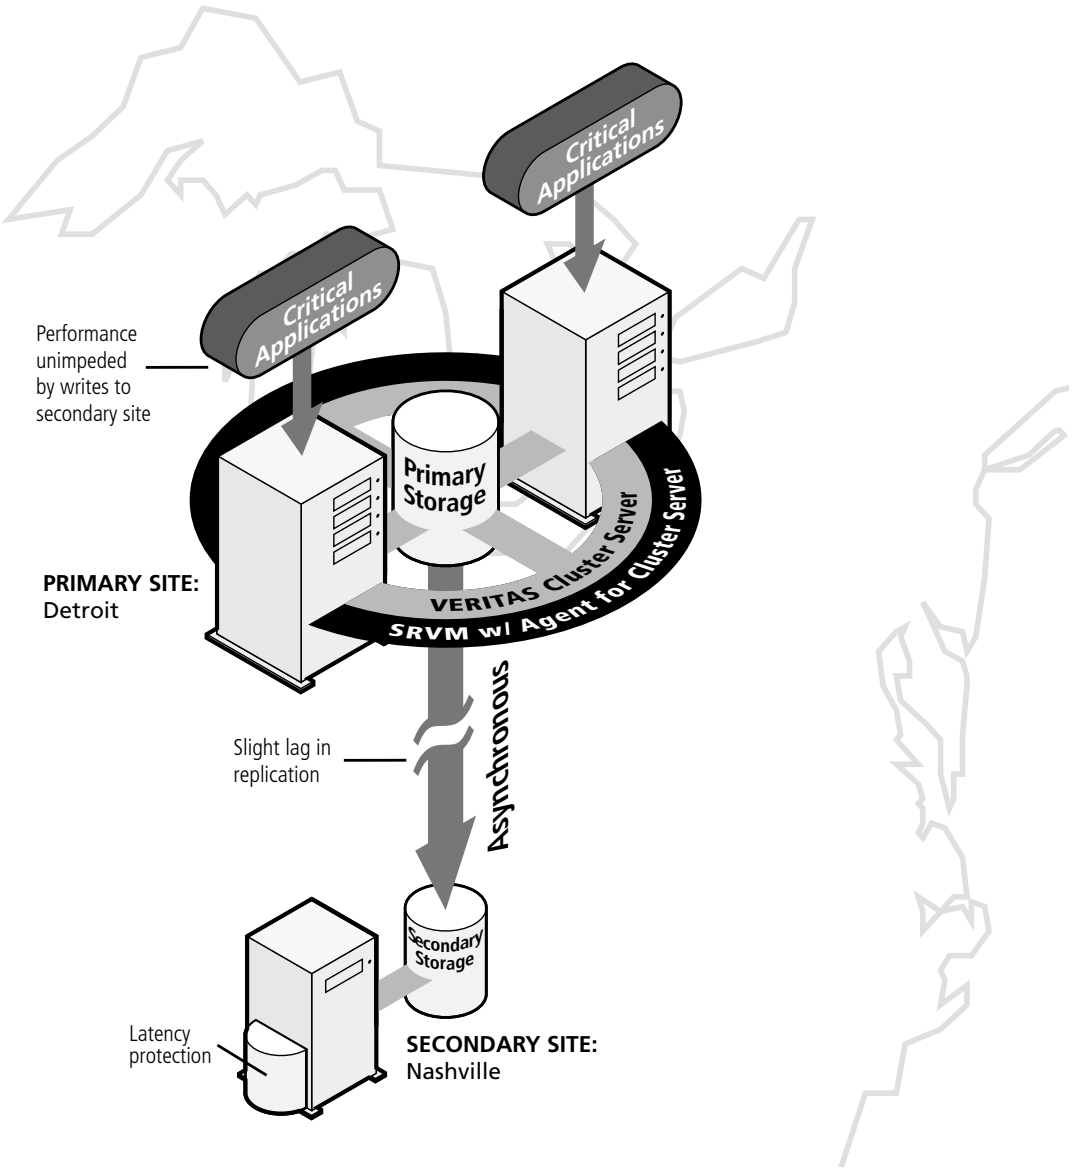

Figure 3. Company B – Asynchronous Replication to a Single Host

One can vary and complicate similar scenarios in many possible ways, and always be able to make a strong case for SRVM's effectiveness. SRVM, together with Volume Manager, is the most flexible, effective and easy to use data replication based disaster recovery tool known to the industry.

Conclusion

Storage Replicator for Volume Manager facilitates data replication with full integrity. In doing so, it fulfills one of the main prerequisites for effective disaster preparedness and recovery for a geographically distributed enterprise. With SRVM, the most current business data will always be available at all locations that need to access them. In its integration with VERITAS Volume Manager, it is the most reliable and highly performing replication solution for mission critical environments on the market today. Together with VERITAS Cluster Server, it allows for high availability at primary and secondary locations and leverages the full strength of VERITAS' experience and vision to enable business without interruption. Its exceptional manageability has significantly helped its worldwide acceptance, resulting in a steadily growing community of satisfied users and receiving outstanding reviews from industry analysts. As VERITAS introduces upgrades and enhancements for this product in its practice of continuous improvement, SRVM will be likely to continue to play a central role in the disaster recovery plans of more and more enterprises.

VERITAS Software
Corporate Headquarters
1600 Plymouth Street
Mountain View, CA 94043

North American Sales Headquarters
400 International Parkway
Heathrow, FL 32746
800-327-2232 or 407-531-7501
407-531-7730 Fax

Global Locations

United Kingdom
0800-614-961 or
44-(0)870-2431000
44-(0)870-2431001 Fax

France
33-1-41-91-96-37
33-1-41-91-96-38 Fax

Germany
49-(0)69-9509-6188
49-(0)69-9509-6264 Fax

South Africa
27-11-448-2080
27-11-448-1980 Fax

Australia
1-800-BACKUP
61-(0)2-8904-9833 Fax

Hong Kong
852-2507-2233
852-2598-7788 Fax

Japan
81-3-5532-8217
81-3-5532-0887 Fax

Malaysia
603-715-9297
603-715-9291 Fax

Singapore
65-488-7596
65-488-7525 Fax

China
011-8610-62638358
011-8610-62638359 Fax

Electronic communication

E-Mail:
sales.mail@veritas.com

World Wide Web:
http://www.veritas.com

90-00858-910 • NT12-SRVMWPR-9900