HITACHI
Inspire the Next

WHITE PAPER

Data Protection: Downtime Is Money

Smarter Approaches to Data Protection and Recovery

By Hitachi Data Systems

June 2017

Contents

Executive Summary

The primary attributes for measuring the effectiveness of a data protection and recovery solution are backup window, recovery point objective (RPO) and recovery time objective (RTO). These points measure the time it takes to perform the protection operation, the interval between protection operations, and the time it takes to perform a recovery operation following an error, failure or disaster.

Each of these assessments also are a measure of downtime: the time data is not available during a backup; the time it takes to recreate the new data that wasn't protected yet; and the time it takes to restore operations. Depending on where the system and data being protected fit within your organization, the amount of downtime caused by data protection and recovery operations will result in some amount of financial impact to the organization.

Andrew Lerner, research vice president at analyst firm Gartner pegs the average cost of downtime across industries at $5,600 per minute, or more than $300,000 per hour[1]. The unavailability of critical data has been shown to cost large organizations millions of dollars per hour.

It would be easy to demand that data, especially business-critical data, be protected continuously against all types of threats, without impacting business operations, and enable instantaneous restoration and recovery in all situations. This is where the fourth measure of effectiveness comes in: cost.

Just as with everything else in information technology, and indeed in life, trade-offs and compromises must be made in designing an effective data protection infrastructure that:

- Provides the minimal levels of service (backup window, RPO, RTO) that the business absolutely requires, for each different set of data, based on its value to the organization.
- Presents the lowest possible cost to the organization.

Data protection investments can be best categorized as an insurance expense. They are known and categorized as costs that do not drive revenue and provide no other value to the organization until something bad happens, and then they only impact the bottom line. However, that isn't really true.

In this paper, we examine the hidden costs and the interrelation of costs and policy choices in the area of data protection and recovery, and how solutions from Hitachi Data Systems can help you avoid these surprises.

[1] http://blogs.gartner.com/andrew-lerner/2014/07/16/the-cost-of-downtime/

Introduction

The information technology infrastructure has grown very complex through a proliferation of server, storage and networking platforms, applications and distributed locations. These developments, combined with annual compound data growth in the 40% to 50% range and ever tighter service level demands by business stakeholders, has made the life of those responsible for data protection exceedingly difficult. Figure 1 outlines some of the complexity, showing some of the combinations of situations that may need to be addressed in a data protection and resiliency plan.

Figure 1. Backup administrators face myriad infrastructure complexities.

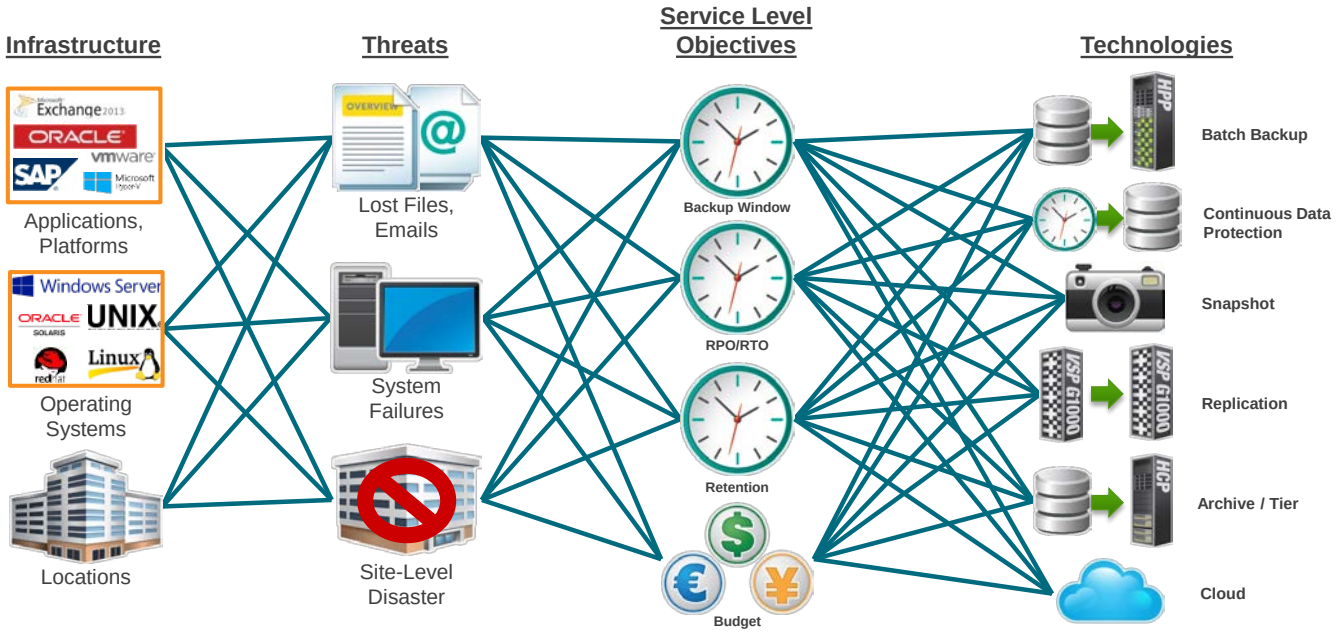

RTO = recovery time objective, RPO = recovery point objective, HPP = Hitachi Protection Platform, VSP = Hitachi Virtual Storage Platform, HCP = Hitachi Content Platform

For each different part of the infrastructure, you need to protect against a variety of threats in a way that supports the service level objectives for each line of business. The service level objectives are the policies that the data protection solution set must deliver against, leading to more than a few technology choices.

This paper focuses on the service level objectives, how they affect each other, and how they impact both obvious and hidden costs.

How Much Does the Backup Window Cost?

Many in the data protection and recovery market refer to “backup window” as the amount of time that it takes to complete backup operations. Technically, that isn’t correct. The real backup window is the amount of time, and the start and stop times, that the organization has allotted for performing the backup for a particular system or dataset. The backup window is one of the four goals that are used to measure the effectiveness of the data protection process; the others are the recovery point objective, the recovery time objective, and the overall cost.

If it takes longer to perform a given backup than the allowed backup window, the potential costs to the business are fairly evident: Either the backup is stopped before completing, leaving important or even critical data at risk of loss, or the backup is permitted to run longer. Each choice impacts the availability of the system being protected, and therefore impacts operations.

In many cases, the backup window is a necessary evil. The business concedes the need for this amount of downtime. Creating the backup is recognized as being important to the survival of the organization, even though it may be at the expense of operational effectiveness and profitability. If you could reduce the amount of time that it takes to perform backups, and reduce or even eliminate the associated backup window, the business would be in position to recoup those lost opportunity costs.

The other time-related cost in backup operations is administrator time. In the “old days,” backup was a manual, labor-intensive operation. The backup administrator needed to:

- Make necessary installation and configuration changes for each new system or user.
- Label the daily tapes, then perform and monitor the backup operations.
- Remove tapes for transport to the off-site disaster recovery site or vault.
- Erase expired backup tapes for re-use.
- Troubleshoot problems and rerun the backup (if the backup window allows it).

Intelligent policy engines, system automation, auto-discovery technologies, the use of disk-based backup storage, and replication to the disaster recovery site save effort and cost. These features and approaches provided the opportunity to eliminate most of the manual effort and reduce the time-related costs of performing data protection operations. However, a large number of organizations are still living in the old world, or have given up and outsourced their backup operations to third-party service providers.

With data protection and recovery solutions from Hitachi, the backup window can be minimized, or even eliminated. These results can be accomplished using fully integrated technologies such as block-level continuous data protection (CDP) and application-consistent, hardware-assisted snapshots. These approaches eliminate the need to scan the file system for incremental changes and reduce the time to copy the data to mere seconds.

Incremental-Forever Backup

Snapshots and CDP capture only new or changed data, using an incremental-forever model that eliminates the need for the periodic full backups (or synthetic full backups) of most other solutions. This approach can reduce the amount of backup data being stored by more than 90%, depending on the data change and retention rates.

For example, let’s assume we have a normal business, school or agency that operates five days per week, 50 weeks per year. We have 100TB of data, and a daily data change rate of 10% (10TB), with a weekly change rate of 50% (50TB). We retain our backups for 12 weeks for operational recovery, and assume that data that needs to be retained longer is archived.

This is an extreme case that illustrates the differences in the traditional and incremental-forever models. A more typical environment would have a total data change rate of 50% per year (50TB), which equates to 1% per week (1TB), and 0.2% per day (200GB).

The old full + incremental model and the HDS incremental-forever model (see Figure 2) each copy 10TB per weekday, but the full + incremental option copies the full 100TB on the weekend while the incremental-forever system takes the weekend off.

Figure 2. Differences in Traditional and Incremental-Forever Models

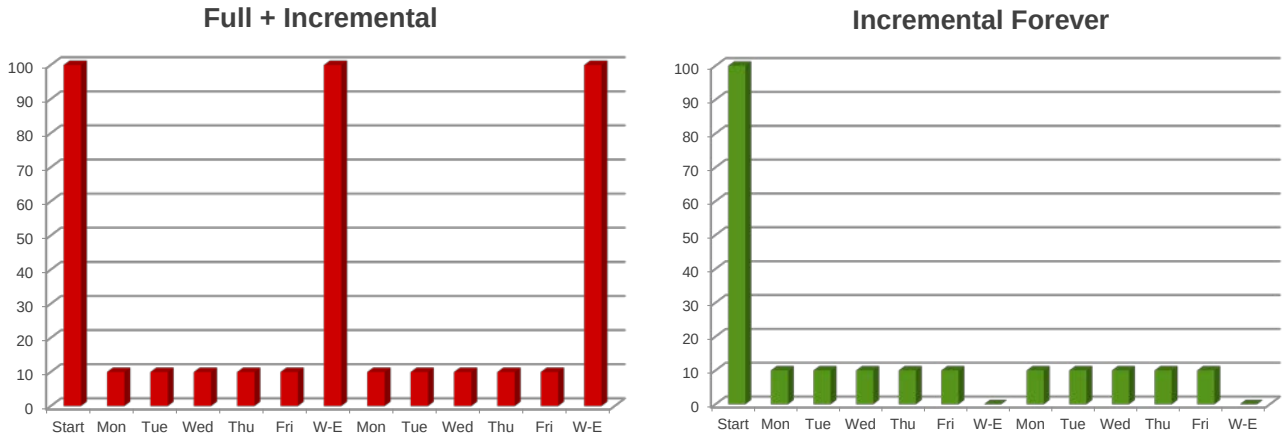

Including the initial full backup (100TB), the total backup storage capacity needed for 12 weeks for each model is:

- Full + incremental: 1,900TB (1.9PB).
- Incremental-forever: 700TB (0.7PB).

Incremental-forever backup reduces capacity requirements by 63%, without spending any money or compromising system performance on data deduplication. How much does 1.2PB of backup storage cost to acquire, manage and maintain? Actually, it is 2.4PB of extra storage, since we'll want to replicate the backup repository for off-site disaster recovery. If the backup data is retained for longer than three months, these savings increase even more.

In the more typical environment with a 50% change rate per year, you would need 1.3PB to store 12 weeks of traditional backups, and only 112TB for incremental-forever backups, for a savings of 91%. It's still a reduction of 1.2PB, as in the extreme example, confirming that almost all the data in the traditional backup repository is redundant.

Taking control of the backup window is the first step in reducing the costs, and risks, of data protection and enabling more effective operational recovery and disaster recovery. Hitachi Data Instance Director helps you meet this challenge through its host-based CDP and storage-based snapshot capabilities.

How Does RPO Impact RTO?

Recovery point objective, or RPO, is the measure of the granularity of previous points in time that you want to be able to recover a particular data set from. An RPO of 24 hours says that a single backup operation per day is "good enough." Other ways to describe it include:

- The frequency of backup operations.
- The amount of new data that you are willing to risk losing.

RPO is totally different than recovery time objective, or RTO. RTO is the goal for how long it should take to restore a system or application, or restore access to a dataset, following an unplanned event such as caused by human error, hardware failure or natural disaster.

RTO defines how much time and, therefore, how much cost, risk or lost revenue opportunity, the organization is comfortable absorbing following an outage or disaster. Often, different recovery time objectives are established for different types data and different types of outages. Consider, for example: two hours to restore a lost file or email, six hours to restore a crashed server, or two days to restore operations following a site-level disaster.

Since RPO and RTO are totally different, does one impact the other? A common reaction to this question is “no”, but, in reality, the way that you reach your RPOs can have a dramatic impact on your ability to meet your RTOs. In fact, as depicted in Figure 3, it is really a tug-of-war between them.

Figure 3. How do you balance RPO and RTO?

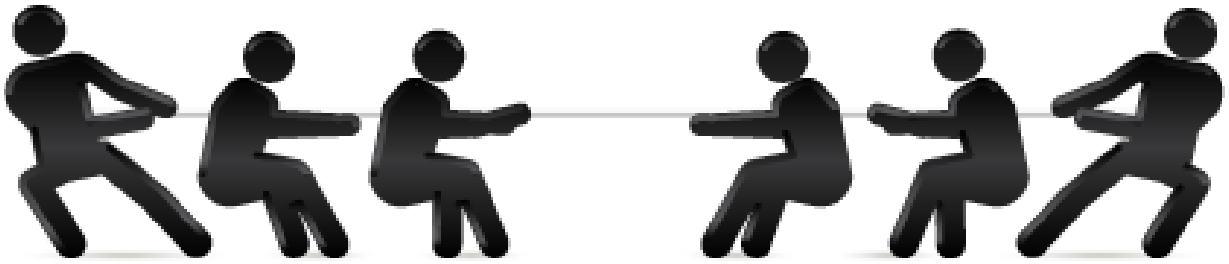

Recovery Point Objectives (RPO) Recovery Time Objectives (RTO)

Let’s say you have a very large database that can only be backed up on long weekends. To meet an RPO of 24 hours, you need to back up the database journals or redo logs on a nightly basis. This way, you can restore the last full database backup and then roll-forward database transactions using the journals or redo logs.

The number and size of these files that need to be restored and applied to the database files can grow dramatically, especially if you have large, clustered database environments such as Oracle Real Application Clusters (RAC).

So, will the time it takes to restore the last full backup and then all the journals meet your established RTO for this large database system? Unless the RTO is measured in weeks or months, the answer is, probably not. The conclusion is that this methodology of database protection can be used to meet reasonable recovery point objectives, but it will not support a reasonable recovery time objective.

A similar situation exists in the traditional full + incremental backup method described earlier. In this model, you typically (hopefully) complete a full backup each weekend, and then perform an incremental capture of each day’s new data during the week. If you suffer a failure on Monday and need to perform a full recovery, no problem: Just restore the full backup from the weekend.

But if your failure happens on Friday, you need to first restore the full backup from the previous weekend and then each of the incremental sets, sequentially, from Monday through Thursday evening. A recovery on Friday will take significantly longer than one on Monday. Does your RTO take this into account? Also, the recovery on Friday is more fraught with risk, since it’s a multistep manual process and you may be overwriting some of the restored data, maybe as many as four times.

Clearly, as data volumes continue to increase in a compound fashion and the IT landscape becomes ever-more complex, better approaches to meeting backup (RPO) and recovery (RTO) requirements are needed.

Hitachi offers a solution to the problem of protecting large databases and critical applications that both enables much better RPOs and meets tight RTOs. This solution has three facets.

- Storage-based snapshot and replication technologies that:
  - Eliminate the load of data protection operations from the database system.
  - Eliminate the need for a backup window and associated downtime.
  - Enable much more frequent backup operations, reducing the amount of new data at risk by 90% or more.

- Database- and application-aware snapshot and replication management software that:
  - Places the database or application into a backup-ready or quiesced state.
  - Executes the storage-based snapshot and then releases the database or application to continue normal operations.
  - Enables, fast, fully application-consistent operational recovery – in minutes, not weeks.
- Assessment and implementation services that define and configure the optimal solution for your unique environment.

The Hidden Cost of RTO May Be Your RPO

What goes into an RTO? It could include any or all of the following, depending on your definition:

- Time to discover and diagnose the event.
- Time to take corrective action: Install a new server, replace a disk, send the offending human to a "timeout," or failover to a disaster recovery site.
- Time to reinstall operating systems and applications, if necessary.
- Time to copy all the needed data from the backup or disaster recovery system.
- Time to start up and test the recovered environment.

That can be a lot of time. Downtime. Time that some part of your business is not being productive. Depending on the people and applications affected, this downtime can have a profound impact on your top line, your bottom line, or both.

But there is one element that is often left off this list, and it also impacts both the time to fully recover and the total cost of the recovery, and that is the recovery point objective. An RPO of 24 hours (typical nightly backup) indicates that you are willing to lose up to one day of new data when something goes wrong.

Often, the RPO is defined for practical reasons, such as you can only take down a certain system during nights or weekends. However, RPOs really should be set by business needs, not because of the compromises that the current backup software forces upon you.

But let's say you do have an RPO of 24 hours and the system crashes at 6 p.m., deleting or corrupting all the data contained in it. You can restore the data from the last backup, but everything created or changed after that time is gone.

Are you really planning to just lose that data? It could include several large orders from the sales system, or a day's worth of design effort, or many other things that are important to the organization. Are you going to just shrug your shoulders and move on? Of course not. You are going to have to recreate that data. That process will take time, and this is time that your staff would have spent on normal activities, further impacting the effectiveness of the business during this recovery period.

The conclusion is that the longer the time between backup operations – the RPO – the more data that will likely need to be recreated following an outage, and the higher the overall cost to the organization. And it could be more than just tangible costs: Imagine going back to your customer and asking them to re-enter that million-dollar purchase order because you had a system failure. Ouch!

The Solution Is Clear

To solve this puzzle, let's jump to the conclusion that we're looking for: Limit the time, and therefore the money, that it takes to recover from any failure. We need to:

- Significantly reduce or eliminate the backup window that constricts the frequency of protection operations (the RPO).
- Significantly increase the frequency of protection operations so that far less data is at risk of loss and requiring re-creation.
- Speed recovery operations, whether locally (operational recovery) or remotely (disaster recovery).

Hitachi offers two approaches to meet these needs, with a single software solution, based on whether Hitachi storage systems are managing the data that is being protected. Hitachi Data Instance Director (HDID) provides unified copy data management for organizations looking to modernize and simplify their data protection, retention and recovery operations. It reduces costs and improves data availability service levels.

Storage-Based Data Protection and Recovery

Hitachi Data Instance Director configures, automates and orchestrates the local and remote replication capabilities of the Hitachi Virtual Storage Platform (VSP) family, including the VSP F series and VSP G series, as well as the Hitachi NAS Platform (HNAS) family. See Figure 4. This integration provides the ability to create fast, frequent copies of production data, with no impact on the performance of the production system. Snapshots, clones and replicas can also be mounted for near-instant access and recovery.

Figure 4. Hitachi Data Instance Director provides unified copy data management.

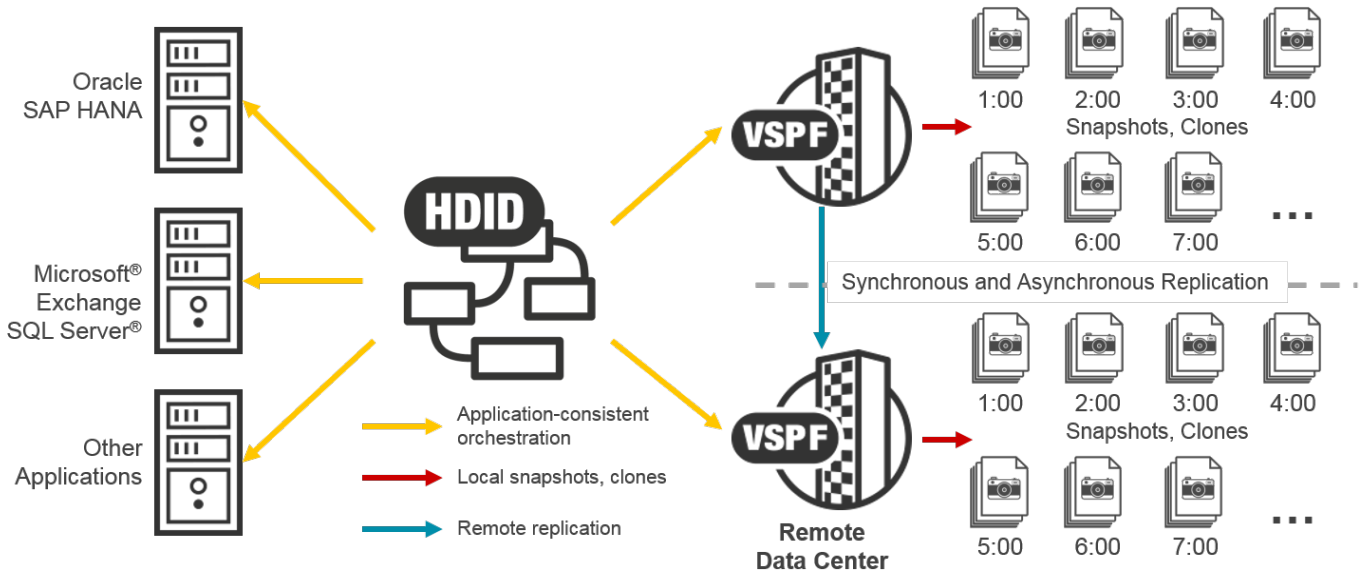

Specifically, HDID supports:

- On VSP family:
  - *Hitachi In-System Replication*, including Hitachi Thin Image snapshot software, which creates a point-in-time, space-efficient, high-performance, incremental copy, and Hitachi ShadowImage software, which creates a local full-copy clone nearly instantly.

- *Hitachi Remote Replication*, including Hitachi TrueCopy synchronous remote replication, which mirrors volumes over metro distances for disaster recovery and high availability, and Hitachi Universal Replicator for long-distance volume replication to support disaster recovery and business continuity.
- Storage clustering via our *global-active device* feature, which eliminates the need for failover and failback across campus distances.
- Application-consistent protection for Microsoft SQL Server and Exchange environments, as well as Oracle and SAP HANA databases. Other applications may be protected using custom scripts.

- On HNAS and the VSP NAS Module:
  - Directory clone quickly creates a writeable snapshot copy of a directory tree.
  - File replication enables fast, granular off-site copy, which is ideal for short-duration disaster recovery and data migration.
  - Application-consistent protection for Oracle databases. Other applications may be protected using custom scripts.

With the unique whiteboard-style graphical user interface in HDID, you can easily combine backup, snapshots and replication of specific data sets into complex workflows that meet the service level objectives of the business.

Host-Based Data Protection and Recovery

When production data is not stored on HDS block or file storage systems, HDID can still provide modern, high-performance data protection using its host-based capabilities:

- Continuous data protection constantly captures application data, providing an up-to-the-second copy of the data. The built-in CDP policy engine allows you to set how often an application-consistent snapshot of this data is sent to the HDID repository. The data is captured incrementally, at the block level, as each block is written to disk, with no performance impact on the production environment.
- HDID also provides scheduled, batch-style backup, again using an incremental-forever data capture model.
- Disaster recovery capability is provided through scheduled, asynchronous replication of the HDID repository to a heterogeneous storage system in another location.

As noted above, Hitachi Data Instance Director is a powerful tool in the fight to increase availability and reduce the data protection costs for critical database and application environments. Hitachi also has a powerful solution for unstructured data with the Hitachi Content Platform (HCP) family, which is built on highly scalable, self-protecting object storage. The family includes a file gateway (Hitachi Data Ingestor) and a robust file sync-and-share solution (Hitachi Content Platform Anywhere) that mobilizes and protects user data.

To learn more about Hitachi solutions for data protection visit www.hds.com/go/protect or contact us at dp-sales@hds.com.

Hitachi Data Systems

Corporate Headquarters
2845 Lafayette Street
Santa Clara, CA 95050-2639 USA
www.HDS.com | community.HDS.com

Regional Contact Information
Americas: +1 866 374 5822 or info@hds.com
Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hds.com
Asia Pacific: +852 3189 7900 or hds.marketing.apac@hds.com

WP-530-C R.Vining June 2017