

CAI2 – Simplified Compute Operations

For the AI ready data center

Chris Merkel – Solutions Engineer – DC Architecture



Agenda

1. Intro
2. Platform Overview
3. Intersight Operations
4. Closing

Platform Overview

Cisco UCS Compute Portfolio – 2009

MAINSTREAM ENTERPRISE SERVERS

UCS 5108 Chassis



UCS B200 M1



UCS C200 M1
4 LFF



UCS C210 M1
16 SFF



UCS C250 M1
8 SFF – Memory Expand



AI SERVERS




Cisco UCS Compute Portfolio

MAINSTREAM ENTERPRISE SERVERS

UCS X-Series
X9508 Chassis

IFM Module



UCS X-Series Direct



UCS X210c M7



UCS X210c M8



UCS X410c M7



UCS X215c M8



New

New

New

UCS C240 M8E3S
36 EDSFF E3.S1T



New

UCS C240 M8SX
28 HDD/SDD/NVMe



New

UCS C240 M8L
16 LFF + 4 SFF



New

UCS C240 M7SN
28 NVMe



UCS C220 M8E3S
16 EDSFF E3.S1T



New

UCS C220 M8S
10 HDD/SSD/NVMe



New

UCS C220 M7N
10 NVMe



UCS C245 M8SX
28 HDD/SDD



New

UCS C225 M8S
10 HDD/SSD



New

UCS C225 M8N
10 NVMe




New

VAST

AI SERVERS

UCS C885A M8
8RU Dense GPU Server



New

UCS C845A M8
4RU MGX Server



New

EDGE COMPUTING

Coming Late 2025!

Agenda

1. Intro
2. Platform Overview
3. Intersight AI operations
4. Closing

X-Series Portfolio

COMPUTE

X210c Compute Node

- 2-Socket, single slot servers
- **Three Generations: M6*, M7 and M8**
- Intel 3rd Gen (Ice Lake) 4th Gen (Sapphire Rapids) 5th Gen (Emerald Rapids) and 6th gen (Granite Rapids) Xeon CPUs



X410c Compute Node

- 4-Socket, dual slot servers
- Intel 4th Gen Xeon CPU
- Up to 64 DDR5 DIMMs



X215c Compute Node

- 2-Socket, single slot servers
- M8 with AMD 4th/5th gen EPYC CPU



* End of sale

FABRIC

4th,5th and 6th Gen FI

- 25/100G ports
- Unified ports: 32G FC 64XX, 6536). 64G FC (66xx)
- Supports VIC 1400, 14000 and 15000 series



UCS X-Series Direct

- Scale at the edge with X-series advantage for 1-16 servers



25/100G IFM

- 8 x 25/100G connectivity



4th and 5th Gen VIC

- 25/100G connectivity for both blades and racks



X-FABRIC AND PCIE NODE

X-Fabric

- Based on native PCIe Gen 4
- Provides GPU acceleration to enterprise application
- No backplane or cables = Easy upgrades



GPU Node and Front Mezz GPUs

- Nvidia A16, Nvidia L40S, Nvidia L4, AMD MI210
- and Nvidia H100-NVL
- GPUs today in various configurations




6th Generation UCS Fabric Interconnects

Redefining Compute Connectivity

- Simplified rack and modular server transition to 64G Fibre Channel
- Densest 25/50G or 100G converged fabric POD for x86 compute
- Comprehensive storage protocol support for compute
- Cloud scale management and visibility



UCS Manager



Cisco UCS X210c M8

2S Compute Node for a wide range of workloads, including virtualization, web, collaboration, cloud, and bare-metal applications

Up to 172 Cores

2 x Intel® Xeon®6 Scalable Processors
6500P/6700P CPUs

Up to 8TB Memory

32x 6400 MT/s DDR5 Up to 256G Per DIMM
Support for MRDIMMs (up to 8000MT/s)



Up to 9 Drives

Up to 9 E3.S drives, or 6
SAS/SATA/NVMe (tri-mode)

25G - 100G
mLOM/Mezz
and M.2 support

GPU Support with X440p & Front Mezz

Nvidia A16, Nvidia L40S, Nvidia
L4, Nvidia H100 NVL

Cisco UCS X215c M8

2S Compute Node for a wide range of workloads, including virtualization, web, collaboration, cloud, and bare-metal applications (DB and analytics)

Up to 320 Cores

2 x AMD 4th/5th Gen EPYC CPUs

Up to 6TB Memory

24x 6000 MT/s DDR5 Up to 256G Per DIMM



Up to 6 Drives

Up to 6 SAS/SATA/NVMe (tri-mode)

**25G - 100G
mLOM/Mezz
and M.2 support**

GPU Support with X440p & Front Mezz*

Nvidia A16, Nvidia L40S, Nvidia L4*, Nvidia H100 NVL

UCS X-Series Direct: Expansion!

FCS RELEASE (Q3CY24)

Up to 8 compute nodes


- X210c M6, X210c M7, X215c M8
- X410c M7
- Support of X-Fabric + X440p
- 4th & 5th Gen VIC
- Full config support of X-series compute Node



Now Available

Up to 16 X-Series compute nodes

- IFM-100G for 2nd Chassis
- Plus: Up to 4 rack server



AI Specific Compute

Cisco Compute AI Portfolio

Address AI workloads with visibility, consistency, and control

Validated solutions for AI with compute, network, storage, and software

Build the model

Training

Optimize the model

Fine-tuning and RAG

Use the model

Inferencing



UCS Dense GPU Servers

Dense compute for demanding AI



UCS Blade (w/GPU Expansion) and Rack

Full stack AI with compute and networking

Cisco UCS C845A Versatile AI Server

Optimized for GenAI


Address a variety of AI use cases
Modular NVIDIA MGX™ architecture enables configuration flexibility
Built on air-cooled enterprise rack design that easily fits in your data center

Scalable AI performance


Choose the number of GPUs that match your use case
Increase the number of GPUs when your workloads demand it
Easily create clusters of AI servers to scale out and meet increasing needs

Consistent management

Cisco Intersight delivers a unified management paradigm across all your UCS servers
Manage your AI server with the same tool as your traditional servers



UCS C845A



2/4/6/8



H100 NVL,
H200 NVL, RTX 6000 Pro, L40S GPUs
350W/400W/600W GPUs


or

2/4/6/8




AMD
MI210 GPUs
300W/GPU

UCS C845A



UCS C885A M8

Massive performance for AI at scale



HGX 8RU 19" EIA Rack
8x Nvidia H100/H200/, AMD MI300X/MI350X GPUs
AMD Genoa 400W/Turin 500W TDP
8x PCIe5 x16 HHHL + 5x PCIe5 x16 FHHL NICs


TRAINING - LARGE / MEDIUM / SMALL MODELS
FINE TUNING
LARGE MODEL INFERENCE
RAG

- Service Providers
- Financial Services
- Manufacturing

- Healthcare and Life Sciences
- Automotive




C885A M8 Specifications



Product Specifications	
Form Factor	• HGX 8U 19" EIA Rack
Compute + Memory	• 2 AMD SP5 Genoa/Turin, 96 Core, TDP 400W, Up to 3.7 GHz • 24 DDR5 RDIMMs • Up to 6,000 MT/S
Storage	• 1 PCIe3 x4 M.2 NVMe (Boot Device) • 16 PCIe5 x4 2.5" U.2 NVMe SSD (Data Cache)
GPU	• 8x H100 or 8x H200 • 8x MI300X or 8x MI350X
Network Cards	• 8 PCIe5 x16 HHHL for E-W NIC ConnectX-7, BF3 B3140H • 5 PCIe5 x16 FHHL for N-S NIC BF3 B3220, B3240 • 2 OCP 3.0 SFF
Front IO	• 2 USB 2.0, 1 ID BTN, 1 Power Button
Rear IO	• 1 USB 3.0 A, 1 USB 3.0 C, mDP, 1 ID BTN, 1 Power Button, 1 USB 2.0 C (for debugging), 1 RJ45 (mgmt.)
Power Supply	• Up to 6 54V 3kW and 2 12V 2.4kW, N+1 redundancy

AI Platform Considerations: UCS GPU Options

GPUs are subject to change as new peripherals are added to the portfolio



* NOTE: GPU Form Factor and GPU model support may vary between AMD and Intel Platforms (i.e. c220/c225, c240/c245, and x210c/x215c). Check the spec sheet for each platform to determine maximum GPU support based on GPU selection

AI Infrastructure requirements by workload type


Workload type	Small* model inferencing / RAG	Large* model inferencing / RAG	Fine Tuning	Small* model training	Large* model training
AI Accelerator (GPU)	Maybe 0-2 GPU per server	Yes 2-4 GPU per server	Yes 4-8 GPU per server	Yes 2-4 GPU per server	Yes 8 GPU per server
Cluster size	0 – 100's Scales horizontally with request rate	0 – 100's Scales horizontally with request rate	20-80 GPUs ~10's of servers	20-80 GPUs ~10's of servers	>1000 GPUs 100s of servers
Intra-host GPU interconnect (e.g. NVLink)	No	Maybe**	Maybe	Yes	Yes
Inter-host GPU interconnect (“Backend” network)	No	Probably Not**	Maybe 200-400G RoCE	Yes 200-400G RoCE	Yes 200-400G RoCE/Infiniband
“Frontend” network	“Any”	100-400G	100-400G	100-400G	100-400G

* “Small” and “Large” here are relative to the size of the GPU. i.e. Small can fit in memory of a single GPU.

** Depends on parallelism strategy, e.g. Tensor Parallel, Pipeline Parallel, etc.

UCS-X GPU Expansion


2nd Gen X580p PCIe Node and X9516 X-Fabric



UCS X-Fabric Technology with PCIe Node

- ✓ PCIe Switching with PCIe Gen 5 connectivity
- ✓ 4x FHFL or HHHH GPUs per PCIe node
- ✓ Intra-host GPU interconnect with NVLink
- ✓ Intersight policy-based Management
- ✓ Inter-host scaling with RDMA over AI Fabric


UCS X580p PCIe Node



- Double wide PCIe node for 4x FHFL GPU and PCIe G5 GPU support
 - Nvidia H200-NVL, RTX PRO 6000 & L40S
- Support multiple vendors: Nvidia, AMD*/Intel*
- NVLink bridge support
- Support up to 600W FHFL GPU
- Managed PCIe node with BMC support
- Policy based GPU management
- Ability to share GPUs across two Compute nodes

* AMD & Intel GPUs support will be post FCS

UCS X9516 X-Fabric




- PCIe Gen5 Switching
- 2x CEM Slots to support HHL NIC cards
 - ConnectX 7 (2x 200GB & 1x 400G)
- Managed XFM Modules with BMC support
- GPU Direct Support over RDMA
- GPU Backend (East-West Traffic) network support

AI Cluster Expansion

GPU-to-GPU connectivity


with XFM external ports

- X-Fabric Module with Gen5 PCIe switch
- SuperNIC Adapter for GPU East-to-West traffic
- 1 or 2 external ethernet ports based on adapter




X580p – GPUs 1/3 Mapped to Server 1 and GPUs 2/4 Mapped to Server 2

(1x NIC mapped to each server)




X580p – All GPUs Allocated to Server 2 w/NVL Bridge

(2x NICs mapped to one server)




Server 1: No GPU
Server 2: 4x GPUs
2x NICs

Note: X580p default in 3/4 or 7/8 and is not supported in slots 2/3 and 6/7. It can be installed in slots 1/2 or 5/6 as well.




X580p – 2x GPUs Allocated to Server 1 and 2x 2 w/NVL Bridge

(1x NIC mapped to each server)



X580p – GPUs 1/3 Mapped to Server 1 and GPUs 2/4 Mapped to Server 2

(1x NIC mapped to each server)





Server 1: 2x GPUs
1x NIC

Server 2: 2x GPUs
1x NIC


Note: X580p default in 3/4 or 7/8 and is not supported in slots 2/3 and 6/7. It can be installed in slots 1/2 or 5/6 as well.

Note: Mixed GPUs in riser support will be post-FCS



Platform Support

X580p & X9516



Compute Node

X210c M7 – EMR – **planning**



X215c M8 – AMD



X210c M8 – Intel



FI & IFM

4th , 5th & 6th Gen FI



4th & 5th Gen IFM



X-Direct Support – planning

PCIe Gen5 X-Fab
Mezz



Management



Intersight SaaS –
At FCS

**Appliance – PVA &
CVA –CY26**

No UCSM Support

Scalable Management

for the AI ready data center

Cisco Intersight


Unified Operating Model

Consistent operational model globally, from DC to edge, at cloud scale


Simplified operation with AI-driven capabilities including Connected TAC, Proactive RMA, and Predictive Insight

Secure operations with built-in advisories and continuous risk mitigation

Automate deployments, configuration, workflows, and day-0 to day-N tasks



Supported Solutions




SaaS US or EMEA Supported Systems
Appliance US or EMEA Supported Systems

Connecting Cisco solutions to Intersight (SaaS example)

Device Connector is built-in:

- Embedded in target management plane
- Delivered with management plane firmware



Device Connector connection:

- Highly secure
- Outbound single-destination
HTTPS 443
- Durable websocket enabling bi-directional communication

***Optional proxy config available**

Connecting 3rd party solutions to Intersight

- Assist via mode selection of Intersight Appliance OVA
- Manages connections to 3rd party solutions via APIs through this device

Claim a New Target

Claim VMware vCenter Target

To claim any on-premises target an Assist Appliance is required. Deploy and claim an Assist Appliance if needed before claiming the target

Assist * ⓘ

houlab-intersight-assist.houston.ciscolabs.com

Hostname/IP Address

Hostname/IP Address

*Assist functionality is also part of the Connected and Private appliance model

Visibility

- ☰
-  **Dashboards**

 Operate

Servers

Chassis

Fabric Interconnects

Nutanix Clusters

Networking

HyperFlex Clusters

Storage

Virtualization

Integrated Systems

 Configure

 Analyze

 Automate

Workflows

Overview  Contracts and Licensing Storage Power & Energy Metrics **New** Fabric Interconnects Metrics **New** Servers Metrics **New** Nutanix Clusters +

 Add Filter

Add Widget

Server Health Summary ⓘ



22
Healthy



2
Critical



3
Warning



6
Healthy

Server Inventory ⓘ



Series

- C 8
- X 6
- B 5
- HX 3
- Redfish 2

12
Unassociated

21
Standalone

3
Managed

Nutanix Cluster Hypervisor Types ⓘ



Types

- ESX 1
- AHV 1

Fabric Interconnect Inventory ⓘ



Models

- 6454 4
- 6332 2

280
Ports

34
Used

246
Available

Server Model Summary ⓘ



Models

- C240 M5SX 8
- UCSX 210C-M6 5

Nutanix Cluster Nodes ⓘ



Cluster size

Fabric Interconnects Bundle Version Summary ⓘ



Versions

© 2025 Cisco and/or its affiliates. All rights reserved.
Cleveland Engage 10/2/2025

40



- 
-  **Dashboards**

 Operate

Servers

Chassis

Fabric Interconnects

Nutanix Clusters

Networking

HyperFlex Clusters

Storage

Virtualization

Integrated Systems

 Configure

 Analyze

 Automate

Workflows

Overview

Contracts and Licensing

Storage

Power & Energy Metrics **New**

Fabric Interconnects Metrics **New**

Servers Metrics **New**

Nutanix Clusters

+



Add Widget

Nutanix Cluster Storage Utilization - Top 5

#	Name	Capac...	Utilization
1	houlab-ntnx-...	37.37 ...	1.4%
2	houlab-ntnx-...	37.37 ...	0.1%

Top 5 Storage Arrays by Capacity Utilization

#	Name	Vendor	Capacity	Utilizati...
1	houlab-pure	Pure Storage	11.17 TiB	4.8%
2	houlab-netapp	NetApp	27.80 TiB	4.2%

Top 5 HyperFlex Clusters by Storage Utilization

#	Name	Capac...	Utilization
1	houlab-hx-ed...	7.23 TiB	4.0%

Nutanix Cluster AOS Versions

Versions

3

Total

6.10 2

2024.3 1

Storage Version Summary

Versions

2

Total

6.3.18 1

NetApp ONTAP 9.12.1RC1P1 1

Top 5 Storage Volumes by Capacity Utilization

#	Name	Vendor	Capacity	Utilizati...
1	houlab_esxi_1	NetApp	2.00 TiB	46.3%
2	houlab-linux-...	Pure Storage	1.00 TiB	32.8%
3	houlab-esxi_1	Pure Storage	4.00 TiB	5.1%
4	houlab_esxi_i...	NetApp	2.00 TiB	4.1%

© 2025 Cisco and/or its affiliates. All rights reserved.
Cleveland Engage 10/2/2025

41



Proactive RMA - example

Servers

All Servers

Health Critical

Add Filter

Export

1 items found

50 per page

1 of 1

Health

1 Critical 1

Power

On 1

HCL Status

Not Listed 1

Bundle Version

1 4.2(3d) 1

Firmware Version

1 4.2(3d) 1

Models

1 B200 M5 1

Configuration

No

Name	Health	UCS ...	S...	H.	U.	L..	F.
houlab-ucs-3-2-1	Critical	68.0	houlab-uc...	h... ✓	100.0	Not List	Adv... 4. ...

Critical

Alarms


1

MemoryUnitUncorrectableError May 2, 2024 10:20 AM

Memory unit houlab-ucs-3/chassis-2/server-1/board/memory-array/memory-unit-DIMM_E1 has encountered an uncorrectable ECC error

Policy Based Configuration

Policy for consistent configuration



Intersight PCIe Connectivity Policy Creation

Note: Subject to Change

Intersight

Q Search

5

2

← Policies

Select Policy Type

Filters

Platform Type

All

UCS Server

UCS Domain

UCS Chassis

HyperFlex Cluster

Kubernetes Cluster

Q Search

Adapter Configuration

BIOS

Boot Order

Certificate Management

Device Connector

Drive Security

Ethernet Adapter

Ethernet Network

Ethernet Network Control

Ethernet Network Group

Ethernet QoS

FC Zone

Fibre Channel Adapter

Fibre Channel Network

Fibre Channel QoS

Firmware

IMC Access

IPMI Over LAN

iSCSI Adapter

iSCSI Boot

iSCSI Static Target

LAN Connectivity

LDAP

Local User

Memory

Network Connectivity

NTP

Persistent Memory

Power

SAN Connectivity

Scrub

SD Card

PCIe Connectivity

Serial Over LAN

Serial Pool Qualification

SMTP

SNMP

SSH

Storage

Syslog

Thermal

Virtual KVM

Virtual Media

↑

Addition of New Policy

Click for to start

↓

Cancel

Start

© 2025 Cisco and/or its affiliates. All rights reserved.
Cleveland Engage 10/2/2025

Intersight PCIe Policy Details

Note: Subject to Change

Intersight

Q Search

5

2

Dashboards

Operate

Configure

Profiles

Templates

Policies

Pools

Analyze

Automate

Optimize

System

Settings

Policies > PCIe Connectivity

Create

1 General

2 Policy Details

General

Add a name, description, and tag for the policy.

Organization *

default

Name *

PCIe-policy

Set Tags

Enter a tag in the key:value format.

Description

Description

0 / 1024

<

Cancel

Next

© 2025 Cisco and/or its affiliates. All rights reserved.
Cleveland Engage 10/2/2025

Note: Subject to Change



Intersight UCS Server Profile Configuration

Note: Subject to Change

Intersight

Search

Refresh

Close

Alerts

Help

Profile

Menu

Dashboards

Operate

Configure

Profiles

Templates

Policies

Pools

Analyze

Automate

Optimize

System

Settings

Policies > PCIe Connectivity

Create

Add End Point Mapping

When creating an endpoint mapping, the CPU socket ID may be set to 'Any', 1, or 2; the GPU quantity must be 1 or 2; and the SmartNIC quantity must be 1 or kept as disabled, while the models for both GPUs and SmartNICs may be set as 'Any' or to a specific type.

CPU

Socket ID

Socket ID

Any

1

2

GPU Model

Any

SmartNICs

Cancel

Add

© 2025 Cisco and/or its affiliates. All rights reserved.
Cleveland Engage 10/2/2025

Intersight UCS Server Profile Configuration

Note: Subject to Change

Intersight

Search

5

2

Dashboard

Operate

Configure

Profiles

Templates

Policies

Pools

Analyze

Automate

Optimize

System

Settings

Policies > PCIe Connectivity

Create

Add End Point Mapping

When creating an endpoint mapping, the CPU allocation may be set to 'Any', 1, or 2; the GPU quantity must be 1 or 2; and the SmartNIC quantity must be 1 or kept as disabled, while the models for both GPUs and SmartNICs may be set as 'Any' or to a specific type.

CPU

Socket ID

Any

GPU

Quantity

2

GPU Model

Any

SmartNIC

Quantity

1

SmartNIC Model

Any

Cancel

Add

© 2025 Cisco and/or its affiliates. All rights reserved.
Cleveland Engage 10/2/2025

Intersight UCS Server Profile Configuration

Note: Subject to Change

Intersight

Search

5

2

Dashboards

Operate

Configure

Profiles

Templates

Policies

Pools

Analyze

Automate

Optimize

System

Settings

Policies > PCIe Connectivity

Create

General

Policy Details

Add Mapping

A maximum of two mappings are supported. To add a new one, please edit or delete an existing mapping.

Search

Filters

2 Results

☐ CPUs	☐ GPUs	☐ GPU Model	☐ SmartNICs	☐ SmartNIC Model	
☐ Any	2	Any	0	Any	...
☐ Any	2	Any	1	Any	...

Rows per page 10 < 1 >

Cancel


Back

Create

© 2025 Cisco and/or its affiliates. All rights reserved.
Cleveland Engage 10/2/2025

Intersight UCS Server Profile Configuration

Note: Subject to Change



Day 2 Monitoring

HCL (Hardware Compatibility List)

← Servers

houlab-ucs-2-1-2

General Inventory UCS

General Inventory UCS Server Prof

Details

HCL Status

⚠ Not Listed

Get Recommended Drivers

Details

HCL Status

⚠ Not Listed

Get Recommended Drivers

Get Recommended Drivers

OS Vendor *

VMware (Current)

OS Version

ESXi 7.0

Recommended Drivers

Download Driver ISO

Add Filter

Model	Firmware Version	Drive...
UCSX-V4-Q25G...	5.3(1.230031)	nfnic
UCSX-V4-Q25G...	5.3(1.230031)	nenic
UCSX-V4-Q25G...	5.3(1.230031)	neni...
UCSX-X10C-RAI...	52.23.0-4782	lsi_...

OS Vendor *

VMware (Current)

CentOS

Citrix

Microsoft

Nutanix

Oracle

Red Hat

Rocky Linux

SuSE

Ubuntu

VMware (Current)

Fabric Interconnects

Nutanix Clusters

Networking

HyperFlex Clusters

Storage

Virtualization

Integrated Systems

 Configure

 Analyze

Explorer New

 Automate

 System

 Settings

New Command Palette ×

Navigate Intersight with Ctrl+K

Explorer

View Explorations

Save

-  This feature is available for Intersight Managed Domains and supports selected sustainability metrics for UCSM Managed Domains. Data is collected at 1 minute for the Advantage license tier in Intersight Managed Domains only. For more information, see [Data Collection Supported Devices](#).
-  Tour Explorer anytime. In the header, click **Help > Tour Explorer**. [Tour Explorer Now](#).

Add Metrics

Metrics

Metrics* ⓘ

Select Metric

INSTRUMENT NAME (11)

Electric Current >

Fan >

Host Power And Status >

Memory Module >

Network Interface >

Power Supply >

Signal Power >

METRICS (57)

Link Status >

Network Bytes >

Network Carrier Sense Errors >

Network Interface Resets >

Network Jabber Errors >

Chart Type

Line

Time Interval

Last 24 Hours

Granularity

10 min



No items available

Display Options

Query Code

Result

Cancel

Select

© 2025 Cisco and/or its affiliates. All rights reserved.
Cleveland Engage 10/2/2025

54



▼

Add Filter

Add Widget

Top 5 Uplink Ports by Peak% RX Bandwidth Utilization

🗕 ⓘ 📅 Last 7D ⋮

Port	Device	Peak Rx	Limit	Peak% Rx ↓
1/48	houlab-ucs-2 ...	3.9 Gbps	10.0 Gbps	38.8%
213	houlab-ucs-2 ...	3.9 Gbps	20.0 Gbps	19.4%
1/47	houlab-ucs-2 ...	814.4 Mbps	10.0 Gbps	8.1%
1/48	houlab-ucs-3 ...	428.8 Mbps	10.0 Gbps	4.3%
216	houlab-ucs-3 ...	833.9 Mbps	20.0 Gbps	4.2%

Top 5 Server Ports by Peak% RX Bandwidth Utilization

🗕 ⓘ 📅 Last 7D ⋮

Port	Device	Peak Rx	Limit	Peak% Rx ↓
1/9	houlab-ucs-2 ...	7.2 Gbps	25.0 Gbps	28.8%
1/10	houlab-ucs-2 ...	5.9 Gbps	25.0 Gbps	23.5%
1025	houlab-ucs-2 ...	11.7 Gbps	50.0 Gbps	23.4%
1/9	houlab-ucs-2 ...	4.6 Gbps	25.0 Gbps	18.4%
1153	houlab-ucs-2 ...	4.6 Gbps	50.0 Gbps	9.2%


Top 5 Port Channels by Peak% RX Bandwidth Utilization

🗕 ⓘ 📅 Last 7D ⋮

Port	Device	Peak Rx	Limit	Peak% Rx ↓
1025	houlab-ucs-2 ...	11.7 Gbps	50.0 Gbps	23.4%
1282	houlab-ucs-2 ...	11.6 Gbps	50.0 Gbps	23.2%
213	houlab-ucs-2 ...	3.9 Gbps	20.0 Gbps	19.4%
1280	houlab-ucs-2 ...	7.2 Gbps	50.0 Gbps	14.4%
1153	houlab-ucs-2 ...	4.6 Gbps	50.0 Gbps	9.2%

Count of Errors by Server Ports ⓘ

Good news — there are no errors
We have checked for IMM domain error types



Closing

1. UCS offers the flexible, policy-based compute for any AI deployment
2. Intersight provides a single place to configure and manage all Cisco Compute
3. Intersight integrates with 3rd party platforms to give a complete view of the DC
4. Intersight offers expanded configuration capabilities through its automation capabilities

Thank you



