



Building AI Data Infrastructure: Lessons from the Field

HPC & AI Summit Lightning Talk

Jonathan Symonds, MinIO
September 2024



As Promised: From the Field

Threat Detection

- 100PB growing to 500PB over the next year
- Colocated in Equinix, repatriated from AWS
- Not just mission-critical, it is the mission of the business.
- Shift to MinIO provided a 2-3% GM improvement.
- Uses $\frac{1}{3}$ of an FTE to manage 100 PB+ of data on MinIO
- 250TB ingest per day

Autonomous Driving

- One of the biggest private cloud AI deployments in the US.
- 33% of the automakers car data. Video, telemetry, log files.
- 1EB of data, doubling this year. 330 Servers, Dual 100 Gbe NICs. 2.2 TiBs per second throughput.
- Bespoke training application across 1K GPUs.

**What Got You Here Will NOT Get You
There.....**




The Scale is Different This Time Around



The PB is the New TB

Every enterprise is working with hundreds, if not thousands of PB.

This stresses older technologies.



Data is Distributed

We are creating 328 EB a day.

But...it gets created everywhere and in every format.

The average enterprise uses 3.4 public clouds and 3.9 private clouds.



Cloud Native is the Only Way

The cloud is an operating model and block and file aren't to be found.

Restful APIs (S3),
Containerization,
Orchestration, Microservices.

Software-defined everything.
Smart software, dumb hardware.



Performance at Scale is Everything



Performance = Workloads

Performance is defined by IOPS AND throughput. Modern object storage offers both but is optimized for throughput.

Small objects to large objects.



Scale Can Destroy Performance

SAN/NAS performance used to be the benchmark - but past 1PB things start to degrade.

Performance needs to stay intact from 100 GB to 1,000 PB - because that is where the data is growing to.



Economics and Performance are Related

No one has an unlimited budget.

If the cost to scale performance is prohibitive - you will make bad decisions.



Other Important Things

1. Go with NVMe SSD
2. Observe, Observe, Observe
3. Site Replication/Bucket Replication
4. Understand Your Data Behavior
5. Automate Everything
6. Security



Summary

00010010
101001101
00010010
111001001
00010010

Build for a Different Class of Scale

We routinely see examples of exabytes of production data.

This is the new normal.

Legacy technologies tap out at 100PiB – we have seen it first hand.



Performance at Scale is a Core KPI

If your technology cannot deliver linear performance and if you don't have proof that it can go past 100 PiB then you should be architecting it out.

If you get it wrong you have to start OVER.



Your Technology Stack Will Have to Change

Appliances don't work in the cloud native world. SAN/NAS are legacy.

Applications are written for S3 now.

Think in terms of scalable units.



Economics Are Part of the Scale Equation

Architect for simplicity.

There is a tendency to focus on cost - but it is simplicity that scales: technically, operationally and economically.

A list of social media and website links is shown in the bottom left. Each link is preceded by a small icon: a Twitter bird for the handle, a GitHub Octocat for the repository URL, a Slack hashtag for the workspace URL, and a globe for the main website URL.

 @minio
 <https://github.com/minio/minio>
 <https://slack.min.io>
 <https://min.io>