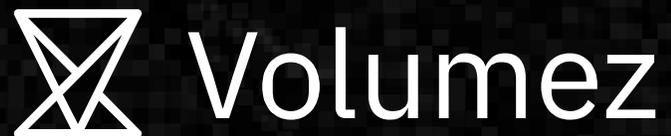




Volumez – Overview

Presented to: **Company Name**

August 13th, 2023



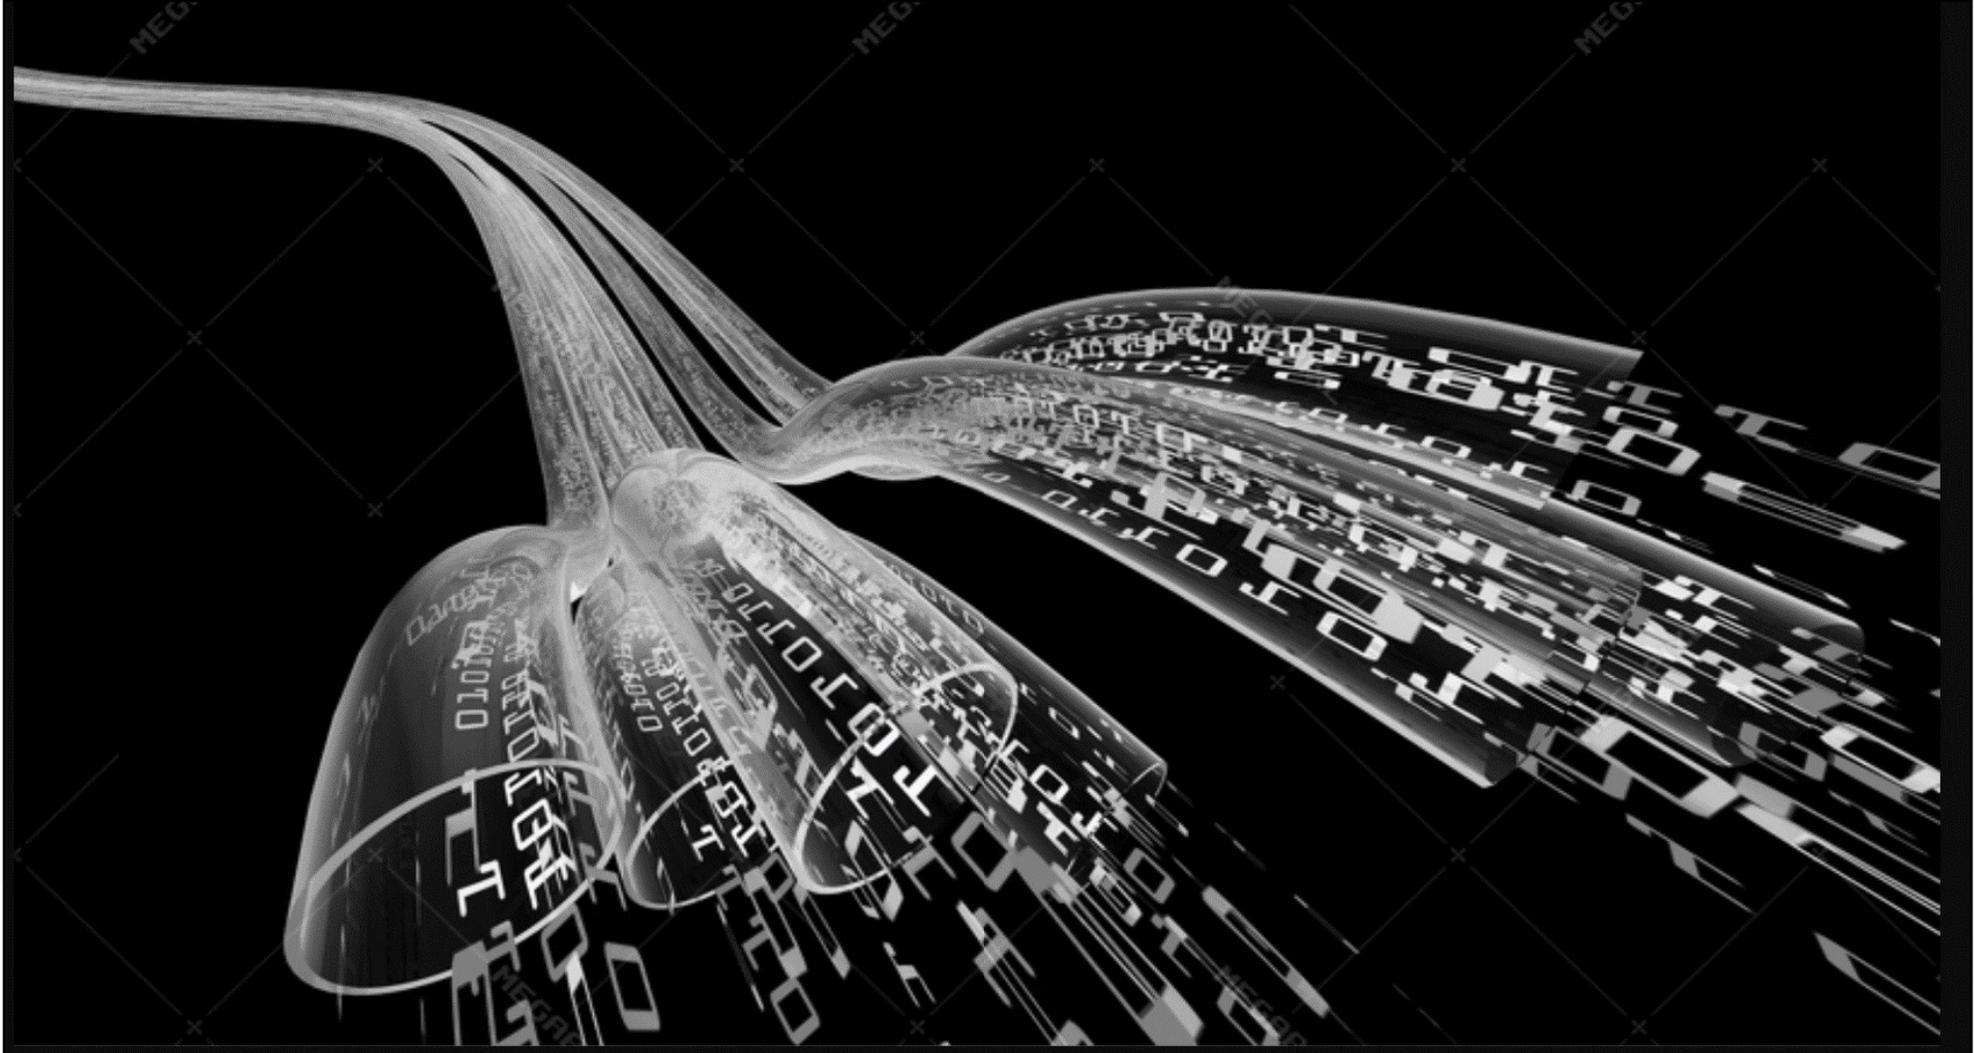
Disclaimer

This Presentation is confidential and proprietary. Viewers of this presentation are hereby advised that the Company's actual results and achievements in the future may differ materially from those presented in this presentation. In addition, forward-looking forecasts and assessments are based on data and information available to the Company at the time the presentation was prepared, as well as third party information, and the Company does not undertake to update and/or modify any such forecast and/or assessment to reflect events and/or circumstances occurring after the date this presentation was created.

Agenda

- Introduction
- Market challenges and vision
- The Volumez service – composable data infrastructure
- Q&A

It's about the data



Introduction / Team



CA

Amir Faintuch
CEO



TLV

Johnathan Amit
Founder & CTO



TLV

Nir Rigai
VP Engineering



CA

John Blumenthal
Chief Product &
Business Officer



TLV

Efri Nattel-Shay
VP Products



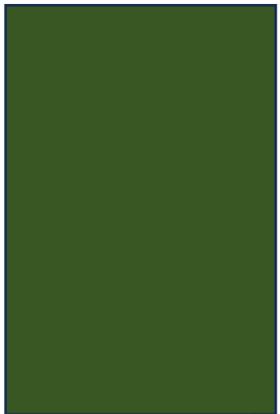
TLV

Carmit Levy
Head of
Operations & HR



CA

Razi Sharir
Head of Strategic
Development



Chief Revenue
Officer

*Starting Sep. 13th.



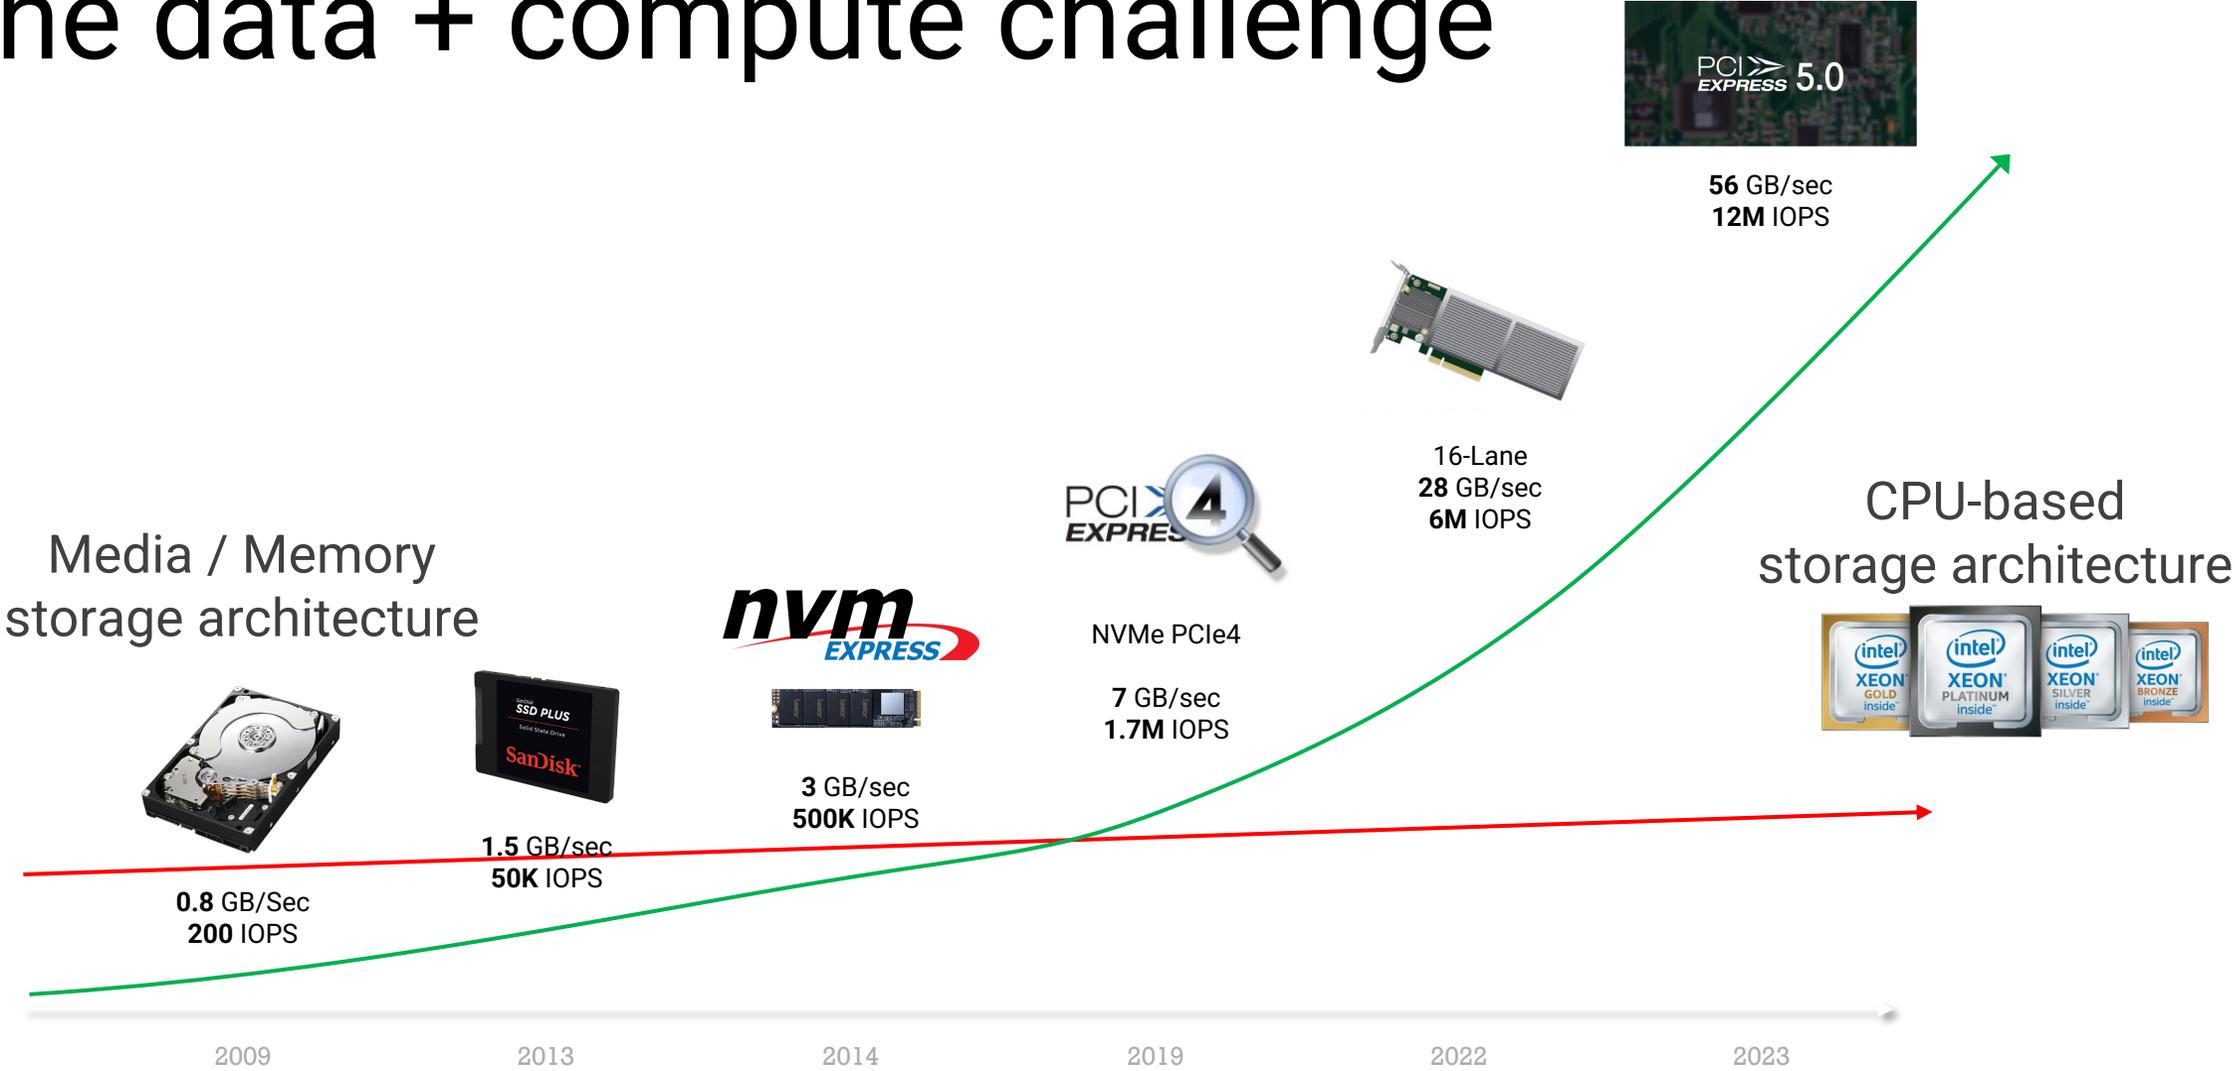
Informatica



Volumez makes composable data infrastructure possible

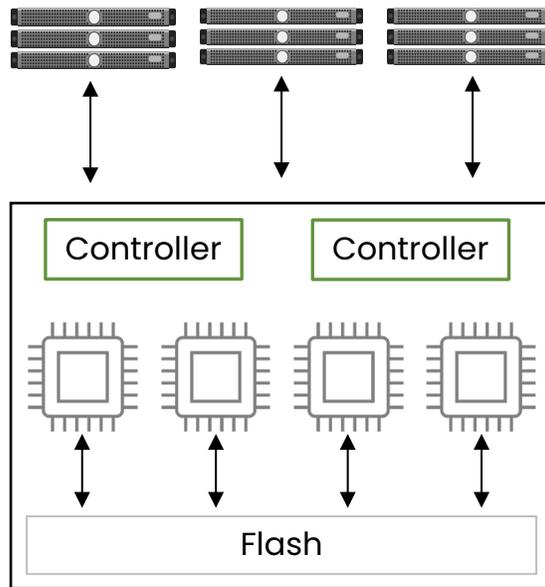
... And the “Compossibilities” are endless

The data + compute challenge

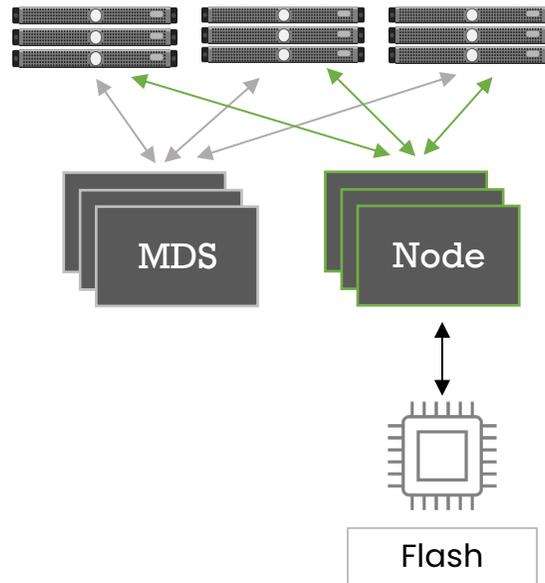


Storage controllers became bottlenecks

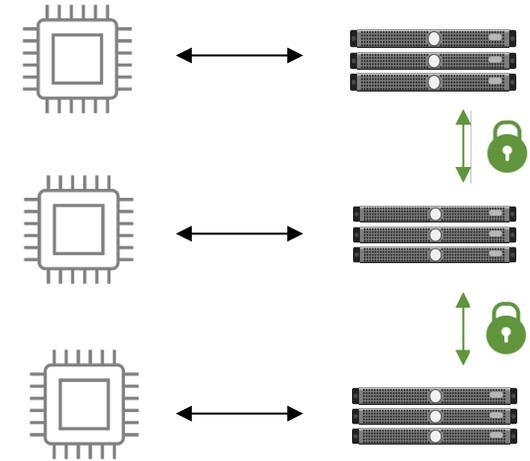
Storage Appliances



Software Defined Storage Meta Data Server



Software Defined Storage Cluster Locks



Today

Compute



Data



Volumez – Essential to data as Kubernetes for compute

Compute



Data



Volumez

Unlocking infrastructure performance

1.5M

IOPS

300 μ s

Latency

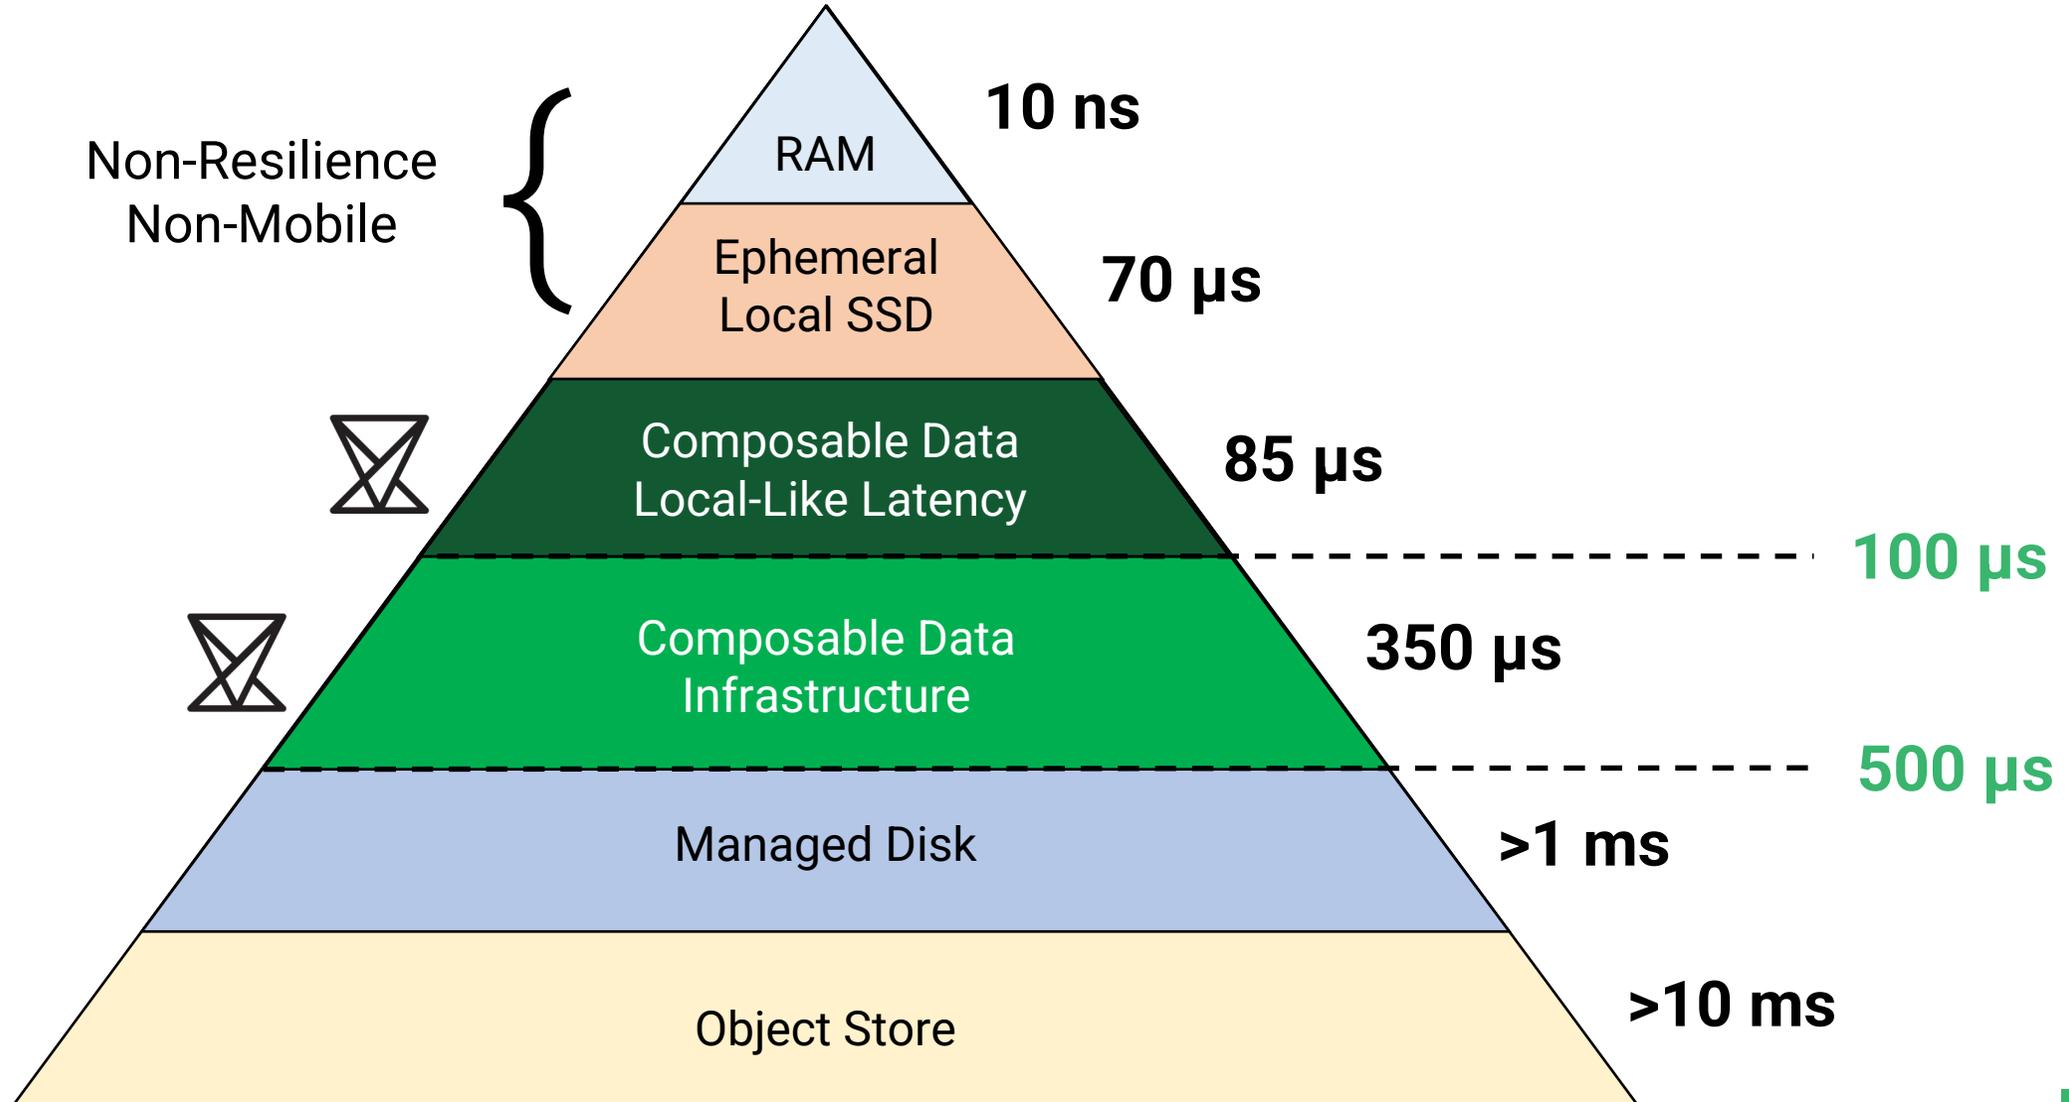
12 GB/s

Bandwidth

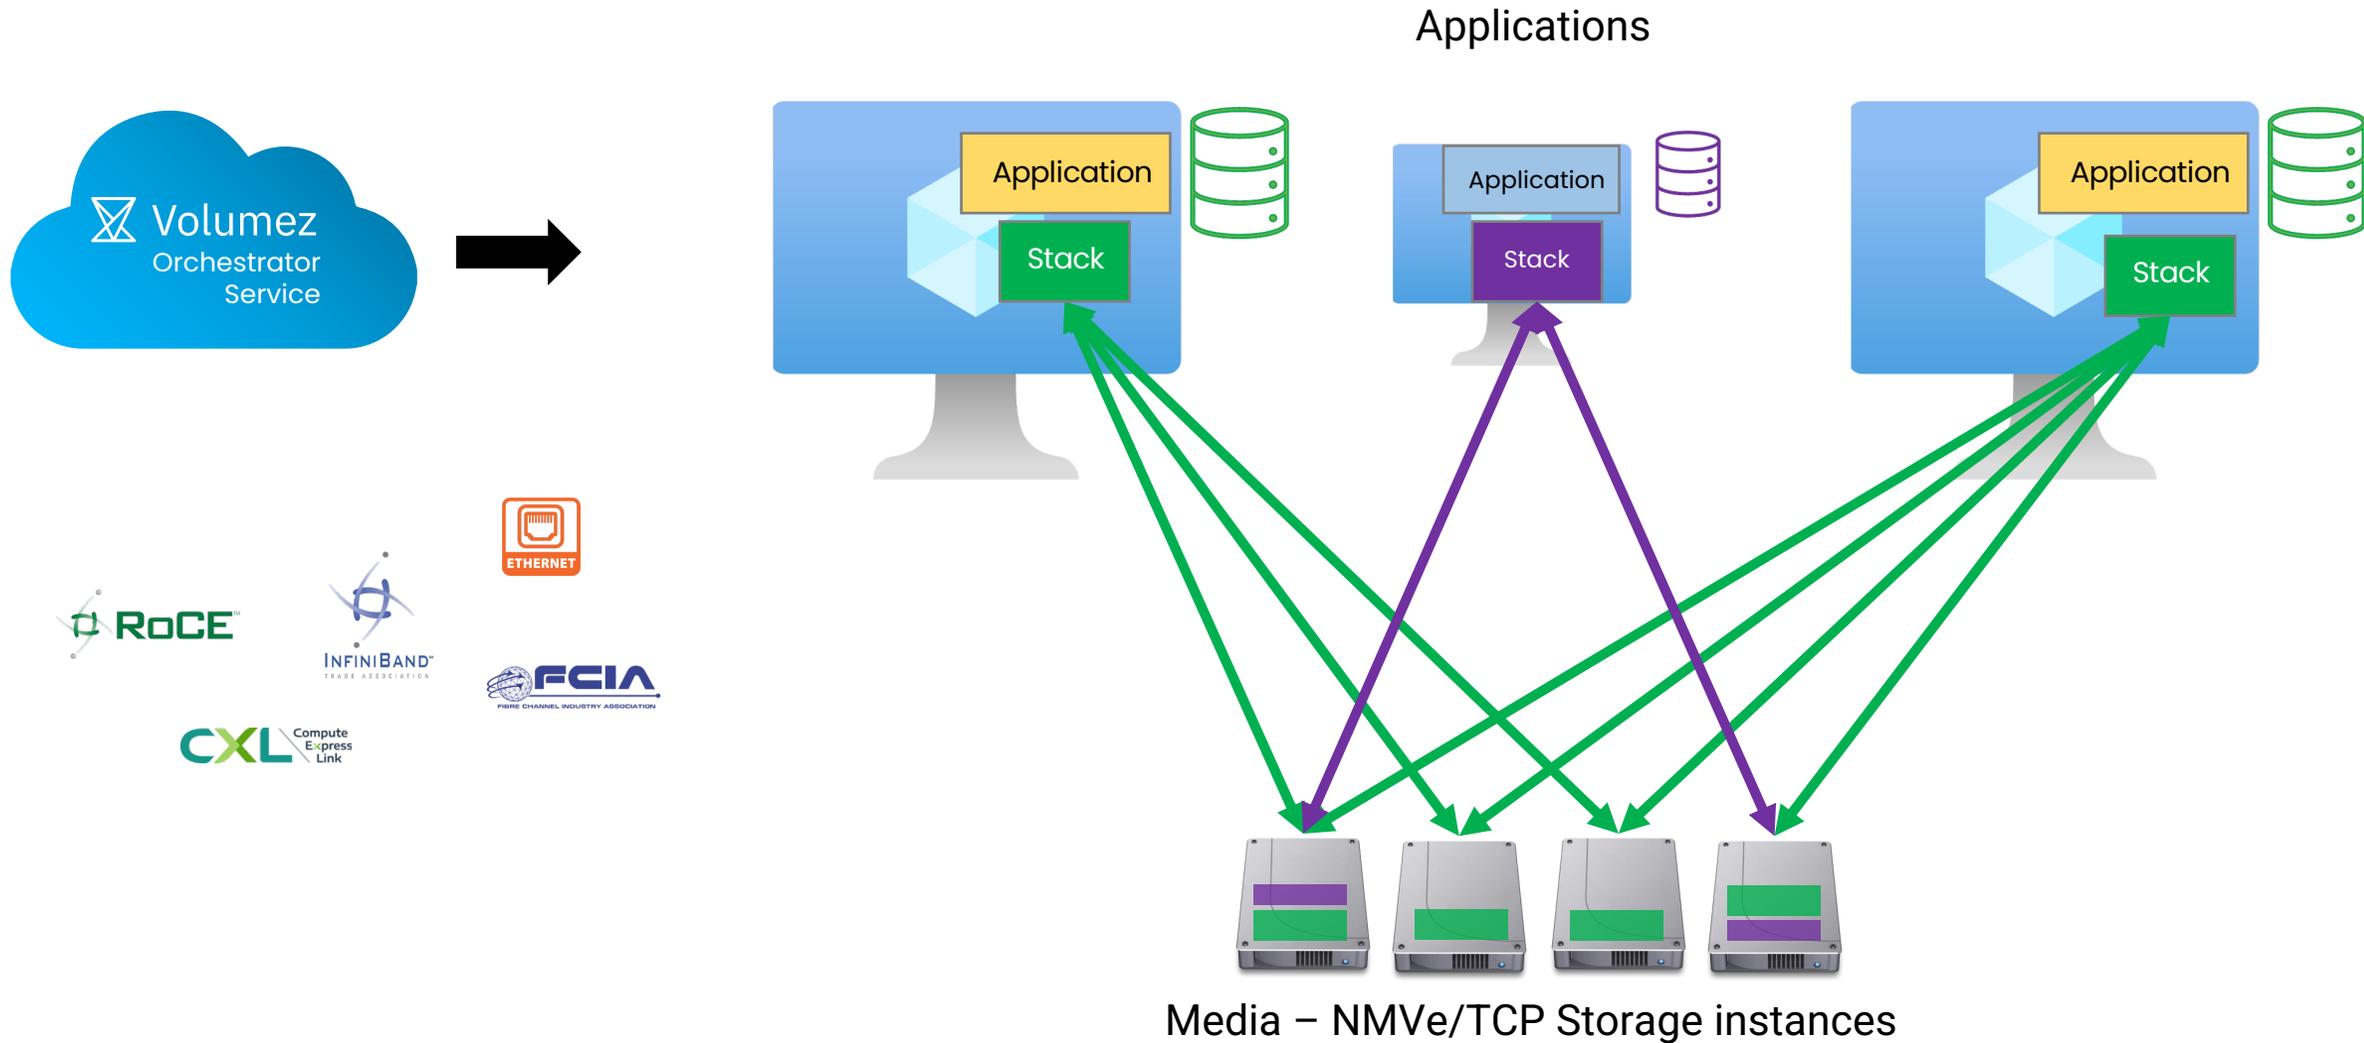
Per volume, guaranteed

- No cache, with mirroring, on AWS, with full data services
- 1.5M IOPS – with mixed 4K random 70%/30% R/W. 2M IOPS for read-only 4K random workloads

Cloud Storage Pyramid



Composable data infrastructure



SaaS orchestration

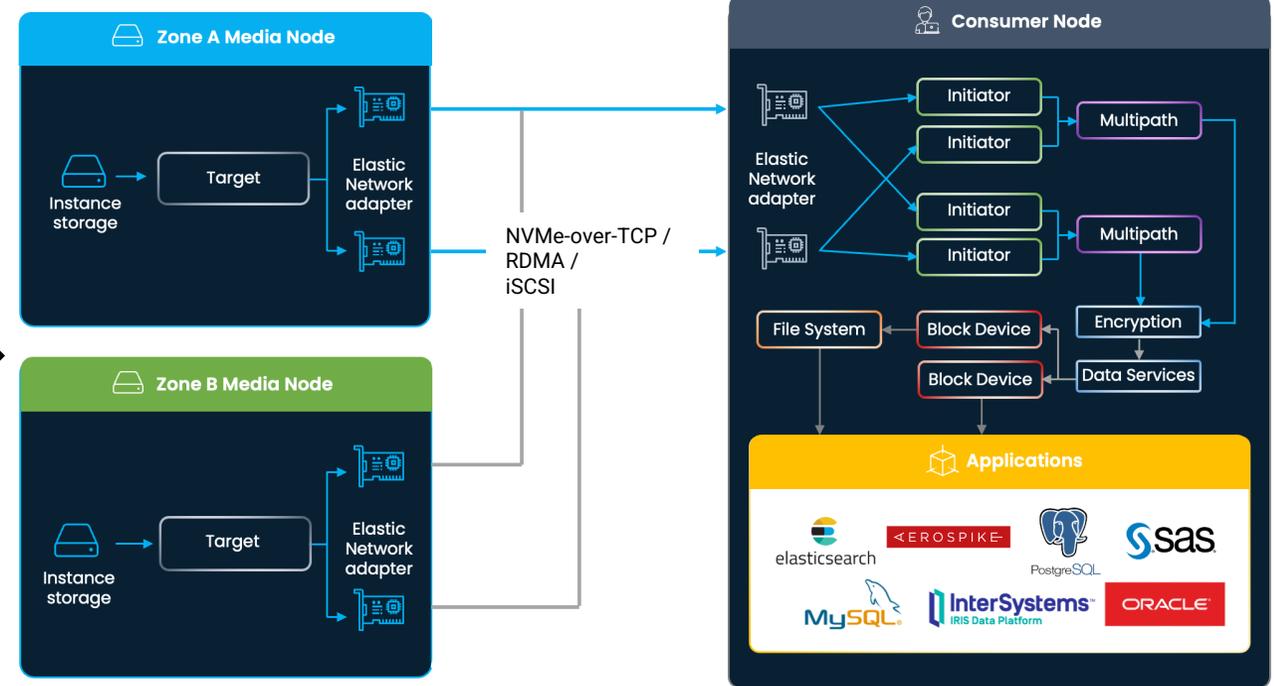
Declarative policy

```
{  
  "name" "production-database",  
  "performance": {  
    "iops" 1000000,  
    "latency_usec" 180  
  },  
  "encryption" true,  
  "resilience": {  
    "media" 3,  
    "nodes" 2,  
    "zones" 1  
  }  
}
```

Declare



Compose

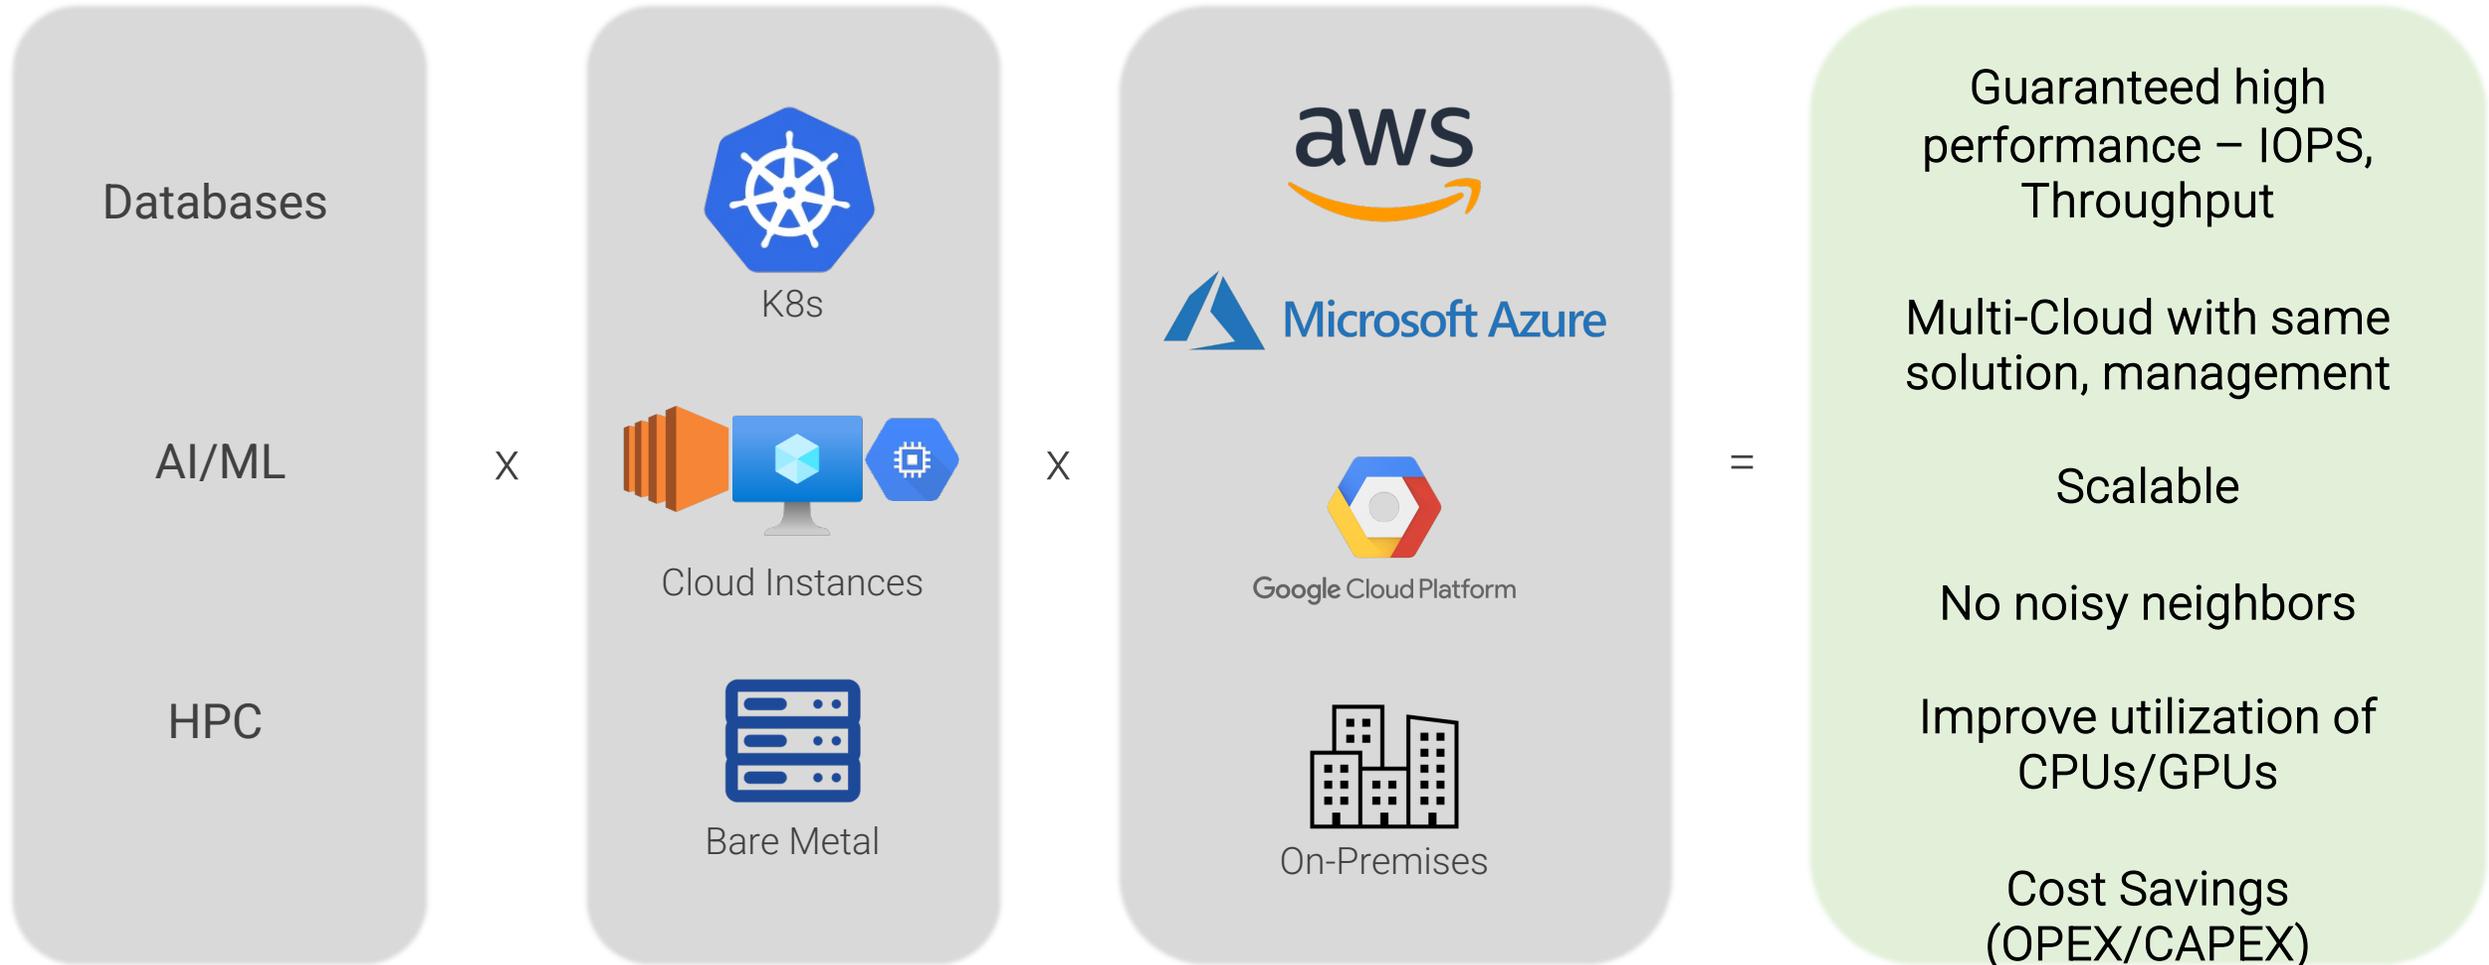


100% Linux Data Path
No Storage Controller



On-Prem

Composable Data Infrastructure – Unleash the power of the cloud



Customer Case Studies

Customer case study – Cloud DBaaS

Customer problem

1. Need to offer DBaaS in multiple clouds
2. Need low latency, high performance, low TCO

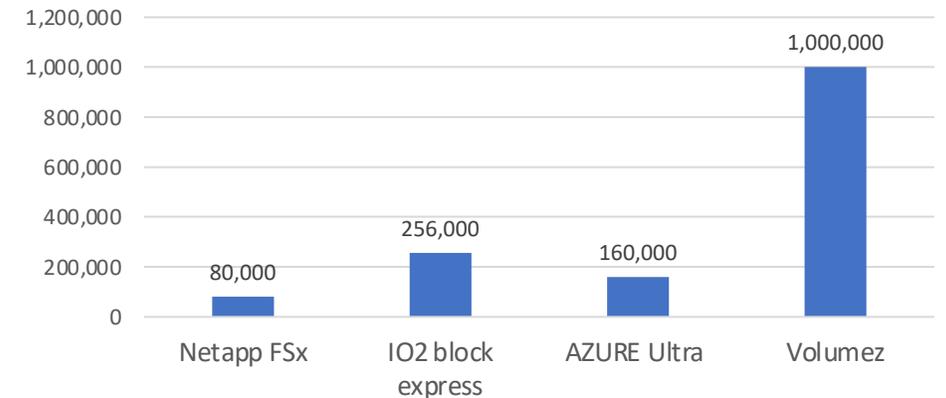
Technical benefits

1. Guaranteed **high performance**: 1M IOPS, 300 microseconds latency (vs. 256K IOPS, 1 millisecond latency in io2)
2. **Single** solution for public and private clouds, using Kubernetes
3. **Resiliency** – survive failures of up to 3 media, 3 nodes, 2 zones
4. Enterprise **data services** – e.g., snapshots, copy to object
5. **Zero** disruption from **noisy neighbors**
6. Linear **scalability**

Business benefits

1. Write new workloads once, run anywhere
2. **20%-50%** total cost savings – allowed launching the service
3. Launch a **better service**

SQLaaS - IOPS
(70%/30%, no cache hit)



Customer case study – leading cloud provider

Customer problem

1. Volumes with significantly lower latency and higher performance than available today
2. Fast provisioning of many volumes in Kubernetes

Technical benefits

1. High guaranteed performance, per volume
2. Create 20 Kubernetes volumes in <40 seconds
3. Volumes with multi-AZ resiliency
4. Fast snapshots

Business benefits

1. Ability to scale Kubernetes applications fast in the cloud
2. Ability to move traditional workloads from on-prem to cloud

3. Performance results per-volume:

1.6M IOPS
500 μ s latency

4K Read-Only

1M IOPS with latencies of
450 μ s (read), 400 μ s (write)

4K Mixed 70%/30%

10 GB/s Read,
4 GB/s Write

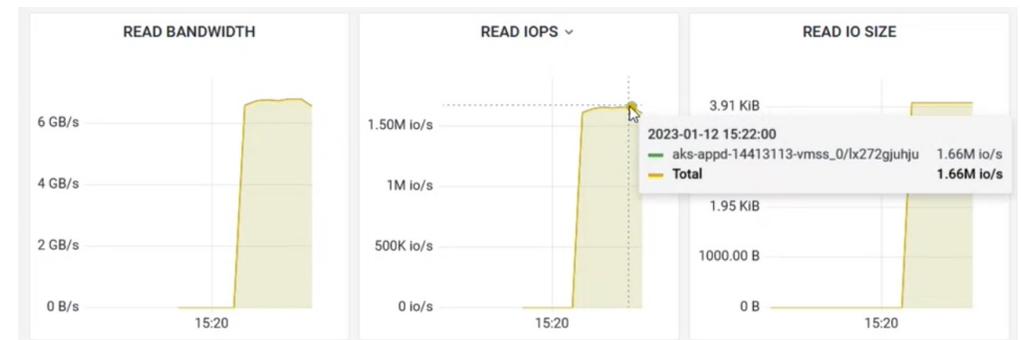
Sequential Reads/Writes

<40 seconds

Create and attach 20 PVCs

The fine print:

- a) R/W were done with no caching
- b) Volumes have resiliency with mirroring (RAID1)



Thank You
