

Complements of StorageIO

This chapter download from the book “Cloud and Virtual Data Storage Networking” (CRC Press) by noted IT industry veteran and Server StorageIO founder Greg Schulz is complements of The Server and StorageIO Group (StorageIO). Learn more about the techniques, trends, technologies and products covered in this book by visiting storageio.com and storageioblog.com and register for events and other promotions. Follow us on twitter @storageio or on Google+ among other social media venues.

Brouwer *Storage Consultancy*
Seminar for Storage Professionals with Greg Schulz
May 7th, 8th, 9th 2012 in Nijkerk Holland
[Click here to learn more](#)



Visit storageio.com/events to see upcoming seminars and activities

Cloud and Virtual Data Storage Networking has been added to the Intel Recommended Reading List (IRRL) for Developers. Click on the image below to learn more about the IRRL.



The Recommended Reading List is a valuable resource for technical professionals who want to thoroughly explore topics such as software threading, wireless technologies, power management, and more. Dozens of industry technologists, corporate fellows, and engineers have helped by suggesting books and reviewing the list.

Learn more about Cloud and Virtual Data Storage Networking (CRC Press) by visiting storageio.com/books

To become a chapter download sponsor or to discuss your other project opportunities contact us at info@storageio.com.

#1 VM Backup

Veeam Backup & Replication for

Hyper-V

is HERE!

[Learn More](#)

VEEAM

StorageIOblog.com site sponsor

Chapter 5

Data Protection: Backup/Restore and Business Continuance/ Disaster Recovery

Information security: To protect, preserve, and serve.

– Greg Schulz

In This Chapter

- The difference between business continuance (BC) and disaster recovery (DR)
- The importance of an effective data protection plan and strategy
- Why it is time to modernize backup and data protection
- How to reduce costs by using tiered data protection and different technologies

This chapter looks at issues, challenges, and opportunities for protecting data in cloud, virtual, and data storage networks. The focus of data protection in this chapter is on maintaining availability and accessibility of both active and inactive data. In the context of this chapter, data protection builds on the previous chapter's subject of security by expanding our focus to information accessibility and maintenance of data integrity. Key themes, buzzwords, and trends addressed in this chapter include high availability (HA), backup and restore, business continuance (BC) and disaster recovery (DR) along with replication and snapshot-related technologies.

5.1. Getting Started

Mention “DP” to people in IT and, depending on their area of interest and their length of experience, you may get answers such as Dual Platter, Dedupe Performance, Data Processing, Double or Dual Parity, or perhaps even Dance Partner from someone more creatively inclined. For the purposes of this chapter, DP is data protection.

“Data loss” can be a misleading idea: If your data is intact but you cannot get to it when needed, is the data really “lost”? There are many types of data loss, including loss of accessibility or availability and complete loss. Loss of data availability means that somewhere—perhaps off-line on a removable disk, optical drive, tape, or at another site on-line, near-line, or off-line—your data is still intact, but you cannot get to it. There is also real data loss, where both your primary copy and backup as well as archive data are lost, stolen, corrupted, or never actually protected.

Protection of data and information services delivery applies to:

- Workgroups, departments, and remote offices/branch offices (ROBOs)
- Enterprise, small to medium-size business (SMB)
- Small office/home office (SOHO) and consumer environments
- Workstations, laptops, and mobile devices
- Physical and virtual servers, workstations and desktops
- Managed service providers, public and private clouds
- Integrated stacks, converged and unified solutions

5.2. Data Protection Challenges and Opportunities

IT organizations of all sizes are tasked with the basic responsibilities of protecting, preserving, and serving information services when needed. Since new data is continuously created while old data must continuously be handled, there is more data to process, move, and store for longer periods of time than there was even yesterday. Consumers of IT services are dependent on applications and data being readily available and protected by BC and DR activities. A challenge for many organizations is how to balance the cost to protect against various threat risks, regulatory and other compliance requirements, and the demand to protect, preserve, and serve more data for longer periods of time in an economical manner.

Data protection trends and challenges include:

- More data to process, move, protect, preserve, and serve
- Shifting data lifecycle and access patterns while retaining data longer
- Continued focus on cost containment or reductions
- Reliance on information services accessible when and where needed
- Increase in mobile-generated and -accessed information services
- Cloud, virtualized, dynamic, and flexible computing
- Outages resulting from human error or design deficiency

There are other challenges related to protecting data and applications in physical, virtual, and cloud environments. For example, in a nonvirtualized server environment, the loss of a physical server impacts the applications running on that server. In a highly aggregated or consolidated environment, the loss of a physical server supporting many virtual machines (VMs) has a much more significant impact, affecting all the applications supported by the virtual servers. Another challenge is protecting the growing amount of structured and unstructured data in primary data centers along with data in ROBOs, workgroups, field offices, and other locations.

Data protection opportunities include:

- Stretch available budgets further to protect and preserve more data longer.
- Maximize return on investment (ROI) in capital and operating expenditures.
- Improve quality of service (QoS), service-level agreements (SLAs) and service-level objectives (SLOs), including recovery-time objectives (RTOs) and recovery-point objectives (RPOs).
- Modernize data protection including backup/restore and BC/DR.
- Reduce cost of services delivered via improved efficiencies.
- Provide protection of cloud, virtual, and physical resources.
- Leverage cloud and virtualization technologies to mask complexities.
- Reconcile and streamline protection frequencies and retention cycles.

5.3. Protect, Preserve, and Serve Information Services

Disaster recovery (DR) can mean different things to different people; however, for the purposes of this chapter it will mean two things. The first is an overall process, paradigm, or set of best practices that spans across different technology groups and organizational boundaries. The second are the steps taken as a last resort to reconstruct or rebuild, reconfigure, restore, reload, rollback, restart, and resume information and organizational services or functionality in the event of a severe incident or catastrophe. Business continuance (BC) and DR are often used interchangeably to mean the same thing. We will treat them separately, with business continuance focused on disaster prevention, surviving a disaster or incident, and keeping the business running, and disaster recovery as the process of putting all of the pieces back together again if HA, BC, and other steps were either not taken or failed.

Threat risks to information services delivery requiring data protection include:

- More data being generated, stored, and used remotely
- Funding constraints coupled with increased demands
- Accidental or intentional deletion and data corruption
- Operating system, application software, server, or storage failure
- Loss of access to site, servers, storage, or networking resources
- Acts of nature or acts of man, headline and nonheadline incidents
- Local site, campus, metropolitan, regional, or global incidents

- Business or regulatory compliance requirements
- Increased awareness of threat risks and reliance on information services
- Technology failure or inappropriate configuration design
- Planned or scheduled and unscheduled downtime
- Network or communications disruptions including cables being cut
- Problems introduced via configuration changes

Table 5.1 shows various situations or scenarios in which information services have been or could be impacted. The scenarios or situations are categorized into different levels that can be used to help determine what type of data protection to apply to counter applicable threat risks.

Table 5.1 Protecting Against Various Levels of Threats and Impending Risks

Level	Description of Incident or Scenario
1	Systems are running alerts warning of potential threat and disruption received
2	Hardware, software, network, or facilities component has failed
3	Single system or application disruption
4	Single major disruption or multiple lower-level incidents
5	Metropolitan or campus disaster
6	Major local or regional disaster

- *Level 1: Systems are running; alerts or advance warning of potential threat and disruption have been received.* Notification or indications of possible threat or service disruption have been received, ranging from a virus or security issue to a system potentially being compromised or a hardware device logging errors or software indicating that consistency checks should be taken. Weather reports might indicate an approaching storm, or acts of civil unrest or other threats may be anticipated. Left unchecked, or not corrected, Level 1 threats may escalate to a higher threat level or, worse, a rolling disaster.
- *Level 2: A hardware, software, or network/facilities component has failed.* Business functionality has not yet been disrupted. Business functions, information services, and their applications remain operational. The incident might be a failure in a component such as a disk drive, storage controller, server, network path, power supply, or other item that is being protected by redundancy and automatic failover. The threat might also be a virus, software, or data correctable error leveraging a translation log or journal rollback. There is vulnerability of a multiple failure during the repair process escalating into a disaster.
- *Level 3: Single system or application disruption.* Overall business or information services remain available, but some functionality is not currently available. An entire system or application (hardware, software, and network) may have failed or been shut down due to a facilities issue such as circuit breaker or zone cooling issue. Some disruption may occur during failover to a standby site if available

or, if the disruption will be extensive in length, restoration from backup media. Failback occurs when resources are ready, safe, and stable. Databases may be read-only until updates can resume.

- *Level 4: Single major disruption or multiple lower-level incidents.* The data center exists and most systems are functional, but some Level 2 or 3 scenarios may be occurring. Performance may be slow due to rebuild, failover, or loss of primary systems placing heavy demand on standby resources. Disruption may be hardware-, applications-, or data-related. Resolution may require failover to a standby system with good data or restoration from a known good copy or snapshot.
- *Level 5: Metropolitan or campus disaster.* The data center, information, and resources are intact, but access to them has been lost for some period of time due to a localized incident. If a standby or failover site is available in a different location, service may resume; otherwise, recovery occurs elsewhere.
- *Level 6: Major local or regional disaster.* Loss or damage to facilities and related infrastructure, including power, water, communications, or personnel, due to acts of nature (flood, earthquake, hurricane) or acts of man, including terrorism. A determination is made that the primary site will not be available/accessible for an extended period of time, resulting in major disruption to business function for any applications not protected via HA or BC.

Different types or levels (Table 5.1) of disasters or incidents can be localized to a given site, campus, metropolitan, regional, or global basis. Understanding the applicable data protection threat risks or scenarios along with the likelihood of their occurrence and subsequent impact to the business is part of technology or service alignment. The importance of technology and data protection service alignment is to make sure that an appropriate level of protection is applied when and where needed to stretch available budgets as far as possible.

Figure 5.1 shows how distance can be part of enabling business or information services survivability to different threat risks for some environments or applications. If applications or services are focused only on a local or metropolitan audience, then regional or global forms of protection may not be required. Granted, they may be nice to have, and if affordable, then practical.

Distance is important for enabling data protection and survivability. While distance is often thought of in terms of physical space, time can also be a function of distance. This means being able to go back to a particular place or point from which data was copied or protected—known as a recovery-point objective (RPO).

Physical distance can be measured in inches, feet or meters, kilometers or miles. How would distance of inches be enough to enable data protection? By having data on two different storage devices located next to each other in case one fails. However, there would still be a point of failure if the server or storage system in which they were installed failed. The next logical step would be to have data on two different storage devices, which might be feet or meters apart in the same facility, to isolate and protect against device failure. Here the single point of failure would be the site or facility; this can be mitigated by having copies of data on different systems spread across a campus, metropolitan area, and region or on a global basis.

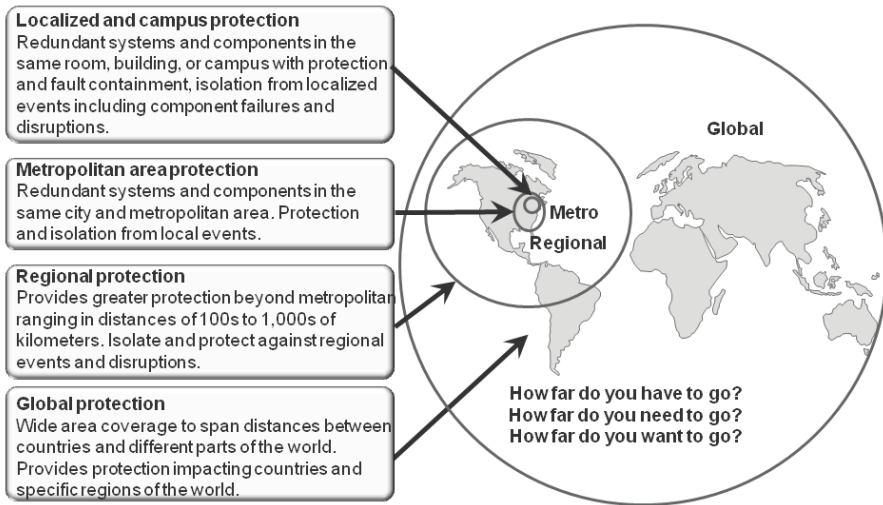


Figure 5.1 Protecting against various threat risks to data and information services.

5.3.1. Basic Information Reliability–Availability–Serviceability (RAS)

As the name implies, basic information services availability means limited or no data protection. This could mean that backups occur now and then with no recurring or regular frequency. Availability may be limited to servers or storage that lack failover or redundancy components, for example, storage that lacks RAID (redundant array of independent disks) data availability capabilities or redundant power supplies and cooling fans. Basic availability can be enhanced by increasing the frequency of backups, ensuring that important information is copied to different locations.

In addition to making copies of data that are stored in different locations (a local copy on disk, another copy on a fileserver, another stored at an off-site cloud or managed service provider site), retention is also important. Retention means how long those copies are kept before being deleted or destroyed. For example, if you have multiple copies of data that all expire after 14 days and you are only making copies of data once a week, if something goes wrong with the last backups, you may be facing a disaster situation. On the other hand, having too many copies for too long adds to the cost of protecting data. Managing threat risks needs to be balanced with available budgets as well as business needs.

Other common constraints for data protection include:

- Growing amount of data to protect and preserve
- Time including backup or protection windows
- Budgets (capital and operating)
- Technology interoperability or interdependencies
- Software license restrictions

- Lack of automation, reporting, or analytics for data protection
- False positives when diagnosing problems
- Staffing and in-house expertise
- Cross-technology ownership issues
- Upper management buy-in, support, or sign-off
- Policies or lack thereof
- Workflow and paperwork overhead

Items that need to be addressed or included in a data protection plan include:

- Facilities—Floor space, primary and secondary power, cooling, fire suppression
- Networking services—LAN, SAN, MAN, and WAN voice and data services
- Security—Physical and logical security including encryption key management
- Monitoring and management—Infrastructure resource management (IRM)
- Diagnostics tools—End-to-end tools for analysis and troubleshooting
- Software—Applications, middleware, databases, operating systems, hypervisors
- Hardware—Servers, storage, networking, workstations, and desktops
- High availability, backup/restore, snapshots and replication, media maintenance
- Best practices—Documentation, communication, change control
- Testing and audits—Review of plans and processes, random testing of activities

5.3.2. High Availability and Business Continuation

Think of high availability (HA) and business continuation (BC) as disaster prevention. Disaster prevention refers to containing or isolating faults from rolling into a larger event or disaster scenario. Essentially, enabling HA and BC means taking adequate steps within reason as well as budget constraints to eliminate or minimize the impacts of various incidents on information services delivery—in other words, enabling information services to actually service in the face of a disaster. Disaster recovery (DR), on the other hand, involves rebuilding, restoring, recovering, restarting, and resuming business after an incident that could not, within reason or budget, be contained.

Enabling HA and BC involves eliminating single points of failure and containing or isolating faults from spreading by using redundant components and failover software. In addition to hardware, software, and networking redundancy on a local as well as remote basis, another important aspect of both IRM in general and data protection specifically is change control. Change control means testing and validating hardware, software, application, or other configuration changes before they are implemented, updating applicable documents as part of workflow management, and having a plan in case the change does not work.

Having a fallback plan or process to back out of the change quickly can help keep a minor incident from escalating. A simple approach to change management is to have multiple copies of the configurations, applications, or data that is being updated, which can be reapplied if needed. Part of change control management should also be a determination of the interdependences of a change and associated remediation.

Not all incidents or outages are the result of a major disaster. As mentioned above, some can be the result of component failures or faults that were left uncontained and therefore expanded into a disaster. There is also the possibility that an IT environment can be reduced to physical ruins by a fire, flood, hurricane, tornado, or explosion caused by an accident or act of man. In other situations, an IT environment may be completely intact but not usable as a result of loss of access to a facility. For example, an area might be evacuated due to a chemical spill from a truck or railroad car. If the site is automated, with intervention available via remote access, the disruption may be minimal to nonexistent unless utilities were also cut. Having on-site standby electrical power and self-contained cooling would mitigate those risks; however, what about communications for networks along with adequate fuel supplies for backup generators and cooling water?

In other, less drastic, incidents, all hardware, networks, and software may be intact but a data corruption or error occurs, requiring rapid restoration to a previous point in time. If a recent snapshot can be rapidly recalled and restored, log or journal files applied, and integrity and consistency checks completed, the outage can be kept to a minimum. If, instead, you have to wait for data to be brought back on-site, reloaded, and then rollbacks along with consistency checks performed, that will take more time. This is where data protection comes back to a balance of cost versus risk to the business and the value of time. Not all applications will have the same time sensitivity, so not all data and applications should be protected the same way. Aligning the data protection strategy with the sensitivity of the data is one way of maximizing budgets and resources.

5.3.3. Disaster Recovery

As mentioned earlier, disaster recovery can be thought of in two ways, one being the overall process of ensuring business and organizational survivability and the other being the activities involved in reconstructing an environment after an incident. Basic RAS (reliability–availability–serviceability), HA, and BC can all be considered part of enabling an overall DR plan and strategy. The last line of defense to various threat levels (Table 5.1) in DR is the process of reconstructing, restoring, and resuming after a major disaster or incident beyond the abilities of HA and BC to cope (Figure 5.2).

What could cause a disaster and what would create only a minor inconvenience to information services delivery? For example, would a short outage of a few minutes result in any data loss, or simply loss of access to data for a short period of time? What would happen if the short outage turned into a day or longer? Figure 5.2 shows examples of normal running states with various forms of data protection occurring at different frequencies and retention lengths to combat various incidents or disaster scenarios.

Part of supporting growth or increasing business demands while reducing costs and maintaining or enhancing quality of service involves aligning the applicable level of data protection to the likely threat risk scenario. What threats or incidents are most likely to occur, and what would be the impact on your organization if they were not remedied? How much protection do you want, how much do you need, and what can

you afford? Put a different way, what can you afford not to do, and what is the subsequent impact to specific information services, applications, functions, or the entire business or organization?

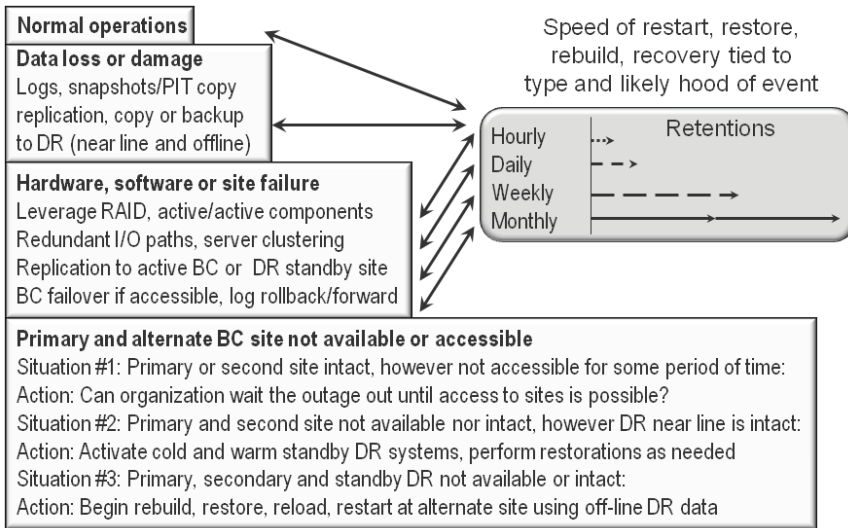


Figure 5.2 RAS, HA, BC, and DR as part of a data protection strategy.

5.3.4. Data Protection vs. Preservation (Backup vs. Archive)

While the two functions are related, backup is focused on protecting data with the intention of it being usable for recovery to a given point in time (the RPO), and archiving is aimed at preserving the state of data or an application for possible future use. They may sound similar, but they differ in retention cycles and in the frequency or interval at which backups are made versus archives.

Archives are usually retained for longer periods of time, such as years, while backups are typically retained for days, weeks, or months. Archives can be established for regulatory compliance purposes as well as to preserve intellectual property (IP) or project data for possible future use. Additionally, archives are used as part of data footprint reduction (DFR) as a means of migrating less frequently used or accessed data off-line or to another medium such as disk, tape, or cloud to reduce online or active storage space needs. The benefit of archiving databases, email, and Microsoft SharePoint or file systems is to free up space while reducing the amount of data that needs to be backed up or protected.

Backups and archives can use the same software and target hardware or service while implementing different policies and business practices. The main difference is that archiving is focused on saving the context of data and applications as of a point in time for long-term retention in case it's needed. Backup, on the other hand, preserves the context of data and applications as of a point in time for routine restoration of a

single file or dataset object or database table. Archiving as a tool to optimize storage capacity will be discussed further in Chapter 8.

5.4. SLO and SLAs: How Much Availability Do You Need vs. Want

Costs associated with data availability need to be understood to determine availability objectives. Vendors use terms such as “five 9s,” “six 9s,” or higher to describe their solutions’ availability. It is important to understand that availability is the sum of all components combined with design for fault isolation and containment. Seconds of downtime per year are shown in Table 5.2. How much availability you need and can afford will be a function of your environment, application and business requirements, and objectives.

Availability is only as good as the weakest link in a chain. In the case of a data center, that weakest link might be the applications, software, servers, storage, network, facilities, processes, or best practices. This means that, for example, installing a single converged SAN and LAN networking switch with “five 9s” or better availability could create a single point of failure. Keep in mind that the failure may be technology-related, a configuration issue, a software update failure, or something as simple as someone unplugging a physical network connection cable. Virtual data centers rely on physical resources to function; a good design can help eliminate unplanned outages to compensate for failure of an individual component. A good design removes complexity while providing scalability, stability, ease of management and maintenance, as well as fault containment and isolation. Design for both maintenance and to contain or isolate faults from spreading, as well as to balance risk, or the likelihood of something happening to required service levels and cost.

Table 5.2 Availability Expressed as a Number of “9s”

Availability (%)	Number of 9s	Amount of Downtime Per Year
99	Two	3.65 days
99.9	Three	8.77 hours
99.99	Four	52.6 minutes
99.999	Five	6.26 minutes
99.9999	Six	31.56 seconds
99.99999	Seven	3.16 seconds
99.999999	Eight	½ second

5.4.1. RTO and RPO: Balancing Data Availability vs. Time and Budgets

Figure 5.3 shows a timeline example that includes a gap in data coverage between where and when data was last protected and where it can be recovered. Also shown are

various component recovery time objectives, such as when hardware becomes available for use for operating systems or hypervisors, data, and applications. While server hardware, hypervisor, and operating system RTOs are important, as are storage and data restoration RTOs, the overall application RTO is what matters to the consumer of the information service or application. Figure 5.3 shows that there are different RTOs that need to be aligned to meet the cumulative service objective for a given class or category of service.

If a given application or information service has, as an example, a 4-hour RTO, it is important to understand what that RTO means. Make sure you know whether the 4-hour RTO is cumulative and when application users or consumers of services can expect to be able to resume work, or whether the RTO is for a given component (Figure 5.3). If the RTO of 4 hours is cumulative, then all other sub-RTOs for data restoration, operating system and hypervisors, database rebuilds or rollbacks, and verification must fit within that 4-hour window.

A common mistake is that multiple groups learn that the RTO is, continuing the example, 4 hours and assume that means they each have 4 hours to complete their required tasks. While some tasks may be done in parallel, some—such as data restoration followed by database verification or rebuild and application verification—are usually done serially; if each team assumes they have 4 hours to complete their task, the 4-hour cumulative RTO cannot be achieved.

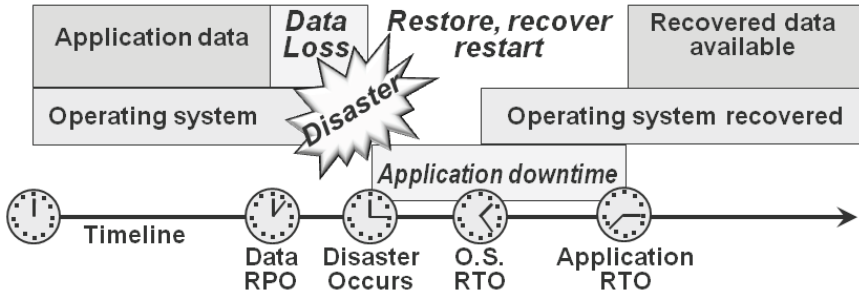


Figure 5.3 End-to-end recovery-time objectives.

5.4.2. Reconciling and Assessing RTO and RPO Requirements

Earlier, we discussed the importance of not treating all applications or data the same so as to do more with what you have while enhancing quality of service. For data protection and availability this is also true, in that an incorrect assumption as to what level of service is desired vs. what is required can increase costs. This means assessing actual availability requirements against what would be nice to have, to be able to align the applicable classes or categories of service and underlying technologies to a given situation.

With a continued industry trend toward using disk-to-disk (D2D) backup for more frequent and timely data protection, tape is finding a renewed role in larger, more infrequent backups for large-scale disaster recovery supporting long-term archiving and

data preservation of project data and compliance data. For example, D2D, combined with compression and de-duplication disk-based solutions, is used for local, daily and recurring backups that have shorter retention but that have more granularities (Figure 5.4). Meanwhile, weekly or monthly full backups are sent to disk at a secondary location, cloud server, or to tape, to free disk space as well as address PCFE concerns. These copies occur less often so there are not as many of them, but they are retained for longer periods of time.

By reconciling and tuning data protection frequencies along with retention cycles (Figure 5.4), the overhead of protecting data can be reduced while increasing survivability in a cost-effective manner. The principal idea is that for more commonly occurring incidents, recovery or restart occurs more often, faster, and with more ease than traditional data protection. D2D data protection combined with data footprint reduction (DFR) techniques means more copies of protected data can be kept closer to where it will be needed at a lower cost. Meanwhile, copies of data that are less likely to be accessed occur in longer cycles and are sent to off-line or cloud facilities.

By not aligning the applicable service level along with reviewing service-level objectives (SLOs) and service-level agreements (SLAs), situations where two parties wrongly assume what the other wanted or needed can be avoided. For example, IT or a service provider assumes that a given application requires the highest level of availability and data protection because that is what the business unit, customer liaison, advocate, or consumers indicated that they would like. However, the consumers or customer representatives thought that they would need the highest level of service without considering the cost ramifications or verifying what they actually needed. Upon review of what is actually required, there is sometimes a difference from the level of service being delivered. When questioned about SLOs or SLAs, business or IT services consumers may want to have the higher level of service, but some due diligence may show that they do not actually need it, and this can help stretch their budget.

The previous is an example of a disconnect between customer/consumer and IT services management. If IT understands their services and costs and works with their customers, the applicable level of capabilities can be delivered. In some cases IT services customers may be surprised to find that IT-provided services are cost effective when compared to cloud and MSP solutions on the same SLO and SLA services basis.




Figure 5.4 Reconciling and tuning data protection intervals and retention cycles.

5.4.3. Tiered Data Protection

Tiered data protection (Figure 5.5) is similar in concept to tiered hypervisors, servers, storage, and networks, in that different types of related technologies exist to meet various needs. The idea of resource tiering is to map the applicable technology or tool in such a way that it meets service requirements in a cost-effective manner. A key theme is to align data protection techniques to meet specific RTOs and RPOs along with other service criteria and cost needs.

Figure 5.5 shows an example of different applications or information services with various service criteria (e.g., RTOs and RPOs). Note that the applications shown along the top of Figure 5.5 do not necessarily correspond one to one with the various data protection techniques shown along the bottom. The important point is that some applications require RTOs and RPOs of zero or close to zero and need technologies such as synchronous replication, data mirroring combined with snapshots, or continuous data protection across multiple sites. Other applications may require clustered servers and highly available storage but can tolerate time delays associated with longer-distance replication or as a means to reduce cost for shorter-distance leveraged asynchronous replication. Applications that need small or zero RPOs will need to have data protection that corresponds to those requirements, while other applications may tolerate longer RTOs and RPOs with longer data protection intervals.

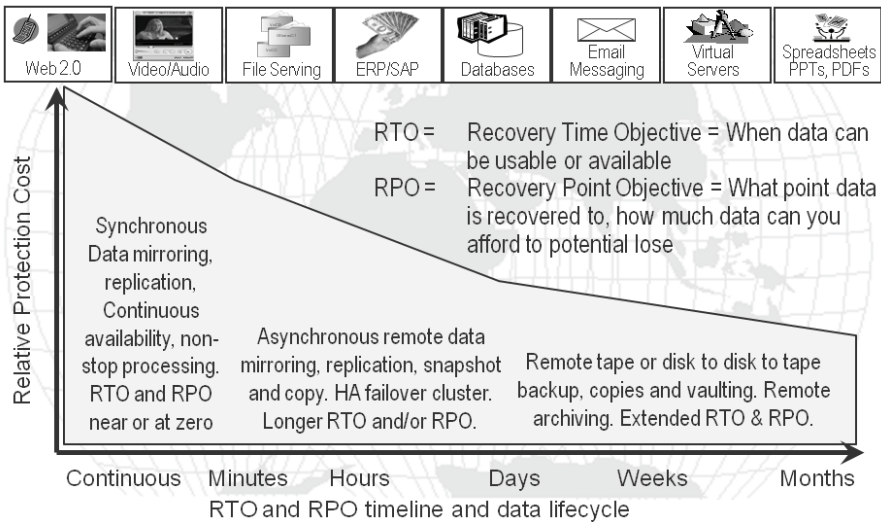


Figure 5.5 Tiered data protection.

5.5. Common-Sense Data Protection

Common-sense data protection (CDP) means complete data protection. Complete data protection means that all data is flushed from applications, databases, file systems, operating systems, and hypervisor buffers to a storage medium before copies (backup,

snapshots, and replication) are performed. The importance of quiescence (quieting) applications and capturing all data and information is to establish and preserve the state or transactional integrity of the data as of a given point in time.

Complete and comprehensive data protection architectures should combine multiple techniques and technologies to meet various RTO and RPO requirements. For example, virtual machine (VM) movement or migration tools such as VMware VMotion provide proactive movement for maintenance or other operational functions. These tools can be combined with third-party data movers, including replication solutions, to enable VM crash restart and recovery or basic availability. Such combinations assume that there are no issues with dissimilar physical hardware architectures in the virtualized environment. It is important to be aware of the motivators and drivers for data protection of a virtual server environment when creating the architecture.

Another aspect of common-sense data protection is that if data is important enough to be backed up or replicated, or if it needs to be archived for planned or possible future use, then the data is important enough to make multiple copies. Making multiple copies also means placing copies in different locations on multiple media—for example, a copy on disk locally for rapid recall, a copy at a managed service provider or cloud provider, and a “gold” or master copy on tape or some other medium, “just in case.” The idea is that, for critical information, reducing risk for what is most likely to occur and having cost-effective options is a top priority. This means having multiple tiers of data protection aligned to various needs, requirements, and threat risks. If all of your data is copied and protected in the cloud or at a managed service provider, what happens when (not if!) you lose access to that data? On the other hand, if your data is simply copied locally or at an alternate site, what happens if you lose access?

Having options that do not cost more than what the threat risk impact would impose on your organization enables the continued delivery of services. Aligning the right level of service to a given application’s or business function’s needs is an aspect of business impact analysis. Also keep in mind that loss of access to data is different than loss of data. For example, there have been incidents where customers are not able to get to their data, which at first may appear or be reported as data loss. In reality, the data is intact on some other medium or in a different location that is not yet accessible by the customer. In that type of situation the RPO may be zero or close to zero, meaning no data loss. However, if you cannot get to your data and your RTO requirements require no loss of access, then you cannot wait and need to initiate a recovery or restart. That recovery or restart may involve going to a backup copy of data, perhaps a recent snapshot or D2D copy or, worst case, to an old archive. The difference between the data from the backup and the data at the time of the loss of access may be an actual data loss. If your environment requires very low RTOs, then additional steps, discussed later in this chapter, should also be taken to avoid having to go to a stale backup or deep cold archive.

5.6. Virtual, Physical, and Cloud Data Protection

There are several approaches to achieve server virtualization, including Citrix/Xen, Microsoft Hyper-V, and VMware vSphere, as well as vendor-specific containers or

partitions. Many of the data protection issues are consistent across different environments, with specific terminology or nomenclature. Virtual server environments often provide tools to facilitate maintenance and basic data protection while lacking tools for complete data protection or BC/DR. Instead, virtual server vendors provide APIs, other tools, or solution/software development kits (SDKs) so that their ecosystem partners can develop solutions for virtual and physical environments. For example, solutions from VMware, Citrix, and Microsoft include SDKs and APIs to support pre- and postprocessing actions for customization and integration with Site Recovery Manager (SRM), or Microsoft Hyper-V Quick Migration.

Additional cloud, virtual, and physical data protection considerations include:

- RTO and RPO requirements per application, VM guest, or physical server
- How much data changes per day, application-aware data protection
- Performance and application service-level objectives per application
- The distance over which the data and applications need to be protected
- The granularity of recovery needed (file, application, VM/guest, server, site)
- Data retention including short-term and longer-term preservation (archive)
- Data usage and access patterns or requirements to meet business needs
- Hardware, network, or software dependencies or requirements
- Focus on doing more with less or doing more with what you have

Another consideration when comparing data protection techniques, technologies, and implementations is application-aware data protection. Application-aware data protection approaches ensure that all data associated with an application, including software, configuration settings, data, and the current state of the data or transactions, is preserved. To achieve true application-aware and comprehensive data protection, all data, including memory-resident buffers and caches pertaining to the current state of the application, needs to be written to disk. At a minimum, application-aware data protection involves quiescence of file systems and open files data to be written to disk prior to a snapshot, backup, or replication operation. Most VM environments provide tools and APIs to integrate with data protection tasks, including prefreeze (preprocessing) and postthaw (postprocessing) for application integration and customization.

5.6.1. Tools and Technologies

The basic tool for enabling data protection is common sense or, if you like jargon, common-sense data protection (CDP), leveraging such ideas as that any technology can fail if humans are involved. Another tenet is designing for maintenance and fault isolation or containment, maintaining data protection as a proactive piece of your data infrastructure strategy rather than just an afterthought. Knowing that technology can fail due to various reasons, the objective is to align different tools, techniques, techniques, and best practices to mitigate or isolate small incidents from cascading into larger disasters.

Another aspect of CDP is to make sure that all data is protected, including whatever is still in application, database, file system, or operating system buffers when a snapshot, CDP, or replication operation is performed; everything needs to be flushed to disk to be protected. Common-sense data protection also means balancing threat risks with the likelihood of a given scenario occurring and its impact to your organization or specific application against the cost. This means not treating all data or applications the same, and applying the applicable level of protection to a given threat risk and cost of availability that is needed.

Various tools and techniques can be used for enabling a flexible, scalable, resilient data infrastructure to support cloud, virtual, and physical environments. Tools include data and application migration or conversion, IT management along with asset discovery, tracking, and configuration management databases. Other tools include physical-to-virtual (P2V) migration, virtual-to-virtual (V2V) migration, and virtual-to-physical (V2P) migration, and automation such as using tape libraries that reduce the need for manual intervention. Policy managers, which can be located at different locations, can also help with automation of common tasks when an event occurs, or with manual intervention taking action involving other tools in a prescribed manner. Automated failover of software and redundant hardware including clustering or path managers for applications, operating systems, hypervisors, servers, storage, and networking are also tools for data protection.

Systems or storage resource management (SRM) and systems or storage resource analysis (SRA) tools are needed for insight and situational awareness, providing reporting along with proactive event correlation analysis. The importance of event correlation and analysis is to be able to identify where actual issues are, to avoid chasing false positives. False positives occur when a given diagnosis points to a technology or configuration that is then repaired, after which it is learned that it was not the real problem, but rather a result of the real issue. Change control and configuration management are also important tools and techniques for enabling resilient environments, to make sure things are configured correctly and tested before being put into production and thus to catch potential errors before they occur.

Additional tools and technologies for enabling data protection include:

- Application plug-ins for backup, snapshots, replication, and failover
- Data protection management (DPM) tools for tracking, alerting, and analysis
- Data protection coverage or exposure analysis tools
- Media tracking and management technologies
- Software management media such as tape read verification analyzers
- Archiving, backup/restore, CDP, snapshots, and replication tools
- Data footprint reduction tools including compression and de-duplication
- Test and diagnostic tools for servers, storage, and networks
- Security and encryption tools and key management
- Automatic rebuild along with RAID
- Dual or redundant I/O networking paths and components
- Change control and configuration management

- Testing and auditing of technology, processes, and procedures
- Routine background data, storage, and networking data integrity checks
- Network Data Management Protocol (NDMP) for protecting NAS devices
- API support including VMware vSphere, Microsoft VSS, and Symantec OST

Source-side tools, which are technologies used for collecting or gathering data to be protected while facilitating recovery or restoration, can reside on clients or on servers. Clients may be workstations, laptops, or hand-held devices as well as virtual and physical servers. In the context of data protection, servers' source-side tools include backup or data protection servers on which data is copied or staged before being sent to a local, remote, or cloud virtual or physical device. Servers in this context serve an intermediary role between the items being protected, such as database or application servers, remote client desktops, or mobile laptops.

Intermediary servers may be referred to as a backup or data protection appliances, gateways, or proxies and help off-load the actual source being protected. Both client and servers can implement data footprint reduction technologies including compression, de-duplication along with encryption, and network bandwidth management for optimization and media retention, and for tracking management. In addition, data protection servers can also use policy management functions to determine what is protected when, where, how, and why. For example, based on scheduled or event-based policies, a data protection server can notify another data protection tool to take some action or tell a VM, its guest, and applications that it is time to quiescence to gain a consistent and complete data protection operation.

Examples of source-side data protection include:

- Backup/restore tools, including agent, agentless, and proxy-based servers
- Application, file system, or operating system tools for snapshots and replication
- Database or other application journal and transaction log file shipping
- Archiving tools for databases, email, and file systems
- E2E DPM and SRA tools for situational awareness

Table 5.3 shows several common nomenclatures for data protection, where the source is usually a disk on a client or server to be backed up directly to a target, to an intermediary backup server or other storage system that, in turn, moves data to another location. For example, simply backing up a server, workstation, or laptop to an attached internal or external disk device would be D2D (disk-to-disk) or to a tape drive as D2T or to a cloud MSP as D2C. An example of a client or application server being backed up to a data protection staging device such as a backup server and then to another disk would be D2D2D.

In addition to source-side protection, which also includes intermediary backup or other data protection servers, the other part of data protection is the target to which the data is sent. Targets can be active in that they may be used for read or read and write by other applications at the same or different locations. Targets can also be passive, with data only stored until needed, when it is either restored or accessed as part of a failover

Table 5.3 Various Source and Target Data Protection Schemes

Acronym	Nomenclature	How Data Is Protected
D2T	Disk to tape	Moves data directly from the source disk to a tape device
D2D	Disk to disk	Copies from one disk to another internal or external disk
D2C	Disk to cloud	Data is backed up, copied, or replicated to a cloud or MSP
D2D2D	Disk to disk to disk	Data is copied to an intermediary disk and then copied to a target disk
D2D2T	Disk to disk to tape	Data is copied to an intermediary disk and then copied to tape
D2D2C	Disk to disk to cloud	Data is copied to an intermediary disk and then copied to a cloud
D2D2C2D D2D2C2T	Disk to disk to cloud to disk or tape	Data is copied to an intermediary disk and then copied to a cloud. Once at the cloud, MSP data is also protected on disk or tape

process. Targets can hold data for short-period data protection such as snapshots, replication, and backup or for longer-term protection, including both active and inactive or cold archives for preservation purposes.

Data protection targets may be local, remote, or cloud-based, leveraging fixed (non-removable) or removable media. Fixed media include solid-state devices (SSDs) and hard disk drives (HDDs) installed into a storage solution, removed only for repair and replacement. Removable media include removable hard disk drives (RHDDs), magnetic tape, optical CD/DVDs or portable FLASH SSD devices. Some targets may also be hybrids, with some media that stays fixed while others are exported or sent to a different physical location, where they may be left in cold storage until needed or placed into warm or hot active storage by being imported into a storage system at the remote site.

You might wonder why with modern data networks anyone would still ship data via portable or removable media. While networks are faster, not only is there also more data to move in a given amount of time, high-speed networks may not be available between the source and destination in an affordable manner. Physically moving data may be a one-time process of initially getting large amounts of data to a cloud, MSP, or hosting facility, or an ongoing process when moving it over a network is not a possibility.

Target or destinations for protecting and archiving data include:

- Backup servers in an intermediary role as both a source and a target device
- Data protection gateways, appliances, and cloud point-of-presence devices
- Shared storage, including primary and secondary storage systems
- Disk-based backup, archives, and virtual tape libraries (VTLs)
- Automated tape libraries (TLs) or automated tape systems (ATS)
- Cloud and MSP solutions supporting various protocols or personalities

Reliability–availability–serviceability (RAS) capabilities include redundant power and cooling, dual controllers (active/passive or active/active), and spare disks with automatic rebuild combined with RAID. Just as there are many approaches and technologies to achieve server virtualization, there are many approaches for addressing data protection in a virtualized server environment. Table 5.4 provides an overview of data protection capabilities and characteristics to address various aspects of data protection in a virtualized server environment.

Table 5.4 Data Protection Options For Virtual Server Environments

Capability	Characteristics	Description and Examples
Virtual machine migration	<ul style="list-style-type: none"> • Move VMs • Facilitate load balancing • Proactive failover or movement vs. recovery 	<ul style="list-style-type: none"> • Vmotion, Quickmigration, and others • May be processor architecture dependent • Moves VM's memory from server to server • Shared-access storage for BC/DR
Failover high availability (HA)	<ul style="list-style-type: none"> • Proactive VM movement • Automatic HA failover • Fault containment • RAID disk storage 	<ul style="list-style-type: none"> • Proactive move of VM to a different server • May require tools for data movement • Low-latency network for remote HA • Replication of VM and application data
Snapshots and CDP	<ul style="list-style-type: none"> • Point-in-time copies • Copies of VM state • Application-aware • In multiple locations 	<ul style="list-style-type: none"> • Facilitate rapid restart from crash event • Guest OS, VM, appliance, or storage based • Combine with other tools • For HA and BC/DR or file deletion
Backup and restore	<ul style="list-style-type: none"> • Application based • VM or guest OS based • Console subsystem based • Proxy server based • Backup server or target resides as guest in a VM 	<ul style="list-style-type: none"> • Full image, incremental, or file level • Operating system and application-specific • Agent or agentless backup • Backup over LAN to backup device • Backup to local or cloud device • Proxy based for LAN- and server-free
Replication	<ul style="list-style-type: none"> • Application based • VM or guest OS based • Console subsystem based • External appliance based • Storage array based 	<ul style="list-style-type: none"> • Application replication such as Oracle • VM or guest OS or third-party software • Application integration for consistency • Replication software or hardware • Storage system controller based replication
Archiving	<ul style="list-style-type: none"> • Document management • Application based • File system based • Long-term preservation 	<ul style="list-style-type: none"> • Structured (database), semistructured • Unstructured (files, PDFs, images, video) • Regulatory and noncompliance • Project data preservation for future use
DPM and IRM	<ul style="list-style-type: none"> • Data protection tools • Analysis and correlation • Backup and replication 	<ul style="list-style-type: none"> • VMware Site Recovery Manager • Data protection advisory and analysis tools • Various aspects of IRM and data protection

5.6.2. Virtual and Physical Machine Movement

Often mistaken, or perhaps even positioned as data protection tools and facilities, virtual machine movement or migratory tools are targeted and designed for maintenance and proactive management. The primary focus of tools such as VMware, vSphere,

VMotion, and Microsoft Hyper-V Quick migration, and others, is to move a running or active VM to a different physical server that has shared access to the storage that supports the VM without disruption.

For example, VMotion can be used to maintain availability during planned server maintenance or upgrades or to shift workload to different servers based on expected activity or other events. The caveat with such migration facilities is that, while a running VM can be moved, those VMs still rely on being able to access their virtual and physical data stores.

This means that data files must also be relocated. It is important to consider how a VM movement or migration facility interacts with other data protection tools including snapshots, backup, and replication, along with other data movers to enable data protection.

In general, considerations pertaining to live movement for VMs include:

- How does the VM support dissimilar hardware (e.g., Intel and AMD)?
- Can the migratory or movement tool work on both a local and wide area basis?
- Will your various software licenses allow testing or production failover?
- How many concurrent moves or migrations can take place at the same time?
- Is the movement limited to virtual file system-based VMs or raw devices?
- What third-party data movers' hardware, software, or network services?

5.6.3. Enabling High Availability

A common approach for high-availability data accessibility is RAID-enabled disk storage to protect against data loss in the event of a disk drive failure. For added data protection, RAID data protection can be complemented with local and remote data mirroring or replication to protect against loss of data access due to a device, storage system, or disk drive failure. RAID and mirroring, however, are not a substitute for backups, snapshots, or other point-in-time discrete copy operations that establish a recovery point.

RAID provides protection in the event of disk drive failures; RAID does not by itself protect data in the event that an entire storage system is damaged. While replication and mirroring can protect data in the event that a storage system is destroyed or lost at one location, if data is deleted or corrupted at one location, that action will be replicated or mirrored to the copy. Consequently, some form of time-interval data protection, such as a snapshot or backup, needs to be combined with RAID and replication for a comprehensive and complete data protection solution.

Techniques and technologies for enabling HA include:

- Hardware and software clustering
- Redundant networking paths and services
- Failover software, pathing drivers, and APIs
- Elimination of single points of failure

- Fault isolation and containment
- Clusters and redundant components
- Active/active and active/passive failover
- Change control and configuration management

Active/passive refers to a standby mode in which one server, storage controller, or process is active and another is in warm standby mode, ready to pick up workload when needed. When failover occurs, the warm standby node, controller, or server and associated software may be able to resume at the point of the disruption or perhaps perform a quick restart, rollback, or roll forward and resume service. Active/active refers to two or more redundant components that are doing work and are capable of handling the workload of a failed partner with little to no disruption. From a performance standpoint, care should be taken to ensure that performance is not impacted if an active/active two-node or dual controller sustains loss and failover of work moves to a surviving member.

For performance-sensitive applications for which service requirements dictate that no degradation in performance occurs during a failover, load-balancing techniques will be needed. If SLOs and SLAs allow short-term performance degradation in exchange for higher availability or accessibility, then a different load-balancing strategy can be used. The key point is that in the course of a storage controller, node, or server failover to a surviving member that provides availability, a subsequent chain of events must not be initiated due to performance bottlenecks, resulting in a larger outage or disruption to service.

Virtual machine environments differ in their specific supported features for HA, ranging from the ability to failover or restart a VM on a different physical server to the ability to move a running VM from one physical server to another physical server (as discussed in the previous section). Another element of HA for physical and virtual environments is the elimination of single points of failure to isolate and contain faults. This can be done, for example, using multiple network adapters (such as NICs), redundant storage I/O host bus adapters, and clustered servers.

Another aspect of HA is when and how new versions of software or configuration changes can be applied. Nondisruptive code load (NDCL) means that new software or configuration can be applied to an application, operation, and hypervisor, networking device, or storage system while in use with no impact. However, the code or configuration changes do not take effect until the next reboot or restart. Nondisruptive code load activation (NDCLA) enables the code to take effect on the fly. If your environment has no planned maintenance windows for scheduled downtime or service interruptions, then NCDLA may be a requirement rather than just a nice-to-have feature. If you do have routine schedule windows when service can be disrupted for a brief period of time for maintenance, then NCDL may be sufficient. Also keep in mind that if you have redundant components and data paths, upgrades can be done on one path or set of components while the others are in use. However, some environments will still schedule updates to one path while the other remains active to occur during scheduled maintenance windows, to avoid the risk of something unforeseen happening.

5.6.3.1. Why RAID Is Not a Replacement for Backup or Time-Interval Protection

RAID provides availability or continued accessibility to data on storage systems, guarding against a device or component—such as a disk drive—failure. Should something happen to the entire RAID storage system, it is a single point of failure if it is not being backed up, snapshots, and replications made to another location, or no other copies of stored data exist.

Another issue is that if the data on the RAID system becomes corrupted, without a point-in-time copy or snapshot, backup copy, or other copy on a different storage system, there is again a single point of failure. Some form of timed recovery-point copy needs to be made and combined with RAID. Likewise, RAID complements time-based copies by maintaining availability in the event a disk drive or component fails, rather than having to go to a backup copy on another disk, tape, or cloud.

Does replication by itself with no time-based copy protect against RAID failure? From the standpoint of maintaining availability and accessibility, the answer can be yes if a complete or comprehensive copy (e.g., all buffers were flushed) maintaining transactional integrity or application state is preserved. However, if data is deleted or corrupted locally, the same operation will occur on the local or remote mirror unless some time-interval copy is introduced. Simply put, combine RAID and replication with some form of time- or interval-based data protection as part of implementing a complete or comprehensive data protection strategy.

5.6.4. Snapshots and Continuous Data Protection

There are a number of reasons why snapshots, also known as point-in-time (PIT) copies, and associated technologies might be utilized. Snapshots are significant in that they create a virtual backup window to enable data protection when a physical backup window is shrinking or no longer exists. Snapshots provide a way of creating virtual time to get essential data protection completed while minimizing impacts to applications and boosting productivity. Different applications have varying data protection requirements, including RTO, RPO, and data retention needs.

Other reasons include making copies of data for test purposes such as software development, regression testing, and DR testing; making copies of data for application processing including data warehouse, data marts, reporting, and data mining; and making copies to facilitate nondisruptive backups and data migration.

Snapshots are a popular approach to reducing downtime or disruptions associated with traditional data protection approaches such as backup. Snapshots vary in their implementation and location, with some being full copies while others are delta (change)-based. For example, an initial full copy is made with deltas or changes recorded, similar to a transaction or redo log, with each snapshot being a new delta or point-in-time view of the data being protected. Another way snapshot implementations can vary is in where and how the snapshot data is stored on the same storage system or the ability to replicate a snapshot to a separate storage system. Space-saving

snapshots leverage redirect on writes to allow copies of snapshot volumes to be made and subsequently modified without the overhead of duplicating data. In the absence of a space-saving snapshot, if a 1-TB snapshot were copied three times, for development, for testing or quality assurance, and for decision support or business analytics, the result could be 4 TB of space required (1 TB original + 3 copies). Space-saving snapshots, which vary by vendor implementation, should reduce the storage space needed to a smaller data footprint consisting of the base amount of data, some small amount of overhead, and any changed data for each of the subsequent copies.

Because snapshots can take place very quickly, an application, operating system, or VM can be quiesced (suspended), a quick snapshot taken of the current state at that point in time, and then resume normal processing. Snapshots work well for reducing downtime as well as speeding up backups. Snapshots reduce the performance impact of traditional backups by only copying changed data, similar to an incremental or differential backup but on a much more granular basis. Snapshots can be made available to other servers in a shared storage environment to further off-load data protection. An example is using a proxy or backup server to mount and read the snapshots to construct an off-line backup.

For virtual environments, snapshots can be taken at the VM or operating system layer, with specific features and functionalities varying by vendor implementation. Another location for snapshots is in storage systems that have integration with the guest operating system, applications, or VM. Snapshots can also take place in network or fabric-based appliances that intercept I/O data streams between servers and storage devices. One of the key points is to make sure that when a snapshot is taken, the data that is captured is the data that was expected to be recorded.

For example, if data is still in memory or buffers, that data may not be flushed to disk files and captured. Thus, with fine-grained snapshots, also known as near or coarse continuous data protection (CDP), as well as with real-time fine-grained CDP and replication, 100% of the data on disk may be captured. However, if a key piece of information is still in memory and not yet written to disk, critical data to ensure and maintain application state coherency and transaction integrity is not preserved. While snapshots enable rapid backup of data as of a point in time (RPO), snapshots do not provide protection by themselves in the event of a storage system failure and need to be backed up to another device.

CDP and snapshot considerations include:

- How many concurrent snapshots can take place, and how many can be retained?
- Where is the snapshot performed (application, VM, appliance, or storage)?
- What API or integration tools exist for application-aware snapshots?
- Are there facilities for pre- and postprocessing functions for snapshots?
- Do the snapshots apply to virtual disks or physical disks?
- What is the performance impact when snapshots are running?
- How do the snapshots integrate with third-party tools including replication?
- What are the licensing and maintenance costs for the snapshot software features?
- When a copy of snapshots is made, is a full copy performed?

Initially, CDP entered the market as standalone products and was slow to be adopted. This is shifting, as is often the case when a new technology feature can be imbedded as additional functionality in currently deployed products. While there are still some purpose-built CDP solutions, the functionality has been added as a feature in many software and storage system products.

Keep in mind that with CDP, it is important to capture all information for a complete or comprehensive data protection plan. This means that for complete data protection with CDP, data from applications, databases, operating and file systems, as well as hypervisors' buffers must be flushed to storage as of a known state for transaction and data and data integrity. For example, for interfacing with VMware vSphere, Microsoft Hyper-V, Citrix Xen, SAP, Oracle Database, Microsoft SQLserver, Exchange, Share-Point, and other applications, it is important for data integrity for them to quiese to a known state and flush their buffers to disk. While capturing 100% of data on disk is good, if only 99.999% of the data to maintain application and data integrity is copied, with the other 0.001% still in a buffer, but that small amount of data is crucial for recovery, than that is a weak link in the recovery and data protection chain.

5.6.5. Backup and Recovery

Backups are a time-tested technique for making a point-in-time copy of data that can be used for many different purposes and have been the cornerstone of many data protection and BC/DR strategies for decades.

Some examples of different types of backups include:

- Full, incremental, or differential backup
- Image backup of physical storage or of VM virtual disk
- Bare-metal restore capable to both PM and VMs
- File-level backup or file-level restore from image backups
- File restoration from traditional backups as well as snapshot or CDP copies
- Application integration, for example, with Oracle RMAN or API tools
- Running on a PM, VM, appliance, or as part of a cloud MSP SaaS model
- Server vs. desktop or mobile device backup

Another form of backup is an object base where all systems or servers, applications, and data associated with a function are backed up. For example, an object backup could be made to ensure that everything associated with manufacturing or accounts payable or customer relations management (CRM) or a website are protected. An object backup can span multiple servers running different operating systems such as a database, unstructured files data, and their applications. The idea behind an object backup approach is that everything associated with that function or information service is protected and can be restored in a coordinated manner.

Replacing tape with disk-based solutions can help address tape-based bottlenecks, but it does not address the root issue involved with many backups. That root issue may be how backup software is configured or used along with the age or capability of the data protection tool. Another consideration is the paradigm shift of leveraging disk-

based backup instead of generating save sets or other backup copies. It is time to update how backups are done, leveraging different tools and technologies either as a near-term tactical solution or as part of a larger data protection modernization—for example, a shift from doing D2T or D2D backups throughout the day to leveraging snapshots and CDP combined with replication and a copy of log or journal files.

For applications or environments that need to retain tape-based backup processes or procedures, virtual tape libraries (VTLs) that emulate tape as well as combining policy-based management, replication, compression, and deduplication while providing disk-based access using NFS or CIFS are a great way to enable a transition to the future. For example, near-term leverage the VTL interface while processes and procedures and backup configurations are modified to use disk based NFS or CIFS access of the device. Then, when the conversion from a tape-based process is completed, that feature can be disabled or retained for legacy support as needed. In essence, the target data protection solution should leverage virtualization to bridge from what you are currently doing to the future while at the same time changing with you by supporting backup and archiving as well as serving as an ingestion point for snapshot copies or other forms of data movement. A primary tenant of virtualization is “abstraction providing transparency along with emulation,” and that has been used for a consolidation focus. This has resulted in a common belief that virtualization equals consolidation and consolidation means virtualization; there is, however, a much larger opportunity for virtualization beyond consolidation. Storage systems that support a VTL interface are examples of leveraging transparency along with emulation by making disks function like tapes to provide a bridge to the current software configuration.

In addition to optimizing the targets where data is sent for protection, another consideration is changing the protection and retention intervals. With disk-based data protection and data footprint reduction techniques, it is possible to keep more frequent backups on-line and readily accessible. The result of having more yet smaller copies of data for fast restoration is the reduction of the number of larger copies retained on other media or in different locations. Data protection and backup modernization then becomes a matter of finding a balance between having better granularity for faster restore of data while having a smaller data footprint without compromising SLOs or SLAs.

Caveats with virtual and physical backups include the fact that backing up many small files can result in lower performance measured in throughput or bandwidth as compared to moving large files. Depending on the type of environment and applications being supported, you could be working with different-sized files. For example, video or photo files can be large, but some applications (energy, seismic, crash test simulation, medical) generate many small time-sliced images (e.g., lots of TIFs, GIFs, or JPGs) that get spliced together for later playback.

Understanding how the size of files impacts data protection is important in order to determine how to configure resources to minimize performance impact or time delays. For example, tape drives work best when data can be streamed to them instead of starting and stopping, which results in “shoe shining” wear and tear. If many small files need to be protected, configure a solution in which those files are staged and able to be streamed directly to tape, disk, or cloud based resources while also using data footprint reduction techniques.

Another dimension to backups involves desktops, workstations, laptops, and mobile devices, including protecting remote office/branch office (ROBO) environments. Some enterprise software provides support for remote or non–data center backups, while other packages are optimized for those functions. Similarly, many cloud and managed service providers have solutions that can be used to protect remote and distributed devices to off-load primary data center solutions. Another approach for protecting remote devices is virtual desktop initiatives (VDIs), centralized applications that are simply presented on workstations or tablets along with terminal services. For some environments, such as call centers or environments with robust networking capabilities, shifting from heavy workstations or desktops to thin or light devices makes sense. For other users, workstations can be protected using tools that back up those systems up to a local or remote server, central location, or a cloud MSP service. The backup can be scheduled, or some solutions can run in a continuous or near-continuous mode, taking routine snapshots and protecting data on an event such as a change or time basis.

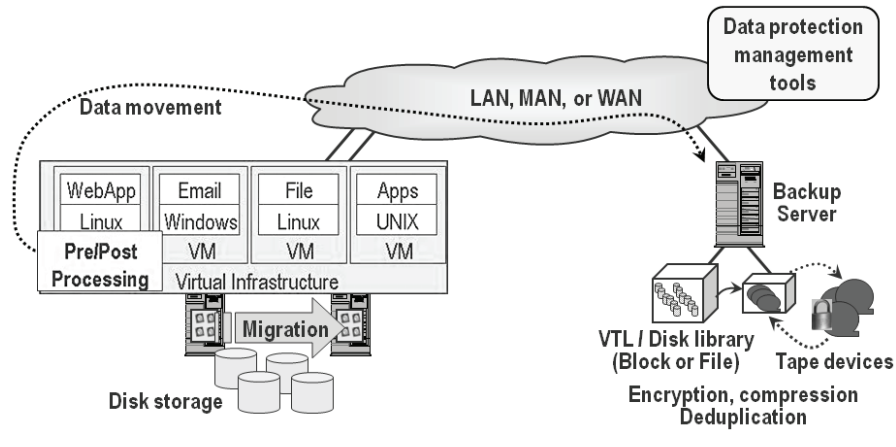


Figure 5.6 Agent-based backup over a LAN.

5.6.5.1. Agent-Based Data Protection

Agent-based backup, also known as LAN-based backup, is a common means of backing up physical servers over a LAN. The term agent-based backup comes from the fact that a backup agent (backup software) is installed on a server, with the backup data being sent over a LAN to a backup server or to a locally attached tape or disk backup device. Given the familiarity and established existing procedures for using LAN and agent-based backup, a first step for data protection in a virtual server environment can be to simply leverage agent-based backup while re-architecting virtual server data protection.

Agent-based backups, shown in Figure 5.6, are relatively easy to deploy, as they may already be in use for backing up the servers being migrated to a virtual environment. The main drawback to agent-based backup is that it consumes physical memory, CPU,

and I/O resources, causing contention for LAN traffic and impacting other VMs and guests on the same virtualized server.

Backup client or agent software can also have extensions to support specific applications such as Exchange, Oracle, SQL, or other structured data applications as well as handling open files or synchronizing with snapshots. One of the considerations regarding agent-based backups is what support exists for backup devices or targets. For example, are locally attached devices (including internal or external, SAS, iSCSI, FCoE, Fibre Channel or InfiniBand SAN or NAS disk, tape, and VTL) supported from an agent, and how can data be moved to a backup server over a network in a LAN-friendly and efficient manner?

Physical servers, when running backups, have to stay within prescribed backup windows while avoiding performance contention with other applications on that server and avoiding network LAN traffic contention. In a consolidated virtual server environment, it is likely that multiple competing backup jobs may also vie for the same backup window and server resources, including CPU, memory, and I/O and network bandwidth. Care needs to be exercised when consolidating servers into a virtual environment to avoid performance conflicts and bottlenecks.

5.6.5.2. Proxy-Based Backup

Agent- or client-based backups running on guest operating systems consume physical resources, including CPU, memory, and I/O, resulting in performance challenges for the server and LAN network during backup (assuming a LAN backup). Similarly, an agent-based backup to a locally attached disk, tape, or virtual tape library (VTL) will still consume server resources, resulting in performance contention with other VMs or other concurrently running backups. In a regular backup, the client or agent backup software, when requested, reads data to be backed up and transmits the data to the target backup server or storage device along with performing associated management and record-keeping tasks.

Similarly, on restore operations the backup client or agent software works with the backup server to retrieve data based on the specific request. Consequently, the backup operation places a demand burden on the physical processor (CPU) of the server while consuming memory and I/O bandwidth. These competing demands can and need to be managed if multiple backups are running on the same guest OS and VM or on different VMs.

An approach to addressing consolidated backup contention is to leverage a backup server and configure it as a proxy (see Figure 5.7) to perform the data movement and backup functions. Proxy backups work by integrating with snapshot, application, and guest operating system tools for pre- and postprocessing. As an example, VMware has replaced the VMware Consolidated Backup (VCB) tool with a set of data protection APIs. The APIs enable a VM or guest operating system and applications to be backed up by a proxy process. The proxy process reduces resource consumption (CPU, memory, and I/O) on the PM where the VMs exist, because the work is done via another server. For its part, Microsoft with Hyper-V, being based on Windows technology,

leverages Volume Shadow Services (VSS) copy and associated application VSS writers for integration.

Rather, it is an interface to VMware tools and enables third-party backup and data protection products to work. To provide data protection using VMware vSphere APIs, Microsoft Hyper-V VSS and DPM, or other hypervisor-based capabilities, third-party tools are leveraged. These third-party tools provide scheduling, media, and data protection management. Third-party tools also manage the creation of data copies or redirecting data to other storage devices, such as VTLs and disk libraries, equipped with data footprint reduction (DFR) capabilities including compression and data deduplication. Virtual machine virtual disk images (e.g., VMDK for VMware or HVDs for Hyper-V), depending on allocation and actual usage, may be sparse or hollow. This means that there can be a large amount of empty disk storage space that has been pre-allocated. The drawback is that extra time may be required to back up those files as well as allocate yet-unused disk space being occupied. The good news is that allocated yet unused disk space can lend itself well for thin provisioning and other DFR techniques, including compression and deduplication.

To help speed up backup or data protection operations, VMware has added change block tracking (CBT) into vSphere, which can be leveraged by third-party tool providers. CBT speeds up backup or data protection copy functions by keeping track of which blocks have changed from within the kernel and using a table that maps to the corresponding size of a VM. When a block changes, the vSphere kernel makes an entry in the corresponding table, which is a disk file. Data protection tools can make an initial copy, and then simply look up blocks that have been changed to reduce the amount of time required and the amount of data to copy.

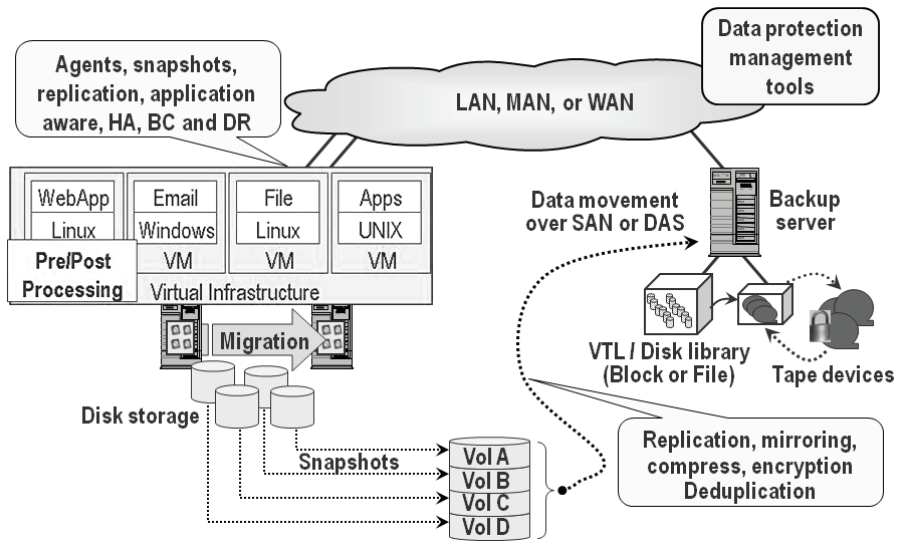


Figure 5.7 Proxy and API-based backup example.

In addition to off-loading the physical server during the proxy backup, LAN traffic is not impacted, as data can be moved or accessed via a shared storage interconnect depending on the specific VM implementation. Third-party backup and data protection software on a proxy server can also perform other tasks, including replicating the data to another location, keeping a local copy of the backup on disk-based media with a copy at the remote site on disk and on a remote off-line tape if needed.

5.6.5.3. Cloud and MSP Backup

For cloud, MSP, and remote backup, there are different options in addition to numerous vendors and service providers. The options are that as an organization you can acquire software from a vendor and deploy your own public, private, or hybrid cloud service, or you can subscribe to a service provider. As a VAR or service provider, you can also acquire software along with hardware resources and provision your own service or work with a larger provider who delivers the capability on your behalf. As a service provider, you can develop your own software running on your or someone else's servers at a hosting, MSP, or cloud site.

Cloud or MSP backups can replace what you are currently doing for backup, or they can complement your existing data protection environment. For example, you could replace your existing backup infrastructure including software and hardware with new tools supplied by the provider and optional on-site staging or caching hardware, if applicable. Some providers allow you to use your existing backup or data protection software, moving the data to the cloud with optional on-site or local copies. Some providers supply software or an appliance that sits at your location to collect information and facilitate transmission to their location or your choice of destinations to be stored.

Some questions and topics to look into regarding cloud and MSP backup providers include how they charge for the service—flat fee for unlimited capacity and bandwidth with no extra fees for access (updates or restores) and views (search, directories, catalogs)? Does the provider charge on a graduated pricing scheme, where there is a base monthly or annual fee plus a fee per gigabyte or terabyte stored? What are the optional or hidden fees in addition to base capacity usage? For example, do they charge for uploading or retrieving files, viewing or generating reports, bulk import and export? Does the provider offer different cloud storage service offerings for parking your data, including the option to select various geographies for regulatory compliance? If there is the option to select different back-end storage service providers, and are the fees the same or do they vary based on SLAs and locations?

Another consideration is how locked in are you to a particular provider. What are your options should you decide or need to switch providers? While your data may reside at a third-party site such as Amazon, is it stored in a format that can be accessible using some other provider's tool? For example, if you start using one service provider to back up your environment and specify Amazon S3 as the target storage pool, what happens if you switch your service to Rack space Jungle disk, which also supports Amazon as a pool?

While Jungle disk gives you the option of storing data at Rack space or Amazon, will it be able to leverage your old backups already stored within the Amazon S3 cloud, or will you have to convert or export and re-import your backups? You may find that to ensure coverage the best option is to maintain access to your old provider until those backups expire and you have sufficient coverage at your new provider. Other considerations include whether you can force your old backups to expire and what is involved in making sure that the provider has taken adequate steps to remove your data and your account profile. During the writing of this book, I switched cloud backup providers and had a period of overlap with additional protection via my normal D2D and D2D2D with a master copy sent off-site to a secure vault.

Additional questions and considerations regarding cloud and MSP backups include:

- In what format are backups stored; can they be accessed without special tools?
- How efficient is the solution in scanning or tracking changes to be backed up?
- What types of data footprint reduction technologies expedite backups?
- How much data can be backed up in a given timeframe with different networks?
- How much data can be restored in a given amount of time with your networks?
- What are your restoration options, including to alternate locations?
- Does the provider offer bulk physical media restoration (e.g., disk or tape)?
- Are there fees for access or bandwidth usage in addition to capacity charges?
- What types of management tools, including automated reporting, are available?
- Can you view your current charges or fees to determine future usage forecasts?
- What is the solution's scaling capabilities, and can it grow to fit your needs?
- Is the solution designed for ROBO or SOHO or mobile or SMB or enterprise?
- What operating system, hypervisors, and application are supported?
- What software requirements are needed; what does the provider supply?
- Do you require any additional hardware at your locations to use the service?
- Can the service support local copies made to existing backup devices?
- What are the security mechanisms, including encryption key management?
- Look beyond the basic cost per gigabyte to understand additional fees and SLAs.

5.6.6. Data Replication (Local, Remote, and Cloud)

There are many approaches to data replication and mirroring, shown generically in Figure 5.8, for local and remote implementations to address different needs, requirements, and preferences. Replication can be done in many locations, including applications, databases, third-party tools, operating systems and hypervisors on host servers (PMs), appliances or networking devices including cloud point-of-presences (cpops) or gateways, as well as primary, secondary, and backup targets such as VTLs or archive systems.

A general caveat is that replication by itself does not provide complete data protection; replication is primarily for data availability and accessibility in the event of a component, device, system, or site loss. Replication should be combined with snapshots and other point-in-time discrete backup data protection to ensure that data can be recovered or restored to a specific RPO. For example, if data is corrupted or deleted on a primary

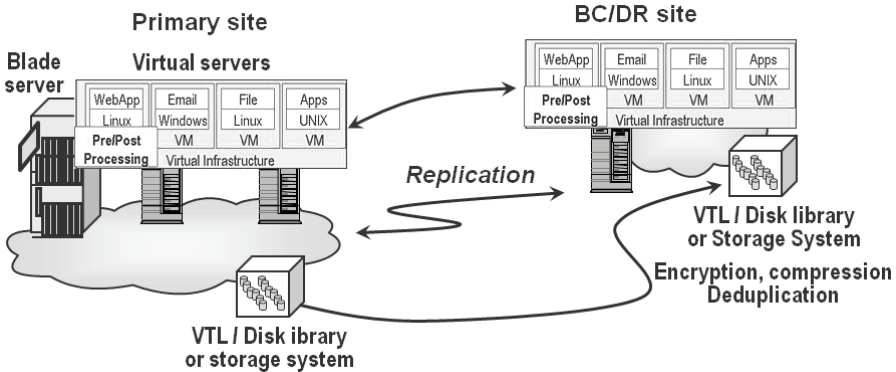


Figure 5.8 Data replication for HA, BC, and DR data protection.

storage device, replication will replicate the corruption or deletion to alternative sites, thus the importance of being able to recover to specific time intervals for rollback.

Considerations for replication include:

- Application integration with snapshots
- Local, metropolitan, and wide area requirements
- Network bandwidth and data footprint reduction capabilities
- Encryption and data security functionality
- Various topologies including one to one, many to one, one to many
- Homogeneous or different source and destination
- Management reporting and diagnostics
- How many concurrent replication streams or copies

Important factors of data mirroring and replication are distance and latency resulting in data delay and negative performance impacts. Distance is a concern, but the real enemy of synchronous data movement and real-time data replication without performance compromise is latency. The common perception is that distance is the main problem of synchronous data movement as, generally speaking, latency increases over distance. The reality is that even over relatively short distances, latency can negatively impact synchronous real-time data replication and data movement.

Distance and latency have a bearing on replication and data movement by impacting decisions on whether to use synchronous or asynchronous data movement methods. The trade-offs beyond costs are performance and data protection. Synchronous data transfer methods facilitate real-time data protection, enabling an RPO of or near zero. However, the trade-off is that, over distance or high-latency networks, application performance is negatively impacted while waiting for remote I/O operations to be completed. Another approach is to use asynchronous data transfer modes where a time delay is introduced along with buffering. By using a time delay and buffering, application performance is not impacted, as I/O operations appear to applications as having

completed. The trade-off with asynchronous data transfer modes is that while performance is not negatively impacted over long distance or high-latency networks, there is a larger RPO exposure window potential for data loss while data is in buffers waiting to be written to remote sites.

Consequently, a combination of synchronous and asynchronous data transfer may be used for a tiered data protection approach, for example, using synchronous data transfer for time-critical data to a reasonably nearby facility over a low-latency network, with less critical data being replicated asynchronously to a primary or alternative location farther away.

5.6.7. Data Protection Management

Data protection management (DPM) has evolved from first-generation backup reporting technology to incorporate multi-vendor and cross technology domain capabilities. In addition, present-generation DPM tools are evolving to manage multiple aspects of data protection beyond basic backup reporting, including replications, snapshot, BC/DR compliance coverage, file system monitoring, and event correlation. Some DPM tools are essentially reporting, status, or event monitoring facilities, providing passive insight into what is happening with one or more data protection IRM focus areas. Other DPM tools can provide passive reporting along with active analysis and event correlation, providing a level of automation for larger environments.

Cross-technology domain event correlation connects reports from various IT resources to transform fragments of event activity into useful information on how, where, why, and by whom resources (servers, storage, networks, facilities) are being used. In virtualized environments, given the many different interdependencies, cross-technology domain event correlation becomes even more valuable for looking at end-to-end IRM activities. The increase of regulatory requirements combined with pressure to meet service levels and 24x7 data availability has resulted in data protection interdependencies across different business, application, and IT entities. Consequently, timely and effective DPM requires business and application awareness to correlate and analyze events that impact service and IT resource usage. Business awareness is the ability to collect and correlate IT assets to application interdependencies and resource usage with specific business owners or functions for reporting and analysis. Application awareness is the ability to relate IT resources to specific applications within the data protection environment to enable analysis and reporting.

Although an environment may have multiple tools and technologies to support IRM activities, DPM tools are evolving to support or coexist with management of multiple data protection techniques including backup (to disk or tape or cloud), local and remote mirroring or replication, snapshots, and file systems. Key to supporting multiple data protection approaches and technologies is the ability to scale and process in a timely manner rapidly increasing large amounts of event and activity log information. At the heart of a new breed of IRM tools, including DPM solutions, are robust cross-technology resource analysis and correlation engines to sift disparate data protection activity and event logs for interrelated information.

and contained using HA techniques such as dual adapters, RAID, and active/active failover. Should a more serious incident occur, failover can occur to the secondary or BC site, where access can continue or resume or where recovery can quickly occur to a given point of time.

Should an even more serious incident occur that results in the primary or secondary BC site and their resources not being available, near-line or off-line data at a warm or cold DR site can be leveraged. In Figure 5.9 an additional step is added in which both the primary and secondary or BC site are actively used, with the production load balanced between them. These two sites complement and protect each other. A third site is a warm or cold site where a minimal number of systems are in place and to which critical data is periodically copied.

The idea is that this third or near-line DR site provides a means for recovery at a location some distance away from the primary and secondary sites. Building on this example, there can be a fourth site where off-line data is copied or sent. This site represents where tape or other removable media are shipped and data either remains on the media or is migrated to a cloud accessible storage environment. This site houses data that will be rarely, if ever, accessed—essentially a last-resort source for data restoration.

5.7.1. Expanding from DR to BC, Shifting from Cost Overhead to Profit Center

BC and DR are often seen as cost overhead items for IT and the businesses that they support. This is because capital and operating budget money is spent on hardware, software, services, networking, and facilities for a capability that it is hoped will never be used other than for testing purposes. On the other hand, if those resources can be safely used for production business, those costs can be absorbed as being able to drive workload, thereby reducing per-unit costs. If your BC or secondary site is also your primary or last-resort DR copy, caution needs to be exercised not to contaminate or compromise the integrity of that environment, by keeping DR, master, or gold copies of data physically or logically isolated from on-line production or active data.

5.7.2. Using Virtualization and Clouds to Enhance Data Protection

Most virtualization initiatives undertaken at present are focused on consolidation of heterogeneous operating systems on underutilized servers. Another aspect is to address physical desktops and workstations with virtual desktop infrastructure (VDI), in part for consolidation but also to simplify management, data protection, and associated cost and complexity. The next (or current) wave of server, storage, and desktop virtualization is expanding the focus to include enabling agility and flexibility. This combines the tenets of consolidation with an emphasis on utilizing virtualization to enable dynamic management of servers and dynamic data protection, for example, using virtualization to support redeployment of servers for workload changes and to provide

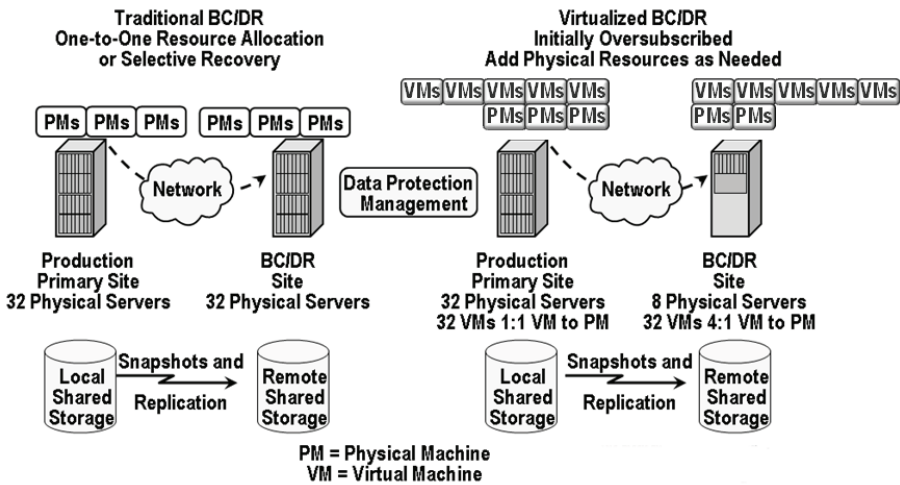


Figure 5.10 Leveraging virtualization to enable HA, BC, and DR.

transparency. In this scenario, consolidation continues to be a driver. However, there is also an emphasis on leveraging virtualization as a tool for applications, servers, and storage that do not lend themselves to being consolidated but can benefit from business and IT IRM-enabled agility, including enhanced performance, HA, DR, and BC.

Virtualization can be used in many ways, including consolidation, abstraction, and emulation, to support load balancing and routine maintenance as well as BC and DR. On the left in Figure 5.10, a traditional BC/DR environment is shown, with dedicated physical resources or selective applications being recovered, or a combination of both.

Challenges include dedicating extra hardware on a one-to-one basis and selecting which servers and applications are recovered to available physical resources. Complexity is involved in maintaining BC/DR plans. This includes testing of configuration changes along with associated costs of hardware, software, and ongoing operational costs of power, cooling, and floor space. Other issues and challenges include difficulties in testing or simulating recovery for training and audit purposes and inefficient use of available network bandwidth, inhibiting the amount of data that can be moved in a timely fashion.

On the right side of Figure 5.10, a solution is shown that leverages virtualization for abstraction and management in which each physical server is converted to a VM. However, the VM is allocated a physical machine, such as a server or server blade, in a one-to-one manner. In the case of a disaster, or for BC or training and testing purposes, multiple VMs can be recovered and restarted on a limited number of PMs, with additional PMs being added as needed to boost or enhance performance to required service-level objectives. To improve data movement and enhance RPO and RTO, reduce the data footprint to boost data movement and data protection effectiveness using a combination of archiving, compression, de-duplication of data being moved or replicated, space-saving snapshots, data replication, and bandwidth optimization. Data protection management tools are used to manage snapshots, replication,

and backup and associated functions across servers, storage, and network and software resources.

Benefits of server virtualization include more efficient use of physical resources; the ability to dynamically shift workloads or VMs to alternative hardware for routine maintenance, HA, BC, or DR purposes; support of planned and unplanned outages; and the enablement of training and testing of procedures and configurations. In addition to supporting HA, BC, and DR, the approach shown on the right of Figure 5.10 can also be used proactively for routine IRM functions, for example, shifting applications and their VMs to different physical servers on-site or off-site during hardware upgrades or replacement of servers or storage.

A variation of the paradigm in Figure 5.10 uses virtualization for abstraction to facilitate provisioning, configuration and testing of new server and application deployments. For example, in Figure 5.11, on the left side of the diagram, multiple VMs are created on a physical machine, each with a guest operating system and some portion of an application. During development and testing and to support predeployment IRM maintenance functions, the various applications are checked on what appears to be a separate server but in reality is a VM.

For deployment, the various applications and their operating system configurations are deployed to physical servers as shown on the right of Figure 5.11. There are two options, one being the deployment of the applications and their operating system on a VM allocated on a one-to-one basis with a physical server as shown. The other option is to convert the VM, along with the guest operating system and application, to run on a physical server without a virtualization layer. In the first example, virtualization is used for abstraction and management purposes as opposed to consolidation. An underlying VM enables maintenance to be performed as well as giving the ability to tie into a virtualized BC/DR scheme as shown in Figure 5.11.

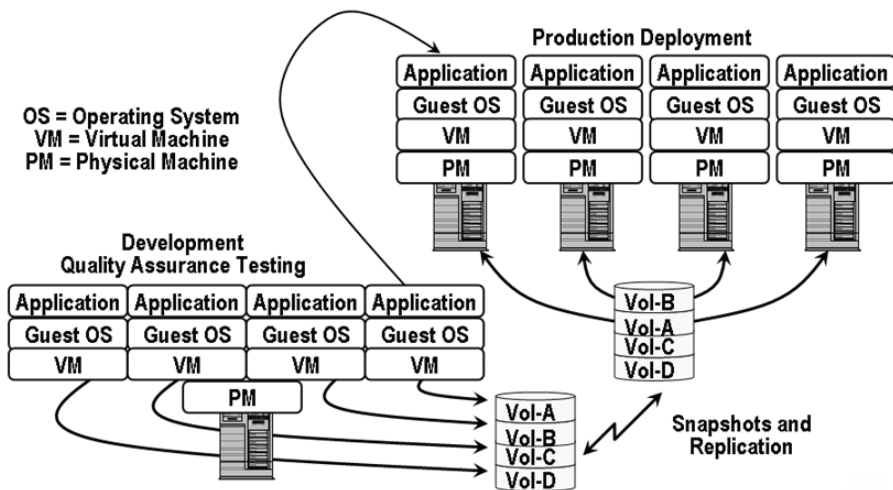


Figure 5.11 Utilizing virtualization for abstraction and server provisioning.

5.8. Data Protection Checklist

There is no time like the present to reassess, re-architect, and reconfigure your data protection environment, particularly if you are planning on or have already initiated a server virtualization initiative. Cloud and virtual server environments require real and physical data protection. After all, you cannot go forward from a disaster or loss of data if you cannot go back to a particular point in time and recover, restore, and restart, regardless of business or organization size.

Some common data protection best practices include the following:

- Some applications, including databases, support automatic log file shipping.
- Most disasters are the result of a chain of events not being contained.
- Leverage common sense and complete or comprehensive data protection.
- Verify support of USB-based crypto for encryption keys with VMs.
- Know what limitations exist for your software license for BC or DR testing.
- RAID is not a replacement for backup; it provides availability.
- Mirroring or replication alone is not a replacement for backup.
- Use point-in-time RPO based data protection such as snapshots with replication.
- Maintain a master backup or gold copy.
- Test restoration of data backed up locally and from cloud services.
- Employ data protection management tools for event correlation and analysis.
- Data stored in clouds needs to be part of a BC/DR and data protection strategy.
- Have a copy of data placed in clouds also in an alternative location for BC/DR.
- Combine multiple layers of protection and assume that what can break will break.
- Determine budget, time frame, tolerance to disruption, and risk aversion.
- Decide which solutions are best for different applications.
- Investigate whether redundant network paths share a common infrastructure.
- Seek out experienced help for assessment, validation, or implementation.

5.9. Common HA-, BC-, and DR-Related Questions

Is tape dead? Tape is alive and continuing to be developed as a technology; however, its role is shifting from routine backup to long-term archiving. D2D backup and data protection combined with data footprint reduction techniques continue to coexist with tape, resulting in more data being stored on tape than in previous history. The key take-away is that the role of tape is shifting to that of long-term or cold data preservation.

What are the barriers to using a cloud or virtualized environment for data protection? There are different constraints associated with the amount of data to be protected, time windows, network bandwidth, RTO and RPO, and budgets. Something else to consider is how your software licenses work or can be transferred to a BC as well as DR site, along with your ability to use those for testing purposes.

Are BC and DR only for large organizations; how can smaller organizations afford them? MSP and cloud-based services along with other technologies that can be installed on servers and workstations are making BC and DR more affordable. As an example, I

have a small business and, practicing what I preach, have implemented a multitier data protection strategy including D2D, D2D2C, and D2D2D on a local, remote, as well as with removable technologies. I leverage a cloud backup MSP where encrypted data gets sent even while I am traveling (I have done backups from commercial aircraft using Gogo WiFi), as well as having local copies on disk. Additionally, I have a master copy off-site in a vault that gets routinely updated using removable hard disk drives. There are many different tools, with more on the way, some which will be available by the time you read this. Check my blog and website for news, announcements, and discussions on related topics, trends, and techniques.

What are some key steps or questions to ask when choosing an on-line or cloud MSP for backup or archive services? Balance the cost or fees of the service with the available functionality, SLAs, hidden fees for accessing your data, import or export charges, options for what locations or regions where your data can be stored, and reporting. For example, I get a daily backup report via email from my service provider that I can check manually or set up a script to scan for exceptions. Also look into how your data is reduced using data footprint reduction techniques before transmission, to either move more data in a given amount of time, or with less network bandwidth capability. Also test how long restores take, to avoid surprises when time may be of the essence, and look into options to get larger quantities of data restored in a shorter period of time.

5.10. Chapter Summary

HA, BC, DR, and backup/restore are changing with evolving and maturing techniques and technologies. Clouds and virtualization need to be protected, but, at the same time, they can be used for enhancing protection.

General action items include:

- Avoid treating all data and applications the same.
- Apply the applicable level of data protection to required needs.
- Modernize data protection to reduce overhead, complexity, and cost.
- Combine multiple data protection techniques in a cost-effective manner.
- Don't be afraid of cloud and virtualization, but have a plan.

Vendors include Acronis, Amazon, Aptare, Asigra, BMC, Bocada, CA, Cisco, Citrix, Commvault, Dell, EMC, Falconstor, Fujifilm, Fujitsu, HDS, HP, i365, IBM, Imation, Inmage, Innovation, Iron Mountain, Microsoft, NetApp, Oracle, Overland, Quantum, Quest, Rackspace, Platespin, Rectiphy, Seagate, Sepaton, Solarwinds, Spectralogic, Sungard, Symantec, Veeam, and VMware, among others.

The bottom line: For now, if you are putting data into any cloud, have a backup or a copy elsewhere. Likewise, if you have local or even remote data, consider using a cloud or managed service provider as a means of parking another copy of backups or archives. After all, any information worth keeping should have multiple copies on different media in various venues.

Complements of StorageIO

This chapter download from the book “Cloud and Virtual Data Storage Networking” (CRC Press) by noted IT industry veteran and Server StorageIO founder Greg Schulz is complements of The Server and StorageIO Group (StorageIO). Learn more about the techniques, trends, technologies and products covered in this book by visiting storageio.com and storageioblog.com and register for events and other promotions. Follow us on twitter @storageio or on Google+ among other social media venues.

Brouwer *Storage Consultancy*
Seminar for Storage Professionals with Greg Schulz
May 7th, 8th, 9th 2012 in Nijkerk Holland
[Click here to learn more](#)



Visit storageio.com/events to see upcoming seminars and activities

Cloud and Virtual Data Storage Networking has been added to the Intel Recommended Reading List (IRRL) for Developers. Click on the image below to learn more about the IRRL.



The Recommended Reading List is a valuable resource for technical professionals who want to thoroughly explore topics such as software threading, wireless technologies, power management, and more. Dozens of industry technologists, corporate fellows, and engineers have helped by suggesting books and reviewing the list.

Learn more about Cloud and Virtual Data Storage Networking (CRC Press) by visiting storageio.com/books

To become a chapter download sponsor or to discuss your other project opportunities contact us at info@storageio.com.

#1 VM Backup

Veeam Backup & Replication for

Hyper-V

is HERE!

[Learn More](#)

VEEAM

StorageIOblog.com site sponsor