

# The Next Tentacle

## Ceph on ARM

Federico Lucifredi

Product Management Director, Ceph Storage Platform, IBM | Red Hat

Peter Pouliot

Senior Developer Evangelist, Ampere Computing

---

June 4, 2025

# me! me! me!



## Things I worked on

Ceph Storage

Ubuntu Server

Landscape

SUSE Studio

SLES

SMT

Ximian Red Carpet

Man (I)



# Peter Pouliot

ppouliot@amperecomputing.com

- I live in Stoneham, MA and am a lifelong New Englander originally from Rhode Island.
- I studied Philosophy for my undergraduate degree (RIC), and Telecommunications Security/Digital Forensics graduate degrees (BU)
- At Ampere Computing, I enable Open Source Software ecosystems with Aarch64 servers in the datacenter, and enabling Aarch64 native OSS server-based cloud computing.
- Previously at Microsoft, I helped to maintain OpenStack integration with Microsoft's virtualization platform Hyper-V and was Microsoft's OpenStack subject matter expert, running the largest CI/CD infrastructure within the OpenStack community for 5+ years.



one more thing

---



## Supported CPU Architectures

RHCS is supported on Intel and AMD x86-64 microprocessors. IBM POWER and S390x CPUs are supported only as clients, but not as clusters.

## Supported Host Operating Systems



**Red Hat**



6/25



## Supported CPU Architectures

RHCS is supported on Intel and AMD x86-64 microprocessors. ARM64, IBM POWER and S390x CPUs are supported only as clients, but not as clusters.

## Supported Host Operating Systems



**Red Hat**





# outlook

## DELIVERY



## PLAN

— SUBJECT TO CHANGE —



# Community

---

# punch list

- Community Lab reopening
- More test automation
- Tuning for ARM
- Optimization for ARM extensions
- ARM Reference architecture



2019...



# SBC distributions

- Armbian
- Ubuntu
- Debian
- Fedora
- RHEL



performance

---

# Ceph Cluster Overview

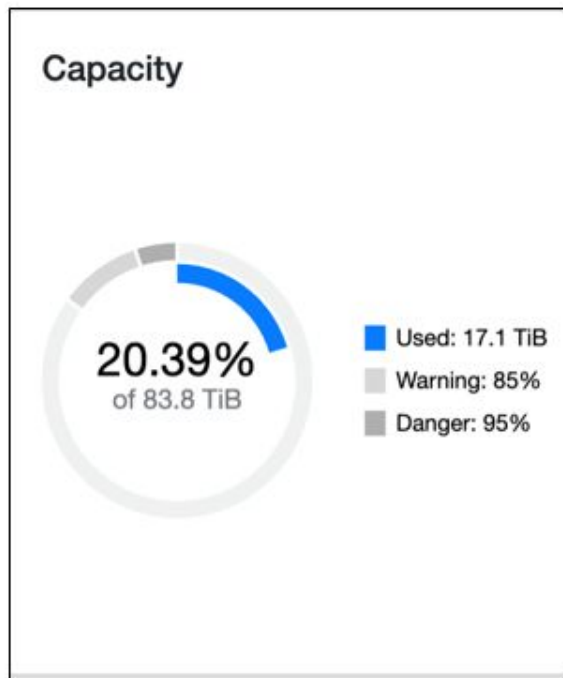


# Ceph Cluster Inventory

Ampere	
CPU	Ampereone
Memory	8 x 64GB DDR5 512 GB per node
Disks	8x 3.5 TB Micron 7450 NVME's per Node
NIC's	1 x 100Gb Private and 1 x 10Gb Public per Node
OS	Ubuntu 24.04
Ceph Version	Squid 19.2.0 with OSD's on Docker Containers
Client Nodes	3 Mt Kims

3 Node Cluster with Ampereone setup for Performance benchmarking

# Ampere 3 Node Cluster – Ceph details



Inventory		
3 Hosts	3	✓
3 Monitors	3	✓
2 Managers	1 ✓ 1 ⓘ	
24 OSDs	24	✓
4 Pools	4	✓
6145 PGs	6145	✓

# Ceph Cluster Block Storage Performance

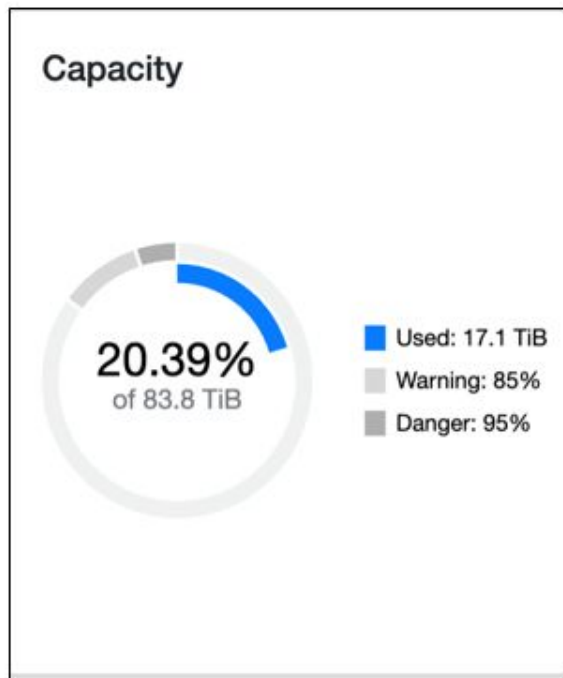


# Ceph Cluster Inventory

Ampere	
CPU	Ampereone
Memory	8 x 64GB DDR5 512 GB per node
Disks	8x 3.5 TB Micron 7450 NVME's per Node
NIC's	1 x 100Gb Private and 1 x 10Gb Public per Node
OS	Ubuntu 24.04
Ceph Version	Squid 19.2.0 with OSD's on Docker Containers
Client Nodes	3 Mt Kims

3 Node Cluster with Ampereone setup for Performance benchmarking

# Ampere 3 Node Cluster – Ceph details



Inventory		
3 Hosts	3	✓
3 Monitors	3	✓
2 Managers	1 ✓ 1 ⓘ	
24 OSDs	24	✓
4 Pools	4	✓
6145 PGs	6145	✓

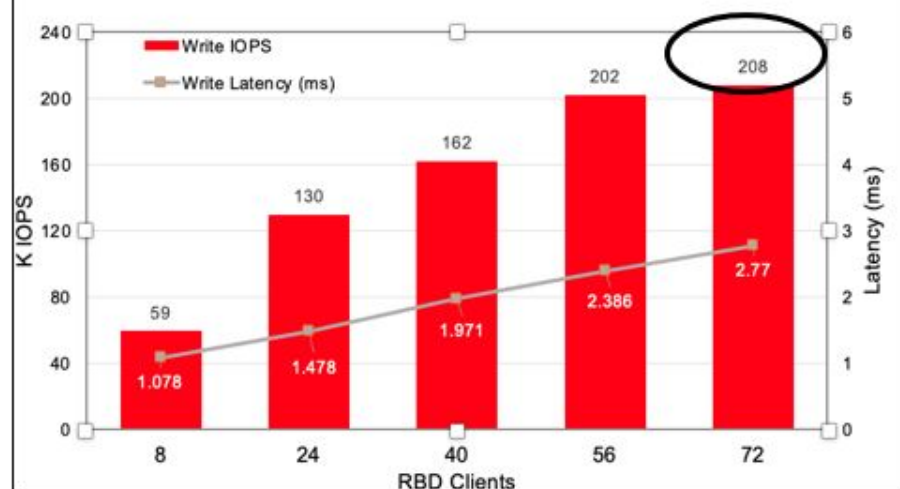
# Ceph Cluster Block Storage Performance



# Block Storage Write Performance Comparison

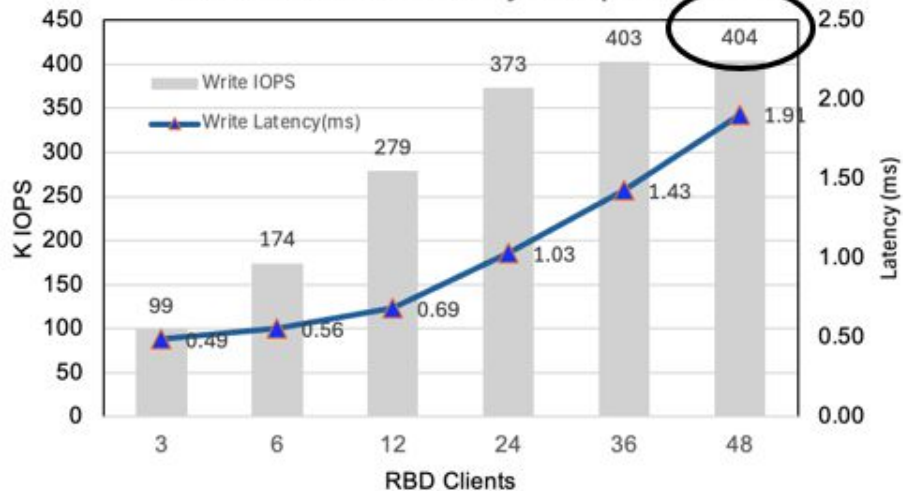
Altra

Write IOPS and Latency - Altra



Ampereone

Write IOPS and Latency - Ampereone



Ampereone Stats from Ceph Dashboard

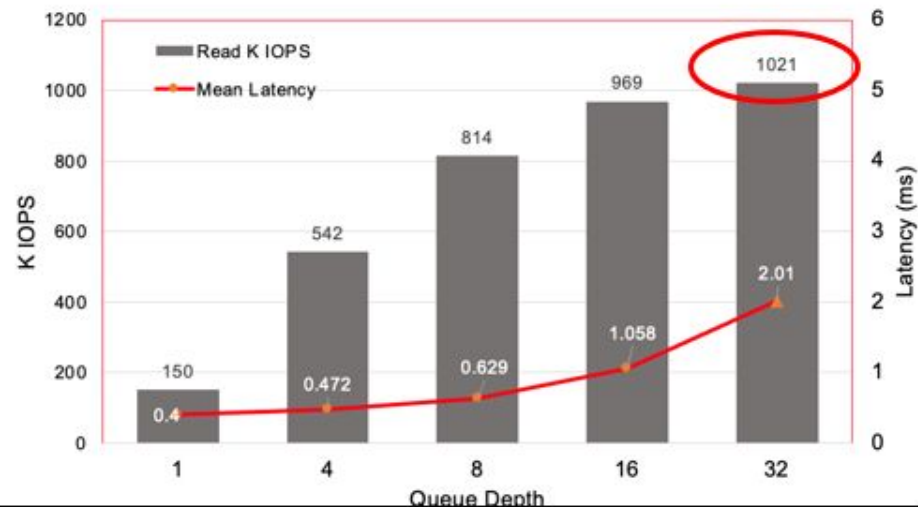


4 KB Random Writes reached a maximum of 400 K IOPS in Ampereone compared to 200 K IOPS achieved with Altra ( Repl = 2 )

# Block Storage Read Performance Comparison

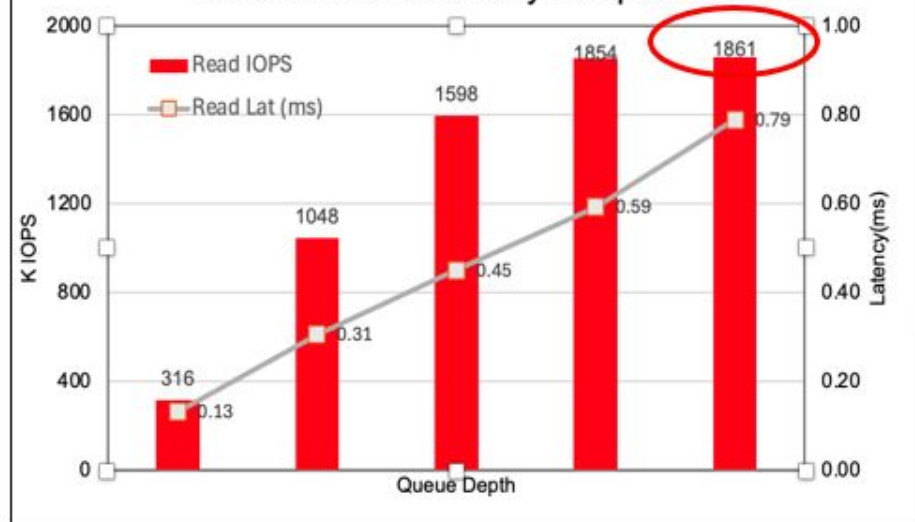
Altra

Read IOPS and Latency - Altra



Ampereone

Read IOPS and Latency - Ampereone

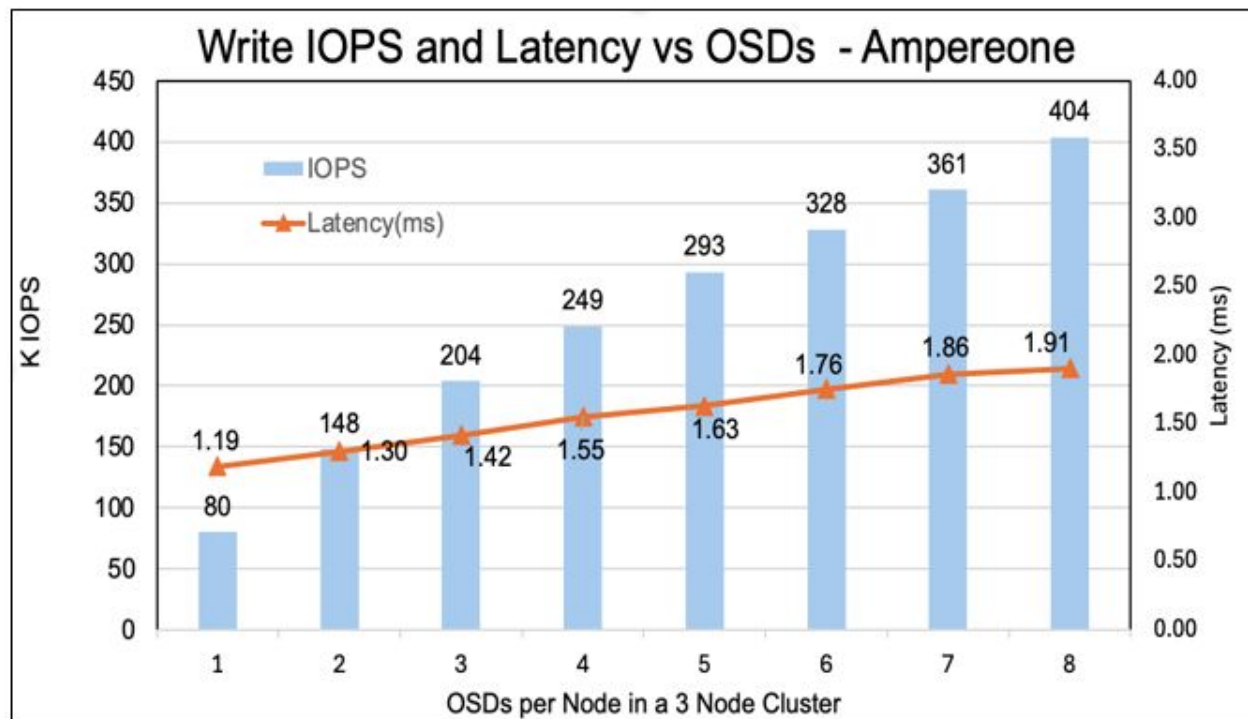


Ampereone Stats from Ceph Dashboard



Disks reaching saturation point in Ampereone

# Ceph Incremental Performance with OSDs



# Compare and Contrast our results

<https://ceph.io/en/news/blog/2024/ceph-a-journey-to-1tibps/>

( Jan 2024 )

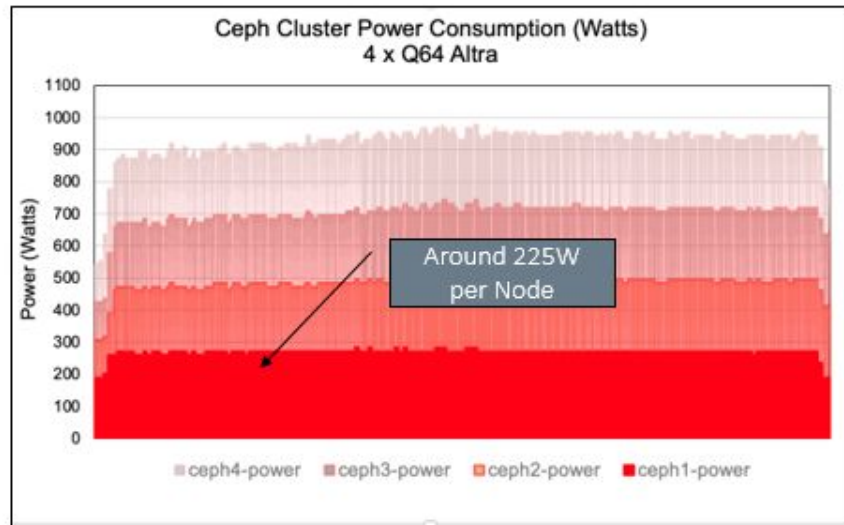
Number of Nodes	3	3
CPU	3 x AMD 9454	3 x Ampereone
Cores/Threads	48 Core/96 Threads	192 Cores
Memory	3 x 192GiB DDR5	3 x 512GiB DDR5
Disks	3 x 10 15TB NVME Drives	3 x 8 3.5TB NVME Drives
Write IOPS	248 K IOPS ( Repl 3 )	404 K IOPS ( Repl 2 ) 277 K IOPS ( Repl 3 )
Read IOPS	1.9 M IOPS	1.86 M IOPS

# Ceph Cluster Power Consumption

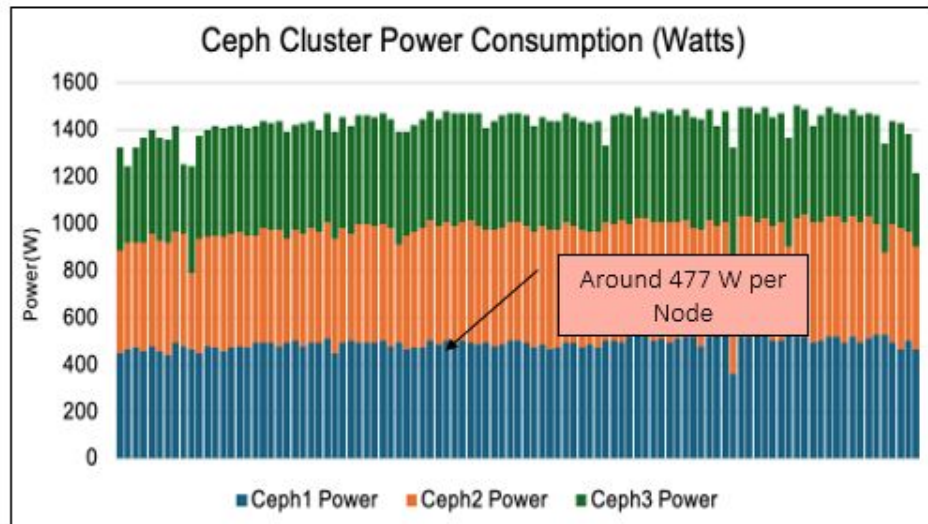


# Ceph Cluster Power Consumption (4k Writes)

Altra



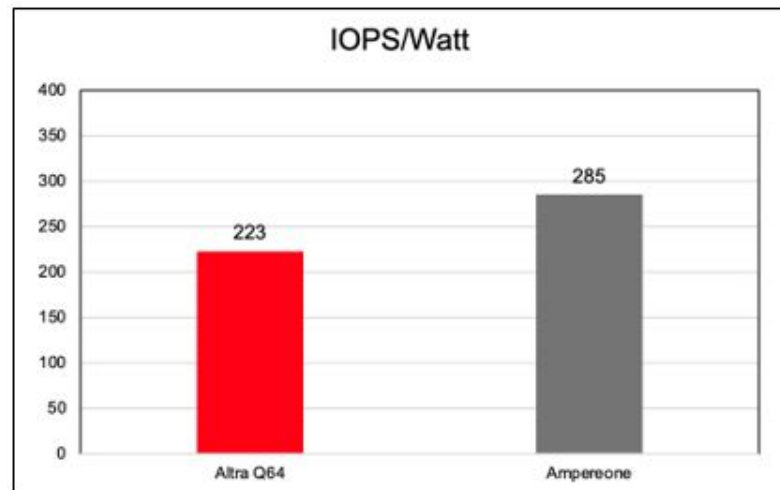
Ampereone



- 4 x Altra nodes delivered around 200 K IOPS while consuming around 900 Watts
- 3 x Ampereone Nodes delivered around 400 K IOPS while consuming around 1430 Watts

# Ceph Back End Performance

	Altra (q64)	Ampereone (192)
CPU Cores/Threads	64 Cores	192 Cores
CPU	56%	55%
Mem	16%	18%
Physical Disk IOPS	40k-45 K	100 K
Total Power	224 W	477 W
CPU Power	84 W	228 W



THANK YOU.

---