
cisco Live!
ALL IN
#CiscoLive

The bridge to possible

# UCS X-Series

## Bring your applications together

Scott Garee
Senior UCS Platform TME
BRKDCN-2794

cisco Live!

#CiscoLive

# Cisco Webex App

## Questions?

Use Cisco Webex App to chat
with the speaker after the session

## How

1. Find this session in the Cisco Live Mobile App
2. Click “Join the Discussion”
3. Install the Webex App or go directly to the Webex space
4. Enter messages/questions in the Webex space

## Webex spaces will be moderated by the speaker until June 17, 2022.

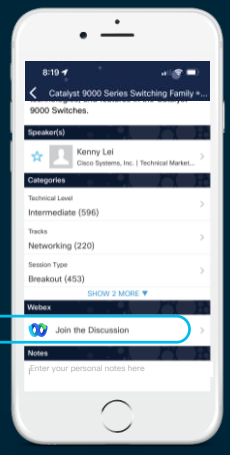


https://ciscolive.ciscoevents.com/ciscolivebot/#BRKDCN-2794

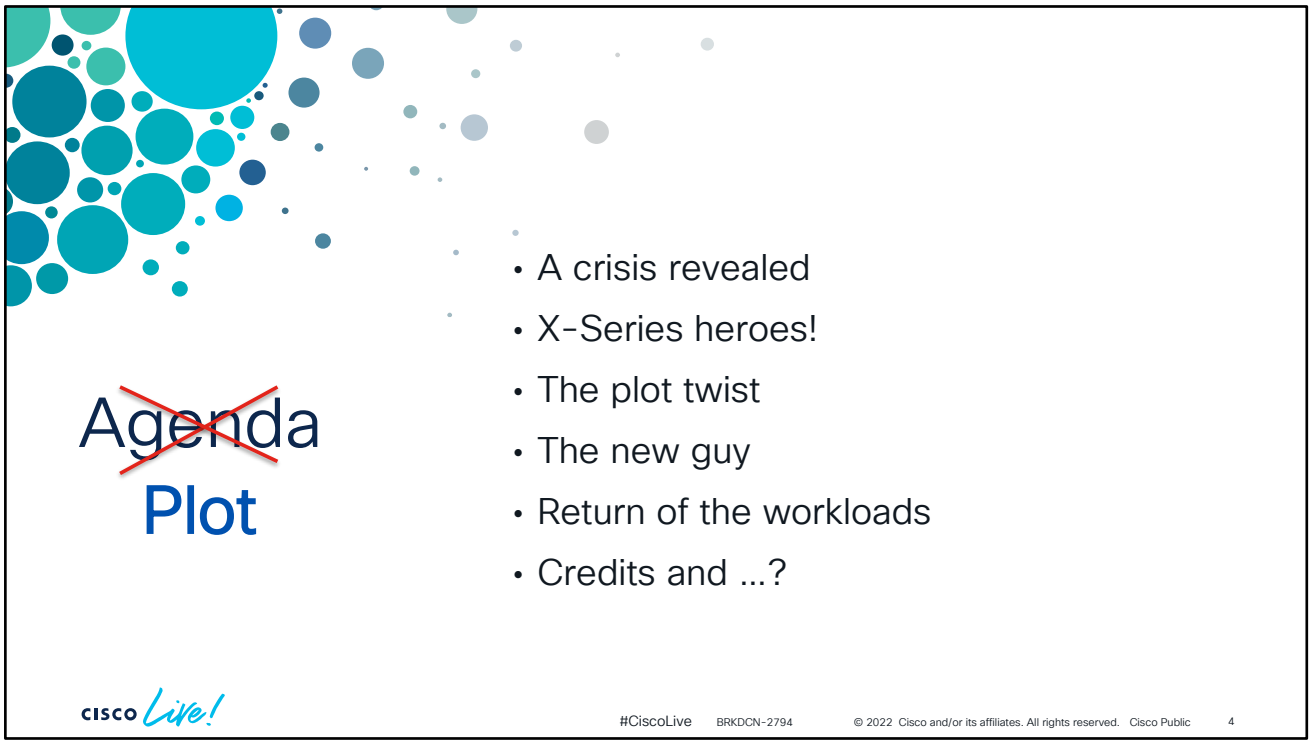


They say a presentation should tell a story. Well, isn't a movie better than just a story?

Applications were always happy with access to some CPU, memory, storage, and networking to talk to their friends.

But as applications matured they began to try new things.

Big data came around, requiting lots of local storage. Databases want crazy amounts of memory. AI/ML wants compute acceleration.

Virtual desktops want faster and faster GPUs to play their YouTube videos.

Then there are these new little ones that don't even want a whole server.

They want to be kept in little containers and run around in the cloud. Now they think they don't even need a server!

All this causes datacenter sprawl. Form factors, devices, management endpoints, aggregation challenges (notifications, logs, …)

The X-Series modular system brings them back together.

Having a more adaptable platform is great, but you are still challenged with management sprawl. Different devices, different clouds, different environments.

Sometimes you want the cloud for good reason, but may want the same application to run on-prem. (dev and prod, data sovereignty.)

Intersight provides a single, flexible control plane for a complete hybrid cloud environment.

Run your workload where it makes the best sense with on-prem covered by the modular X-Series system.

UCS-X offers a long-term roadmap of continuing innovation without critical choke points leading to a limited life.


Should now be obvious why I went to engineering school, not art.


A lot has changed and will continue to change in the computing landscape:

Several years ago, technologies and innovations around software defined storage and accelerated computing with GPUs created architectural shifts in infrastructure and spawned new workloads (AI/ML) and solutions (HCI).

As we look just over the horizon, new technologies will shape infrastructure for the next several years. PCIe and CXL will enable the disaggregation of resources and the custom reassembly of those resources to best serve applications. In addition, high performance CPUs and GPUs need to be accommodated and 100/200G networking to bring next gen performance and bandwidth.

We believe there are a few of key trends happening – hybrid cloud is driving macro-orchestration across resources, on-premise and in the public cloud, to drive simplicity and agility. Micro-segmentation and the disaggregation of technologies/resources to meet the growing application diversity with system designs that can be molded or adapted to modern applications. Lastly, operations at scale to provide visibility, optimization, orchestration, automation and support across a hybrid cloud landscape.

## Architectural silos drive complexity

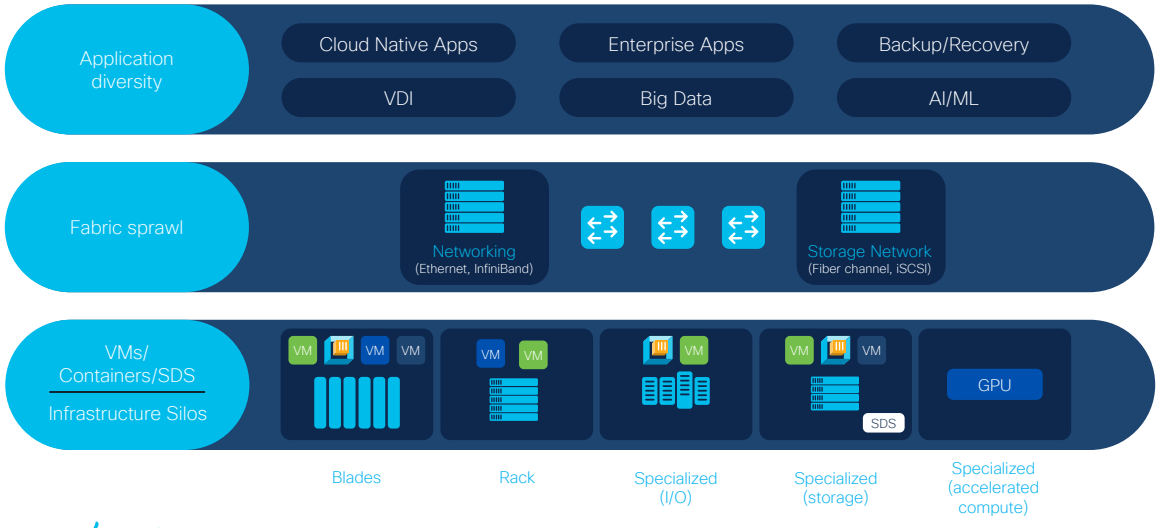






**Traditional Infrastructure will continue to fail**

Workloads are expanding beyond just leveraging virtualization and cloud for legacy applications. Big Data and analytics are requiring new approaches to infrastructure and the IoT is blurring the lines of where data is generated and processed. IT teams are choosing the optimal infrastructure architecture to accommodate this application diversity but that can create silos, complex management models, and a lack of visibility and control.

Complexity continues to be the data centers Achilles heel. In the face of tremendous application growth, IT teams still choose unique systems and solutions

for each application. The overhead to plan, procure, operate, optimize, a portfolio of non-standard architectural choices (i.e. blades, racks and specialized systems) can require unique management tools and make it difficult to adapt, scale, and maximize efficiency, resulting in higher power and cooling costs, forklift upgrades to accommodate new technologies, and the inability to quickly respond to dynamic changes in application requirements. For example, in traditional rack environments, scale may come in small increments but at the expense of operational efficiency resulting in higher power and cooling costs. In most instances, automation and agility remain out of reach.

In an era where orchestration and automation is the key to agility, complex integration if disparate architectures is complexity that needs to be avoided.

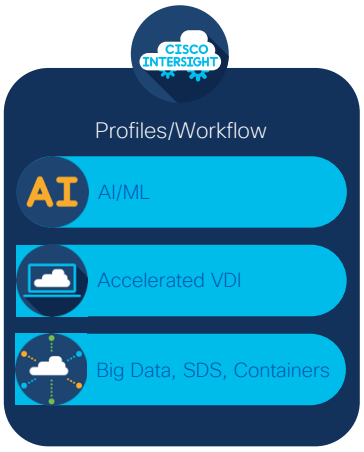


## UCS X-Series With Intersight

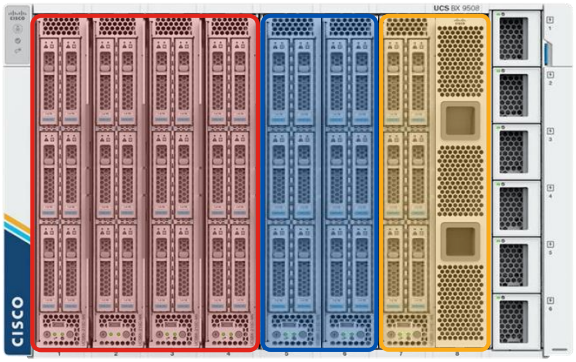


UCS X-Series – 7 RU, 8 Nodes

| 600+ | ~100 TB | ~1 PB<br>48 Drives | 48M IOPs | ~1600G | ~100G | 24 |
|---|---|---|---|---|---|---|
| Cores | Memory | Storage | In a Chassis | Agg N/W in a Chassis | End-to-end | GPUs per Chassis |





X-Series has greatly expanded the applicability of a modular system to workloads that have been tied to rack systems.

Tailor the chassis, or parts of the chassis, to the needs of the workload.

Provision only what's needed.

Cooling scales with the nodes, so capacity isn't wasted when high performance options aren't needed.

# X-Series Heroes (and a sidekick)

## X-Fabric
## 5th Gen Networking

# UCS X-Fabric and PCIe Nodes with GPU

**GPU support**
- 4X T4
- 2x A16
- 2x A40
- 2 x A100

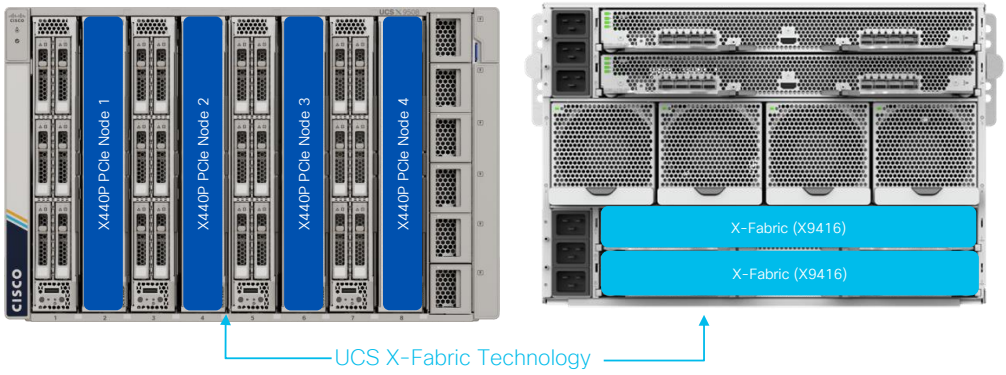






1. Overview of the season 2 X-Fabric innovation.
2. GPUs on PCIe nodes connected to the compute nodes through the X-Fabric modules.

## X9416 X-Fabric XFM Modules

- The X9416 XFM provides fixed PCIe Gen4 x16 links between each pair of odd (N) and even (N+1) slots
- Replaces lower fabric module pair
- Can be hot swapped in a chassis with the original XFM blanks
- Connects to all chassis slots
- **No configuration required**

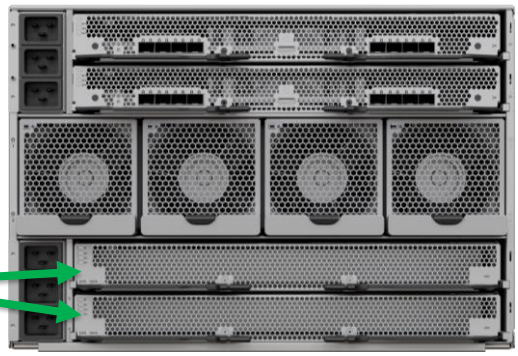

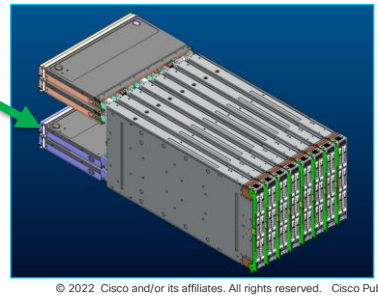

## XFM Components

- No active fabric components on the system board
- Three replaceable fan modules
- Same fans as IFM

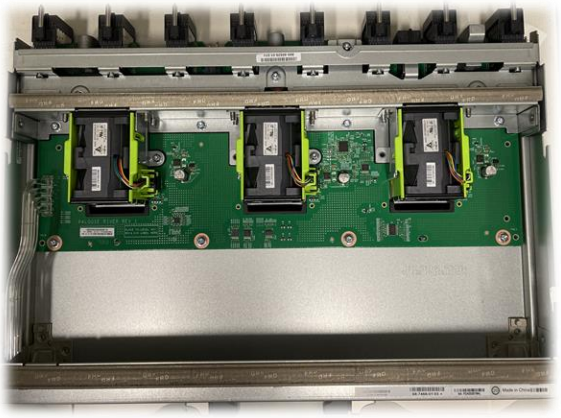

## X-Fabric Mezz Cards

- All X-Series compute node mezzanine cards provide connectivity to the X-Fabric
- 100Gb VIC 14825 Mezz
- PCIe pass-through Mezz
- One PCIe Gen4 x16 link per CPU

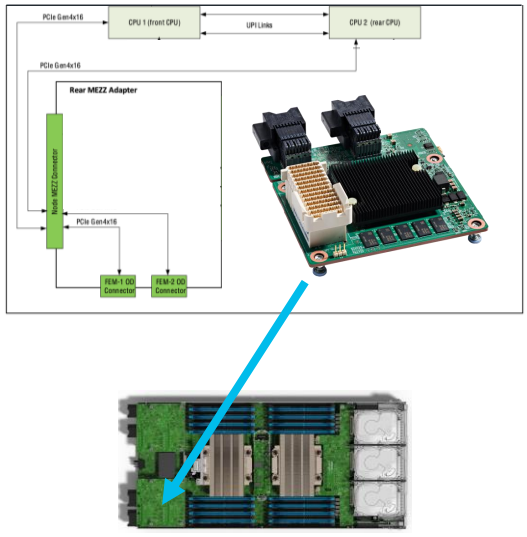






Just like the VIC mLOMs the physical connectors from the compute node to the XFM are located on the Mezzanine card, so the fabric is independent of the compute node.

# X440p PCIe Node

- The X440p PCIe node provides 2 or 4 PCIe slots connected to an adjacent compute node
- Includes two riser cards, type A or B
- Risers include GPU power cables
- Riser A – Up to two dual width GPUS
  - A16, A40, or A100
- Riser B – Up to four single width GPUs
  - T4
- No mixing of GPU models on a server

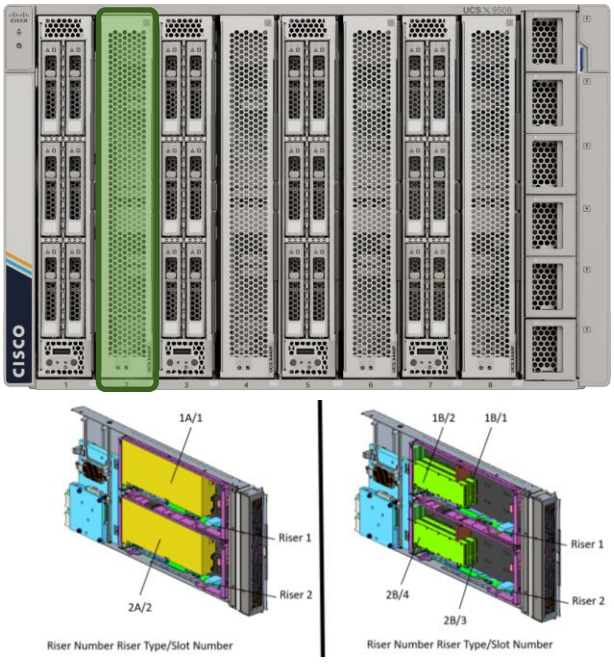






X-Fabric connectivity is integrated on the PCIe node. No separate mezzanine required.

Different GPUs may be used on separate PCIe nodes, as long as they are connected to separate servers.

Compute and PCIe node position are interchangeable. E.g. PCIe node in slot 1 and X210c in slot 2 is also valid.

NVIDIA does not support mixed GPU models on a single server.

# Intersight PCIe Node inventory

- GPUs will show up under the compute node inventory
- Includes PCIe node details

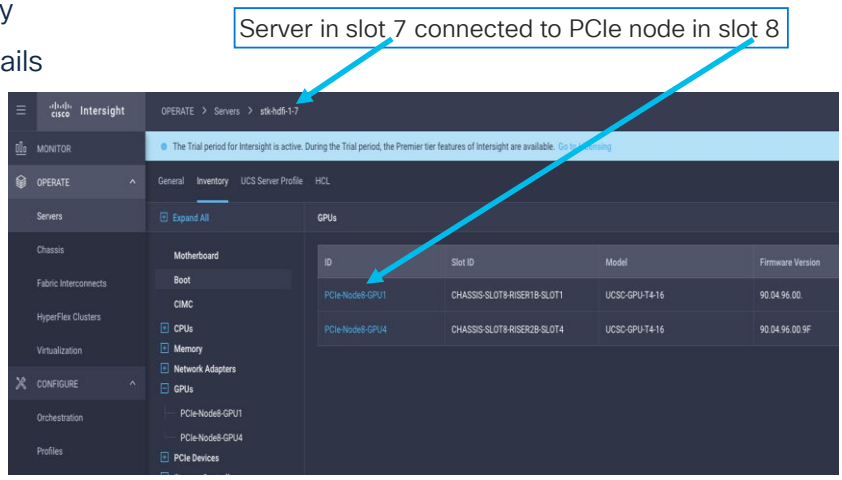

# UCS X-Series GPU Portfolio






UCS-X with the A16 is a fantastic accelerated VDI solution. Up to 8 A16 on 4 compute nodes in 7RU, providing a total of 32 GPUs for amazinf density and performance.

A40 delivers a new top-end experience for virtual engineering workstations.

Up to 24 T4 in a single chassis with 6 on each of 4 compute nodes for an excellent accelerated knowledge worker experience.

For raw AI/ML up to 8 A100-80G on 4 compute nodes.

# UCS X-Fabric - X210c and X440p PCIe Node

- An X-Fabric Mezz card (VIC or Pass-through) on the compute node connects it to the XFM pair
- One Gen4 x16 link per CPU is routed to the XFM pair (CPU1 to XFM1 and CPU2 to XFM2)
- The X440p connects the XFM1 link to Riser 1 and XFM2 to Riser 2
- GPU firmware is updated with the attached server firmware upgrade

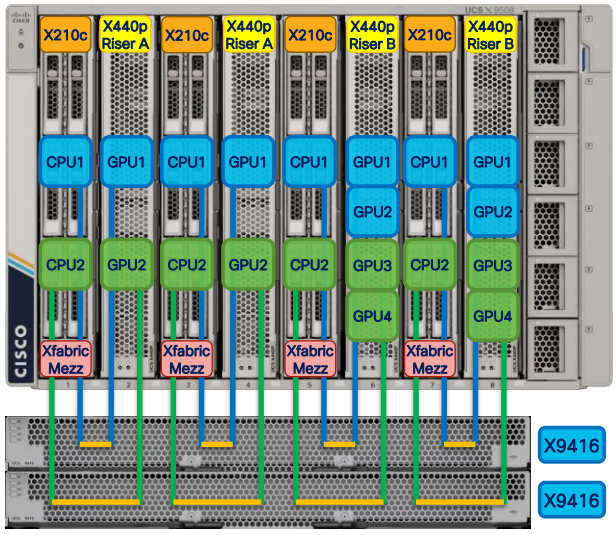






No failover or redundancy of the PCIe fabric. The two fabrics provide higher total bandwidth.

# UCS X-Fabric, X440p PCIe Node, and GPU Mezz

- Compute nodes can continue to exist in adjacent slots where the PCIe node is not needed (example slots 1 and 2)
- PCIe links through the X-Fabric between compute nodes will not come up, even with an X-Fabric Mezz card installed on the compute node (example slots 3 and 4)
- Compute with the GPU Front Mezz can be used with the PCIe node and Riser B to support up to six GPU per node (slots 7 and 8)

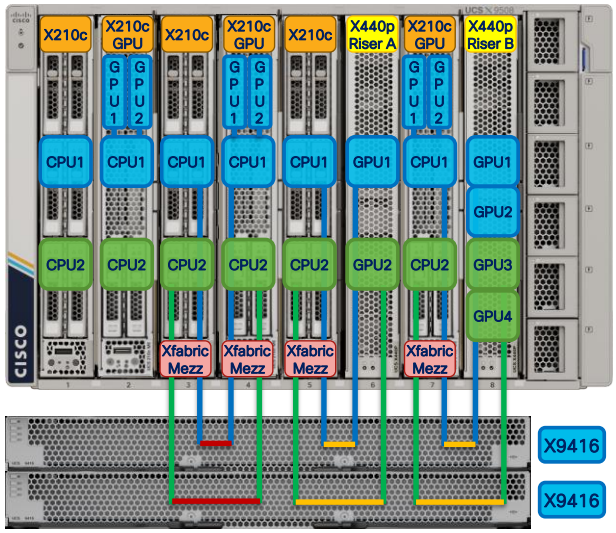






Using X-Fabric does not force slots to be dedicated to a specific node type.

Compute can still be populated in any slot.

Both GPUs in slots 2, 4, and 6 must be the same.

In the slot 7 and 8 example all 6 GPUs must be the same type (T4 in this case.)

## X10c GPU Front Mezz

- UCSX-X10C-GPUFM
- Two PCIe Gen4 x8 HHHL slots
- One or two NVIDIA T4 GPUs
- Requires PID UCSX-GPU-T4-MEZZ
- Up to two U.2 NVMe Drives
- 4x PCIe Gen4 Per Drive
- No SAS/SATA Support
- Intel VROC Capable with Intel Drives

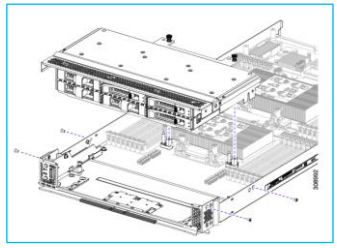





These T4's have a different heatsink to allow the proper airflow with the vertical GPU orientation, thus a unique PID.

# 5th Gen Fabric

cisco Live!

# 6536 5th Generation Fabric Interconnect

- 7.4 Tbps bandwidth per FI in 1 RU form-factor
- 36x QSFP Ports
  - 32x 100G Ethernet ports (1 – 32)
  - 4x Unified ports (33 – 36)
    - 100G Ethernet or
    - 16x 8/16/32G FC ports via 128G-FC QSFP28 break-out

# FI 6536 – Rear and front view

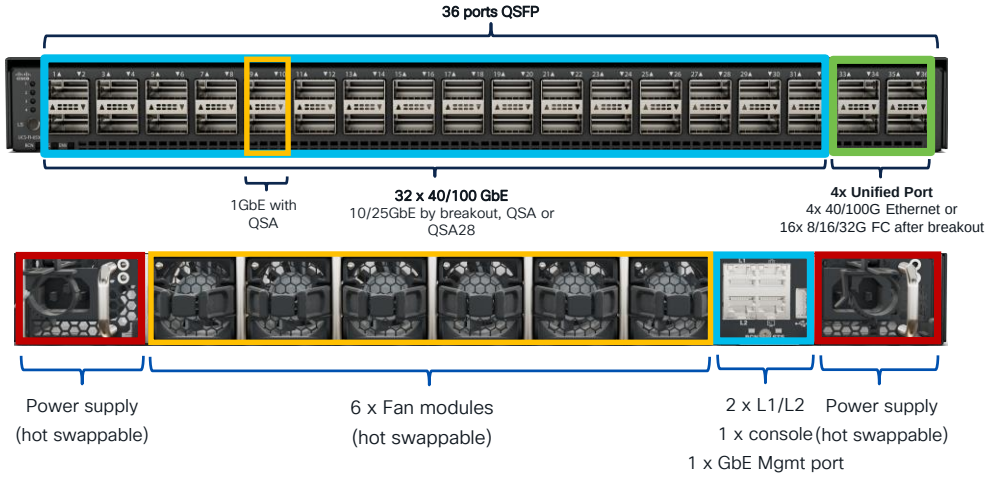





16x 8/16/32G FC ports using 128G-FC QSFP28 break-out

### Front-to-back cooling

- Fan-side intake, port-side exhaust

### hot-swappable fans and power supplies

- Each fan module consists of two fan rotors. Redundancy of fan is implanted in rotor level. With a total of 12 rotors, the system continues to operate with 9 fan rotors.
- Helps enable high availability in multiple configurations
- Increase serviceability
- Provides uninterrupted service during maintenance

### Power supply

- Two power supplies (AC or DC)

## Cisco UCS 9108-100G Intelligent Fabric Module

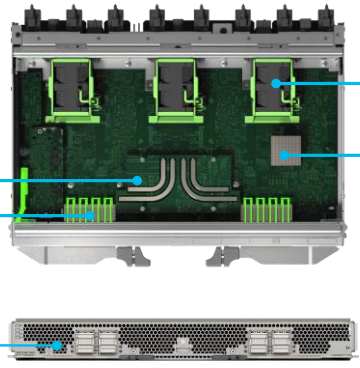

**Cisco ASIC**
- 2 Tbps throughput per IFM
- 8 x 100G or 32 x 25G ports per server

**FAN**
- Three internal fans

**EXTERNAL PORTS**
- 8x 100G QSFP ports
- Can cable from 1 to 8 ports

**CMC Processor**
- Chassis management





When you first look at the IFM, you see three fans that snap in and snap out. Although they are not hot-swappable, they are easily replaceable because the IFM can be removed from the chassis to replace the fans taking down only one side of the fabric. On the back, each IFM has eight ports, and each port is capable of 100 gigabits. In this architecture, the fabric interconnects (FIs) are 100 gigabit. The connections supporting the transceivers and the copper or optical cables are 100 gigabit.

- Three fans snap in and snap out
- Not hot-swappable
- High reliability
- Each IFM has eight ports, each port is capable of 100 gigabits
- 1-8 ports cabled to FI, numbered 1-8, port groupings are based on space, no functional difference
- In this architecture, FIs are 100 gigabit
- 100-gigabit connections supporting transceivers, copper or optical cables

# VIC 15231 mLOM

- 2 x 100G with 5th Gen FI
- 100Gb Ethernet vNIC
- 100Gb FC vHBA
- PCIe Gen 4 x16
- NVMeoF: FC-NVMe, RoCEv2

cisco Live!

# Return of the workloads


Does the system really deliver what applications need and have been getting from rack servers.

# Consistent raw storage performance

- Performance Summary. ----- Chassis level performance across individual blades
  - 8 X210c Nodes each with 6 intel P5600 NVMe Drives
  - FIO Mix of Reads / Writes - 4K 100 % Reads to 100% Writes
  - Sustained and reproducible performance for longer duration (4– 8 hours)
  - Near identical performance from 1 Blade to 8 Blades – Power, Cooling and Efficiency
  - Total 38 million Read IOPs in a single Chassis - Each Blade ~4.8 million with less than 400 micro-Seconds latency

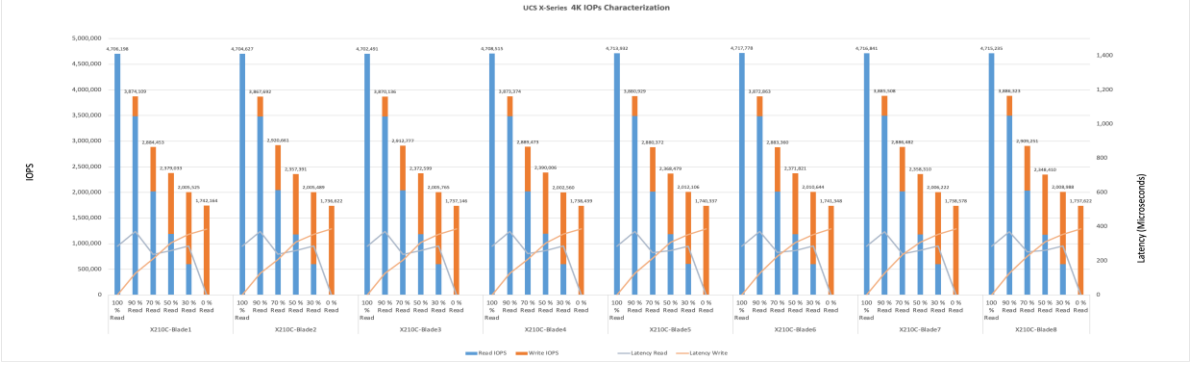






Chart shows performance from 100% read to 100% write across all 8 compute node slots.

Higher level storage constructs are doomed if basic storage performance is compromised.

No hot spots, no weak power corners.

No concern for where or how much storage is provisioned.

# TCP End to End 100G : VIC 15231 performance

**iPERF Traffic** :
Fabric A : Blade3 vnic1 (client) -> Blade2 vnic1 (Server)
Fabric B : Blade3 vnic2 (client) -> Blade2 vnic2 (Server)

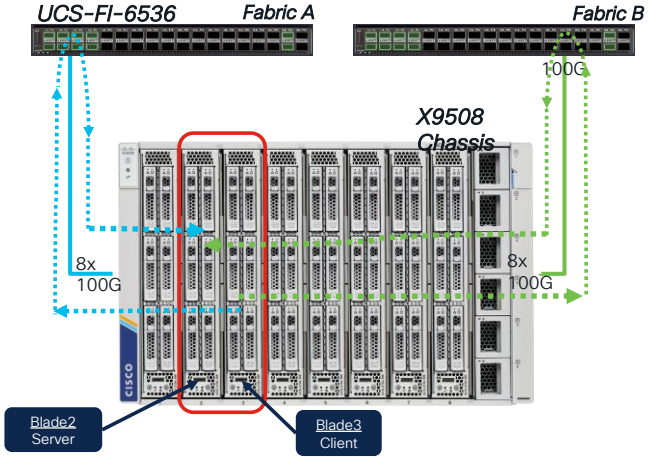


**iPerf results**

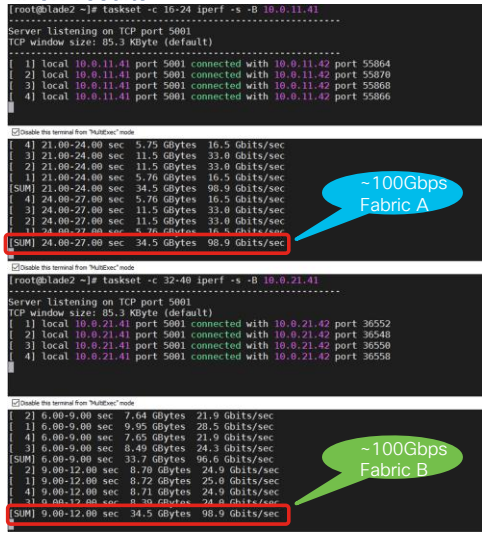

# End to End 32G FC: VIC 15231 performance

## 32G FC connectivity to a Storage Target:

Chassis-Blade2 vHBA1 <-- HBA1 Storage (100% read)
Chassis-Blade2 vHBA1 <-- HBA5 Storage (100% read)

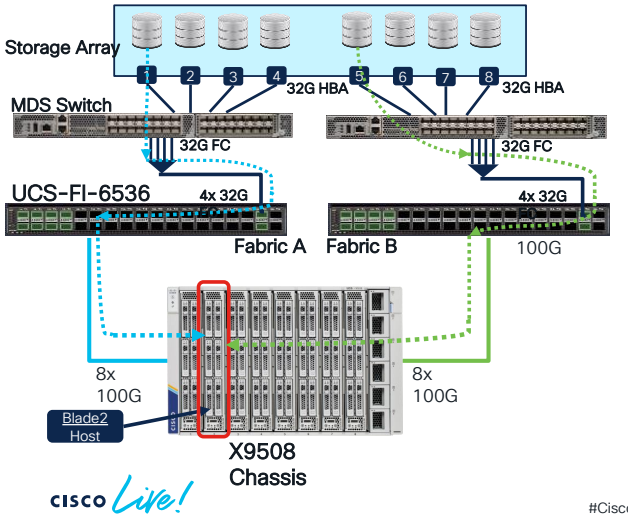


## FIO results per x210c

```
[root@ora5g-node1 ~]# cat /sys/class/fc_host/host3/speed
100 Gbit
[root@ora5g-node1 ~]# fio --filename=/dev/mapper/5g_vol1:/dev/mapper/5g_vol2 --direct=1 --rw=r
w --ioengine=libaio --bs=512k --rwmixread=100 --iodepth=8 --numjobs=8 --runtime=120 --group_re
porting --name=seqreadwrite
seqreadwrite: (g=0): rw=rw, bs=(R) 512KiB-512KiB,                    512KiB-512KiB, ioengi
ne=libaio, iodepth=8
...
fio-3.19
Starting 8 processes
Jobs: 8 (f=16): [R(8)][11.6%][r=3141MiB/s] r=6281 IOPS][eta 01m:47s]
```

~32G FC Fabric A

```
[root@ora5g-node1 ~]# cat /sys/class/fc_host/host7/speed
100 Gbit
[root@ora5g-node1 ~]# fio --filename=/dev/mapper/5g_vol9:/dev/mapper/5g_vol10 --direct=1 --rw=
rw --ioengine=libaio --bs=512k --rwmixread=100 --iodepth=8 --numjobs=8 --runtime=120 --group_r
eporting --name=seqreadwrite
seqreadwrite: (g=0): rw=rw, bs=(R) 512KiB-512KiB,                    512KiB-512KiB, ioengi
ne=libaio, iodepth=8
...
fio-3.19
Starting 8 processes
Jobs: 8 (f=16): [R(8)][12.4%][r=3142MiB/s] r=6284 IOPS][eta 01m:46s]
```

~32G FC Fabric B

## Result from Storage Array

```
FlexPod-A800::>
FlexPod-A800::> qos statistics performance show
Policy Group                 IOPS      Throughput     Latency Is Adaptive? Is Shared?
-------------------- ---------- ---------- ---------- ------------ ----------
-total-                  108565   6342.20MB/s      4
User-Best-Effort         101469   6342.12MB/s      5
System-Work                7096     74.24KB/s     38.
-total-                  101496   6341.45MB/s   5.00
User-Best-Effort         101455   6341.44MB/s      5                false true
System-Work                  41     15.06KB/s      1.00us false true
-total-                  101503   6342.33MB/s   5.00ms - -
User-Best-Effort         101465   6342.31MB/s   5.01ms false true
System-Work                  38     19.36KB/s  105.00us false true
```

Total 64G across Fabric A & B

# 100G FC per Fabric : VIC 15231 performance

## 100G vHBA performance in a x210c:

Chassis Blade2 <-- 4x 32G FC storage target (100% read)

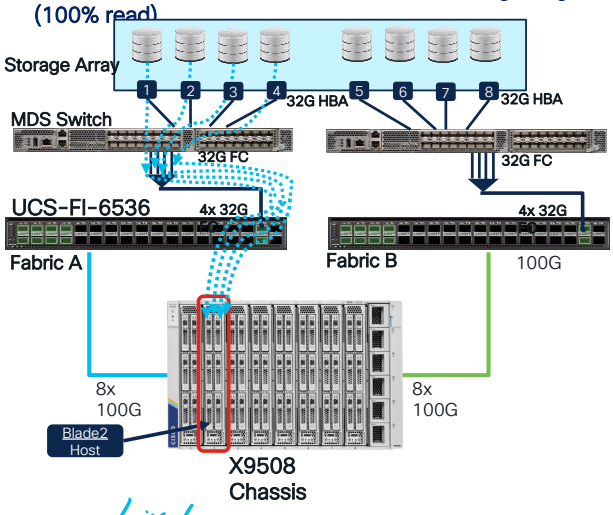


## FIO results per x210c

```
[root@ora5g-node1 ~]# fio --filename=/dev/mapper/5g_vol1:/dev/mapper/5g_vol2:/dev/mapper/5g_vol3:/dev/
mapper/5g_vol4:/dev/mapper/5g_vol5:/dev/mapper/5g_vol6:/dev/mapper/5g_vol7:/dev/mapper/5g_vol8 --direc
t=1 --rw=rw --ioengine=libaio --bs=512k --rwmixread=100 --iodepth=32 --numjobs=64 --runtime=30 --group
_reporting --name=seqreadwrite
seqreadwrite: (g=0): rw=rw, bs=(R) 512KiB-512KiB, (W) 512KiB-512KiB, (T) 512KiB-512KiB, ioengine=libai
o, iodepth=32
...
fio-3.19
Starting 64 processes
Jobs: 64 (f=512): [R(64)][64.5%][r=11.1GiB/s][r=22.7k IOPS][eta 00m:11s]
```

```
FlexPod-A800::> q s p s
  (qos statistics performance show)
Policy Group             IOPS      Throughput    Latency Is Adaptive? Is Shared?
------------------------ --------- ------------- ---------- --------- ----------
-total-                  183591    11479.40MB/s     4.9
User-Best-Effort         183529    11479.38MB/s     4.
_System-Work                 62       22.08KB/s
-total-                  183511    11474.38MB/s     4.60
User-Best-Effort         183489    11474.38MB/s     4.                false true
_System-Work                 22           0MB/s       .00us false true
-total-                  183569    11478.01MB/s     5.09ms - -
User-Best-Effort         183529    11478.00MB/s     5.09ms false true
_System-Work                 40        5.34KB/s   200.00us false true
```

92G FC (~100G) per Fabric

# 100G NFS: VIC 15231 performance

## 100G NFS connectivity to a Storage Filer:

Chassis-Blade2 vNIC1 <-- 100G NIC on Storage (100% read)

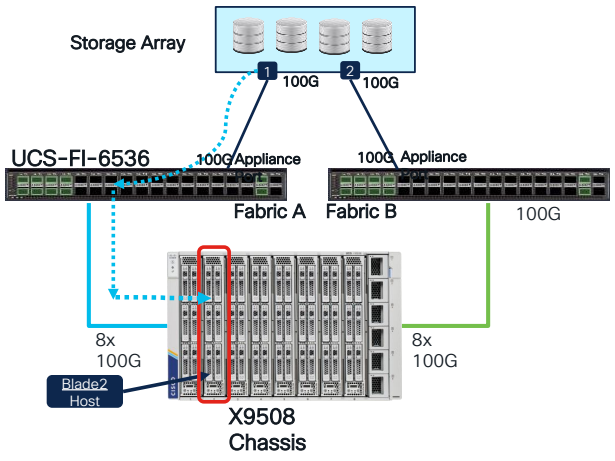


## NFS results per x210c

```
root@oraSg-node1 ~]# fio --name=test1 --filename=/fiovol1/data1:/fiovol1/data2:/fiovol1/data3:/fiovol1/data4:/fiovol2
/fiovol2/data22:/fiovol2/data23:/fiovol2/data24:/fiovol3/data31:/fiovol3/data32:/fiovol3/data33:/fiovol3/data34:/fiovo
l:/fiovol4/data42:/fiovol4/data43:/fiovol4/data44 --rw=read --direct=1 --ioengine=libaio --bs=512k --numjobs=16 --iode
runtime=300 --time_based --group_reporting
test1: (g=0): rw=read, bs=(R) 512KiB-512KiB, (W) 512KiB-512KiB, (T) 512KiB-512KiB, ioengine=libaio, iodepth=64
...
fio-3.19
Starting 16 processes
Jobs: 16 (f=256): [R(16)][100.0%][r=11.3GiB/s][r=23.2k IOPS][eta 00m:00s]
test1: (groupid=0, jobs=16): err= 0: pid=82192: Fri Apr 22 13:53:52 2022
  read: IOPS=23.3k, BW=11.4GiB/s (12.2GB/s)(3400GiB/300020msec)
    slat (usec): min=26, max=5329, avg=56.33, stdev=25.13
    clat (usec): min=1957, max=531941, avg=43929.16, stdev=9451.20
     lat (usec): min=1997, max=531987, avg=43995.59, stdev=9450.83
    clat percentiles (msec):
     |  1.00th=[   25],  5.00th=[   28], 10.00th=[   31], 20.00th=[   38],
     | 30.00th=[   41], 40.00th=[   44], 50.00th=[   46], 60.00th=[   47],
     | 70.00th=[   49], 80.00th=[   51], 90.00th=[   53], 95.00th=[   56],
     | 99.00th=[   62], 99.50th=[   66], 99.90th=[  103], 99.95th=[  131],
     | 99.99th=[  205]
   bw (  MiB/s): min= 5977, max=12232, per=100.00%, avg=11646.89, stdev=19.82, samples=9584
   iops        : min=11947, max=24465, avg=23293.77, stdev=39.65, samples=9584
  lat (msec)   : 2=0.01%, 4=0.01%, 10=0.01%, 20=0.06%, 50=78.42%
  lat (msec)   : 100=21.39%, 250=0.10%, 500=0.01%, 750=0.01%
  cpu          : usr=0.22%, sys=0.45%, ctx=6244500, majf=0, minf=
  IO depths    : 1=0.1%, 2=0.1%, 4=0.1%, 8=0.1%, 16=0.1%, 32=
     submit    : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=
     complete  : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=
     issued rwts: total=6980099,0,0,0 short=0,0,0,0 dropped=0,0,0,0
     latency   : target=0, window=0, percentile=100.00%, depth=64

Run status group 0 (all jobs):
   READ: bw=11.4GiB/s (12.2GB/s), 11.4GiB/s-11.4GiB/s (12.2GB/s-12.2GB/s), io=3408GiB (3660GB), run=300028-300020msec
```

~100G NFS

```
FlexPod-A800: cluster.cluster: 4/22/2022 13:50:07
 cpu  cpu   total                 fcache   total    total data     data     data cluster  cluster  cluster     disk     disk     pkts     pkts
 avg busy     ops  nfs-ops cifs-ops    ops    recv     sent busy     recv     sent    busy     recv     sent     read    write     recv     sent
 ---- ---- ------- -------- -------- -------- -------- -------- ---- -------- -------- ------- -------- -------- -------- -------- --------
  18% 43%  187117   187117        0        0   38.1MB   11.4GB  97%   37.7MB   11.4GB      0%   338KB   336KB   478MB    1.7MB  230535  1371110
  18% 42%  186775   186775        0        0   38.4MB   11.5GB  98%   38.1MB   11.5GB      0%   285KB   286KB   481MB    3.9MB  236859  1375341
  18% 43%  187603   187603        0        0   39.0MB   11.5GB  98%   38.7MB   11.5GB      0%   286KB   285KB   488MB   23.9KB  245848  1376171
  18% 43%  186620   186620        0        0   38.0MB   11.5GB  98%   37.8MB   11.4GB      0%   266KB   266KB   609MB    1.1MB  231885  1371855
  18% 43%  186775   186775        0        0   38.5MB   11.5GB  98%   38.2MB   11.5GB      0%   301KB   301KB   643MB   17.9KB  237274  1374307
  18% 43%  186634   186634        0        0   38.4MB   11.4GB  97%   38.1MB   11.4GB      0%   251KB   251KB   613MB    1.4MB  239374  1370362
  18% 43%  185568   185568        0        0   37.7MB   11.4GB  97%   37.4MB   11.4GB      0%   278KB   278KB   610MB    3.9MB  226770  1369212
  18% 43%  187085   187085        0        0   38.3MB   11.4GB  97%   38.0MB   11.4GB      0%   307KB   307KB   582MB   23.9KB  235274  1365179
  18% 43%  186116   186116        0        0   38.4MB   11.4GB  98%   38.2MB   11.4GB      0%   269KB   269KB   577MB    1.1MB  237298  1369754
  18% 43%  186856   186856        0        0   38.1MB   11.4GB  98%   37.8MB   11.4GB      0%   270KB   270KB   579MB   29.9KB  234293  1372264
  18% 43%  186380   186380        0        0   38.5MB   11.4GB  97%   38.2MB   11.4GB      0%   281KB   281KB   590MB    646KB  238783  1371946
  18% 43%  186480   186480        0        0   38.6MB   11.4GB  97%   38.3MB   11.4GB      0%   309KB   309KB   577MB    2.2MB  241873  1366584
  18% 43%  185636   185636        0        0   38.2MB   11.4GB  97%   37.9MB   11.4GB      0%   284KB   282KB   560MB   23.9KB  235542  1368954
  18% 43%  185936   185936        0        0   37.6MB   11.3GB  97%   37.3MB   11.3GB      0%   280KB   282KB   594MB    1.1MB  225957  1368059
  18% 43%  187402   187402        0        0   38.3MB   11.5GB  98%   38.0MB   11.5GB      0%   268KB   268KB   587MB   29.9KB  234891  1375568
```

# Speeds and Feeds – Oracle SLOB performance

- Performance Summary
  - SLOB – Silly Little Oracle Benchmark
    - SLOB is an Oracle I/O workload generation tool kit
    - Tests physical random single-block reads (db file sequential read), single block writes (DBWR flushing capacity) and REDO logging I/O capacity
  - Near linear User Scale with a mix of Database 8K Reads and Updates at 100 %, 70% and 50% Reads/Updates
  - 6 x 2.9 TB NVMe Disks as single "DATA" Disk Group
  - 2 TB Schema, 16 GB SGA

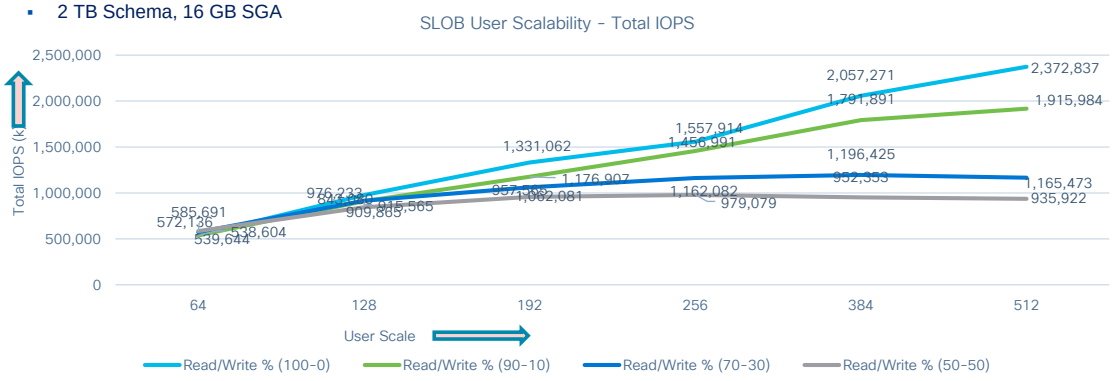






Now laying an Oracle DB on the local storage using Oracle ASM (Automatic Storage Management)

And using SLOB to measure performance.

As users scale the IOPs also scale. The limit is the read/write bandwidth of the physical drives.

# Database Performance Test

## Swingbench – Users and TPM Scale

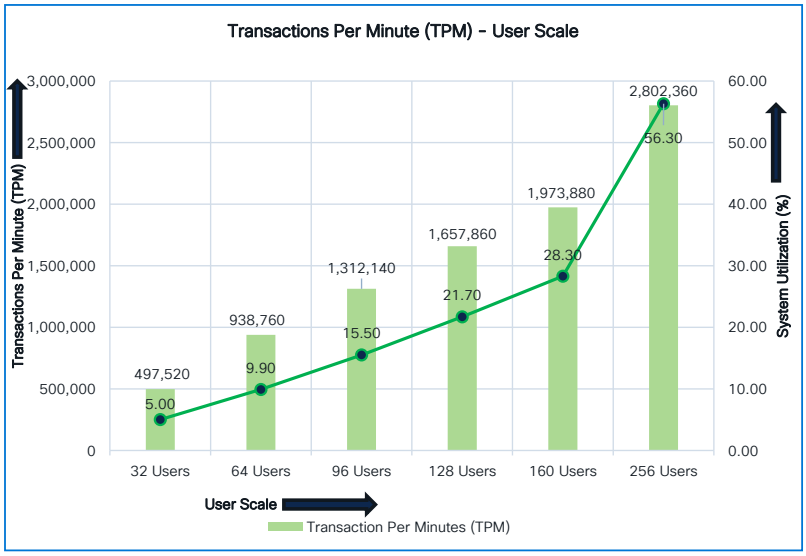


### ✓ Workload

- Tool Used – Swingbench
- Oracle Database – 5 TB
- OLTP Workload – SOE
- User Scale – 32 to 256
- Each Test Run for 24 Hours

### ✓ Observation

- Linear Scalability TPM
- Latency under 0.4 milliseconds

## Oracle AWR Latency Sample





Similar testing using an alternate benchmark (SwingBench))

The limiting factor is increasing CPU utilization as user count reaches maximum.

# Database Performance Test

- Swingbench – Users and IOPs Scale

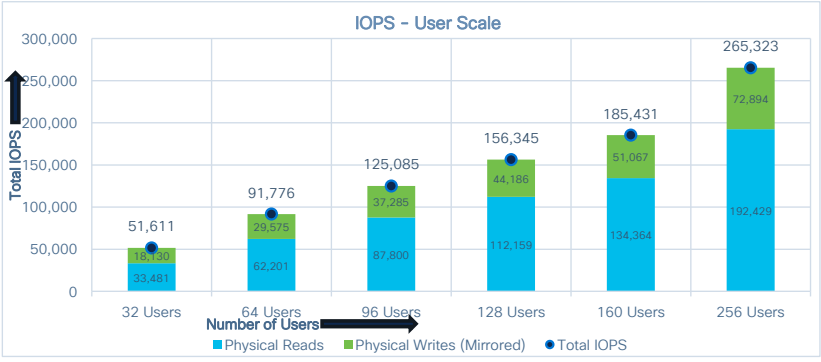


- iostat 300 user test output – Even IOPs on all drives

| Device | r/s | w/s | rMB/s | wMB/s | rrqm/s | wrqm/s | %rrqm | %wrqm | r_await | w_await | aqu-sz | rareq-sz | wareq-sz | svctm | %util |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| nvme3n1 | 32136.53 | 33370.60 | 251.10 | 264.75 | 0.00 | 6.60 | 0.00 | 0.02 | 0.12 | 0.06 | 0.32 | 8.00 | 8.12 | 0.02 | 99.83 |
| nvme4n1 | 32165.57 | 33612.77 | 251.32 | 264.22 | 0.03 | 6.60 | 0.00 | 0.02 | 0.12 | 0.06 | 0.30 | 8.00 | 8.05 | 0.02 | 99.83 |
| nvme0n1 | 32272.23 | 32910.87 | 252.15 | 265.11 | 0.07 | 5.80 | 0.00 | 0.02 | 0.12 | 0.07 | 0.39 | 8.00 | 8.25 | 0.02 | 99.87 |
| nvme2n1 | 32158.27 | 33337.97 | 251.27 | 263.98 | 0.00 | 5.83 | 0.00 | 0.02 | 0.12 | 0.06 | 0.33 | 8.00 | 8.11 | 0.02 | 99.78 |
| nvme1n1 | 32235.40 | 33527.93 | 251.86 | 265.31 | 0.00 | 5.47 | 0.00 | 0.02 | 0.12 | 0.07 | 0.37 | 8.00 | 8.10 | 0.02 | 99.79 |
| nvme5n1 | 32179.27 | 33424.87 | 251.40 | 263.69 | 0.03 | 5.50 | 0.00 | 0.02 | 0.12 | 0.07 | 0.40 | 8.00 | 8.08 | 0.02 | 99.86 |

## ✓ Workload

- Tool Used – Swingbench
- Oracle Database – 5 TB
- OLTP Workload – SOE
- User Scale – 32 to 256
- Sustainable Test results over 24 hours
- IOSTAT sample for 300 Users

## ✓ Observation

- Linear Scalability IOPs by adding Users
- 50k to 265k IOPs
- Latency under 300 Microseconds
- IO evenly distributed across all NVMe drives





Again performance scales until CPU is maxed out.

No hot spots across the available drives.

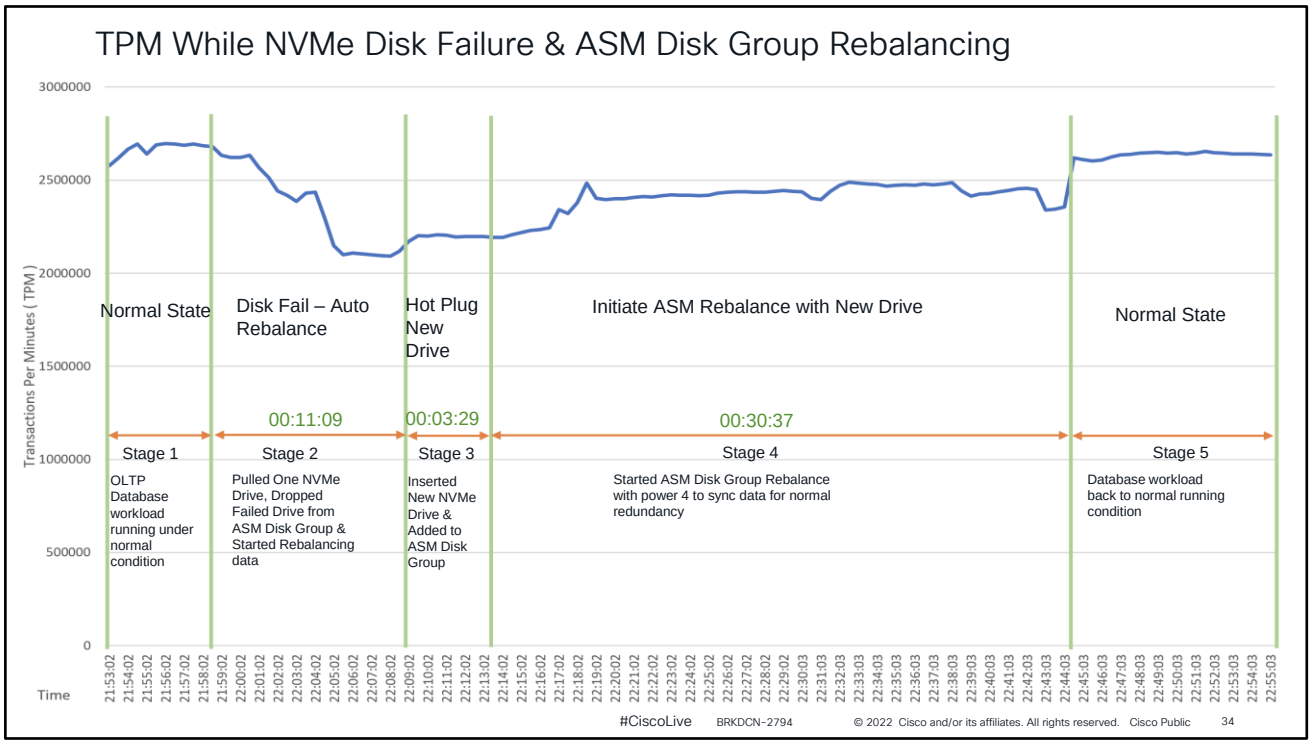


Performance across the full test remains consistent and high.

# Hyperflex on X-Series
## Coming Q4

### All Flash and All NVMe options

**#1 Performance**
Proprietary, data-distributed architecture

**Seamless Scale**
Edge to Datacenter, perf & capacity scaling options

**Ease of Use**
Simplicity from eval, sizing, Day 0/1 to support

**Cloud Integrations**
Multicloud ready, Intersight for mgmt,analytics,support

cisco Live!

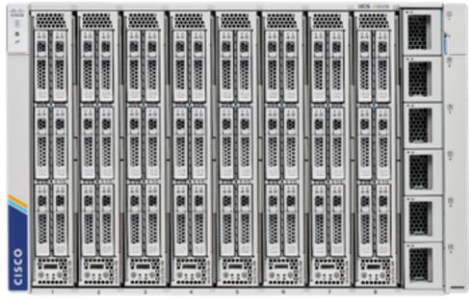

Cisco Hybrid Cloud Platform
HyperFlex Data Platform + Intersight






Hyperconverged infrastructure requires high, consistent performance from the storage and network fabrics.

X-Series provides compute with SSD or NVMe drives.

# UCS-X vSAN ReadyNode


| AF-4: UCSX-210c-M6 | | |
|---|---|---|
| Generation: Intel Xeon Scalable (Ice-Lake-SP) - Profile: AF-4 Series - Raw Storage Capacity: Up To 7.68 TB | | |
| Components | Details | QTY |
| SKU | UCS-VSAN-X210c-M6N | |
| ESXi Pre-Installed? | No | |
| SYSTEM | UCSX-210C-M6 | 4 |
| CPU | Intel Xeon Ice lake 6348 2.60Ghz 28Core | 8 |
| MEMORY | 32GB DDR4-3200-MHz | 64 |
| CachingTier | Cisco SSD DC P5600 Series SSDPF2KE016T9K (1.6TB, 2.5" U.2, PN: UCS-NVMEI4-I1600) | 8 |
| CapacityTier | Cisco SSD DC P5500 Series SSDPF2KX038T9K (1.92TB, 2.5" U.2, PN: UCS-NVMEI4-I1920) | 16 |
| Controller | N/A | 0 |
| NIC | VIC 14425 | 4 |
| BootDevice | UCS-M2-HWRAID M.2 Raid controller with 2x240GB SATA M.2 | 8 |
| Supported Releases | ESXi 7.0 U3 (vSAN 7.0 Update 3) | |





Hyperconverged infrastructure requires high, consistent performance from the storage and network fabrics.

X-Series provides compute with SSD or NVMe drives.

NVMe nodes coming soon.

# Now playing in the World of Solutions

PARTNER
SOLUTIONS





cisco Live!

# The plot twist

It’s not always about the platform.  Our reality is hybrid cloud.  No one wants to be forced to use a different management tool depending on where an application is deployed.

## Public Cloud Integrations for Container and Virtual Machine Operational Consistency across Hybrid Cloud


We can't forget it's all about hybrid clouds now. Not every application will run on-prem.

Intersight is ready for hybrid cloud management.

For containers:

IKS now bridges to the public cloud

We're introducing an open-source hypervisor to lower cost in container environments where customers are encapsulating with VMs

VM integration with AWS

For VMs (AWS only)

- Start/Resume
- Stop
- Suspend
- Reset/Restart
- Launch VM Console

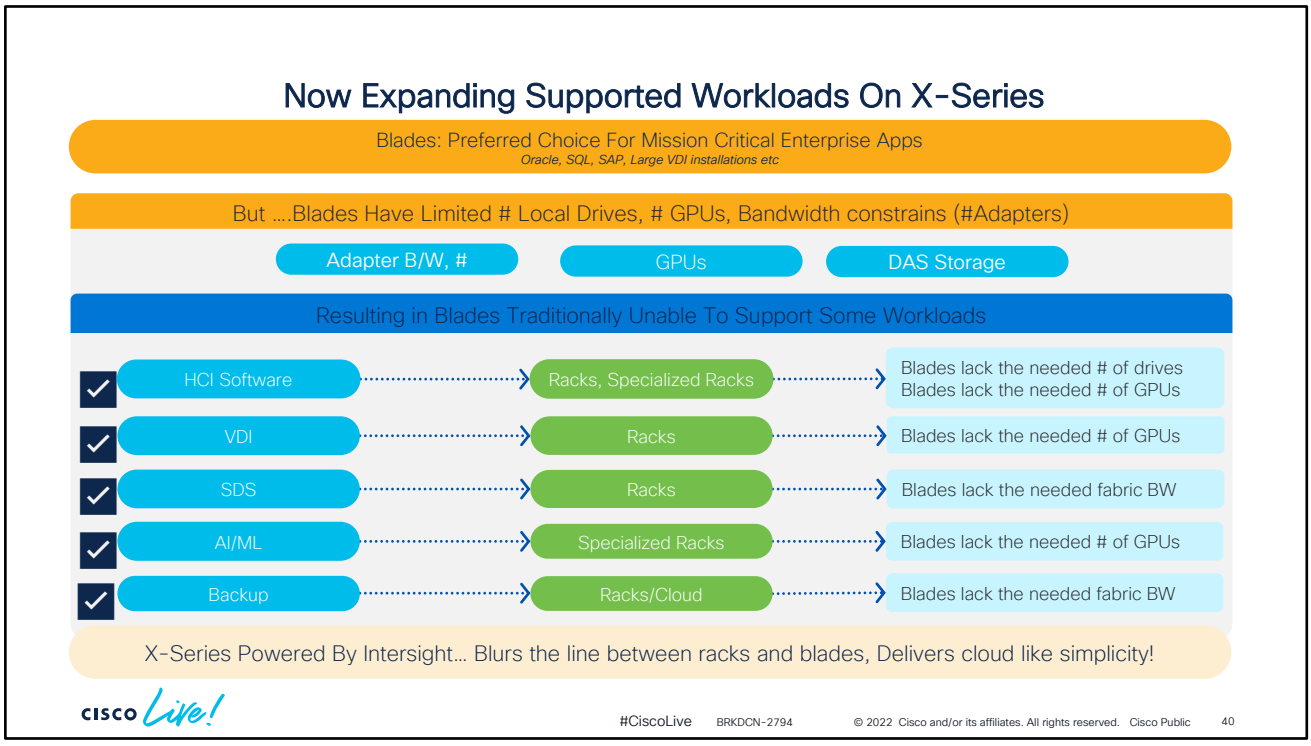


Closing it out, we see we're able to check of the needs of even more workloads using the X-Series modular system.

X-Series will continue to innovate to bring new technology into the datacenter in a disaggregated model.

Intersight will continue to innovate to bring unified management of traditional and cloud native applications.

# Also playing in the World of Solutions

X-SERIES &
FABRIC DEMOS

**Product Innovation On Display**

- X-Series
- X-Fabric
- 5th Gen Fabric Interconnect
- 5th Gen VIC
- GPU Node

**On Demo**

EXPERIENCE

### Simplicity with X-Series

Deploy X-Series Live with just a few clicks (Power of Intersight)

Flexpod: Define FlexPod with X-Series in Intersight

FlashStack: Pervasive visibility across X-Series FlashStack with Intersight

Deploy Red Hat OpenShift Container Platform with confidence (Ceph)

Automation with X-Series leveraging Intersight Ansible

EXPERIENCE

### Innovative X-Fabric: Blurring line between Rack and Blade

Simplicity of adding GPUs to Compute nodes...in few minutes!

EXPERIENCE

### Amazing Performances on X-Series

Performance out of the box

Oracle Swing bench Performance

EXPERIENCE

### New 5th Gen Unified Fabric: Simplify, Accelerate and Reduce Components

100G/200G Ethernet and NFS with a single VIC 15231

100G FC aggregate and 32G E2E FC per vHBA on a single VIC 15231

cisco Live!

# X-Series white papers

- Cisco UCS X-Series Quick Start Guide
- Cisco UCS X210c M6 Compute Node Disk I/O Characterization
- Deploy Cisco UCS X210c Compute Node with Cisco Intersight Management Mode for VDI
- FlashStack with Cisco UCS X-Series and Cisco Intersight
- FlexPod Datacenter with Cisco UCS X-Series and Cisco Intersight
- Power SAP HANA with the Cisco UCS X-Series Certified by SAP
- Deploy SAP HANA Scale-Up Appliance with UCS X-Series
- Cisco UCS and Intel SGX with Fortanix Confidential Computing Manager
- Deploy a High-Performance Standalone Oracle Database Solution: Oracle 19c on Cisco UCS X-Series
- FlexPod Datacenter with Citrix VDI and VMware vSphere 7 for up to 2500 Seats
- FlexPod XCS Solution with Cisco Intersight Platform Tech Preview
- Red Hat OpenShift Container Platform with OpenShift Data Foundation on Cisco UCS X-Series
- Cisco UCS X-Series Servers with Intel Optane Persistent Memory for Virtual Desktop Infrastructure White Paper
- Get Answers from Your Data with Cisco UCS Integrated Infrastructure for Splunk Enterprise

# Coming soon!

Accelerated VDI Papers and Solutions

X-Series Back from the Rack Episode 3

cisco Live!

# Catch up on the prequels!

BRKCLD-3010
UCS X-Series: Blurring the Line Between Rack and Blade for Modern Applications

BRKDCN-2587
Best Practices for Cloud and Compute Connectivity with the 5th Generation UCS Fabric and Intersight

cisco Live!

# Cisco Learning and Certifications

From technology training and team development to Cisco certifications and learning plans, let us help you empower your business and career. www.cisco.com/go/certs

**Pay for Learning with Cisco Learning Credits**

(CLCs) are prepaid training vouchers redeemed directly with Cisco.

## Learn

**Cisco U.**
IT learning hub that guides teams and learners toward their goals

**Cisco Digital Learning**
Subscription-based product, technology, and certification training

**Cisco Modeling Labs**
Network simulation platform for design, testing, and troubleshooting

**Cisco Learning Network**
Resource community portal for certifications and learning

## Train

**Cisco Training Bootcamps**
Intensive team & individual automation and technology training programs

**Cisco Learning Partner Program**
Authorized training partners supporting Cisco technology and career certifications

**Cisco Instructor-led and Virtual Instructor-led training**
Accelerated curriculum of product, technology, and certification courses

## Certify

**Cisco Certifications and Specialist Certifications**
Award-winning certification program empowers students and IT Professionals to advance their technical careers

**Cisco Guided Study Groups**
180-day certification prep program with learning and support

**Cisco Continuing Education Program**
Recertification training options for Cisco certified individuals

Here at the event? Visit us at **The Learning and Certifications lounge at the World of Solutions**

# Continue your education

- Visit the Cisco Showcase for related demos
- Book your one-on-one Meet the Engineer meeting
- Attend the interactive education with DevNet, Capture the Flag, and Walk-in Labs
- Visit the On-Demand Library for more sessions at www.CiscoLive.com/on-demand

The bridge to possible

# Thank you


Cisco Live!

#CiscoLive


cisco Live!
ALL IN
#CiscoLive