Clipper Notes™

VENDOR-AGNOSTIC EXPLANATIONS AND ADVICE
FOR THE INFORMATION TECHNOLOGY BUYER

Navigating Information Technology Horizons℠

Data Backup May Not Be Enough —
The Key Is Recovery Management

Analyst: Michael Fisch

Management Summary

Is your business ready to recover in the event of a computer system failure, data corruption, operator error or a site outage? Consider these scenarios from the perspective of an IT manager:

Scenario #1: E-mail Virus Attack

The phone rings and it is the CEO. He wants to know why e-mail is down and how quickly you can restore it. After all, e-mail is a business critical application these days. You discover that a potent computer virus has infected the Microsoft *Exchange* server and must restore the data to just prior to the virus attack. *Tick tock, tick tock...*

Scenario #2: Order Processing Is Down

A hardware fault on a database server causes the processing of customer transactions to cease. The business is unable to take orders, check status, or ship products. For every hour the database is unavailable, significant revenues are lost or deferred. In fact, if it is not restored within a few hours, it will have a material effect on this quarter's financials. *All eyes are on you to restart the database on a failover server at a secondary data center.*

Scenario #3: Business Is Flooded

A flash flood soaks your enterprise headquarters, which unfortunately includes the primary data center in the basement. The executives want you to make sure all data has been appropriately backed up. *Moreover, since the flood is not expected to recede for days, they want you to restart key applications being replicated to a remote site ASAP, so the business can continue operating.*

Recovery Management and Replication

While it is human nature to prefer not to think about these kinds of worst-case scenarios, they do happen and prudence suggests it is better to be prepared than just to hope for good luck. The profitability and continuity of the business depend on it. For these and other data unavailability scenarios, traditional tape backup may not be enough to ensure a timely and full recovery. The good news is that complementary and affordable technologies are available to fill out an enterprise's data protection strategy. These include backup to disk, snapshot copies, continuous data protection, replication, and monitoring and analytics. In particular, replication offers a fast and complete recovery from system failures as well as local and regional disasters.

The key to all of this is *recovery management.* This approach embraces the whole picture of data protection to create a solution that more specifically and consistently meets an enterprise's recovery requirements. Read on for details.

Challenges of Data Protection

One could debate whether money or information is the true currency of modern business. Ask yourself which is worth more: One million dollars or...

- A field analysis of an untapped oil deposit?
- A patent for a new cancer treatment?
- The current customer list and accounts receivable for your enterprise?

The point is that information is valuable and, like money, needs to be protected. However, this is a challenge for several reasons.

Data is a moving target because it never stops growing. This phenomenon is a by-product of the Information Age. All enterprises face it, whether their data growth rate is 30%, 50%, or 100% per year. It presents a real data protection challenge.

Enterprises face stricter requirements for information protection and retention. The primary reason is businesses are more dependent on IT. Workers cannot get along without e-mail, Web access, and enterprise applications. A company's operational and strategic secrets are wrapped up in its data. Regulatory compliance adds legal mandates and accountability to the equation.

Protection requirements also vary because not all data is of equal value. For instance, archived financial records from ten years ago are worth less than customer transactions from the last month. The challenge is that it is too costly to give all data the highest level of protection, on one hand, while it is too risky to forego rock-solid protection for the most valuable jewels of data, on the other. To address this, enterprises are turning to the strategy of information lifecycle management (ILM). It lets them classify data and apply different and appropriate levels of protection, recovery, and performance.

Adding another dimension of complexity are the multiple technologies now available for data protection. They include backup to tape, backup to disk, snapshot copies, replication, continuous data protection (CDP), and monitoring and analytics. Each technology has different strengths and capabilities and meets different requirements. (*See sidebar for brief descriptions.*) Traditional backup to tape, while important, may not be enough.

Finally, the data protection infrastructure

Data Protection Technologies

Backup

The most common way to protect data is daily backups. Tape has traditionally been the preferred media, though more and more companies are backing up to low-cost disk, at least for the initial target, because it speeds up both backups and restores.

Replication

Replication generates a real-time data copy at a remote site that can be used for a rapid failover in the event of a local disaster or system failure. It is also useful for data migration and centralized backup.

Snapshot Copies

These data copies are made directly on disk and can be created more frequently than backups. Snapshots are useful for quick recovery to a recent point in time, non-disruptive backups, and data repurposing.

CDP

If snapshots are like still photos, then CDP is like video. It provides granular recovery to virtually any prior point.

Monitoring & Analytics

These software tools give insight into the data protection process. For instance, it can help spot and resolve backup problems.

---

can be complex to manage. An enterprise might have multiple management tools that do not communicate with each other. The tools may not provide clear insight into what is happening, such as if backup jobs are failing and why. Too often, the emphasis is on creating backups and copies, with little attention paid to the ability to recover and testing. The result is high management costs and insufficient reliability.

Recovery Management

So, where does one start amidst all of these challenges? **Recovery management is a comprehensive strategy and approach to data protection that starts with the end in mind – with recovery itself.** It views data protection as:

- A single process,
- Whose goal is recovery, and
- Which uses different technologies to meet multiple service levels more fully, easily, and reliably?

This stands in contrast to the view of data protection as a set of uncoordinated activities, whose goal is producing backups, being reactive, and hoping for the best. In short, recovery management is a new mindset that goes beyond the status quo.

Recovery is restoring access to data after system failure, data corruption or accidental deletion, or a site outage. It is defined in terms of how quickly and fully data is re-covered after an interruption. These are known as *recovery time objective (RTO)* and *recovery point objective (RPO)*, respectively. These requirements are determined by busi-ness processes and applications. In other words, they have to fit the business. And, most importantly, recovery enables business continuance – keeping an enterprise going, de-livering goods and services to its customers, and staying financially sound.

Recovery management classifies data into service level tiers and employs multiple technologies to meet their different require-ments. The technologies include established ones like backup, replication, and snapshot copies, as well as emerging ones like con-tinuous data protection (CDP) and monitoring and analytics. Some vendors have begun inte-grating these at the management and/or func-tional layers, making them easier to use together.

Replication

Replication is one of the key recovery management technologies, because it offers distance and immediacy to data protection. It maintains a complete physical copy of data on disk at a remote site in real time. Like the spare tire in the trunk, if the source data becomes unavailable for some reason, a current or nearly current copy is available to resume operations in seconds or minutes. Replication protects from computer system failures as well as local and regional disasters, such as fires, floods, or electricity outages. It can is also useful for consolidating data from branch offices for centralized backup, or for sending data in the reverse direction for distribution.

The two fundamental modes of replication are synchronous and asynchronous, which differ in their levels of performance, distance, and protection. Synchronous provides a fully current copy at a remote site, so no data is lost when a local failure occurs. The tradeoff is potentially slower application performance as well as a distance limitation on the order of 100 km. Asynchronous provides a *nearly* current copy over virtually unlimited distances without affecting application performance. The tradeoff is a lag of seconds to minutes between when writes are committed to the source and target. If a disaster occurs, this data lag is exposed to loss. Advanced, three-site configurations can do both – synchronous replication to a metro-area site and asynch-ronous to a site in a different region or country.

Replication solutions may run on a host server, storage array, or network platform. The most sophisticated solutions are array-based, though server-based solutions stand out for affordability and flexibility. Which you choose depends on your environment, budget, and recovery needs.

Conclusion

Recovery management is a smart approach because it enables faster, easier, and more reliable recovery. Talk to your vendors and strategic partners about it, and ask how their data protection technologies fit into the broader picture of recovery management. It is worth taking the time to do well since the success and continuity of the business depend on it.

And, remember:

- Focus on recovery and the whole picture of data pro-tection.
- Look at it as a contributor to business continuity and productivity.
- Classify data and map it to service level tiers.
- Use the right set of tech-nologies to meet your full spectrum of requirements.

About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

➢ *The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.*

About the Author

Michael Fisch is Director of Storage and Networking for The Clipper Group. He brings over ten years of experience in the computer industry working in sales, market analysis and positioning, and engineering. Mr. Fisch worked at EMC Corporation as a marketing program manager focused on service providers and as a competitive market analyst. Before that, he worked in international channel development, manufacturing, and technical support at Extended Systems, Inc. Mr. Fisch earned an MBA from Babson College and a Bachelor's degree in electrical engineering from the University of Idaho.

➢ *Reach Michael Fisch via e-mail at mike.fisch@clipper.com or at 781-235-0085 Ext. 211. (Please dial "211" when you hear the automated attendant.)*

Regarding Trademarks and Service Marks

The Clipper Group Navigator, **The Clipper Group Explorer**, **The Clipper Group Observer**, **The Clipper Group** *Captain's Log*, **The Clipper Group Voyager**, **Clipper Notes**, and *"clipper.com"* are trademarks of The Clipper Group, Inc., and the clipper ship drawings, *"Navigating Information Technology Horizons"*, and *"teraproductivity"* are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.