White Paper
Public

The bridge to possible

Epic ODB Data Protection with Cohesity on the Cisco UCS X-Series Modular System

Contents

This performance study is the outcome of an Epic operational database (ODB) data protection proof-of-concept (PoC) implementation conducted for a large healthcare customer using Cohesity on a Cisco® system.

Executive summary

As cyber-attacks continue to be common and ransomware attackers target healthcare organizations, the ability to effectively defend against threats and recover rapidly from a data breach has become critical. Keeping patient information available to healthcare practitioners and preventing data loss is imperative, especially when millions of patients can be affected in the event of a disaster or cyber-attack. To mitigate risk from an outage, healthcare organizations need to ensure secure backup of data and systems and the ability to access a clean copy of data for rapid recovery. With patient data residing across many data centers that may be in different states and countries, the migration of data also needs to be seamless to prevent emergency outages.

This white paper documents the outcome of a proof-of-concept (PoC) implementation undertaken for a healthcare customer using Cohesity on the Cisco UCS® X-Series Modular System. The PoC demonstrated that the joint solution can perform backup and recovery of patient data in record-breaking time. The organization stores and manages patient data using Epic electronic health record (EHR) software running on InterSystems operational databases. Because these databases are rapidly growing and vary in size from 30 to 80 terabytes (TB), speed crucial when protecting and restoring them. Designed with multilayered protection to safeguard data against cyber-attacks, Cohesity on the Cisco UCS X-Series Modular System radically simplifies backup and recovery across a broad set of data sources including network-attached storage (NAS) and offsite data centers.

Epic deployments use operational database (ODB) products from InterSystems. Many InterSystems customers have migrated or will be migrating from InterSystems Cache' to the InterSystems IRIS platform product. Additionally, many customers and vendors opt to consolidate these products under the generic term "Epic ODB." In this document, we simply refer to "ODB," using this term to refer to either product. We feel this is a fair approach because nothing, in terms of our findings, is substantively different whether using an underlying InterSystems Cache' or InterSystems IRIS database.

Cisco UCS X-Series Modular System cloud-managed with the Cisco Intersight™ platform is built to meet the high demands of modern applications and improve operational efficiency, agility, and scalability.

The healthcare customer is the PoC study wanted to achieve backup and restore performance for a 60-TB ODB in about 12 hours. The primary database was provisioned from a high-performing all-flash storage subsystem. Cohesity Data Cloud and the backup host (Cohesity Linux Agent) were deployed across four to five nodes of all-Non-Volatile Memory Express (NVMe) Cisco UCS X210c Compute Nodes installed in the latest Cisco UCS X-Series Modular System. The main results of this study are summarized here:

- With the ease of provisioning and managing X210c nodes in a Cisco UCS X-Series Modular System through the Cisco Intersight platform and with automation playbooks available from Cisco DevNet Code Exchange, the Cohesity cluster was operational in just a few hours.
- A four-node Cohesity cluster deployed across four all-NVMe X210c nodes was provisioned. A Cohesity Linux Agent was also installed in another X210c node deployed in the same Cisco UCS X-Series Modular System.
- A 20-tebibyte (TiB) ODB data set was generated, and a full backup of the entire data set was completed in 2 hours and 11 minutes, achieving backup performance of 10 TB per hour (TB/hr),

twice the required backup performance. The addition of a fifth X210c node to the Cohesity cluster increased backup performance to 11.24 TB/hr.

- The same 20-TiB data set was restored in just 2 hours and 5 minutes, achieving restore performance of 10.5 TB/hr, again twice the required restore performance. The addition of a fifth X210c node to the Cohesity cluster increased backup performance to 11 TB/hr.
- The solution was tested for the resiliency of the restore process during failure of one of the nodes in the Cohesity cluster. The Cohesity cluster was expanded to five nodes deployed across five X210c nodes. During the restore process, a node was caused to fail. This study showed that the restore process still was successful without any increase in restoration time, demonstrating the a customer's data would be resilient in the event of mishaps and continuously available when needed, especially in an emergency.
- The Cohesity Linux Agent installed in one of the X210c nodes in the Cisco UCS X-Series system demonstrated read-write network throughput of from 3.1 to 3.5 GBps, as measured during the backup and restore process.
- The Fibre Channel logical unit number (LUN) storage mounted on the backup host (Cohesity Linux Agent) was identified as the primary bottleneck that restricted performance beyond about 11 TB/hr.

This document helps demonstrate to Cisco® customers and business partners an all-in-one data management solution with Cohesity. In partnership, Cisco and Cohesity have captured data to show how easy data management can be, as well as how fast data can be backed up, restored, scaled out, and secured. The detailed test procedure and results are presented in the following sections.

Audience

The intended audience for this document includes sales engineers, field consultants, professional services, IT managers, IT engineers, partners, and customers who are interested in learning about the process of backup, recovery, and data protection. Overall, anyone looking for data protection in a timely manner should look at the performance achieved described at the end of this document.

Introduction

Year over year, the accumulation of data is increasing, and data protection across backup and recovery is more critical than ever. This document shows that with Cisco and Cohesity systems, an Epic customer can back up and restore data in a timely manner no matter how big the database currently is or will become in decades to come. A typical customer's database is between 30 and 80 TB and can grow substantially when consolidated with other sites as a result of business consolidation, new site build-outs, etc. To keep up with demand, organizations must have a simple way to manage data, and with Cohesity and Cisco UCS X-Series systems, the solution is clear.

Cisco UCS X-Series for Cohesity Data Cloud

The Cisco UCS X-Series with the Cisco Intersight platform is a modular system managed from the cloud. It is designed to be shaped to meet the needs of modern applications and improve operational efficiency, agility, and scale through an adaptable, future-ready, modular design. The Cisco UCS X-Series Modular System supports both blade and rack servers by offering computing density, storage capacity, and expandability in a single system.

The Cisco UCS X-Series Modular System is certified with Cohesity Data Cloud. The solution uses all-NVMe X-Series nodes (X210C) equipped with two third-generation (3rd Gen) Intel® Xeon® Scalable processors and 91.8 TB of all-NVMe storage per node, providing both computing and storage resources with exceptional backup and recovery performance (Figure 1).

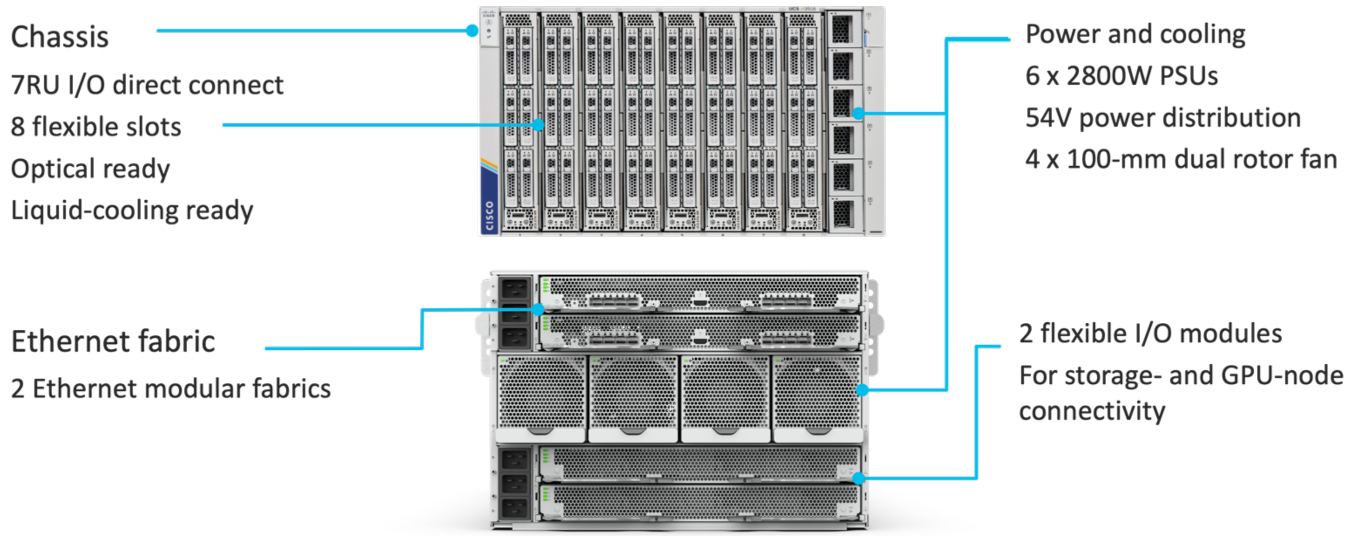

Figure 1.
Cisco UCS X-Series Modular System with Cohesity Data Cloud

Epic operational database protection

Epic is the leading provider of electronic health record (EHR) software. Hundreds of millions of patients have a current electronic record in Epic. This technology supports the entire end-to-end patient care experience: registration and scheduling systems, clinical systems for medical personnel, systems for pharmacists and radiologists, and billing systems for insurers.

The main components of an Epic EHR system are a database server, an Epic reporting server, a read-only server, and an Epic disaster-recovery server. Each of these components uses the ODB. The storage for this database is hosted on a high-performing storage array.

Epic EMR workloads create and use ODBs with:

- A limited number of data files
- Data files often over 1 TB in size
- Files that often may change daily
- Files that are highly compressible, in the range of 12 to 20x.

Two main methods are used to back up the InterSystems ODBs used by Epic:

- Mount the host or backup host using agents from backup software vendors.
- Integrate the storage array and use an API and snapshot differences to find changes blocks.

In this solution, the ODB is protected through the mount host or backup host installed with the Cohesity Linux Agent. Database backup and recovery is conducted by exposing the volume of the associated databases from the disaster-recovery site to a Linux proxy server or AIX server. Subsequently, the file system hosted on the agent server is backed up through certified data protection software such as Cohesity Data Cloud.

In the present solution, the backup host server (Cohesity Linux Agent) was deployed on a high-performing Cisco UCS X210c node hosted on a Cisco UCS X-Series modular chassis.

Figure 2 shows the backup and recovery process for an ODB using the Cisco UCS X-Series with Cohesity Data Cloud.

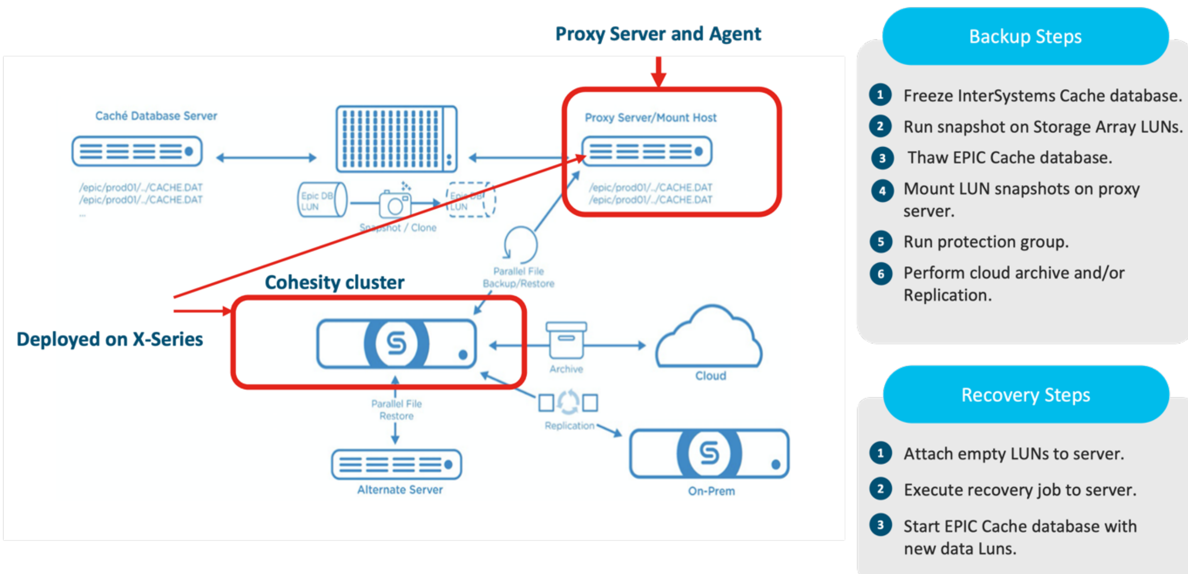

Figure 2.
ODB backup and recovery using Cisco UCS X-Series with Cohesity Data Cloud

Solution deployment details

Figure 3 shows the main hardware and software components deployed to test backup and recovery performance for the ODB.

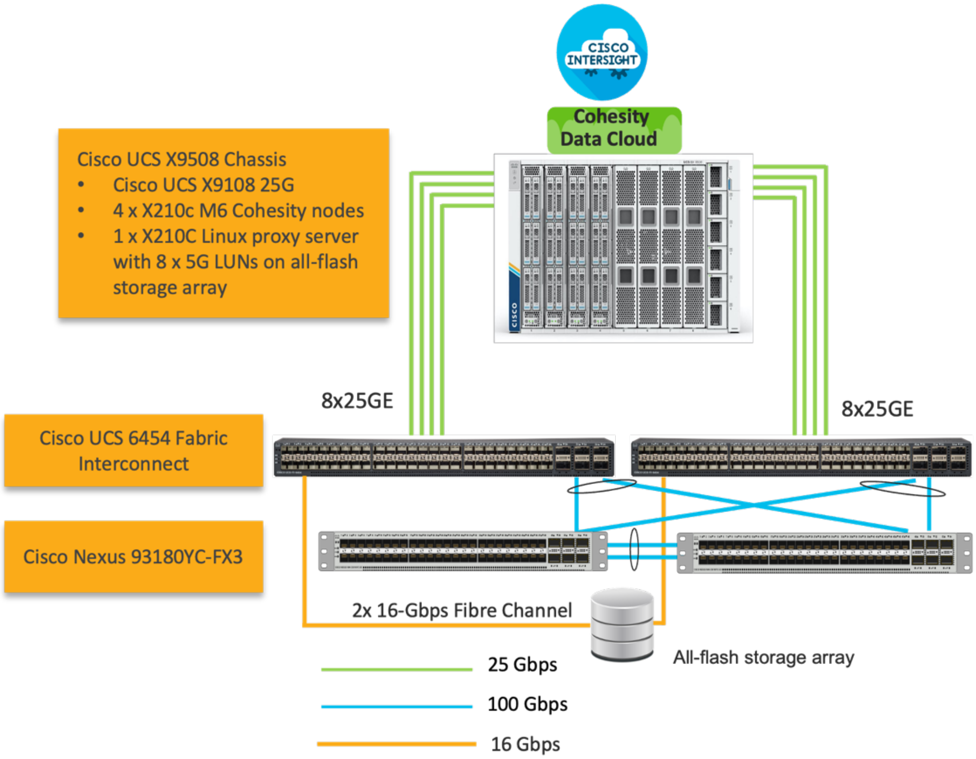

Figure 3.
Main solution components

The main aspects of the deployment configuration are described in the following sections.

- Cohesity Data Cloud was deployed across five all-NVME X210c nodes. Each node was equipped with two 3rd Gen Intel Xeon Scalable processors, 384 GB of memory, and 91.8 TB (6 x 15.3-TB NVMe) of storage per node, providing both computing and storage resources.
- The X-Series modular system was deployed through the Cisco Intersight platform, a software-as-a-service (SaaS) infrastructure lifecycle management solution delivering simplified configuration, deployment, maintenance, and support services with ease of management across a hybrid cloud environment.
- The backup host (Cohesity Linux Agent) was deployed on an X210c node deployed in the same Cisco UCS X-Series chassis hosting Cohesity Data Cloud.
- The all-flash storage array hosting the ODB was directly attached to Fibre Channel ports configured on the Cisco UCS 6454 Fabric Interconnect.
- Epic EMR data was generated to populate the associated ODB through the Epic iO simulator, a propriety data-generation tool used by Epic.

Cisco UCS X-Series configuration

The Cisco UCS X-Series Modular System was deployed through the Cisco Intersight platform, a SaaS infrastructure lifecycle management solution that delivers simplified configuration, deployment, maintenance, and support services with ease of management across a hybrid cloud environment.

Figure 4 illustrates the Cisco UCS X-Series Modular System deployed and managed through the Cisco Intersight platform.

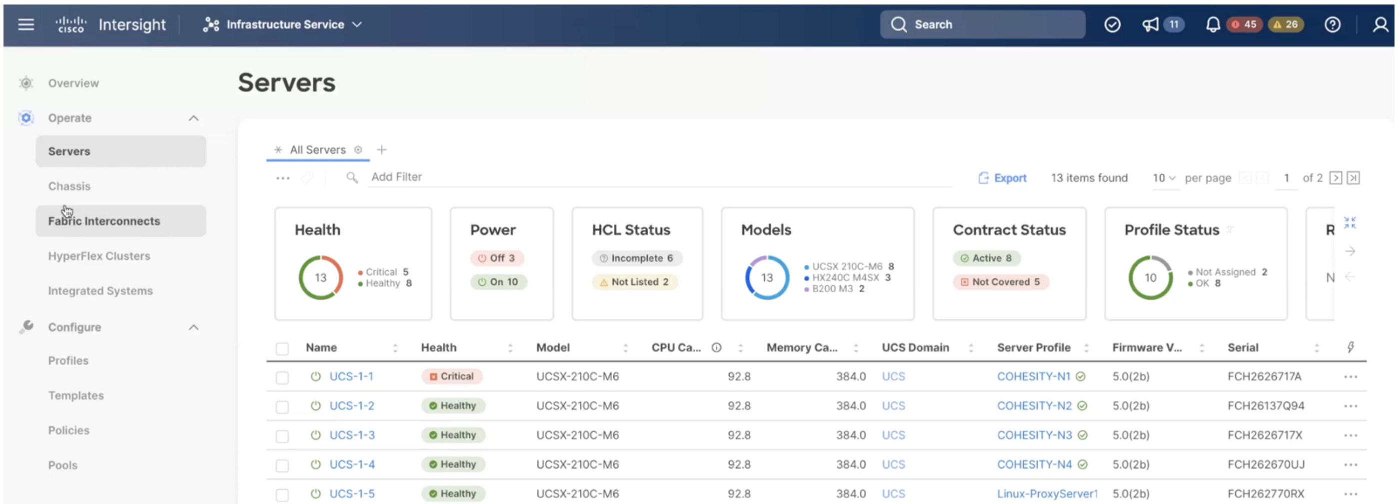

Figure 4.
Cisco UCS X-Series Modular System managed through the Cisco Intersight platform

Cohesity Data Cloud was deployed across four all-NVME X210c nodes. Each node was equipped with four all-NVMe X210c nodes with two 3rd Gen Intel Xeon Scalable processors and 91.8 TB of storage per node, providing both computing and storage resources.

Figure 5 shows the X210c node configured through the Cisco Intersight platform and managed through a server policy-based model.

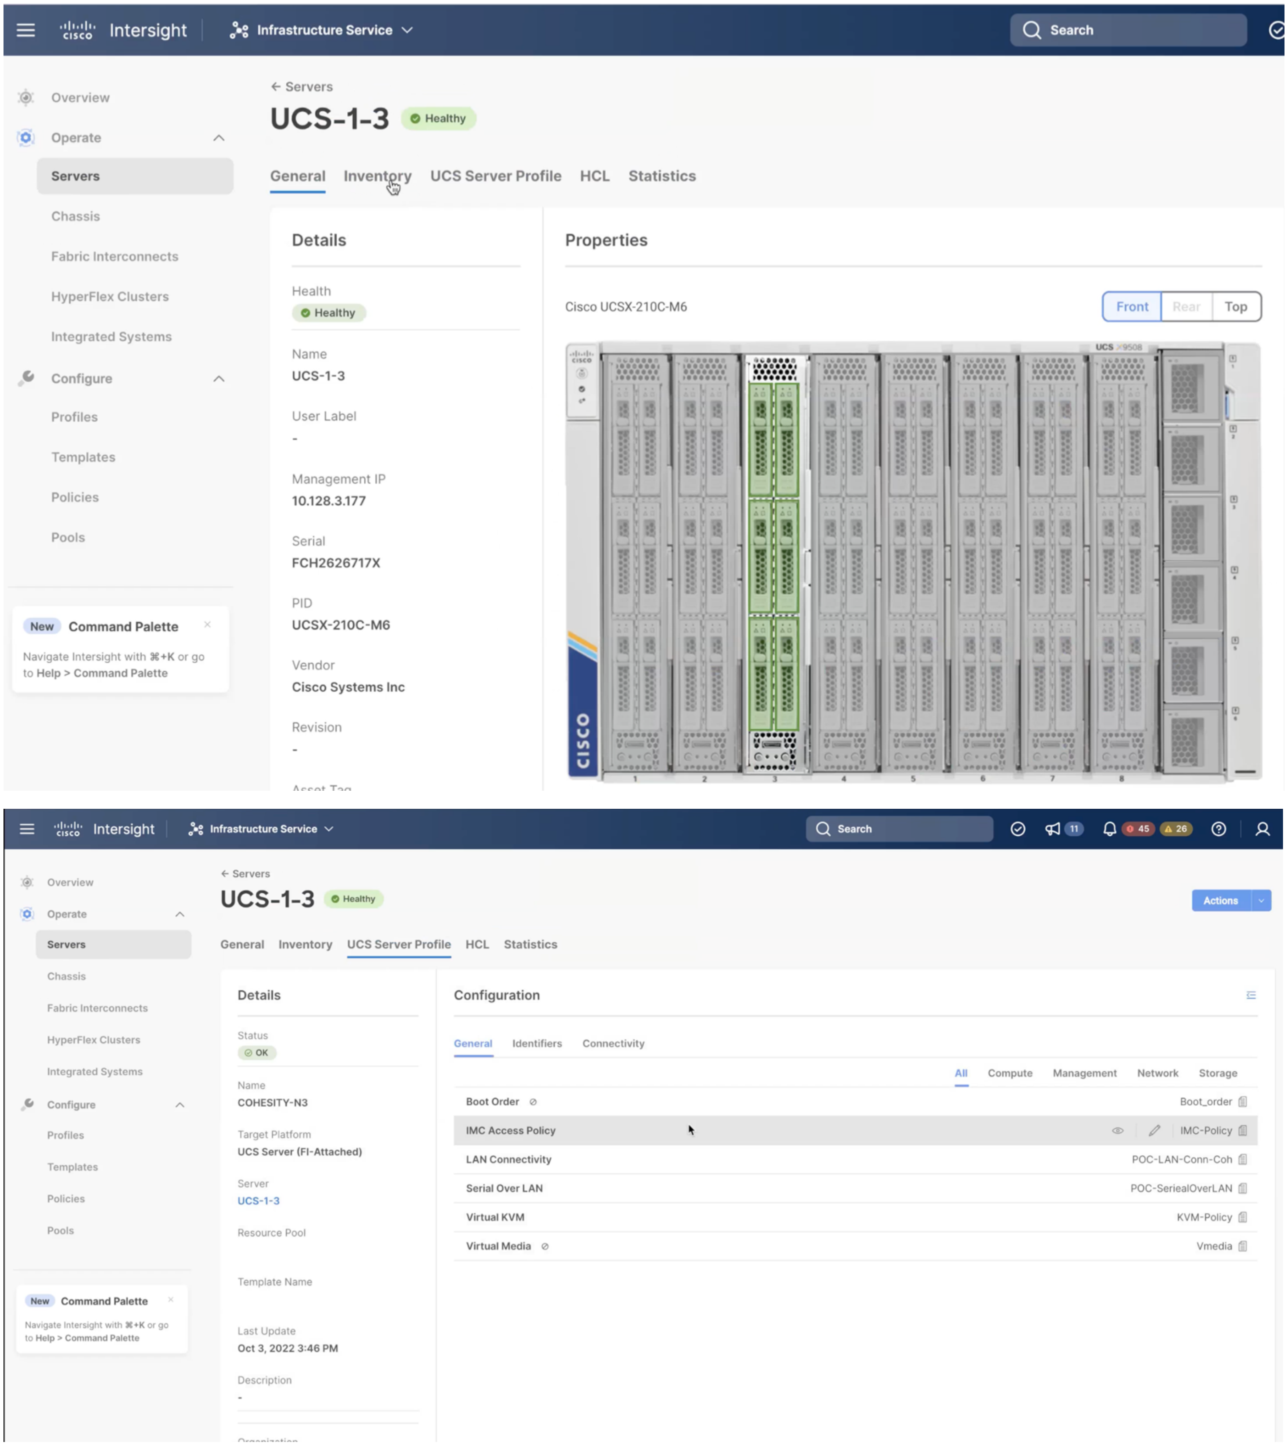

Figure 5.
Cisco UCS X210c node configured through the Cisco Intersight platform

The all-flash storage array hosting the ODB was directly attached to Fibre Channel ports configured on the Cisco UCS 6454 Fabric Interconnect.

Figure 6 shows the Fibre Channel ports configured on the Cisco UCS 6454 and directly attached to an all-flash storage array.

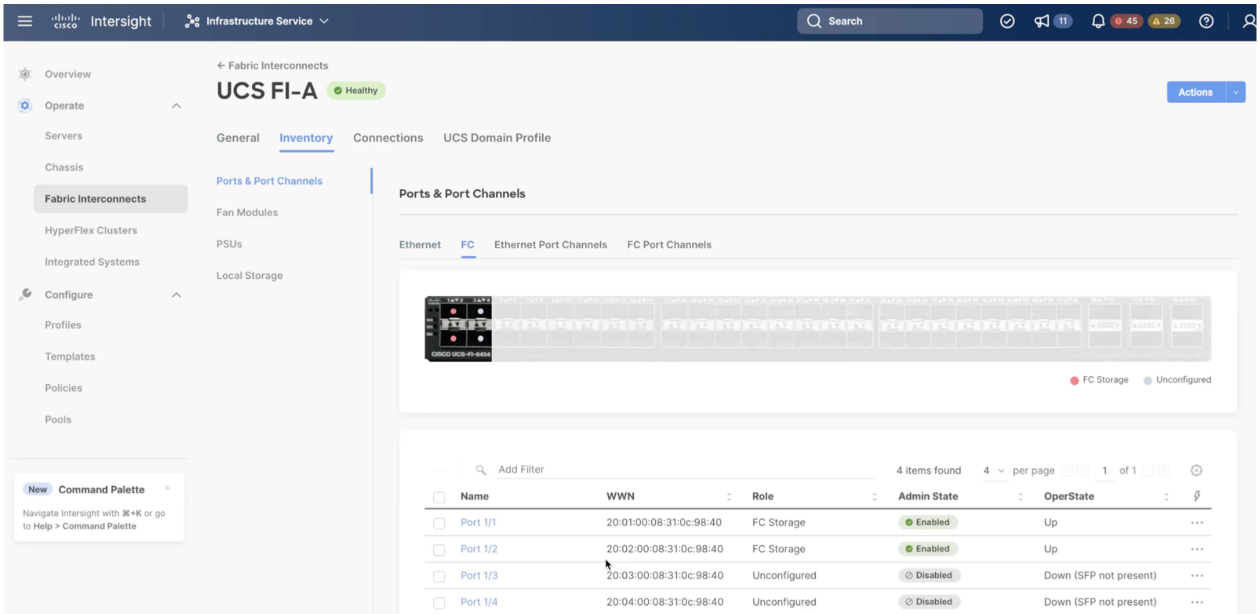

Figure 6.
Fibre Channel ports configured on the Cisco UCS 6454 and directly attached to an all-flash storage array

The backup host (Cohesity Linux Agent) was deployed on an X210c node deployed in the same X-Series chassis hosting Cohesity Data Cloud. The backup host was mapped with eight 5-GB LUNs carved on the all-flash storage array.

Figure 7 shows the Fibre Channel LUNs carved on an all-flash storage array.

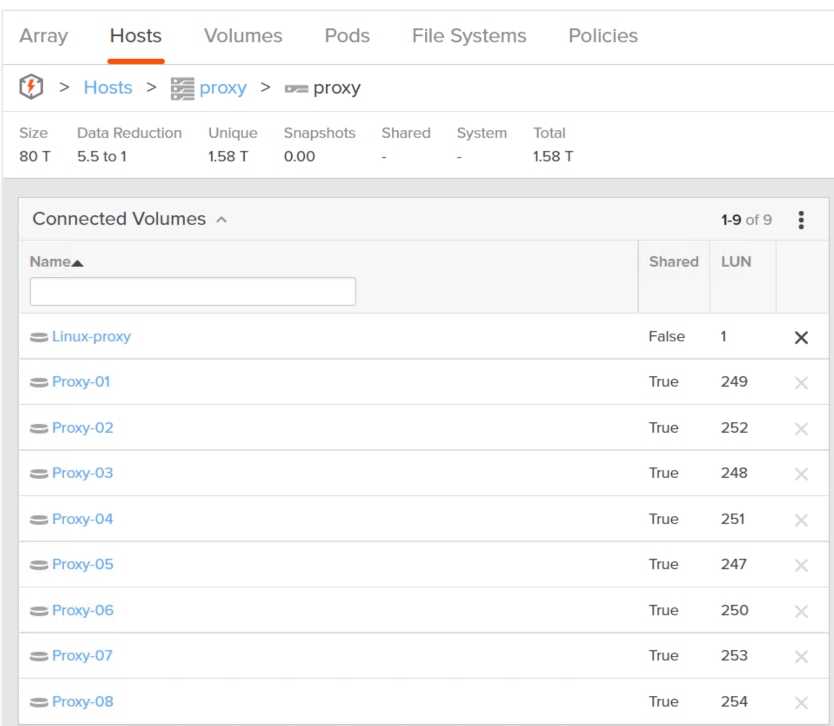

Figure 7.
Fibre Channel LUNs carved on an all-flash storage array

Figure 8 shows the logical volume created across the eight Fibre Channel LUNs hosted on the Cisco UCS X210c proxy node with the Linux OS.

```
[root@proxy ~]# lvm
lvm> pvs
  PV                                            VG      Fmt  Attr PSize    PFree
  /dev/mapper/3624a937072d6942b333f4aa3000117d0 data-vg lvm2 a--    <5.00t 12.79g
  /dev/mapper/3624a937072d6942b333f4aa3000117d1 data-vg lvm2 a--    <5.00t 12.79g
  /dev/mapper/3624a937072d6942b333f4aa3000117d2 data-vg lvm2 a--    <5.00t 12.79g
  /dev/mapper/3624a937072d6942b333f4aa3000117d3 data-vg lvm2 a--    <5.00t 12.79g
  /dev/mapper/3624a937072d6942b333f4aa3000117d4 data-vg lvm2 a--    <5.00t 12.79g
  /dev/mapper/3624a937072d6942b333f4aa3000117d5 data-vg lvm2 a--    <5.00t 12.79g
  /dev/mapper/3624a937072d6942b333f4aa3000117d6 data-vg lvm2 a--    <5.00t 12.79g
  /dev/mapper/3624a937072d6942b333f4aa3000117d7 data-vg lvm2 a--    <5.00t 12.79g
  /dev/sda3                                     rhel    lvm2 a--  <222.31g     0
lvm> lvs
  LV    VG      Attr       LSize   Pool Origin Data%  Meta%  Move Log Cpy%Sync Convert
  tst01 data-vg -wi-ao----  39.90t
```

Figure 8.
Logical volume created across eight Fibre Channel LUNs hosted on the Cisco UCS X210c proxy node

Cohesity Data Cloud

Cohesity enables a Health Insurance Portability and Accountability Act (HIPAA)–compliant platform for Epic EHR environments on ODBs.

Figure 9 shows Cohesity Data Cloud deployed on the Cisco UCS X-Series Modular System.

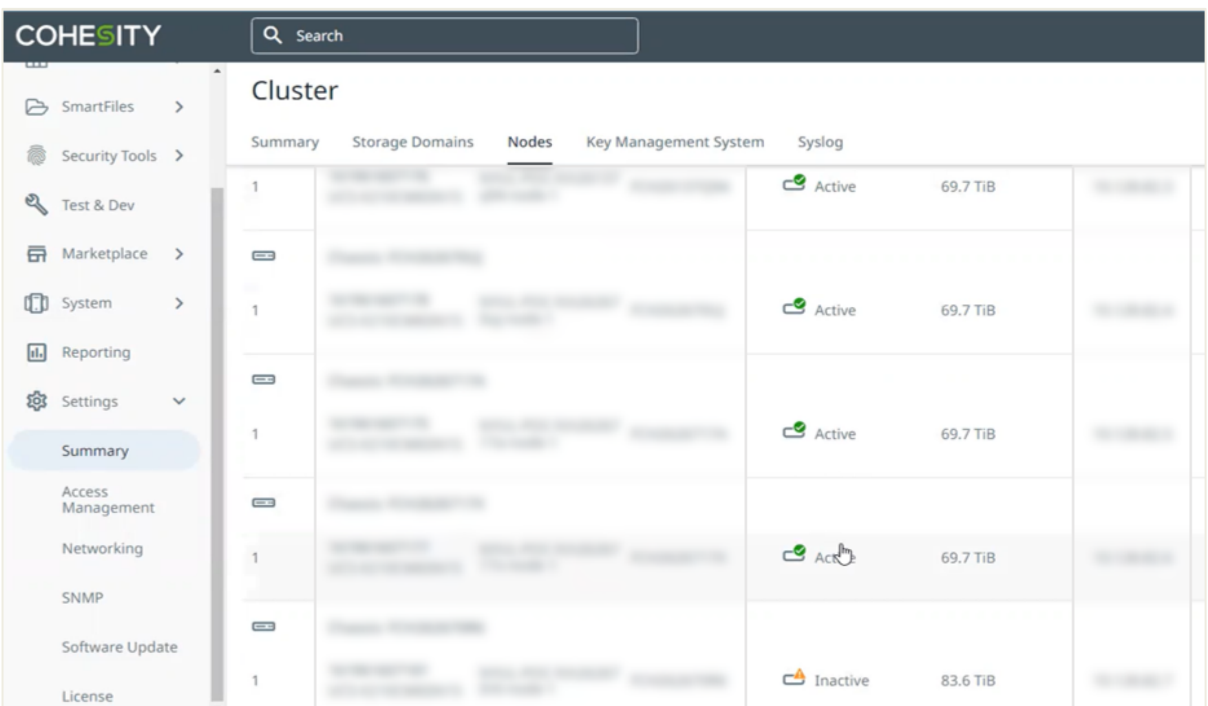

Figure 9.
Cohesity Data Cloud deployed on Cisco UCS X-Series Modular System

The Cohesity Agent for backup of the ODB was deployed on the Linux proxy server hosting the ODB backup data.

Figure 10 shows the deployment of Cohesity Agent.

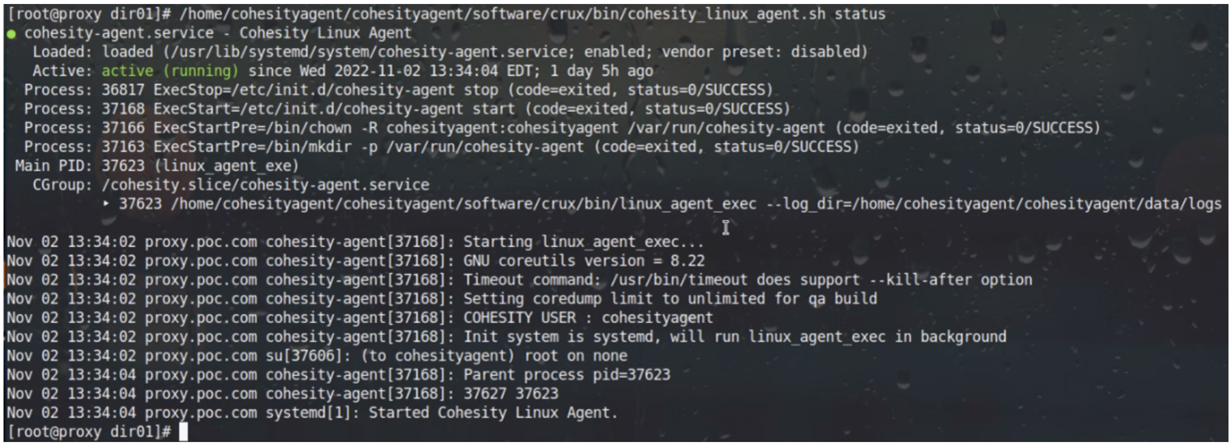

```
[root@proxy dir01]# /home/cohesityagent/cohesityagent/software/crux/bin/cohesity_linux_agent.sh status
● cohesity-agent.service - Cohesity Linux Agent
   Loaded: loaded (/usr/lib/systemd/system/cohesity-agent.service; enabled; vendor preset: disabled)
   Active: active (running) since Wed 2022-11-02 13:34:04 EDT; 1 day 5h ago
  Process: 36817 ExecStop=/etc/init.d/cohesity-agent stop (code=exited, status=0/SUCCESS)
  Process: 37168 ExecStart=/etc/init.d/cohesity-agent start (code=exited, status=0/SUCCESS)
  Process: 37166 ExecStartPre=/bin/chown -R cohesityagent:cohesityagent /var/run/cohesity-agent (code=exited, status=0/SUCCESS)
  Process: 37163 ExecStartPre=/bin/mkdir -p /var/run/cohesity-agent (code=exited, status=0/SUCCESS)
 Main PID: 37623 (linux_agent_exe)
   CGroup: /cohesity.slice/cohesity-agent.service
           ‣ 37623 /home/cohesityagent/cohesityagent/software/crux/bin/linux_agent_exec --log_dir=/home/cohesityagent/cohesityagent/data/logs

Nov 02 13:34:02 proxy.poc.com cohesity-agent[37168]: Starting linux_agent_exec...
Nov 02 13:34:02 proxy.poc.com cohesity-agent[37168]: GNU coreutils version = 8.22
Nov 02 13:34:02 proxy.poc.com cohesity-agent[37168]: Timeout command: /usr/bin/timeout does support --kill-after option
Nov 02 13:34:02 proxy.poc.com cohesity-agent[37168]: Setting coredump limit to unlimited for qa build
Nov 02 13:34:02 proxy.poc.com cohesity-agent[37168]: COHESITY USER : cohesityagent
Nov 02 13:34:02 proxy.poc.com cohesity-agent[37168]: Init system is systemd, will run linux_agent_exec in background
Nov 02 13:34:02 proxy.poc.com su[37606]: (to cohesityagent) root on none
Nov 02 13:34:04 proxy.poc.com cohesity-agent[37168]: Parent process pid=37623
Nov 02 13:34:04 proxy.poc.com cohesity-agent[37168]: 37627 37623
Nov 02 13:34:04 proxy.poc.com systemd[1]: Started Cohesity Linux Agent.
[root@proxy dir01]#
```

Figure 10.
Cohesity Agent deployment

Epic iO simulator

The Epic iO simulator is a proprietary Epic data-set generation tool that creates the data files that simulate ODB files. It generates data in a format very similar to that of the real ODBs and creates files stored within a set of directories. In the present test, a total data set of about 20 TB was generated through this tool.

Figure 11 shows the ODB data generated through the Epic iO simulator tool.

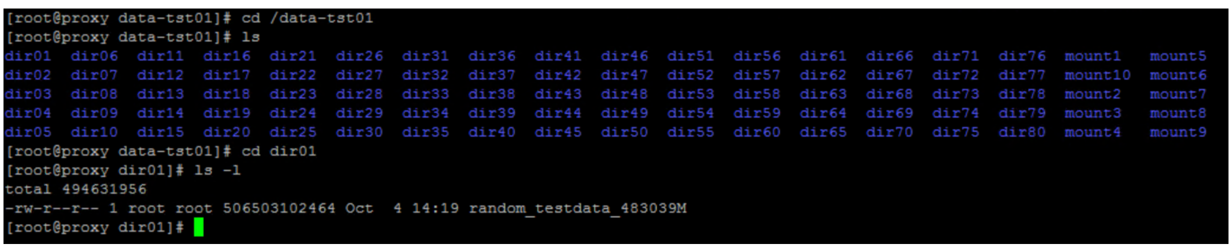

```
[root@proxy data-tst01]# cd /data-tst01
[root@proxy data-tst01]# ls
dir01  dir06  dir11  dir16  dir21  dir26  dir31  dir36  dir41  dir46  dir51  dir56  dir61  dir66  dir71  dir76  mount1   mount5
dir02  dir07  dir12  dir17  dir22  dir27  dir32  dir37  dir42  dir47  dir52  dir57  dir62  dir67  dir72  dir77  mount10  mount6
dir03  dir08  dir13  dir18  dir23  dir28  dir33  dir38  dir43  dir48  dir53  dir58  dir63  dir68  dir73  dir78  mount2   mount7
dir04  dir09  dir14  dir19  dir24  dir29  dir34  dir39  dir44  dir49  dir54  dir59  dir64  dir69  dir74  dir79  mount3   mount8
dir05  dir10  dir15  dir20  dir25  dir30  dir35  dir40  dir45  dir50  dir55  dir60  dir65  dir70  dir75  dir80  mount4   mount9
[root@proxy data-tst01]# cd dir01
[root@proxy dir01]# ls -l
total 494631956
-rw-r--r-- 1 root root 506503102464 Oct  4 14:19 random_testdata_483039M
[root@proxy dir01]#
```

Figure 11.
ODB data generated through the Epic iO simulator

Performance test

This section describes the performance results achieved for backup and restoration of the ODB with the Cohesity Linux Agent installed on an X210c Linux node and Cohesity Cloud Edition deployed across four all-NVMe X210c nodes, all installed on a Cisco UCS X-Series modular chassis.

The ODB was generated through the Epic iO simulator with a total size of 20 TiB, or about 22 TB.

Backup performance

Table 1 shows the full backup results for the ODB on a four-node all-NVMe X210c.

Table 1. Backup results for the ODB on a four-node all-NVMe X210c

Full backup of 20-TiB (22-TB) ODB	**2 hours and 11 minutes (10 TB/hr) with 4-node Cohesity cluster**

Figure 12 shows the successful full backup captured from the Cohesity dashboard.

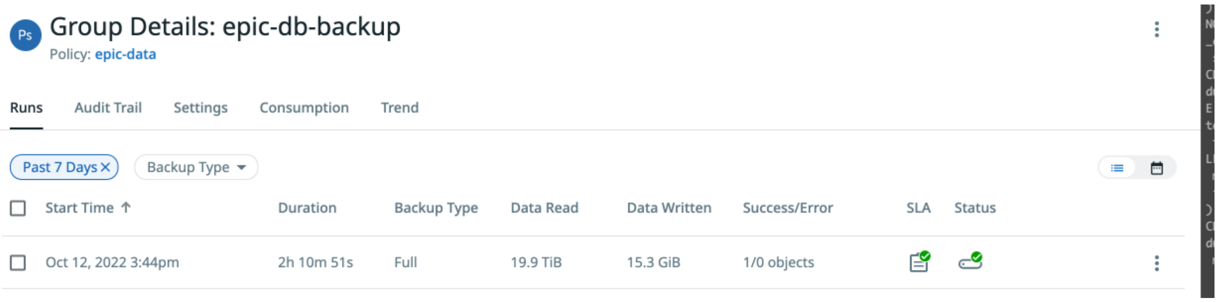

Figure 12.
Cohesity dashboard showing successful backup

During full backup of the ODB, all the read operations occur through the Cohesity Linux Agent installed on a Cisco UCS X210c node. The read operations are spread across multiple Cohesity Agent child processes, and the data is compressed, deduplicated, and written on Cohesity Data Cloud installed on a four-node all-NVMe X210c cluster.

Figure 13 shows average read operations of about 3.2 GBps on the Linux agent.

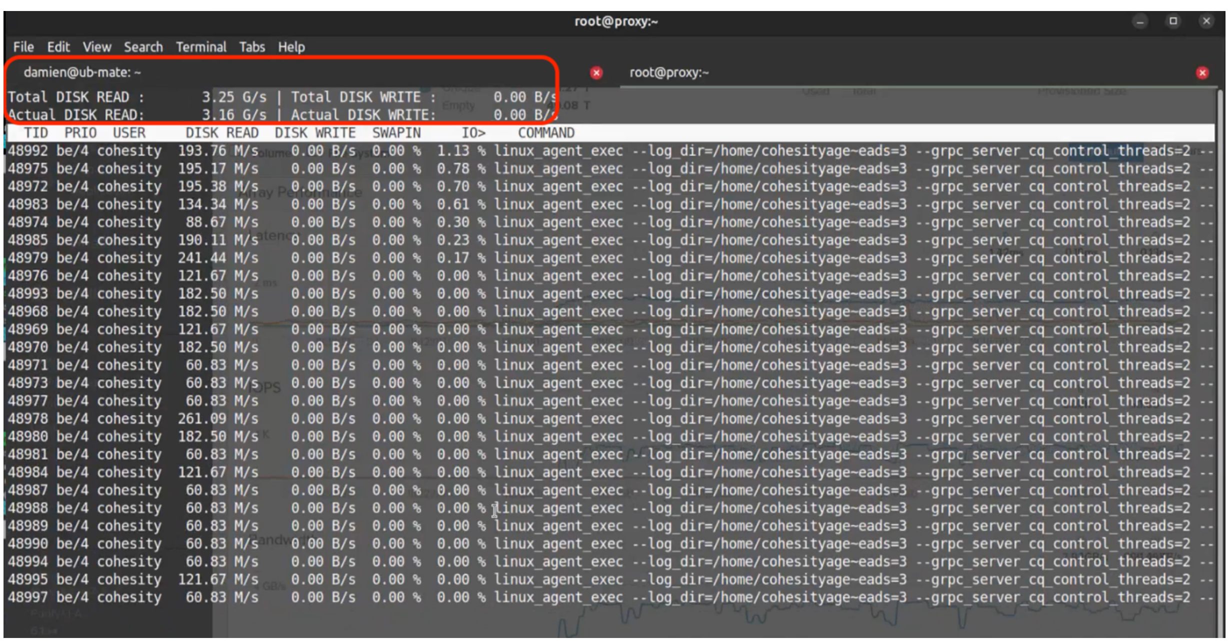

Figure 13.
Average read operations per second

Restore performance

Table 2 shows the full restore process results for the ODB from a four-node all-NVMe X210c.

Table 2. Restore process results for the ODB from a four-node all-NVMe X210c

4-node Cohesity cluster	Full restoration of 20-TiB (22-TB) ODB	**2 hours and 5 minutes (10.5 TB/hr)**

Figure 14 shows the successful full restoration captured from the Cohesity dashboard.

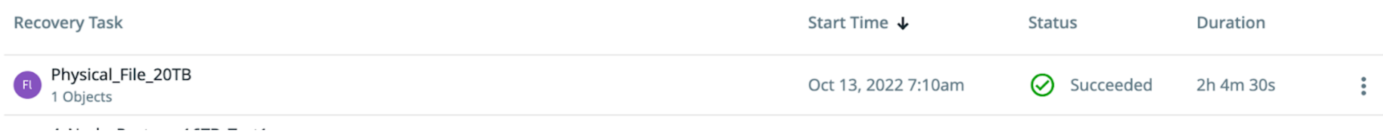

Figure 14.
Cohesity dashboard showing successful restoration

During a full restoration of the ODB, all the write operations occur in the Cohesity Linux Agent installed on the Cisco UCS X210c node. The write operations are spread across multiple Linux agent child processes and read from the Cohesity Data Cloud installed on a four-node all-NVMe X210c cluster.

Figure 15 shows average write operations of about 3.2 GBps on the Linux agent.

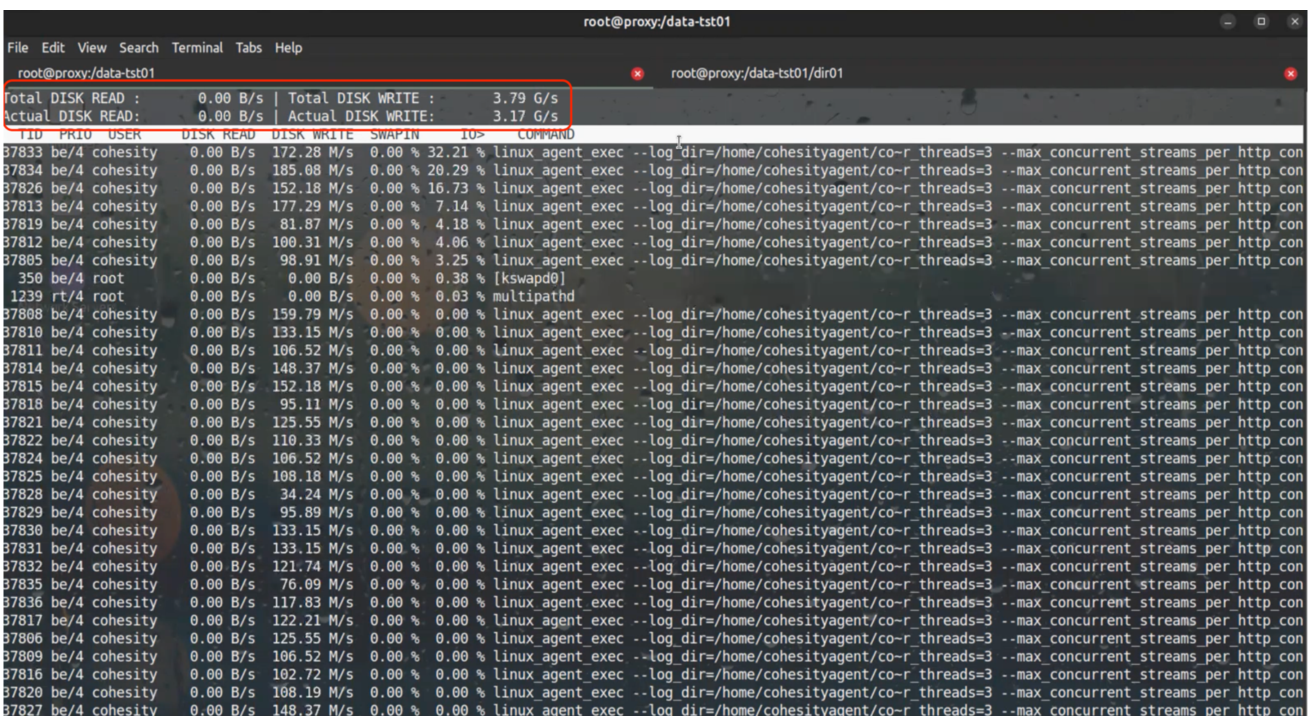

Figure 15.
Average write operations per second

Scalability

The backup and restore performance of workloads protected with Cohesity Data Cloud on Cisco UCS X-Series Modular Systems and the capacity of the Cohesity cluster scale with the addition of nodes to the existing cluster. In the present solution for protection of ODBs, a Cohesity Agent was installed on the backup host deployed on a Cisco UCS X210c node.

The performance of the backup and restore process depends primarily on three aspects of the deployed architecture protecting the database:

- How quickly can the data be read or written from the storage array? In the present setup, the speed is directly proportional to the Fibre Channel bandwidth between the host and the SAN storage array.
- How quickly can the data be transferred from the backup host (Cohesity Agent) to the Cohesity cluster? This speed is proportional to the network bandwidth between the backup host (Cohesity Agent) and the Cohesity cluster.
- How quickly can the data be written or read from the Cohesity cluster? This speed is proportional to the write and read throughput of the Cohesity cluster and on how the cluster manages the large multi-terabyte files.

Figure 16 shows these three aspects of the architecture.

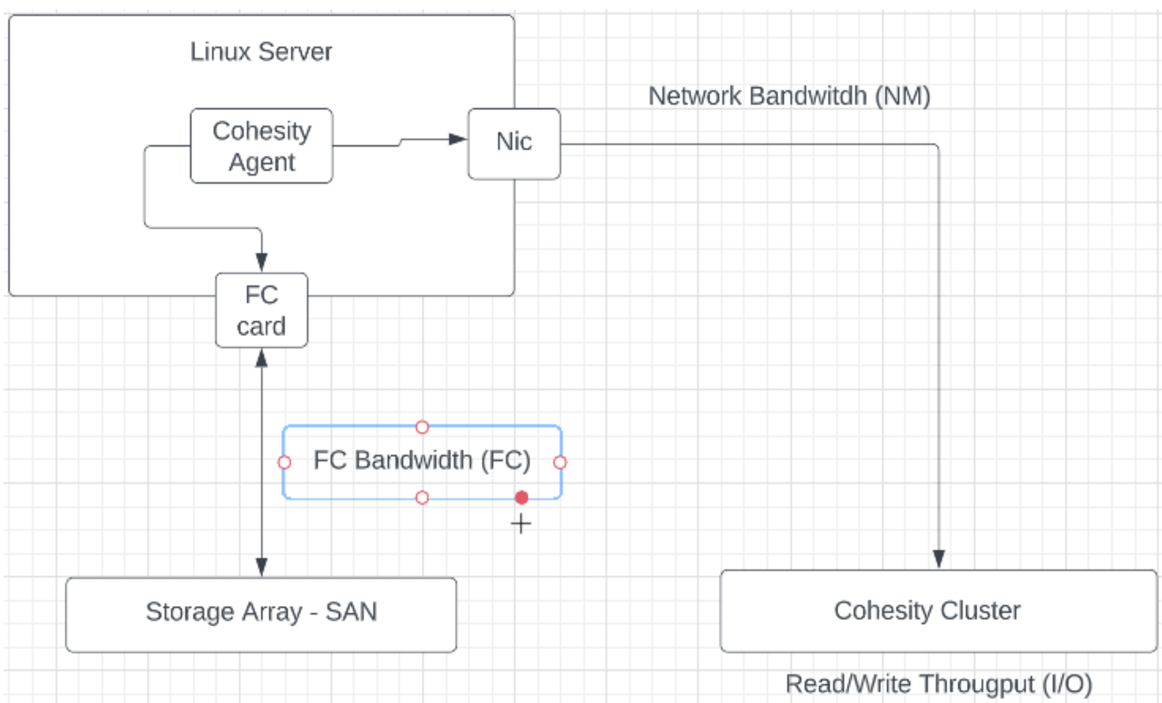

Figure 16.
Read-write throughput (I/O)

In summary, to scale the performance of the backup and restore process, you need to scale the Fibre Channel bandwidth of the LUNs mounted on the backup host (Cohesity Linux Agent) or the network bandwidth of the backup host, or you need to add more nodes to the existing Cohesity cluster.

In the present setup, backup and restore performance was analyzed with the addition of a fifth node to the existing four-node Cohesity cluster deployed on the Cisco UCS X-Series all NVMe modular system.

Table 3 shows the main results.

Table 3. Main backup and restore performance results

Database size (TiB)	Number of nodes in Cohesity cluster	Backup performance	Restore performance
20 TiB (about 22 TB)	4 x X210c all-NVMe nodes	10 TB/hour	10.5 TB/hour
20 TiB (about 22 TB)	5 x X210c all-NVMe nodes	11.24 TB/hour	11 TB/hour

The main findings from the scalability testing are listed here:

- The Fibre Channel connectivity for the backup host (Cohesity Linux Agent) to the all-flash storage array was configured with two 16 Gbps from each side of the storage array controller (A and B). This connection was the primary bottleneck that restricted backup and restore performance to about 11 TB/hour (31.3 GBps).
- The Cohesity cluster configured in the Cisco UCS X-Series was easily scalable. Because X210c nodes provide both computing and all-NVMe storage resources, the existing Cohesity cluster is scaled simultaneously with the addition of a node. Just add an X210c node to the existing X-Series system, derive a server profile, and install the Cohesity OS through the Cisco Intersight platform. Thereafter, expand the Cohesity cluster through Cohesity Helios. No cabling or network configuration is required.
- The Cisco UCS X-Series Modular System is ready for the future and easily resolves the bottleneck caused by network limitations in the backup host (Cohesity Linux Agent). The present study uses 25-GB Cisco UCS X-Series Intelligent Fabric Modules (IFMs), providing 2 x 50-GB network bandwidth to each X210c node. Customers can use a 100-GB IFM without changing any of the configurations in the existing X210c node deployed for the Cohesity cluster. The X-Series Modular System can be upgraded to a 100-GB IFM with a 100-GB Cisco virtual interface card (VIC) installed in the X210c backup host (Cohesity Linux Agent). The Cohesity nodes on the X210c configured with 4 x 25-GB Cisco VICs and Cohesity Linux Agent with 2 x 100-GB Cisco VICs can be placed in the same X-Series Modular System. This placement would lead to twice the network bandwidth for the backup host (Cohesity Linux Agent), overcoming the network contention without any alteration in the existing Cohesity cluster deployed on the Cisco UCS X-Series Modular System.
- The Cohesity Linux Agent could completely saturate and use the Fibre Channel bandwidth between the SAN storage and the backup host by performing multithreaded and parallel read and write operations across single or multiple files. The main aspects of the Cohesity Linux Agent that enabled high backup and restore performance are listed here:
    - Use of Fibre Channel bandwidth: The Cohesity Linux Agent can completely saturate and use the Fibre Channel bandwidth between the SAN storage and the backup host by performing multithreaded and parallel read and write operations across single or multiple files. The Cohesity MegaFile feature breaks any file larger than 64 GB into chunks of 16 GB to increase the parallelism and thereby reduce backup and recovery times to use the complete infrastructure fully.
    - Reduced network bandwidth utilization: The Cohesity Linux Agent reduces the use of network bandwidth by computing SHA1 fingerprints and transferring the data to the backup system during incremental backup only if those fingerprints are not present. Typically, a 4 to 5x reduction is observed in the data sent over the network during incremental backups. The backup daily window and overall network consumption are reduced without increasing the resources on the backup host, such as its CPU or memory requirements, because the SHA1 fingerprinting process is scalable and lightweight.

- Read and write throughput: The Cohesity file system performs variable block-level deduplication and reduces the data even further and checks across the entire data set in the cluster. Because Cohesity has a scale-out architecture, you can seamlessly scale the number of nodes on the cluster as your needs dictate. The Cohesity MegaFile feature provides similar read and write throughput regardless of whether you have a single large multi-terabyte file or multiple files.

Failure testing

Another important aspect of a data protection solution is the resiliency of the configured system. To test the resiliency of the solution, customers should know the effects on backup and restore operations during failure of a node configured across a distributed Cohesity cluster on a Cisco UCS X-Series Modular System. Because restoration is crucial to data protection, the following steps were performed to test a failure during restoration of ODBs:

1. The four-node Cohesity cluster on the Cisco UCS X-Series system was expanded to five nodes. The process was simple: insert another X210c node in the Cisco UCS X-Series system, clone the node, attach a server profile, install the Cohesity OS through the Cisco Intersight platform, and add the new node to the existing cluster. Because a single Cisco UCS X-Series system supports up to eight all-NVMe nodes, no cabling was required. Moreover, to expand beyond eight nodes, customers could install a new Cisco UCS X-Series system under the same fabric interconnect domain.

2. A restore job for the ODB was started.

3. After the restore process was in a steady state, one of the nodes was shut down.

4. After the node shutdown, performance and any failure of the restore process were monitored.

Figure 17 shows a five-node Cohesity cluster on the Cisco UCS X-Series Modular System.

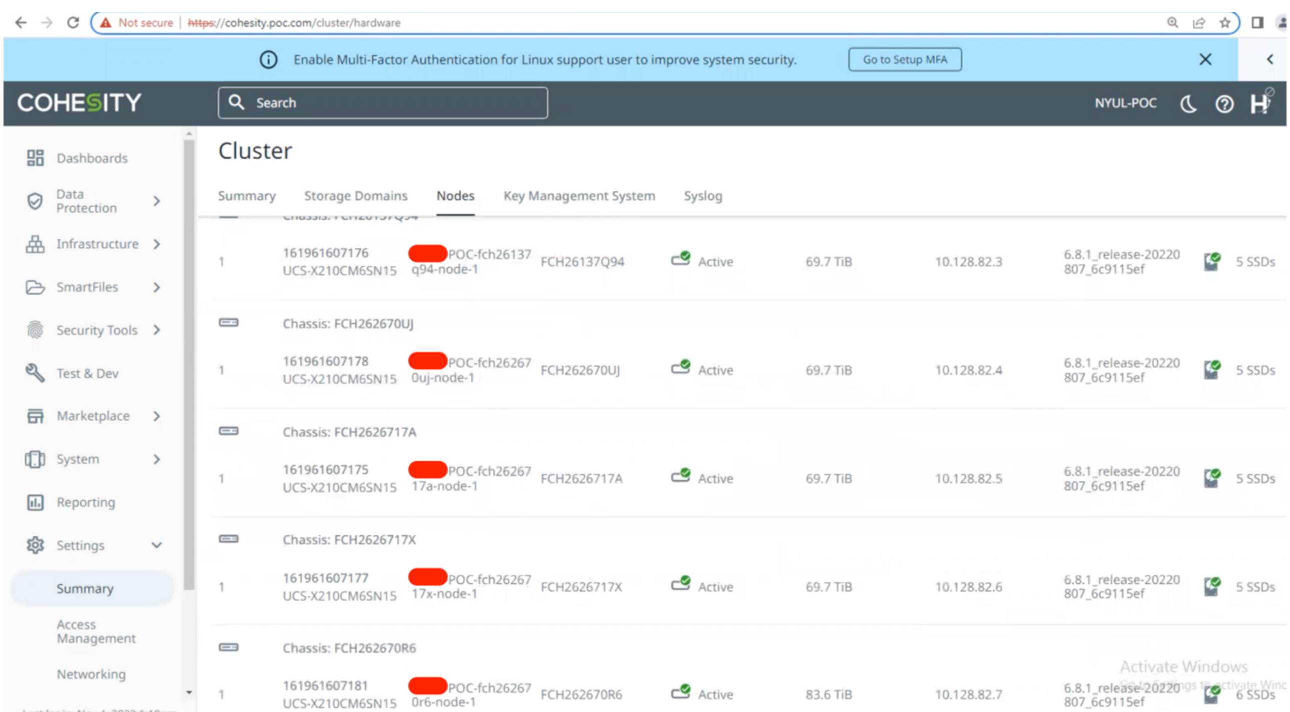

Figure 17.
Five-node Cohesity cluster on the Cisco UCS X-Series Modular System

Figures 18 and 19 show the shutdown of one of the X210c nodes in the Cohesity cluster.

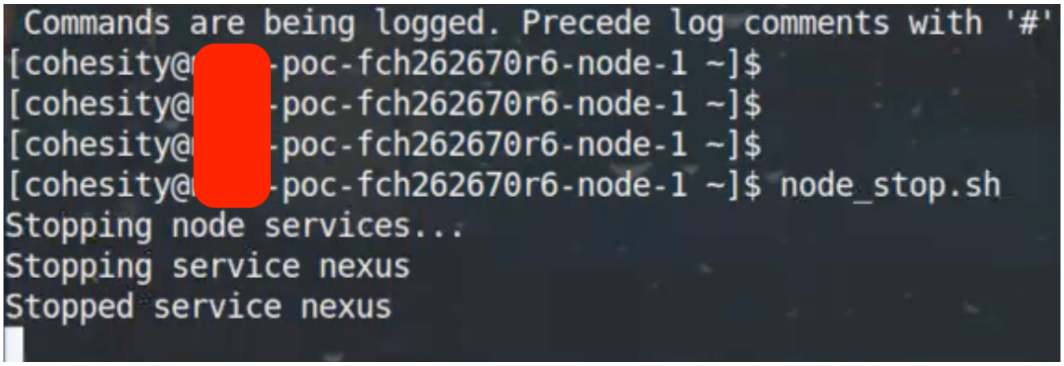

Figure 18.
Shutdown of one of the X210c nodes in the Cohesity cluster

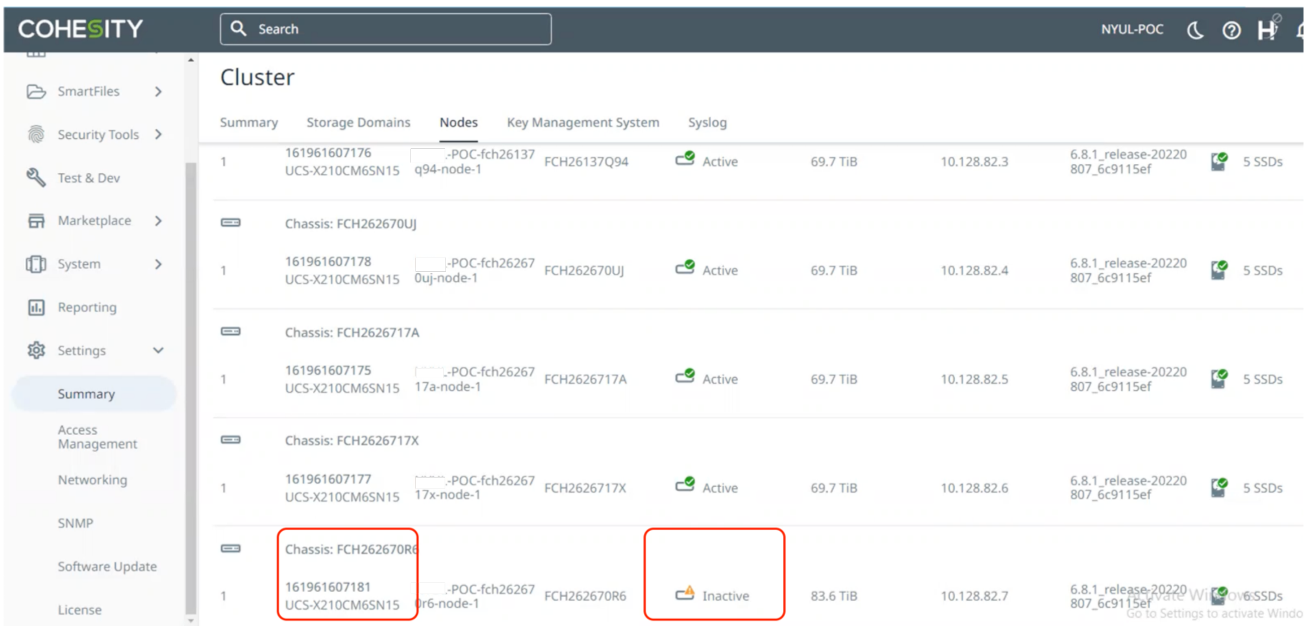

Figure 19.
One of the nodes is shut down

The Cohesity Linux Agent writes the restored data to the Fibre Channel LUNs mounted from the storage array. The write operations during the node failure process were monitored on this storage array.

No write failures were seen during the entire restore process before or after the node was shut down in the Cohesity cluster.

Figure 20 shows the write operations during the node failure as measured on the storage array.

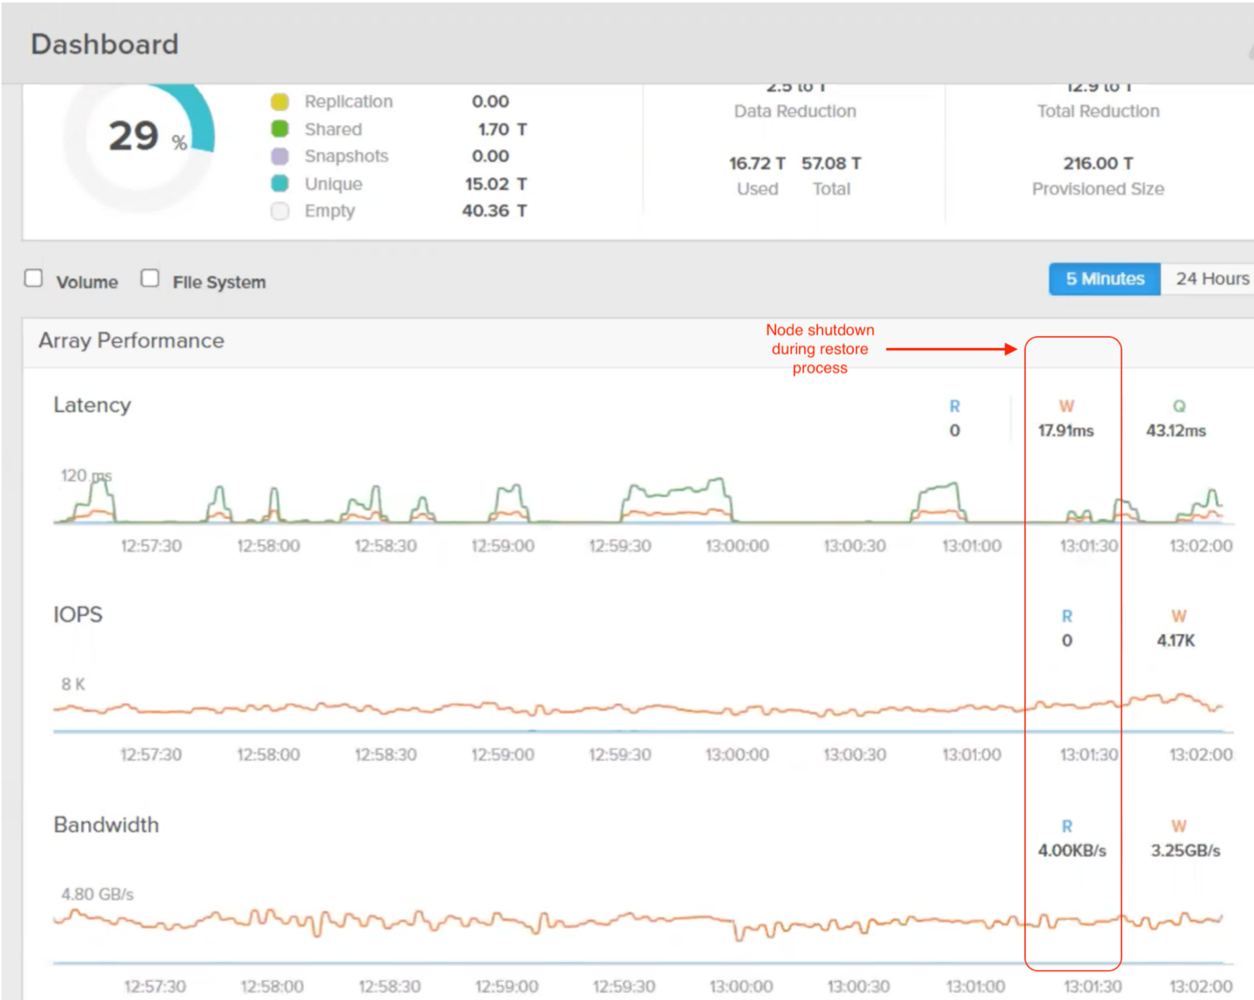

Figure 20.
Write operations during node failure

For more information

For additional information, see the following resources:

- Cisco Intersight Help Center: https://intersight.com/help/saas/home
- Cisco UCS X-Series product installation guide: https://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-x-series-modular-system/products-installation-guides-list.html
- Cohesity on Cisco
  - https://www.cisco.com/c/en/us/solutions/global-partners/cohesity.html
  - https://www.cohesity.com/solutions/technology-partners/cisco/
- Cohesity on Cisco UCS X-Series installation guide: https://docs.cohesity.com/hardware/PDFs/SetupGuideCiscoXseries.pdf
- Ansible automation for Cohesity server profile for Cisco UCS X-Series: https://developer.cisco.com/codeexchange/github/repo/ucs-compute-solutions/cohesity_xseries_ansible

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 **https://www.cisco.com/go/offices**.

Printed in USA 223053.3 04/23