IBM Spectrum Scale Software with Cisco UCS S-Series Storage Servers

Solution Brief
March 2017

Manage Your Structured and Unstructured Data with Cisco UCS S3260 Storage Servers and IBM Spectrum Scale Software

Cisco UCS with IBM Spectrum Scale Software

IBM Spectrum Scale

Highlights

Deploy Integrated Infrastructure for Software-Defined Storage Solutions

- Using Cisco UCS® Integrated Infrastructure for Big Data and Analytics and IBM Spectrum Scale, you can quickly deploy massively scalable software-defined storage.

Simplify Data Management at Scale

- Scale your solutions up—independently upgrading or adding computing, storage, and networking resources—or scale out from a small number to a large number of nodes—adding multiple systems

Protect Investments

- Optimize your budget with a modular architecture that allows you to independently upgrade computing, storage, and network components as needed, protecting your long-term investment as technology advances.

Deliver End-to-End Data Integrity, Availability, and Reliability

- Trust your information to a system with no single point of failure in both small- and large-scale storage infrastructures alike

Our software-defined storage approach makes it easy to store, access, and manage your structured and unstructured data for big data analytics.

Massive volumes of data make it difficult to store and retrieve vital information. Unless your storage solution can seamlessly scale and perform, you won't be able to support large-scale content repositories, run technical computing workloads, or complete demanding data analyses. That's why enterprises use Cisco UCS® Integrated Infrastructure for Big Data and Analytics with IBM Spectrum Scale software. With this innovative solution, you can easily deploy software-defined storage infrastructure and manage your structured and unstructured data with increased performance, visibility, and control.

The Solution

Optimizing the storage, access, and management of your structured and unstructured data requires underlying storage infrastructure that can be deployed, scaled, and managed in an agile way. Cisco UCS Integrated Infrastructure for Big Data and Analytics with IBM Spectrum Scale integrates computing, network, storage, and management resources into a cohesive programmable infrastructure. Our solution consists of Cisco UCS S3260 Storage Servers, Cisco UCS 6300 Series Fabric Interconnects, Cisco UCS Manager, and IBM Spectrum Scale software.

Cisco UCS S3260 Storage Server

The Cisco UCS S3260 Storage Server is a high-density, modular storage server designed to deliver efficient, industry-leading storage for data-intensive workloads. A modular chassis supports dual-server nodes (two servers per chassis) and up to 60 large-form-factor (LFF) drives in a 4-rack-unit (4RU) form factor. Server, storage, and network components can be upgraded independently as technology

advances. You don't need to replace the entire server; you can simply upgrade an individual component.

Each server node uses dual Intel® Xeon® processor E5-2600 v4 series CPUs and supports up to 512 GB of main memory. The system includes a Cisco UCS Virtual Interface Card (VIC) 1300 platform chip onboard the system I/O controller, offering high-performance bandwidth with dual-port 40 Gigabit Ethernet and Fibre Channel over Ethernet (FCoE) interfaces per system I/O controller.

Designed for storage-intensive environments, the Cisco UCS S3260 supports a range of hard-disk-drive (HDD) and solid-state-disk (SDD) options, as well as a pass-through controller or a RAID controller with a 4-GB cache and a host bus adapter (HBA). With a capacity of 60 drives, you can configure the storage that fits your application needs.

- Two 480-GB, 6-Gbps, 2.5-inch internal SSDs are used as boot drives.
- Up to 56 top-loading LFF HDDs are supported per chassis, with a maximum capacity of 10 TB per HDD. These drives can be mixed and matched with up to 28 SSDs (up to 3.2 TB per SSD).

A PCIe expansion slot can be used to adapt the server to your business, IT, and application requirements. You can use this PCIe slot for an additional server node; four additional LFF HDDs with up to 10 TB of capacity per HDD; or an additional PCIe expansion tray with up to two x8 half-height, half-width PCIe slots that can use industry-standard PCIe cards, including Fibre Channel and Ethernet cards.

Cisco UCS 6300 Series Fabric Interconnects

Cisco UCS 6300 Series Fabric Interconnects provide high-bandwidth, low-latency connectivity for servers, with Cisco UCS Manager providing unified management for all connected devices. The Cisco UCS fabric interconnects are a core part of the Cisco Unified Computing System™ (Cisco UCS) platform, providing low-latency, lossless 40 Gigabit Ethernet, FCoE, and Fibre Channel functions with management capabilities for systems deployed in redundant pairs. These innovative interconnects offer the full active-active redundancy, performance, and exceptional scalability needed to support a large number of nodes.

Cisco UCS Manager

The unified management capabilities of Cisco UCS Manager simplify your deployment and provisioning processes and provide the automation you need to be efficient. Using roles and polices, your IT staff can provision servers in minutes rather than the days or weeks required in traditional environments. Ongoing maintenance activities, such as firmware updates across the entire cluster, are automated. And advanced monitoring allows the system to raise alarms and send notifications about the health of the solution.

IBM Spectrum Scale Software

IBM Spectrum Scale allows you to combine flash-memory, hard-disk, and tape storage into a high-performance, low-cost, scale-out storage solution that can support your virtualization, analytics, and file and object use cases. This software-defined storage solution offers many advanced capabilities to support data-intensive applications, including content repositories, technical computing, and big-data analysis.

- **Information lifecycle management:** IBM Spectrum Scale helps lower data management costs significantly using multiple tiers of storage. Storage policies allow data to be automatically tiered, compressed, and migrated to the right storage platform. You can group your storage devices (flash memory, SSD, and HDD) based on your latency, performance, locality, and cost requirements.
- **Global file sharing:** By using Active File Management (AFM) capabilities, you can deliver the right data to the right user at the right time regardless of location. This distributed disk-caching technology expands the name space across geographic locations while delivering accelerated read and write performance to every user.
- **Scalability:** With the flexibility to scale storage capacity, performance, and protocols, you can start with a small configuration and grow to petabytes of capacity and support billions of files. The capability to prioritize I/O operations per second (IOPS) helps ensure that your IT infrastructure scales to deliver the right level of performance to users, applications, and services.
- **Performance:** Using inexpensive SSDs or flash memory located in your servers for your local cache helps accelerate I/O performance. CPUs spend less time waiting for data, and the load on your network and storage resources is significantly

reduced, allowing other applications to benefit from available bandwidth.

- **Data security:** Native encryption capabilities can help you protect your data from unauthorized access, theft, loss, and inadvertent or improper deletion. When you want to be sure that your data is deleted, cryptographic erase features provide fast and secure file deletion.
- **Data reliability, availability, and integrity:** With no single point of failure, even in large-scale deployments, you can have confidence that the system will automatically recover so that your data remains available in the event of a node, storage, or other infrastructure failure.
- **Management simplicity:** With a single name space and a single point of management, you can easily manage very large quantities of file and object data. Your administrators can monitor multiple installations from a single interface, improving visibility, control, and productivity and accelerating the provisioning, configuration, and monitoring of the cluster. In addition, integration with IBM Spectrum Control allows you to monitor multiple installations without the need to learn new tools.

Delivering Performance

IOzone was used to validate the performance characteristics of Cisco UCS Integrated Infrastructure for Big Data and Analytics with IBM Spectrum Scale. The tests used six Cisco UCS S3260 Storage Servers, each configured with twenty-eight 8-TB drives, and were connected to two 40-Gbps fabric interconnects (Figure 1).

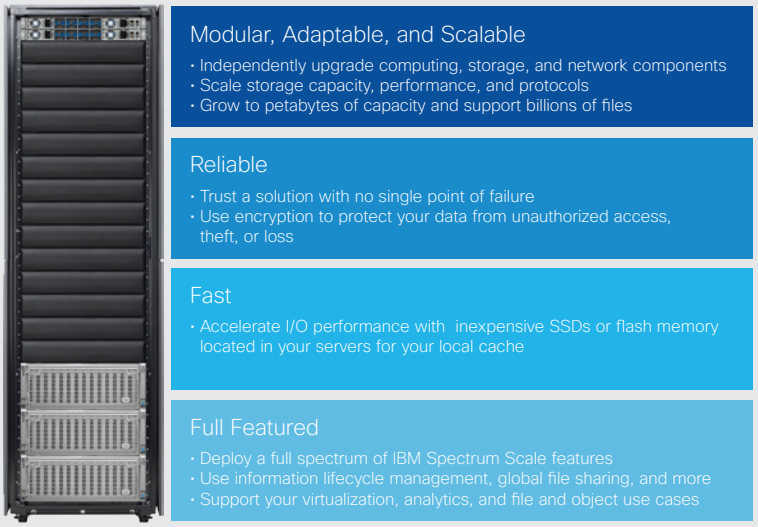

Figure 1 Solution Configuration, Features, and Benefits

Tests with Different Block Sizes

The solution was tested using multiple transfer block sizes to determine the I/O throughput characteristics of the solution. To fully saturate the servers, IOzone cluster mode was used to start multiple IOzone test threads on multiple client nodes. This procedure allowed the aggregate throughput of the server to be gauged.

Cisco UCS C240 M4 Rack Servers were used as client nodes to simulate data access patterns. Four IOzone threads ran on each of the eight client nodes, resulting in a total of 32 IOzone threads. The sequential read and write test was performed multiple times using transfer block sizes ranging from 4 KB to 16 MB (Figure 2). These results show that the system performs consistently and can adapt to many application scenarios regardless of the application's I/O block size. Because the system used two data replicates, the total throughput measured is twice that shown in Figure 2.

Tests for Scalability

Several tests were run to demonstrate the scale-out performance of the solution, with the number of storage nodes increased in each test iteration. IOzone was used to perform a sequential read and write test on multiple nodes. This test was run multiple times over a different number of server nodes. The solution scales linearly with the number of server nodes, validating that Cisco UCS servers running IBM Spectrum Scale can be used in clusters ranging from

a few nodes to thousands of nodes to support your applications and growing data volumes (Figure 3).

Trust a Proven Solution

If you need software-defined storage, consider Cisco UCS Integrated Infrastructure for Big Data and Analytics running IBM Spectrum Scale software. This enterprise platform can help you store and process your valuable data and deploy and manage your content repositories. You can simplify your data workflows, improve service levels, manage risk, and deliver business results today while positioning your data center for growth.

For More Information

For more information about Cisco UCS S3260 Storage Servers, visit http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-s-series-storage-servers/index.html.

For more information about Cisco UCS big data solutions, visit http://www.cisco.com/go/bigdata.

For more information about Cisco UCS Integrated Infrastructure for Big Data, visit http://blogs.cisco.com/datacenter/cpav4.

For more information about IBM Spectrum Scale, visit http://www.ibm.com/systems/storage/spectrum/scale.

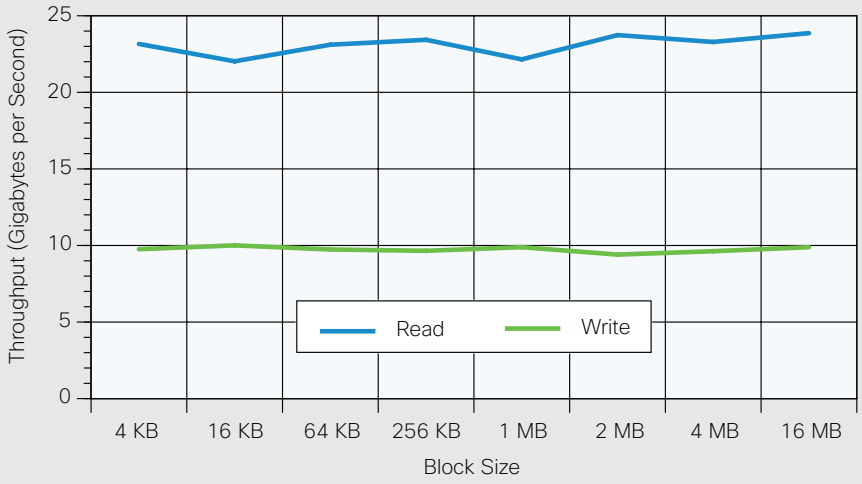

Figure 2 Test Results Using Different Block Sizes

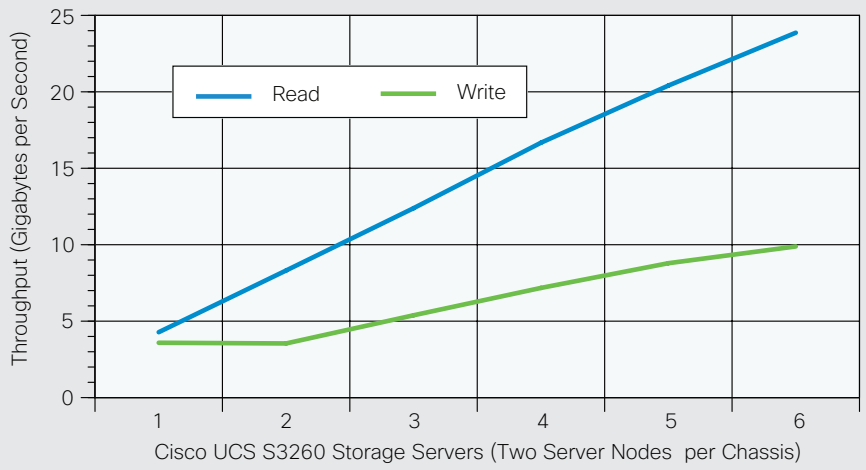

Figure 3 Scalability Test Results

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices**.

LE-58904-00 03/17