



Hitachi Data Systems File and Content Solutions Integrate Services for Better Efficiencies

The Economics of Hitachi Data Systems File and
Content Solutions

By Hitachi Data Systems

July 2011

Table of Contents

Executive Summary	3
Introduction	4
File Services for Unstructured Data	4
An Integrated Approach for File, Block and Content	5
The Role of Hitachi Data Systems Storage Economics	6
What Are the Tools for Change?	10
Hitachi NAS Platform	11
Hitachi Content Platform	14
Hitachi Data Ingestor	17
Hitachi Data Discovery Suite	19
Hitachi Data Discovery for Microsoft SharePoint	21
Work Together to Reduce Cost	22
A Single Integration Platform	22
The Finer Points of Cost Efficiency	23
Conclusion	24
Appendix A: Case Sidebars	26
Hitachi NAS Platforms Host Media Delivery Network for WebTV and Internet Provider	26
Hitachi Content Platform Facilitates Turnkey Operations without Backup	26
Appendix B: Cost Family Reduction Matrix for Hitachi Data Systems File and Content Solutions	27

Executive Summary

As the amount and variability of unstructured data make their collective mark on an organization's ability to effectively manage its digital content, so do Hitachi Data Systems file and content solutions. With a keen emphasis on helping customers understand how to integrate file and content services for better efficiencies, Hitachi Data Systems delivers a cohesive approach to getting in front of costs, complexities and risks often associated with unstructured data. The file and content solutions portfolio includes Hitachi NAS Platform, powered by BlueArc®, as well as Hitachi Content Platform, Hitachi Data Discovery Suite, Hitachi Data Ingestor and Hitachi Discovery Suite for Microsoft SharePoint.

In this paper, Hitachi Data Systems Storage Economics first examines the cost-reduction benefits of each of these products. Then it looks at the cost-reduction benefits of seamlessly integrating them into one unified platform that can be managed through a single interface. Storage Economics attempts to uncover how the file and content solutions contribute to savings in both capital expenditures (CAPEX) and operational expenditures (OPEX), ultimately creating a lower total cost of ownership (TCO) for the modern data center.

Introduction

One of the storage industry's most pressing data growth issues is how to cost-effectively manage unstructured data. Digital images, contracts, video, presentations, web pages and many other content types so prevalent across the modern business organization are common unstructured data types. The rapid growth rates for unstructured data are unprecedented and unpredictable. For many organizations, being able to gain control over the spiraling costs, complexities and risks associated with this data has reached critical mass.

IT administrators continuously look for ways to manage the persistent growth of unstructured data with their already strained resources, limited staff and aging equipment. They are trying to figure out how to maintain a sense of coherency while scaling data environments to house hundreds of millions of documents or objects and ensure accessibility for indefinite lengths of time. These objects can range from highly accessed, active content to intermittently retrieved records to mostly archived, inactive data. Because the importance of unstructured data to the organization fluctuates and can quickly change, being able to cost-effectively store content on the appropriate level of disk may become problematic and expensive. Administrators must joust with accelerated storage consumption and poor utilization of assets, and how best to manage tiered storage, all of which are collectively draining available capacity and budgets. Unintended storage architectures are getting built up over time to accommodate all the unexpected data growth, leading to more complicated and disparate systems and tools to manage. Floor space, power and cooling costs continue to increase. And as distributed IT environments, such as cloud deployments, ratchet up and data retention legalities march on, so must an organization's abilities to separate active from stale content, and rapidly archive and retrieve data. Failure to locate information can result in penalties or fines, again racking up the costs of managing this storage.

File Services for Unstructured Data

With so many technologies and products promising better cost efficiencies across the data center, it's hard to know what to try or what will work. The Enterprise Strategy Group cites that the need for file-based services, i.e. storage infrastructure capable of supporting unstructured data, is growing at a compound annual rate of 79%.¹ Yet, the challenge of managing costs and complexities of file-based data growth is one of the top 5 most difficult issues facing Global 5000 companies today.² While file-based technologies such as file servers and network attached storage (NAS) solutions can work wonders for unstructured data, they have spawned other issues around management, backup and resilient scalability.

One of the biggest kinks with file-based storage, as well as block-level storage, is that it still more often than not lacks awareness of the content within the data that is being stored. What is this and why is it important? Content awareness describes a set of capabilities that dynamically classify information and assign policies to the data files; it creates an understanding of the actual content while the data is at rest (in storage), in use (during operations) and in motion (moving across networks).

Content-aware systems allow organizations to turn unstructured data stores into valuable business

¹ Source: <http://www.hds.com/assets/pdf/hitachi-lab-validation-report.pdf>

² Source: <http://www.busmanagement.com/article/The-Challenge-of-File-based-Data-Storage/>

intelligence to help enforce best practices, make decisions faster and protect sensitive company data. Depending on the nature of the content, data files could be of very high value to competitors or criminals; they could be of critical importance to e-discovery or legal activities; or, they might simply be old but retention-worthy. By not having the proper tools in place to dynamically assess data, organizations are missing valuable insights and could be spending too much to manage it. For these reasons, content awareness has become an imperative attribute for sifting through unstructured data as cost-efficiently as possible to create a "total customer voice."

Storage vendors have made strides to improve the content awareness of file-based services. By adding structure to data through a litany of mining and analytic methods, some software and hardware solutions and industry protocols have been able to classify, index and authenticate content. The shortcoming here tends to be that the solutions are not entirely integrated. Subsequently, data centers end up with a proliferation of NAS and file servers, manual integration requirements and the management of silos and separate interfaces. Plus, they face the added expense and difficulty of managing them.

Instead, IT organizations are calling for file-based systems that provide a more native and automated awareness about the data being housed. Legal, regulatory and business stipulations are driving the need for a true assimilation of file and content services, a more savvy recognition of the relevance of the data stored. Organizations need to be able to perform federated search: single-query searches across multiple data resources using a single interface. Federated search allows management of content-triggered migrations and other integral data activities.

The focus is on more advanced solutions capable of delivering integrated block, file and content services. These services support policy-based management, migration, virtualization, archiving, thin provisioning, storage clustering and dynamic file tiering, to name a few. As IT budgets and administrative resources succumb to staggering demands for storing and managing unstructured data, the necessity for cost-effective, integrated, file-and-content-aware storage has moved front and center.

An Integrated Approach for File, Block and Content

So how do we sort out what is important in considering file and content solutions? The Hitachi Data Systems strategy manages file-based data by predefining storage tiers, indexing the data and providing policies to automatically migrate the data to the appropriate tier. This strategy is about making file and content smart, efficient and scalable. It starts with a highly virtualized, tightly coupled and dynamically tiered storage architecture; then it applies file intelligence that allows administrators to manage content to the needs of the business. By directing files to the right platform at the right time, organizations begin to emerge with greater efficiencies and agility across the enterprise. It's about integrating compliance and federated discovery across a heterogeneous environment, reducing storage waste, maximizing system performance and being able to scale with growth demands. By increasing utilization and flexibility, administrators can improve data protection, increase the amount of data manageable per person and, ultimately, control storage transparently to deepen productivity and business value.

The Hitachi Data Systems file and content solutions strategy comprises five key areas:

- **Consolidation.** By creating larger pools of storage and leveraging virtualization across platforms, organizations gain better utilization and can reinvigorate or extend the useful life of existing assets.
- **Tiering.** After creating that common pool of storage, organizations can better leverage existing assets to enable cost-effective data copies and match price to performance. By dynamically tiering active data to the right platform, organizations tend to right-size disk drive requirements, too [i.e., more data on SATA, less on Fibre Channel, SAS or solid state disk (SSD)].
- **Reclaiming.** By intelligently tiering data to the appropriate storage, IT is able to corral previously allocated but unused disk space, leverage data dispersion technologies and facilitate consistently better performance.
- **Archival.** At this point, organizations are able to nondisruptively move data between active storage and archives using policy-based automation in SAN and NAS environments. By dynamically moving inactive data to a content repository with replication capabilities, organizations will need to store less NAS and file server data on primary storage and, consequently, have less to back up.
- **Backup.** For every terabyte or petabyte of content, organizations typically will have 5 to 10 times that amount in backup. By using policy-based movement of aged data to cost-effective, longer term mediums, organizations will end up with a smaller volume of active data to back up and less of a need to back up inactive data. Down go backup costs as well.

This strategy, in turn, reshapes future growth of data. By centrally pooling storage and scaling intelligence across systems, independent of capacity and location, there are fewer points of management and data is directed by policy-based automation. With a smaller amount of active data on primary storage, such as Fibre Channel and SAS, and a greater amount of data shifting to economically appropriate tiers such as SATA, organizations begin ticking off the savings. Costs decline for administration, backup, capacity planning, power and cooling consumption, resulting in a general footprint economy.

Hitachi Data Systems file and content solutions are designed to help organizations quickly get control of unstructured data environments with unified, scalable management platforms that consolidate, simplify, archive, protect and search. By leveraging content services at the storage infrastructure level, file-based services can better address performance, costs, content awareness, transparency, chain of custody, discovery and more.

The benefits of integrating file and content services infrastructure include more supple movement and management of unstructured data, and, consequently, operational efficiencies, cost reductions and risk mitigation. To better understand the value of having integrated file and content services in the data center, Hitachi Data Systems recommends taking a look at the true costs associated with their infrastructure investments and operations. Guiding organizations through this labyrinth is Hitachi Data Systems Storage Economics.

The Role of Hitachi Data Systems Storage Economics

Getting in front of storage efficiency challenges means having the right data on the right storage device at the right time. And it also means knowing the value of that efficiency. Hitachi Data Systems Storage Economics has an extensive and successful record of helping IT administrators identify and capitalize on the long-term value of storage purchases by first defining a framework for assessing the true TCO. Storage Economics uses financial metrics and calculable techniques to illuminate cost efficiencies and determine the accurate hard and soft costs associated with storage decisions. Ultimately, Storage

Economics aids IT leaders with proven ways to reduce TCO and improve the return on assets (ROA).

Storage Economics helps organizations to swiftly zoom in on cost efficiencies and minimize the unpredictability of managing unstructured data assets. For organizations striving to manage budget and resource constraints amid growing, morphing IT requirements, Storage Economics provides invaluable business acumen. Assessing the efficiencies, the technologies and integrative capabilities of Hitachi Data Systems file and content solutions is relevant to understanding how they can meet high-priority business and IT drivers, such as reducing costs.

Some of the key metrics that Storage Economics will use to evaluate the cost efficiencies of file and content solutions include:

- **Return on Investment (ROI).** This is a method for calculating the benefits of a particular investment. ROI is effective when challenging the status quo with a proposition to replace an existing solution by analyzing financial pros and cons of a purchasing decision. The ROI analysis would illustrate how much is to be invested, how quickly the investment is to be recouped and the amount of net savings to be expected.
- **Return on Assets (ROA).** This is a method for calculating how profitable an organization is relative to its total assets, or how an individual asset impacts profitability. ROA can be a key metric to justify investments that improve aggregate utilization of IT assets in general, and storage specifically, even beyond the depreciation life of those assets.
- **Capital Expenditure (CAPEX).** This is the cost of purchasing or extending the useful life of a fixed or physical asset, such as equipment or property.
- **Operating Expenditure (OPEX).** This covers the ongoing expenses for managing, supporting, maintaining and upgrading a system over its lifespan. For example, OPEX includes costs of electricity to power the system and administrative labor and related day-to-day costs incurred by the system to support the business.
- **Total Cost of Ownership (TCO).** This is a method for calculating all costs that will be incurred over the asset's useful life. TCO analysis is used when expansion is anticipated and the financial benefits of two or more proposed solutions must be assessed. The analysis places the total lifetime operating and purchasing costs of the assets side by side for comparison.

To ultimately calculate these economic metrics, Storage Economics first examines and considers 34 various cost elements of storage TCO. These are not necessarily equal in weight, importance or time relevance, but they are important for better understanding the true cost of storage infrastructure. IT planners usually select between 6 and 10 of these hard and soft costs to create their own TCO baseline, based on business needs. The critical value of citing cost elements is to help the IT organization develop and implement a plan for reducing TCO.

Storage TCO cost modeling is likely to be inclusive of the entire storage infrastructure, not just storage devices or systems only. So many of the cost elements can be derived from backup, SAN, remote and local circuits, tape systems, and management or monitoring consoles. In fully probing the value of the Hitachi Data Systems file and content solutions portfolio, Storage Economics methodology will list applicable costs from the general list of 34 that are addressed and reduced by each product. Storage Economics provides the following descriptions for costs relevant to the Hitachi Data Systems file and content solutions portfolio:

Procurement Impact

- Capital costs, lease expense for hardware and software
 - Storage hardware (and software), when depreciated, incurs a capital cost at the time of purchase. Each year after (during the depreciation term), the depreciation expense continues to be recognized on "the books" or the accounting system. Capital costs are up front, in that the payment is made one time, but the annual depreciation costs are still a real factor for IT accounting; this is true even if the cost is not seen as a real expense in that year.
- Cost of waste
 - The cost of waste is a sub-element of capitalization expense or lease expense, but it is the portion of the total asset that is not being used. The reasons for underutilization vary, but the fact is that poor utilization exists and needs to be addressed. Sources include:
 - Usable but un-allocated capacity (waste or reserve)
 - Allocated but unwritten to capacity (white space)
 This can be due to:
 - Underlying architecture
 - Multiple systems, tiered islands and segmented IT solutions
 - Increased waste equates to lower ROA
 - Cost of waste (as a percentage of the capital expense) can be 40% to 50% or higher.
- Cost of duplicate data
 - Duplicate data cost is different from the cost of copies in that an individual piece of data may be propagated or copied within a file system, whereas volume copies are done for production databases and applications. This category focuses on data files, and it is a more granular and problematic part of data growth.
- Cost of growth
 - Most storage infrastructures are in a constant state of capacity growth. Data is not static. Growth is not zero. Therefore, a key cost component requires the inclusion to grow, or show the growth (at the margin) for the next period of time.
 - Some architecture is inexpensive to initially purchase, but is not so flexible in growth or cost effective in the fluid nature of growth.
 - This cost needs to be correlated with the cost of waste. Having 20% to 40% reserve on-hand does not constitute an effective growth program. Reducing the cost of waste and the cost of growth are mutually inclusive strategies.

Management, Operations, Lifecycle

- Storage management labor
 - Labor to manage, engineer and support the storage infrastructure can often be one of the largest cost components over time. It is common to see labor cost make up roughly 20% to 30% of storage TCO.

Environment

- Power consumption and cooling
 - Powering and cooling storage infrastructure components take energy, and energy costs money.
 - For most of the world, the power and cooling costs over 4 years will be as much or more than the original purchase price of the storage asset.
 - On a cost/terabyte/year basis, power and cooling tends to represent 15% to 30% of the storage TCO.
- Data center floor space
 - Floor space costs tend to be 4% to 8% of the storage TCO.

Business Impact

- Cost of performance
 - The cost of good or poor performance can be an attribute for most storage architectures. Including the performance metric in the cost modeling can help justify additional features and capabilities that help deliver the needed performance levels. Examples of storage features that assist with high-performance computing and storage include SSD, RAID mirroring, high-speed Fibre Channel or Ethernet, controller cache, etc.
 - Performance can be more than a qualitative feature or function. Performance can often be converted to revenue or business loss.
 - For some, increased performance equals increased revenue or increased profits, where time is money.

General Risk

- Data loss
 - Previous cost categories have characterized outage risk, outage time and missing opportunities to recover data and systems to a full operational level in the case of a significant event. This cost area deals with the impact when data is lost or corrupted (beyond use) and is not recoverable.
 - Data loss can be the results of numerous events. The root cause can include poor backup coverage, sabotage, cyber attack, double disk failure, etc.
- Litigation, discovery risk
 - When data is backed up or archived for long-term retention, some businesses will depend on access to this older data for discovery purposes. The time it takes to locate and recover this data can present a risk cost of risk, depending on the business.
 - This cost has at its core litigation discovery risk. The length of time required to find and compile data needed for a lawsuit or court filing time can have a direct impact on the outcome and legal implications of a court case.
- Noncompliance risk (archive, data retention)
 - Noncompliance can incur fines, negative publicity and criminal prosecution.

- Banking and Finance
 - Sarbanes-Oxley
 - Basel, Basel II
- Healthcare
 - Health Insurance Portability and Accountability Act (HIPAA)
 - International Classification of Diseases (ICD10)

Data Protection

- Backup infrastructure
 - The infrastructure required to protect data with backups can be large and complex. Typically, it will include several hardware components:
 - Data servers, media servers
 - Tape libraries, robotics
 - Virtual tape library disk
 - It is common to see backup infrastructure occupying 15% to 20% of storage TCO.
- Backup media
 - The cost of backup media is constantly changing with the change in media density format.
 - It is necessary to separate media costs from the backup infrastructure costs because:
 - Different media (paper, fiche, tape, CD, disk) is used for long-term data retention.
 - Media costs are directly proportional to the data capacity stores.
 - Backup media costs may also include transportation and offsite media storage, which are common in many backup architectures. It is important not to double-count these costs.
- Security, encryption
 - Security has costs, and can be applied to data (active or at rest) while moving throughout the infrastructure.
 - Authentication and real-time protection is within the infrastructure.

What Are the Tools for Change?

Albert Einstein is credited with a quote about how doing the same thing over and over again but expecting different results is insanity. IT administrators understand this sentiment very well with the grueling back-and-forth between storage needs and storage resources. To get in front of the storage lunacy requires not only a transformation in the tools but also in the thinking.

Hitachi Data Systems file and content solutions help launch an important renovation of today's data center into dynamic information centers where access to blocks, files and content is seamless and exists in a more fluid and virtualized environment. The Hitachi technologies behind file and content services help organizations make better use of existing assets and reduce storage spend, while controlling unstructured data terrain.

The key to exploiting these technologies is through a unified approach that enables integrated file, block and content, intelligent file tiering, content-aware searches and reduced backup efforts. When

organizations are able to perform sweeping consolidations and pooling of errant storage silos, and consequently have centralized functionality to dynamically move data, the results are often dramatic. Hitachi Data Systems file and content solutions are intended to do just that. Whether data centers need to handle high-volume transactional operations or enterprise-class application demands, these file and content solutions cohesively support Oracle and VMware, as well as Microsoft Exchange, SQL Server and SharePoint environments for shaving storage-related expenses and challenges.

The mainstay business benefits of Hitachi Data Systems file and content solutions include:

- Lowered storage cost
- Simplified management
- Increased ROA
- Storage optimization and utilization
- Risk mitigation

Realizing these benefits help to amplify:

- CAPEX and OPEX savings
- Better business intelligence
- Compliance and legal discovery

A Note about Virtualization

As IT professionals consider options for synthesizing unstructured data resources, it is important to illuminate the role that virtualization plays in integrated file and content services. Virtualization comes in many forms. Storage systems might have built-in virtualization (internal) that enables tiering inside the box for dynamic volume movement and less disruption to applications. External virtualization allows tiering outside the box to dynamically move volumes across heterogeneous systems connected to that box. LUN virtualization is commonly known as thin provisioning, which is a method for creating a large number of volumes of varying sizes that can each draw from the same pool of capacity. Storage system resource virtualization enables file, block and content services to use common storage resources.

When storage services are not virtualized, it typically implies that servers and storage operate in silos that must be management individually. Nonvirtualized systems cannot share capacity, and data movement activities between storage frames may require disruption to applications.

Virtualization wields its clout across the data infrastructure to supply common management and common storage services so that storage is no longer stranded but pooled for heterogeneous storage systems. Using virtualization as one of the ways to manipulate unstructured data helps support across-the-board savings of time, resources and costs, as well as providing compliance assurances.

Hitachi Data Systems file and content solutions capitalize on virtualization technologies. Let's examine these file and content tools for change and how they work together.

Hitachi NAS Platform

Hitachi NAS Platform is a high-performance integrated network attached storage (NAS) family

of systems. This platform manages active data, file sharing, enterprise-class consolidation, data protection and general-purpose NAS workloads using a single management interface. Its advanced hardware-accelerated architecture provides organizations with new levels of storage performance, scalability, reliability, productivity and value.

For demanding data center environments, Hitachi NAS Platform models 3100 and 3200 deliver accelerated productivity gains and processing cycle times while minimizing the time to access files. Ideal for large enterprise data sets and high-performance needs, Hitachi NAS Platform models 3080 and 3090 create a no-nonsense, pay-as-you-grow approach to allow organizations to easily scale and manage unified NAS and SAN environments. All NAS Platforms integrate with Hitachi Command Suite and Hitachi Data Discovery Suite, and work with VMware, Microsoft and Oracle environments.

Key Features and Differentiators

- **Native Intelligent File Tiering** facilitates policy-based migration of data to an appropriate tier of storage, known as hierarchical storage management. It also supports dynamic read caching, to optimize performance and cost of storing data on Fibre Channel, SAS, SATA or archive media. Data can be promoted or demoted easily through the policy manager menu.
- **Data Scalability and Performance** deliver up to 16PB of usable storage capacity, up to 1600MB/sec throughput, and up to 200,000 IOPS per node to reduce processing cycle times of sequential workloads. These features also ensure active-active clustering of up to 8 nodes, with near-linear performance scalability to scale workloads and capacity and to consolidate and simplify the IT infrastructure with less file system overhead.
- **Cluster Namespace** grows data storage within a single namespace up to 2PB for unified directory structure across file systems, global data access and more efficient management of physical storage pools.
- **Advanced Virtualization Support** provides up to 64 virtual servers and role-based security, supporting high-performance storage for virtual machine environments and computing applications, dynamic storage pool expansion and thin provisioning capabilities.
- **Built-in Data Protection** ensures nondisruptive rolling upgrades and faster failover to minimize downtime; it also supports block-and-file remote replication and high-speed pointer-based snapshots to significantly increase agility and utilization across NAS environments.
- **Advanced Capacity Utilization** with dynamic provisioning allows administrators to grow or shrink capacity on the fly. It supports file cloning (BlueArc JetClone) to make duplicate snapshots of data stores without making physical copies, yet allow writing to the clone. It also supports object-based replication (BlueArc JetMirror) to provide a high-speed replication capability over WANs to efficiently and prudently save capacity and cost.

Mapping Product Benefits to Categories of Cost Reduction

The Hitachi NAS Platform product benefits can be mapped to Storage Economics cost categories, as shown in Table 1.

TABLE 1. HITACHI NAS PLATFORM — STORAGE ECONOMICS

Product Benefit	Storage Economics Costs Reduced	How?
NAS consolidation enabled through hardware acceleration	<ul style="list-style-type: none"> • Capital costs, lease expense for hardware and software • Cost of performance • Cost of waste • Cost of growth • Storage management labor • Power consumption and cooling • Data center floor space 	<p>Consolidation of existing file servers and NAS devices onto fewer nodes allows organizations to do the same or more work with fewer devices and less overhead. This reduces costs related to administration, management and file server proliferation.</p> <p>Hitachi NAS Platform uses less than half the nodes required by competitive solutions to reach critical performance and scalability levels, again reducing the quantity of devices required to meet storage requirements. The resulting reduction in storage capacity purchases (CAPEX) helps to free up or reduce required floor space costs and lower energy costs for power and cooling (OPEX).</p>
Higher scalability and capacity	<ul style="list-style-type: none"> • Capital costs, lease expense for hardware and software • Cost of waste • Cost of growth 	Hitachi NAS Platform scales significantly better than comparable models and, together with a more efficient file system structure, enables users to get more usable capacity (24% to 46% higher). With more storage on fewer devices, NAS Platform helps reclaim wasted space and overhead.
Built-in virtualization of third-party NAS and/or de-duplication devices	<ul style="list-style-type: none"> • Cost of waste • Cost of growth • Storage management labor • Power consumption and cooling • Data center floor space 	Virtualization improves storage utilization rates, allowing organizations to centralize and use more available storage capacity at the disk and system levels. Improved storage utilization helps delay capital spending and improves ROA.
Hitachi NAS Platform thin provisioning	<ul style="list-style-type: none"> • Cost of performance • Cost of waste • Cost of growth • Storage management labor 	Thin provisioning works to boost utilization rates for capacity, compounding the profitable ROA.
Native intelligent file tiering	<ul style="list-style-type: none"> • Cost of performance • Cost of waste • Cost of growth • Storage management labor 	Intelligent file tiering provides dynamic and automated data movement across tiers for greater administrative efficiencies in managing assets. Also, file tiering improves ROA by storing data in the appropriate place relative to its business value. For example, “stale” data is stored on archive or lower performance drives, while frequently accessed (active) data is stored on high-performance storage.
Multinode clustering, enabling scaling out or scaling down	<ul style="list-style-type: none"> • Capital costs, lease expense for hardware and software • Cost of performance • Cost of waste • Cost of growth 	The scaling features of Hitachi NAS Platform allow for efficient growth in all directions, whether from a capacity standpoint or from a performance standpoint. This technology reduces the need and cost typically needed to accommodate infrastructure growth.
Secure multitenancy, intelligent caching and policy-based automation	<ul style="list-style-type: none"> • Capital costs, lease expense for hardware and software • Cost of performance • Cost of waste • Cost of growth • Storage management labor 	Secure multitenancy, intelligent caching and policy-based automation provide the ability to deploy an efficient, pay-per-use cloud or chargeback model and to reduce management steps involved in keeping data on the right media at the right time.

Improvement in Economic Metrics

The Hitachi NAS Platform is a premier component of the Hitachi Data Systems file and content solutions portfolio, enabling extensive integration and consolidation opportunities for the data center.

More than 40% of organizations that purchased Hitachi NAS Platforms used reduced costs as a primary rationale for those purchases.³

Hitachi Content Platform

The Hitachi Content Platform avoids the limitations of traditional file systems by intelligently storing content in far larger quantities and in a much more efficient manner. Hitachi Content Platform provides for the new demands imposed by the explosion of unstructured data and its growing importance.

Hitachi Content Platform treats file data, file metadata and custom metadata as a single object that is tracked and stored among a variety of storage tiers. With secure multitenancy and configurable attributes for each tenant, the platform can be divided into a number of smaller virtual content platforms that present configurable attributes to support different service levels. This allows the object store to support a wide range of workloads, such as archive, data protection and even cloud from a single physical infrastructure. One infrastructure is far easier to manage than disparate silos of technology for each application or set of users.

By integrating many key technologies in a single storage platform, Hitachi Content Platform provides a path to short-term ROI and significant long-term efficiency improvements. Hitachi Content Platform helps IT evolve to meet new challenges and stay agile over the long term and to address future change and growth.

Key Features and Differentiators

- **Scalable Architecture** supports near-bottomless scalability to manage 40PB of capacity and billions of objects per cluster, for nimble and cost-efficient expansion of unstructured data without service interruption.
- **Multitenant Platform Virtualization** employs configurable attributes for tenants and namespaces, role-based management and monitoring, and built-in chargeback capabilities to uniquely segregate and secure multitenancy.
- **Data Integrity, Protection and Backup** bolster business continuity, compliance and data placement regulations with high availability and fault tolerance with no single point of failure. They use multiple data protection formats, including RAID-6, replication, data integrity checks, encryption, "write once, read many" (WORM), audit logging and version awareness. Together, these enable automated policy enforcement for retention, disposal, shredding and lifecycle management. Hitachi Content Platform's internal protection and replication capabilities reduce or eliminate the need for tape-based backup. One in 3 Hitachi Content Platform customers has successfully eliminated the need to back up to tape.
- **Multiprotocol Access** allows connectivity to a wide range of applications via HTTP, REST, NFS, CIFS and other protocols, with no custom APIs to integrate. This allows seamless data sharing and mobility, as well as support for other Hitachi storage and third-party storage for single-cluster operations and multiple applications.

³ Source: Survey of Hitachi NAS Platform users by TechValidate. TVID: 848-E57-36F

- **Storage Efficiency and Utilization** comprise state-of-the-art reduction technologies, such as single instancing (de-duplication) that eliminates on-disk data redundancy, data compression that condenses on-disk data and disposition services for regaining expired objects. This provides maximum utilization and reclamation of wasted space, and faster and more complete recovery.
- **Simplified Administration** enables a single management interface to control resources and multitenancy on the same physical storage infrastructure. IT organizations are better able to consolidate resources, lower administrative costs and scale storage on demand to meet business needs, while mitigating much of the risk and costs of controlling data growth. Hitachi Content Platform helps IT administrators manage more terabytes per staff (up to 600TB per administrator), provide e-discovery support and transparently refresh infrastructure components.
- **Integrated Search** indexes the content, system metadata and custom metadata as objects are ingested. This enables rapid retrieval and federated queries based on content and metadata across a vastly searchable repository of information. Data remains in its native format, rather than being renamed.

Mapping Product Benefits to Categories of Cost Reduction

The Hitachi Content Platform product benefits can be mapped to Storage Economics cost categories as shown in Table 2.

TABLE 2. HITACHI CONTENT PLATFORM — STORAGE ECONOMICS

Product Benefit	Storage Economics Costs Reduced	How?
Active archive	<ul style="list-style-type: none"> • Capital costs, lease expense for hardware and software • Cost of waste • Cost of growth • Storage management labor • Backup and disaster recovery labor • Backup infrastructure • Backup media • Power consumption and cooling • Data center floor space 	<p>The cost and time spent on tape-based backups is a significant drag on resources. When storing data in a replicated Hitachi Content Platform environment, tape-based backups are no longer necessary; this may result in immediate and long-term savings. This helps reduce hardware purchases and labor related to backup management (OPEX).</p> <p>By implementing active archiving with different tiers of storage, organizations can move unused or infrequently accessed data to lower tiers. Moving primary content off Tier 1 to Hitachi Content Platform helps lower the cost per terabyte for that capacity. Consequently, Tier 1 storage has been freed up and may be able to serve read and write requests more quickly, extending its useful life and improving overall application performance.</p> <p>Nearly half of all Content Platform customers have moved more than 40% of primary storage content to the platform and reclaimed capacity on primary storage systems.</p>
Compliance and e-discovery support	<ul style="list-style-type: none"> • Cost of disaster risk, business resumption • Data loss • Litigation, discovery risk • Noncompliance risk (archive, data retention) 	<p>The costs of meeting compliance requirements and providing e-discovery support are reduced, and net savings are more quickly realized, when consolidating data from existing silos and traditional systems onto Hitachi Content Platform.</p> <p>Content Platform has content-aware intelligence that reduces the manual administrative efforts (and associated costs) required to meet growing demands for always-accessible data. Content Platform is a dynamic, content-aware platform with built-in intelligence that enables complete and rapid search and retrieval of content, thereby contributing to improved data access for greater productivity of verification processes and e-discovery activities.</p>
Compression and single instancing	<ul style="list-style-type: none"> • Capital costs, lease expense for hardware and software • Cost of waste • Cost of growth • Storage management labor • Cost of duplicate data • Power consumption and cooling • Data center floor space 	<p>More than 40% of Hitachi Content Platform customers have saved 30% or greater in total capacity compared to traditional storage alternatives, speeding the time to recoup purchase costs.</p> <p>Reduction technologies such as compression and single instancing, as well as replication and protection capabilities enable organizations to reduce or eliminate tape-based backup equipment and services from core to edge. Thus, they reduce overall capital equipment purchases for backup media and footprint.</p>
Hitachi Dynamic Provisioning	<ul style="list-style-type: none"> • Cost of waste • Cost of growth • Storage management labor • Power consumption and cooling • Data center floor space 	<p>Hitachi Dynamic Provisioning allows IT administrators to delay capacity purchases by improving storage utilization across existing assets. Improving storage utilization contributes to improved ROA.</p>
Flexibility and massive capacity (scale) potential	<ul style="list-style-type: none"> • Capital costs, lease expense for hardware and software • Cost of performance • Cost of waste • Cost of growth • Storage management labor • Power consumption and cooling • Data center floor space 	<p>Hitachi Content Platform readily adapts to changes of system components (servers, storage, software version, etc.), ensuring it will stay active both during the process and long into the future. Also, Content Platform has the ability to scale significantly without incurring new infrastructure costs; it helps provide ultimate investment protection and an improved return on investment metric.</p> <p>Organizations can save an average of US\$400,000 in CAPEX costs for higher tiered storage upgrades, expansion of capacity or net-new platforms by using Content Platform (see note *)</p>
Comprehensive indexing and content preservation attributes	<ul style="list-style-type: none"> • Cost of performance • Cost of waste • Cost of growth • Storage management labor 	<p>By storing and leveraging both content and metadata together, Hitachi Content Platform and its indexing tools help unlock the value of unstructured data by creating a searchable repository of information that can be used to find new opportunities</p>
Hitachi Content Platform's advanced tool set, interface and granular control	<ul style="list-style-type: none"> • Cost of performance • Cost of waste • Cost of growth • Storage management labor 	<p>Hitachi Content Platform software helps administrators cost-efficiently grow the IT infrastructure horizontally to support multiple applications and content types while scaling vertically to support continued data growth.</p> <p>The costs associated with managing and maintaining data resources go down when administrators can centrally unify and control resources through a single interface, and ultimately manage more terabytes per administrator.</p>

TABLE 2. HITACHI CONTENT PLATFORM — STORAGE ECONOMICS (Cont.)

Product Benefit	Storage Economics Costs Reduced	How?
Multitenancy, configurable attributes and policy-based automation	<ul style="list-style-type: none"> • Cost of waste • Cost of growth • Storage management labor 	Organizations can now economically spread the cost of operating object storage across numerous uses cases and workloads to better leverage cloud and other distributed environments. Chargeback capabilities further streamline accounting activity to automatically allocate costs by user groups.
Data retention, audit tracking and disposal services	<ul style="list-style-type: none"> • Cost of disaster risk, business resumption • Data loss • Security, encryption • Litigation, discovery risk • Noncompliance risk (archive, data retention) 	Automatically marshal content as prescribed by legal and organization policies, reducing the expense, risk and complications often associated with conventional methods for keeping data to meet regulatory, legal and economic mandates.

** Note: Source: TechValidate TVID: 213-3BB-FE4*

Improvement in Economic Metrics

With its agile and robust infrastructure, Hitachi Content Platform fosters one platform for all data to eliminate the need for a siloed approach to storing unstructured content. IT organizations are better equipped to address growing demands and scope changes, thanks to "the more data, the greater the savings" dynamic of Hitachi Content Platform. More than 40% of IT organizations chose Content Platform for its ability to support different levels of service and multiple applications on a single cluster.⁴

Hitachi Data Ingestor

Hitachi Data Ingestor is a dedicated on-ramp optimized for Hitachi Content Platform. This integrated edge-to-core solution provides distributed consumers of IT, such as branch offices or cloud storage users, with a seamlessly scalable, backup-free storage solution. Deployed as a minimal-footprint or virtual appliance, Hitachi Data Ingestor sends data from the edge to a core infrastructure, employing advanced storage and data management capabilities. With this solution, organizations can greatly reduce the cost and complexity of providing IT services to geographically dispersed locations or cloud consumers.

Key Benefits and Differentiators

- **Minimal-footprint or Virtual Appliance** uses HTTP/HTTPS to move data across the LAN, WAN, Internet or virtual private network into Hitachi Content Platform, and requires no application recoding.

⁴ Source: TechValidate TVID: F21-40F-1E3

- **Native File System Access** uses common Internet file system (CIFS) and network file system (NFS) access to a distributed object store (Hitachi Content Platform). This allows file system clients instant remote access and sharing between nodes on the network, as well as preserving Microsoft Active Directory investments and LDAP authentication policies.
- **Simplified Migration** automatically replicates all resident data to Hitachi Content Platform and moves all changes once daily using automated policies that can be adjusted by an administrator to meet recovery point objectives. All content on the Data Ingestor is migrated while maintaining a local link to that migrated data via a persistent data structure that preserves previous versions when modified for immutability.
- **Bottomless and Backup-free**, Data Ingestor can recover all of its files at any time. All edge content is replicated to and centralized through Hitachi Content Platform to support the bottomless backup-free edge solution. Similar to Content Platform, Data Ingestor uses high-availability architecture with replication-aware failover and no single point of failure to ensure data availability at the edge.
- **Stubbing** allows applications to write a file to Data Ingestor, which in turn is replicated to Hitachi Content Platform. Initiated when the capacity of Data Ingestor reaches 90%, this automated process creates 4KB stubs for all files above that 90% threshold and backs up the stubs just as it would files. These stub files are transparent to users and are simply recalled.
- **WORM for Compliance** supports full read-write access to files and provides administrators with flexibility to specify retention data and WORM restrictions that best meet governance and regulatory requirements. Once a file's retention date is determined, Data Ingestor goes into compliance mode, allowing retention dates to be extended but not shortened using an internal clock and enforcement features, and also transfers these files to Hitachi Content Platform. Deletion of WORM files, renaming of WORM files and directories containing WORM files are all prohibited until the retention period expires.

Mapping Product Benefits to Categories of Cost Reduction

The Hitachi Data Ingestor product benefits can be mapped to the following Storage Economics cost categories as shown in Table 3.

TABLE 3. HITACHI DATA INGESTOR — STORAGE ECONOMICS

Hitachi Content Platform on-ramp device (bottomless, backup-less storage)	<ul style="list-style-type: none"> • Capital costs, lease expense for hardware and software • Backup and disaster recovery labor • Storage management labor • Data center floor space • Backup infrastructure • Backup media 	<p>Hitachi Data Ingestor helps organizations reduce administrative overhead at the edge by alleviating the need for IT expertise to properly manage, protect and back up storage residing at distributed locations.</p> <p>In addition, Data Ingestor provides users and applications with seemingly endless storage and a backup-less environment.</p> <p>As a result, Data Ingestor can replace local backup and archiving devices and associated hardware, software and footprint expenses.</p> <p>This helps IT organizations more accurately plan for capacity and thereby save on unnecessary edge devices to support growth. This also speeds ROI and increases the net savings.</p>
Dynamic movement of data to the most appropriate tier of storage	<ul style="list-style-type: none"> • Capital costs, lease expense for hardware and software • Cost of performance • Storage management labor • Data center floor space 	<p>Hitachi Data Ingestor reduces the total cost of storage within IT organizations by automatically moving data to the appropriate tier of storage relevant to its business need.</p>
Highly scalable extension to Hitachi Content Platform	<ul style="list-style-type: none"> • Capital costs, lease expense for hardware and software • Storage management labor • Data center floor space 	<p>Hitachi Data Ingestor helps streamline and accelerate cloud deployments and distributed IT environments. This allows for lower storage costs, as a result of eliminating the need to purchase local high-performance NAS systems, reduced complexity, lower operating expenses and capital expenditures. This ultimately equates to a lower TCO.</p>
Centralized data protection, replication and disaster recovery	<ul style="list-style-type: none"> • Backup and disaster recovery labor • Storage management labor • Backup infrastructure • Backup media 	<p>Hitachi Data Ingestor is able to provide secure, transparent data access for users at the edge with automated policies that replicate all local "resident" files to Hitachi Content Platform, enabling easy self-service options for retrieving and recovering data in the event of failure.</p> <p>Replication is automated via customizable policies, stubbing and persistent data structure to ensure data and metadata remain immutable. This saves administrative resources from manually migrating data from edge to core.</p>

Hitachi Data Discovery Suite

The Hitachi Data Discovery Suite is a scalable, index and search, appliance-based software solution. The suite sits on the internal network to facilitate uncomplicated, multiplatform end user search across multiple platforms and content types, and content index and search activities. It is designed to improve productivity and business competitiveness, and permit proactive discovery and enforcement of legal and regulatory requirements. Hitachi Data Discovery Suite provides these services primarily for Hitachi NAS Platform and Hitachi Content Platform, as well as NFS file server, Microsoft Windows file servers and NetApp devices.

Data Discovery Suite provides a single interface to drastically simplify and expedite searches across these storage systems. By using Hitachi Content Platform search tools, the scalable architecture Data Discovery Suite allows content search indexes to be built for the data stored on Hitachi NAS Platform, Hitachi Content Platform, Windows and NFS files servers as well as NetApp devices. End users can seamlessly search across the indexes with a federated query to bring back one consolidated result set. In this way, the indexes allow end users and IT staff to perform true federated searches. Users can take various actions on data depending on business needs, such as exporting, tiering or saving information. The Data Discovery Suite can also grant tier information from a result set or through a scheduled policy from Hitachi NAS Platform to lower tiered storage, and allows for self-service file restoration and file-based tiered storage.

Key Benefits and Differentiators

- **Intelligent File Tiering** upholds hierarchical storage management practices. This ensures that data is transparently moved between Hitachi NAS Platform to Hitachi Content Platform or other third-party servers, based on administrator-defined policies or search results.
- **Advanced Storage Management** employs a flexible and extensible file metadata method for assigning data to the most appropriate media as online, nearline or archive. This helps to trim backup times, retrieval times and data duplication.
- **Integrated Access Control** grants access to users based on what content they are authorized to view by supporting Microsoft Active Directory and LDAP. It also supports mapping Active Directory groups to specific roles.
- **Full Content Search** supports search across all data and metadata to bring back one consolidated result list.
- **Information Repurposing** gathers data and converts it into pertinent information for another use.
- **Archiving** uses content-aware and metadata-aware query results to trigger data migration and data movement between tiers for long-term storage and retention policy enforcement.
- **Export Information with Manifest** exports native files from either a result list or collection. This information can be imported into a third-party product, such as a search case assessment product.

Mapping Product Benefits to Categories of Cost Reduction

Of the 34 cost elements identified by Storage Economics, those in Table 4 can be mapped to and reduced by the Hitachi Data Discovery Suite.

TABLE 4. HITACHI DATA DISCOVERY SUITE — STORAGE ECONOMICS

Enterprise search	<ul style="list-style-type: none"> • Storage management labor 	<p>Hitachi Data Discovery Suite enables IT administrators to search, manage, restore and repurpose data from individual storage devices. This allows administrators to improve productivity, reduce risk and reduce costs associated with enterprise search.</p> <p>Through transparent updates, searches and the scalability to index billions of objects from one interface, Data Discovery Suite contributes to improved performance of storage assets and reduces OPEX.</p>
Intelligent file tiering	<ul style="list-style-type: none"> • Cost of waste • Cost of growth • Storage management labor • Data center floor space • Cost of performance 	<p>Through intelligent file tiering, Hitachi Data Discovery Suite helps ensure that data resides on the most appropriate storage tiers to meet business relevancy.</p>
Proactive e-discovery	<ul style="list-style-type: none"> • Storage management labor • Litigation, discovery risk • Noncompliance risk (archive, data retention) 	<p>Hitachi Data Discovery Suite reduces costs associated with the e-discovery process by defining the scope of a global pool of content. The content has been indexed for federated search and discovery across both production and archive data on the network, without manual intervention, IT expertise or additional management software.</p>
Integration with file and content solutions and existing authentication and access control frameworks	<ul style="list-style-type: none"> • Storage management labor • Litigation, discovery risk • Noncompliance risk (archive, data retention) 	<p>Hitachi Data Discovery Suite helps organizations reduce risk associated with unauthorized access of confidential or restricted data, and unintended data loss. This saves on the expense of having to recreate file security hierarchy or protocols.</p>

Hitachi Data Discovery for Microsoft SharePoint

Hitachi Data Discovery for Microsoft SharePoint is software that simplifies file and content tiering and archiving while improving performance, scalability and backup across Microsoft environments. Available for single or multiple server clusters, Data Discovery for Microsoft SharePoint connects with the Hitachi Content Platform or Hitachi NAS Platform via a familiar Microsoft SharePoint interface. This allows the data burden from SharePoint and Microsoft SQL Server systems to be shifted to Hitachi Content Platform or Hitachi NAS Platform without additional hardware or middleware.

With Data Discovery for Microsoft SharePoint, users can span and even consolidate multiple Microsoft environments and securely create, manage and share content. An optional interface is available for Hitachi Content Platform Search, which allows users to conduct structured and advanced queries of non-SharePoint content stored on Hitachi Content Platform, to search and recover individual files stored there, and to handle shredding and other retention rules.

Key Benefits and Differentiators

- **Deep SharePoint 2010 Support** is provided for single or multiserver farms. This enables seamless integration with Hitachi Content Platform and Hitachi NAS Platform, versioning of files to Content Platform, easy installation and usage, and multitenancy and namespace support on Content Platform.
- **Simplified Microsoft Searches** make use of the SharePoint search bar for simple queries and retain the searches for recurring search needs.
- **Data Relocation** uses file pointers to allow SharePoint files and bulk document libraries to be easily relocated to Hitachi Content Platform or Hitachi NAS Platform to speed performance and backup.
- **Stubbing** creates a stub when content is moved to the storage platform, based on retention rules; then it verifies the stub to preserve access pathways that reduce workload demands for the SharePoint database.
- **Microsoft Active Directory Integration** enables users to find and retrieve only the data they are authorized to access.
- **Compliance Support** is provided with Hitachi Content Platform Search. This allows organizations to create a searchable, scalable and tamperproof archive using audit trails, authoritative copies, automate shredding, and content lock-down based on retention values and chain-of-custody requirements.

Mapping Product Benefits to Categories of Cost Reduction

Of the 34 cost elements identified by Storage Economics, items listed in Table 5 can be mapped to and reduced by the Hitachi Discovery Suite for Microsoft SharePoint.

**TABLE 5. HITACHI DISCOVERY SUITE FOR MICROSOFT SHAREPOINT
— STORAGE ECONOMICS**

SharePoint integration	<ul style="list-style-type: none"> • Capital costs, lease expense for hardware and software • Cost of growth • Storage management labor 	<p>Hitachi Data Discovery for Microsoft SharePoint is fully integrated with Microsoft SharePoint. This provides for a simple, yet powerful interface to speed data relocation, and directory access across SharePoint and SQL environments. It also reduces the need to deploy and learn new software.</p>
SharePoint SQL Server data relocation	<ul style="list-style-type: none"> • Capital costs, lease expense for hardware and software • Cost of growth • Storage management labor • Cost of performance • Backup infrastructure 	<p>With Hitachi Data Discovery for Microsoft SharePoint, data is dynamically relocated from SQL Server to the appropriate tier of storage relative to the business value of the content.</p> <p>As a result, organizations are better able to cost-effectively utilize storage assets and reclaim expensive, high-performance storage capacity. This also allows for:</p> <ul style="list-style-type: none"> • Improved scalability • Improved performance and productivity • SharePoint consolidation reducing administrative spend • Reduced SQL Server backup and capacity requirements • Capacity purchases deferral • Reduced management cycles and improved productivity <p>Overall, the TCO can be significantly reduced in SharePoint environments with the adoption of Data Discovery for Microsoft SharePoint.</p>

Work Together to Reduce Cost

At this juncture, after examining the cost categories and cost savings of each product within Hitachi Data Systems file and content solutions, there is an opportunity to learn more about the compound benefits of moving to a single integrated platform for all file, block and content data. By going beyond the industry status quo of just storing data, Hitachi Data Systems has focused on enabling organizations to unlock the business potential of their data and save up to 40% in their storage costs. Getting past the normal storage mindset means understanding a file's content and metadata, and that file's requirements. It is about adding business intelligence and capitalizing on the value of the data.

A Single Integration Platform

By combining products from Hitachi Data Systems file and content solutions, which are designed to work in concert, organizations are able to swiftly and harmoniously unify data storage. At the apex is a single integration platform. This can be illustrated by integrating Hitachi NAS Platform, Hitachi Content Platform and Hitachi Data Discovery Suite, which together allow content-triggered tiering and federated search across Hitachi and third-party file-based and active archive storage infrastructure. In this way, all the data activity across solution components is orchestrated from a central interface. IT is then able to perform common management, common replication, performance monitoring and tuning, nondisruptive dynamic data movement, storage virtualization and thin provisioning from one unified point of management.

The integration continues seamlessly for organizations wanting to provide on-ramps with Hitachi Data Ingestor for intelligent core-to-edge collaboration, or when IT uses Hitachi Data Discovery

Suite for Microsoft SharePoint to lighten its SharePoint or SQL overhead burden. The plug-and-play integration of Hitachi Data Systems file and content solutions allows organizations to reap ongoing business benefits of unifying file, block and content for flexible, scalable and cost-effective growth patterns.

Collectively, the crux of the Hitachi single integration platform provides the following key capabilities:

- **Integrates File, Block and Content** by converging these services and using a single management interface for all data.
- **Intelligently Tiers** using policy-based, content-aware data movement across storage and consolidates multiple NAS systems to a Hitachi Content Platform.
- **Performs Content-aware Search** with federated discovery and retrieval across Hitachi NAS Platform and Hitachi Content Platform, as well as NetApp and other file servers, and allows search-result triggers to dynamically move data.
- **Reduces Backup** by archiving inactive data and reducing active data storage, thereby cutting backup volume and cost.

The Finer Points of Cost Efficiency

The Hitachi Data Systems file and content solutions, when working as a single integrated platform, further extend storage savings. Storage Economics considers the following cumulative cost efficiencies:

NAS (Head) Consolidation

By consolidating existing file servers or NAS into fewer Hitachi NAS Platform heads, IT can accomplish the same or greater amounts of work with less, to improve operational efficiency:

- Reduces power and cooling costs
- Condenses rack space requirements
- Minimizes points of management
- Run more data and applications with higher IOPS or throughput

File Tiering

By automating data movement according to pre-defined policies, file tiering becomes intelligent and extends flexibility, cost savings and productivity across the enterprise. Now data anywhere throughout the integrated storage nucleus can be fluidly moved and directly stored to whichever tier or location best suits the data at any given time, such as:

- Directly to Tier 1 NAS via SAS or Fibre Channel
- Directly to less expensive NAS onto SATA disk drives
- From primary storage or higher cost NAS to Hitachi Content Platform
- Directly to FAST NAS, then later migrated to SATA drives on NAS or to Hitachi Content Platform

Capacity Reclaim

When IT organizations reach this level of storage efficiency through an integrated platform, utilization

rates go up, wasted overhead goes down and storage previously allocated by unused space can be reclaimed and streamed back into the storage pool.

- Storage capacity is better managed and used.
- File systems can safely be grown beyond 70% capacity.
- Disks can reach greater population, reducing purchase costs.
- Storage capacity waste is reduced by 30% to 40%.

Archive

Across an integrated platform, IT is now able to transform efficiencies into a significantly agile data center for greater savings and lower total cost of ownership. By compressing and single instancing files and employing policy-based active archiving to move inactive files out of primary NAS or file server storage, organizations quickly attain space-efficient storage and cost efficiency.

- There is continuous dynamic movement of files between NAS and archives.
- Data migrated to archives requires less disk space to store.
- Primary NAS space is released for future years' growth.
- Purchases for more NAS can be deferred.
- Data takes on all the characteristics of archive, such as chain of command and authenticity verification requirements.
- Less total space is required to store the same or more data in the environment.

Backup

At this point, IT organizations are able to obtain remarkable savings in backup requirements, recovery capabilities and overall efficiencies across the data infrastructure.

- By archiving static content from multiple tiered storage layers to Hitachi Content Platform, archive environments remain instantly accessible, replicated and protected.
- Reducing the active working set of data on NAS also reduces the amount of backup required because archives are backed up by second replica copies.
- Replication of the archive provides data protection and reduces the need for tape-backup media and library upgrades.
- Now IT only needs to back up active production data, saving time, cost and space associated with backup environment.
- With less data to back up, the search, retrieval and recovery times are faster, and organizations are better able to meet service level agreements and dependencies.

Conclusion

As Storage Economics has examined, Hitachi Data Systems file and content solutions include products designed to swiftly improve cost and management efficiencies across data infrastructure. The separate components each provide demonstrated savings in capital and operating expenses, as well as contribute to better returns on investments and assets, ultimately lowering the TCO. When

working together as a wholly integrated platform, these file and content components can be managed centrally from a single interface for amplified cost cutting and simplification.

Perhaps the most visible improvements to the bottom line are seen with the CAPEX savings due to fewer backup requirements, system and data consolidation, and data reduction with compression and single instancing. Often less visible yet no less critical to the ledger are the soft costs associated with OPEX, as seen through the economic efficiencies of file and content solutions, including:

- Better recovery times and recovery points
- Alleviation of migration and backup
- Automation of storage policies to self-authenticate, self-manage, self-heal and self-configure unstructured data
- Provide cluster namespace and multitenancy capabilities
- Deliver content awareness for simpler, more complete data search and retrieval and records management

Hitachi Data Systems file and content solutions help to solve the difficulties of managing unstructured data in cost efficient ways. IT organizations are now able to gain that elusive control over costs, complexities and risks, while revamping how unstructured content is managed, preserved, accessed and stored. In these ways and more, Hitachi Data Systems remains intrinsically focused on customers' abilities to nimbly and cohesively address the inexorable demands of both today's and tomorrow's data.

Appendix A: Case Sidebars

Hitachi NAS Platforms Host Media Delivery Network for WebTV and Internet Provider

Internet users worldwide rely on nacamar GmbH, a leading streaming and media hosting provider in Germany, to provide the best audio and video streams live and on demand globally. To meet the growing demands for its WebTV and IPTV, nacamar required a resilient storage solution that could handle extremely high data throughput, performance, availability and scalability requirements. At the same time the solution needed to consolidate isolation environments, lower space and energy consumption, and reduce the total cost of ownership.

This tall order was fulfilled with two Hitachi NAS Platform 3100 models, powered by BlueArc®. The NAS Platforms delivered the required data throughput and IOPS, extensive scalability and connectivity, NFS and CIFS protocols support, and the capabilities to support virtualized memory pools and multi-client environments. Using the NAS Platforms, nacamar was able to create, for the first time, a media data center with over 20TB of video-on-demand content in Flash format. The systems are also able to support expansion of nacamar's memory environment to more than 4PB with linear performance scaling and without disruption. Future customer requirements, such as parallel downloads and hosting for new media customers, can be served simultaneously.

"Our business and customer demands are rapidly developing, and with the help of Hitachi storage we can not only meet these demands but also act in an environmentally conscious way," says Uwe Schnepf, managing director of nacamar, GmbH.

Hitachi Content Platform Facilitates Turnkey Operations without Backup

Payformance Corporation manages claim settlement processing for the healthcare industry. Insurance companies and medical providers rely on Payformance to handle electronic transmittals, payments and processing. A brisk data growth rate and a continual generation of new data objects every day were hampering the company's ability to effectively manage data and support high availability requirements. Payformance needed to replace its existing database and tape backup infrastructure with an active archive solution to channel data wherever and whenever it was needed, and support any number of unstructured data formats. The company wanted to eliminate backup issues, support continuous growth, replicate to a secondary site, simplify administration and lower costs.

Payformance is now able to manage and replicate more than 13 million objects in near real time between two Hitachi Content Platform units, one at the data center and the other at Payformance's disaster recovery site. The company has realized an 80% increase in administrative efficiencies since the Hitachi deployment, eliminating lengthy tape backup and manually labor-intensive management efforts, and making data readily available in various formats to customers.

"With the Hitachi Content Platform, I can set the retention features the way I want, and I just know that the data is going to be available, without worry. The bottom line for us was always better manageability, better data protection, and highest performance and availability. The Hitachi Content Platform delivers on that promise every day," says Jason Beckham, director of infrastructure at Payformance.

Appendix B: Cost Family Reduction Matrix for Hitachi Data Systems File and Content Solutions

Cost Reduction	Hitachi NAS Platform, powered by BlueArc®	Hitachi Content Platform	Hitachi Data Ingestor	Hitachi Data Discovery Suite	Hitachi Data Discovery for Microsoft SharePoint
Capital costs, lease expense for hardware and software	✓	✓	✓		✓
Cost of waste	✓	✓		✓	
Cost of duplicate data		✓			
Cost of growth	✓	✓		✓	✓
Backup and disaster recovery labor		✓	✓		
Storage management labor	✓		✓	✓	✓
Power consumption and cooling	✓	✓			
Data center floor space	✓	✓	✓	✓	
Cost of performance	✓		✓	✓	✓
Cost of disaster risk, business resumption		✓			
Data loss		✓			
Litigation, discovery risk		✓		✓	
Noncompliance risk (archive, data retention)		✓		✓	
Backup infrastructure		✓	✓	✓	✓
Backup media		✓	✓		
Security, encryption		✓			

Hitachi Data Systems Corporation

Corporate Headquarters

750 Central Expressway
Santa Clara, California 95050-2627 USA
www.HDS.com

Regional Contact Information

Americas: +1 408 970 1000 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

Hitachi is a registered trademark of Hitachi, Ltd., in the United States and other countries. Hitachi Data Systems is a registered trademark and service mark of Hitachi, Ltd., in the United States and other countries.

All other trademarks, service marks and company names in this document or website are properties of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, expressed or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems Corporation.

© Hitachi Data Systems Corporation 2011. All Rights Reserved. WP-405-B DG July 2011