From Tray to Training Job: L10–L12 Manufacturing Test

Meghan Kachhi *Technical Leader – AI Networking, Cisco*

Sunny Dugar *Manager Solutions Management - AI, Supermicro*

Cisco Confidential

Executive summary



Expanding AI use cases

- Large-scale inferencing
- Explosive agentic workloads
- Trillion-parameter model training



AI infrastructure (re)consideration

- Density, power efficiency, thermal
- Supply chain constraints
- Faster time to first token

OCP Members - Partnership to simplify AI adoption

- Portfolio expansion (rack-scale and liquid-cooled)
- Expanded full-stack offer for building AI factories
- New use cases in neoclouds, sovereign clouds, large enterprises

L10 proves a component. L11 proves a rack. L12 proves a cluster.

L1–9 build components; these are the integration levels that determine whether AI infrastructure actually works.

L10
Component



Full server assembly + test

- POST, OS boot, GPU & NIC initialization
- Sustained-load burn-in and SBOM / firmware check
- GB200/300: Grace + 4 GPUs + 2 BF3 NICs

OUTPUT

A server that boots into an OS

L11
Rack



Rack integration + intra-rack validation

- Cabling, TOR, PDU and CDU thermal validation
- NVLink switch domain + NCCL across 72 GPUs
- Catches copper cartridge seating / bundle errors

OUTPUT

A complete, cabled working rack

L12
Cluster



Multi-rack assembly + full software stack

- IB / Ethernet fabric and cross-rack management
- NCCL inter-rack + GPUDirect RDMA validation
- BCM, Slurm/K8s, UFM, DCGM + training JCT

OUTPUT

A workload-ready cluster

Cisco Confidential

L10 Compute-tray acceptance

L10

Can the tray boot, initialize every expected device and sustain load under its qualified fixture?

Test family	Checks for the declared tray profile	Required evidence
Power-on and boot	BMC event log, UEFI POST, OS boot and required services	Timestamped boot and event logs
Inventory and topology	2 Grace CPUs, 4 GPUs, 4 compute HCAs and 2 BF3 DPUs in the example profile	FRU, PCIe and device manifest
GPU and network initialization	Driver bind, expected link state, RDMA mapping and BF3 operating mode	Enumeration, link and RDMA output
Firmware and SBOM	Approved versions, signed manifest, secure state and software inventory	Manifest hash and drift report
Sustained load	Mixed CPU/GPU/memory load with DCGM, ECC, XID, AER and throttle monitoring	Time-aligned health telemetry

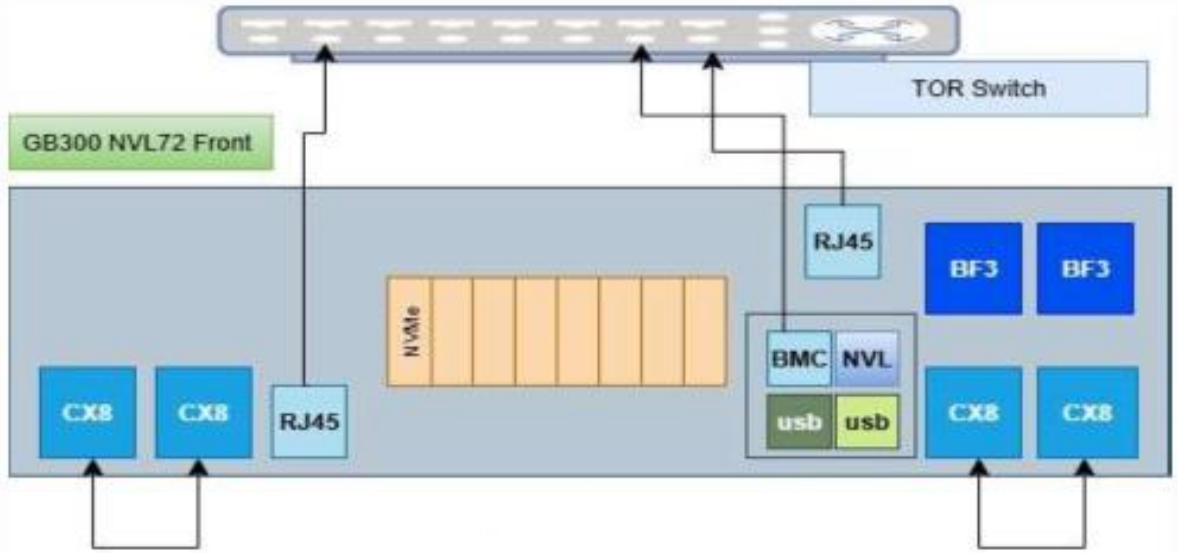
PROVES The tray is known, bootable, driver-ready and stable under the declared local workload.

DOES NOT PROVE Cable-cartridge seating, the 72-GPU NVLink domain, rack facilities or cross-rack behavior.

Cisco Confidential

L10 Every Compute Tray Proven Before Rack Build

Grace CPU, Blackwell Ultra GPUs, NVLink, ConnectX-8 and storage exercised tray by tray



Compute tray I/O under test: 4 × ConnectX-8 800G (loopback), 2 × BlueField-3, BMC, NVMe

Grace CPU & memory
TegraCpu · TegraMemory · MemorySweep · STREAM · SAT

GPU & HBM
Gpustress · Gpumem · DCGM Level 3 · GPU burn

NVLink & NVLink-C2C
NvlBwStress · TegraClink · C2C

PCIe & network
PCIe · CX-8 800G loopback · port link

Power transients
GPU pulse 1 & 4 kHz · CPU-GPU sync 1–500 Hz

Storage
SMART · DD · FIO stress

Test Item	Validation Criteria Status - OK
Dumping inforom	
Testing Checkinforom	
Testing Inventory	
Testing CxPcieProperties	
Testing SsdPcieProperties E1S	
Testing TegraCpu	
Testing TegraMemory	
Testing CpuMemorySweep	
Testing TegraClink	
Testing Gpustress	
Testing Gpumem	
Testing Pcie	
Testing Connectivity	
Testing NvlBwStress	
Testing NvlBwStressBg610	
Testing C2C	
Testing C2CGpuPulsePower1kHz	
Testing C2CGpuPulsePower4kHz	
Testing CpuGpuSyncPulsePower1Hz50duty CpuGpuSyncPulsePower50Hz50duty CpuGpuSyncPulsePower500Hz50duty	

Exit gate: every item OK → tray released to rack build

- ✓ Hardware Inventory Validation
- ✓ NVIDIA MFG Diagnostic Test
- ✓ NVIDIA-DCGM Level 3
- ✓ Mellanox Port Link Status Check
- ✓ CX-8 800G Loopback Test
- ✓ Disk Smart Data Check/Disk DD Stress /Disk FIO Stress
- ✓ System Idle Test
- ✓ GPU Burn Test
- ✓ STREAM Memory Test and StressAPPTest (SAT)

Following OCP ORv3 MGX 19” rack specifications



Delivered as full length tray with DC power supplies and busbar connection kit assembled for Cisco N93108TC-FX3 and ready to be inserted into the MGX rack

Includes rail kit for the tray

https://www.opencompute.org/documents/mgx-accelerated-computing-rack-and-trays-specification-101024-pdf





Cisco Confidential

L10 proves a component. L11 proves a rack. L12 proves a cluster.

L1–9 build components; these are the integration levels that determine whether AI infrastructure actually works.

L10

Server

Full Component assembly + test

- POST, OS boot, GPU & NIC initialization
- Sustained-load burn-in and SBOM / firmware check
- GB200/300: Grace + 4 GPUs + 2 BF3 NICs

OUTPUT

A server that boots into an OS

L11

Rack



Rack integration + intra-rack validation

- Cabling, TOR, PDU and CDU thermal validation
- NVLink switch domain + NCCL across 72 GPUs
- Catches copper cartridge seating / bundle errors

OUTPUT

A complete, cabled working rack

L12

Cluster

Multi-rack assembly + full software stack

- IB / Ethernet fabric and cross-rack management
- NCCL inter-rack + GPUDirect RDMA validation
- BCM, Slurm/K8s, UFM, DCGM + training JCT

OUTPUT

A workload-ready cluster

Cisco Confidential

L11 rack acceptance

L11

Does the complete NVL72 rack operate as one physical, thermal and 72-GPU NVLink domain?

Domain	What the rack test checks	Evidence
Physical	18 compute trays, 9 NVLink switch trays, cable cartridges and declared copper map	Rack manifest and seating record
Management	Two TOR switches, all BMC/HMC endpoints and rack services	Reachability and interface counters
Power	PDU or power-shelf telemetry, step load, alarms and redundancy state	Power trace and event log
Cooling	Approved pressure/flow profile, supply and return temperatures, leak sensors	Commissioning and telemetry record
NVLink domain	72-GPU membership, expected link map, switch health and NCCL all-reduce	Topology, errors and collective JSON

PROVES The rack is fully seated, reachable, thermally stable and able to run a correct 72-GPU collective.

DOES NOT PROVE Cross-rack transport, provisioning, scheduler placement or training job completion time.

L11 The Integrated NVL72 Rack, Tested as One System

18 compute trays + 9 NVLink switch trays · 72 GPUs in one NVLink domain · power shelves and CDU under load



- ✓ NVIDIA Compute Tray Diagnostic Test
- ✓ NVIDIA Switch Tray Diagnostic Test
- ✓ NVIDIA Cable Cartridge Diagnostic Test
- ✓ NCCL Node and Rack Level Stress Test
- ✓ GPU Burn Test
- ✓ FIO Test

1. **Firmware alignment**
   Compute tray, NVLink switch tray, power shelf and BMC set to the approved matrix, then verified
2. **NVIDIA Field Diagnostics**
   All 18 compute and 9 switch trays; full 72-GPU NVLink connectivity
3. **Stress & stability**
   GPU stress, NVIDIA Diagnostic Test and FIO with live thermal and power tracking
4. **Network**
   RDMA perftest (latency, bandwidth, message rate), NCCL, OSU
5. **Power**
   Repeated DC power cycles; power-shelf telemetry
6. **Liquid cooling**
   CDU flow, pressure, supply/return temperature, leak, humidity, dew point

Sample L11 Report: What Ships With Every Rack

Auto-generated HTML · PDF · JSON, traceable to rack, tray and component serial numbers



Rack report: cover and table of contents

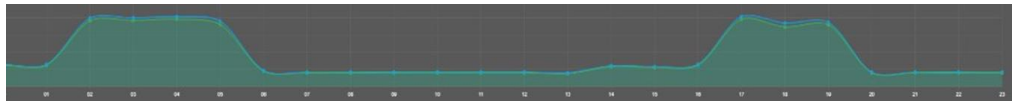
Firmware update result, all 18 compute trays

TOTAL SYSTEMS	UPDATING	PASSED	FAILED	PASS RATE
18	0	18	0	100.0%

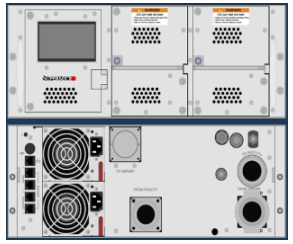
NVLink switch firmware

ITEM	STATUS
system_image	Latest
ASIC	Latest
BIOS	Latest
BMC	Latest
CPLD1	Latest
CPLD2	Latest
CPLD3	Latest
CPLD4	Latest
EROT	Latest
EROT-ASIC1	Latest

Rack power across burn-in cycles



CDU health during the run



Coolant temp ✓ | Coolant flow ✓

Loop pressure ✓ | Leak/humidity ✓

Report sections

Rack image | Inventory | BMC | Firmware | Field Diag | NCCL | FIO | Power shelf | CDU

Cisco Confidential

L10 proves a component. L11 proves a rack. L12 proves a cluster.

L1–9 build components; these are the integration levels that determine whether AI infrastructure actually works.

L10

Component

Full server assembly + test

- POST, OS boot, GPU & NIC initialization
- Sustained-load burn-in and SBOM / firmware check
- GB200/300: Grace + 4 GPUs + 2 BF3 NICs

OUTPUT

A server that boots into an OS

L11

Rack

Rack integration + intra-rack validation

- Cabling, TOR, PDU and CDU thermal validation
- NVLink switch domain + NCCL across 72 GPUs
- Catches copper cartridge seating / bundle errors

OUTPUT

A complete, cabled working rack

L12

Cluster



Multi-rack assembly + full software stack

- Ethernet fabric and cross-rack management
- NCCL inter-rack + GPUDirect RDMA validation
- BCM, Slurm/K8s, DCGM + training JCT

OUTPUT

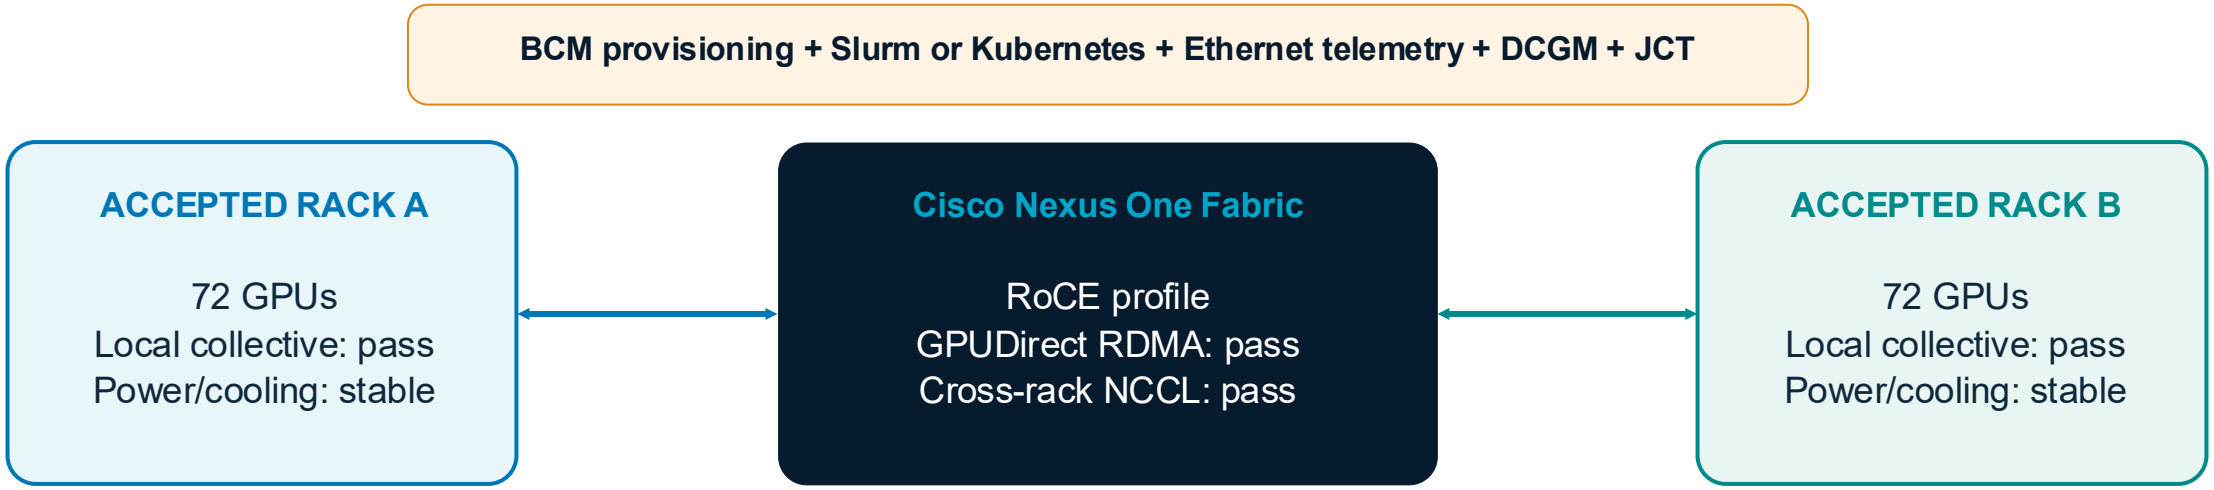
A workload-ready cluster

Cisco Confidential

L12 cluster acceptance

L12

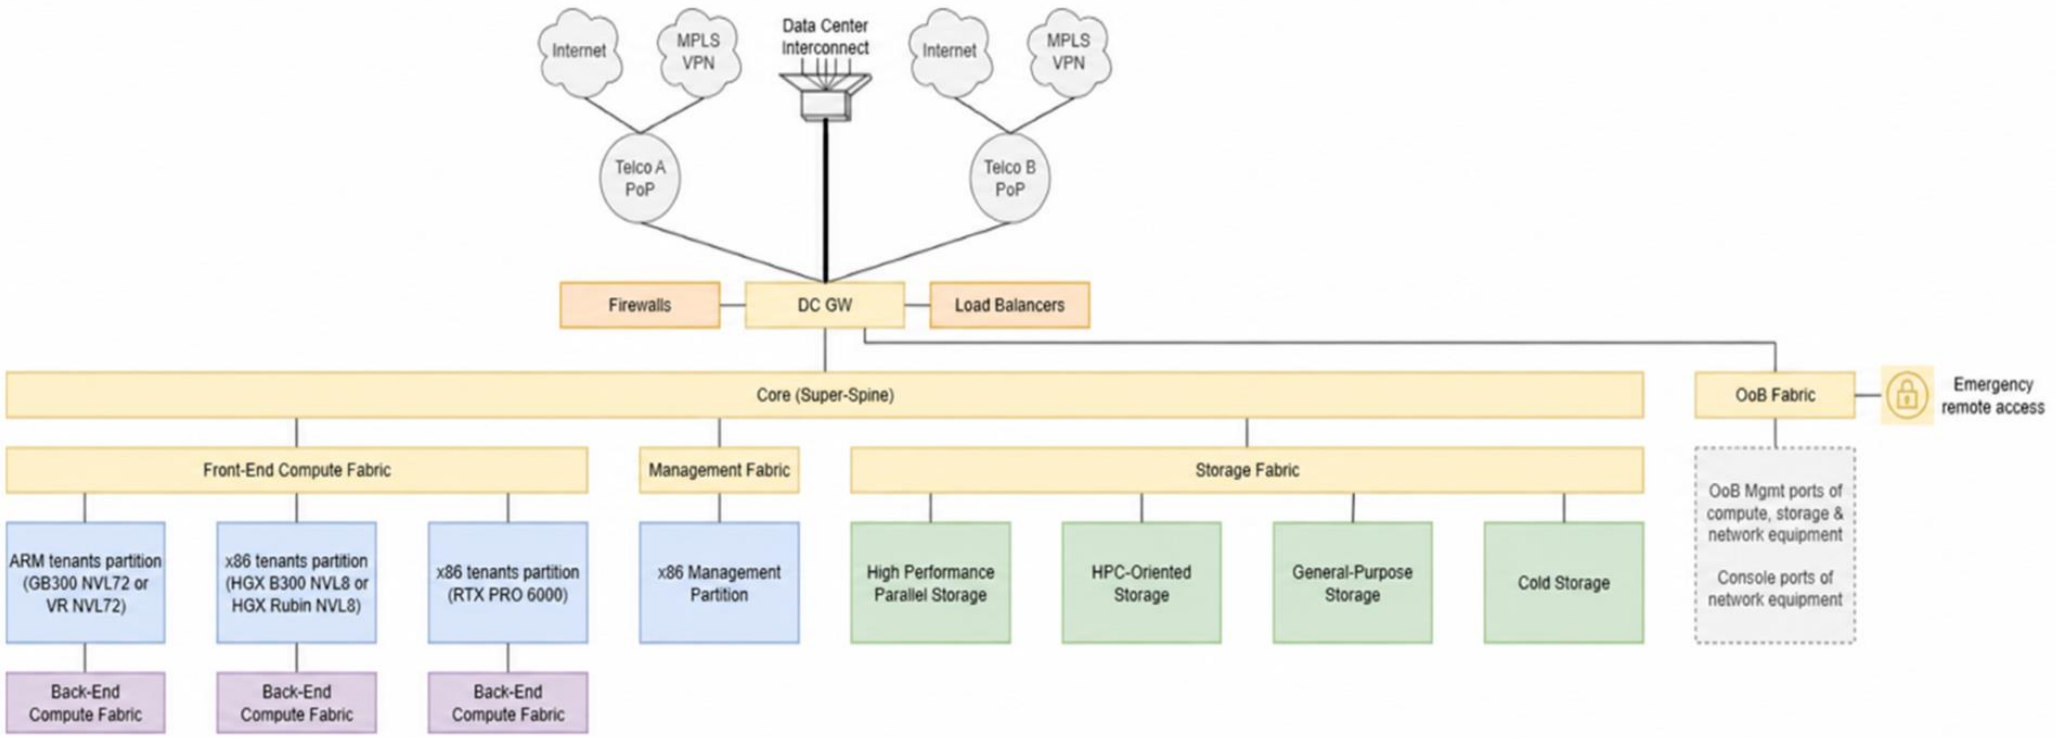
Can accepted racks communicate through the production fabric and complete a scheduled training canary?



PROVES The tested topology can provision, communicate, observe and complete the defined training gate.

DOES NOT PROVE Long-term availability, every workload, or performance beyond the validated topology and profile.

L12 Cluster Validation: Hardware Stack Overview



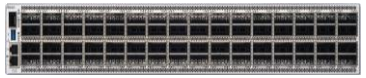
Cisco Confidential

Cisco Nexus One - AI Networking Portfolio (SONiC or NX-OS)

Frontend

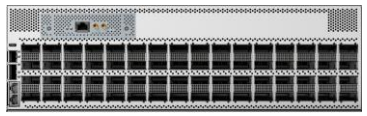
Air Cooled

N9364E-SG2X-O



Silicon One G200X – 2RU - 64x800Gbps OSFP

N9364F-SG3-EA



Silicon One G300 – 3RU - 64x1.6Tbps OSFP

Liquid Cooled

N9364F-SG3-RL

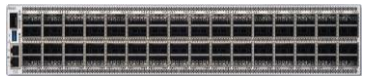


Silicon One G300 – 3OU (ORV3) - 64x1.6Tbps OSFP

Backend

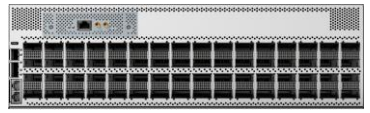
Air Cooled

N9364E-SG2X-O



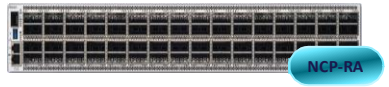
Silicon One G200X – 2RU - 64x800Gbps OSFP

N9364F-SG3-EA



Silicon One G300 – 3RU - 64x1.6Tbps OSFP

N9164E-NS4-O



Spectrum-4 – 2RU - 64x800Gbps OSFP

Liquid Cooled

N9364F-SG3-RL



Silicon One G300 – 3OU (ORV3) - 64x1.6Tbps OSFP

N9164F-NS6-ML



Spectrum-6 – 2RU (MGX) - 64x1.6Tbps OSFP

Storage

N9K-C9332D-GX2B



CloudScale – 1RU – 32x400Gbps QSFP-DD and 2x1/10G SFP+

Management

N93108TC-FX3-MGX



CloudScale – 1RU - 48x100M/1/10GT and 6x40/100G QSFP28
Integrated MGX tray with busbar connector

ND-CLUSTERG5S



Nexus Dashboard 3 Nodes Cluster

Nexus Hyperfabric (SaaS)

Many more…

https://www.cisco.com/c/en/us/products/collateral/data-center-networking/n9000-gb300-nvl72-cra.html

Sample IBPerf Test

Server-1

```
ib_write_bw -a -d mlx5_1 --report_gbits
-q 16
```

Server-2

```
ib_write_bw 192.168.101.25 -a -d mlx5_2
--report_gbits -q 16
```

-a: Runs the test across all available message sizes to provide a comprehensive performance profile.

-d mlx5_1: Specifies the RDMA-capable network device to use, in this case, the second Mellanox/NVIDIA adapter found by the system (mlx5_1).

--report_gbits: Changes the output unit from the default Megabytes per second (MB/s) to the more network-centric Gigabits per second (Gbps).

-q 16: Sets the number of Queue Pairs (QPs) to 16. Using multiple QPs enables parallel data transfers, which is necessary to saturate high-speed links.

```
---------------------------------------------------------------------------------------
                    RDMA_Write BW Test
 Dual-port       : OFF          Device         : mlx5_1
 Number of qps   : 16           Transport type : IB
 Connection type : RC           Using SRQ      : OFF
 PCIe relax order: ON           Lock-free      : OFF
 ibv_wr* API     : ON           Using DDP      : OFF
 CQ Moderation   : 100
 CQE Poll Batch  : 16
 Mtu             : 4096[B]
 Link type       : Ethernet
 GID index       : 3
 Max inline data : 0[B]
 rdma_cm QPs     : OFF
 Data ex. method : Ethernet
---------------------------------------------------------------------------------------
 local address: LID 0000 QPN 0x07c5 PSN 0xf04146 RKey 0x203d00 VAddr 0x007fffefac3000
 GID: 00:00:00:00:00:00:00:00:00:00:255:255:01:01:101:25
remote address: LID 0000 QPN 0x0880 PSN 0xc8558f RKey 0x1fff00 VAddr 0x007ffff7226000
 GID: 00:00:00:00:00:00:00:00:00:00:255:255:01:01:201:25
---------------------------------------------------------------------------------------
 #bytes     #iterations    BW peak[Gb/sec]    BW average[Gb/sec]   MsgRate[Mpps]
 8388608    80000            391.53             391.53               0.005834
---------------------------------------------------------------------------------------
```

Snipped output for better readability

Sample NCCL Communication Test Result

```
# Collective test starting: all_reduce_perf
# nThread 1 nGpus 1 minBytes 524288 maxBytes 17179869184 step: 2(factor) warmup iters: 1 iters: 50 agg iters: 1 validation: 1 graph: 0
#
# Using devices
#  Rank  0 Group  0 Pid 7563 on   rack-a-node00 device  0 [0008:06:00] NVIDIA GB300
#  Rank  1 Group  0 Pid 7564 on   rack-a-node00 device  1 [0009:06:00] NVIDIA GB300
<Truncated>
#  Rank 142 Group  0 Pid 3623 on  rack-b-node36 device  2 [0018:06:00] NVIDIA GB300
#  Rank 143 Group  0 Pid 3624 on  rack-b-node36 device  3 [0019:06:00] NVIDIA GB300
<Truncated>
#
#                                                     out-of-place                       in-place
#       size         count      type   redop    root     time   algbw   busbw #wrong     time   algbw   busbw #wrong
#        (B)    (elements)                               (us)  (GB/s)  (GB/s)            (us)  (GB/s)  (GB/s)
<Truncated>
  8589934592    2147483648     float     sum      -1    23161  353.24  709.93      0    23123  353.68  710.18      0
 17179869184    4294967296     float     sum      -1    46159  374.93  739.46      0    46091  375.11  741.99      0
# Out of bounds values : 0 OK
# Avg bus bandwidth    : 300.901
#
# Collective test concluded: all_reduce_perf
```

Sample Output: Actual results may completely vary

Metric	Definition	What to Look For
algbw (Algorithm BW)	message_size / time — raw throughput	Useful but not the key metric
busbw (Bus BW)	Adjusted for collective algorithm efficiency. For AllReduce: algbw × 2(n-1)/n	This is the metric that matters — it represents actual NIC utilization
#wrong	Correctness check failures	Must be 0 — non-zero means data corruption (bad NIC, bad cable, lossy fabric)

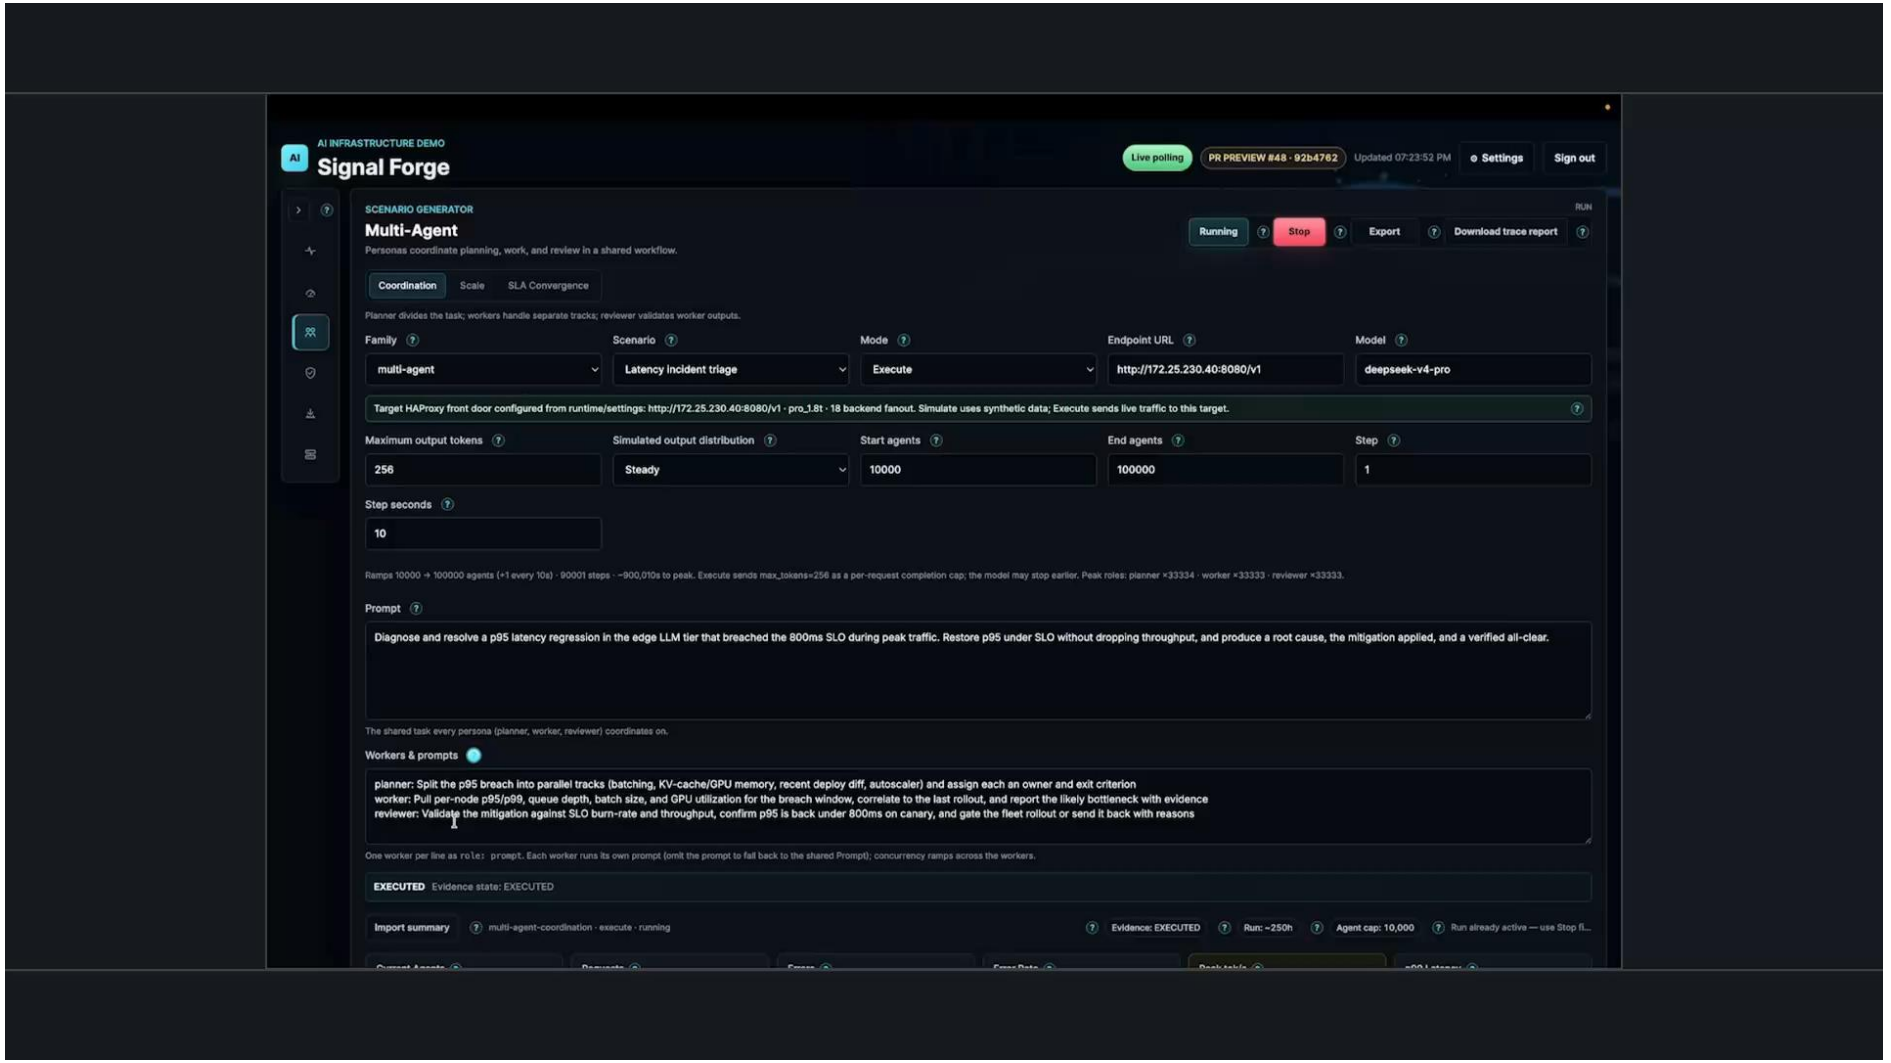
Cisco Confidential

Demo

Deploy Agentic AI in Rackscale Systems to fully populate and observe signals from rack

Cisco Confidential

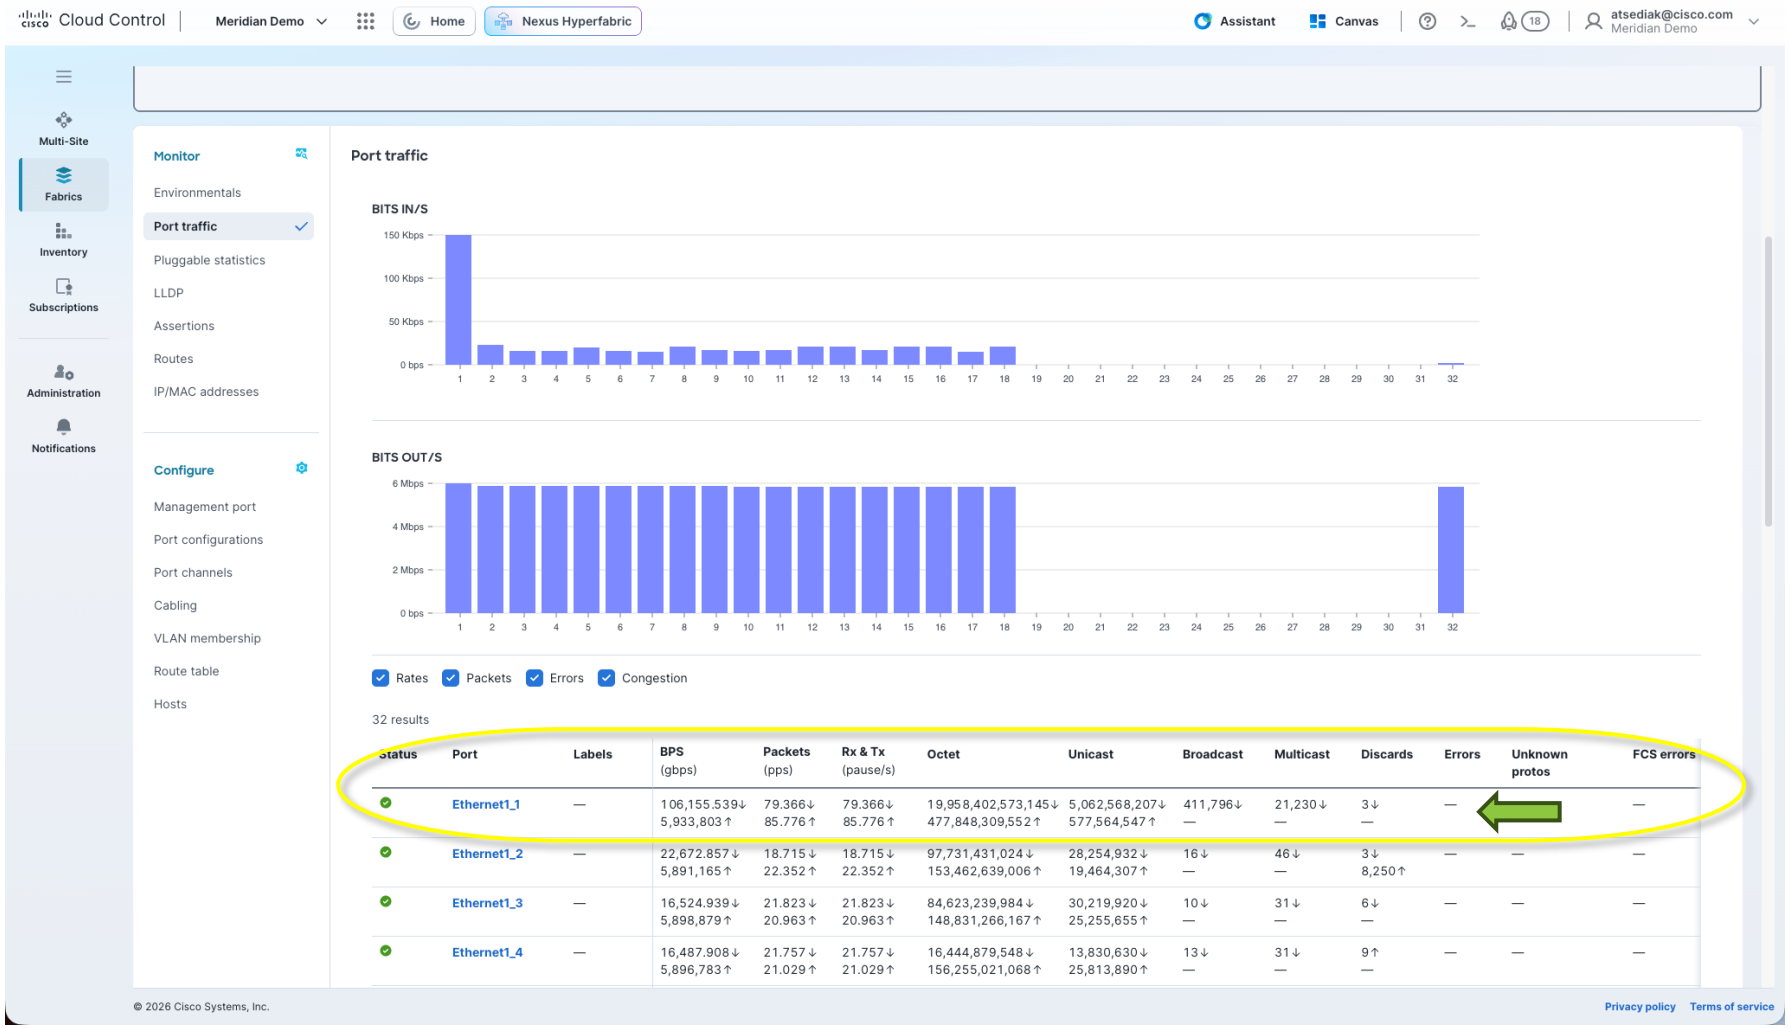
Cisco Nexus One and Splunk Observability Cloud



Cisco Confidential

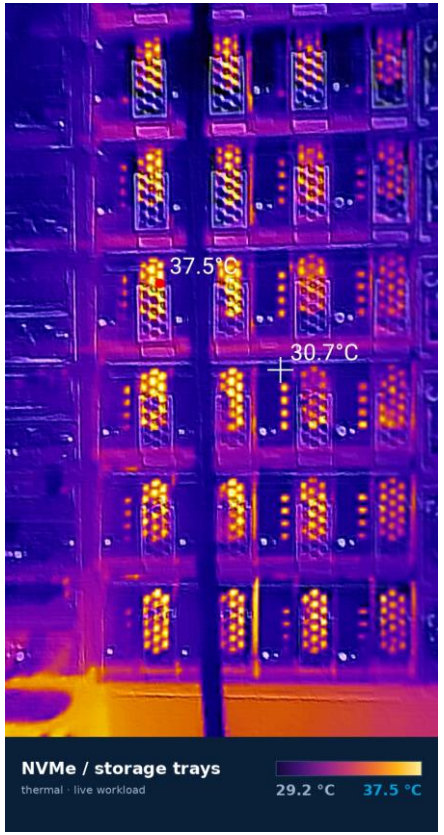
Cisco Nexus One: Observing Fabric under load

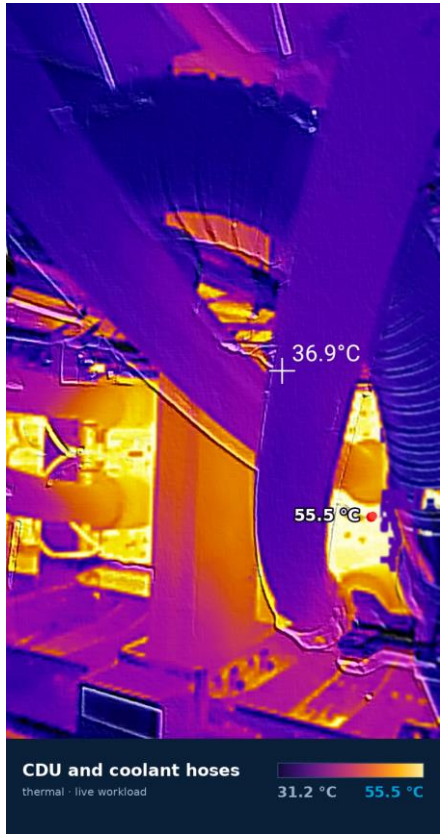
- 18 x 100G BlueField-3 on ns-leaf001.
- PFC + ECN. RoCE v2 on Ethernet you already operate.
- A 1.6T model loads over 100G in minutes. That is why north-south fabric matters.
- GPU-to-GPU communication inside the rack uses the NVLink.

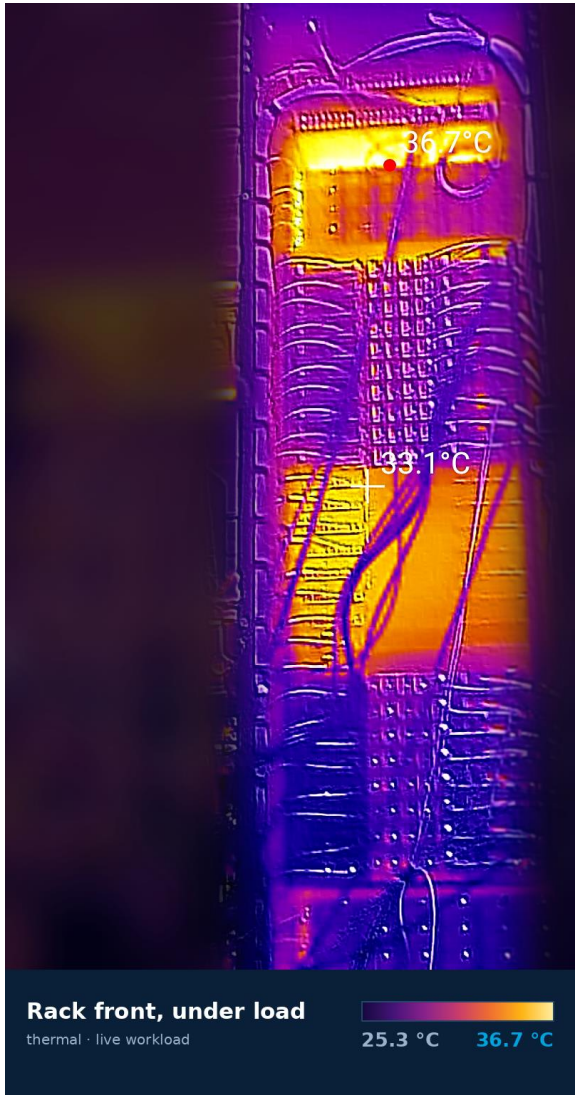


Observing the raw signals simplified with Cisco Nexus One

Every event on the NVL72, correlated.
Observing the rack's raw signals — GPU, power, thermal, fabric — into one timeline you can actually read.







Call to Action

- Visit us at Cisco Booth #C59 and Supermicro Booth #B4
- Open Compute Project Community: https://www.opencompute.org/community
- Open Compute Project Products: https://www.opencompute.org/products

Cisco Confidential

Thank You!



Cisco Confidential