



Addressing the Broken State of Backup with a New Category of Disk-Based Backup Solutions



Vol 2 Issue 1

- 2 Addressing the Broken State of Backup with a New Category of Disk-Based Backup Solutions
- 6 About ExaGrid
- 7 From the Gartner Files: Best Practices for Repairing the Broken State of Backup

Featuring research from

Gartner

Addressing the Broken State of Backup with a New Category of Disk-Based Backup Solutions

Today, IT organizations are faced with a number of challenges when managing backup processes, including the need for faster backup, restore, tape copy, and recovery times. The first challenge is ensuring that backups occur in the shortest amount of time and are not disruptive to IT operations. Another challenge is maintaining a timely backup window regardless of data growth. Yet another challenge is ensuring that data is readily available for rapid restore, tape copy, and recovery.

Gartner, in its August 2013 report titled, “Best Practices for Repairing the Broken State of Backup,” agrees. “Growth of data, proliferation of new applications and expanding production usage of virtualization, combined with the users’ need for continual data access, impact how storage managers can use traditional backup methods.”

Also in the report, Gartner states, “A legacy backup approach does not scale well, as it cannot handle increased demands for protecting today’s data sources and volumes to meet the backup window.”

Therefore, the challenges of effectively and efficiently managing unrelenting data growth and extended backup windows require a break from traditional backup storage architecture design. New technology advancements in backup solutions are among the leading ways to improve the data protection infrastructure. These advancements include a next-generation scale-out architecture, data deduplication for backup storage, WAN-efficient replication for disaster recovery, and improvements in virtualized server backup. This report looks at the challenges and best practices in use today, and examines the modern backup solution that will solve these challenges forever.

Migrating from the Legacy Data Backup World to a Modern Architecture with ExaGrid

ExaGrid addresses all of today’s challenges with its scalable disk-based backup with data deduplication appliances. Its approach shortens the backup window up front and maintains a fixed-length backup window over time. ExaGrid provides for the fastest restores, tape copies, and instant recoveries. It works with incremental, incremental forever, synthetic full, and virtual machine backups. ExaGrid’s deduplication enables long-term retention using a minimal amount of disk. Its product is simple to use, tackling the complexity issue head on. With its award-winning customer support, ExaGrid offers the best support in the industry.

How ExaGrid Does It

Deduplication of data is compute intensive; as data grows, so does the corresponding workload. To accommodate that growth, both compute resources and capacity must be added to keep the backup window from expanding. ExaGrid’s disk-based backup appliances with data deduplication all come with disk capacity, processor, memory, and bandwidth. Additional appliances are simply added into the GRID as the volume of data increases. This unique approach ensures that as data (and the corresponding need for more deduplication power) grows, both compute resources and capacity needs are met. The results

New England Law | Boston

New England Law | Boston had concerns about the evolution of its existing backup strategy, which was becoming less than optimal from both a disaster recovery and operational perspective. This, as well as the ability to support increasing service demands, led the institution to rethink its practice and investigate a more appropriate strategy.

“We were facing operational limitations with our previous strategy, such as capacity management and lengthy backup times which ran anywhere from 24 to 30 or more hours.

“The ExaGrid system was cost effective and more scalable than the EMC Data Domain product, which for a similar price point, limited us in total capacity and required significant labor and planning just to extend us to the next level. We also liked ExaGrid’s approach to the deduplication process which focuses on backing up the data as quickly as possible, helping to meet and exceed our backup windows.”

Derek Lofstrom
Senior Network Engineer

are fast backups for a short backup window and a fixed backup window over time.

Cost-Effective Scalability in the Age of Big Data

The runaway growth of data in organizations not only creates cost and complexity in primary storage and production applications, but it also has a profound impact on backup and recovery operations. As a result, how disk-based backup architectures scale should be a key consideration when selecting a vendor and a product. Many vendors have great products in many categories, but most have missed some key elements in the way disk backup should scale. Adding deduplication to disk is just the beginning. But for many vendors, that is where the innovation stops.

Combining disk and deduplication creates a backup and recovery appliance that is cost effective as it uses the least amount of disk. However, simply combining disk and deduplication assumes that backup and recovery are only a storage problem. But is that really the case?

The answer is no. Backup and recovery are more than just a storage problem. In fact, these challenges are a:

- Data movement problem. Significant data volumes must be moved within a predefined backup window.

- Data processing problem. Data needs to be processed to be stored in deduplicated form.

As for the challenge of data storage, deduplication enables more backup data to be stored in far less disk space.

Both the data movement and the data processing problems only get worse with continued data growth. Because most vendors view backup as purely a storage problem, they implement scale-up architectures for growth. Scale-up architectures (shown on the left in Figure 1) provide a fixed amount of compute resources up front but enable only capacity to be expanded. Naturally, if the capacity of a system involved in data movement and data processing is expanded but its compute resources are not, all processes will take longer. The result is an expanding backup window and increasing time to deduplicate and replicate data. Ultimately, these architectures mean forklift upgrades and technology obsolescence.

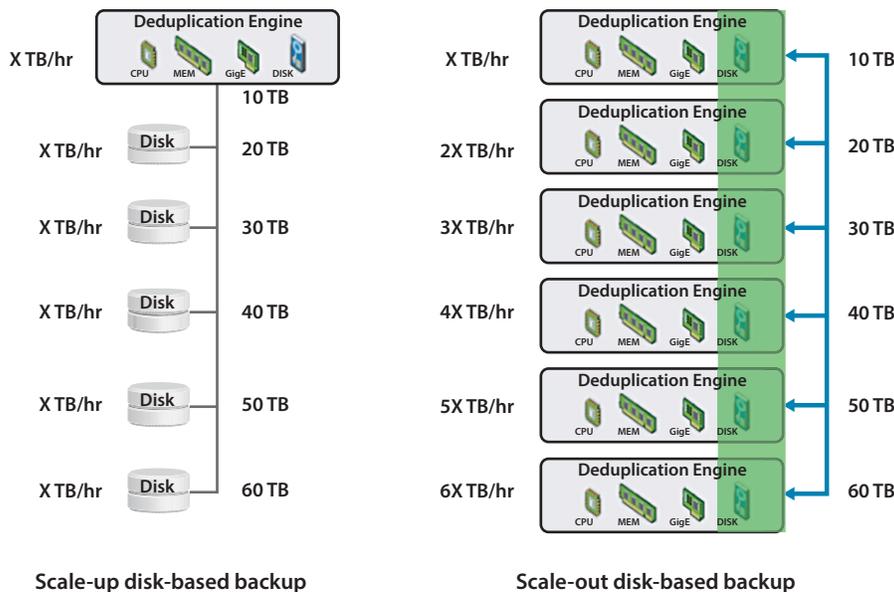
Alternatively, scale-out architectures (shown on the right in Figure 1) expand by adding not only capacity but also compute resources for each increase in workload. Therefore, all critical processes remain short, despite data growth. Further, a scale-out architecture eliminates forklift upgrades and technology obsolescence.

Salvation Army

“We liked ExaGrid’s approach to data deduplication. Because the deduplication process is performed after the data is backed up, the network and backup servers don’t get bogged down and backups run as quickly as possible.”

Michael Levine
Technology Research and Assessment Manager

FIGURE 1 Scale-up vs. Scale-out Architecture



Source: ExaGrid

NSM Insurance Group

NSM Insurance Group stores 58TB of data in 1.4TB of disk. In the past, they had been backing up to tape, but rapid business growth and a move to virtualization translated into an explosive increase in data, and as a result, a search for a new, more reliable backup method. NSM installed an ExaGrid system in its main datacenter and backs up nearly 40 virtual servers to the ExaGrid system through Veeam. The firm sees dedupe ratios as high as 41:1 for SQL data and 58.5TB of data, consuming 1.375TB of disk space using ExaGrid. NSM is currently retaining eight months of SQL data.

“Prior to installing the ExaGrid system, we were dealing with long backup times and reliability issues due to the vast amount of data Veeam was pushing to the unit. We considered upgrading the NAS unit, but in the end, we decided that a new solution with built-in data deduplication would serve us well.

“Our backup speed is fantastic. Our weekly full backups used to run from Friday night to Monday morning (60+ hours), but now they finish in approximately 17 hours. The most impressive part is that a lot of our backup jobs run simultaneously, and we haven’t even started using the additional network ports on the ExaGrid, which would undoubtedly improve our performance even more.

“We’re in the insurance business, and ExaGrid’s GRID architecture acts as a policy against having to ‘rip and replace’ our backup infrastructure when our backup needs increase ... [ExaGrid and Veeam] work extremely well together and have essentially solved our backup issues.”

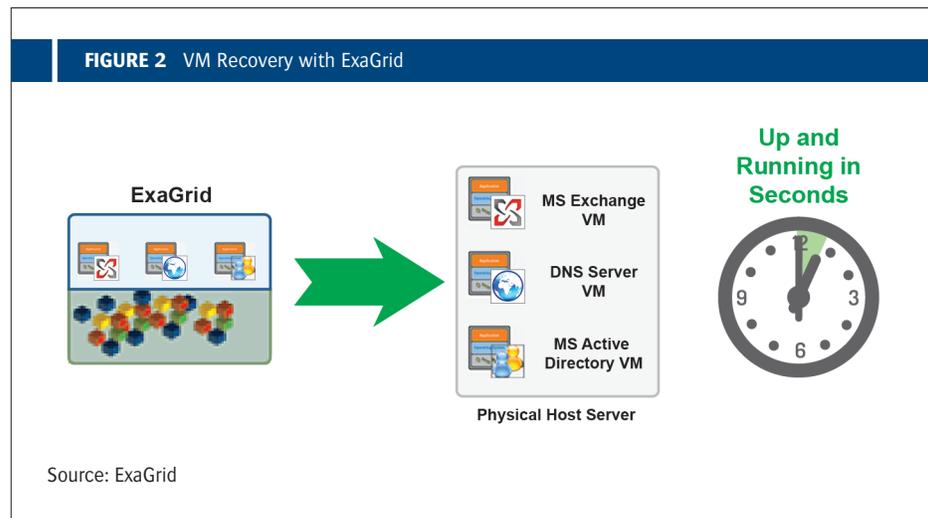
Todd George
Technology Manager

ExaGrid is the only scale-out disk-based backup architecture on the market today. As such, ExaGrid is the only vendor that solves the data movement, data processing, and data storage challenges all in one product. In addition, ExaGrid has a unique landing zone that stores the most recent backups in their full un-deduplicated form so that data is ready for fast restores, fast tape copies, and fast instant recoveries. All other disk-based backup solutions deduplicate the data inline or on the fly which means they only store deduplicated data that needs to be put back together or rehydrated for every restore, tape copy, or instant recovery request which greatly impacts restore and recovery times.

The bottom line: If all three of these challenges are not resolved, the host of issues organizations face—including backup window expansion, lack of scalability, and technology obsolescence—will not be fixed.

Optimized Recovery in a Virtual World

Today’s datacenters show the tremendous growth in virtual server and storage infrastructure use. As a result of this growth, a host of both mature and new backup software applications bring to market features that can take full advantage of the optimized backup and recovery methods virtual infrastructure vendors offer. These features include critical recovery techniques made possible by virtual environments such as instant virtual



machine (VM) recovery. Now found in a variety of backup applications, instant VM recovery involves the direct use of the VM's backup copy to instantly recover from a production outage.

However, not all disk-based backup architectures enable organizations to fully exploit these features. In many cases, the method of deduplication can interfere with the function of these critical optimizations. This is especially true with instant VM recovery.

Because many organizations are backing up VMs to disk rather than tape, the VM can simply be initialized immediately from disk (shown in Figure 2). This dramatically reduces recovery time over traditional restore methods.

For this type of restoration to work successfully—delivering on the promise of recovery in seconds to minutes—how the backup copy of the VM is stored on the appliance is critical. ExaGrid's landing zone relies on the capability to land the un-deduplicated VMs in native form, and then keeps the most recent versions in an optimized cache for instant recovery (shown in Figure 2).

Unlike ExaGrid, most vendors of disk-based backup with deduplication appliances deduplicate these virtual images inline/on the fly. This means that only deduplicated copies of the VMs exist on disk; there are no intact, easy-to-read copies. Booting a VM that has been significantly deduplicated dramatically increases the time it takes that machine to be production-ready and reduces the machine's performance when it is ready; in some cases, the machine outright fails. The reason is simple: The virtual image has to be rebuilt (or rehydrated) from potentially millions of small

blocks of data that are scattered all over the appliance (shown in Figure 3).

Regarding the challenge of complexity, Gartner's "Best Practices for Repairing the Broken State of Backup" report says, "Organizations also complain about the complexity of their backup solutions. The feeling is that too much time, expertise, and effort are spent keeping the current recovery systems afloat. Organizations would like backup to be a process that requires far less supervision and administrative attention, and for higher backup and restore success rates to be easier to achieve."

Support

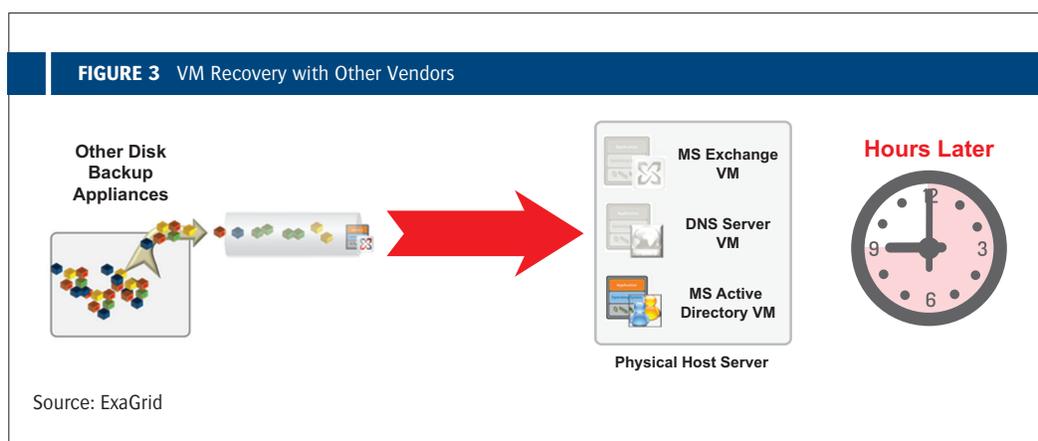
ExaGrid has reinvented support by offering local technical support with highly experienced level 2 technicians who are assigned to each customer, so the same technician works with the same customer the majority of the time. This approach assigns techs with software expertise that matches each customer's environment and backup application, and enables technicians to get to know the customers and their needs.

Ameris Bank

Ameris Bank has over \$3 billion in assets and over 60 banking locations. Ameris had been primarily backing up its data to tape, but managing and troubleshooting backups had become increasingly time consuming, and recoveries were error prone. The Bank decided to install a two-site ExaGrid system as part of its datacenter upgrade and relocation project.

"The main reason we chose the ExaGrid system was its tight integration with our existing backup application, Veeam Backup & Replication. We are ninety percent virtualized and after doing some research, we learned that the combination of ExaGrid and Veeam would deliver the fast restores, data deduplication, and backup performance we were looking for ... [It's] a really solid, rugged system."

Shan Venable
Systems Architect



ExaGrid offers the following options to provide the best and the lowest-cost technical support in the industry:

1. Yearly support as a percentage of what is paid for ExaGrid's appliances
2. All-inclusive fees, so there are no hidden or future costs, that include:
 - a. Assigned level 2 technicians
 - b. Health monitoring, including the option to have ExaGrid receive all alerts and alarms
 - c. All point and full version releases
- d. Next-business-day delivery of any failed component (disk drives, power supplies, etc.). Each appliance includes:
 - i. RAID6 with a spare, so the system continues to run even if there is a simultaneous multiple drive failure.
 - ii. Redundant power supplies, so the system continues to run if a power supply fails.

Summary

In the world of virtualization and big data, architecture matters. With the wrong disk-based backup architecture, an organization will not be able to keep up with its growing data and the need for fast restores to keep users productive.

ExaGrid eliminates all of the challenges that Gartner has identified in its report, "Best Practices for Repairing the Broken State of Backup."

ExaGrid's unique approach of a landing zone and scale-out architecture enables ExaGrid to make five commitments to its customers:

1. You will have the shortest backup window.
2. Your backup window will not increase as your data grows.
3. You will have the fastest restores, tape copies, and recoveries.
4. Your virtual machine instant recoveries will occur in minutes.
5. You will have the lowest cost solution up front and over time, with world-class support.

Source: ExaGrid

The Energy Authority

"I worked with our ExaGrid customer support engineer to install the system and we were able to get it up and running fairly quickly. It really is a 'set it and forget it' type of technology. I get a daily report with details on the state of each backup job and ExaGrid reaches out and notifies me if there's a problem with the system. I'm not manning or managing the device every day – it just runs. We also have a good relationship with our support engineer. He's proactive and knowledgeable and is a good resource for us.

"We've expanded the ExaGrid system at our primary site, and we're planning to expand it in our disaster recovery site within the next 30 days. It's incredibly simple to scale the system. Once the unit is racked up and we assign an IP address, ExaGrid support takes over and finishes the setup. It takes only a few minutes."

Scott Follick
IT Manager, Service Delivery and Support

About ExaGrid

Backup without compromise – solving the backup and restore challenge forever

ExaGrid Systems, Inc. has revolutionized how organizations back up and protect their data. Its cost-effective disk-based backup systems help IT organizations solve two of the most pressing issues they face today: how to protect and manage an ever-increasing amount of data and how to do so at less cost.

For more information:

www.exagrid.com

1.800.868.6985

44.1189.497.051



From the Gartner Files:

Best Practices for Repairing the Broken State of Backup

Storage managers and IT leaders face significant challenges when managing backup processes for their data center infrastructure. We offer best practices to help manage data growth, improve service levels and contain costs.

Key Challenges

- Growth of data, proliferation of new applications and expanding production usage of virtualization, combined with the users' need for continual data access, impact how storage managers can use traditional backup methods.
- A legacy backup approach does not scale well, as it cannot handle increased demands for protecting today's data sources and volumes to meet the backup window.
- The primary backup application switching triggers for data center personnel remain: (1) cost — including acquisition and maintenance; (2) perceived vendor capabilities; (3) management complexity; (4) completeness and scale of solution; and (5) customer support issues.

Recommendations

- Start with the end in mind by developing a comprehensive backup or recovery plan that includes service-level requirements, data growth and network architecture changes over the next three years.
- Deploy proven enhancements in available backup products — such as incremental forever backups, synthetic or virtual processing, compression or deduplication, file scan improvements, enhanced server virtualization support, and snapshot and replication integration.
- Implement a tiered recovery plan with multilayered SLAs that match recovery techniques to business requirements, and that will optimize your availability strategy. A "one size fits all" backup strategy is no longer feasible.

- Select a unified recovery management solution in which all, or at least most, backup techniques are managed from a centralized point.
- Perform backup recovery tests at least once a year on a data subset to ensure the backup strategy can meet the stated protection SLAs.

Strategic Planning Assumptions

Through 2015, disk-to-disk-to-tape (D2D2T) backup will remain the predominant strategy for 40% of large enterprises.

By the end of 2016, 40% of large enterprises, up from 20% at year-end 2012, will have eliminated tape for operational recovery.

Between 2012 and 2016, one-third of organizations will change backup vendors due to frustration over cost, complexity or capability.

By 2016, at least 20% of large enterprises, up from less than 7% in 2013, will abandon backup or recovery solutions using a traditional methodology, and adopt solutions that employ only snapshot and replication techniques.

Introduction

Backup and recovery are some of the oldest and most frequently performed data center operations. Yet they remain frustrating, as the need for effective recovery challenges storage and data center managers. Gartner's client inquiry call volume regarding backup remains high on a year-to-year basis. We see logjams in the data backup and recovery process, such as:

- Amount and proliferation of data
- More-stringent SLAs
- Antiquated backup models (that is, scale issues)

The Three Cs of Backup Concerns

Cost

While the upfront acquisition cost remains a concern, the annual maintenance fee associated with backup software is the issue for larger enterprises. As more applications have been deployed, and more machines (physical and virtual) have been installed, the costs of backup have risen. Often coupled with this is the desire to add protection to remote offices, and sometimes to protect desktops and laptops. All combined, these things can result in a larger backup bill, and much larger annual maintenance and service costs. In some situations, newer pricing models such as pricing based on capacity (in terabytes) has helped, while in other cases, they have become an increasing concern.

Capability

From end-user inquiries, conference polling and worldwide surveys, Gartner clients' top complaints about backup capability include not meeting the backup window, not being able to restore data fast enough to meet SLAs, point solutions not being able to fully protect all data types or locations, daunting administrative interfaces, and incomplete reporting. In particular, they have concerns that their vendors or products are not keeping pace with the level of innovation of competing solutions.

Complexity

Organizations also complain about the complexity of their backup solutions. The feeling is that too much time, expertise and effort are spent keeping the current recovery systems afloat. Organizations would like backup to be a process that requires far less supervision and administrative attention, and for higher backup and restore success rates to be easier to achieve.

The Fourth and Fifth Cs of Backup Concerns

Over the past three years, Gartner clients have mentioned additional backup frustrations, such as completeness and scale, and customer support.

Completeness and Scale

This is a variation on the capability issue above. Although a particular function may be available, organizations have concerns about how robust it is, how much additional effort is required to employ it, and especially, how effective the overall solution performs at broad scale. An example of this is

server virtualization support, which may have been recently improved and deemed workable initially, but at scale, has exhibited issues.

Customer Support

Many organizations are basing their vendor renewal considerations on the quality of support that they receive. Breakdowns in support systems can lead to lost confidence in the product or vendor. Organizations increasingly do not want to rely too much on the heroics of their in-house teams to deliver an effective backup practice.

Analysis

1. Fully Implement Current Backup Solution Capabilities

For a variety of reasons, some organizations have yet to embrace and deploy data protection approaches and techniques that are already available. The expanded use of disks for backup — while a marketing slogan for some time — is now the norm for the industry. Client- and target-side data deduplication solutions are offered by many providers. Most backup suppliers have delivered significant improvements in server virtualization and SharePoint recovery in their most recent product releases. Support for endpoint backup and cloud applications is becoming more robust.

New capabilities can take three to five years or more to gain widespread adoption, as most organizations are risk-averse. However, many companies delay implementation because they do not know their vendor options and capabilities.

Action Item:

- Before making plans to jettison your current backup product, ensure that your vendor has provided feature updates from the past three years to fully leverage the investment.

2. Implement Archiving and Improved Data Management Practices

The vast majority of the backup methodologies use a model whereby frequent full backups of all data are taken. While it is often standard procedure to configure the application for nightly incremental backups for six consecutive days and then once a week (usually on a weekend day) to take a full backup, many organizations opt for a full, nightly backup for email and critical databases, as that practice can minimize the amount of restore processing that needs to occur. While this

approach has worked well in the past, many now find they cannot contain backup activity in the available time (the backup window). While newer backup approaches can help address this, many backup implementations still rely on the “full plus” incremental concept.

Removing files, email and application data from primary storage can drastically reduce the amount of backup processing required during each full backup, known as reducing the “working store.” Organizations could perform a garbage collection process, using a storage resource management (SRM) tool. However, most use data identification tools in an archive solution or emerging stand-alone file analysis solutions to identify archive candidates, or they outright delete unneeded and duplicate data.

Implementing an archive solution that moves data to lower-cost storage devices can reduce the backup window. Archiving also can:

- Provide faster restore times for a complete recovery of all backed-up data (since there is less data to bring back)
- Reduce the cost of storing and potentially transporting the backup media off-site
- Decrease backup retention periods

Gartner typically recommends a 90-day backup retention. This results in lower exposure for e-discovery during a litigation activity, and can help contain the significant labor issue of scanning backup tapes for required legal materials.

Action Items:

- Implement an archiving solution as part of an overall information governance strategy to improve backup and restore times, and reduce primary and backup storage costs.
- Evaluate SRM and especially file analysis tools as a way to implement a “defensible deletion” policy in your organization.
- If possible, reduce backup retention to 90 days to reduce costs, accelerate overall backup or recovery processing, and reduce e-discovery exposures.

3. Evaluate Newer Backup Pricing Models

Backup products are traditionally priced on a per-server basis — with add-on costs for advanced features, and newer capabilities commanding a premium price. Over time, most vendors continue to collapse the number of charged items into their base products, or add features into an expanded, extended or enterprise offering. As a result, current deduplication charges should be expected to decline, perhaps as a result of competitive pressures.

Nearly every backup vendor has introduced capacity-based pricing. For organizations that deploy server virtualization (which was sometimes a backup option separately charged for), many application agents or advanced disk features, the capacity-based bundle can represent a more attractive overall initial cost. For organizations that have a fewer number of servers and a larger amount of data — especially when a single server is used as a proxy for many terabytes of network-attached storage (NAS) Network Data Management Protocol (NDMP), or when few advanced features have been implemented — capacity-based pricing may cost more than traditional server-based licensing.

For larger enterprises, the maintenance costs of the typical three-year backup software purchase and renewal agreements loom as a greater concern than the initial acquisition cost, as these represent a future spending commitment. This future-oriented concern is amplified by capacity-based models since the growth rate of data is faster than the growth rates of physical and virtual servers. Vendors can differ as to where the capacity is measured — on the “front end” for the data being backed up, or the “back end,” measuring the amount of data that the backup solution generates, which is typically after compression and deduplication.

Clients tell us they are concerned about capacity-based pricing, noting that some leading vendors are expensive when deployed at scale. In July 2013, Asigra announced a new pricing model that has a smaller base charge, but charges on successful recoveries performed.¹ Gartner expects that the industry will begin to pressure backup software vendors for lower costs at larger backup volumes to avoid the similar issue where backup appliances become viewed as cost-prohibitive when broadly deployed.

Action Items:

- If you are under a maintenance agreement, investigate whether you are entitled to free upgrades. Consider newer versions when the vendor bundles previously charged-for features.
- When evaluating new backup solutions or extending a maintenance agreement with your current backup vendor, look for new pricing and packaging plans, such as capacity-based licensing, product versions with a collapsed parts list that include additional features in the base product, and current and upcoming bundles that combine features at a lower overall cost.
- When negotiating with vendors, first, understand which features are additionally charged for and what capabilities are included at no additional expense in all products that are on your shortlist to ensure an accurate TCO and for use in pricing negotiations.

4. Fully Implement Data Reduction

The value of data reduction technologies, such as deduplication, cannot be overstated. Deduplication improves the economics of disk-based backup and recovery approaches by reducing data — resulting in significantly lower disk requirements and cost, and providing more-efficient and faster replication of backup data to off-site locations. Gartner believes that data reduction, such as compression and deduplication, is a “must have” capability for backup solutions.

The benefits of deduplication are in resource savings. Potential savings can occur on many levels. The primary benefit is in substantially decreasing the amount of disk space required to store a given amount of data. Gartner clients typically report deduplication ratios of 7-to-1 to 20-to-1. Actual ratios vary depending on the amount of data redundancy, the type of data (encrypted and previously compressed files often do not further compress or deduplicate with backup deduplication), the data change rate, and the backup methodology (for example, full, full plus incremental, full plus differential or incremental forever).

The more often that full backups are conducted, the higher the deduplication ratio is. Depending on the deduplication implementation, there can

be bandwidth savings in the amount of data transferred over the network. Deduplication can decrease power and cooling needs and the physical size of storage devices, as this practice uses less physical capacity and lowers acquisition costs. The benefits of data reduction increase as more data is processed and stored, and as a larger history is available for comparison and reduction.

Action Items:

- When evaluating backup software or disk-based hardware solutions, consider data reduction (such as compression and data deduplication) a must-have feature and an essential part of assessment criteria. Many backup vendors have released new or expanded deduplication features. Understand the latest capabilities so as not to be incorrectly swayed by a vendor’s positioning of the competitive capabilities of other solutions.
- Re-evaluate the applicability of deduplication to more workloads. The use of solid-state drives (SSDs) for portions of the backup infrastructure (especially for the deduplication index) — or more-refined algorithms that split processing requirements across backup clients, the media or master servers — have made the use of deduplication more appropriate for more data types.
- Deduplication, in particular, offers the potential for many cost savings (decreased disk, bandwidth, power and cooling requirements). However, ensure that any premium paid for the capability does not offset the economic savings. Even if deduplication costs more than a nondeduplicated solution, consider the performance and operational benefits.

5. Implement Unified Recovery Management (Snapshot and Replication Support)

The environment that needs to be protected is expanding. There are not only newer types of applications and an increasing amount of virtualization deployed in the data center, but also new workloads such as test and development, remote office and branch office (ROBO), and endpoints (laptops, tablets and smartphones), generating data that needs to be protected. At the same time, hypervisor solutions, and snapshot and replication solutions that are based on servers, storage arrays and networks,

are becoming pervasive at attractive cost points. Enterprise backup vendors are responding with leading backup solutions, adding capability to protect more data center and non-data-center workloads at the file, application, virtual machine and volume levels.

This has led to the notion of a single administrative console, catalog and reporting engine for all data capture mechanisms. This can enable application and virtualization administrators to perform backups and restores by utilizing the primary backup application. Backup vendors are integrating with storage systems to catalog snapshots (such as, in order of the breadth of support, CommVault, HP, IBM, EMC, Symantec, Asigra and Veeam), and some products offer integration with replication engines.

Traditional backup products will eventually transform into recovery management solutions that may not own all of the data capture and data transfer techniques. In addition to traditional backup and recovery (application-, file- and image-based, and so on), stronger support for solutions for server-based replication, storage-array-based replication, or intelligent-switch-based or network-based replication will become more important.

The notion of copy data management — which reduces the proliferation of secondary copies of data for backup, disaster recovery, testing and reporting — is becoming increasingly important to contain costs and to improve infrastructure agility.

There will also be a “manager of managers,” a common and established concept in the networking and system management domains. A hierarchy of federated management tools feed into one another, becoming a unified recovery manager. This allows for simplified implementation of several tiers and service levels, offering centralized monitoring, reporting and control.

Action Item:

- Before making additional investments in backup software, push incumbent and prospective recovery vendors for current and committed road maps for their manager-of-managers support — especially snapshot and replication management and integration, copy data management, and single-console capabilities.

6. Implement Tiered Recovery

Directionally, disk usage, data deduplication, replication for electronically vaulting off-site copies of data, and snapshots for taking more-frequent copies of data are all on the rise. Yet the same tools, technologies and backup techniques from decades ago are also typically still implemented. This expanded menu of options, techniques, cost structures and service levels has changed the way that organizations deliver backup services.

In the past, backup was very much a “one size fits all” endeavor. Often, the only major backup decisions being made were whether something should be backed up and, if so, how long it should be retained. Today, new techniques and technologies have led to an expanded menu of choices, and one current or emerging recovery approach does not always win out over another. Administrators will have more flexibility, including differentiated levels of cost and service, in providing recovery solutions. Just as the concept of tiered storage provides a variety of cost and performance levels for storage capacity, tiered recovery provides differentiated levels of backup and recovery services. Unlike tiered storage, the tiered recovery model may be additive, with an organization using multiple techniques together to achieve the needed overall level of data availability and recovery characteristics, and to ensure that business risk and continuity requirements are met.

To implement tiered recovery, organizations should conduct a business impact assessment (BIA) to categorize the criticality of the IT services. Any recovery architecture must include an understanding of the IT services supporting business processes and associated service levels. Service levels affect the capabilities, cost, architecture and complexity of the backup solution. Gartner recommends performing a BIA for backup data to contain cost and to deliver the most appropriate service levels. Most organizations specify three to five tiers of criticality, with the highest tier having the most stringent service levels.

Action Items:

- Implement tiered recovery to optimize the balance between cost and recovery service levels, recognizing that the exploitation of storage device snapshot and recovery techniques are likely to be leveraged to reduce backup windows and improve restore times.
- Conduct a BIA, and review it annually to determine the criticality of your business systems and their data. Implement tiered recovery by using the BIA results, and devise three to five tiers. Associate recovery service levels to each tier, including recovery time objective (RTO), recovery point objective (RPO), retention and off-site copies.

7. Perform Regular Restore Testing

Backups may be unrecoverable for many reasons. Some of the more common issues are server configuration and application deployment updates, user or operator error in the backup process, and hardware and software failures. In most organizations, backups are initially set up and then automatically run. Backup verification tends to be only a review of the backup logs, with a quick scan for failures and error messages.

This process may be acceptable for determining whether data was successfully written to the backup media. However, it doesn't provide information about whether the data is recoverable, and does not validate that the data is logically consistent and usable by the application, nor that the right information was backed up. Some businesses have instrumented backup reporting tools to better understand how backups trend over time, and to get more visibility into backup successes and failures.

Actual recovery is the only way that a data center can be certain that data is fully recoverable. Backup or restore testing has become a dying practice in most data centers, and as a result, organizations could be far less resilient than they believe they are.

Action Item:

- Perform data recovery testing at least once a year on a subset of data to ensure that the backup strategy can effectively meet the stated protection SLAs. More-frequent testing of the most mission-critical data may be warranted.

8. Ensure That the Cloud Backup Has a Local Copy

Public cloud options are becoming increasingly considered for server workloads, especially for remote-office and departmental computing environments. The best practices discussed above apply equally as well in a cloud backup deployment. While most organizations cite concerns over security as their top cloud issue, the greater issue is often latency, as data encryption and key management are well-established methods for protecting off-site backup data. This means that an on-premises, local copy of the data, or at least the most recent backup of the most critical data, is best practice. Thus, a disk-to-disk-to-cloud (D2D2C) model is emerging.

Action Item:

- Ensure that all servers with a restore data payload of 50GB or more, and an RTO of one day or less, have local disk copies of the backup data for first-level operational restore requirements. This would only be for protection against logical errors, such as accidental deletion and data corruption, and a limited number of physical threats, and the cloud copy of the data would be used for larger disaster recovery remediation.

Evidence

This research is based on more than 1,250 client interactions relating to backup and storage with Gartner analysts on an annual basis. Gartner also conducts conference polls and surveys worldwide that offer similar evidence.

¹ Asigra's press release for its Recovery License Model, 10 July 2013.

Gartner Research Note G00252768, Dave Russell, 15 August 2013