Market Report

Augmenting Data Backup and Recovery with System-level Protection

By Lauren Whitehouse

March, 2011

Contents

Introduction

File-level recovery on its own, while an important mainstay in data protection best practices, is an incomplete strategy when it comes to meeting today's stringent recovery time objectives (RTOs) for complete system recovery. Whether it's a physical or virtual machine, the time necessary to "rebuild" the server stack and re-establish configuration settings can delay actual data recovery.

System-level recovery can enhance or even eliminate manual, operator-driven recovery and deliver significant benefits to IT organizations. This paper will investigate what's at risk, where system-level recovery fits relative to current data protection approaches, the impact system-level recovery can have on IT's ability to meet RTOs and RPOs, and the potential of system-level recovery to reduce costs.

The Pressure is On

Digital information is fundamental to conducting business. Whether it's e-mail communications, supply chain management, sales order processing, content management, customer relationship management systems, or financial applications, data is the backbone of business. The more reliant a company is on digital data, the lower its tolerance for any interruption in application or data availability.

ESG defines a simple three-tier data model with the following definitions:

- **Tier-1:** Mission-critical data and applications with the highest requirements for both availability and performance.
- **Tier-2:** Data and applications requiring good performance and reliability, but not at the level of mission-critical data. This can include a mix of less-critical application data and older mission-critical data.
- **Tier-3:** Typically archived data, so performance is less essential, but the data must be retrievable.

For each class of application/data, risks must be assessed. Once the implications of downtime and data loss are understood, IT organizations can use these factors to determine application availability requirements and the protection mechanisms that should be applied to ensure compliance.

As shown in Figure 1, ESG research found that midmarket organizations (companies with 100 to 999 employees) cannot endure disruption for too long. More than half (52%) of midmarket companies can tolerate only one hour or less of downtime for their most mission-critical data (tier-1); nearly half (49%) can tolerate three hours or less of downtime for the next tier of mission-supporting data (tier-2); and 43% can tolerate nine hours or less for all other data (tier-3).[1] That's why considerable investments are made to protect all tiers of data: making backup copies helps to minimize a company's risk should data be lost or corrupted.

[1] Source: ESG Research Report, *2010 Data Protection Trends*, April 2010.

Figure 1. Downtime Tolerance

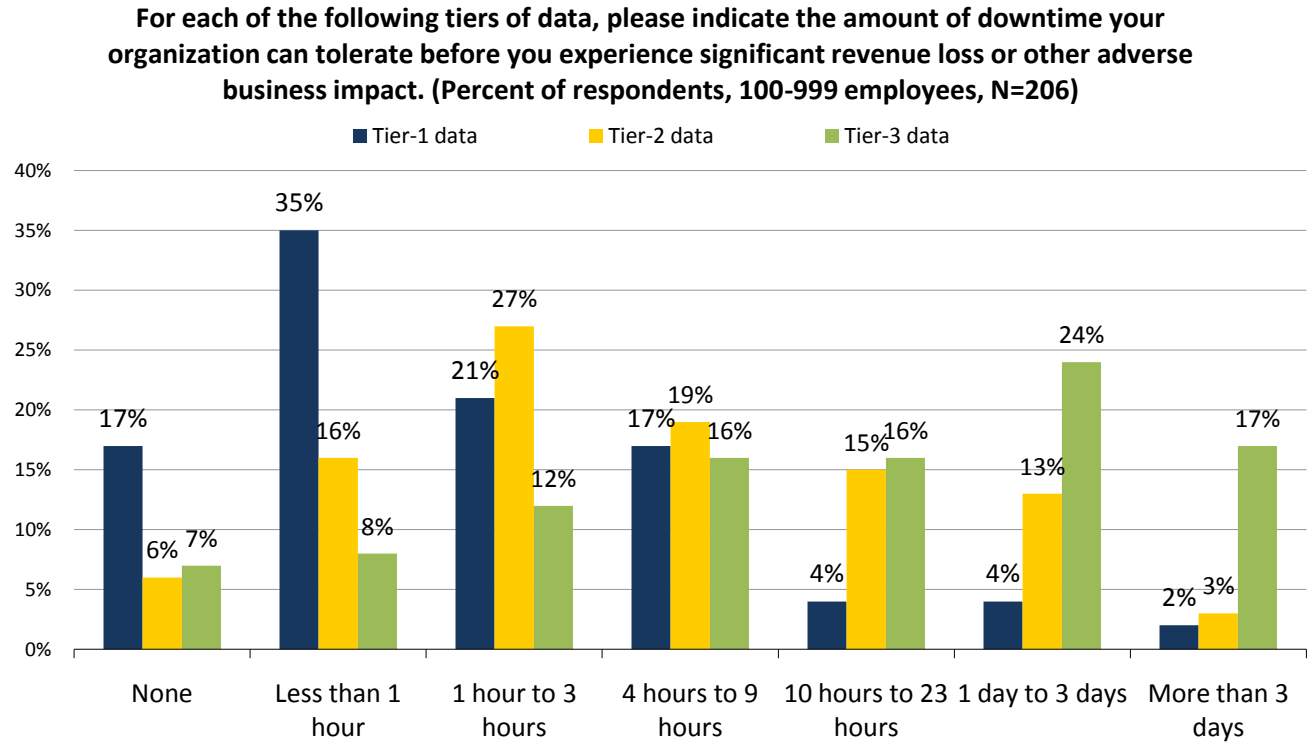

Source: Enterprise Strategy Group, 2010.

Managing risk is only one part of the equation. Preventing downtime has to be done within a budget—and IT budgets have been increasingly scrutinized over the past few years. While budgets have greatly improved since the economic downturn, cost reduction initiatives top the list of IT spending priorities. In a recent survey of IT professionals regarding spending this year, ESG found that 39% of survey respondents cite that their organization is in major or moderate cost reduction mode while another 49% are in cost containment mode (see Figure 2).[2] Furthermore, reduction in operational expenses (labor, utilities, materials, and other ongoing costs associated with operating an application, system, or business process) ranks first with 49% of respondents citing it as the most important justification for IT investments, putting considerable pressure on IT staff to do more with less.

[2] Source: ESG Research Report, *2011 IT Spending Intentions Survey*, January 2011.

Figure 2. IT Organizations Pursuing Cost Reduction Initiatives

Which of the following statements best characterizes your organization's change in IT spending from 2010 to 2011? (Percent of respondents, 100-999 employees, N=218)

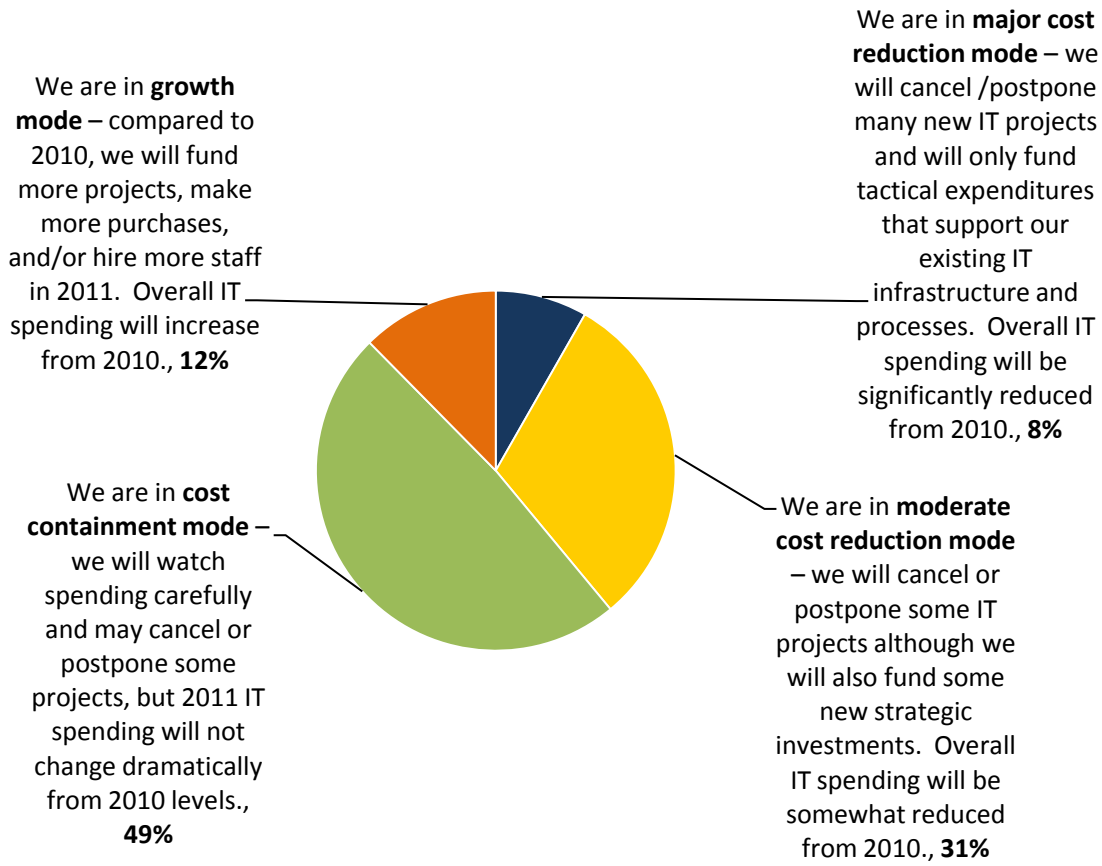

Source: Enterprise Strategy Group, 2011.

Doing more with less seems like an impossible task, especially when there is more—not less—data to manage. Midmarket organizations report they are most challenged with keeping pace with the amount of data they need to protect (see Figure 3).[3] Over 40% of the midmarket organizations ESG surveyed cite a more than 20% growth rate in data volume per year. Not only does this pose the obvious repercussions of requiring more storage capacity, bandwidth, and staff to deal with growing data volumes, but the rate of growth is impacting backup and recovery times. Since time is often a fixed variable, IT organizations are finding it increasingly more difficult to complete backup processes within the designated backup "windows" and meet RTOs.

[3] Source: ESG Research Report, *2010 Data Protection Trends*, April 2010.

Figure 3. Top Ten Data Protection Challenges

Which would you characterize as the primary challenge for your organization?
(Percent of respondents, 100-999 employees, N=206, top 10 responses)

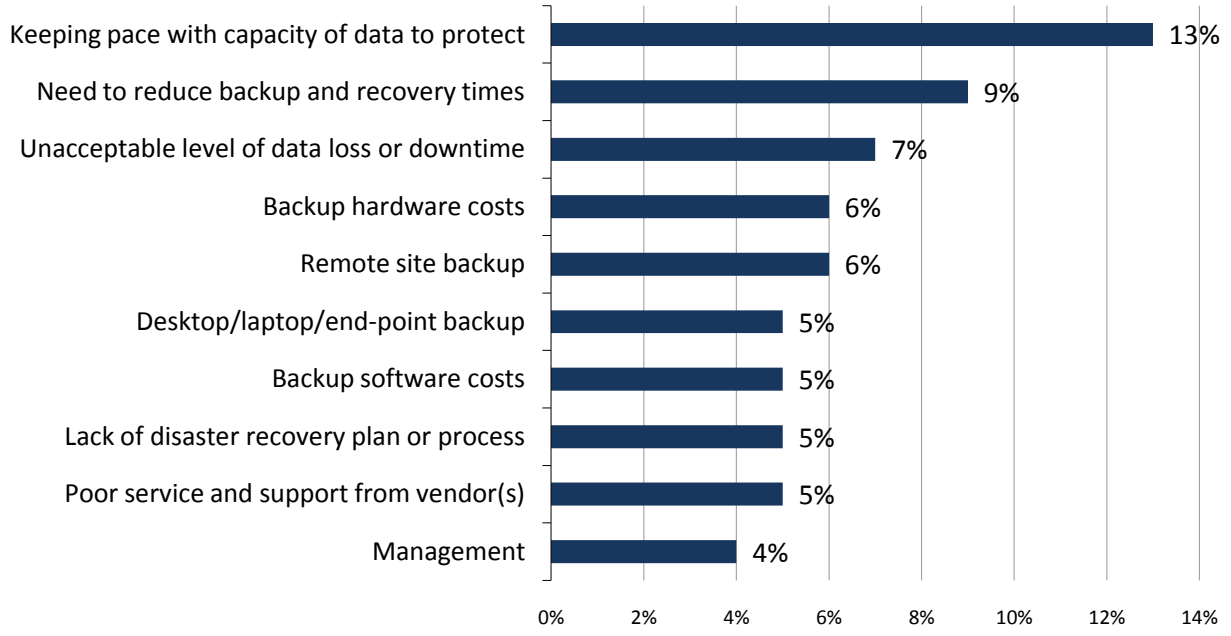

Source: Enterprise Strategy Group, 2010.

As if these issues were not enough, there's another pressure point wreaking havoc on data protection strategies: the deployment of server virtualization. Server virtualization has taken root in many organizations: among ESG's survey respondents, 55% percent are currently using server virtualization and 34% have not yet deployed the technology but plan to do so (see Figure 4).[4] Adoption is having a big impact on the IT environment, touching everything from IT infrastructure to IT management and organizational issues.

Figure 4. Virtual Server Adoption

Is your organization currently utilizing a server virtualization solution?
(Percent of respondents, 100-999 employees, N=131)

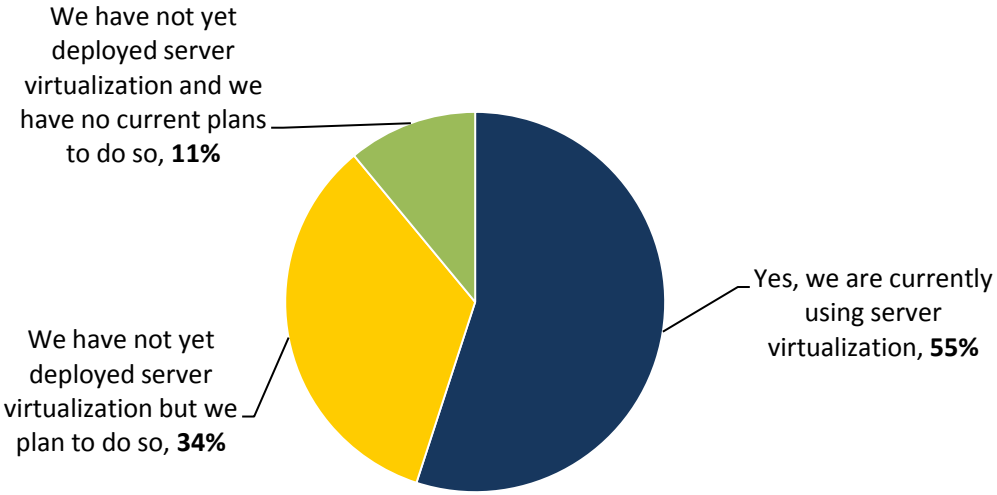

Source: Enterprise Strategy Group, 2010.

[4] Source: ESG Research Report, *The Impact of Server Virtualization on Data Protection*, September 2010.

Backup and recovery is especially affected since virtualization changes backup and recovery requirements. Why? The consolidation ratio of virtual machines running on a single physical server introduces risk. If a physical server fails, it will affect all of the virtual machines and application interdependencies running on that particular piece of hardware. There are also implications for multiple virtual machines sharing physical resources as well. A physical server environment is characterized by low resource utilization, but plenty of available bandwidth for backup. The new paradigm, on the other hand, flips that model, with higher resource utilization and less bandwidth for backups. Simultaneous resource-intensive processes (such as file- or application-level backup processes running in more than one virtual machine) occurring on one physical host could cause resource contention—potentially impacting the pool of application workloads sharing common physical resources and causing performance issues.

These pressures are causing IT organizations to re-examine aspects of backup and recovery to better optimize processes. For many, standard data-only backup processes are being revised to include system-level protection.

Mitigating Risk

Whether caused by a true "disaster" (fire, flood, or a natural disaster like an earthquake or hurricane) at the primary production site or something as mundane as a server failure, disk drive failure, software error, data corruption, computer virus, or "pilot" error, recovering a failed system using manual, operator-driven methods takes a significant amount of time. Given the aforementioned intolerance for downtime for any tier of application in the environment, it's imperative to have rapid, streamlined recovery processes in place—especially since suffering the loss of a critical business application could have a negative impact on business, such as revenue and employee productivity losses, lower customer satisfaction, damage to a company's reputation, or loss of a customer, partner, or supplier.

Loss also impacts the IT organization, since IT staff has to race against the clock to get downed systems operational again. This may entail acquiring replacement hardware, rebuilding the system, recovering the data, and enabling application access to users—which could take multiple hours to days. This process could be slowed significantly if the recovery process is not documented, staff is not trained, or required components are off site and have to be shipped to the recovery location.

While most organizations have daily data protection plans in place, fewer focus their efforts on system-level backup and recovery. So, what's the difference between data- and system-level protection?

Protecting Data

Data protection is defined as any process or technology that makes a secondary copy of a given *data* set so that business operations may continue in the event of the loss of the primary copy of that *data*. Typical best practices involve once per day data copies and the process of collecting data can be time-consuming, especially when data stores are large. Similarly, recovering data can take hours or longer depending on the type of media the data copy is stored on and where it resides.

If the whole production system is lost, then manual, operator-driven system recovery involves reinstalling the operating system and applications, re-establishing configuration settings, and then recovering data. This also includes applying operating system and application service packs and patches as well as updating drivers. If the operator makes a mistake, he or she has to start over from scratch. For tier-1 data RTOs of one hour or less and tier-2 data RTOs of three hours or less, significant risk can be introduced without a better level of protection than daily *data* backup copies.

Protecting Systems

System-level protection is defined as any process or technology that captures a complete image of the machine, including the operating system and application software, system state, and data. It's designed to allow the system to be rapidly recovered from "bare metal." The process of capturing a system image is greatly accelerated versus data backup where a file-level inspection must be performed—potentially reducing backup time requirements.

As more organizations standardize on running business applications in a virtualized environment, offering protection against virtual machine failure, application failure, resource failures, or wide area disasters can be even more important. The "all your eggs in one basket" vulnerability in virtualized environments, where multiple virtual machines share physical resources, has implications for downtime being even more significant. A virtual machine encapsulates all of the files an individual machine may contain, so considerations should be given to safeguarding the virtualized environment against downtime. System-level protection solutions capture and transfer active virtual machine disk files as a whole, enabling rapid local or remote recovery.

The System-level Recovery Difference

So what is the right approach? Augmenting *data* backup and recovery with *system-level* backup and recovery delivers more comprehensive protection of the environment. By "doubling up" on protection at the *data* and *system* levels, IT removes vulnerability, mitigating the risk of downtime and data loss. Now, instead of a one-size-fits-all-approach to recovery, data can be recovered with a data protection solution and systems can be recovered with one for systems protection.

The benefits of adding system-level protection can be sweeping in addressing:

- **Aggressive backup time and recovery service level agreements (SLAs).** By capturing system images non-disruptively and continuously, system-level recovery eliminates backup windows. System-level recovery addresses the rapid recovery needs of organizations by restoring systems to their pre-failure state and includes the ability to perform recovery to dissimilar hardware. Since "improving backup and recovery times" is the second-highest-ranked data protection challenge, this will be meaningful to every IT administrator.
- **Cost reduction initiatives—especially OPEX**. At 29%, operational support staff costs make up the largest percentage of the data protection budget.[5] Automation and process improvements realized by implementing system-level recovery that can drive down requirements for operator intervention not only keep staff costs down, but also make it easier to justify the investment in system-level backup and recovery.
- **New requirements for virtual environments.** By enabling capture of whole virtual machines (rather than just the data associated with them) and providing conversion capabilities from physical to virtual (P2V), virtual to physical (V2P), and even virtual-to-virtual (V2V), a physical or virtual machine can be restored rapidly and easily.

By augmenting data protection with system protection, the top challenges IT organizations are faced with today get tackled. Risk is managed with a dual-pronged strategy intended to close the gaps of traditional approaches. IT agility is maintained in spite of time constraints and cost control is addressed by focusing on reducing operator intervention in backup and recovery processes.

[5] Source: ESG Research Report, *2010 Data Protection Trends*, April 2010.

The Bigger Truth

While natural disasters and the havoc they can cause for business continuity grab the headlines, it's often the more mundane issues inside the data center, such as the loss of a server, that tend to occur with greater frequency. Traditional file-level backup of data is most often employed to enable data recovery, but a series of steps must be performed to recover the system before data can actually be restored. With downtime tolerance in the hours to minutes for a business's most critical data, a data-centric protection strategy will leave the company at risk.

System-level backup and recovery is the perfect complement to file-based data protection strategies. Taking a more holistic approach by capturing a whole system rather than only the data layer will enable IT organizations to be compliant with RTOs, insulating the company from the impact of business interruption. The potential to save time and money, while improving the quality of IT service delivery, is difficult to ignore.

Enterprise Strategy Group | **Getting to the bigger truth.**

20 Asylum Street | Milford, MA 01757 | Tel:508.482.0188 Fax: 508.482.0218 | www.enterprisestrategygroup.com