



Getting the Most from your Data Protection Solution

A practical roadmap for
comprehensive data protection

*Bob Baird, Senior Solutions Architect
Symantec Corporation*

Getting the Most from your Data Protection Solution

Contents

Introduction	4
Data protection problems and challenges	5
Data protection problem symptoms, and the repercussions of doing nothing	8
Anatomy of a comprehensive data protection solution	10
Benefits of a comprehensive solution	10
Symantec data protection services	13
The Symantec technical solution	17
Conclusion	19
Glossary	20

Getting the Most from your Data Protection Solution: A practical roadmap for comprehensive data protection

Introduction

This paper describes why companies need a new data protection approach and describes the approach being adopted by Symantec. Companies frequently need help with some parts of their IT environment but do not necessarily need full outsourcing. There are several ways to support companies in a flexible, cost-effective manner. This paper presents the Symantec view of data protection services, discusses why data protection services are important, lays out general characteristics for a comprehensive data protection solution, and outlines Symantec's approach.

Data protection is preparation for and recovery from data emergencies. Data emergencies, in the context of this discussion, are corruption and damage mainly resulting from operational mishaps or disastrous events. Backup, monitoring, and replication contribute to recovery. Each task is only part of a data protection solution. A data protection solution also involves best practices, services, and technology. Figure 1 illustrates the relationship of data protection to major operational tasks. As might be implied, data protection is an underlying foundation for disaster recovery and high availability. High availability solutions mitigate single points of failure at a given site. Disaster recovery solutions mitigate multiple points of failure at a given site or the complete outage of an entire site. However, data protection solutions also mitigate the impact of operational mishaps and intrusions that damage data or make data recovery necessary.

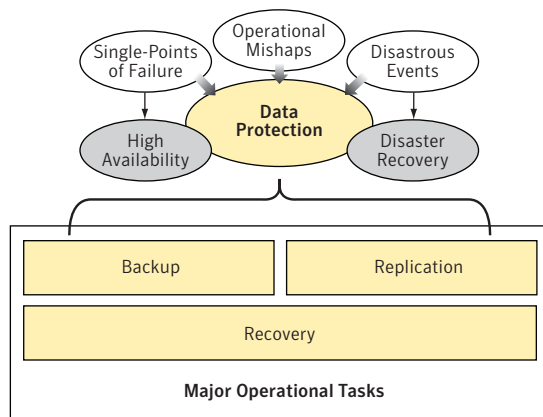


Figure 1. Major operational tasks related to data protection

Getting the Most from your Data Protection Solution: A practical roadmap for comprehensive data protection

Data protection problems and challenges

To understand data protection in today's business environment, it is necessary to place it in the context of business trends. At one time; online production operations ran during first shift; batch update and reporting operations ran during second shift; and backup and maintenance operations ran during third shift. As more applications share the same body of data for different purposes and from locations scattered across numerous time zones, off-hours is becoming a vestige of the past. Essentially, it is daytime somewhere all the time. Old data protection approaches can barely solve today's problems and will not solve tomorrow's problems. Figure 2 illustrates the many opposing forces challenging data protection organizations.

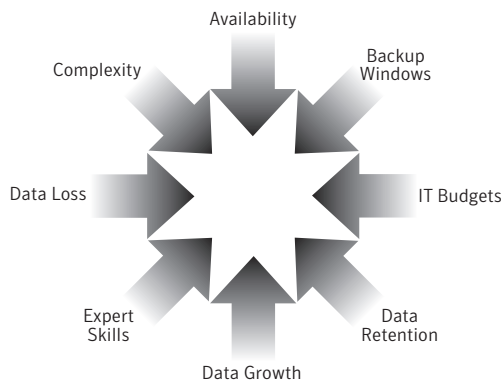


Figure 2. Data protection challenges

Some problems are closely related. For example, when it comes to complexity and expert skills, expert skills are more closely related to complexity than to simplicity. Similarly, data growth and data retention are closely related to IT budgets, whereas backup window and data loss problems both impact availability.

The skills problem is of such importance, it is worth more discussion. During a data emergency, people unexpectedly face problems sometimes not seen for months or years. It makes very little business sense for an organization to employ full-time staff simply to respond to very infrequent events. Such responsibilities are often given to IT staff that normally performs other duties. When facing emergencies, people experience a period of abnormal stress, which often results in the compounding of problems. Therefore, response to a data emergency requires a very high level of expertise and practice. However, utilizing these experts to respond to such episodes would require them to be reassigned from other important projects at a moment's

Getting the Most from your Data Protection Solution: A practical roadmap for comprehensive data protection

notice. People with such a high level of expertise are usually over-committed and are as likely to be traveling as they are to be onsite when a high-impact incident or data emergency occurs. A technical solution without skilled people who are rehearsed and ready to respond is of minimal value. Finding, training, and retaining skilled data protection professionals are ongoing concerns for IT management.

As backup windows disappear—a phenomenon that is already occurring in most large international enterprises—simply performing backup operations faster is not a sufficient response. The disappearing backup window is symptomatic of a business environment that demands greater data availability. Various replication and online backup approaches seemingly avoid backup windows, but they often have such an appetite for system resources that they contend with production workloads. Therefore, new approaches must be adopted that do not contend with production workloads for resources and are not bound by narrow backup windows.

The volume of data requiring protection continues to grow while recovery demands become more aggressive. Growth in the amount of data requiring protection is driven by two factors: (1) The volume of data continues to grow 40 to 60 percent per year for most large enterprises, and (2) new regulations require companies to retain more data for longer periods of time. Recovery requirements are becoming more aggressive as (1) business dependencies on data become more critical and (2) more and more critical applications depend on the same body of data.

When multiple applications access the same body of data, data recovery objectives are determined by the most aggressive objective for any of the applications. Since it has become commonplace for multiple applications to share a body of data for different purposes, recovery of shared data is the most critical data protection problem to solve. Most data sharing is accomplished by client/server data access, which delegates a server to provide access to and management of a given body of data. Applications access the body of shared data through a given data server, which acts as a broker. Therefore, recovery of shared data depends on recovery of its broker—its data server. Figure 3 illustrates a hypothetical three-tier architecture where multiple applications access a common body of data. The tiers are:

1. Workstation applications (A1–An) populate the presentation tier.
2. Departmental (B), mail (C), Web (D), and application (E) servers populate the *application tier*.
3. Database (F) and file (G) servers populate the *data tier*. In this example, the application server could also be considered part of the *data tier*. The differentiation between the application server and the database servers is not always orderly.

Getting the Most from your Data Protection Solution: A practical roadmap for comprehensive data protection

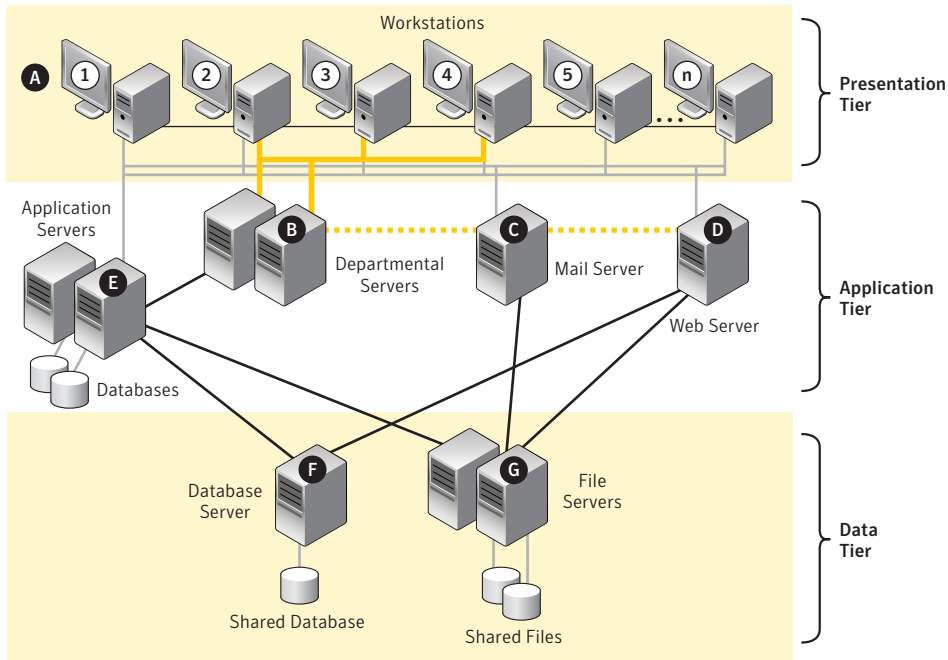


Figure 3. Typical three-tier architecture

As denoted by direct connections between the left-most workstation (A1) and application servers (E), it is possible to bypass tiers. Figure 3 shows application servers (E) both encapsulating their own databases and connecting to a shared database through a database server (F), which illustrates that tiers might not be entirely disjointed. The ability to bypass and combine tiers requires the bottom-most tier to encapsulate all shared data. This encapsulation places the burden of recovering shared data on database (F) and file (G) servers. Therefore, recovery objectives for shared databases and files tend to be the most aggressive in a multi-tier topology, which is the most common topology in large enterprises.

For example, if three applications at the presentation or application tier had 10-minute, 2-hour, and 4-hour *recovery time objectives* (RTO) respectively, the RTO at the data tier would be 10 minutes. Assuming that half the recovery time at the data tier is related to tasks other than data recovery, the RTO for data recovery would be 5 minutes. Obviously, the most aggressive RTOs are associated with data.

Getting the Most from your Data Protection Solution: A practical roadmap for comprehensive data protection

Data protection problem symptoms, and the repercussions of doing nothing

Doing nothing about your data protection problems may have severe consequences, such as:

- Lost business or productivity
- Liability for losing data or failing to keep data private
- Fines for regulatory violations, or inability to defend lawsuits (in addition to and beyond prosecution)

Of course, not all disruptive events have severe, immediate consequences, but frequent small losses and excess IT costs will drain company profits over time. Since doing nothing can have severe consequences, you should ask the following questions when assessing your business needs for data protection:

Question: (4) Strongly agree, (3) Agree, (2) Neutral, (1) Disagree, (0) Strongly disagree	Response
Have you protected yourself from data emergencies that might result from operational mishaps, hostile intrusions, and disastrous events?	
Are you satisfied that money you have invested to protect your critical data has been well spent?	
Have you implemented best practices that prepare you for data emergencies?	
Do you test your state of preparedness and restoration processes regularly?	
Can you retrieve data on demand from online and offline sources in audit situations or data emergencies?	
Is data at your remote/satellite offices protected?	
Is your investment in data protection solutions and skills adequate to maintain a state-of-the-art environment?	
Are you able to find, retain, and replace skilled data protection professionals?	
Score: Excellent 29–32; Good 24–28; Fair 20–23; Warning 16–19; Trouble 0–15	

Getting the Most from your Data Protection Solution: A practical roadmap for comprehensive data protection

Of course, few organizations can answer all questions affirmatively. However, if you are dissatisfied with too many negative responses, further exploration is likely required. Rate yourself. What score is high enough to protect your company and your career? Following are some of the many symptoms that might signal data protection problems:

- Backup
 - Frequent failure to complete backup jobs on the first attempt
 - Backup jobs encroach on OLTP or batch job periods
 - Backup jobs delayed until the next available backup window
 - Critical data not backed up or backed up too infrequently
 - Backup tapes not sent offsite in a timely manner or not protected from destruction
- Data restoration and operational recovery
 - Excessive data recovery incidents caused by operational mishaps
 - Compounding of data recovery problems due to operational mistakes
 - Frequent failure to restore data from the backup on the first attempt
 - Too long to retrieve backup tapes from offsite storage
 - Regular failure to meet recovery objectives
- Replication and disaster recovery
 - Insufficient or no disaster recovery plans
 - Excessive time elapses before recovery can begin
 - Cannot support aggressive recovery objectives for time-critical data
 - Inadequate or nonexistent recovery for database servers, file servers, and mail servers
 - Replication falling too far behind to achieve recovery objectives
 - Demand on replication network resources too high to be cost-effective
 - Replication processes impact OLTP applications or batch job throughput
 - Difficult to recover data protection infrastructure in a disaster
- Testing, troubleshooting, and solution maintenance
 - Testing of data recovery procedures known or perceived to be a potential disaster
 - Resolution of data protection problems drag on for weeks to months
 - Root-cause analysis takes too long or is not done at all
 - Patch levels are far out-of-date or backup/replication products are out-of-date by more than two release levels

Getting the Most from your Data Protection Solution: A practical roadmap for comprehensive data protection

Anatomy of a comprehensive data protection solution

Data protection has both technical and nontechnical aspects. Although technical aspects are of major importance, people and processes are of equal or greater importance. Technical aspects of data protection are most closely related to solution design, solution implementation, and operational tasks. People and process aspects of data protection are most closely related to planning, best practices, and ongoing testing. A comprehensive data protection solution combines technology and services into a cost-effective solution with benefits described in the following table.

Benefits of a comprehensive solution

- Reduces the amount of application downtime caused by data emergencies
- Meets recovery objectives that support even the most critical data
- Cost-effectively backs up and retains massive amounts of noncritical data
- Raises backup and recovery success rates well above industry standards
- Mitigates constraints imposed by tight or disappearing backup windows
- Proactively prepares for operational mishaps and disastrous events
- Minimizes the gap between the current environment and a state-of-the-art environment
- Mitigates the attrition/loss of skilled data protection professionals
- Makes the overall cost of disaster recovery and replication more affordable
- Minimizes and manages IT operational risks

The major technology goal is to design, implement, and maintain a state-of-the-art data protection environment, which displays the following functional characteristics:

- Protects data in a manner that enables recovery within a time window specified by the *recovery time objective* (RTO) and with no more data loss than specified by the *recovery point objective* (RPO). The solution might employ various combinations of replication and backup to this end.
- Supports recovery of all classes of data (e.g., database, flat files, and email). Recovery means restoring data to its normal operational state and verifying success.
- Employs the most cost-effective protection methodology for each class of data without compromising recovery objectives. For example, flat files with very relaxed recovery objectives would be handled differently than mission-critical databases.

Getting the Most from your Data Protection Solution: A practical roadmap for comprehensive data protection

- Recovers database servers, file servers, and mail servers to an operational state. Recovery is only deemed successful when data can be safely accessed by applications through a data server.
- Automates recovery while leaving critical recovery decisions to IT staff. Automation reduces the incidence of operational mishaps and prevents problems from compounding during stressful recovery situations.
- Avoids contention with applications for server, network, and storage resources. Ideally, offloads CPU-, memory-, and I/O-intensive operations from application servers to management servers. The potential that backup jobs might disrupt production workloads forces backup operations to be performed during off-hour backup windows. Avoiding or eliminating backup windows increases backup success rates and provides flexibility in backup schedules.
- Uses tape media and tape drives efficiently to contain costs. Costs associated with tapes are a major factor in a backup solution.
- Eliminates or greatly reduces time to retrieve backup media from offline storage. Delays associated with retrieving backup media can add many hours to end-to-end recovery time.
- Uses the replication network thriftily to contain network costs. Studies and experience have shown that the network is a major cost factor in employing a replication solution.
- Integrates data protection software with IT management frameworks such as HP OpenView and IBM® Tivoli® TME® 10. Leverages your investment in IT management frameworks.
- Requires no more than a skeletal onsite staff for data protection operations. Eliminating the need for a large staff at production locations saves onsite staffing costs, leverages offsite services for multiple production locations, mitigates the impact of a local area disaster, and increases outsourcing opportunities.

Getting the Most from your Data Protection Solution: A practical roadmap for comprehensive data protection

There is no single technical approach that both delivers aggressive recovery and meets low cost objectives. Therefore, any solution must be a hybrid, where each technical approach addresses part of the problem. Figure 4 illustrates various technical approaches and the problem area each addresses.

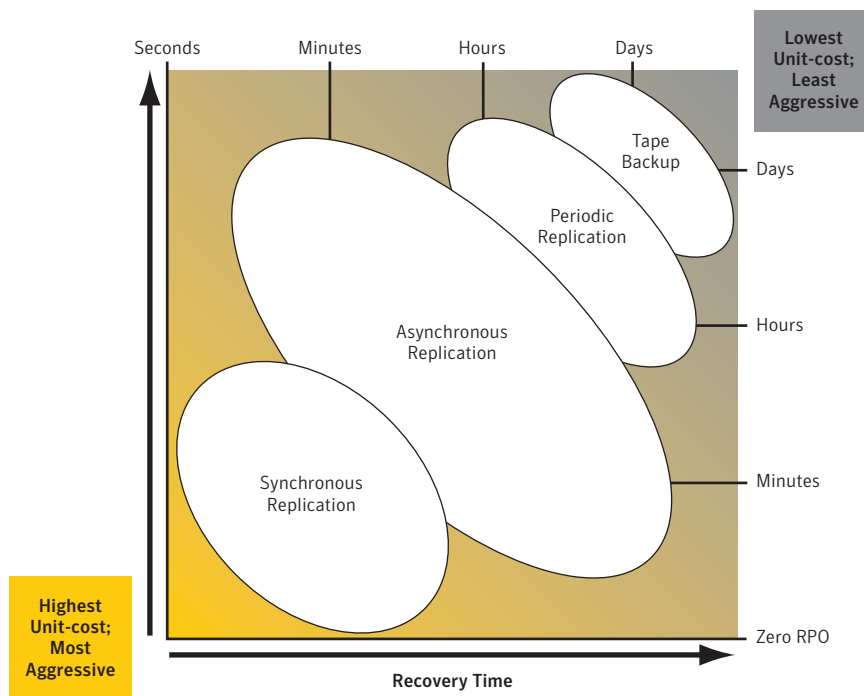


Figure 4. Technical approaches aligned to recovery objectives

Backup addresses the low-cost problem where recovery objectives are least aggressive: RTO and RPO of hours to days. Backup is the data protection approach most commonly employed as the last line of defense when all other recovery approaches fail. Also, disk-based backup approaches can drive down recovery time and make it possible to back up more often than daily. A backup approach might be sufficient to address recovery of application databases in the earlier example.

Getting the Most from your Data Protection Solution: A practical roadmap for comprehensive data protection

Periodic replication approaches can drive RTO and RPO lower than can be done by traditional tape backup approaches but at a higher unit-cost than tape backup. However, periodic replication and disk-based, incremental backup produce very similar results. A periodic replication approach or a disk-based, incremental backup approach might be a good way to protect file servers in the earlier example.

Where periodic replication can drive RTO and RPO down to a few hours, asynchronous replication can drive RTO and RPO to few minutes but at a higher unit-cost than periodic replication.

Synchronous replication is the only way to achieve zero RPO (i.e., no data loss). It might also be used to drive RTO down to a few minutes. As expected, purely synchronous replication approaches have the highest unit-cost, especially over long distances. Skillful hybrids of synchronous and asynchronous approaches can deliver zero RPO over long distances, in all but the most extreme cases, at unit-costs nominally higher than purely asynchronous approaches.

An overriding problem is that a cost-effective, integrated solution to data protection problems does not come ready to use, out of a box. It requires skillful integration of technical approaches, rigorous testing, knowledgeable application to various classes of data, and a long-lasting commitment to maintenance. Anything less does not produce the desired result: protection of the data.

Symantec data protection services

Symantec has launched data protection operational and residency services to help customers manage their IT environments more effectively and help them extract maximum value from the data protection products they have deployed. This service offering adopts a flexible approach to customer needs, does not need a large onsite staff, and does not require prohibitive service definitions to administer. Symantec services vary from light touch support through fully managed operational services.

Operational and residency services comprise Symantec's data protection services and complement each other by meeting different data protection needs. Operational services focus on daily, periodic, and episodic tasks required to protect and recovery data. Residency services supplement an IT organization's staff in areas where they need additional skills or lack specific skills. Figure 5 summarizes the Symantec data protection operational and residency services.

Getting the Most from your Data Protection Solution: A practical roadmap for comprehensive data protection

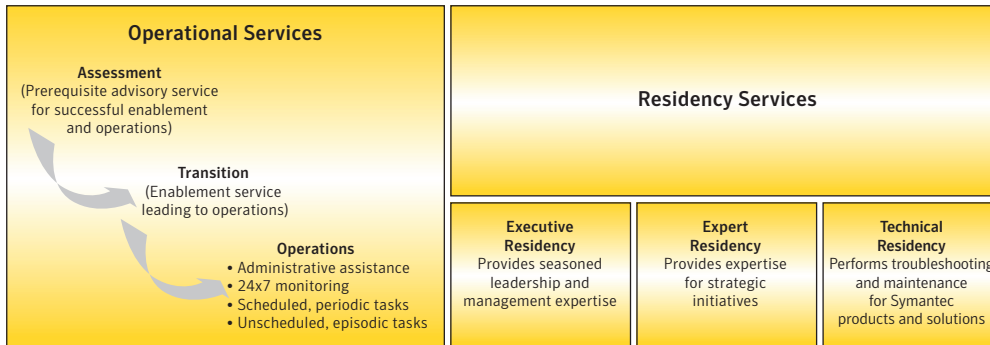


Figure 5. Symantec data protection services

The intent is for customers to win through improved service delivery, increased efficiency, reduced risk, and lower costs. Whether winning is through operational or residency services, the goal is to help Symantec customers achieve their IT objectives.

Symantec data protection operational services address the following situations, depending on customer needs:

- Recovery from data emergencies caused by operational mishaps at data centers through (1) basic data center backup and restoration of files and raw volumes and (2) backup and recovery within a given production site or at a nearby site
- Recovery from data emergencies caused by disastrous events at data centers through (1) replication of critical data sources to a disaster recovery site and (2) backup and recovery at a nearby or distant disaster recovery site
- Recovery of databases following data emergencies at data centers through (1) physical database recovery and (2) recovery of databases to the operational state specified by RPOs
- Monitoring and recovery of backup processes, replication processes, database servers, file servers, and mail servers
- Backup/restore of files on remote servers and always-on workstations located at satellite offices

Getting the Most from your Data Protection Solution: A practical roadmap for comprehensive data protection

The following tasks make up the Symantec data protection operational services:

Assessment tasks

The assessment is designed to reduce the risk of developing the wrong solution and of setting unrealistic service-level objectives. Symantec must perform a thorough assessment to assure the success of operational services, especially since the success of operational services is normally tied to a service-level agreement (SLA).

Transition tasks

These tasks transition a customer from the current state to the recommended state. The transition addresses both the technical solution and operational practices. Briefly, a transition involves detail solution design, implementation, stabilization, and training.

Operational tasks

Operational tasks include (1) administrative assistance as requested, (2) 24x7 monitoring (3) scheduled periodic tasks, and (4) unscheduled episodic tasks. A customer may request Symantec assistance, as needed, with any or all administrative tasks. A customer may also engage Symantec to perform 24x7 monitoring, scheduled period tasks, and unscheduled episodic tasks or retain some or all of that responsibility themselves. In all cases, Symantec will train the customer's staff to operate the new or upgraded data protection solution.

Administrative guidance involves:

- Assistance in scheduling backup jobs
- Assistance in restoring files from backup media
- Assistance in recovering file-systems and databases
- Assistance generating reports

Symantec will normally provide administrative assistance remotely but onsite coverage can be arranged when required.

24x7 monitoring involves continuous monitoring of backup/replication processes and management servers plus generation of trouble alerts, as needed.

Getting the Most from your Data Protection Solution: A practical roadmap for comprehensive data protection

Scheduled, periodic tasks help prepare an IT staff for data emergencies and prevent the solution from becoming prematurely obsolete. Preparatory tasks are most often neglected until a serious disruption becomes a catalyst for action. In one way or another, customers must either pay for preparedness or for lack of preparation. Symantec builds preparation into its operational services offerings. Even if an IT organization chooses to perform daily administrative tasks, it would be advisable to have Symantec perform periodic tasks. Scheduled, periodic tasks include:

- Ongoing maintenance of data protection software
- Reassessment of the data protection environment and definition of improvement projects
- Monthly or quarterly operational recovery rehearsals
- Annual or biannual disaster fire drills and health checks

Unscheduled, episodic tasks are performed in response to events that occur at inconvenient times, sometimes after not occurring for months or years. It makes business sense for an IT organization to engage Symantec to cover serious but infrequent episodic events. Unscheduled, episodic tasks include:

- Root-cause analysis of high-impact incidents
- Response to data emergencies caused by operational mishaps and disastrous events
- Off-hours support

Symantec offers *Residency Services* to supplement a customer's IT staff in one or more critical areas where they are understaffed or not staffed at all. Residency Services fall into three categories:

- *Executive residents* are experienced principal consultants who provide seasoned leadership and enterprise data protection management expertise. Executive residents focus on the strategic planning related to the way data protection (or lack of it) can impact a customer's business.
- *Expert residents* are experienced consultants and project managers who assume key data protection roles to deliver ongoing expertise for data protection-related strategic initiatives. Expert residents focus on requirements, architecture, planning, and projects designed to help improve the data protection environment.
- *Technical residents* are qualified data protection consultants who provide administration, maintenance, and basic troubleshooting of data protection solutions. Technical residents focus on the Symantec data protection products and technical solutions.

Getting the Most from your Data Protection Solution: A practical roadmap for comprehensive data protection

The Symantec technical solution

Point solutions may address parts of the data protection challenge (e.g., backup windows) but fail in others (e.g., automated recovery) and add complications (multiple management domains, multiple backup and recovery processes). Of course, the answer is a unified data protection solution, which provides a single interface for an enterprise's data protection operations.

The Symantec data protection technical solution integrates backup, recovery, replication, monitoring, and reporting technology to produce a reliable technical solution and set the stage for evolving that solution. Furthermore, the Symantec solution addresses total cost of ownership (TCO) as well as RTO and RPO. Any given implementation might integrate several components, such as:

- A backup component: Veritas NetBackup™, Symantec Backup Exec™, Symantec LiveState™ Recovery, PureDisk™
- A replication component: Veritas™ Volume Replicator, disk-hardware replication, transactional replication
- A monitoring component: NetBackup Operations Manager optionally integrated with HP OpenView or IBM Tivoli
- A reporting component: Veritas Backup Reporter
- Other components: Veritas Cluster Server, Veritas Patch Manager

The following discussion highlights a few of the many technical capabilities embodied in the Symantec solution:

- Disk-based backup/restore capabilities of NetBackup allow recent backup images to be kept online. Since the most recent backup image is the most likely candidate for restoration, the time to retrieve backup images from online storage is instantaneous in most cases, and restoration from online disk is much faster than from offline tape. Also, backing up to disk is faster and permits greater parallelism than backing up to tape. The combined effect results in much higher throughput of backup jobs.
- Staging of backup images from disk to tape, as backup images age, minimizes the cost of relatively expensive disk media; old backup images reside on relatively inexpensive tape media. NetBackup provides a convenient, policy-driven interface to manage staging of backup images from disk to tape.

Getting the Most from your Data Protection Solution: A practical roadmap for comprehensive data protection

- Integration of replication into the Symantec data protection solution supports a wide range of moderate to aggressive recovery objectives. Veritas Volume Replicator, disk array, or third-party transactional software technology may provide the replication capability for the Symantec data protection solution.
- The synthetic backup capability of NetBackup eliminates the need to take full backups on an application client, thereby reducing the impact of backup on production workloads and reducing the time to back up a body of data during a backup window. After the first full backup, only incremental backups are required. Incremental backup images can be processed offline to create cumulative backup images or a new full backup image.
- Archival capabilities of NetBackup reduce the amount of inactive data that needs to be backed up, stored, and restored. Archiving inactive files and logs can reduce the cost and time of backup operations.
- Snapshot technologies found in Veritas Storage Foundation™ or in disk arrays allow NetBackup to move resource-intensive backup operations from application clients to media servers. NetBackup also exploits underlying third-party copy technology to help eliminate costly data movement from application servers. Consequently, restrictive backup windows can become a vestige of the past, and backup operations can occur nearly anytime.

The Symantec solution advances and supports the idea of continuous data protection (CDP) as a long-term strategy. The idea of CDP involves collecting changes to data in a manner that allows rolling data backward or forward to any point in time, at any time, under all circumstances. Execution of the CDP strategy will occur in phases. Likewise, CDP will apply differently to different classes of data and not at all to some classes of data. Some data is critical and warrants the expense of CDP. Less critical data does not warrant CDP. With Symantec's recent acquisition of Revivio's CDP technology, Symantec customers can look forward to major advances toward CDP.

Getting the Most from your Data Protection Solution: A practical roadmap for comprehensive data protection

Conclusion

Symantec is a leader in data protection. Symantec products have captured a significant share of the backup software market, and more data is protected by Symantec than by any other software vendor. You can benefit from our integrated technical solution and our operational and residency services. The Symantec data protection solution integrates industry-leading technology to produce a reliable, state-of-the-art technical solution. Our operational and residency services can help solve your immediate operational problems and guide you along the way as the state-of-the-art advances toward continuous data protection.

Since ignoring data protection problems can have severe consequences for your company, it makes sense to do something about it. Forward-thinking IT organizations should be planning how they will go about creating or acquiring a complete data protection solution that both meets current needs and will evolve to meet future needs. The business environment is changing, and IT must change with the business.

Getting the Most from your Data Protection Solution: A practical roadmap for comprehensive data protection

Glossary

This glossary provides brief definitions for key terms used in this white paper, along with terms closely related to data protection and disaster recovery.

Archiving is the act of copying data from online storage (usually disk) to offline storage (usually tape) for the purpose of saving money on relatively expensive online storage. Archival solutions usually have the ability to automatically recall data from offline storage and may also provide query capabilities.

Backup is an operation performed periodically to create an image of some body of online data that can later be restored.

- *Full backup*: Creating a complete backup image of some unit of data
- *Incremental backup*: Creating a partial backup image where only changes to a body of data have been included
- *Off-host backup*: A backup operation that takes place on a host other than on the system hosting an application that reads or writes the data being backed up
- *Online backup*: A backup operation (sometimes called a hot backup) that proceeds in parallel to applications that read and write the data being backed up

Backup window refers to a time period in which backup jobs have been scheduled to execute and are expected to complete.

Continuous Data Protection (CDP) is a technology that continuously collects changes to data in a manner that allows rolling data backward or forward to any point in time, at any time, under all conditions.

Data protection is the preparation for and recovery from data emergencies caused by operational mishaps and disastrous events. Data protection solutions and services support disaster recovery.

Data recovery is an operation that recovers some body of data to a fully operational state from an online or offline image.

Database recovery is an operation that recovers a database to a fully operational state using the database log.

Getting the Most from your Data Protection Solution: A practical roadmap for comprehensive data protection

Data replication is an ongoing act of creating two or more copies of the same data so that any copy can be used by applications but not necessarily at the same time.

- *Synchronous replication*: Writing one or more secondary copies of some primary unit of data before the process that has written the primary copy proceeds
- *Asynchronous replication*: Writing one or more secondary copies of some primary unit of data, which has been written safely and while the process that writes the primary copy proceeds
- *Periodic replication*: Asynchronous replication that allows changes to secondary copies to accumulate and be written in batches

Data restoration is the act of restoring the online image of some body of data from a backup image.

Data server refers to a computer system that manages a body of data and provides data access to an application. A data server acts as a broker for accessing data.

- *Database server*: A data server that provides access to databases
- *File server*: A data server that provides access to flat files
- *Mail server*: A specialized data server that provides access to email

Disaster refers to any disruptive event that causes the prolonged interruption of service at a given site and might result in loss of data. Technically, a disaster is the prolonged outage of multiple components at a given site that hinders recovery of services at the site.

Disaster recovery (DR) is the act of restoring application and information services to a business following site-wide outages caused by a disastrous event. Disaster recovery services support business continuity requirements and depend on high availability services.

File restoration is the act of restoring one or more files from a backup image or from an alternate file server. Generically, file restoration is called file recovery.

High availability (HA) refers to the maintenance of access to applications and information despite single server, storage unit, or network outages within a single data center. High availability services support business continuity requirements and lay the foundation for disaster recovery.

Getting the Most from your Data Protection Solution: A practical roadmap for comprehensive data protection

Recovery point objective (RPO) is a measure of how old data can be following a recovery incident. RPO generally specifies the acceptable quantity of recent changes to a body of data (in the period immediately prior to a service interruption), which may be lost due to a disastrous event.

Recovery time objective (RTO) is a measure of how long an organization can be without access before experiencing severe consequences. RTO can either include or exclude the delay between a service interruption and declaration of a disaster, depending on whether RTO is seen from the perspective of the business or the IT organization. RTO can apply separately to business processes, applications, and data.

Service-level agreement (SLA) is a contract between a service provider and a client that specifies success criteria for each service and consequences of not meeting criteria.

Service-level objective (SLO) is a quantitative measurement that determines whether or not a service has been delivered successfully.

Snapshot is an instantaneous copy of some body of data, created using storage software or hardware technology that represents some atomic point in time. A snapshot may be dense (a complete copy of the entire unit of data) or sparse (a copy containing only modified data). Nearly all snapshots in use today are dense.

Third-party copy is a technology that copies data without the participation of an application host. Third-party copy is usually associated with storage system hardware, but appliances can closely approximate third-party copy with software.

About Symantec

Symantec is a global leader in infrastructure software, enabling businesses and consumers to have confidence in a connected world.

The company helps customers protect their infrastructure, information, and interactions by delivering software and services that address risks to security, availability, compliance, and performance. Headquartered in Cupertino, Calif., Symantec has operations in 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 © 2007 Symantec Corporation. All rights reserved. Symantec, the Symantec logo, Backup Exec, LiveState, NetBackup, PureDisk, Veritas, and Veritas Storage Foundation are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. IBM, Tivoli, and TME are trademarks of International Business Machines Corporation in the United States, other countries, or both. Other names may be trademarks of their respective owners. Printed in the U.S.A.
04/07 11887862