HITACHI
Inspire the Next

Simplify Data Protection, Retention and Recovery

Use the Right Tool for Each Job, but Manage Them From One Place

By Hitachi Data Systems

July 2016

Contents

Executive Summary

Over time, the information technology landscape has become exceedingly complex. Each new breakthrough ushers in not only new opportunities to accelerate your business, but also new layers of systems, applications, processes, network protocols and skills. Sometimes these new solutions replace aging legacy infrastructure, but often they simply add more complexity.

With complexity comes cost and risk. This is especially true when considering how to protect, retain and recover the data that is created, stored, manipulated and exploited by the new systems.

There are many examples of this phenomenon in the history of computing: the transition from mainframe to client/server, virtualized servers, remote and branch offices, mobile workers, big data and analytics, in-memory databases, converged and hyperconverged infrastructure, containers, and private, public and hybrid clouds. Each of these advances, and many more, have created the need for new solutions for the protection of data, whether it be installing a new agent or module for your existing software, or a totally new approach. Either way, it adds complexity.

Your legacy backup solutions are usually not capable of handling the newest technologies. If you are using a solution from one of the big vendors in the backup market, it will probably be years before they develop, test and release a solution for your new system. This opens the door for start-up vendors to enter your data center. They are in business solely to solve your new, specific need, and they often do it well. But they just add one more thing to find, acquire, learn, manage, monitor and report on, maintain, upgrade and eventually migrate away from and decommission.

It gets worse. The new data protection “point solutions” don’t integrate with your existing backup tools, and they often cover only one of the aspects of the data management paradigm:

- Operational recovery.
- Disaster recovery.
- Long-term retention.

At some point, you will step back and look at all the different tools you’ve deployed to keep all of your data safe, and when you do, you will likely say something like, “This is CRAZY!” Some data sets are protected by multiple tools. Each creates its own copy of the data, with its own retention and ownership policies. Who has control and knowledge of all this copy data? Does any of it meet corporate governance and security policies?

Better questions to ask are: When something bad happens and you experience a data loss, will the right person, with the right training, be able to log into the right system, find and retrieve the right data? Will they be able to restore it to the right place, do it within prescribed service level agreements, and not break anything else along the way?

If you cannot answer these questions with an emphatic “Yes,” you should probably take action before something bad does happen.

This white paper describes the various challenges, variables and risks that need to be considered. It goes on to describe how Hitachi Data Systems is working toward a unified approach to protecting, retaining and recovering data across a complex IT environment. And it looks at the results, which can include dramatic reductions in costs, complexity and risk, all leading to a better night’s sleep.

Introduction

The world is changing faster than ever, and most organizations are struggling to keep up. This is especially true for IT departments, as new technologies, applications and solutions burst onto the scene. It is increasingly difficult to maintain the quality of services the business demands, due to a combination of factors:

- It is difficult for IT staff to maintain currency on all the new and improved technologies.
- The amount of data continues to grow, straining the integration and capabilities of all types of systems and processes.
- Operations and staff are becoming increasingly distributed and mobile, and therefore outside of the physical and logical boundaries of IT.

As the IT department struggles, the business is forced to find new ways of accomplishing its mission and keep up with the competition. This leads to the outsourcing of many common IT responsibilities, such as sales-force automation, content management, project management, communications and even general data storage services. This shift does not come without risk to the organization. As your data is handled by third parties and out of the control of central IT, the door is open to security breaches, data loss, litigation and even failure to comply with data retention regulations.

To combat these threats, IT leadership should be looking at all possible opportunities to restore control, while providing the services and flexibility today's business leaders require. A key part of this effort is to simplify IT operations and processes to provide maximum agility while maintaining service level objectives, and above all reducing costs and risks.

Hitachi Data Systems can help in a number of ways. Options include providing private cloud converged infrastructure with Hitachi Unified Compute Platform and simplifying the storage infrastructure with Hitachi Storage Virtualization Operating System (SVOS). We can also reduce primary storage requirements through archiving or tiering to Hitachi Content Platform (HCP), as well as providing a number of other opportunities.

Additionally, Hitachi Data Systems Global Services Solutions (GSS) has helped many organizations to simplify their data protection, retention and recovery environments. We help examine the depth of complexity in each IT environment that is driving the need for a more unified approach to managing copy data processes and workflows.

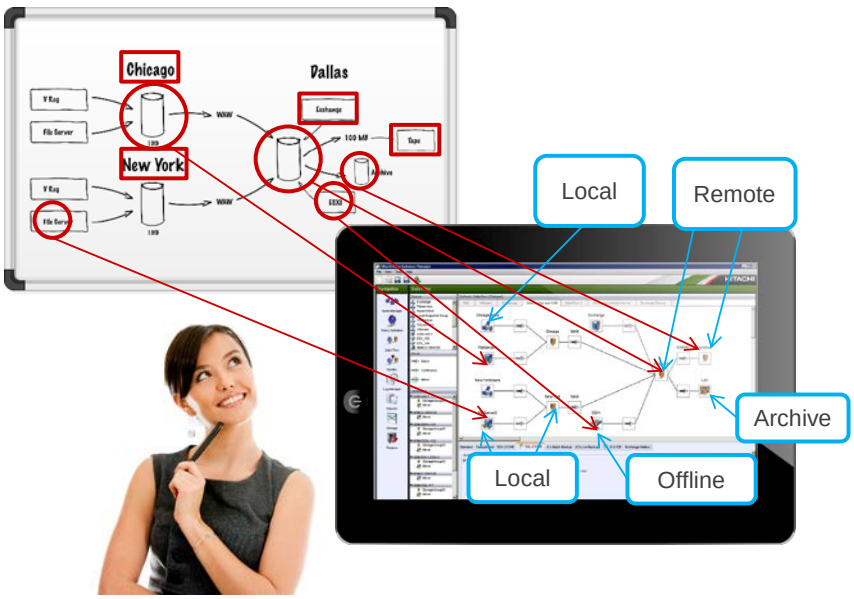

Data Protection Is Complicated

In the modern enterprise, the data protection (also known as backup and recovery) infrastructure has been cobbled together over time, adding features and point solutions as new needs arise. If this has resulted in a cost and management nightmare at your organization, it's time to perform a thorough review of your needs, and assess what's working and what isn't. It's time to design and adopt a simpler approach that covers what you need now and grows, both in terms of breadth and capacity, over time.

In the discussion below, we examine data protection complexity in the following areas:

- Infrastructure.
- Threats.
- Service-level objectives.
- Technology choices.

By the end of this discussion, you see that there may be a need for hundreds, if not thousands of individual policies and workflows to protect, retain and recover data in the ideal manner to sustain your business.

Infrastructure

In this section, we examine what it is you are trying to protect. See Figure 1. We start with applications and platforms. Each application is at least somewhat unique in the way it creates, processes and stores data. These differences lead to best practices for protecting that data. For example, a database application, such as SAP HANA or Oracle, needs to be put into a backup-ready state in order to capture all of the components of a current transaction. Otherwise, the backup application creates an inconsistent copy, which may crash when being restored. Each has its own interface for accomplishing this, and, therefore, the backup application needs a specialized agent or other means of communicating with each of them, such as an Application Programming Interface, or API. For example, SAP has BR*Tools; Oracle has RMAN; and Microsoft® Exchange and SQL Server® have Volume Shadow Copy Service (VSS).

Figure 1. Areas to Protect

Applications, Platforms

Operating Systems

Locations

Even standard file systems need special attention, to ensure that any open files do not get corrupted in the backup process.

Virtualization platforms such as VMware vSphere and Microsoft Hyper-V® create their own challenges. Do you protect the virtual machine (VM) from within using a standard application agent? Do you protect the entire VM as you would protect a file? Or, do you work through a vendor-specific interface, such as VMware's vSphere APIs for Data Protection (VADP), to take advantage of advanced functionality. The chosen path could have an impact on the backup infrastructure, performance and availability of the VM and its applications, and overall maintenance and administrative overhead.

Having different operating systems in the environment further complicates the deployment of a cohesive data protection solution. Not only are there different flavors of Linux, UNIX and Microsoft Windows® operating systems, but there are also different versions of each, and the ones supported by the backup software change over time. Installing the wrong version of a backup agent causes nothing but grief.

We also need to consider the location of the data that needs to be protected. These whereabouts tend to move in waves. Originally, it was all centralized with dumb terminals accessing the data from a mainframe or minicomputer system. Then the client-server model moved much of the data to the desktop computer. Large amounts of data started being created and stored in remote and branch

offices, and then in home offices. Mobile devices, virtualization technologies and faster Internet connections have ushered in the more recent move to virtual desktops.

The bottom line is that corporate data can now be found everywhere, including public cloud storage services, thus leading to copies of data everywhere. If the data is important to your business, you need to assess where it is, where the copies are, and how to properly and efficiently protect and retain it as prescribed by corporate and government policies.

Threats

There are many things that can happen that affect the availability and usability of your data. See Figure 2. Losing data can have a profound impact on your business, depending on the importance of the data, the volume of the data, and the time or money it takes to recreate the data (if that's even possible).

Figure 2. Challenges to Availability

Lost Files, Emails

System Failures

Site-Level Disaster

The most common data loss is a lost file, folder, email, database or other data object. The most common cause, by far, is human error, followed by software and hardware errors, malicious behavior and virus attacks, hackers and other types of data corruption. Recovering a single file may seem pretty easy, but it depends on when and how it was last backed up, and where the backup is stored, and how much data needs to be sifted through to find the right file, or the right version of the right file. The last thing you want to do is restore an entire system, which can take many hours or days, just to retrieve a single file.

System-level failures are the next most common cause of data loss, although this has become much more rare as hardware designs have become more reliable and resilient to failure. But even in systems that are designed to be highly available, bad things do happen, again most frequently because of human error. A case in point was the massive loss suffered by 27 agencies of the Commonwealth of Virginia in 2010, when a technician pulled the wrong redundant memory board and corrupted not only the local array, but also the mirrored disaster recovery array. The process for restoring an entire system as quickly as possible is different from restoring a single file, so you need to account for that in your plans.

The least common data loss incidents, the major site-level or regional disasters, are the ones that seem to get the most attention. While the odds are low that your company will ever face one, there does seem to be a report almost every week of a major flood, earthquake, hurricane, power grid failure or man-made disaster somewhere in the world. If your business cannot tolerate a lengthy rebuilding process, or damaged reputation, you want to have a disaster recovery plan in place. And, you want that plan to allow you to restart operations in, or from, another location that is outside the potential disaster zone. Many organizations opt for a 3 data center topology to assure both fast failover following a site-level event and resiliency from major regional events.

Figure 3. Infrastructure Threats

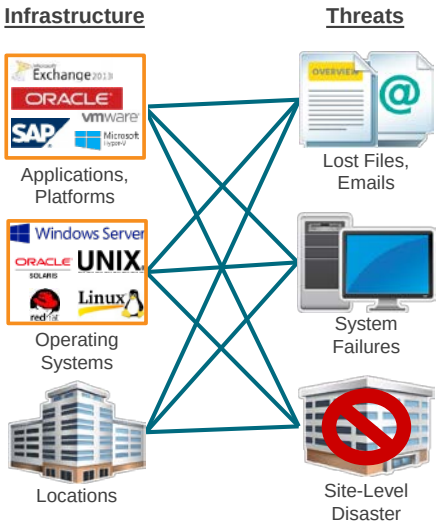

Now, consider that for each application, operating system and location that we discussed in the last section, you need to define, test and implement a solution for each of the different kinds of threats that you might face. See Figure 3. Add up the data instances and multiply by the number of threats, and you've already got a large number of possibilities to address with your data protection, retention and recovery plans.

Don't worry, we're not done. It gets worse.

Service-Level Objectives

All data does not have the same value or importance to the organization. For example, losing the order entry database would be far more costly than losing a purchase order stored on a sales rep's laptop. Losing either would cause pain and monetary loss, but the former may put you out of business.

One way to look at this would be the pyramid in Figure 4, dividing the data set into critical, important and standard tiers.

Figure 4. Data Tiers

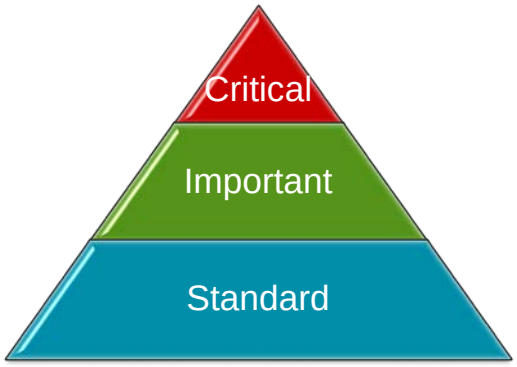

- The critical data, such as the design files for your key product, is often the smallest section by volume, but it would cause the most harm if lost.
- Important data would be very painful to lose, but you would probably recover, at least to some extent.
- Losing standard data, such as that found on an employee's workstation, would cause a limited loss of productivity while the affected individuals get back up and running, but the impact to the business would be small.

Figure 5. Factors That Determine Data Protection Needs


Backup Window


RPO/RTO


Retention


Budget

Several factors determine data protection service levels. See Figure 5. Each of these factors are specified to meet the backup, recovery and retention policies deemed necessary by the business, and they may differ greatly depending on the type and importance of the data.

The ***backup window*** is the maximum amount of time that a given backup or copy operation should take. Another way to look at it is, how much time is the organization willing to pause its access to the data while it is backed up. Traditionally, this may be several hours each night to perform an incremental backup, or longer periods on weekends to perform a full backup. But as businesses become more global and interconnected, this amount of downtime is not ideal for many applications. A modern approach, using snapshots, can eliminate the need for a backup window.

Recovery point objective (RPO) defines the frequency of the previous points in time from which data can be restored, and comes down to how often the backup is performed. A nightly backup results in a 24 hour RPO, meaning that up to 24 hours of your most recently created data is at risk of loss. That is often fine for standard data, but probably not for important and critical data.

The ***recovery time objective*** (RTO) is the amount of time in which a system, application or process must be restored following an outage. This measure could include the time to troubleshoot the problem, apply a fix, restart and test. It is very common to have a different RTO for each failure type, as described in the threats section above. For example, you may have an RTO of two days following a major disaster, 30 minutes to restore a single file or email, or less than a minute to restore a business-critical application.

Retention defines how long the copy of the data object needs to be stored, and in which manner. This service-level objective (SLO) can be applied for point-in-time recovery purposes, as in how long to keep a backup set, or for longer-term requirements which may be set by government regulations and corporate governance mandates. The other side of retention is ***expiration***, which specifies if, and when, to delete the copy data, and whether to do so in a manner that prevents its future discovery.

A final objective, which ties all the others together, is the available ***budget***. Since data protection is, in essence, an insurance policy that only adds value when something bad happens, the organization seeks to minimize its costs, across hardware, software, services and personnel resources. However, the service levels chosen for each type of data are a huge factor in determining these costs. Setting the minimum service level objectives carefully is very important to keeping costs down.

Adding the various service-level objectives to the infrastructure and threat variables takes the complexity to an entirely different level. See Figure 6.

Consider that for each instance of each application in each location, you should define policies for backup window, RPO and RTO requirements to recover from each of the various things that can go wrong. Also, define how long each data set should be retained and how it should be deleted.

Figure 6. Service-Level Objectives: More Requirements for Data Protection

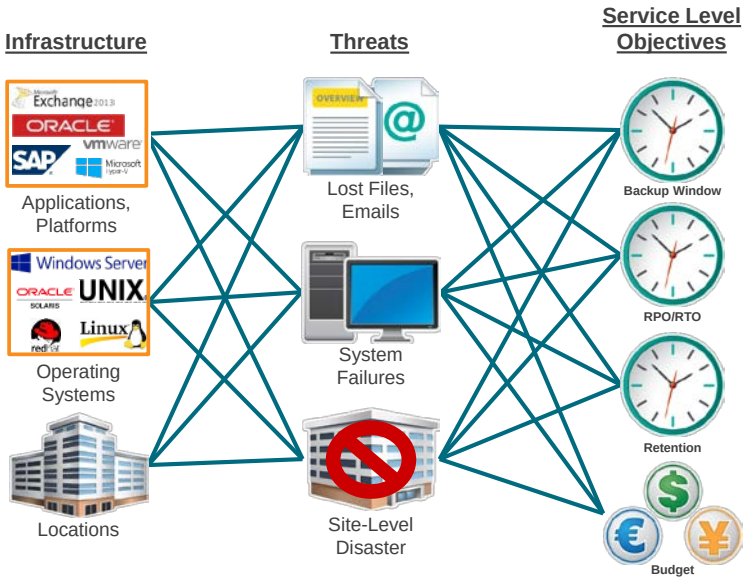

The policy definitions drive the budget requirements, and, therefore, the discussions on trade-offs to lower SLOs to meet resource availability.

Figure 7. Data Protection Solution Types

VSP G1000 = Hitachi Virtual Storage Platform G1000, HPP = Hitachi Protection Platform, HCP = Hitachi Content Platform

Once you have all of these requirements defined, you can start to look at which technology solutions are available to help you meet them.

Technology Choices

As information technology has evolved and made data protection more difficult and complex, the IT vendor community has responded by developing a range of technologies and processes that tend to address specific needs, but also add to the overall cost, complexity and risk. It all started with the classic backup process, which makes a point-in-time copy of your data and stores it in some defined location until it is either needed to perform a recovery, or discarded to make room for newer copies.

Over the past 20 years, traditional backup vendors have done plenty to help make their solutions more complex, by introducing new agents and licensing schemes and repository targets and software upgrades after each new data source has been introduced to the market. However, the larger vendors in this space (Veritas, IBM, EMC and HP) are typically pretty slow to roll out any new support, and this opens the door for specialized solutions to enter the market.

In addition to the traditional full + incremental batch backup process, below are just some of the available technologies to address specific infrastructure, threat or service level requirements.
See Figure 7.

- **Continuous data protection (CDP)** captures every change that is written to disk, either at the file or the block level. This provides near-zero backup windows and RPOs, but can also consume a lot of backup storage capacity. CDP is suitable for short-term recovery of highly critical data sources that have a low tolerance for data loss.
- **Snapshot** technologies capture the pointers to new data to create a point-in-time view of a data volume. The snapshot can be performed by host-based or storage-based software, and, like CDP, a snapshot can provide much faster and more frequent copies than traditional backup. Snapshots can also enable much faster restores to improve RTO. The downside of snapshots is that the captured data remains on the protected storage system, so they do not provide protection against system-level or site-level disasters.
- **Replication**, also known as mirroring, sends an exact copy of data from one storage system to another, either locally or across great distances, to provide failover and disaster recovery capabilities. The replication can occur synchronously (over shorter distances) or asynchronously (over longer distances), and can be performed by the storage system software, or by host-based data protection software. The former typically requires the same hardware on both sides of the replication, while the latter can use heterogeneous storage. The limitations of replication include a lack of point-in-time recovery capabilities (if the source gets corrupted, so does the replica) and a lack of application awareness.
- **Archive or tiering**, moves the data from the production system to a secondary system for long-term retention, protection, discovery and access. Archiving is suitable for data that isn't accessed frequently but needs to be retained to meet regulatory or corporate compliance requirements. For example, financial records may need to be kept for seven years, medical records for 30 years, and product designs or test data for the supported life of the product. The requirements are different for every industry, with additional variances by country.
- **Cloud backup.** Back up or archive to cloud storage services, also known as backup as a service (BaaS) and archive as a service (AaaS), respectively, are among the latest data protection models to hit the market. They provide more flexibility and mobility, and may offset management costs. In the cloud model, instead of purchasing storage capacity to host your copy data, you rent the capacity from a third-party service, on a cost-per-capacity-per-month basis. This shifts high initial capital expense (capex) to lower ongoing operational expense (opex). In addition to pricing structures to account for varying service levels, you may also pay for additional copies of your data in case the service provider experiences a failure. This approach can save the organization a lot of money, but it brings in questions of performance, reliability, long-term viability and data security. It is essential to choose a trustworthy vendor for cloud-based services, such as Hitachi Cloud Services.

When everything is considered, we end up with potentially hundreds or even thousands of policies and procedures being executed by any number of individual tools, scripts and manual processes. See Figure 8. Do you end up with specialists for each of the tools, or for each platform, or for each recovery scenario? Can any one person really know and understand what's going on at any given time, and react appropriately when a disaster strikes?

Figure 8. More Complexity With Various Data Protection Technologies

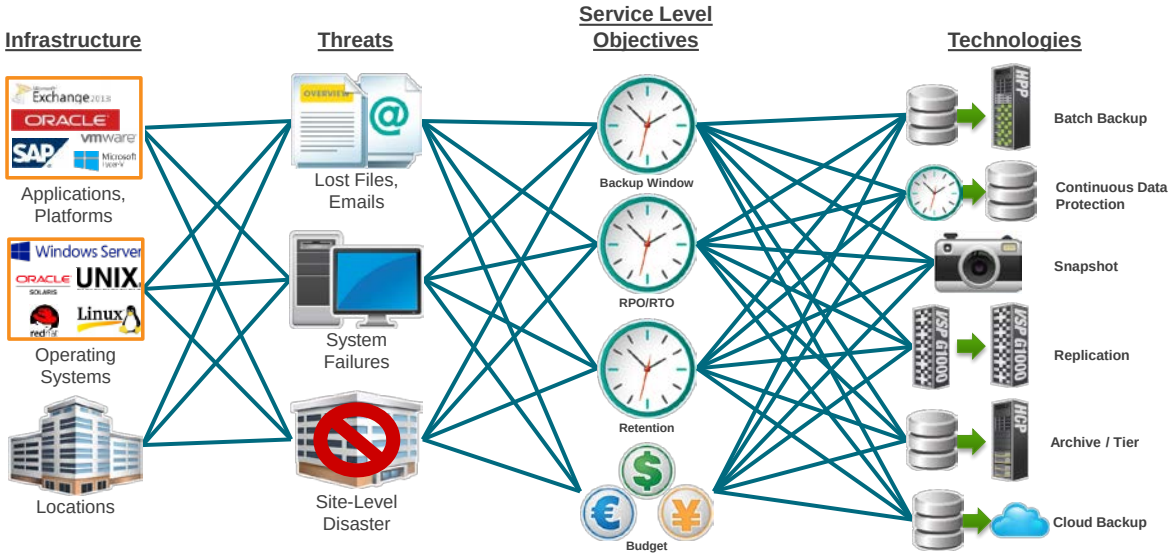

Unified Data Protection and Recovery Management

There is no single answer to all these challenges:

- Backup is too slow and cannot be run frequently enough for important and critical applications.
- CDP consumes more storage capacity than may be necessary, and by itself does not offer application consistency.
- Snapshots are stored on the same physical system as the source data, providing no protection from a full system failure or site-level disaster.
- Replication does not provide application consistency or recovery from a previous point in time.
- Archiving simply moves the source data; it does not, in itself, provide a recovery copy.
- Cloud-based backup and archive adoption is slowly becoming mainstream based on the points mentioned above as these services mature.

To satisfy all of the combinations of complex requirements, and to do it in a way that does not bankrupt the budget, you need some combination of these technology choices. The challenge then centers on how to simplify this chaos to limit the costs and the risks. Deploying a new point solution to address every new requirement is not the answer, at least over the long term.

Hitachi Data Systems believes that the answer is in unifying all of these disparate tools under a single management framework. Yes, use the right tool for each job, but manage them all from one place. And, where needed, have them work in tandem to support sophisticated policies and workflows.

HDS offers 3 data center protection, retention and recovery solutions that provide unified management: Hitachi Content Platform (HCP), Hitachi Data Instance Director (HDID) and Hitachi Data Protection Suite (HDPS), powered by Commvault. ,HCP is ideal for storing and self-protecting user data. HDID and HDPS enable you to take advantage of a range of data protection technologies, but simplify the administration of them, and orchestrate the movement of data between the technologies. Both offer policy-based backup, application-consistent snapshot management, replication of the backup repository, archiving and access to private and public cloud storage services.

Hitachi Content Platform

HCP is a highly scalable object storage system with a robust set of data protection and security capabilities. Data stored on HCP does not need to be externally backed up. See Figure 9. It is an ideal repository for unstructured data, such as user data that is composed of the typical files that we all use in our jobs: spreadsheets, presentations, documents and so forth. Typically, losing all such data impacts individual productivity but does not cause a devastating loss for the entire organization.

Figure 9. HCP Self-Protection Features

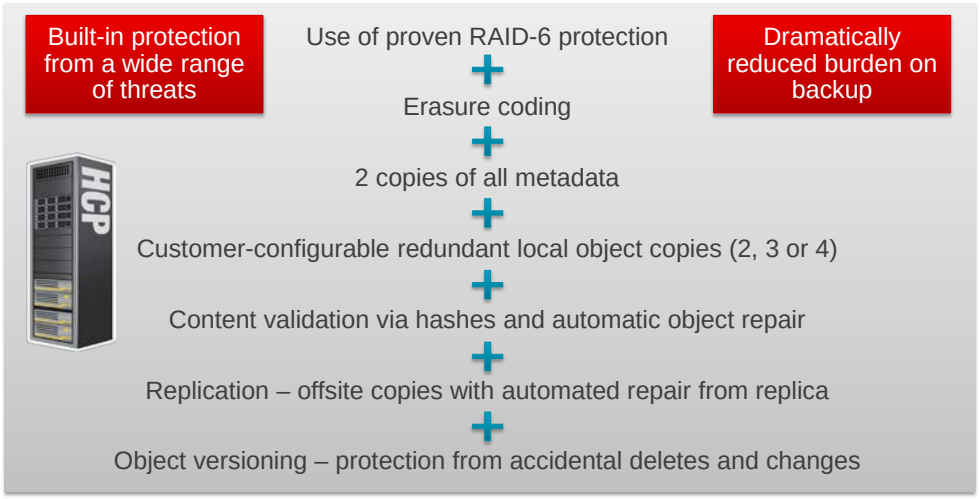

HCP can scale from 1 to 80 compute nodes, with each node supporting up to 5.7PB of storage. And it supports a number of data transfer protocols, including NFS, CIFS, REST, HTTP, HTTPS, WebDAV, SMTP and NDMP. For remote office users, data can be transferred using Hitachi Data Ingestor (HDI) and all users can benefit from the HCP-based private file, sync and share platform, Hitachi Content Platform Anywhere.

HCP is also supported by a large set of independent software vendor (ISV) partners, enabling a wide range of long-term retention and other use cases. See Figure 10.

Figure 10. HCP Ecosystem

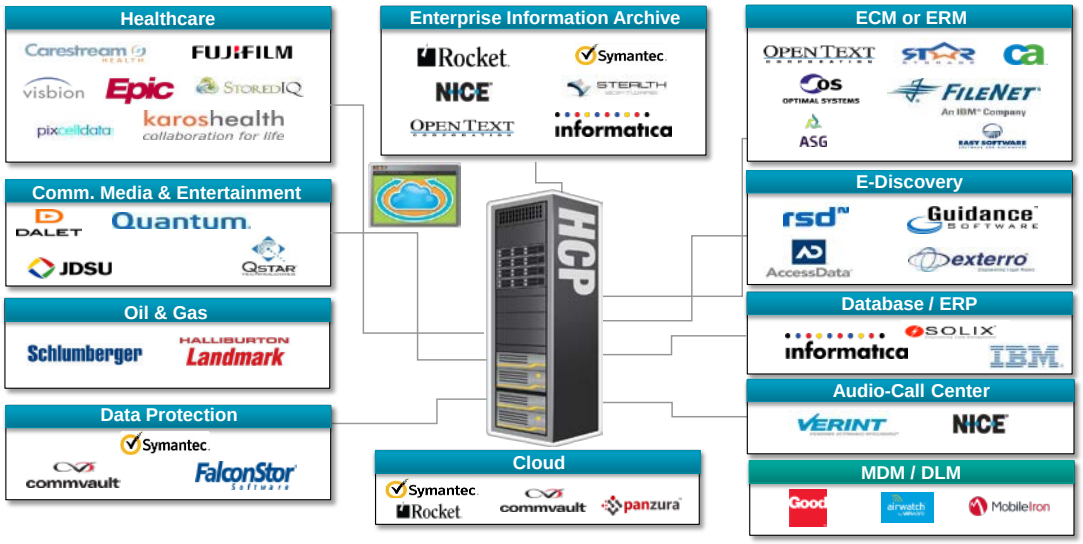

HCP = Hitachi Content Platform

Hitachi Data Instance Director

HDID represents a new paradigm in data protection software. It has been designed with the administrator in mind, using modular concepts that allow new functionality to be easily added and integrated with existing capabilities. An example of this is in HDID v5, where the orchestration and automation of storage-based snapshot, clone and replication management was added. This approach is exemplified by the unique, whiteboard-like user interface that makes creating and managing complex workflows a very simple drag-and-drop process. See Figure 11.

Figure 11. The HDID Whiteboard Interface

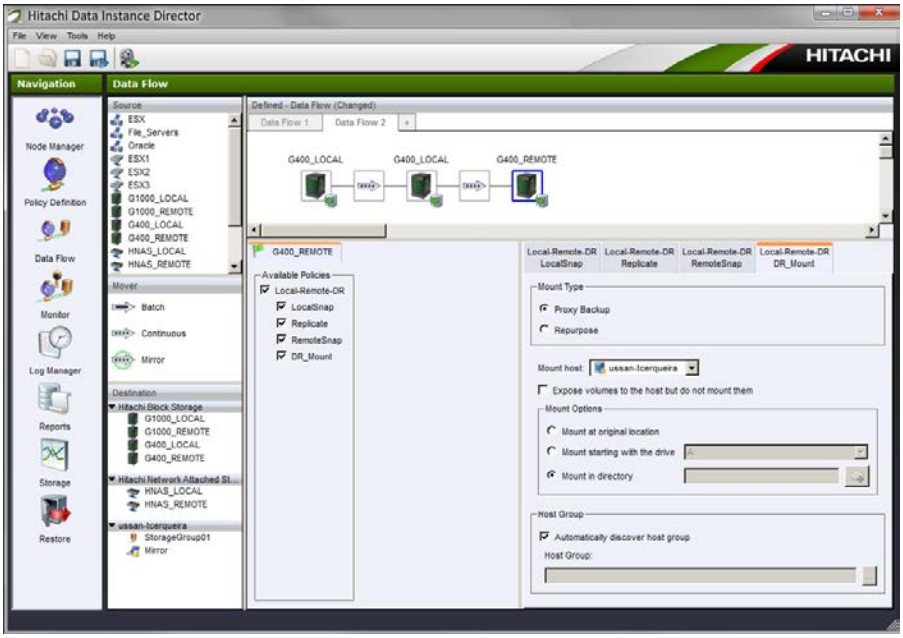

For example, a large health insurance provider estimated it would take them at least two days using their existing backup software to design, implement and test the required protection and recovery processes for a new health information system. The same task was completed in 10 minutes using HDID.

The process is strikingly simple. Create or edit an existing policy with a few clicks to define what to protect, when and how often. Drag that policy to the whiteboard, then drag a data mover (such as backup, CDP or snapshot), and then drag a repository. You can add sources with different policies to the same repository, or additional repositories with different retention characteristics.

Unlike other enterprise-class data management solutions, HDID is easily user-installable and offers very fast time to value. With its ability to quickly define and re-use policies and workflows, the larger the environment the more it can save in deployment time and expenses.

HDID is ideal for midsized-to-large environments, potentially protecting terabytes of data. HDID can provide CDP on Microsoft Windows systems, and batch backup of Windows and Linux systems. And it efficiently archives Windows file and email objects to Hitachi Content Platform and Microsoft Azure®.

The real strength of HDID is in automating and orchestrating application-consistent storage-based snapshots and clones to protect Microsoft Exchange and Microsoft SQL Server, as well as Oracle and SAP HANA databases. Other application environments can also be protected using these modern methods through pre- and post-process scripts. HDID also automates storage-based remote replication in 2- or 3-data center topologies.

Hitachi Data Protection Suite

HDPS delivers a broad range of data protection, retention and recovery capabilities, across most enterprise platforms and applications, across locations and at any scale, from a few terabytes to many petabytes. It supports management through a single console. You can protect complex, heterogeneous environments, physical and virtual servers, remote offices and even information on laptops and desktops with one solution. See Figure 12.

Figure 12. HDPS Data Protection, Retention and Recovery Capabilities

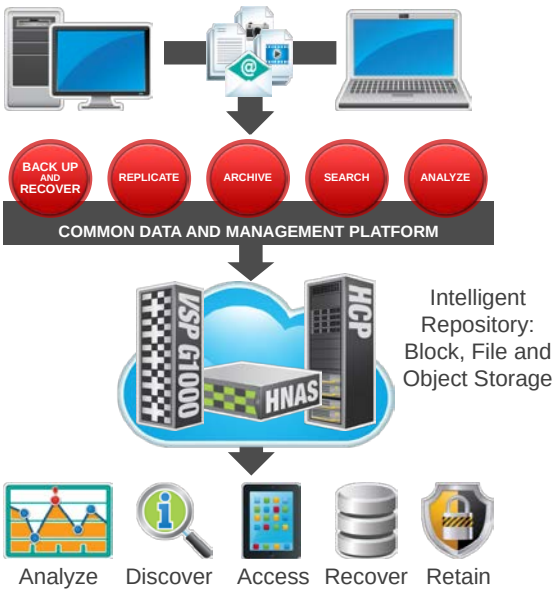

VSP G1000 = Hitachi Virtual Storage Platform G1000, HNAS = Hitachi NAS Platform, HCP = Hitachi Content Platform

All of the protected data becomes an asset by securely opening it up for self-service access through a single searchable index. You can maximize efficiency by leveraging whatever storage best meets your needs for performance, reliability and cost, including disk arrays, virtual tape, object storage, tape and the cloud.

HDPS is ideally suited for large, complex, distributed and heterogeneous environments. Its capacity-based licensing model allows you to deploy the right tools for each job without worrying about purchasing the correct modules. Or, you can select individual solution sets to get started, without replacing the legacy software where it is still meeting your needs, and then add additional functionality later, as needed. These specialized offerings provide:

- Basic backup and recovery.
- Intelligent snapshot and clone management.
- Email protection and archiving.
- Protection of virtual and cloud computing environments.

Table 1 shows the capabilities of Hitachi Data Instance Director and Hitachi Data Protection Suite. However, these are high-level yes/no evaluations. Consult your HDS or Hitachi TrueNorth Partner representative for further details.

Table 1. Comparison of HDID and HDPS

Operational Recovery	HDID	HDPS
Batch backup (full, incremental):		
■ **Microsoft® Windows® and other Microsoft apps**	✓	✓
■ **Linux**	✓	✓
■ **UNIX**	IBM® AIX®	✓
■ **Incremental-forever**	✓	VMs and synthetic full
Continuous data protection (Windows)	✓	
Application-aware hardware snapshot, clone:		
■ **Hitachi Virtual Storage Platform (VSP), Hitachi Unified Storage VM (HUS VM) (Hitachi Thin Image, Hitachi ShadowImage)**	✓	✓
■ **Hitachi NAS Platform (HNAS) (Hitachi NAS File Clone, and directory clone feature)**	✓	
■ **Other storage vendors**		✓
Disaster Recovery	**HDID**	**HDPS**
Hardware Replication:		
■ **VSP or HUS VM (Hitachi TrueCopy, Hitachi Universal Replicator)**	✓	
■ **HNAS**	✓	
Initiate remote hardware snapshots	✓	
Software replication of repository	✓	✓
Bare Machine Recovery (Windows)	✓	✓
Long-Term Recovery	**HDID**	**HDPS**
Archive (move)	✓	✓
Tier (stub)	✓	✓
Single ingest (backup → archive)		✓
Hitachi Content Platform integration	✓	✓
Archive to cloud	Microsoft Azure™	Multiple
Index and search (e-discovery)	✓	✓

Summary

Basic data protection can be considered little more than an insurance policy. Like an automobile or homeowners policy, it is a cost to the organization that doesn't add any value at the top line. It does, however, offer tremendous value to the bottom line when something bad happens. You don't know what that bad thing might be, so you have to protect against a wide range of threats. Each data asset, application or system requires slightly different technologies and processes to protect it against each different threat.

It has been understandable that as each new asset and each new threat has come into your environment, the natural tendency has been to find and deploy a new solution to that challenge. However, that has now led to most organizations implementing multiple tools under the umbrella of data protection, retention and recovery. These tools don't work together, and they usually aren't managed by the same people.

The analyst firm Enterprise Strategy Group recently conducted a survey on this subject and found that only 31% of respondents are using a single data protection tool. We believe that number is very high, especially in larger, more diverse enterprises. Even if the tools are from the same vendor, it is likely that there are multiple tools for different applications, virtualized servers, remote offices, workstations and more, as well as different tools for operational recovery, disaster recovery and long-term retention.

Obviously, there is a need to simplify the data protection infrastructure, just as there has been a need to simplify the compute, storage and networking infrastructures. Hitachi Data Systems is an expert in all of these areas, providing solutions and services that drive down complexity, costs and risks, and quickly deliver positive business outcomes.

To learn more about HDS data protection solutions, please visit hds.com/go/protect, or contact our expert team 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-528-B Hitachi Data Systems July 2016