Symantec™ Replicator Option

Data replication across multivendor storage for cost-effective disaster recovery

Data Sheet: Disaster Recovery

Overview

Symantec™ Replicator Option provides organizations with a comprehensive solution for heterogeneous data replication. When combined with Symantec Storage Foundation, Replicator Option enables cost-effective replication of data over IP networks. This gives organizations an extremely flexible, storage hardware–independent alternative to traditional array-based replication architectures. Replicator Option provides the flexibility of block-based continuous replication with Volume Replicator and file-based periodic replication with File Replicator. This flexibility lets organizations choose virtually any combination of storage devices on any major operating system, providing a consistent, easy-to-manage disaster recovery solution.

In addition to providing one of the most flexible replication architectures available, Volume Replicator is tightly integrated with the industry-leading application availability software—Symantec Cluster Server. This combination provides an automated, reliable disaster recovery solution that reduces recovery time for data and applications alike.

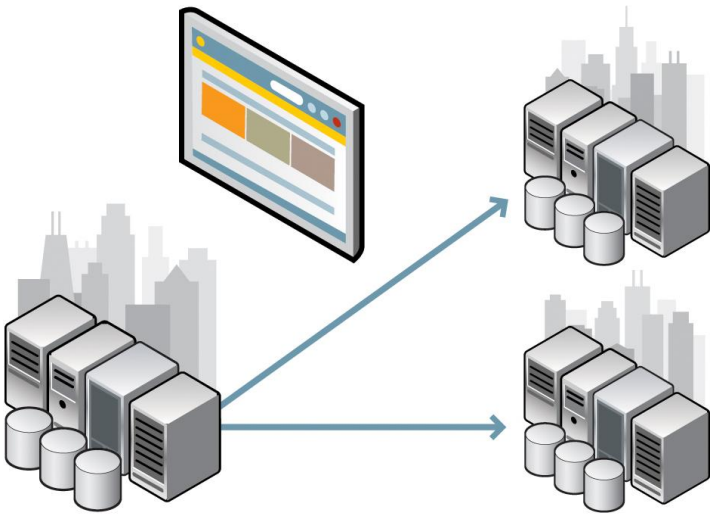

Figure 1. Volume Replicator enables high – performance data

Key Features

- **Replication over any distance**—Replicates data over any distance without performance impact to applications
- **Zero recovery point objective (RPO)**—Ensures zero data loss, without sacrificing performance, using Bunker Replication
- **Storage-independent replication**—Replicates between heterogeneous storage devices, enabling tiered storage strategies as part of a disaster recovery plan
- **Multiple site replication—**Replicates data simultaneously from primary site to multiple secondary sites with each site having its own RPO
- **Cross-platform and thick-to-thin data migration**—Enables administrators to migrate data easily across platforms or data centers and reclaim unused thin storage at primary and secondary sites
- **Efficient use of network bandwidth**—Minimizes bandwidth utilization through compression and bandwidth throttling, and maximizes replication throughput through bulk transfer of data
- **Efficient use of processor—**Minimizes processor utilization by batching replication updates and secondary logging of updates
- **Centralized management and reporting**—Allows administrators to centrally monitor and manage multiple replicated data sets across multiple sites with Veritas Operations Manager
- **Data and database consistency protection**—Protects data consistency at all times through the use of persistent disk-based replication logs
- **Automated disaster recovery**—Automates site-to-site failover for quick, reliable recovery of critical applications

Replication over any distance

Disaster recovery plans require data to be available at another site in case the production site goes down. Volume Replicator

enables synchronous and asynchronous data replication over IP networks to provide disaster recovery capabilities over any distance, without compromising performance or data consistency. With support for up to 32 secondary targets per application or server, Volume Replicator makes it possible to concurrently replicate the same data volumes in synchronous and asynchronous mode, giving organizations the ability to adapt replication strategies to match any number of data center locations.

Zero RPO

Volume Replicator enables a unique Bunker Replication configuration that combines the zero data loss advantages of synchronous replication with the long-distance performance advantages of asynchronous replication. This gives organizations maximum protection for critical applications over any distance and provides a cost-effective alternative to multi-hop or STAR configurations from array vendors. Bunker Replication is integrated with Cluster Server to provide fully automated, zero-data-loss recovery during application failures.

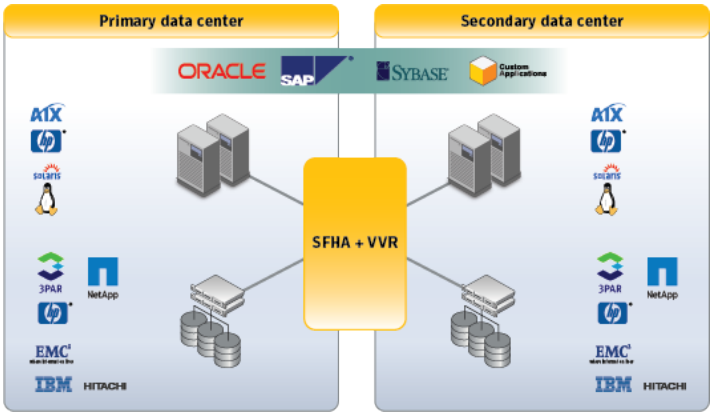

Figure 2. Volume Replicator is a proven replication solution

Storage-independent replication

Building a data replication strategy can be expensive, as most array-based data replication requires that you use the same array at the disaster recovery site. Unlike proprietary, inflexible storage hardware replication approaches, Volume Replicator offers customers greater flexibility in choosing any mix of storage area network (SAN)-based architectures, including thin storage. Replicator Option replicates data over existing IP networks and enables replication among the major storage hardware platforms, eliminating vendor-specific limitations and allowing organizations to choose appropriate storage investments based on application priority, not storage compatibility. This means users may choose any major storage vendor at the disaster recovery site, thereby reducing hardware costs.

Cross-platform and thick-to-thin data migration

Volume Replicator offers simplified cross-platform data migration that enables automated, transparent, online replication of data between heterogeneous server architectures (for example, Oracle® Solaris™ to IBM® AIX®). The prospect of a data center migration can be overwhelming, considering the incompatibilities between storage systems and server technologies and the demand for around-the-clock availability. Whether for permanent migration to a new environment, or for the operational convenience of “off-host” processing, the ability to move data between computing platforms of different types increases the flexibility of enterprise IT operations significantly. Volume Replicator, coupled with Storage Foundation Portable Data Container technology, reduces the time and resources required to migrate data between dissimilar platforms.

Replicator Option is also an extremely cost-effective solution with which to manage the risks and complexity of a data center migration, and dramatically minimizes application downtime. Using Replicator Option, customers can keep applications online as data is automatically migrated to a brand-new site. Replicator Option is a proven replication solution for heterogeneous hardware environments, allowing large enterprises to perform complex data center migrations or hardware refreshes, while keeping applications online and optimizing overall costs.

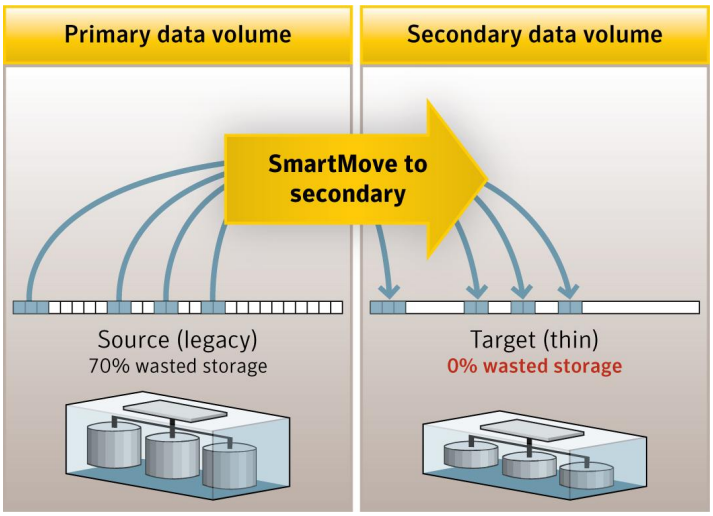

Figure 3. Using host file system knowledge, SmartMove for Volume

Replicator Option automatically reclaims all unused storage when using replication to migrate online from thick to thin storage. SmartMove for Volume Replicator leverages the host file system's knowledge of used and unused blocks to streamline the amount of resources spent during volume replication and copy only used blocks to thin storage. By replicating only the used blocks, SmartMove for Volume Replicator also significantly reduces CPU, bandwidth, and storage needs for the initial synchronization operation. Volume Replicator also offers the flexibility of automatically reclaiming the unused thin storage at the secondary site through integration with the Thin Reclamation capabilities of Storage Foundation.

Efficient use of network bandwidth and processor

A challenge with most data replication technologies is the network resources that are required for effective data transfer. Volume Replicator includes controls to reduce the impact that replication can have on scarce network resources. Through efficient volume-level replication based on actual application writes, Volume Replicator keeps wide area network (WAN) traffic to a minimum by replicating only the data that actually changes. This solution increases existing bandwidth efficiencies through asynchronous replication with robust logging capabilities, allowing organizations to model bandwidth requirements based on average application activity rather than peak activity.

Bandwidth efficiencies include compression of the replicated data to reduce consumed bandwidth and differential-based resynchronization to reduce the time and bandwidth required to migrate back to a primary site following a disaster. For the most critical environments, Volume Replicator also includes bandwidth-throttling capabilities on a per-application basis to reduce application contention for limited network resources. This ensures that critical applications have the necessary network resources, even in a bandwidth-constrained environment.

To increase asynchronous replication, throughput data is replicated to a disaster recovery (DR) site in bulk, reducing Volume Replicator processor overhead and roundtrip network latency.

Centralized management and reporting

For organizations that require the replication of a large number of critical applications, Symantec provides centralized management of all replication instances to ensure that replication manageability scales with data center environments. Leveraging Operations Manager, Replicator Option can be configured, monitored, and managed across multiple data sets and multiple operating systems through a single interface. This capability reduces the time required for initial replication configuration and deployment and improves efficiencies in the ongoing management of a large number of replicated applications.

Data and database consistency protection

Through the use of persistent disk-based replication logs, Volume Replicator maintains data consistency between primary and secondary data sets in synchronous and asynchronous modes of replication. By maintaining write-order fidelity, Volume Replicator ensures consistent restarts of critical applications and databases in virtually any operating environment. Tightly integrated with the database functionality of Storage Foundation, Volume Replicator maintains the consistency of Oracle, Oracle RAC, IBM DB2®,

Sybase®, Microsoft® SQL Server®, Microsoft Exchange, and other enterprise database management systems throughout replication. Volume Replicator even protects data consistency during temporary or extended network outages, which is an absolute requirement for long-distance replication over WANs. With uninterrupted consistency protection, organizations can be confident that replicated data sets will provide the high levels of business continuity required for critical operating environments.

Automated disaster recovery

Replicating data is only one aspect of disaster recovery. How do you recover the applications themselves? Manual recovery often leads to errors and requires personnel who may not be available during a disaster. The full integration of Volume Replicator and Cluster Server provides the most powerful disaster recovery automation available for data center applications. This solution enables organizations to monitor all applications and associated replication jobs in a multisite framework, as well as automate the process of failover/failback between sites. In the event of a failure at any of the monitored sites, Cluster Server will automatically alert administrators, control the shift of replication roles to the secondary site, mount data volumes, restart critical applications, and redirect client traffic, drastically reducing total recovery time for maximum business continuity. In addition to automated recovery, Cluster Server and Volume Replicator offer noninvasive disaster recovery testing. Using the Fire Drill feature enables disaster recovery testing without ever bringing primary production systems offline, allowing organizations to test disaster recovery more frequently than ever before.

Additional features

- **Replicate over IPv6 links**—Offers the flexibility to run replication in mixed environments; a single Replicator Option host can have IPv4 links as well as IPv6 links
- **Integrated snapshots—**Uses in-band control messaging and initiates remote snapshots based on local replication controls
- **Replication of Oracle RAC—**Supports replication of shared storage resources in conjunction with Oracle RAC and Cluster File System implementations
- **Simple configuration**—Enables easy configuration and administration with a distributed CLI framework or through Operations Manager

Related products

- **Symantec Storage Foundation**—Storage Foundation maximizes storage efficiency, availability, agility, and performance across operating systems, virtualization technologies, and storage hardware.
- **Symantec Cluster Server**—Cluster Server delivers high availability and disaster recovery for business-critical services running across your physical and virtual infrastructure. It includes automated application failover to disaster recovery sites, and a Fire Drill feature for noninvasive disaster recovery testing.
- **Symantec File Replicator**—File Replicator enables cost-effective periodic replication of data over IP networks, giving organizations an extremely flexible storage-independent data availability solution for disaster recovery and off-host processing. With flexibility in scheduling the replication intervals to match business requirements, File Replicator tracks all updates to the File System and replicates these updates consistently at the end of the configured time interval. File Replicator leverages data deduplication provided by the file system to reduce the impact that replication can have on scarce network resources. File Replicator 6.0 is supported on Red Hat® Linux.

Supported operating systems

- HP-UX®
- IBM AIX
- Microsoft Windows®
- Oracle Enterprise Linux
- Oracle Solaris
- Red Hat Linux

More Information

Visit our website

http://enterprise.symantec.com

To speak with a Product Specialist in the U.S.

Call toll-free 1 (800) 745 6054

To speak with a Product Specialist outside the U.S.

For specific country offices and contact numbers, please visit our website.

About Symantec

Symantec protects the world's information, and is a global leader in security, backup, and availability solutions. Our innovative products and services protect people and information in any environment – from the smallest mobile device, to the enterprise data center, to cloud-based systems. Our world-renowned expertise in protecting data, identities, and interactions gives our customers confidence in a connected world. More information is available at www.symantec.com or by connecting with Symantec at go.symantec.com/socialmedia.

Symantec World Headquarters

350 Ellis Street
Mountain View, CA 94043 USA
+1 (650) 527 8000
1 (800) 721 3934
www.symantec.com

21213903-4 12/13