

Experiences with INAM on SDSC Systems

Mahidhar Tatineni

Director, User Services, SDSC

SC23 OSU Booth (#1581) Talk

Nov 14, 2023



Outline

- Overview
- INAM on Comet
 - System architecture
 - Network and Job Level Views
 - Notifications
 - Use case illustrating value of INAM
- Future work
 - Planned installation on Expanse
 - Expanse architecture
- Summary

Overview

- OSU InfiniBand Network Analysis and Monitoring (INAM) tool provides an open-source option to monitor IB networking
- On Comet, CW3E workloads can involve both IO intensive and communication intensive workloads. Sometimes this leads to performance issues. The goal is to use both historical and live data to troubleshoot such issues.
- Monitoring of network fabric health, utilization with notification thresholds set for errors and utilization.
- Enable Lustre traffic monitoring to get a handle on the IO aspect of workloads
- Use MPI information from jobs in combination with Lustre info to evaluate impact on network

Outline

- Overview
- INAM on Comet
 - System architecture
 - Network and Job Level Views
 - Notifications
 - Use case illustrating value of INAM
- Future work
 - Planned installation on Expanse
 - Expanse architecture
- Summary

Comet Supercomputer

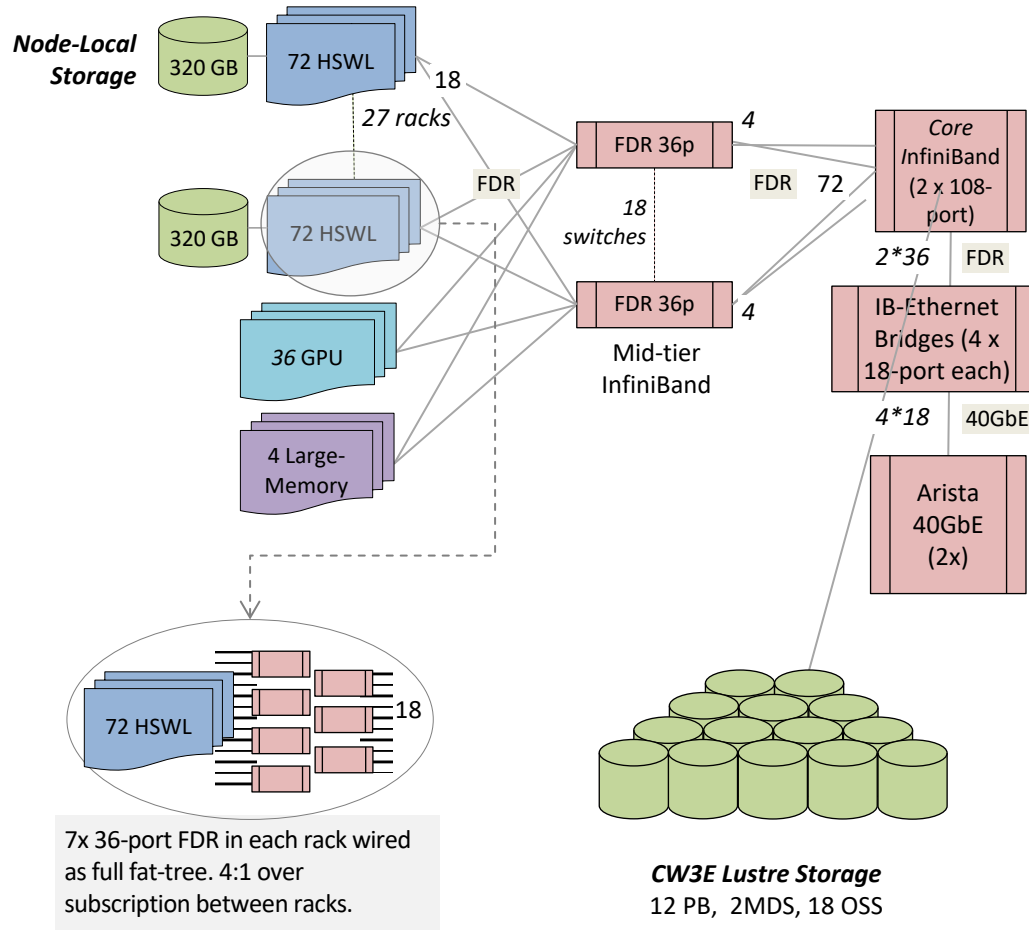
- *Center For Western Weather and Water Extremes (CW3E)* has made exclusive use of Comet starting *July 2021* after it was retired from NSF XSEDE/ACCESS service
- Comet is being used for West-WRF ensemble runs during wet-seasons for near real-time forecasts. In addition, it used by CW3E researchers for several research projects.
- *Managed by the San Diego Supercomputer Center (SDSC)*
- Represents over 1 billion core hours (SUs) of computing over the period of 2.5 years from July 2021 through December 2023



Highlights: ~440 M CPU hours / year, ~1.2 M GPU hours / year, >12 PB of storage

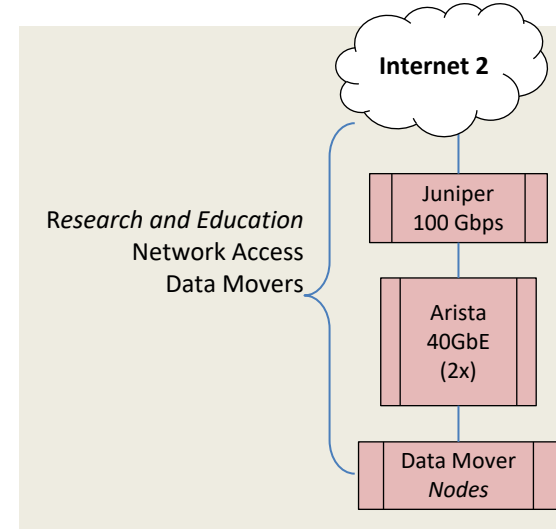
Comet Network Architecture

InfiniBand compute, Ethernet Storage



7x 36-port FDR in each rack wired as full fat-tree. 4:1 over subscription between racks.

CW3E Lustre Storage
12 PB, 2MDS, 18 OSS




Additional Support Components
(not shown for clarity)

- Ethernet Mgt Network (10 GbE)
- NFS Servers for Home Directories
- Virtual Image Repository
- Gateway/Portal Hosting Nodes
- Login Nodes
- Rocks Management Nodes


INAM Experience on Comet

- OSU INAM group helped build a custom INAM RPM for Comet OS and OFED stack.
- There was an issue parsing the switch map file. Worked ok after we removed all of the blank lines in our map file. Suggestion to fix as the map file can be parsed as is by other tools.
- At present the install is setup for internal use. Will need user/auth setup additions to enable end user access.
- PhantomJS took some time to build up cached info initially, but web service has been good since.
- We will be installing INAM for Expanse as well. Plan to evaluate use of ClickHouse DB backend.

Full System Network View from INAM

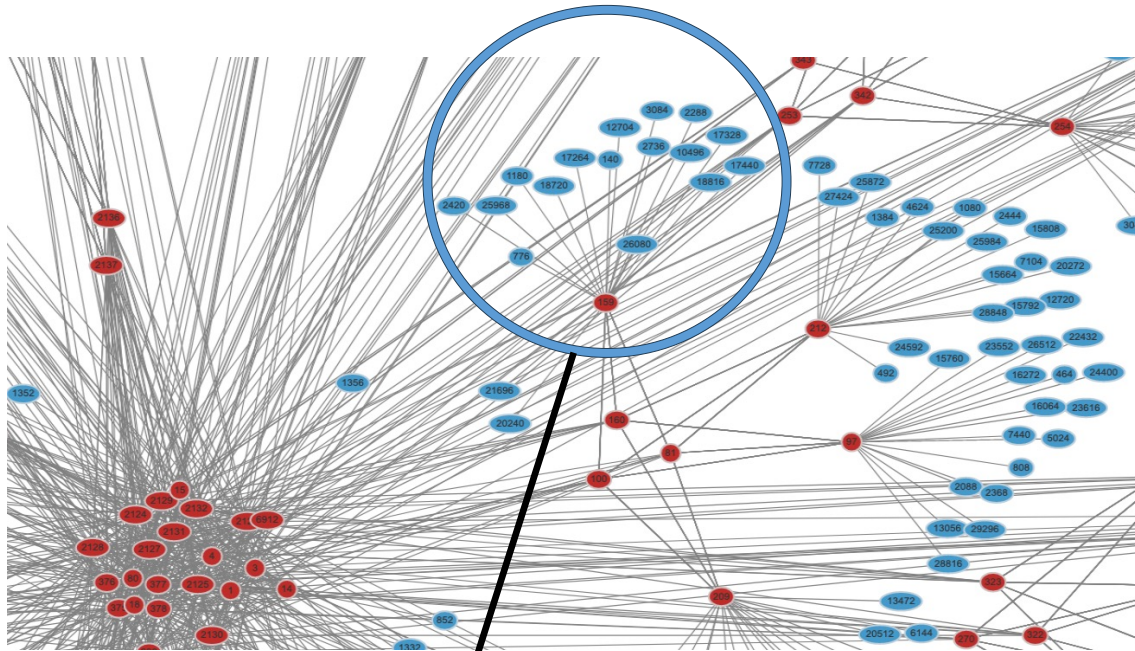


Rack Level Network View from INAM

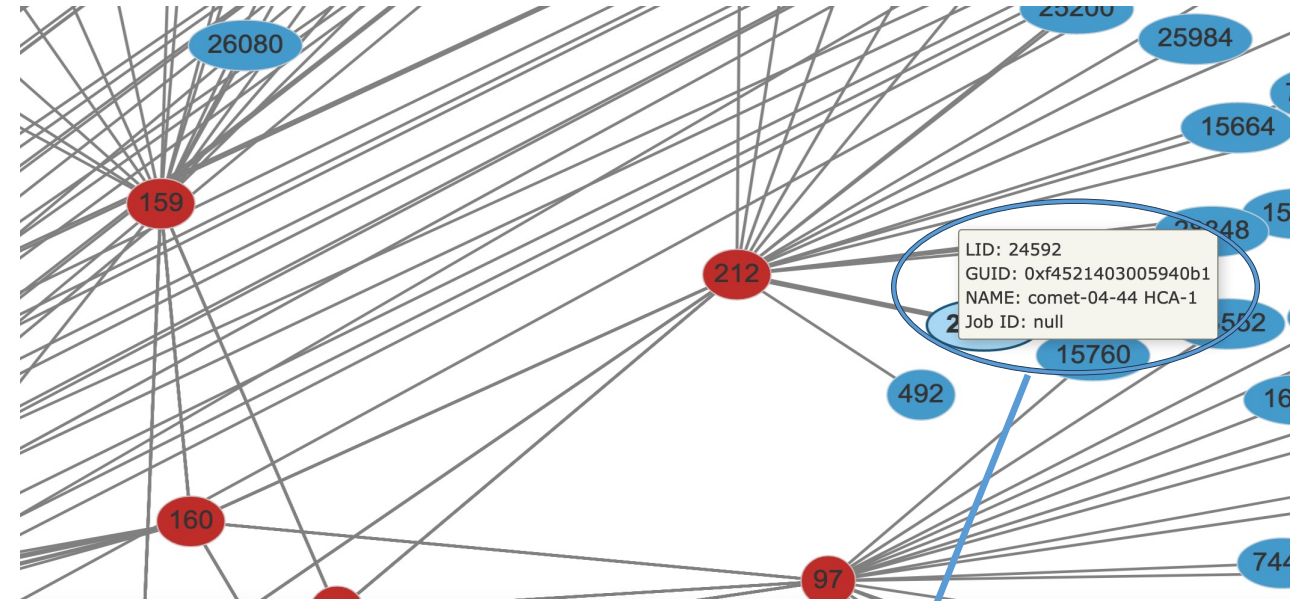


- *18 nodes per switch*
- *4 switches at first level that feed into 3 switches in the next tier*

Network Views from INAM



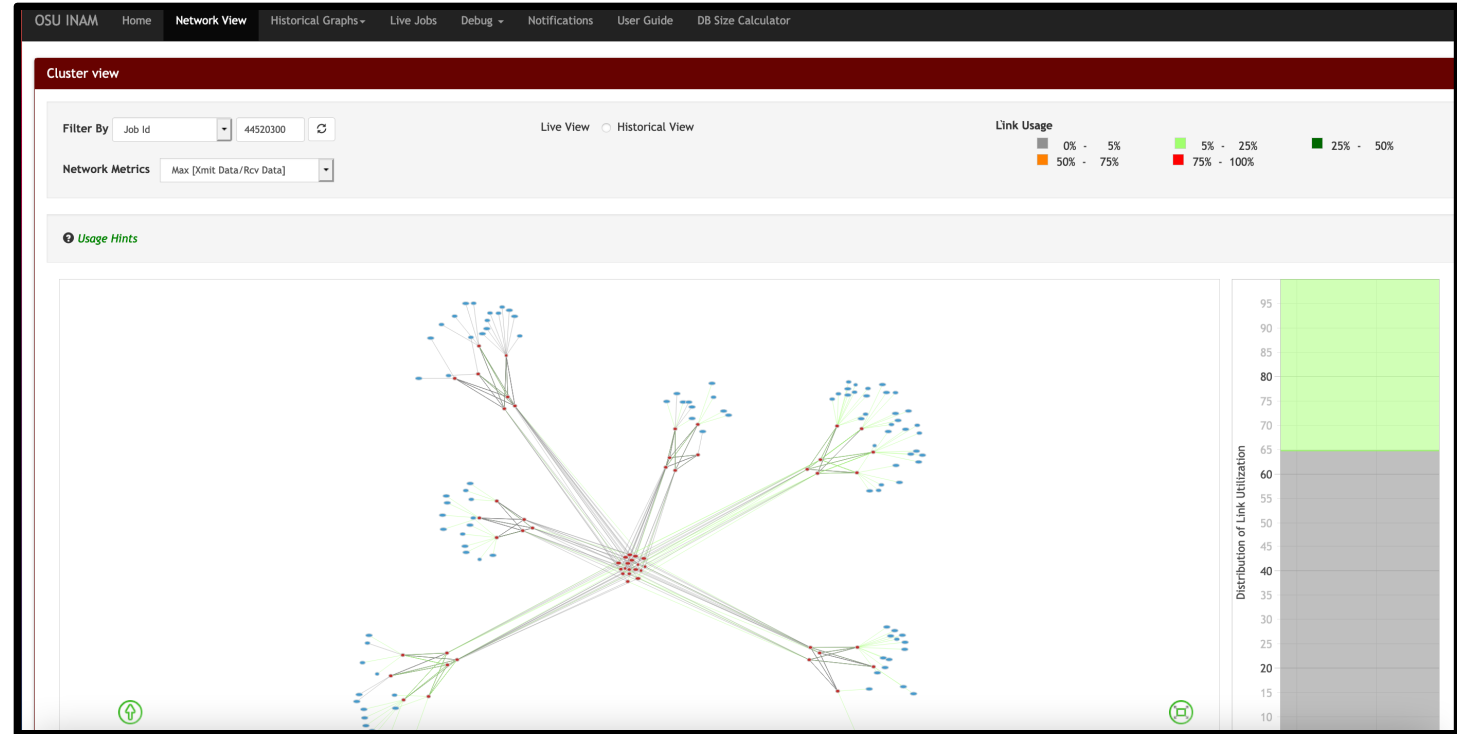
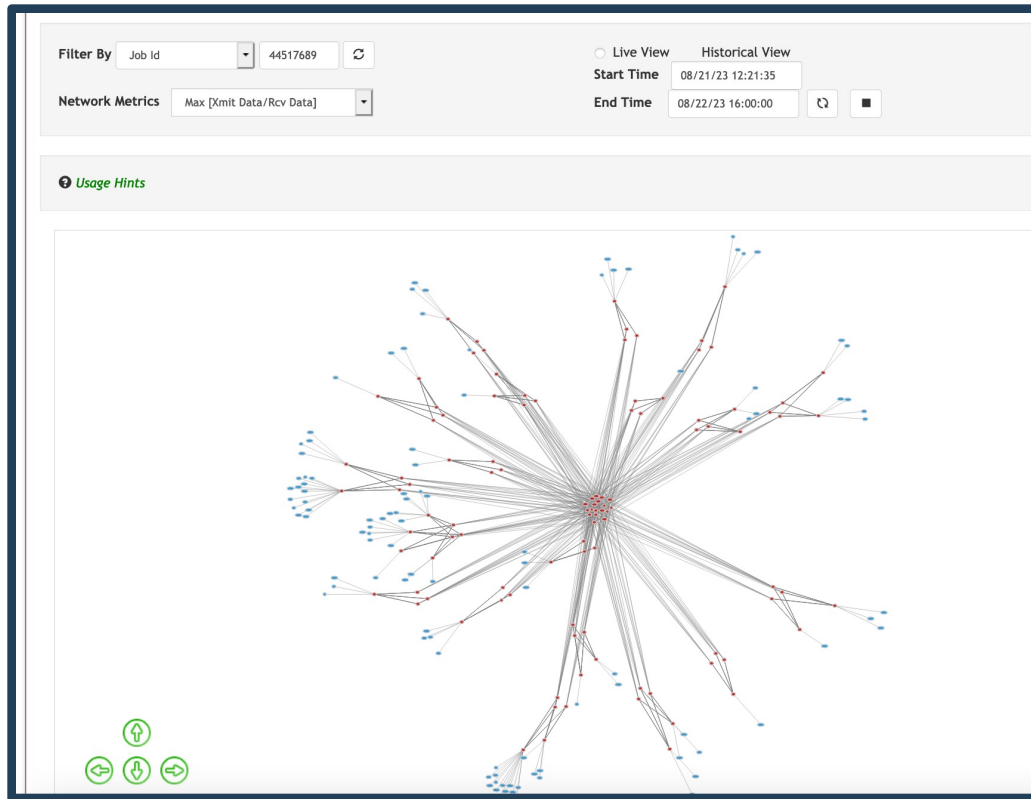
Single Switch with 18 nodes



Node LID, GUID, Name info

LID: 24592
GUID: 0xf4521403005940b1
NAME: comet-04-44 HCA-1
Job ID: null

Job level view from INAM



Notifications from INAM

OSU INAM Home Network View Historical Graphs Live Jobs Debug Notifications User Guide DB Size Calculator

Notifications & Criteria

Notifications

Search [] [] [] []

Notification ID	Criteria ID	Category	Metric	Condition	Threshold Value	Time	Additional Information	Action
2	1	Port Counters	Symbol errors	Greater than	50	22 Aug 2023 15:33:50	MF0;switch-928eee:5X6036/U1 (0x248a070300f25580):11 - comet-34-08 HCA-1 (0x248a0703005c39e1), MF0;sw...	
1	3	Port Counters	Link utilization (percentage)	Greater than	95	22 Aug 2023 09:25:57	comet-dr1-L01 (0xf4521403006ae330):7 - cw3e-oss-1-5 HCA-1 (0x0c42a1030079f00c)	

Showing 1 to 2 of 2 rows

Notification Criteria

Add Notification Criteria +

Search [] [] [] []

Criteria ID	Category	Metric	Comparison	Threshold Value	Is Recurring	Action
1	Port Counters	Symbol errors	Greater than	50	✓	
2	Port Counters	Link downed	Greater than	1	✓	
3	Port Counters	Link utilization (percentage)	Greater than	95	✗	

Showing 1 to 3 of 3 rows

Notifications from INAM

The screenshot shows the OSU INAM web application interface. The browser address bar displays `comet-fe5.sdsc.edu:8080/notifications`. The navigation menu includes **OSU INAM**, Home, Network View, Historical Graphs, Live Jobs, Debug, Notifications, User Guide, and DB Size Calculator. The main content area is titled "Notifications & Criteria" and contains a "Notifications" table. A modal window titled "Additional Information" is open, displaying details for Notification ID 2.

Notifications & Criteria

Notifications

Notification ID	Criteria ID	Category	Metric
2	1	Port Counters	Symbol e
1	3	Port Counters	Link utili (percenta

Showing 1 to 2 of 2 rows

Additional Information

Notification ID: 2

Additional Information

- MF0;switch-928eee: SX6036/U1 (0x248a070300f25580):11 - comet-34-08 HCA-1 (0x248a0703005c39e1)
- MF0;switch-928eee: SX6036/U1 (0x248a070300f25580):12 - comet-34-07 HCA-1 (0x248a0703005c39f1)

Showing 1 to 2 of 2 rows

Close

Notifications can identify issues

In Port Counters , linkusage exceeded threshold value - 95.0.

Link information

IB-R24-L3-01 (0xf452140300f635a0):36 - comet-dr1-L04 (0xf4521403006ae2d0):2
comet-dr1-L04 (0xf4521403006ae2d0):23 - comet-dr1-S01 (0xe41d2d0300002500):23
MF0;comet-dr1: SX6506/L06/U1 (0xf4521403008b7120):16 - cw3e-oss-3-5 HCA-1 (0x0c42a10300825a84)
MF0;comet-dr1: SX6506/L06/U1 (0xf4521403008b7120):17 - cw3e-oss-3-3 HCA-1 (0x0c42a1030079c610)
comet-dr1-L01 (0xf4521403006ae330):17 - cw3e-oss-2-1 HCA-1 (0x0c42a10300825ad4)
comet-dr1-L02 (0xf4521403006ae090):13 - cw3e-oss-2-5 HCA-1 (0x0c42a10300825a64)
comet-dr1-L01 (0xf4521403006ae330):11 - cw3e-oss-1-3 HCA-1 (0x0c42a10300825ac8)
comet-dr1-L02 (0xf4521403006ae090):7 - cw3e-oss-2-2 HCA-1 (0x0c42a10300825a60)
comet-dr1-L01 (0xf4521403006ae330):1 - cw3e-oss-1-1 HCA-1 (0x0c42a10300825a94)
IB-R09-L3-02 (0xf452140300f74d50):35 - comet-dr1-L04 (0xf4521403006ae2d0):3
comet-dr1-S02 (0xe41d2d0300002520):2 - comet-dr1-L01 (0xf4521403006ae330):26
comet-dr1-L04 (0xf4521403006ae2d0):28 - comet-dr1-S02 (0xe41d2d0300002520):22
...
...
comet-dr1-L02 (0xf4521403006ae090):9 - cw3e-oss-2-3 HCA-1 (0x0c42a103006b1810)
comet-dr1-L01 (0xf4521403006ae330):13 - cw3e-oss-1-4 HCA-1 (0x0c42a10300825a74)
comet-dr1-L01 (0xf4521403006ae330):9 - cw3e-oss-1-2 HCA-1 (0x0c42a10300825a4c)
comet-dr1-L03 (0xf4521403006ae390):23 - comet-dr1-S01 (0xe41d2d0300002500):17
MF0;comet-dr1: SX6506/L06/U1 (0xf4521403008b7120):13 - cw3e-oss-3-6 HCA-1 (0x0c42a1030079f000)
MF0;comet-dr1: SX6506/L06/U1 (0xf4521403008b7120):14 - cw3e-oss-3-4 HCA-1 (0x0c42a103006b1804)

Connections from mid-tier to director switch

Connections from leaf to spine in director switch

Connections from director switch to Lustre OSSs

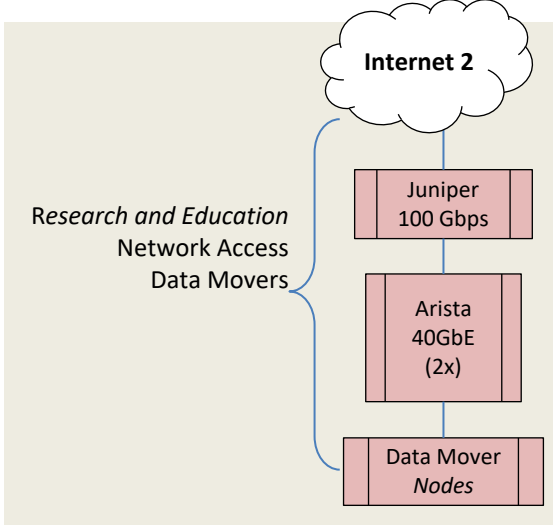
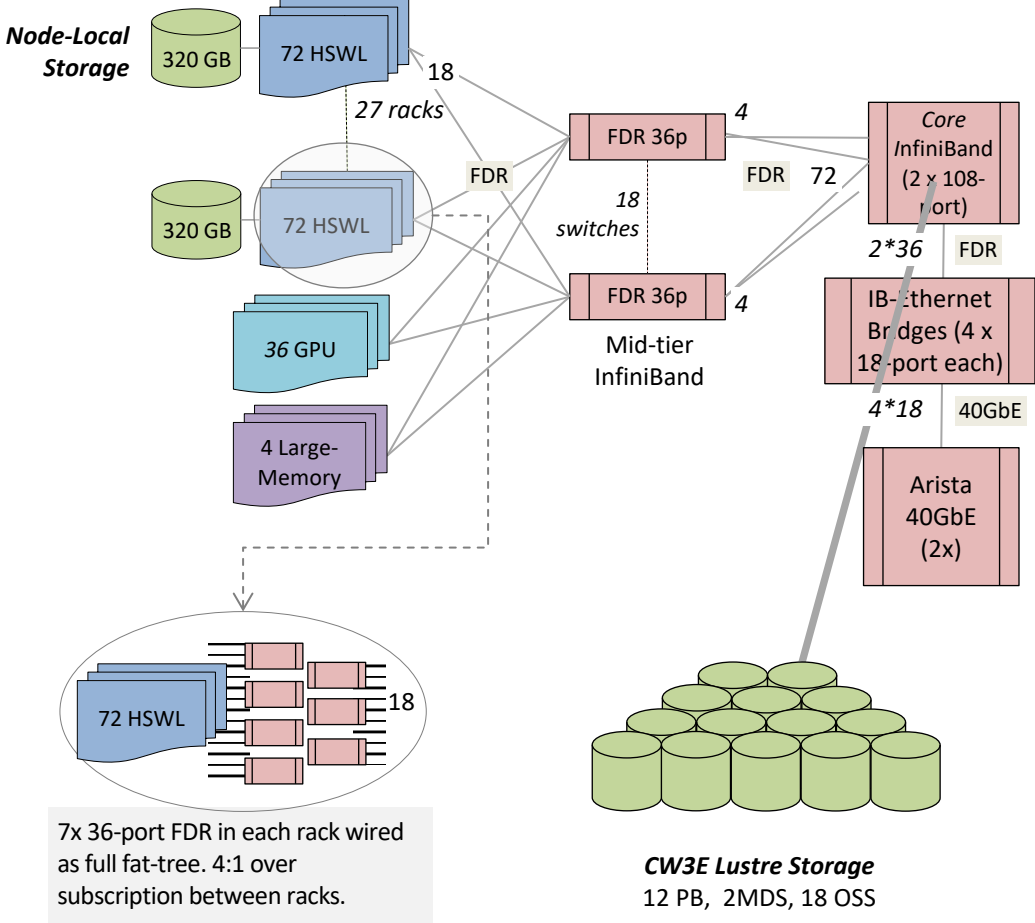
-OSU INAM

Notifications from INAM

- These notifications were received when the intensive near real time (NRT) forecast simulations were being run.
- The I/O patterns and load were saturating leaf to spine links in the director switch and also some of the director switch to OSS links were saturating.
- This was impacting I/O performance for the NRT runs and in some cases that lead to missed targets (time bound).
- Based on information from the INAM notifications and network monitoring data, decision was made to move the OSS links into the mid-tier switches.
- This removes the leaf to spine bottlenecks and also balanced out the traffic on the OSS links.

Comet Network Architecture (original)

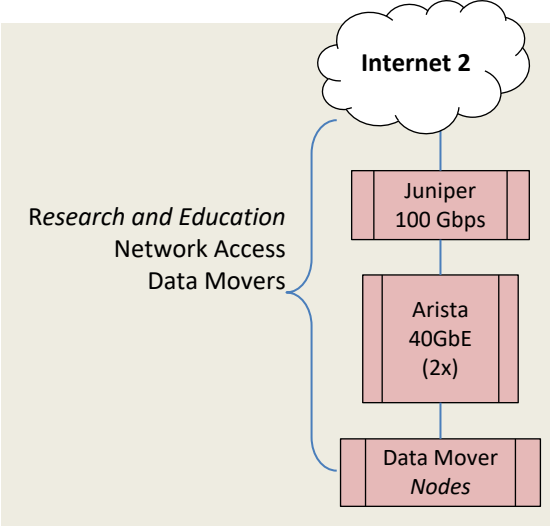
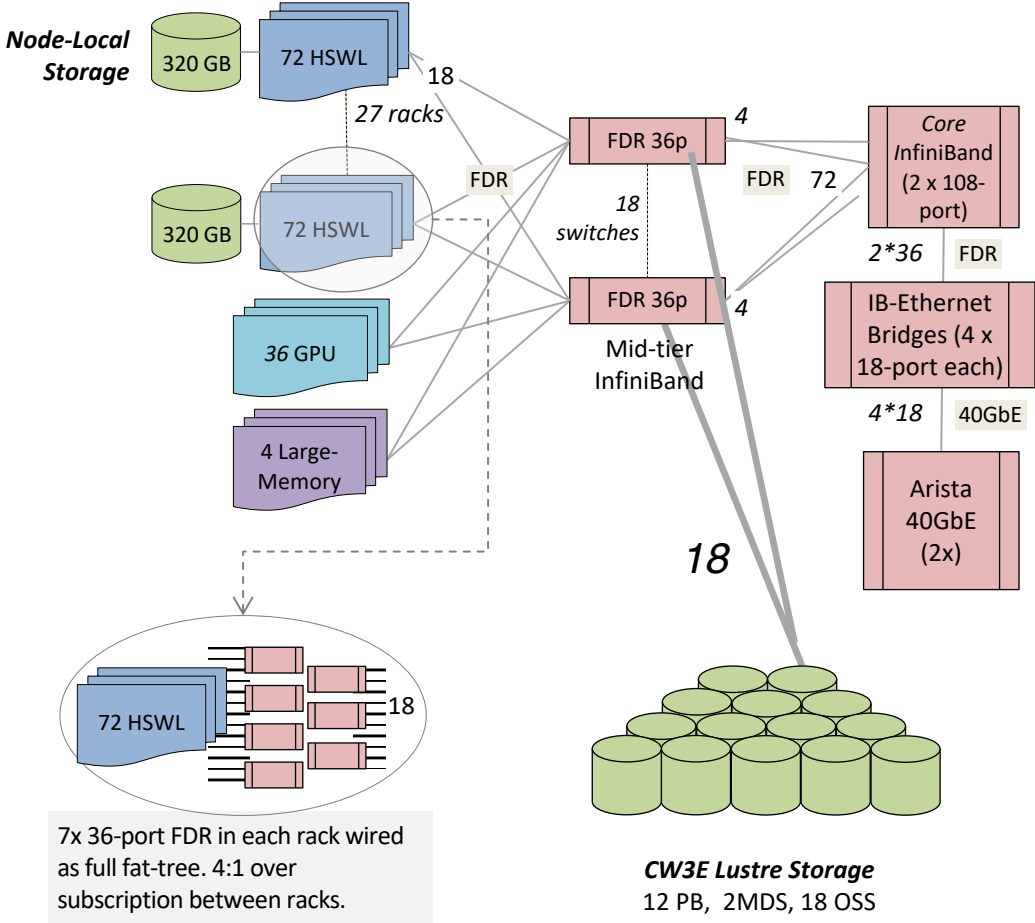
InfiniBand compute, Ethernet Storage



- Additional Support Components**
(not shown for clarity)
- Ethernet Mgt Network (10 GbE)
 - NFS Servers for Home Directories
 - Virtual Image Repository
 - Gateway/Portal Hosting Nodes
 - Login Nodes
 - Rocks Management Nodes

Comet Network Architecture (New)

InfiniBand compute, Ethernet Storage



- Additional Support Components**
(not shown for clarity)
- Ethernet Mgt Network (10 GbE)
 - NFS Servers for Home Directories
 - Virtual Image Repository
 - Gateway/Portal Hosting Nodes
 - Login Nodes
 - Rocks Management Nodes

Notifications after recabbling

Link information

IB-R16-L1-02 (0xf4521403008fcbf0):2 - IB-R16-L2-05 (0xf452140300918b40):8
IB-R16-L1-02 (0xf4521403008fcbf0):36 - comet-16-37 HCA-1 (0xf4521403005bff01)
IB-R20-L1-02 (0xf452140300969400):2 - IB-R20-L2-05 (0xf452140300969100):8
IB-R28-L1-02 (0xf4521403006db860):7 - IB-R28-L2-06 (0xf4521403008dce90):7
IB-R28-L1-04 (0xf4521403008dcf10):15 - IB-R28-L2-07 (0xf4521403008dca90):21
IB-R11-L1-02 (0xf45214030075bfe0):26 - comet-11-47 HCA-1 (0xf4521403005c0c01)
IB-R28-L1-02 (0xf4521403006db860):4 - IB-R28-L2-05 (0xf4521403008d7980):10
B-R16-L1-02 (0xf4521403008fcbf0):13 - IB-R16-L2-07 (0xf4521403009010f0):7
IB-R20-L1-02 (0xf452140300969400):13 - IB-R20-L2-07 (0xf452140300969800):7
IB-R28-L1-04 (0xf4521403008dcf10):26 - comet-28-11 HCA-1 (0xf4521403005936f1)
IB-R08-L2-06 (0xf452140300968400):34 - IB-R24-L3-07 (0xf452140300f67570):10
IB-R19-L1-02 (0xf452140300953000):9 - IB-R19-L2-06 (0xf452140300952e00):9
...
...
IB-R19-L1-02 (0xf452140300953000):1 - IB-R19-L2-05 (0xf4521403009528a0):7
IB-R07-L1-01 (0xf452140300f631a0):6 - IB-R07-L2-05 (0xf452140300f65ef0):6
IB-R23-L1-03 (0xf4521403006f1f00):31 - comet-23-24 HCA-1 (0xf452140300596971)
IB-R19-L1-02 (0xf452140300953000):22 - comet-19-51 HCA-1 (0xf452140300592531)
IB-R07-L1-01 (0xf452140300f631a0):12 - IB-R07-L2-06 (0xf452140300f66570):6
IB-R12-L1-01 (0xf45214030075c0e0):18 - IB-R12-L2-07 (0xf45214030075c960):6
IB-R24-L3-05 (0xf452140300f65df0):36 - comet-dr1-L04 (0xf4521403006ae2d0):10

-OSU INAM

Connections within rack level switches

Connections from nodes to in rack switches

Connections from mid-tier to director switch

The OSS links and director switch leaf/spine links no longer show up in the saturation lists

Outline

- Overview
- INAM on Comet
 - System architecture
 - Network and Job Level Views
 - Notifications
 - Use case illustrating value of INAM
- **Future work**
 - **Planned installation on Expanse**
 - **Expanse architecture**
- Summary

Future Work: INAM on Expanse

- Expanse features a similar architecture to Comet with 3:1 oversubscription of links between racks.
- Single switch at rack level (unlike Comet which had 7).
- Will be using the ClickHouse database server and the C++ client. Expected to perform better than the setup on Comet.
- Expanse system is managed using Bright Cluster Manager (BCM). Look into install of INAM using BCM or develop a compatible process.

EXPANSE

COMPUTING WITHOUT BOUNDARIES
5 PETAFLOP/S HPC and DATA RESOURCE

HPC RESOURCE

13 Scalable Compute Units
728 Standard Compute Nodes
52 GPU Nodes: 208 GPUs
4 Large Memory Nodes

LONG-TAIL SCIENCE


Multi-Messenger Astronomy
Genomics
Earth Science
Social Science

DATA CENTRIC ARCHITECTURE

12PB Perf. Storage: 140GB/s, 200k IOPS
Fast I/O Node-Local NVMe Storage
7PB Ceph Object Storage
High-Performance R&E Networking

INNOVATIVE OPERATIONS

Composable Systems
High-Throughput Computing
Science Gateways
Interactive Computing
Containerized Computing
Cloud Bursting



REMOTE CI INTEGRATION



Heterogeneous Resources



Open Science Grid

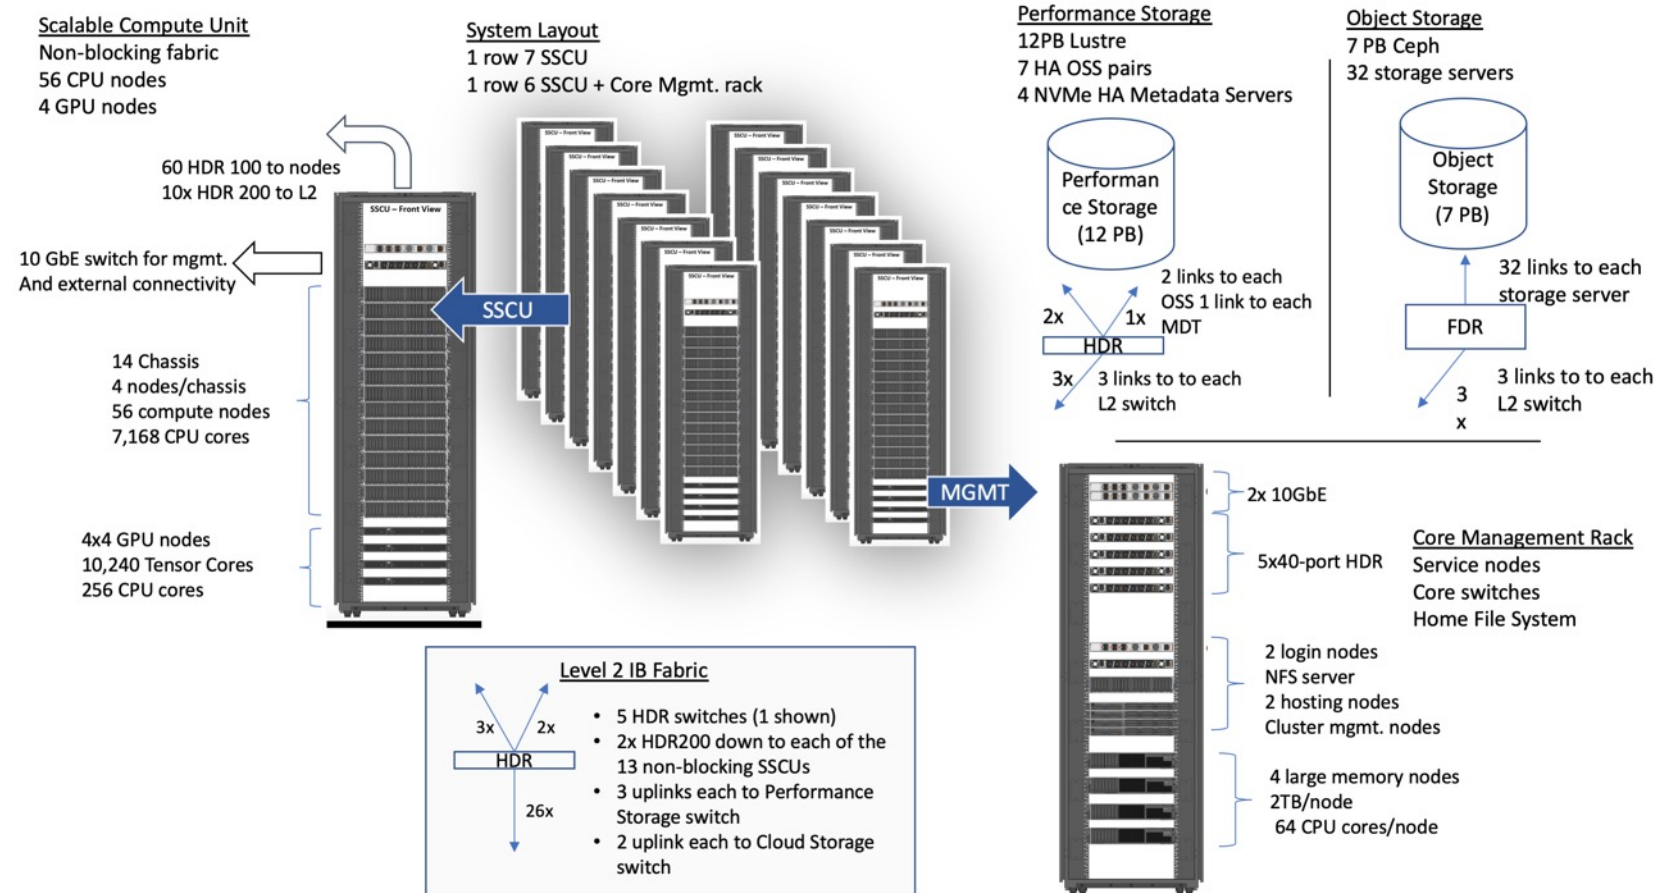
NSF Award # 1928224

PIs: Mike Norman (PI), Ilkay Altintas, Amit Majumdar, Mahidhar Tatineni, Shawn Strande

Expanse is a heterogeneous architecture designed for high performance, reliability, flexibility, and productivity

System Summary

- 14 SDSC Scalable Compute Units (SSCU)
- 784 x 2s Standard Compute Nodes
- 100,352 Compute Cores
- 200 TB DDR4 Memory
- 56x 4-way GPU Nodes w/NVLINK
- 224 V100s
- 4x 2TB Large Memory Nodes
- HDR 100 non-blocking Fabric
- 12 PB Lustre High Performance Storage
- 7 PB Ceph Object Storage
- 1.2 PB on-node NVMe
- Dell EMC PowerEdge
- Direct Liquid Cooled



The SSCU is Designed for the Long Tail Job Mix, Maximum Performance, Efficient Systems Support, and Efficient Power and Cooling

Standard Compute Nodes

- 2x AMD EPYC 7742 @2.25 GHz
- 128 Zen2 CPU cores
- PCIe Gen4
- 256 GB DDR4
- 1 TB NVME

GPU Nodes


- 4x NVIDIA V100/follow-on
- 10,240 Tensor Cores
- 32 GB GDDR
- 1.6 TB NVMe
- Intel CPUs

SSCU Components



- 56x CPU nodes
- 7,168 Compute Cores
- 4x GPU nodes
- 1x HDR Switch
- 1x 10GbE Switch
- HDR 100 non-blocking fabric
- Wide rack for serviceability
- Direct Liquid Cooling to CPU nodes

Non-blocking Interconnect

1 HDR Switch/SSCU



5 Level 2 switches



Summary

- INAM installed and in production use on Comet.
- Used for network health/performance monitoring and to identify sources of congestion.
- Successfully identified hotspots and bottlenecks that were causing Lustre filesystem performance issues under heavy NRT simulation loads.
- Recabled and confirmed performance improved significantly and I/O bottlenecks mitigated.
- Expanse install in planning phase. Will be using the ClickHouse database server and the C++ client.