

IT Infrastructure that Accelerates Digital Transformation

Companies proceeding with digital transformation have recently been experiencing the need to rapidly rebuild their business applications to handle changes in the business environment. Container technology is attracting interest as an execution platform for these business applications. This article presents a scalable data store that rapidly provides containers with persistent volumes. Also presented is a refinement of this data store in the form of a development concept for a next-generation data platform for Hitachi digital infrastructure. This platform was conceived in response to a rise in software-defined IT infrastructure. It enables optimum operation of business applications and securely stores their data. It should help Hitachi's Lumada business grow in the global market.

Shinichiro Seki
Mitsuo Hayasaka, Ph.D.
Kenta Shiga

1. Introduction

IT companies have recently been experiencing the need to harness digital technologies and digitalized information to create previously unattainable new value by reforming their business and operations. These companies are trying to find out how they can innovate their business through careful trial-and-error study, so they need to rapidly create low-cost prototypes. Meanwhile, since analyzing these prototypes often involves handling information that drives their ability to compete in the market, many companies are resistant to having the working systems output data externally. So an essential requirement for achieving the desired business innovation is a platform that will enable rapid and on-premises development of applications supporting many repeated trial-and-error cycles over a short period.

To enable fast trial-and-error cycles, the developed applications need to be deployable in increments as small as possible and in short cycles. The platform also needs to support container technology^{*1} and container orchestration tools^{*2}.

Hitachi enables these types of systems to be deployed rapidly and efficiently by providing common functions in the form of Lumada's Digital Innovation Platform. This platform uses the Docker^{*3} open source software (OSS) as its container technology, and Kubernetes^{*4} (also OSS) as its container orchestration tool. So a data store for highly

^{*1} A technology that creates logical compartments (containers) in the host operating system. All the libraries, programs, and other components needed to run applications are gathered together in a single location by each container, so that each container can be used as a separate server. Containers have less overhead than virtual servers, so they offer the benefit of lightweight, rapid operation.

^{*2} Tools enabling integrated management of containers (operations such as starting/stopping containers, allocating run hosts, and alive monitoring) when containers are operated in a cluster configuration composed of multiple hosts.

^{*3} Docker is a trademark or registered trademark of Docker, Inc. in the United States and other countries.

^{*4} Kubernetes is a trademark or registered trademark of The Linux Foundation in the United States and other countries.

reliable and scalable storage of the data used by the applications or middleware running in these containers was needed.

The first part of this article describes the method Hitachi used to implement a highly reliable and scalable data store in Lumada. The second part presents Hitachi’s product development concept for the future.

2. Data Store for Lumada Solution Hub

Lumada brings together a catalog of various solutions for digital transformation (DX) in the Lumada Solution Hub (LSH). Just by selecting a solution from the catalog, LSH can use container technology to construct an entire infrastructure as a service (IaaS)-based environment and rapidly start verifying the solution. Since Kubernetes was used as the container orchestration tool in the LSH container technology, Hitachi needed to provide a storage device to supply persistent volumes (PVs) to containers in association with Kubernetes.

2.1 Requirements

The LSH data store for containers needed to satisfy the following four requirements:

- (1) To be a scalable data store enabling supply of PVs to containers
- (2) To be able to supply readable/writable PVs from multiple containers
- (3) To enable PVs to be created/deleted on demand from the data store
- (4) To be highly reliable

2.2 Measures for Handling Requirements

The OSS-based distributed file system GlusterFS⁵ was used to handle Requirements (1) and (2). GlusterFS is used to build ecosystems together with orchestration systems such as Kubernetes, and has an extensive operation track record.

The OSS-based GlusterFS management service Heketi⁽⁴⁾ was used to handle Requirement (3). Heketi enables Kubernetes to dynamically generate, modify, or delete GlusterFS volumes, and supply them as PVs to containers.

To handle Requirement (4), the data itself was stored in the Hitachi Virtual Storage Platform (VSP) family of storage equipment, enabling the use of Hitachi’s highly reliable, high-performance data protection.

The scalable data store for containers shown in **Figure 1** was developed with these measures incorporated into it. Hitachi began providing it in July 2019.

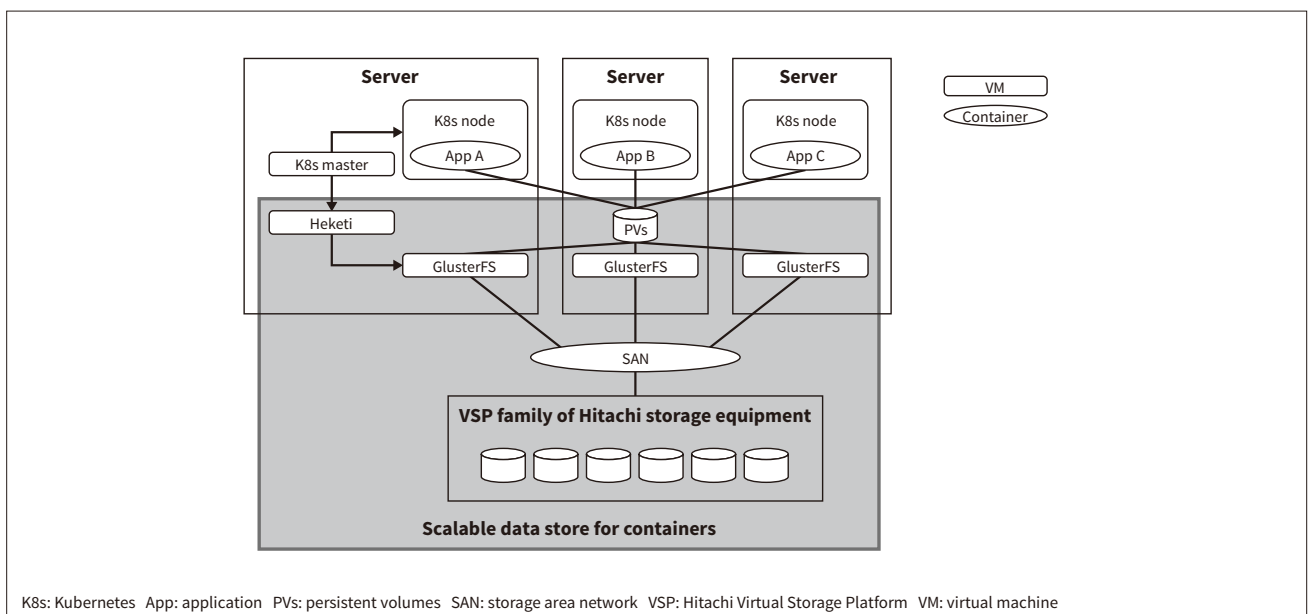
3. Product Development Concept Adapted to Software-defined Trend

To adapt to radical changes in the business environment, companies have recently been experiencing the need for IT infrastructure enabling rapid and flexible modification. The use of software-defined IT infrastructure is on the rise as a way to satisfy this need. Software-defined IT infrastructure is a framework that uses virtual versions of the server, network, and storage resources that make up IT infrastructure.

⁵ GlusterFS is a registered trademark of Gluster, Inc. in the United States and other countries.

Figure 1 — Scalable Data Store for Containers

GlusterFS and Heketi were used to create a scalable data store that works with a K8s master to dynamically provide PVs to containers. To create high reliability, the data itself was placed in VSP (Hitachi storage equipment).



It uses software instead of human operators to control the deployment of these resources or to change their configuration. This approach lets companies increase the agility and flexibility of their in-house resource deployment, change configuration, and IT infrastructure adaptation to levels rivaling a public cloud.

Hitachi is working on developing a next-generation data platform for digital infrastructure that will serve as the future model for a software-defined version of the scalable data store for containers described in the previous section. **Figure 2** provides an overview of Hitachi's next-generation platform. A conventional IT infrastructure configuration is shown on the left side of the diagram. In a conventional configuration, the central processing unit (CPU), memory and storage area resources needed by business applications and distributed file systems are tied to hardware such as servers and storage devices. When additional resources are needed, the administrator needs to decide which and how many of these resources are needed, and manually perform the installation or configuration work. And since various different types of hardware are needed, it takes time and effort to procure each item separately.

Hitachi's next-generation data platform will virtualize resources by using virtual platforms such as hypervisors^{*6}

or containers. An operation management software application called an orchestrator^{*7} will then predict the resources needed in the future and deploy the required resources to the servers in just the required amounts. This approach will let companies increase the agility and flexibility of their in-house resource deployment, change configuration, and IT infrastructure adaptation.

Virtual desktop infrastructure (VDI) is one possible application of Hitachi's next-generation data platform. Demand for VDI is rising as telecommuting becomes more widespread amid work style reforms and efforts to prevent COVID-19 transmission. The number of employees using VDI changes on a daily basis, but Hitachi's next-generation data platform will be able to rapidly and flexibly adapt to the changes.

The platform will use software running on general-purpose servers to provide storage functions—a software application known as software-defined storage (SDS). Hitachi has developed its own SDS application (Hitachi software-defined storage) that offers the following benefits:

(1) Agility of adaptation

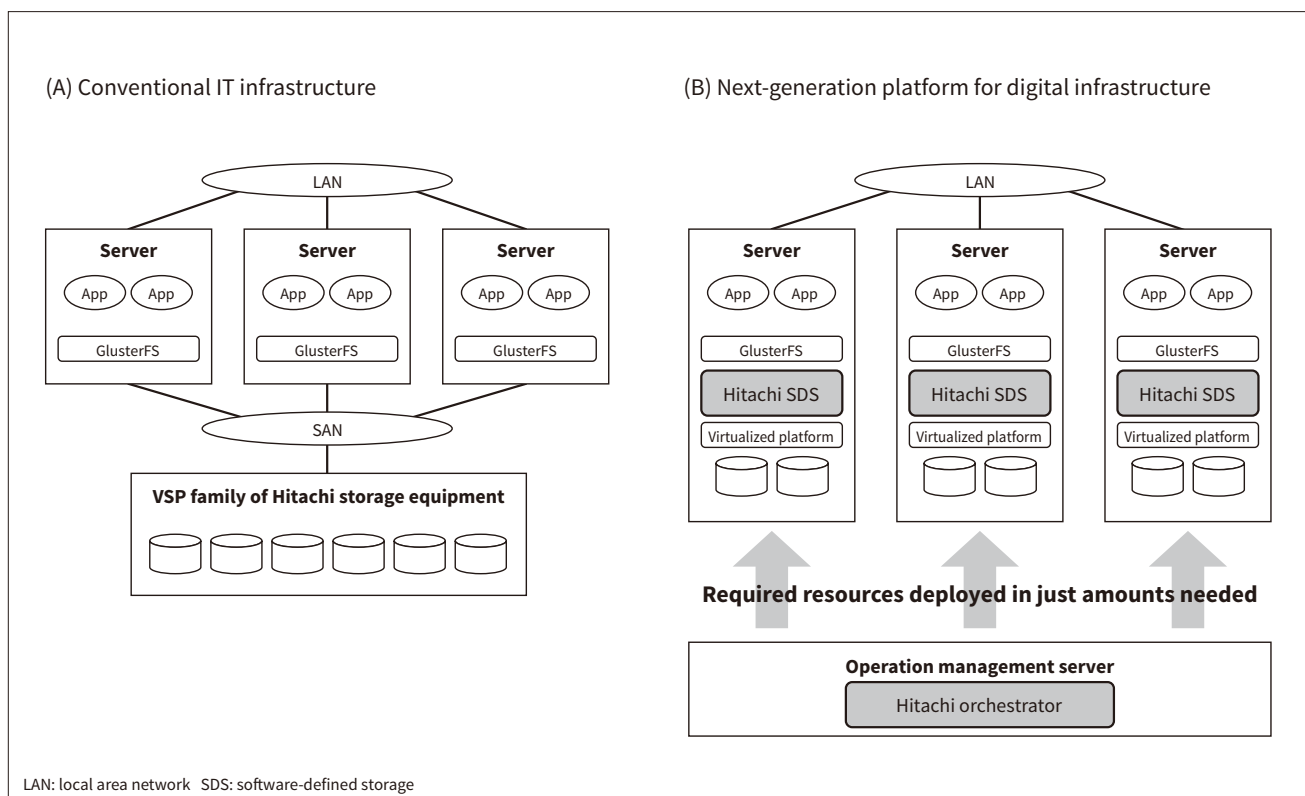
The data read/write performance or storage capacity of Hitachi SDS can be expanded just by upgrading the servers. This feature lets the user increase the resources they need

^{*6} Control programs used to create virtual machines (a computer virtualization technology).

^{*7} A software application or system used to make the configuration and management of a complex computer system automated or autonomous.

Figure 2 — Next-generation Data Platform for Digital Infrastructure

The resources of conventional IT infrastructure are tied to hardware. Hitachi's next-generation data platform virtualizes the resources and deploys the required resources to the servers in just the amounts needed, resulting in more rapidly adaptable and flexible IT infrastructure.



in just the amounts needed to adapt to radical changes in the business environment.

(2) Hitachi’s own highly efficient data protection technology

Hitachi’s next-generation data platform will use servers as the hardware, and will need to continue providing data access when a server becomes inaccessible. Conventional SDS applications satisfy this requirement by using mirroring (two sets of identical data on two servers) or triplication (three sets of identical data on three servers). However, these methods have the drawback of consuming two or three times the storage capacity needed for the original data. A technology called erasure coding is becoming a common way to solve this problem. Erasure coding generates erasure correction codes from multiple items of original data and stores the original data items and the erasure correction codes over multiple servers in a distributed manner. One benefit of this method is its ability to provide the same error tolerance (data redundancy level) as mirroring or triplication while consuming less storage capacity. Meanwhile, problem with the conventional erasure coding method is that it spreads the original data over multiple servers, resulting in the need for frequent communication between servers and a drop in data reading performance.

To solve this problem, Hitachi has developed its own erasure coding method called multi-stage erasure coding (MEC)⁽⁵⁾. Figure 3 shows the differences between conventional erasure coding and MEC. MEC stores all the original data on the local drive, speeding up reading and writing.

It also generates erasure correction codes using local server data and data sent from other servers. These features create a good balance between performance and error tolerance.

Hitachi SDS comes with MEC along with functions enabling maintenance work and drive upgrades or other configuration changes during operation. These features give Hitachi SDS higher availability and data reading performance while using general-purpose servers as hardware.

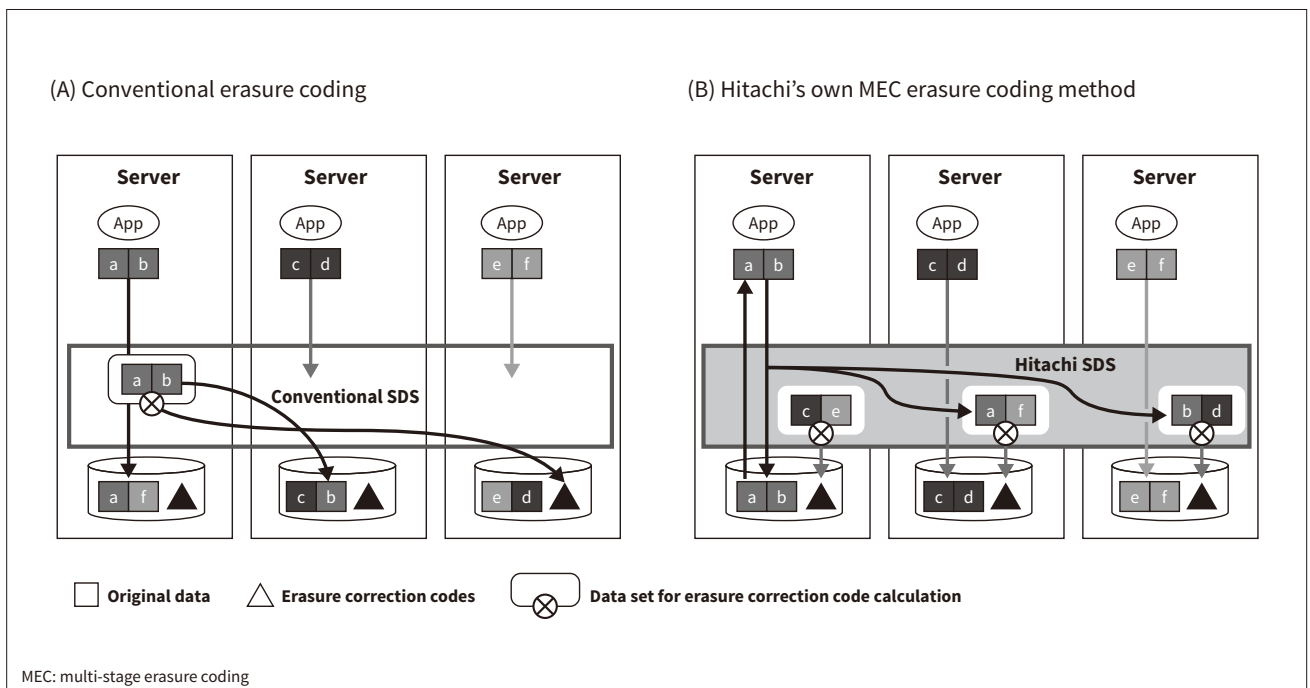
By giving Hitachi’s next-generation data platform rapid agility, high availability, and high performance, the benefits described here let the platform provide software-defined IT infrastructure possible only from Hitachi.

4. Conclusions

This article has presented a scalable data store for the containers that underpin LSH, along with the development concept for a next-generation data platform for Hitachi digital infrastructure that refines this scalable data store in response to the rise in software-defined IT infrastructure. Hitachi’s next-generation data platform will enable optimum operation of the applications and middleware used to create solutions in Digital Innovation Platform. The next-generation data platform will also be refined into hyper-converged infrastructure that will enable secure storage of the data of these applications and middleware, helping Hitachi grow its Lumada business in the global market.

Figure 3 — Hitachi’s Own MEC Erasure Coding Method

Conventional erasure coding methods suffer from lower performance by generating communication between nodes when reading data. MEC stores all the original data on the local drive. It also generates erasure correction codes from the data on the local server and other servers. This approach provides a good balance between performance and error tolerance.



References

- 1) Docker, <https://www.docker.com/>
- 2) Kubernetes: Production-Grade Container Orchestration, <https://kubernetes.io/>
- 3) Gluster, <https://www.gluster.org/>
- 4) Github: Heketi, <https://github.com/heketi>
- 5) H. Akutsu et al., "MEC: Network Optimized Multi-stage Erasure Coding for Scalable Storage Systems," 2017 IEEE 22nd Pacific Rim International Symposium on Dependable Computing, pp. 292–300 (Jan. 2017).
- 6) Hitachi, Ltd., "Lumada Solution Hub" in Japanese, https://www.hitachi.co.jp/products/it/lumada/about/lumada_hub/
- 7) Hitachi Storage Solutions in Japanese, <https://www.hitachi.co.jp/products/it/storage-solutions/>

Authors



Shinichiro Seki

Cloud and Products Service Development Department, IT Platform Products Management Division, Hitachi, Ltd. *Current work and research:* Development of solutions and services.



Mitsuo Hayasaka, Ph.D.

Data Storage Research Department, Center for Technology Innovation – Digital Technology, Research & Development Group, Hitachi, Ltd. *Current work and research:* Research and development of distributed data storage and data management. *Society memberships:* The Information Processing Society of Japan (IPSJ).



Kenta Shiga

Foundation Software Development Department, IT Platform Products Management Division, Hitachi, Ltd. *Current work and research:* Development and interoperability testing of control software of storage products. *Society memberships:* IPSJ.