Enterprise Strategy Group

IT Knowledge • Business Results

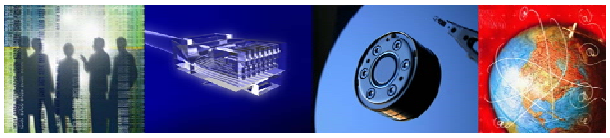

White Paper

Recovery-focused Data Protection:
Research Shows Your Future Depends On It

By Heidi Biggar
Storage Analyst, Data Protection
Enterprise Strategy Group

January, 2007

Table of Contents

Introduction

The Rising Costs of Downtime

Just a couple of years ago the "magic number" for recovering applications and data was four.

In a 2004 research report, ESG found that if applications were up and running and critical data were recovered within four hours of an outage or a system failure, odds were favorable the business impact would be manageable.[1] Five or more hours and the odds quickly soured. Users told ESG that the potential negative business consequences of this length of downtime would likely be significant and, in some cases, insurmountable.

However, as the value of data within organizations has increased and regulatory and corporate governance requirements have mounted, the "magic number" for application and data recovery has decreased by orders of magnitude, from hours to mere minutes or, in many cases, to near zero or zero. Not being able to recover data quickly can cost organizations big dollars in terms of lost revenue, potential fines and reputation or brand damage.

This is not to say that all applications and data are mission-critical (i.e., that they cannot withstand any downtime or loss, respectively) and therefore necessitate a true "recovery window" of zero. Some do and some don't. But there is a "range of tolerability" for all applications and data, and, as this paper will demonstrate, this range continues to narrow.

Falling above or below the uptime tolerability range for a given application can have significant business consequences: Fall above the range and you may be putting your organization's applications and data at risk in outage situations; fall below it, and you may be overprotecting these resources.

The key is to determine what this range is for your organization's data and then match data protection resources accordingly (see Figure One). ESG refers to this type of data protection infrastructure as a data protection continuum, or ecosystem. Build this type of environment and you'll be positioned to meet – or even exceed – service level agreements (SLAs) for application and data availability, uptime, and performance.

Figure One: VTL and the Data Protection Continuum

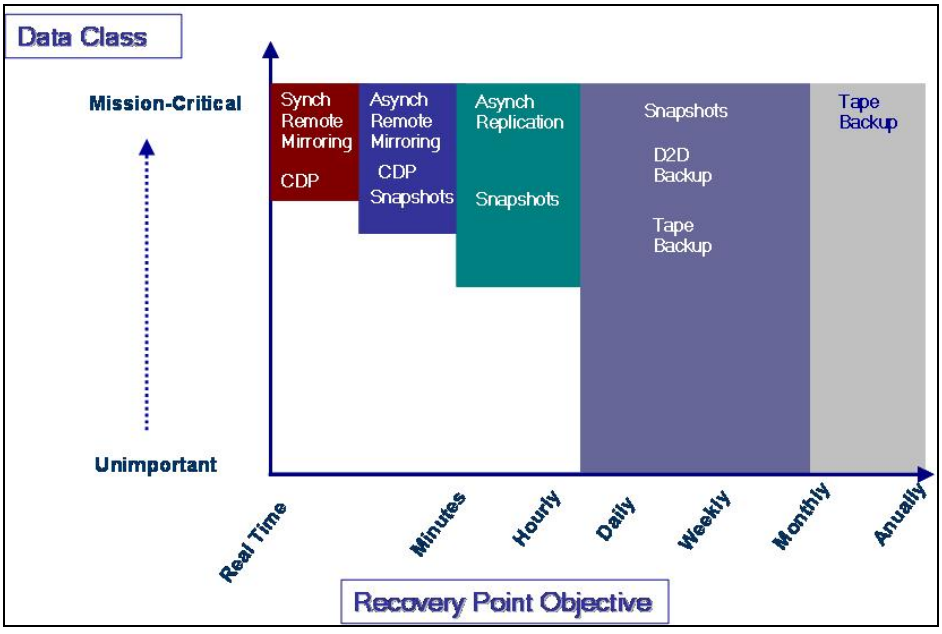

[1] ESG Research Report: *The Evolution of Data Protection*, December, 2003.

End-User Perspective: Failure to Adhere to SLAs Has Costly Results

"SLAs are what we live and breathe by. When a system or application falls outside of its respective range, the results can be significant. We recently experienced an outage with our outsourced e-mail ISP that resulted in a full day of unproductive time for our 56,000 employees. That outage has been estimated to have cost the company $1M."

And the Survey Says . . .

As a barometer of recovery practices today, ESG recently surveyed 156 storage professionals and IT managers at a cross-section of large ($500m or greater in annual revenue) organizations in North America. The findings were conclusive:

- Users are thinking more and more in recovery terms.
- Users see disk-based technologies as a means to better recovery performance.

Users who switched from tape to disk reported a dramatic improvement in both backup and recovery performance compared to those who still relied solely on tape. In fact, ESG's research revealed a direct correlation between an organization's implementation of disk-based, versus tape-based, data protection technologies and its ability to meet backup and recovery SLAs (see Figure Two).

Figure Two: Improved Data Recovery Performance by Backing Up to Disk

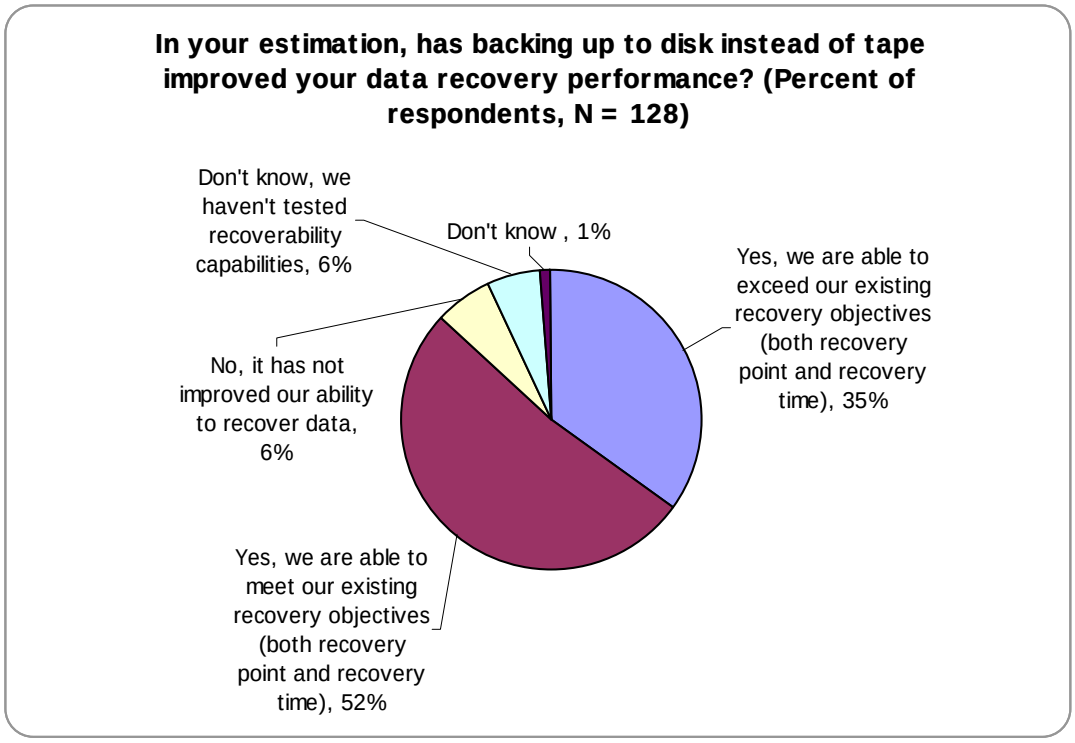

Specifically, 87% of survey respondents said that backing up to disk had improved their recovery performance, enabling 52% of respondents to "meet" recovery objectives and 35% of respondents to "exceed" them

Respondents who "backed up to disk and then to tape" also reported a significantly higher frequency of meeting recovery point and recovery time objectives (RPOs and RTOs) than those who backed up to "tape only."

Key Findings

A Case for Recovery Management

A powerful case can be made for a "recovery management" approach to data protection -- one that leverages a continuum of disk-based technologies [e.g., continuous data protection (CDP), snapshots, replication, disk as backup target, virtual tape library (VTL), etc.] to satisfy user requirements.

Consider the following proof points supported by ESG's recent research:

1. **SLAs are becoming more commonplace – and more aggressive.** This reality is forcing organizations to think "outside the box" when it comes to data protection.

   In ESG's most recent data protection survey, an SLA was defined "as the part of a service contract in which one party agrees to provide another party [i.e., an internal customer] with a certain level of service." IT SLAs are typically measured in terms of availability, uptime and performance of applications and IT services. The ability to – or more specifically, frequency of – meeting defined RTOs and RPOs is one measurement of SLA adherence.

   According to survey results, SLAs for e-mail/messaging, transaction processing and web-based applications are commonplace among organizations of various types and sizes today and are getting increasingly aggressive. While more respondents (69%) said they had SLAs in place for transaction-processing applications than for either e-mail/messaging or Web-based applications (58% and 51%, respectively), nearly one-third of respondents in each application category reported an average SLA of less than one hour (see Figure Three).

Figure Three: Applications Supported by Formal SLAs

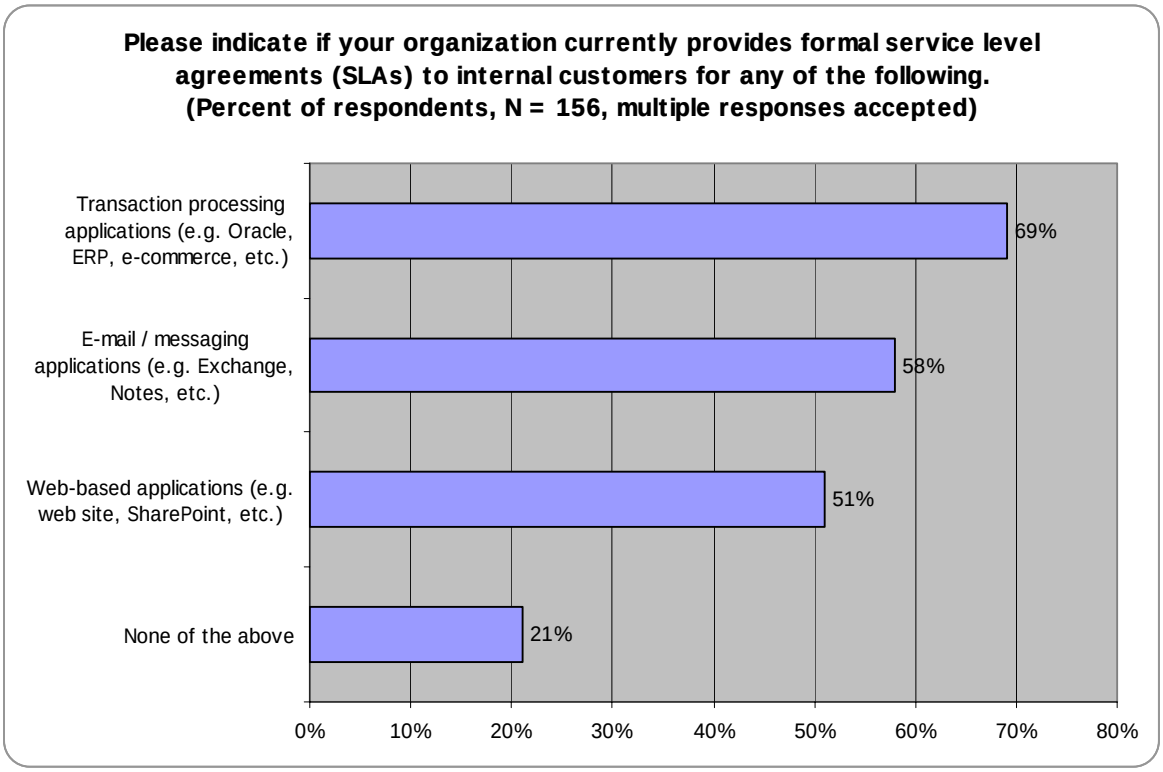

That said – and this is important to note – about half of the survey respondents also reported SLAs of two hours or more for these applications and about a third of respondents reported SLAs of four

hours or more (see Figure Four). The bottom line: While SLAs are getting increasingly demanding – due to the potential negative business consequences of protracted outages – there are still significant gaps in exposures. Backup-to-disk technologies can help users improve RTOs and RPOs, which in turn can help them better meet SLAs.

Figure Four: Average RTOs and RPOs for Most Important Applications

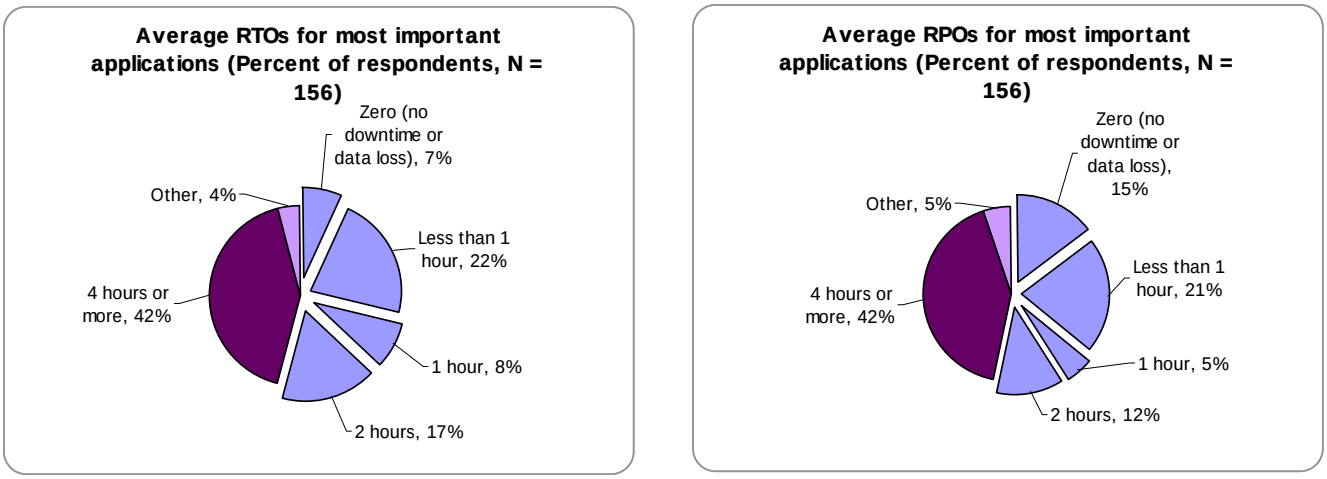

2. **Improving recovery performance (i.e., the ability to recover data to specific points in time and in specific time frames) is key to meeting – and potentially exceeding – SLAs**. Our research shows that users are looking for – and, importantly, are implementing – data protection technologies that quicken recovery times (i.e., RTOs) and minimize potential data loss (i.e., RPOs). These technologies include products like continuous data protection (CDP), replication, snapshot, virtual tape libraries (VTL), data de-duplication and backup and analytic software.

   ESG believes organizations should strive to meet recovery SLAs 100% of the time, whatever the RPO and RTO requirements. A typical user environment will likely have a range of RTOs and RPOs, to which a continuum of technologies will need to be matched.

   From a high level, survey respondents report good success at meeting existing RTOs and RPOs, with 16% of respondents saying they are able to meet RTOs and RPOs 100% of the time and 42% and 38% saying they were able to meet them 95% to 99% of the time, respectively (see Figure Five). ESG attributes this progress largely to the increased use of the backup-to-disk technologies listed previously. Users who rely solely on tape-based backups reported significantly lower success at meeting RTO and RPO objectives.

   However, while users are making great strides in improving recovery performance by using disk-based data protection technologies, ESG believes there is still room for improvement both in terms of improving the frequency with which they meet RTOs and RPOs and the manner in which they do it, especially as SLAs become more aggressive. Why? Because just one failure could lead to significant business consequences. In other words, an organization could meet RTOs/RPOs 99% of the time, but the 1% of the time it doesn't could be very costly. A critical application could go down and a huge business opportunity could be missed or hefty fines could be imposed.

Figure 5: Frequency of Meeting RTOs and RPOs

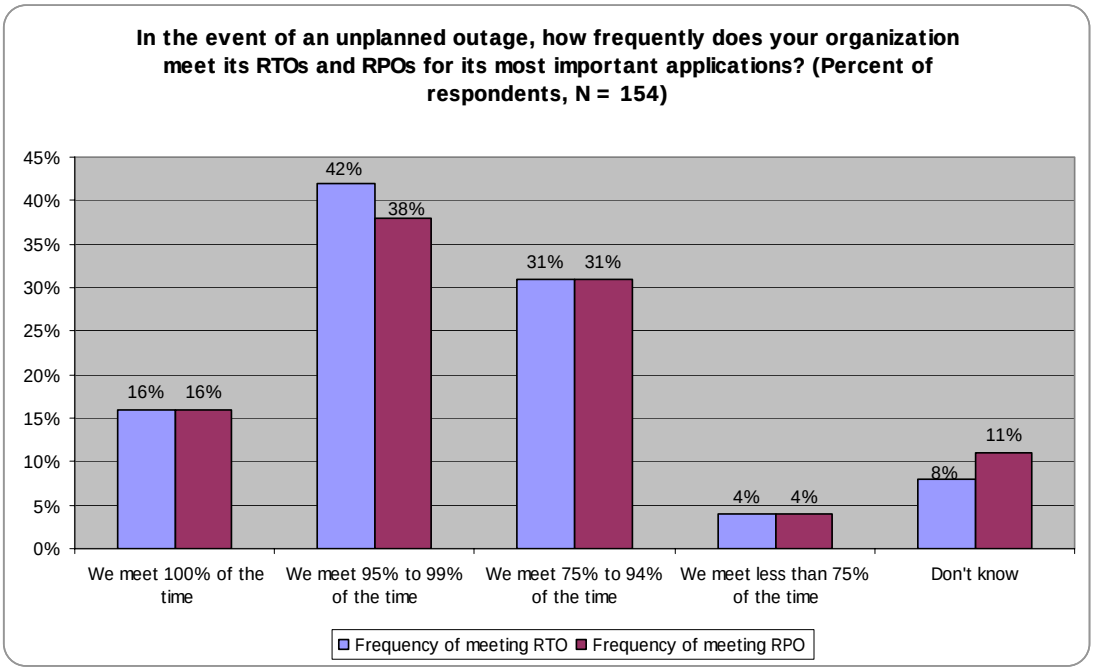

End-User Perspective: How Good Is Good Enough

"RTO and RPO objectives do vary by application and in some case by physical location. The 95% to 100% numbers also vary by acceptability. There is a continuous improvement program within our organization whose goal is to move those numbers closer to 100%. The program is enterprise-wide and therefore carries great incentives for realization."

3. **Users see disk-based data protection technologies as a way to boost backup and recovery performance**. It's a fact of which users are well-aware: Among the many benefits of backup-to-disk technologies is their ability to improve backup and recovery performance as well as backup and recovery reliability (see Figure Six). And, of course, they can lessen or eliminate tape management headaches, depending on the mix of technologies used.

   According to ESG's research, the number one reason respondents said they had deployed or planned to deploy backup-to-disk was "the need to reduce recovery times" to meet SLAs. The need to reduce backup times and improve backup and recovery reliability were significant secondary factors for making the switch from tape to disk for data protection.

Figure Six: Reasons Why Organizations Deployed Backup-to-Disk Technologies

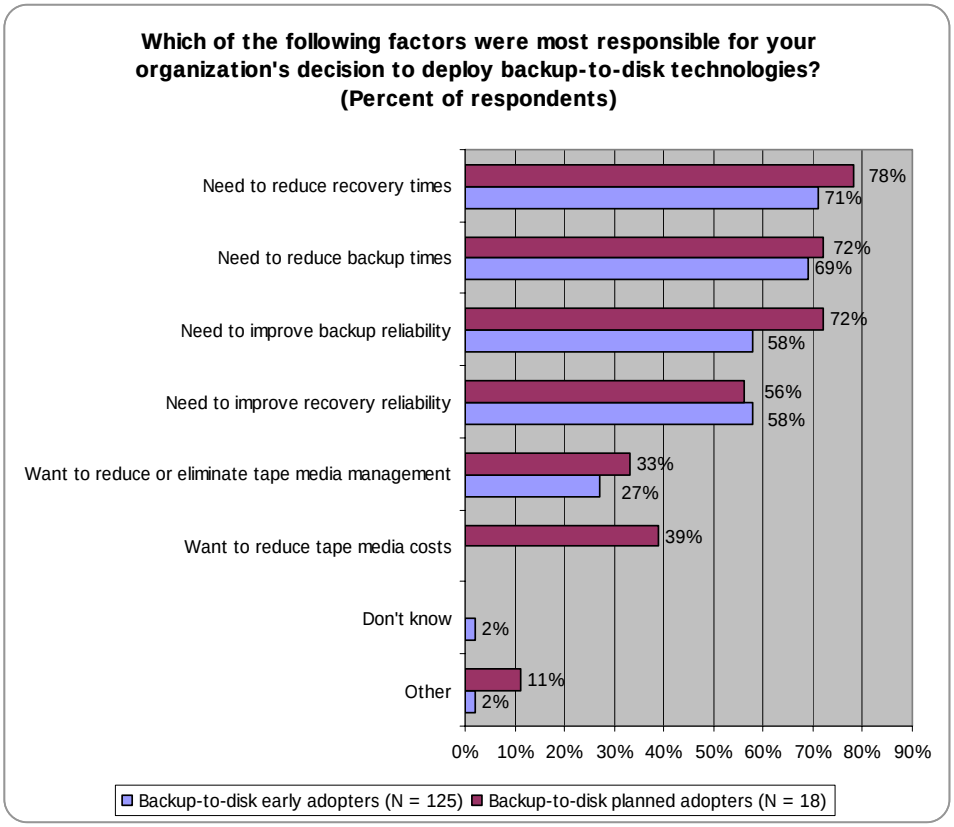

ESG's research indicates that when it comes to recovery objectives, disk-based technologies have lived up to their promise. As shown previously in Figure Two, 52% of respondents who have deployed backup-to-disk technologies say they are now able to meet recovery objectives (as dictated by SLAs) and more than one-third (35%) say they've been able to exceed them.

What is important to note here is that the failure to meet RTOs and RPOs can (as mentioned previously) have a significant negative business impact, resulting in application downtime, customer and user dissatisfaction, lost data, and in 20% of cases, direct revenue loss. In some cases, these effects can be crippling (see Figure Seven). This need to improve recovery performance will become increasingly important as SLAs become stricter over time. Again, having the right disk-based technologies in place will determine how successful users will be in adapting to and meeting evolving objectives. Existing and emerging disk-based technologies, such as CDP, VTL, remote replication, and data de-duplication, will play key roles going forward.

Figure Seven: Issues Experienced by Not Meeting RTOs and RPOs

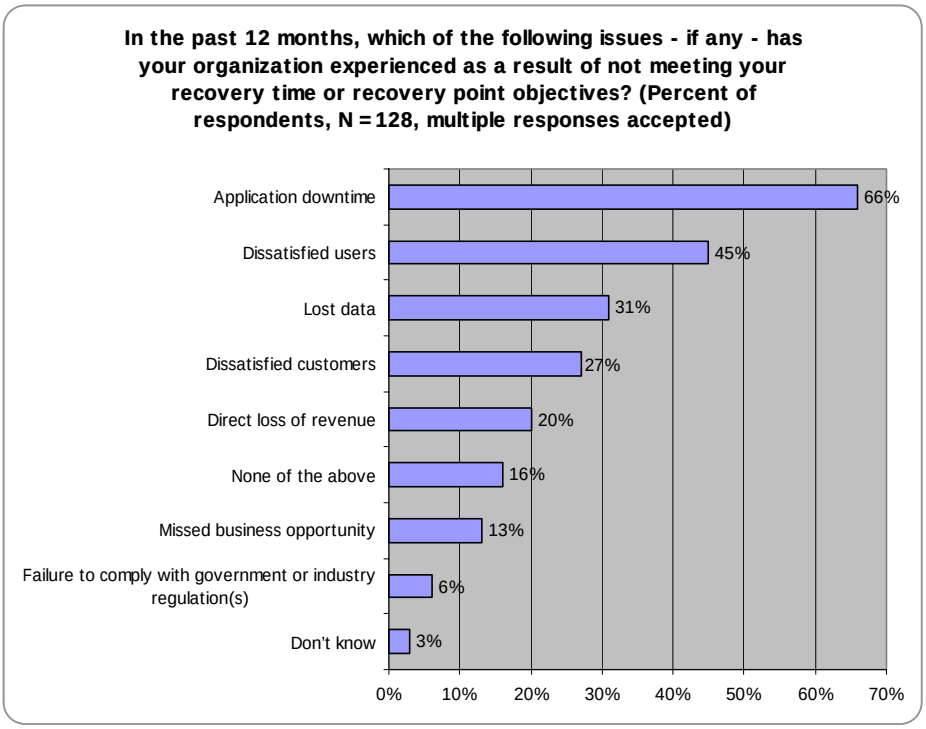

4. **Backup analytic software and a single management interface are on the end-user "wish list."** Users are ready and willing to adopt new data protection technologies, but still cite cost, complexity and management as ongoing challenges. As a result, end-users are looking for recovery-focused data protection technologies that don't just address RTOs and RPOs but also make the management of these technologies easier. Backup analytic software and a single management interface are two examples.

   By capturing and correlating information from backup and recovery operations, backup analytic software can help users better predict, identify and resolve potential data protection issues, including helping organizations better meet RPOs and RTOs – hence SLAs – and reduce instances of lost or unrecoverable data. ESG's research has found that end-users increasingly see the value in this type of technology: Seventy-three percent of survey respondents familiar with backup analytic software believe analytic tools would help their organization meet its SLAs more effectively.

   End-users are also interested in streamlining the management of the backup and recovery process. Why? Thirty-five percent of respondents to ESG's survey stated that having "too many application interfaces to manage" was one of their organization's greatest challenges in recovering applications, systems or data to meet application SLAs. Consequently, it is not surprising that 60% of survey respondents believed that having a single management interface for the recovery process would help reduce recovery times and 72% of respondents believed that such an interface would help increase recovery success rates.

So, what does this all mean? ESG believes end-users use of and interest in technologies designed to ease the recovery management process speak to:

- The increasing complexity of today's data protection environment.
- Increased internal pressure to meet and exceed SLAs.
- Better awareness among end-users about new data protection technologies and their capabilities.

Conclusion

The days of traditional backup practices being "good enough" are long gone. As the value of data has increased and corporate governance requirements have mounted, the time to get critical applications up and running after an outage has decreased by orders of magnitude, from hours to minutes or even near zero.

ESG's research clearly demonstrates that data protection is all about recovery. It's about being able to recover data in appropriate time frames and to acceptable recovery points. Just what is acceptable and what isn't is determined less and less by the IT staff and more and more by business departments.

Users who switched from tape to disk reported a dramatic improvement in both backup and recovery performance compared to those who still relied solely on tape. 87% of survey respondents said that backing up to disk had improved their recovery performance, enabling 52% of respondents to “meet” recovery objectives and 35% of respondents to “exceed” them.

A powerful case can be made for a **recovery management** approach to data protection -- one that leverages a continuum of disk-based technologies [e.g., continuous data protection (CDP), snapshots, replication, disk as backup target, virtual tape library (VTL), etc.] to satisfy user requirements.

ESG believes that organizations that recognize these trends, think about applications and data in recovery terms, and put disk-based recovery-focused data protection strategies in place that are aligned with business objectives will be well-positioned to not only tackle SLA demands as they arise (or before they do) but also raise the value of IT within their organizations.