

DATA DRIVEN GLOBAL VISION CLOUD PLATFORM STRATE  
ON POWERFUL RELEVANT PERFORMANCE SOLUTION CLO  
VIRTUAL BIG DATA SOLUTION ROI FLEXIBLE DATA DRIVEN

WHITE PAPER

# In the Age of Unstructured Data, Enterprise-Class Unified Storage Gives IT a Business Edge

## 7 Key Elements to Look for in a Multipetabyte-Scale Unified Storage System

By Hitachi Data Systems

April 2014

## Contents

Executive Summary .....	3
Introduction.....	4
Rapid Rise of Unstructured Data Magnifies Storage Challenges .....	5
3 Industry Trends Highlight Growing Importance of Unstructured Data .....	6
Trend 1: Ethernet Performance Is Outpacing Fibre Channel.....	6
Trend 2: Analysts Predict 50% of New VMware Deployments Will Be Over NFS .....	6
Trend 3: Oracle Over NFS Architecture Continues to Proliferate .....	6
7 Key Elements of an Ideal Unified Storage Solution for Unstructured Data .....	6
Key Element 1: Virtualized Storage .....	6
Key Element 2: Dynamic Storage Provisioning .....	7
Key Element 3: Intelligent File Tiering .....	7
Key Element 4: Unified Management Interface .....	7
Key Element #5: Flash Enablement and Deduplication .....	7
Key Element #6: Data Protection and Disaster Recovery .....	7
Key Element #7: Designed for Big Data and the Cloud.....	7
Hitachi Delivers Enterprise-Class Unified Storage Without Compromise .....	8
Benefits of Hitachi VSP G1000 and Hitachi Storage Virtualization Operating System .....	8
Summary: The Time Is Right for Enterprise-Level Unified Storage .....	9

## Executive Summary

Change is coming fast to enterprise storage. It's not just the increasing flood of data pouring into data centers; it's the fact that most of the data is unstructured from sources such as audio, video, emails, digital images, social networking data and more. The traditional method of managing file, block and object storage independently is expensive and complex, especially given the cascade of this unstructured data.

To deal with this deluge, some enterprises replacing their separate file and block solutions with unified storage systems that can manage structured and unstructured data on a common platform. While this approach is good news for organizations, there is a caveat. Not all unified storage systems are created equal. To make the transition to unified storage successful, and worth the investment and risk, the solution needs to be designed for the enterprise. It also needs to combine virtualized storage, intelligent file tiering, nondisruptive migration, a unified management interface and a near-100% availability guarantee.

This paper will help you understand 3 trends that reflect the industry's rapid movement toward accommodating unstructured data. It also contains 7 key elements of the ideal unified storage system. It depicts the features for enterprises to look for to ensure that their storage solution provides the flexibility, investment protection and cost savings they need. Transitioning now to unified storage can put you ahead of the curve. It can enable you not only to survive the influx of unstructured data but also to use that data to innovate and drive competitive advantage.

## Introduction

Change is coming fast in enterprise storage, placing greater-than-ever stress on IT environments. It is not just the increasing flood of data pouring into data centers: an estimated 40 zettabytes by 2020, representing 50-fold growth from 2010.<sup>1</sup> It is the fact that by 2015, 77% of that data will be unstructured, including audio, video, emails, digital images, social networking data and more.

The looming question is: How should enterprises handle this unprecedented growth? The traditional method of managing file, block and object storage independently is expensive and complex, especially given the cascade of unstructured data pouring into enterprise data centers. That's why enterprises are beginning to replace their separate file and block solutions with unified storage systems that can manage structured and unstructured data on a common platform.

But, not all unified storage systems are created equal. To make the transition to unified storage successful, and worth the investment and risk, the solution must be designed for the enterprise. It also needs to combine virtualized storage, intelligent file tiering, nondisruptive migration, a unified management interface and a near 100% availability guarantee.

In this paper, we take a closer look at the challenges plaguing today's storage environments, largely due to the influx of unstructured data. We examine 3 trends that reflect the industry's rapid movement toward accommodating unstructured data. And we propose 7 key elements of the ideal unified storage system. We examine the features enterprises should look for to ensure that their storage solution provides the flexibility, investment protection and cost savings they need.

---

<sup>1</sup> IDC, "The Digital Universe in 2020: Big Data, Bigger Digital Shadows, Biggest Growth in the Far East," 2012.

## Rapid Rise of Unstructured Data Magnifies Storage Challenges

Traditionally, enterprises have relied on separate point storage solutions to handle their structured and unstructured data. When storage demands increase, you simply add more storage systems. But today, the amount and variety of data is rapidly making such solutions untenable.

Enterprises still have to store massive amounts of structured and captive data in traditional database applications. But now, you also have to store unstructured data from mobile devices, electronic sensors, social media, websites and more. In fact, the amount of unstructured data is growing nearly twice as fast as the amount of structured data.

The torrent of unstructured data that is pouring into enterprise data centers is making it too costly and complex to manage file, block and object storage independently. Specific problems with the siloed storage approach include:

- **Capital costs are too high.** Continually adding storage systems to handle unprecedented data growth is cost-prohibitive.
- **Operating costs are too high.** Server sprawl leads to rising space, power and cooling costs.
- **Management is too complex.** Managing heterogeneous point solutions with multiple management systems is highly complex, leading to high labor and training costs.
- **Performance is declining.** Application performance issues are increasingly common as storage demands increase; system migrations are difficult and time-consuming.
- **Scalability is insufficient.** Traditional storage solutions are not capable of handling the massive influx of structured and unstructured data that enterprises will be seeing over the next decade.



### Is Your Enterprise Ready for Big Data? Answer These 9 Questions

The right infrastructure can make or break a company in today's competitive markets. Use this checklist to gauge the readiness of your infrastructure for big data in general, and unstructured data specifically:

- 1) **Global virtualization.** Do you have a vendor-agnostic, continuous storage infrastructure that delivers the agility of an IP-based system complementing Fibre Channel?
- 2) **Integrated active mirroring.** Are you keeping TCO down by enabling system and data center extensibility through provisioning and management of active-active volumes?
- 3) **Data center and mobility resilience.** How difficult are your data migrations? Are they self-managed and efficient, or time-consuming and costly?
- 4) **Unified management.** Do you have a unified management solution in place, or are you using multiple point solutions with high administration costs and complexities?
- 5) **Workload consolidation.** How scalable is your storage solution? Will you be able to meet growing demands at petabyte scale?
- 6) **Block and file tiering.** Do you have an automated tiering solution in place for internal and external data that achieves ideal storage economics?
- 7) **Accelerated Flash.** Do you have a flash solution to handle your most demanding, enterprise-class workloads?
- 8) **Flexible placement.** Can you optimize your data center floor space and eliminate isle hotspots with flexible placement of controllers and storage racks?
- 9) **Server virtualization.** Can you integrate your storage solutions with leading virtual server platforms for end-to-end visibility from virtual machines to the storage logical units?

## 3 Industry Trends Highlight Growing Importance of Unstructured Data

Enterprises have traditionally focused on block-based storage protocols (Fibre Channel and iSCSI), while file-based protocols (NFS and SMB) that accommodate unstructured data have played a secondary role. However, as unstructured data becomes more prevalent and important, the storage industry is shifting toward NFS and SMB protocols. Here are 3 trends that highlight this shift:

### Trend 1: Ethernet Performance Is Outpacing Fibre Channel

A key trend that highlights the industry-wide shift toward file-based protocols is the move from Fibre Channel to Ethernet. While the Fibre Channel roadmap as a network medium tapers off at 32Gb/sec, 40Gb/sec Ethernet is already well in use. The Ethernet speed that the NFS file protocol runs over will reach 100Gb/sec and its delivery is expected in 2015.

The future  
of Big Data  
White Paper

DOWNLOAD

Pundits have long predicted the demise of Fibre Channel as a network medium. Ethernet equipment is highly available, easy to deploy, less expensive and benefits from a larger knowledge workforce. It is also ideally suited to data transfers over longer distances. In the years to come, the shift to Ethernet will continue, which will further drive down costs and accelerate adoption. Ideal storage architectures will enable coexistence of the 2, as they are highly complementary in the right use cases. So, using a unified architecture will protect your investment and help soften any investment dilemmas.

### Trend 2: Analysts Predict 50% of New VMware Deployments Will Be Over NFS

VMware is also moving toward file-based protocols. The leader in virtualized servers traditionally connected its servers to storage via SAN. Now, industry analysts are predicting that more than half of its deployments will soon be over NFS.

Using an IP-based network (over Ethernet) and NFS file sharing protocol provides agility and flexibility within a virtual machine (VM) or virtual desktop infrastructure (VDI) environment. Storage resource allocation can be accomplished in fewer steps without having to worry about SAN IDs or network addressing. Therefore, VMware vSphere administrators can provision capacity on demand from a common storage pool, rather than from specific arrays or even disks.

### Trend 3: Oracle Over NFS Architecture Continues to Proliferate

Oracle is known as a structured database architecture that runs over SAN or servers. But the Oracle over NFS version is gaining more attention. The adoption of file-based implementations is more compelling than ever because of performance increases in Ethernet and IP-based protocols, and adoption of VMware over NFS.

These trends of adoption are not surprising: Deploying Oracle over NFS protocols is easier and quicker because administrators don't have to deal with RAID or LUN configuration or capacity provisioning tasks. Administrators can simply point to an NFS mount and create an Oracle folder. If the Oracle server has to be relocated, unmounting and remounting of the NFS-based data store is all that is needed. Adding new capacity is also painless relative to SAN deployments.

## 7 Key Elements of an Ideal Unified Storage Solution for Unstructured Data

Once you recognize the need to shift from siloed, block-oriented storage solutions to unified storage designed to accommodate file, block and object storage, it is time for the next step: Identify the key elements of the ideal enterprise storage solution.

### Key Element 1: Virtualized Storage

A virtualized common storage pool is necessary to manage the massive amounts of structured and unstructured data entering the data center. Global storage virtualization is virtualization applied across heterogeneous storage systems,

across VSP G1000 clusters. It can enhance capacity utilization, boost performance and simplify storage management, even over diverse platforms. It can also make it easier to expand storage as needed, without taking systems down or disrupting users. The bottom line: Global storage virtualization extends the life of 3rd-party storage systems and simplifies data migration when needed.

### Key Element 2: Dynamic Storage Provisioning

Each time an application needs more storage capacity, administrators should not have to add physical disks. Instead, the storage system should automatically allocate the necessary physical storage from a virtual pool. Doing so improves usage rates, increases application availability and reduces the time administrators spend provisioning new storage.



### Key Element 3: Intelligent File Tiering

To keep costs down, it is essential for enterprises to store data in appropriate tiers. Reserve high-performance tiers for business-critical data and maximize low-cost tiers for less critical data that is rarely accessed. Enterprise storage solutions must enable policy-based, automated movement of data to the appropriate tier at a granular level. This tiering ensures that even within a volume or file, highly referenced data is available on the highest-performing devices.

### Key Element 4: Unified Management Interface

One of the most significant challenges with traditional point storage solutions is that they require multiple management interfaces. To simplify management and reduce ongoing costs, centralized management software must administer all physical and virtualized storage solutions, from any vendor, from a single pane of glass.

### Key Element #5: Flash Enablement and Deduplication

In enterprise environments, performance is always a critical concern. The ideal storage system must provide flash storage to enable rapid access to information on the order of 1 million or more transactions per second. Database, analytics, virtual desktop and virtualized server environments all benefit from superior performance and improved response times. And since flash is still expensive storage real estate, having technologies such as primary deduplication is important to effectively store the most data in that space.

### Key Element #6: Data Protection and Disaster Recovery

Nobody likes to lose data. While every vendor claims that their data protection is best, the key is to allow the storage system to execute and manage data protection. With this approach there is less administrative overhead, and lower equipment, licensing and maintenance costs. Data protection must also deliver a performance advantage since data is being protected at the source, and not through another device over the network.

### Key Element #7: Designed for Big Data and the Cloud

Enterprises are increasingly moving toward a private cloud infrastructure, and a virtualized, automated storage system is a key component in that move. The storage system must support enterprises that are moving to and capitalizing on the private cloud IT model. The system must also offer a highly scalable virtualized infrastructure and enable solutions designed to support analytics and intelligence.



## Hitachi Delivers Enterprise-Class Unified Storage Without Compromise

Hitachi Data Systems recognizes the ever-increasing need for unstructured data storage at the enterprise level. We also know that data capacity, applications and virtual server environments are growing at exponential rates, while IT budgets are not. Data centers must become more efficient to meet these challenges. They must transition from traditionally separate file and block solutions to unified storage platforms that provide greater flexibility, investment protection and cost savings.

Hitachi Virtual Storage Platform (VSP) G1000 is the most intelligent and powerful enterprise storage system in the industry, providing top storage resource efficiencies, productivity and scalability for today's enterprises. Designed to accelerate performance of enterprise business applications, Hitachi VSP G1000 optimizes support for critical applications, cloud-ready infrastructure and data center consolidations, all through a single intuitive interface.

Hitachi VSP G1000 provides several unique features and capabilities for the enterprise market. It is the only system available today that:

- Enables central consolidation of file, block and object data across mainframe and open systems environments, as well as storage from other vendors.
- Provides a 3-D scaling storage platform designed for all data types.
- Flexibly scales for performance, capacity and virtualization of multivendor storage to optimize return on storage assets.
- Combines unsurpassed performance and capacity with the industry's lowest power and cooling requirements.
- Features a unified management interface to learn and be trained on one environment and save time and costs for retraining.



### Benefits of Hitachi VSP G1000 and Hitachi Storage Virtualization Operating System

Hitachi VSP G1000 is the latest version of the industry-leading Virtual Storage Platform. It provides even more speed and scalability for a truly enterprise-grade storage solution through its rich suite of software, including Hitachi Storage Virtualization Operating System (SVOS).

SVOS delivers an enterprise-ready global storage virtualization environment. It abstracts information from storage systems, virtualizes and pools available storage resources, and automates key data management (configuration, mobility, optimization and protection) functions. This unified virtual environment maximizes the utilization and capabilities of your storage resources. It also significantly reduces operations overhead and operations risk.

SVOS is open standards compatible for easy integration. Its industry-leading virtualization and management capabilities enable the utmost agility and control required to build infrastructures that are available, automated and agile.



**Key Updates in Hitachi VSP G1000**

- Global storage virtualization spans and integrates multiple virtualized storage platforms with active-active clustering and spanned virtualized volumes. Global active devices are the spanned virtual volumes for local and remote physical storage systems.
- Hitachi Universal Volume Manager virtualizes heterogeneous external storage. It enables older external storage with Hitachi Storage Virtualization Operating System (SVOS) capabilities, such as thin provisioning, heterogeneous replication and support for advanced server environments.
- Hitachi Dynamic Tiering virtualizes and automates mobility between tiers for maximum performance and efficiency.
- Hitachi Tiered Storage Manager enables advanced control of online volume migrations (moving data for performance or operational reasons). It provides control within a storage system or between virtualized storage systems, without server or application disruption.
- Hitachi Dynamic Provisioning enables both block and file environments to use both internal and virtualized storage more efficiently.
- Hitachi Dynamic Link Control Advanced provides multipathing and centralized management of distributed SAN multipathing configurations.

**Real-World Examples of Hitachi Virtual Storage Platform**

A leading insurance provider's storage environment consisting of several brands had become siloed, unevenly utilized, and overloaded by inactive and stale unstructured data. Older data had to be stored in a secure, tamperproof archive for several years. The company wanted to migrate its legacy data from old applications to the archive in a manner that was seamless to end users. The end users needed readily available access to the archived data in its original format.

The insurance provider chose Hitachi Data Systems to help virtualize their existing storage systems to migrate and store the legacy data. The Hitachi solution was composed of Hitachi Virtual Storage Platform, Hitachi NAS Platform and Hitachi Content Platform. The solution supported multiple application archives in an open, scalable and high-performing manner. It also enabled the company to automatically delete objects when the retention period ended. Today, the solution has been expanded to support remote locations and to create a core object repository for long-term retention.

## Summary: The Time Is Right for Enterprise-Level Unified Storage

As the amount of unstructured data increases, enterprises are reaching a turning point. Staying with traditional, siloed solutions for file and block storage will lead to increasing complexity and costs over time, leaving you in catch-up mode for many years to come. Transitioning now to unified storage puts you ahead of the curve. You will not only survive the influx of unstructured data but also use that data to innovate and drive competitive advantage.

Hitachi Virtual Storage Platform, the leading enterprise-class unified storage solution, consolidates file, block and object data, and storage from other vendors. It enables accelerated performance and optimal infrastructure growth. VSP G1000 extends those benefits, delivering the fastest and most scalable platform available today, all fully backed by the expertise of Hitachi Data Systems Global Solution Services team.

For enterprises that are feeling the pressure of unstructured data growth, as well as those organizations that want to stay ahead of the inevitable growth to come, Hitachi Data Systems is an ideal partner. As the enterprise storage leader for 2 decades, we can help you identify and deploy a unified storage solution that satisfies the resilience and availability needs of even the most demanding applications. We can help you transform your data center to be more agile, more powerful and more efficient.



**Hitachi Data Systems**

**Corporate Headquarters**

2845 Lafayette Street  
Santa Clara, CA 96050-2639 USA  
[www.HDS.com](http://www.HDS.com) [community.HDS.com](http://community.HDS.com)

**Regional Contact Information**

**Americas:** +1 408 970 1000 or [info@hds.com](mailto:info@hds.com)  
**Europe, Middle East and Africa:** +44 (0) 1753 618000 or [info.emea@hds.com](mailto:info.emea@hds.com)  
**Asia Pacific:** +852 3189 7900 or [hds.marketing.apac@hds.com](mailto:hds.marketing.apac@hds.com)

© Hitachi Data Systems Corporation 2014. All rights reserved. HITACHI is a trademark or registered trademark of Hitachi, Ltd. All other trademarks, service marks, and company names 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.