



An Approach to Enterprise Backup and Recovery Service Levels

by Brett Foy
Technology Architect
Ahead

April 7, 2011

Executive Summary

Often a request is made for a backup and recovery SLA for services described as > 99.9% successful. Is this reasonable or even achievable? It depends on what is actually desired.

- Does this refer to a 99.9% success rate for daily backups?
- Is the SLA referring to restores? Does it refer to a 99.9% success rate for requested restores? Based on what measurement of success?

Specific to backup success, in any enterprise-level organization if you are looking at the entire server infrastructure footprint, a > 99.9% backup success rate is very unlikely to be achievable. If a subset of critical servers and applications is identified that truly requires this level of backup success, it is possible to deploy such an infrastructure, given appropriate investment and limited scope/scale. However, it may not be practical. An attempt to expand this requirement to hundreds or thousands of servers would require an extraordinarily high investment in both finance and resources.

Specific to recovery success, the SLA request is murkier. If someone demands a 99.9% success rate on a restore, is a restore of two-week-old data really a “Success?” By language yes, but realistically a two-week-old data restore would most likely fail the standard of being relevant and within a specific recovery point objective.

This paper outlines general backup and recovery concepts, describes the infrastructure and process components involved in backup and recovery, and provides insight into focus areas that are critical to defining and accepting SLAs. (While using backup and recovery infrastructure for disaster recovery may be critical in many instances, this paper only briefly touches on disaster recovery.) Throughout this paper, the term “customer” is intended to be the end user of the IT services organization. The IT services organization can either be an Enterprise IT shop serving its business unit customers and/or an IT service provider serving an end user/customer.

In summary, IT infrastructure can be built to any level of required availability—at a trade-off of cost and complexity. When reviewing options for high-availability backup infrastructure, decisions should be made about realistic requirements for SLAs that provide adequate protection for the business, but can be done at reasonable costs on a tiered service catalog basis. Establishing a services catalog for backup and recovery availability requirements on a tiered basis allows for cost-effective deployment of advanced recovery capabilities without the extraordinary expense associated with building an availability metric too high for most of the infrastructure that does not require it.

Building a Measurable Backup/Recovery SLA

To define a success metric for restores, RTOs and RPOs must be defined on an infrastructure, application, or server basis. Depending on customer demands for RTO and RPO, different technologies are deployed that will vastly change the economic investment required for the solution. Additionally, depending on scope and complexity, simple operations of the backup and recovery process and the controls in place around that process will also change depending on requirements.

RTO is the recovery time objective: if a data restore is required, how much time can the customer wait for the data to be restored and made available to the application owner for recovery of the system? It is very unlikely that any customer infrastructure would have a single RTO for all systems. Usually key or tier-one applications and core infrastructure will have a relatively short RTO, mid-tier backup departmental applications will have medium range RTOs, and other systems, such as development, will have very long RTOs.

RPO is the recovery point objective: how much data can be lost? For example, say a company relies on tape as the only means of backup/restore. If tape backups occur at 11 p.m., but the system fails at 7 p.m. the next evening, the RPO would effectively be 20 hours. The data would be brought back to the point of the last good backup and data or changes

after that point in time would have to be manually recreated. Usually key or tier-one applications and core infrastructure will have a relatively short or near-instant RPO, mid-tier backup departmental applications will have medium range RPOs, and other systems, such as development, will have very long RPOs or even no RPO at all because they are easier and cheaper to recreate than to back up and restore.

What's being backed up? What does the SLA apply to?

Typically a service catalog will be required in which tiers of backup and recovery service can be defined with RTO and RPO characteristics and have an SLA assigned to each tier. A service catalog approach allows a relatively expensive, low-RTO/RPO solution to be designed and deployed for only the most critical systems, with less expensive solutions deployed to support applications with more tolerance for higher RPO/RTOs.

Within most organizations, successful backups are reported as the number of “successfully completed backup jobs” during a given backup cycle or window. With adequate reporting tools and capabilities, this metric provides the most accurate snapshot of pure backup success. It is also the most likely indicator of future restore capability because a successful backup indicates that all of the required supporting infrastructure and processes are functioning correctly. A report of a failed backup provides a logical starting point for troubleshooting the failed client. Other metrics exist, but looking at non-host-specific environmental factors may provide over-generalization or fail to supply the granularity required to assist in the troubleshooting process.

What constitutes a very high backup success rate?

Once you ask for 99%+ success in any endeavor, it is always good to level set what that means. First, let's look at the use case demanding a very high backup success rate, based on successful vs. failed jobs. A 99.95% success rate for backups on a nightly basis means you can have *no* failures in a 1,000-server environment, as one failure for any reason immediately takes you to 99.9%.

As an infrastructure scales in size into thousands of servers, it becomes progressively more difficult to stay in the 99% range. In fact, you need 2,000 or 4,000 servers with only one or two failures per night to attain and maintain a 99.95% rate. These metrics are illustrated in Table 1.

# of Servers	# of Failures	Percent Success
1,000	1	99.90%
1,000	5	99.50%
1,000	10	99.00%
1,000	30	97.00%
.	.	.
.	.	.
.	.	.
4,000	2	99.95%
4,000	20	99.50%
4,000	40	99.00%
4,000	120	97.00%

Table 1: Relative failures vs. percent success metrics

Factors to consider

To understand what it takes to achieve high success rates in backup infrastructure across an entire data center, as opposed to an application suite or specific server, requires understanding two areas:

- all of the infrastructure and application components/complexities
- process, controls, inventory, staffing

Backup Infrastructure Components

A 97% success rate as outlined in Table 1 would allow 30 failures in a 1,000 server infrastructure on any given night. This number seems high, but look at the components and integration involved to make sure a backup works successfully.

Server: tape software, server OS, file system, network (connectivity, bandwidth, throughput, IP, FC, etc.)

Infrastructure: network port, network routing, network congestion, network throughput /capacity—or—SAN, SAN configuration, etc.

Tape: physical tape drives, physical robots/libraries, physical tape media—or—virtual tape emulation/backup to disk

Operations: scheduling, physical resource availability (drives, bandwidth), software version control, monitoring, reporting, and alerting, etc.

Software: reporting, RMAN, snaps, clones, hot-backup modes, scripting, exports, dumps, etc.

Every one of the components listed above must work seamlessly all the time for a backup to work. For example, take a typical Microsoft Windows server with locked files that won't back up: the backup software will report a "server backup failure" if even a single locked file was not backed up.

In most tape infrastructures, errors do happen. They happen every night, and they relate to one or more of the points listed above, oftentimes masked by other issues or other points. The true key to managing the success of backups is in responding to a failure. Failing one backup for one night, while not good, can be tolerated as long as the problem is researched and resolved prior to the next backup window, if not in the same window. Neglecting to address failures in a system over a period of time is a tremendous risk and exposure point. Lack of process, control, and staffing will weaken an organization's ability to respond in a timely manner, causing inability to execute both backups and restores.

Process, Controls, Inventory, and Staffing

Establishing an application- or infrastructure-wide SLA necessarily dictates that the organization actually knows what needs to be backed up. Processes must be in place and routinely followed for each of the following.

Server life-cycle management

- Backup must be part of whatever official provisioning and deprovisioning processes exist.
- Controls must be in place to prevent unmanaged server sprawl and deployment; e.g., no servers should be entered into service (physical or virtual) without following a change/implementation process that includes provisioning backup software, scheduling, and media.
- If a general infrastructure SLA is established for “all servers,” it is critical that these controls are in place so all servers are added or removed from the inventory as required and that a server failure with backup never happens.
- General maintenance of servers and images must be maintained.
 - Drivers/tape clients on servers must be maintained at level.
 - Images and server deployment process must be maintained at level as part of image and server deployment process.

Problem management and resolution

- Appropriate staffing levels must exist within an organization to ensure ongoing operations.
 - Problem detection/resolution must occur within one missed backup window.
 - Inability to respond only compounds backup problems over time. If enough resources do not exist to resolve problems as they occur, catching up will be difficult, if not impossible. This will have a long-term negative effect on SLA attainment.
- Appropriate tools must be deployed.
 - Tools need to be automated, not only reporting on backup/restore success, but assisting in problem identification and resolution. They must be considered authoritative.
 - Error conditions should not be suppressed. It is common to suppress some types of error conditions, such as locked files. However, suppressing errors can mask bigger issues, such as thousands of locked files.

Infrastructure Options for RTO/RPO

When infrastructure components and process/people are working well together, it is possible to build and manage a backup-success SLA. Still, as with any technology implementation, many technical options must be considered.

As previously mentioned, one aspect of measuring a backup SLA is to look at the backup software reporting. Without valid and successful backups, there is no way to have a successful and valid restore.

However, running a successful backup infrastructure is only half the picture. The other half is the recovery process. A backup previously listed as successful that fails during a recovery operation would certainly be considered a failure. A backup that provides a recovery point in time that is too old or is not useful doesn't work for most organizations.

When a service catalog is developed, infrastructure assumptions need to be made. The first is based on the RTO requirements for the application. The decision of whether or not an application must be returned to service within an hour, 4 hours, 8 hours, or even within 24 hours will most definitely have an impact on the type of backup or backups required, as well as the infrastructure required to support those backups.

The second assumption is based on the RPO requirements for the application. The decision as to "how much data loss is acceptable" is critical and also plays an important role in determining the type of backup process and infrastructure.

Table 2 shows a hypothetical organization with 1,200 servers, possible tiers, and backup SLAs, as well as appropriate RTOs and RPOs as defined by the organization.

Table 2 illustrates two basic principles. The first, based on the backup-success SLA, is the number of failed backups that can occur within the SLA, based on server install base. The second, for RTO and RPO, provides metrics for how quickly the application service must be restored and to what point in time, from a data loss perspective.

Tier	# of Servers	Backup SLA	# of Permissible Failures	RTO	RPO
Tier 1	100	99%	1	< 4 hours	Near-zero loss
Tier 2	200	98%	4	< 8 hours	< 4 hour loss
Tier 3	700	97%	21	< 24 hours	24 hour loss
Dev/Test	200	95%	10	> 24 hours	> 24 hour loss

Table 2: Example service catalog/tiering approach

There are a number of technology decisions required to enable the types of RTOs and RPOs listed in Table 2. The next section outlines concepts for data protection and recovery that should be considered to meet different RTO and RPO points. To determine the best technical recommendations to meet the above, as well as to support RTO and RPO requirements, many different technologies may need to be deployed, working with all the infrastructure and server components listed earlier in this document.

Data Replication

- For an application with near 0 recovery time or recovery point objectives, **data replication** is often chosen to ensure another copy of the data “as/is” is always available either on-site or off-site to quickly re-attach existing or disaster equipment to that data to immediately bring it up with a minimum RTO and RPO.
 - Synchronous replication provides an absolute guarantee that data is in two locations, but can have an impact on application performance. Committed writes must take place in two physical locations prior to a write commit being returned to the host.
 - Asynchronous replication provides “almost” real-time replication of data—either on-site or off-site—to ensure that data exists in a near-time state online in multiple locations. The amount of “out of synch” between primary and target is a function of distance, latency, and the technology solution chosen.
 - One problem inherent to both synchronous and asynchronous replication is that any online data corruption or logical application data issue will replicate with the data and cause both copies of the

data to be affected. For this reason, it is recommended that additional backup technology (clones, backup to disk, or tape) always be deployed to protect against data corruption.

- **Snaps** provide a point-in-time snapshot of the data that can be used to recover—either to the primary application server or to a secondary host.
 - Snap technology uses minimal space to track changes from the time of the snap. Snaps are not full volume copies. They rely on a good copy of the original data always being available and a mechanism to track changes to that data over time.
 - Snaps are convenient, but have disadvantages based on growth, space, and “impact of use” on the original data volumes if they are used. Specifically, if a snap is mounted to another host for use or backup, there is a very real risk that the additional host I/O against this snap will impact performance of the production volume, since both hosts are actually reading from the same production volume.
 - Snaps will protect against logical data corruption, because they can be restored to the original data source. However, because snaps rely on original source data, hardware data corruption (where drives are damaged or a source LUN is physically corrupted) will cause a snap to become inoperative. For this reason, it is recommended that additional backup technology always be deployed, usually by mounting the snap and backing up to disk or tape during a non-critical production window when the read I/O for the backup will not affect production performance.

- **Clones** provide a point-in-time snapshot of the data that can be used to recover, either to the primary application server or to another host.
 - A clone is a full copy of the original data set, which may be expensive, as it is another copy of the source data. The clone method does not rely on the original volume to remain healthy or online, and use of a mounted clone will not typically impact the production volume.
 - Architected correctly, a clone will protect against both logical and hardware corruption. When a clone is placed on separate physical disks within an array, it will not be affected by a physical corruption event, and as a point-in-time copy, it will protect against logical

corruption. Many organizations rely on a combination of clones and replication to provide multiple copies of protected data in multiple locations. They may or may not decide to then back up this data to disk or tape.

Continuous Data Protection (CDP)

- CDP provides real-time data journaling and logging, which enables rollback and RPO to a specific I/O or data transaction.
 - CDP provides the most granular protection and can be implemented either for local or remote recovery capability.
 - Because it stores every I/O and also journals or provides a “track” of every write to support the point-in-time rollback capability, CDP can be very expensive. For this reason, most CDP implementations retain a small amount of online CDP capability and rely on other types of backup/recovery for data loss or recovery past a certain recovery point.

Backup to Disk (B2D)

- Over the last five years, backup to disk has become, in many organizations, the primary means to accomplish nightly backups within a specific time window with high success.
 - During this time the cost of B2D technology has decreased. At the same time, storage capacity and specialized technology, such as data deduplication, has provided a massively scalable backup architecture that, in most cases, cannot be matched by pure tape infrastructure.
 - Many enterprises use different forms of B2D to provide a quick online backup to increase speed of both backup and restore. Most organizations then move backup data off-site for DR and regulatory requirements via B2D replication or physical tape.
 - There are four primary types of backup to disk—two newer technologies based on deduplication, and two more traditional solutions—that emerged when the price point of B2D began to compete with physical tape infrastructure. Most new solutions on the market today tout some type of cost-saving deduplication features. These are typically source- or target-based, but some have hybrid

solutions using both functions or an intermediary client to perform the global deduplication prior to writing to a target device.

- **Source-based deduplication technology** uses client (host)-based software to send only unique data to a central store.
 - In most implementations, there is a global store of data that contains one copy of each of the many blocks of data found in the enterprise. The central store maintains a catalog of these blocks, and when a distant client encounters new data blocks, it inexpensively checks with the global store to determine if the data block is unique and does indeed require transmission for backup.
 - Source-based data deduplication has many specific use cases that provide much value to the Enterprise IT organization.
 - **Remote Office/Limited Bandwidth:** Source-based deduplication allows an entire remote office backup to take place over a very small WAN infrastructure due to reduction in total backup capacity transmitted to the central store.
 - **NDMP/File Serving:** Many enterprise organizations have network attached storage CIFS/NFS infrastructure in place, with tens of millions of files that are ill-suited to be backed up by most other backup technologies. Using source-based deduplication enables quick backup of only changed and unique data on a nightly basis. This is unlike other solutions that require full volume image-based backups or complete file system scan and backups.
 - **Virtualized Infrastructure:** Using source-based deduplication in a virtualized server infrastructure will typically provide a significant reduction of hypervisor resources (network, CPU, memory) and will dramatically shorten time for backup windows to complete. Reducing the impact of backup workload and time required for backups of virtualized infrastructure will generally allow most organizations to achieve higher server consolidation ratios, which will have a direct impact on the cost, TCO, and ROI of any server virtualization initiative.
 - Most source-based deduplication technologies provide additional data protection methodology through replication of the global content

store to another off-site deduplication appliance. In some cases, limited tape-out support exists, but for organizations with an absolute requirement for full physical tape cloning of the entire backup infrastructure, source-based-only solutions can pose some challenges.

- **Target-based deduplication technology** primarily uses a physical appliance approach to transmitting data from a client to a target device in which the data is processed and stored in a deduplicated fashion. There are two primary types of target-based deduplication: in-band and postprocess.
 - **Post-processing target-based deduplication** transmits data from the client to the target array at very high speeds, stores it on disk, and then deduplicates it via post-backup processing. This requires more disk than other methods for two reasons.
 - To accommodate the I/O workload of multiple full backup streams, many disks are required to serve I/O.
 - Because deduplication happens after the backup, significant buffer space is required to hold the full backup prior to deduplication and space reclamation.
 - Additionally, in most implementations of post-processing target-based deduplication, any off-site replication for data protection will be delayed until all post-processing is complete, which is a potential data protection issue to be taken into consideration.
 - **In-band target-based deduplication** analyzes data as it is transmitted from the client to the target array and deduplicates the data prior to writing it to disk.
 - This requires significantly fewer disks as there is considerably less disk I/O and no buffer requirement for non-deduplicated storage.
 - With advancements in CPU capabilities over the last two to three years, many in-band disk appliances can process deduplication at line speed without slowing down or otherwise limiting the speed with which data can be streamed to the backup appliance. Today in-band disk appliances take advantage of high-speed 10 GB networking infrastructure cores.

- Target-based deduplication is poised to become the de facto standard for enterprise data centers. The ability to handle massive volumes of data streamed into the appliance and deduplicated prior to being written to disk will allow organizations to scale their backup capacity to match the explosive growth of enterprise data without significant requirements for massive additional spend in infrastructure, floor space, and physical tape media.
- Some vendors have combined source-based deduplication with target-based deduplication and can now deliver the best of both worlds specific to use case and data workload in any given infrastructure.
- Most source/target-based deduplication technologies provide additional data protection methodology through replication of the global content store to another off-site deduplication appliance. In almost all cases, full tape-out capability is available for organizations with an absolute requirement for full physical tape cloning of the entire backup infrastructure. Because a target-based appliance is usually written to by a traditional tape catalog software package, that package can natively “clone” data to tape for off-site removal as required.
- **Virtual tape appliance technology** came to market in the mid-'00s, changing the game because it enabled the introduction of disk-based backup by emulating physical tape infrastructure. Virtual tape could be implemented without requiring any significant changes to the existing backup software or infrastructure.
 - Virtual tape offered enterprise organizations unprecedented scale and growth capability to support an ever-expanding data workload requirement.
 - Virtual tape typically offered no features to reduce capacity requirements other than tape-based compression capabilities, usually gaining a 2x advantage over pure backup to disk (that is, backup to disk pools without emulating tape).
 - Virtual tape solutions offer significant throughput and can handle almost any demanding enterprise backup workload given an appropriate configuration and enough spindles to handle I/O.

However, with the advent of 10 GB Ethernet and target-based deduplication appliances, this use case for virtual tape becomes less of a differentiator.

- Because pure backup to disk and virtual tape solutions do not offer single instance or other deduplication technologies, they continue to grow massively in scale, size, and expense. They will ultimately be phased out or integrated into deduplication appliances.

Physical tape backup

Physical tape backup has been the data center standard for as long as the distributed data center has existed. Tape backup relies on physical tape drives and physical tape media to store backup data and often includes large tape libraries and even robotic installations to handle tape mobility in a massive configuration requiring dozens or more tape drives and thousands of tapes.

- Over the years, physical tape infrastructure has scaled to meet the demands of a massive workload. The latest technology in tape drives and tape media provides significant capacity and tape speeds to ensure that tape is still, in many cases, considered a viable platform. However, over the next decade it is uncertain whether physical tape cartridge media can scale to match the unprecedented data growth being felt by most organizations.
- There can be many reasons for an organization not to shift from physical tape to virtual tape:
 - regulatory statutes or lack of clearly defined regulatory permission to move away from a physical tape media
 - industries that require unalterable off-site media for litigation, discovery, or other purposes
 - sunk cost in existing technology
 - primary backup may be B2D, but there is no DR site for online replication, so off-site tape cloning is required
 - primary backup may be B2D, but organizational policy demands physical tape off-loading

- However, many organizations have either made or are in the process of making the switch to a pure tapeless infrastructure for powerful reasons.
 - Backup and restore failures are significantly more common with physical tape than with backup to disk solutions.
 - Physical tape media ages with use, becomes stale, and does not undergo data integrity checks over time (e.g., while in storage).
 - Failure of a single tape in a large tape pool, required for a significant restore, will require fallback to older tapes. This means more potential data loss
 - Physical tape drives, tape libraries, and tape robots all have mechanical components that are more prone to failure with less redundancy than physical disk drives in B2D appliances.
 - Going tapeless eliminates security risk for data loss due to lost tapes. Many companies have lost millions of dollars, had significant brand damage, experienced customer trust issues, and accrued government fines due to loss of tapes.
 - With continued advancement of technology capability and an ever-decreasing cost/GB price point, any cost advantages that physical tape ever had over backup to disk are already at breakeven or close to it.
 - In many organizations, physical tape simply cannot manage the bandwidth and throughput requirements to provide backup within back windows or even full tape cloning with a 24-hour window.

To understand the data protection mechanisms and physical infrastructure required to support an RPO, Figure 1 illustrates the best use case of the technologies discussed above for three RPO use cases: near 0, same day, and up to 24 hour.

Recovery Point Objective

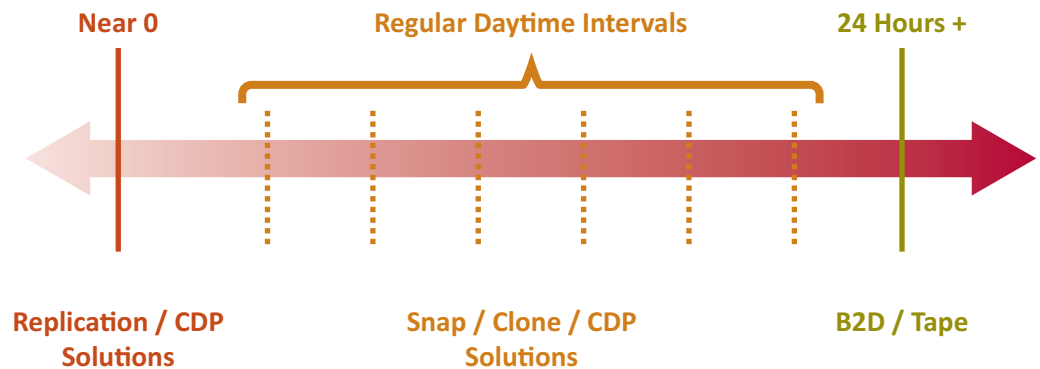


Figure 1: Understanding technology application to achieve RPOs

To understand the data protection mechanisms and physical infrastructure required to support an RTO, Figure 2 illustrates the best use of technology to support required RTOs.

Recovery Time Objective




Figure 2: Understanding technology application to achieve RTOs

These figures illustrate types of recovery point and time objectives and the relative durations/tolerances for recovery. Both of these figures assume the complete server restore of a sizable system. Individual file level recovery

and system results may vary. (This is intended to be a representative example of the types of solutions available and the general use case of each.)

Looking back at the hypothetical organization infrastructure and service catalog identified in Table 2, an application of technologies as outlined might look something like this:

Tier	RTO	RPO	Technology Considerations	
Tier 1	< 1 hour	Near-zero loss	Data replication and CDP	Replicated target-based deduplication
Tier 2	< 8 hours	< 4 hour loss	Data replication	Replicated target-based deduplication
Tier 3	< 24 hours	24 hour loss		Target-based deduplication, physical tape cloning
Dev/Test	> 24 hours	> 24 hour loss		Target-based deduplication, physical tape cloning, or no backup

Table 3: Understanding technology application to build a service catalog

In any production data center infrastructure, a combination of technologies will be required to meet most application backup and recovery requirements. It is possible to look at the problem of providing backup and recovery on a two-dimensional basis focused simply on the actual ability to back up and restore a client within a specific RTO or not-granular daily RPO. However, it is unlikely that this approach by itself will achieve the desired results. Looking at a third dimension of capabilities around local and remote replication—potentially including snaps, clones, or CDP—will be required in many use cases to create a consistent or point-in-time data set to be presented in a state that the data can be successfully backed up and restored. The ultimate question relates to the data and application recovery requirements, from which all design considerations can be taken into account either on an application basis or on a service catalog level basis.

Summary and Recommendations

Prior to defining, recommending, or agreeing to an SLA in the backup/recovery stack, many considerations must be taken into account that will drive what types of SLAs should be agreed to and what can reasonably be achieved. Key decision points include:

- organizational change and asset management processes and maturity
- status and use of application/infrastructure service catalog
- determination of appropriate RTO and RPO per catalog tier
- design and application of technology to specific use cases, based on business drivers tied to availability, recovery, and data sensitivity
- existing tools and reporting processes

In summary, it is unlikely that a single SLA can be created that has any meaning for an organization's ability to measure success for backup and recovery of the data center infrastructure. Establishing a very high 99.5%+ SLA will be unreachable for an entire multi-thousand server infrastructure based on the essential nature of backup and recovery technologies and organization capabilities. On the other hand, simply setting a "low" SLA as a general measure of the entire infrastructure is also inappropriate. This would introduce risk to critical applications, potentially not affording them the protection they need and enabling them to be lost in the shuffle of the rest of the infrastructure.

The responsible course is to develop a tiered backup for the organization's applications based on its application service catalog. If no application service catalog exists, then one should be created for this purpose. Developing tiers and measuring each tier based on application or infrastructure criticality will help keep focus and tight measurement on critical systems, while still allowing an appropriate level of management and response with more reasonable measurement for less critical systems. Establishing a service catalog with backup SLAs and RTO/RPOs will enable creation of a measurable, defensible, and ultimately cost-effective infrastructure that can be charged back as appropriate, while guaranteeing data recovery capabilities as required.

Brett Foy is a data center technology architect at Ahead, where he designs and architects advanced data center technologies around virtualization, computing, networking, storage, and backup/recovery. Prior to joining Ahead, Brett worked at a large IT manufacturer where he was responsible for the divisional Virtualization Data Center Practice as a practice manager. He also spent five years in a technology consulting and architecture sales support role. Prior to moving into the consulting and architecture role, Brett was an AVP for a multinational global insurance company, responsible for managing the North American data center infrastructure, focused on data center virtualization and consolidation around many technologies, including server, backup, desktop, email, security, web hosting, and SAP.

About Ahead

Ahead is purpose built to enable companies to understand and take advantage of today's revolutionary technologies. These technologies drive the next-generation data center, which improves agility, availability, and responsiveness, and takes out cost. Our technologists, consultants, and sales team help clients evaluate the potential technical, operational, and financial benefits of these innovative solutions. Where the benefit is adequate, we work to properly architect, procure, and deploy these solutions in a timely and effective manner.

*Ahead employs a unique methodology called **THINK | LOOK | PLAN | MOVE**.*

*In the **THINK** phase Ahead equips clients with knowledge of current technologies and trends, and helps clients understand the state of their current environment.*

*In the **LOOK** phase Ahead helps identify the client's technology and business objectives, architects alternatives, and quantifies total cost of ownership.*

*In the **PLAN** phase Ahead develops strategic and tactical plans, with a focus on reducing costs by compressing deployment time frames.*

*Finally, during the **MOVE** phase Ahead works with clients to execute the plans to deploy the future-state.*

For more details, visit www.ThinkAheadIT.com.

The material in this document is the proprietary information of Ahead, LLC.

All products, trademarks, and copyrights herein are the property of their respective owners.

©2011 Ahead, LLC. All rights reserved. Revised 4/7/11.