WHITE PAPER

Hitachi Data Systems Introduces Its Software-Defined Infrastructure Play

Sponsored by: Hitachi Data Systems

Eric Burgener
April 2015

IDC OPINION

One of the key challenges IT organizations face today is how to accommodate the needs of new applications in mobile computing, social media, big data and analytics, and the cloud while continuing to cost effectively manage legacy IT infrastructure and applications. As IT organizations evolve to take advantage of newer technologies, they must provide continuity for their customers running on legacy applications and equipment. As enterprises seek to deploy more agile IT infrastructures, more and more functionality is migrating to software, creating more "software-defined" infrastructures. Software-defined infrastructures allow a wide range of storage functionality to be consistently applied to different types of hardware platforms, easing the introduction of different storage architectures that may be better suited to specific types of application environments without causing disruption or putting the ability to meet service-level objectives (SLOs) at risk. This increasing emphasis on software will transform IT infrastructure in three key areas:

- **Automation.** Increased automation is required to enable administrators to simultaneously deal with data growth and SLOs and continue just to "keep the lights on" in a massively scalable datacenter environment that is becoming increasingly heterogeneous.
- **Access.** The IT infrastructure of the future must deal with structured, unstructured, and semistructured data that must be accessible through a variety of different methods – block, file, and object – to drive business insights, and software can enable access with minimal disruption and optimum efficiency.
- **Abstraction.** By moving from hardware-defined to software-defined storage functionality, software-defined infrastructures open up the agility to use different storage architectures and provision and allocate storage resources with the maximum amount of freedom to meet dynamically changing business requirements.

IN THIS WHITE PAPER

This IDC white paper discusses the key role that software-defined infrastructures will play in bridging the transition from the older, client/server-based infrastructure and applications model to the new 3rd Platform computing-based infrastructure that is dominating IT today. This document then quickly reviews the Hitachi Data Systems (HDS) Software-Defined Infrastructure announcement, assessing the new functionality in light of enterprise customer requirements demanded by the migration to 3rd Platform computing.

SITUATION OVERVIEW

Driven by the increasing importance of IT services to business success, the IT industry is in the middle of a transition to a new computing model that IDC calls 3rd Platform computing. A much more flexible computing paradigm, built around virtual infrastructure, flash, and cloud, is required to accommodate all of the new applications that are being driven by the increasing importance of mobile computing, social media, and big data and analytics. At the same time, a number of legacy workloads — relational databases, messaging and collaboration systems, and home directories — must be maintained as well. By seeking improved efficiencies, IT organizations are moving these older workloads to virtual infrastructure even as they are bringing up the newer workloads on the same infrastructure. A key challenge for IT organizations is how best to manage this transition while meeting requirements for performance, scalability, availability, data resiliency, security, and manageability in a cost-effective manner.

While scale-up storage architectures dominated in the client/server era, the new requirements of 3rd Platform computing have prompted the rise of additional architectural options. Among these newer options is what IDC refers to as software-defined storage (SDS). IDC tracks SDS as a separate market that will grow at a 22.5% compound annual growth rate (CAGR) to reach $2.8 billion in 2017 and defines SDS as follows:

- Software platforms that deliver the full suite of storage services via a software stack that uses (but is not dependent on) commodity hardware built with off-the-shelf components
- Cannot contain any proprietary hardware components like custom ASICs
- Should be able to run on multiple physical or virtual hardware instances that do not have to be factory configured by a supplier
- Must be autonomous or standalone systems

A key concept in virtual computing is the idea of the "software-defined datacenter" (SDDC), where resource allocation and functionality are logically defined in software that runs on commodity hardware. We'll delve more into the implications of this for the storage layer as we discuss the SDS infrastructure attributes that must underlie the SDDC, a concept that is largely synonymous with 3rd Platform computing.

Accommodating the needs of newer workloads like mobile, social, and big data requires new levels of scalability, agility, and manageability. Exploding growth, driven by the collection of a number of different data types (structured, unstructured, semistructured) from a variety of new sources, will drive a CAGR of 44% for data over the next five years. Regulatory, compliance, and analytical requirements encourage enterprises to store more of this data than ever before, and even midsize enterprises are often managing hundreds of terabytes of data. IT administrators need to be able to allocate IT resources for new projects literally within hours and then quickly and easily return those resources to a common pool once projects complete. Workloads need to be efficiently moved not only within but also across datacenters for performance, data protection, and recovery reasons. As administrators struggle to meet these new demands, they still need to be able to effectively manage performance, availability, security, and data protection to meet evolving business requirements.

New strategies to manage platform abstraction, data access, and automation are emerging to help storage administrators — who increasingly are IT generalists with strong virtualization skills rather than storage skills — meet these requirements. Administrators must continue to manage older scale-up storage platforms even as they add newer software-defined solutions built around scale-out designs as well as cloud-based storage. Hyperconverged architectures place compute and storage services in

each node in massively scalable scale-out architectures, providing additional performance options for certain workloads like big data and analytics. Block, file, and object data types must be managed. A software-defined management layer that can automate provisioning, performance optimization, data protection, and workflow orchestration while providing relevant and meaningful analytics across all these platforms and data types is required as nonstorage-savvy administrators are asked to manage these environments with ever-increasing spans of control.

Placing the functionality that enables abstraction, access, and automation in a software-defined layer is key to meeting the requirements in bridging the old and the new. Software-defined infrastructure introduces IT agility and flexibility that enable faster innovation and business outcomes, simplify the lives not only of administrators trying to manage these environments but also of end users trying to access services, and expedite business insight that leverages data assets for competitive advantage. SDS is key to a shift to more IT-as-a-service-based management, enabling self-service catalogs that offer automated, application-based provisioning, cloud infrastructure provisioning, and content mobility for secure workgroup sharing. Standard workflows become more intuitive, reliable, and efficient, and APIs at the different layers allow storage platforms of all types to easily integrate into preexisting or new workflows as required. Dynamic tiering capabilities manage performance optimization in real time, leveraging newer storage technologies like flash very efficiently, while quality-of-service (QoS) controls ensure that application performance meets SLOs regardless of what else is going on with other applications.

The HDS Software-Defined Infrastructure Announcement

HDS has a long history of providing high-performance, highly available, and feature-rich enterprise storage solutions, and HDS has proven to be a trusted provider across tens of thousands of customers worldwide. Like the customers of other enterprise IT providers, HDS customers have a large installed base of legacy applications — relational databases, messaging and collaboration applications, and file-based data stores — even as they are building out a wide range of newer, 3rd Platform computing workloads in mobile computing, social media, big data analytics, and cloud. The strong capabilities of HDS in enterprise storage virtualization have helped the company's clients manage heterogeneous environments and easily evolve to newer storage technologies without disrupting application services. These latest announcements build on that heritage, providing a comprehensive set of software-defined infrastructure management capabilities to help HDS customers bridge installed IT infrastructure and the newer, more agile software-defined infrastructure that is required for 3rd Platform computing environments.

The HDS Software-Defined Infrastructure announcement includes:

- Bringing all midrange and high-end enterprise storage offerings under a single software-based management environment — Storage Virtualization Operating System (SVOS) — which effectively extends HDS enterprise-class data services offering down into the midrange
- Introducing a new suite of software tools that deliver the abstraction, access, and automation necessary to effectively manage storage to meet business objectives while bridging older, more hardware-defined infrastructure and newer, more software-defined infrastructure; newly announced tools include Hitachi Automation Director, Hitachi Infrastructure Director, Hitachi Data Instance Director, and a new version of Hitachi Unified Compute Platform Director
- Releasing storage platform enhancements, including an aggressive refresh of the Virtual Storage Platform (VSP), three new Hitachi Unified Compute Platform models, and a new hyperconverged offering called the Hitachi Hyper Scale-Out Platform (HSP)

Customers are looking to software-defined solutions to make their IT infrastructures more agile and easier to manage. With this announcement, HDS is making several moves in the software-defined infrastructure area. First, HDS is extending SVOS to all of its midrange VSP models. HDS now supports a single operating environment (SVOS) that provides a consistent set of storage management capabilities across a newly invigorated line of midrange and high-end VSP models. SVOS supports proven enterprise-class data services for snapshots, clones, QoS and replication, external storage virtualization applicable to a wide range of heterogeneous storage platforms, and continuous availability that spans datacenters with its Global Active Device support. An extremely resilient platform, VSP delivers "five-nines plus" availability with hot plug everything, online drive and controller firmware upgrades, and nondisruptive data migration. SVOS also supports all-flash configurations at high densities (19.2TB/U with Hitachi Accelerated Flash) as well as hybrid configurations that leverage self-tuning automated tiering and a new feature called Hitachi Dynamic Tiering, active flash to maintain optimized performance regardless of evolving workloads. Other key enterprise features include data-at-rest encryption and storage efficiency technologies like thin provisioning and primary file data deduplication.

HDS is also introducing a new set of software management tools that help customers more easily manage a mixed environment that includes both older and newer workloads. These tools implement many of the agility and automation advantages of SDS for use with all HDS storage platforms, helping simplify provisioning, enable cross-platform data mobility, and deploy centralized monitoring and policy-based storage management for more automated and reliable operations. The goal of HDS with these products is to set up a software-defined, application-led storage provisioning and deployment model, knowing that this provides a more intuitive management paradigm well understood by the virtual administrators who are increasingly managing storage. Hitachi Automation Director provides service catalog–based administration for seasoned storage administrators to manage traditional HDS storage platforms, leveraging intelligent automation with the flexibility to create and customize infrastructure services. Hitachi Automation Director enables storage provisioning from predefined, application-specific templates and allows administrators to select and associate different levels of services (i.e., bronze, silver, gold) with particular applications when they are deployed. Hitachi Automation Director also enables workflow integration with external service portals through a REST-based API, another important capability to enable reliable self-service capabilities.

Hitachi Infrastructure Director is targeted for use more by the IT generalists who are increasingly taking over storage management responsibilities, providing a higher-level but more intuitive cross-platform management interface for midtier HDS storage environments. Infrastructure Director incorporates guided, recommended configuration practices, reducing the number of steps needed for system configuration and storage management operations. It offers the ability to configure and provision storage without in-depth knowledge of underlying infrastructure resource details. The new version of Hitachi Unified Compute Platform Director focuses on integrating cloud services orchestration and includes new automation features and other features that make the newly expanded platform portfolio easier to manage and broaden access.

Hitachi Data Instance Director provides centralized management of all snapshot-related workflows, including the creation and retention of application-consistent snapshots and clones in both block- and file-based environments, as well as associated replication-based workflows. This product provides centralized management of these related features for data protection, recovery, data mobility, and test and development purposes, among others.

New hardware platforms include the new G200 to G800 VSP models, three new Hitachi Unified Compute Platform models (UCP 1000, UCP 2000, and UCP 6000), and a new hyperconverged platform targeted for big data and analytics environments (HSP). All of these platforms sport embedded REST-based APIs to provide access for integration purposes, with the ultimate goal being a suite of device, infrastructure, and services APIs that enable full access to all platform functionality in a programmatic manner. The new VSPs have increased cores, connectivity, and cache to drive 3x the performance of the older models, supporting up to 5.8PB of raw capacity. The expansion to the Hitachi UCP portfolio includes the UCP 1000, a new low-end EVO RAIL offering running VMware Virtual SAN; the UCP 2000, based on the new VSP G200 and Brocade converged networking; and a new high-end UCP 6000, based on the larger new VSP models and either Brocade or Cisco networking. The UCP 1000 and 2000 integrate a new, lower-cost server, while the UCP 6000 includes the feature-rich Hitachi Compute Blade 2500 server.

The Hitachi HSP is a new hyperconverged, highly available platform built around a scale-out design that is targeted for big data and analytics workloads, such as Hadoop clusters, and comes with broad OpenStack compatibility. Future releases will be available with prebuilt analytics applications and solutions for targeted verticals that leverage technology from Pentaho, a company that HDS is in the process of acquiring. The Hitachi HSP accelerates solution delivery to HDS customers in the telco, healthcare, surveillance, oil and gas, automotive, and other verticals where fast analysis of massive data sets directly drives competitive advantage. The value proposition for this new scale-out platform in these environments is to provide an automated, self-managed storage environment with dramatically reduced setup time that allows customers to run virtualized applications (on KVM) at the data source, speed data ingest, and accelerate time to results for data-in-place analysis.

FUTURE OUTLOOK

The need to make storage management easier and more intuitive is driving storage management product evolution in the industry. Separating the control and data planes with software-defined infrastructure for enterprise storage solutions will become increasingly common, helping deliver the agility and cost structure needed in 3rd Platform computing environments. Particular areas of evolution have included storage provisioning, performance optimization, and cross-platform workload definition. More platforms will implement VM-aware storage management that allows storage operations like provisioning, snapshots, clones, and replication to be more easily administered at the application level by the IT generalists who are increasingly managing storage. Hypervisor-level APIs like VMware's Virtual Volumes (VVOLs) will help move the industry in this direction, enabling this level of management for legacy platforms, but increasingly newer software-defined platforms will implement VM-aware storage management as a native capability.

Performance optimization will come to include more controls to define service levels, such as storage latencies, minimum and maximum IOPS and throughput, and (for hybrid systems) cache hit rates and other metrics to evaluate the efficiencies of automated tiering approaches. Newer technologies must be efficiently leveraged to provide flash performance where needed and to help lower effective cost per gigabyte for primary storage with storage efficiency features like inline data reduction, thin provisioning, space-efficient snapshots and clones, and snapshot-based replication that leverages data differentials. With hybrid systems in particular, effective automated tiering capabilities that operate in real time at a granular level will quickly become part of the baseline feature requirement. Systems will also increasingly become self-managing to hit SLOs and will feature automatic data rebalancing in the event of system expansion or reconfiguration to help support consistent performance and rapid recoveries from failures.

Most datacenters have large investments in preexisting workflows that help them manage data protection, recovery, and other routine tasks more effectively. The more easily a system can be integrated into these workflows, the better. APIs and other scripting interfaces make this integration easier, leading to shorter deployment time frames. APIs that enable broader access to enterprise data sets, regardless of how they are actually stored, open up opportunities to use more scalable, cost-effective object-based data stores as common repositories and yet still share that data across a variety of different application types. As IDC expects more IT organizations to become hybrid cloud environments, with some on-premise IT infrastructure for certain applications and more and more applications running in the cloud, the ability to integrate both of these environments under the same business management framework will be critical.

CHALLENGES/OPPORTUNITIES

As an established enterprise IT supplier with a large installed base, HDS faces a different set of challenges than start-ups that offer only newer storage architecture products. HDS, throughout its history, has focused on allowing customers to take advantage of newer storage technologies while preserving existing investments — all the while continuing to meet extremely high requirements for reliability and availability. The software-defined infrastructure announcement is another example of this focus, providing the abstraction, access, and automation benefits of SDS across the company's installed base and newer platforms. The introduction of a hyperconverged scale-out platform (Hitachi HSP) provides a more targeted platform for big data and analytics workloads, showing that HDS is not reticent to bring new architectures into its product portfolio as those architectures prove themselves ready for enterprise deployment.

The challenge for HDS is in how quickly the company moves newer technologies into its portfolio. With a reputation for mature offerings, the company cannot afford to move too soon, but this may lead HDS to lose some business among customers that may desire earlier adoption of these technologies. The disruptive technologies of the past decade — virtualization, data deduplication, converged infrastructure, and flash — all have a solid place in the HDS enterprise storage portfolio, and this announcement brings two more of them — software-defined storage and hyperconvergence — into the fold. The strength of HDS is the company's proven track record of mature, enterprise-class storage solutions that cover a wide range of primary and secondary requirements. The challenge for HDS is in maintaining that reputation as it continues to provide new technology options for its customers in a timely manner without forcing a break with installed platforms.

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-insights-community.com
www.idc.com