Invest in Storage Professional Services, Not More Hardware

Gartner RAS Core Research Note G00165133, Adam W. Couture, 27 February 2009

Storage assessment services are now commonly promoted by storage OEMs and independent software vendors (ISVs) as well as value-added resellers (VARs), system integrators and independent storage consultants. A typical $50,000 to $100,000 investment in a storage professional services engagement can free up wasted storage capacity and improve storage efficiency, avoiding millions of dollars in hardware acquisition expenses. Additionally, storage service provider tools, methodologies and managed services can make IT staffs more productive and better able to cope with head count restrictions during a down economy.

Key Findings

- Based on end-user inquiry, Gartner has observed that most storage infrastructure is managed inefficiently. By improving storage utilization, IT departments can delay incremental storage acquisitions.
- Storage professional services are available through storage system vendors, ISVs, resellers, system integrators and VARs as well as independent service companies.

Recommendations

- Unless your storage utilization rate is above 70%, immediately begin planning how you can stave off hardware purchases through better storage management.
- Honestly assess your internal capabilities to determine if you require external assistance to extend the useful capacity of your existing infrastructure. For example, based on Gartner inquiry, many IT organizations are ill-equipped to address the nuances of eDiscovery and compliance.
- Although capacity reclamation services will free up storage space, growth will eventually require new storage purchases. Make certain your storage assessment includes a road map for cost-effective acquisition strategies as well as short-term savings.
- If you are asking the service provider to benchmark savings from staving off hardware acquisitions because of recovered capacity, make certain that Gartner hardware benchmark reports are used as the basis of the calculations.

Gartner

ANALYSIS

You think you need a new car. The odometer on your aging sedan reads 150,000 miles, and you're expecting to pay at least $30,000 to replace it. But a friend tells you he has a mechanic who, for $500, will make improvements to your car and teach you some driving techniques guaranteed to wring another 50,000 miles out of your vehicle. Would you consider it? That's just the sort of value proposition storage professional services providers bring to the table. In many instances, they hold out the promise of freeing up capacity and extending the useful life of storage arrays for the price of a professional services engagement.

A professional services organization can be a critical adjunct to your IT organization, bringing fresh tools, methodologies, efficiencies and insight into your infrastructure and staff. Assessments that a typical IT department may undertake once every five years are daily events for storage service providers. What's more, they have invested heavily in the assessment tools, training and methodologies that would be uneconomical for enterprises that would use them only rarely.

Most storage vendors and ISVs can deliver a variety of professional services. Additionally, many larger resellers and system integrators, such as Forsythe Solutions, Datalink and Champion Computer, have developed consulting and even management practices, supplementing product sales and implementation services. A handful of pure-play consultancies, such as GlassHouse Technologies and Contoural, eschew product sales and come to market with a vendor-neutral approach to professional services.

Figure 1 shows the Gartner decision framework for assessing the impact of cost optimization projects related to storage assessments. Infrastructure managers should use it to plan for configuration-related cost-saving initiatives in their organizations. For each of the six elements of the decision framework shown in the figure, we provide our expectations of the impact of the storage assessment cost reduction initiatives, including the low, medium or high benefit.

Figure 1. Decision Framework for Storage Assessment Cost Optimization

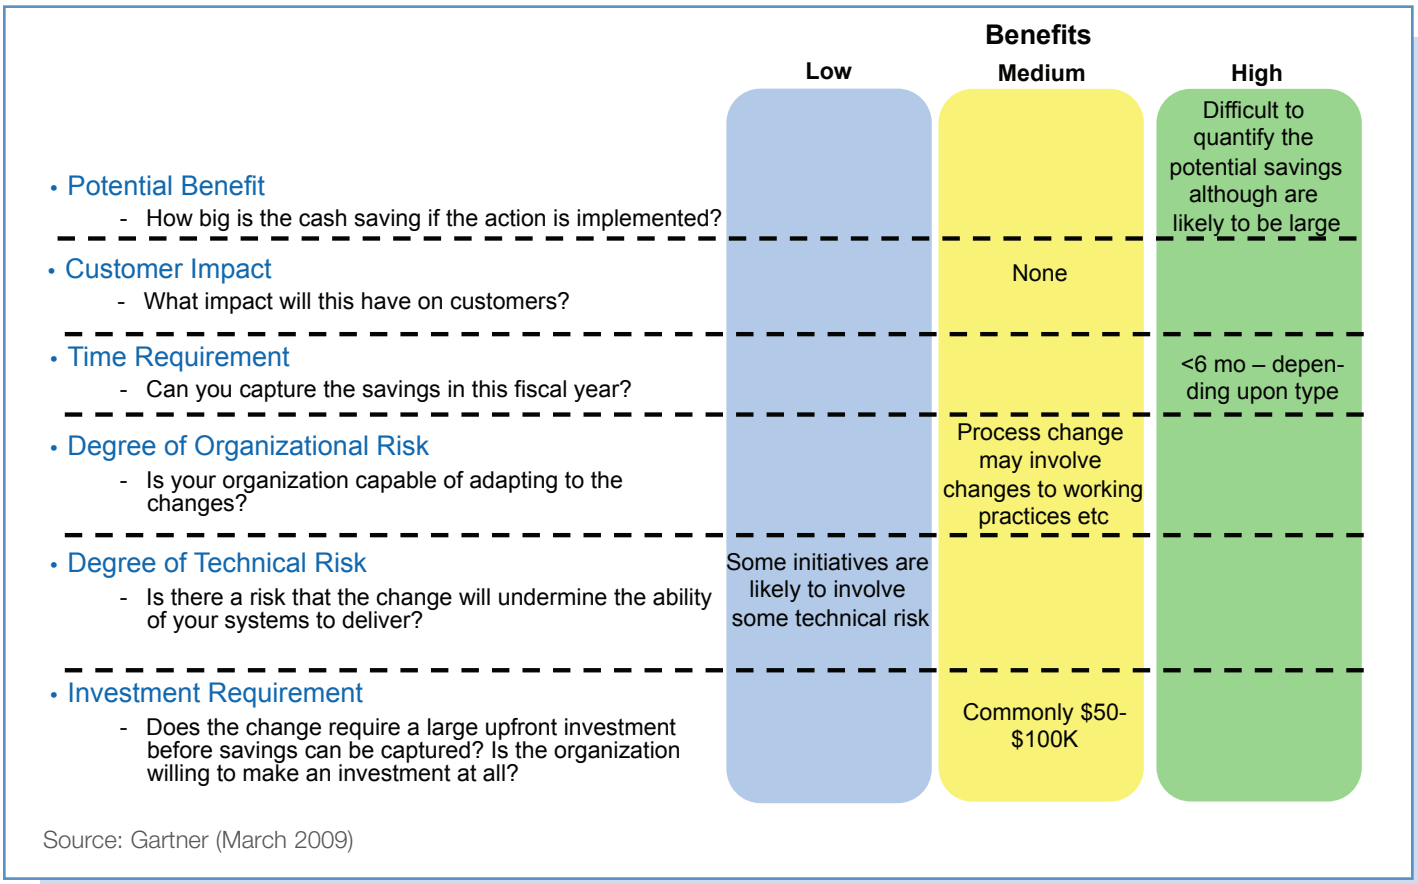

Source: Gartner (March 2009)

1.0 Data-Reducing Professional Services

Although storage service providers offer dozens of professional and managed services each, for the purposes of this report, Gartner focuses on those services designed to improve storage infrastructure efficiency, free up usable capacity and increase utilization rates, thus staving off immediate needs for incremental major hardware acquisitions. Professional and managed services that fall into this category include the following:

- Backup assessments
- Unnecessary data identification and elimination
- "Right-sized" storage allocation
- E-mail and file system archiving
- Litigation hold process improvement
- Deduplication assessments and implementations
- Data classification
- Process efficiency and automation
- Managed services

1.1 Backup Assessments

Professional services organizations are usually called in to perform backup assessments when customers are struggling with failed backups and/or restores, shrinking backup windows, and aging backup infrastructure. But in addition to resolving these problems, backup assessments frequently identify unnecessary backups, such as backing up archived data or multiple backups of the same data. During a storage assessment engagement, one service provider discovered that its client was backing up the same database eight times. Another discovered a 40-1 ratio of replicated data to production data. Obviously, the client reclaimed an enormous amount of wasted disk space. Moreover, eliminating unnecessary backups not only frees up backup media, but it can also reduce the cost of off-site tape vaulting while shrinking backup window requirements. As a consequence, IT staffs become more productive.

1.2 Unnecessary Data Identification and Elimination

Unnecessary data is at the heart of storage bloat. For example, e-mails and attachments that will never see the light of day again and have no retention requirements for compliance or eDiscovery can comprise 30% or more of mail server capacity. Similarly, stale databases, unneeded image and graphic files, and even unused applications are the "bosom buddies" of storage hardware salespeople. Storage assessments can identify data that serves no purpose and can be safely removed from the data store. Service providers can also teach employees how to identify and reduce the amount of unstructured data they preserve.

1.3 Right-Sized Storage Allocation

Database administrators need one terabyte of storage for their application. To ensure against running out of space, they request two terabytes from the storage administrator. To ensure against running out of space, the storage administrator allocates four terabytes. That means three terabytes of storage have been allocated to a database that doesn't require it and can't be used by other applications that really do.

Sound familiar? Unfortunately, it's all too common a practice. Cisco went so far as providing prizes to database and storage administrators who returned unused allocated storage. A storage professional services engagement can discover instances of storage overallocation and recommend right-sized storage allocation for each application or virtualized server instance. Unfortunately, service providers can't enforce proper storage allocation. That responsibility will remain with internal IT departments.

Problems associated with overallocation can also be addressed through thin provisioning, which allocates storage on a just-in-time basis, effectively allowing IT departments to subscribe far more storage than is physically on the array. However, unless companies already have arrays accommodating thin provisioning, it is not something a service provider could accomplish without new hardware expenditure.

1.4 E-Mail and File System Archiving Managed Services

After 90 days, e-mail messages are rarely ever accessed. Just the same, they chew up disk capacity on primary storage arrays as readily as production data. Archiving e-mails and their attachments means that they can be removed from the e-mail server, freeing up space for contemporary e-mail messages. Moreover, because the data has been archived, it no longer needs to be backed up, thus freeing up backup media capacity and shrinking backup windows. Granted, an on-site e-mail archiving implementation may require users to purchase a low-cost storage tier for archived data. However, a managed e-mail archiving service can free up capacity without requiring investing another dime in storage infrastructure.

1.5 Litigation Hold Process Development and Improvement

When companies implement litigation holds (litigation is increasing in this economy), they often simply stop rotating backup tapes and retain all backups for the duration of the litigation, if not forever. This is expensive in terms of media costs, and extremely expensive if they need to search through tapes. Companies can save significant money if they develop more-refined, intelligent hold processes, saving only relevant data. Also important is a process for releasing holds, which can free up significant space. While professional services providers have varying degrees of expertise in litigation-related processes and practices, they are often far better informed than many internal IT departments in identifying which data is truly case-related and must be preserved and which can be safely scrubbed.

1.6 Deduplication Assessments and Implementations

In recent years, the term "single instance store" (SIS) has been used to describe some data reduction approaches. SIS means that only one copy of a file will be retained, even if there are dozens, or thousands, of the same file in the data store. When using SIS, only one copy of the actual duplicate file would be stored; the other occurrences would point to the single copy, thereby reducing capacity requirements.

Data deduplication takes the data reduction concept of SIS further. Data deduplication uses data identification and comparison algorithms to dramatically reduce the space requirements of individual objects. It does this by storing only unique "chunks" of data, thereby eliminating redundancy. Although SIS detects and eliminates duplicates on a file basis, block-level deduplication examines redundancy at a subfile level, comparing previously stored data.

When some people think of the term "compression," the concept of making a single file or backup stream smaller comes to mind. What makes deduplication transformational is that it is not applied to a small amount of data; rather, it is operating at a global level, reducing data size across a much-broader environment, and the knowledge of previous data chunks is retained and used as new data is examined in the future.

Data deduplication has profound implications on customer storage infrastructure acquisitions. Vendor claims range from capacity savings of 20-to-1 to 400-to-1. Gartner clients with deduplication technology in production report data reductions ranging from 18-to-1 to almost 300-to-1 for backup data. As a consequence of this freed space, customers can potentially stave off some storage infrastructure acquisitions for years.

Initially, service providers may be able to provide deduplication services on a one-time basis, without the expense of investing in deduplication technologies, because the service provider will use its own or licensed intellectual property for the service. In the future, customers will have to invest in deduplication technologies to prevent duplicate files resurfacing in data stores.

1.7 Data Classification and Information Life Cycle Management

Information life cycle management (ILM) is essentially an enhanced version of hierarchical storage management. It attempts to classify various categories of data according to their importance and business value and establish policies to migrate the data to different storage media accordingly. Needless to say, companies promoting ILM tend to define it around their own hardware, software and service capabilities, which generates considerable customer confusion about what ILM is and isn't. Moreover, ILM isn't usually a stand-alone storage initiative but a foundation for crafting storage architectures and solutions. Despite the differences in how individual companies define ILM, one profound commonality exists among them: Services such as data classification policy development and architecture tiering are key to the success of any ILM implementation.

Data that is properly classified against a proper tier of storage has the potential to produce enormous storage savings. One major storage OEM estimates that its customers now have 60% or more of their data on their most expensive storage tier and further estimates that half of that data could effectively be moved to a less expensive tier.

But what if you have only one tier of storage? How can data classification help there, without making new storage investments? A service provider can help define multiple storage tiers even from a single array with identical hard drives. For example, different applications could be tiered according to backup schedules and retention policies. Of course, achieving meaningful long-term cost savings may ultimately require incremental storage infrastructure purchases. But the road to ILM can begin with existing hardware.

1.8 Process Efficiency and Automation

When internal IT departments are tasked to "do more with less," CFOs aren't just looking to curtail hardware and software spending. Frequently, IT hiring requisitions get frozen or eliminated. Worse, a reduction in force will require tough decisions regarding which IT staff members are retained and which are let go. Storage professional services organizations have invested heavily in developing highly efficient processes, methodologies and technologies to make their employees as efficient and productive as possible, and, in turn, increasing profitability per employee. These can range from simple scripts to Information Technology Infrastructure Library-based processes and methodologies to proprietary process automation software. During a professional services engagement, storage service providers not only employ these processes and tools themselves, but they also train IT staffs in how to use them to become more efficient and productive — in essence, doing more with less.

1.9 Managed Services

Managed services can also alleviate IT staffing pressures. Traditionally, staff augmentation services, such as on-site residencies, have been the stalwart of managed services. But more and more often, storage service providers are delivering managed services remotely and at costs far below what an on-site resident would cost. Some of these managed services include e-mail archiving, remote managed backup and remote storage administration. Services such as remote backup and archiving can shift infrastructure requirements from the customer to the service provider, freeing up capacity or even eliminating some infrastructure altogether. Additionally, many CFOs prefer the ongoing operational expenses of a service provider to investing in their own capital assets that must be managed and depreciated.

Figure 2. Energy Cost Savings by Drive Type

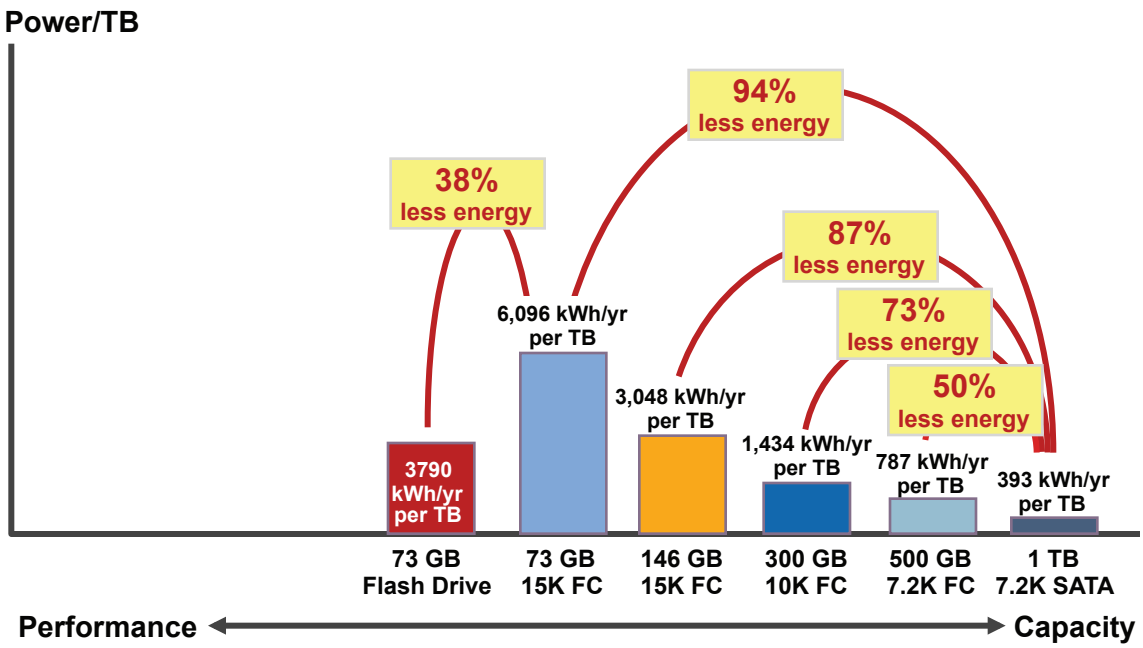

Source: EMC

2.0 Other Cost-Reducing Services

While the focus of this report is on storage professional services associated with extending the useful capacity of existing storage infrastructure without making investments in incremental hardware and software, eventually storage growth will exceed the incremental capacity freed up through these savings, and new investments will be required. The good news is that the same storage professional services deployed at short-term capacity reclamation can also deliver a strategy for cost-effective storage growth with significant longer-term savings.

After classifying data, they can architect multiple tiers of storage so that less-demanding applications aren't devouring the most expensive storage infrastructure. They can architect multiple disk tiers, reconciling expensive Fibre Channel disks against low-cost Serial Advanced Technology Attachment drives.

As shown in Figure 2, this not only impacts the initial purchase price of spinning media, but it can also have a long-term benefit because of reduced power consumption requirements of different drives.

3.0 Taking the Next Step

No matter which storage professional services you choose to engage, chances are that the cost of the engagement will need to be weighed against quantified cost avoidance. If you are asking the service provider to benchmark savings from staving off hardware acquisitions because of recovered capacity, make certain that Gartner hardware benchmarks are used as the basis of the calculations.

Services that extend beyond storage can also positively impact costs. Green IT assessments strive to reduce carbon footprints but also have a cost-saving benefit from reduced power consumption. Similarly, data center consolidation projects gain efficiencies while reducing real estate, personnel and power costs. While these projects can be exponentially more expensive than a short-term storage assessment, the payoff can be exponentially higher.

Finally, a storage assessment, like any professional services engagement, can be fraught with pitfalls if the objectives aren't properly defined, the required savings aren't specified, or the scope of the project isn't bounded.