

DDN | HPC User Forum

DDN Introduction

Technical Panel

April 9th, 2024



DDN Powers the World's Toughest Data Challenges

The Complete Cloud, AI, Virtualized, Containerized, HPC, SW-Defined Portfolio

EXAScaler *AI & HPC*



Cloud and All NVMe flash & hybrid scale-up & scale-out parallel storage

IntelliFlash *Unified*



All NVMe flash & hybrid scale-up SAN, NAS & S3 storage appliances

VMstore *VM Aware*



Virtual Machine aware all flash, scale-out intelligent storage for VMWare, Hyper-V, CiTRiX & Containers.

Infinia *SW-Defined*



Composable Storage Multi-tenant Ground Up Block(CSI), Object, File

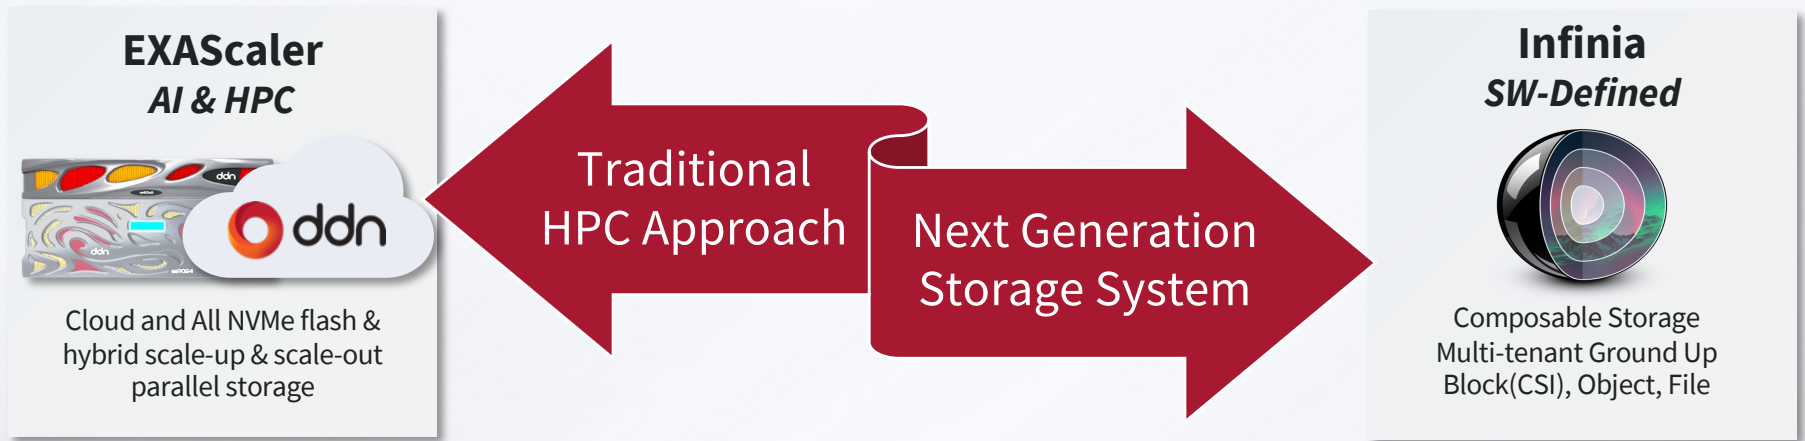
Data Management SW

dataflow

Software solution for backup, archival, sync and migration from any to any storage

DDN Powers the World's Toughest Data Challenges

The Complete Cloud, AI, Virtualized, Containerized, HPC, SW-Defined Portfolio





Safer with DDN

- ✓ End-to-End Encryption
- ✓ End-to-End Data Integrity
- ✓ Multi-Tenancy
- ✓ No SPOF Server Architecture

Accelerated Computing Full Stack



Faster with DDN

- ✓ 30X Faster Single Thread
- ✓ 15X Faster Checkpoints
- ✓ 4X Faster Data Load
- ✓ Higher per-Port Throughput

More Efficient with DDN

- ✓ End-to-End Data Reduction
- ✓ Up to 15X Better/Watt Power
- ✓ Lowest Compute Overhead
- ✓ Bluefield-3 Acceleration
- ✓ GPUDirect Speedup

More Scalable with DDN

- ✓ Standard with NVIDIA SuperPOD
- ✓ No Backend Network Layer
- ✓ Intelligent Client Network & Compute
- ✓ Linear Performance Scaling
- ✓ Scale with TLC, QLC & HDD




**10X
Speed**

**1/4
Power**

**10X
Scale**

DDN: World's Most Performing and Flexible Storage At Any Scale



Best IOPS & Throughput per rack

Up to 70M IOPs in a single rack
1.8 TB/s of read throughput
1.4 TB/s of write throughput
16PB of Flash

Best Price per Flash TB

Up to 14M IOPs in a single rack
360 GB/s of read throughput
240 GB/s or write throughput
Up to 26PB of Flash

Best Price per TB


20PB per Rack
90 GB/s of read throughput
65 GB/s of write throughput

DDN Infinia will Enable the Next Generation of AI

The 4 big data challenges of the new Era of AI: **Efficiency** | **Metadata** | **Sharing** | **Moving**



Data services



**Containerized
MicroServices**



**Scalable Data &
Unlimited Metadata**



**Native Tenancy
Architecture**



**Move Data; Edge,
Cloud, Datacenter**

DDN Infinia is 100% Fit Into Modern AI Development

- **Fully Containerized Architecture** → **Elastic Scalability. Always On**
- **Multi-Tenancy that is 100% Dynamic** → **Easily change Tenant Allocation**
- **Massive Unstructured Data** → **Tag & Search without limits, collapse complexity**



**TAG, SEARCH,
STORE**



**TENANTS &
SUBTENANTS**



DYNAMIC QoS



CLOUD NATIVE



SCALABLE

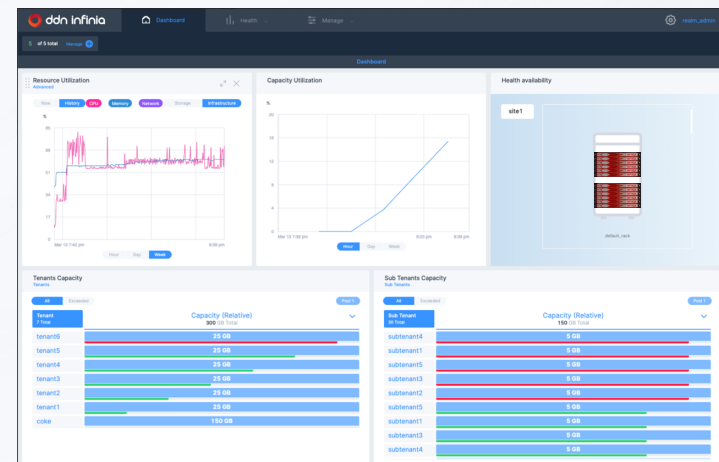
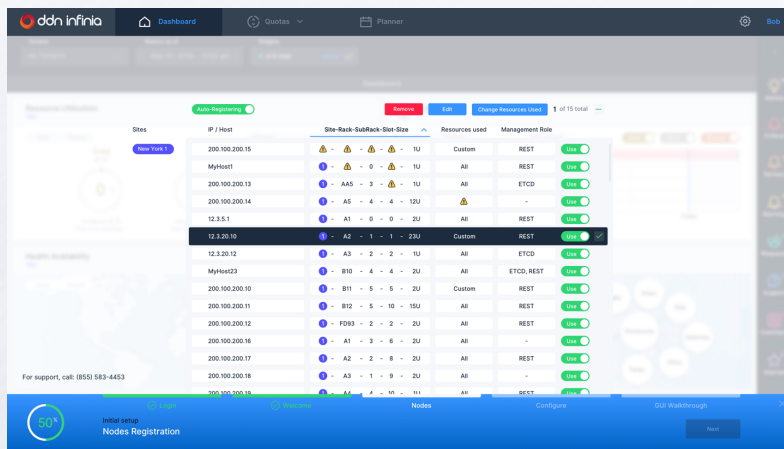


ELASTIC

Easy to Deploy. Easy to Upgrade. Easy to Manage

- ✓ Auto-Installation, **10 minutes to Deploy**
- ✓ **Zero Downtime** Upgrades, Elastic Expansion
- ✓ Deploy new services and tenants in just 4 minutes

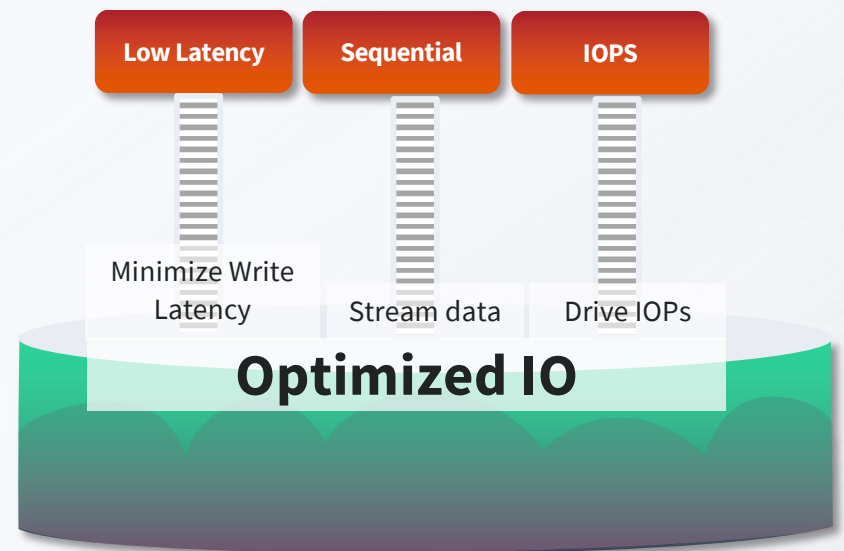
- ✓ **Hands-Free QoS**
- ✓ Multi-Tenanted **Self Service**
- ✓ **No LUNs/No Tuning**



Keeping Every Application Optimized Regardless of Type

Hybrid IO Engines Handle Diverse Workload Needs with High Efficiency

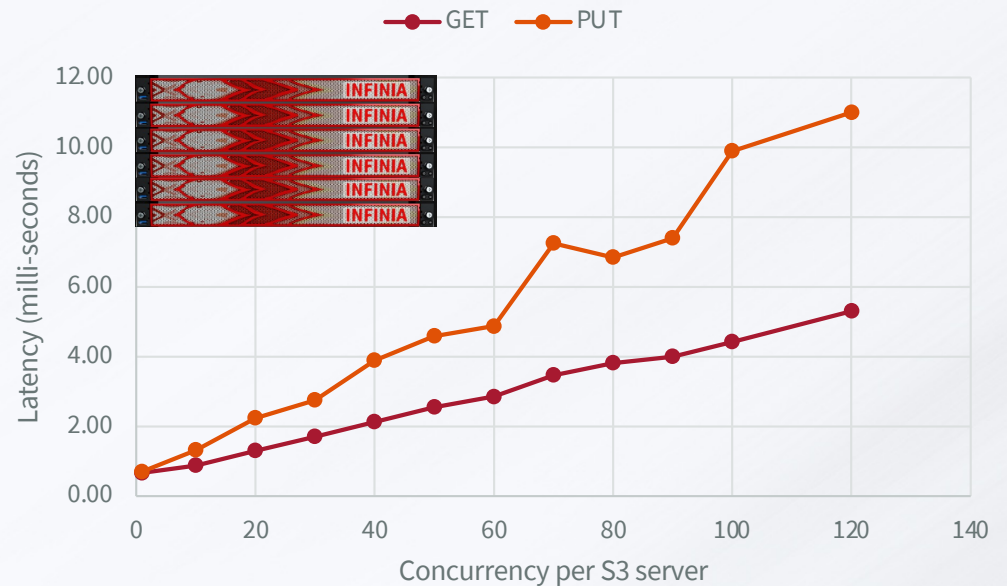
- Autonomous optimization for **low latency and IOPs and throughput and metadata intensive workloads**
- No read/modify/write overhead due to log structured operations
- Byte addressable with wide striping for large I/O blocks
- No Complexity for Administration - Elastic space for metadata and data



Accelerate AI Workloads - 100x Lower Latencies

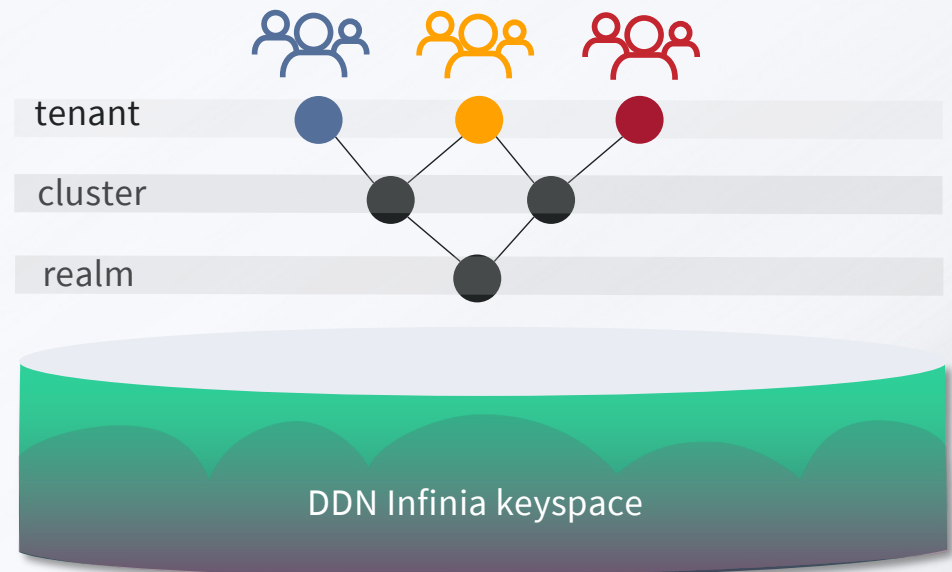
- Infinia is the first Object Storage to deliver **sub-millisecond latencies** for PUTs and GETs
- Over 200K GETs per second from 6 Nodes
- Improve website load times, database response times, queries etc improving end user satisfaction
- Accelerate Enterprise Analytics workloads like Apache Spark, Starburst Presto/Trino, Clickhouse

Infinia Latency - 6 Nodes



Organizational Multi-Tenancy with Maximum Efficiency

- ✓ 100% **Software-based, secure, distributed** keyspace
- ✓ Keyspace is sharded across the entire infrastructure for **maximum resource sharing and highest efficiency**
- ✓ The keyspace tree structure divides to track the site, clusters, tenants, subtenants and datasets – allows tagging, search, advanced AIOPs



DDN Infinia Automatically Delivers The Right Data Resilience

Specify Site, Rack, Subrack, Slot so that Infinia can manage data protection with knowledge of fault domains and geo-locality

The screenshot shows the DDN Infinia management interface. At the top, there are navigation tabs for Dashboard, Health, Performance, and Manage. The 'Servers' section is active, showing a list of 12 servers. A red line highlights the 'Site-Rack-SubRack-Slot-Size' column in the table, which is linked to the text box on the left. The table contains the following data:

IP / Host	Cluster	Site-Rack-SubRack-Slot-Size	Resources used	Management Role
200.100.200.14	production	A5 - 4 - 4 - 12U	All	redagent, redapi, redclient, redui, etcd
12.3.5.1	test	A1 - 0 - 0 - 2U	All	redagent, redapi, redclient, redui, etcd
12.3.20.10	test	A2 - 1 - 1 - 23U	Custom	redagent, redapi, redclient, redui, etcd
12.3.20.12	test	A3 - 2 - 2 - 1U	All	etcd
MyHost23	test	B10 - 4 - 4 - 2U	All	redapi, redclient
200.100.200.10	production	B11 - 5 - 5 - 2U	Custom	etcd
200.100.200.11	production	B12 - 5 - 10 - 15U	All	redapi, redclient
200.100.200.12	production	FD93 - 2 - 2 - 2U	All	etcd
200.100.200.16	production	A1 - 3 - 6 - 2U	All	etcd
200.100.200.17	production	A2 - 2 - 8 - 2U	All	redapi, redclient
200.100.200.18		A3 - 1 - 9 - 2U	All	redapi, redclient
200.100.200.19		A4 - 4 - 10 - 1U	All	redapi, redclient


DDN and NVIDIA BlueField-3 DPU

- ✓ **Accelerates AI Frameworks**
Over 10x faster data to applications gives you up to 50% more output from your Infrastructure
- ✓ **Simplifies Networking**
Create Ethernet Clouds without complexity, Scale to wherever your business needs
- ✓ **Keeps Data Safe**
EXAScaler Optimizes Network traffic and Secures Tenants allowing quick and Safe Enablement of Shared Infrastructure



Transforming the DataCenter With NVIDIA Bluefield

- DDN Infinia **Merlin** Systems leverage NVIDIA Bluefield exclusively to reduce power/improve performance and efficiency
- DDN Infinia Microservices deployment allows a dramatic re-architecture of the datacentre, leveraging 100x scale of NVIDIA Bluefield to:
 - Reduce latency
 - Create a more efficient datapath
 - Scale applications





DDN Infinia Cloud Services

S3 Storage Service
scalable object storage

Elastic Block Storage
Consistent, low-latency for apps, K8s, DBs

Elastic File System
fully managed file storage

Archival
low-cost storage for archiving and backup

Storage Gateway
enables on-prem apps to use cloud storage.

EXAScaler PFS
high-performance file system for AI/HPC



QoS | Performance | Multi-Tenancy | Encryption | Data Reduction | Sub Millisecond Latency | Elastic Scaling

DDN is Built for the Scale and Complexity Challenges of AI

DDN EXAScaler solves the challenges of

- Extreme scale
- Maximum efficiency
- Maximum BW and IOPs

DDN Infinia solves the challenges of

- Data Distribution
- Model Governance
- Data Sharing
- Multi-Tenancy
- Data Management



