



# DDN – UPDATE for 2020

**DDN**<sup>®</sup>  
AI • BIG DATA • HPC

2019/12/12 @ PCCC19

DataDirect Networks Japan, Inc.

橋爪信明

[nhashizume@ddn.com](mailto:nhashizume@ddn.com)



At Scale Data Services  
for AI, Big Data and  
HPC



Unlimited Performance for  
your most Demanding  
Workloads



Optimized AI Platforms For  
Every Use Case

---

# DDN<sup>®</sup> AT SCALE ENTERPRISE

---



High Performance Flash &  
Hybrid Unified Storage



SW Defined Storage for  
Telco 5G, IoT and the  
Enterprise



Simplicity and Control for  
Virtualized Environments



# DDN EXA5

EXASclaer 5.0

# DDN EXA5 X-APPLIANCES



## ES200NVX

**24r/20w**  
**1.5M IOPS**

**UP TO 24 DRIVES**

**ALL NVME**



## ES400NVX

**48r/30w**  
**>3M IOPS**

**UP TO 384 DRIVES**

**+0 +1 +2 +4**



## ES7990X

**24r/20w**  
**800K IOPS**

**UP TO 450 DRIVES**

**+0 +1, +2, +4**

**HYBRID**



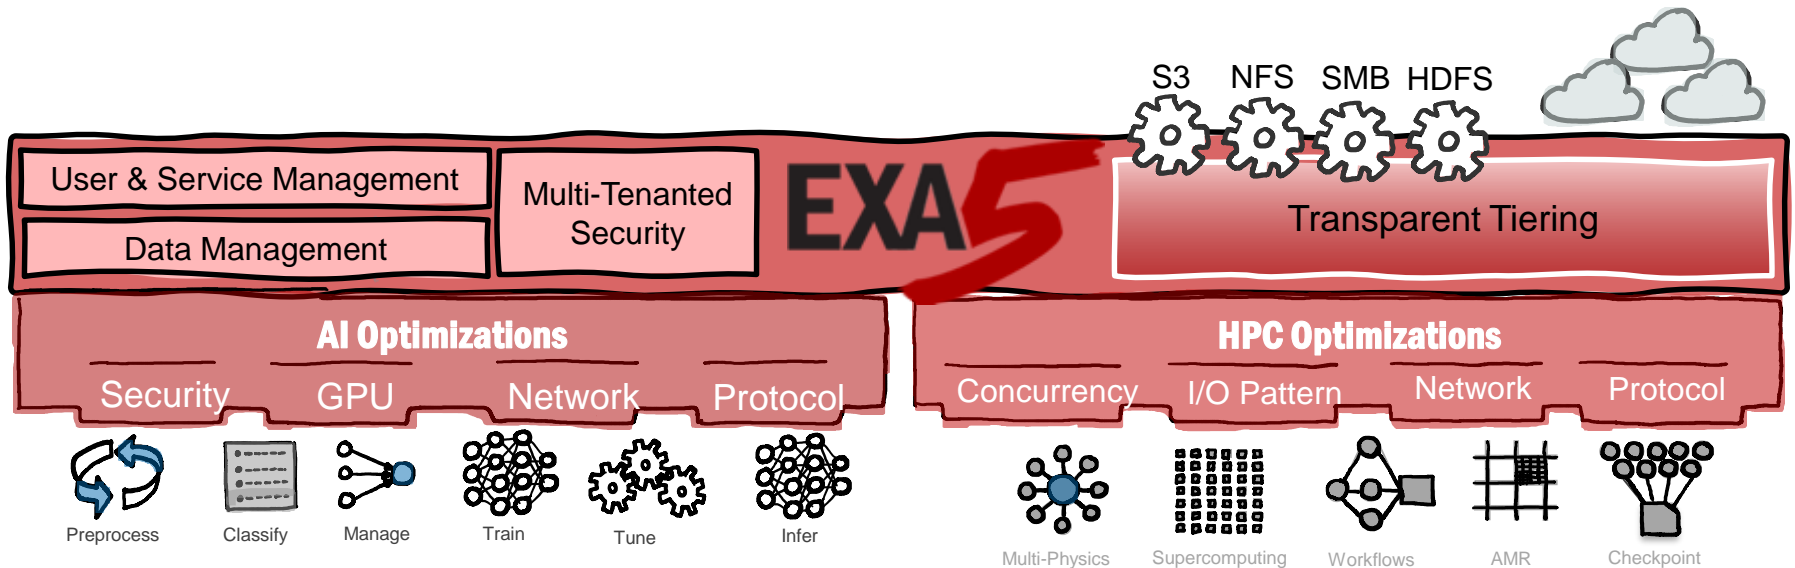
## ES18KX

**70r/70w**  
**>3M IOPS**

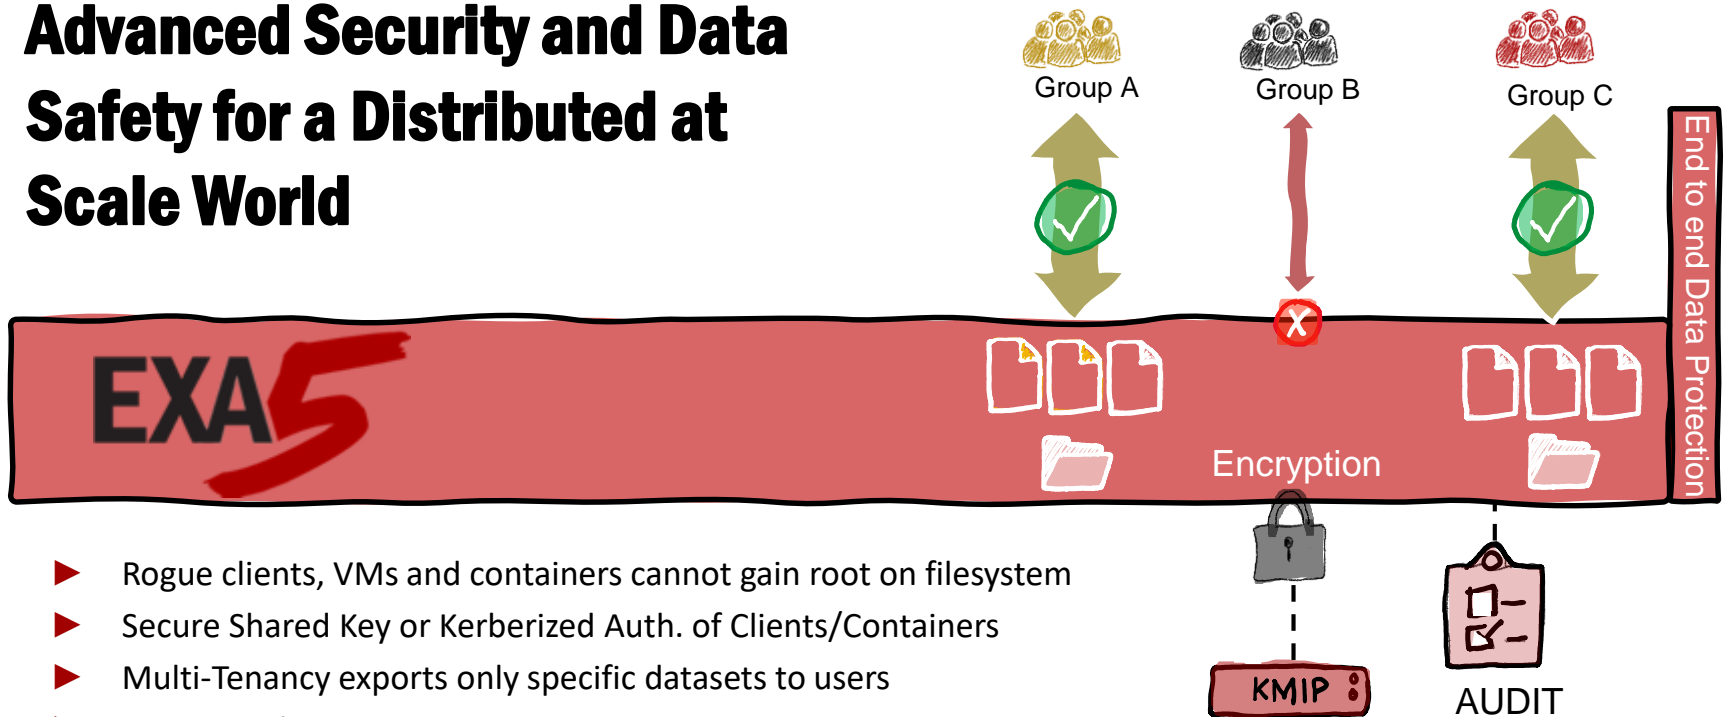
**UP TO 1872 DRIVES**  
**+0 +5, +6, +8, +10, +16, +20**

# The Intelligent, Optimized Environment for AI and HPC

- ▶ Deep Optimizations for both AI and HPC delivers for the highest efficiencies and the right capabilities
- ▶ *Your data in the right place at the right time*



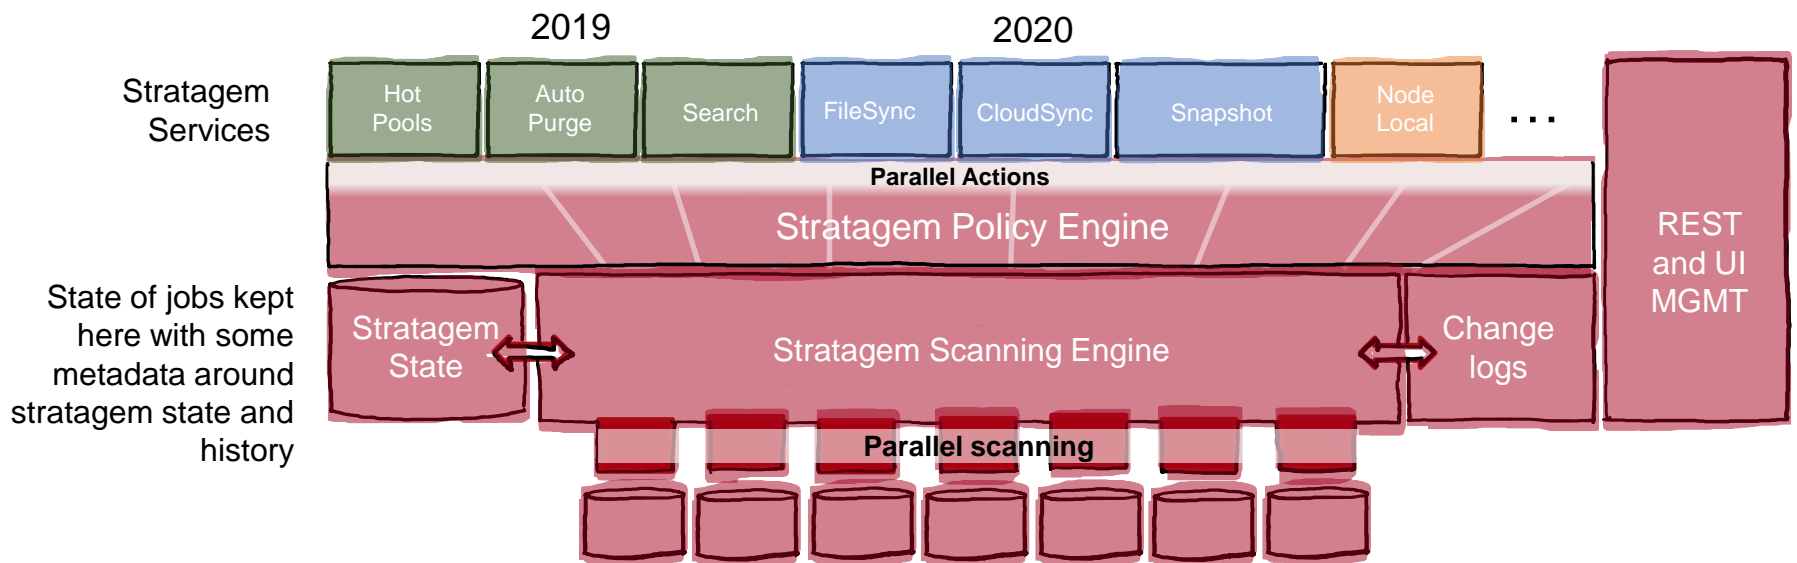
# Advanced Security and Data Safety for a Distributed at Scale World



- ▶ Rogue clients, VMs and containers cannot gain root on filesystem
- ▶ Secure Shared Key or Kerberized Auth. of Clients/Containers
- ▶ Multi-Tenancy exports only specific datasets to users
- ▶ Secure Audit Logging
- ▶ Data at Rest Encryption for Drives and NVMe\*


\*NVMe SED available Q1 2020

# EXA5 Stratagem Data Management Architecture



# EXA5 HOT POOLS: At Scale Transparent Tiering

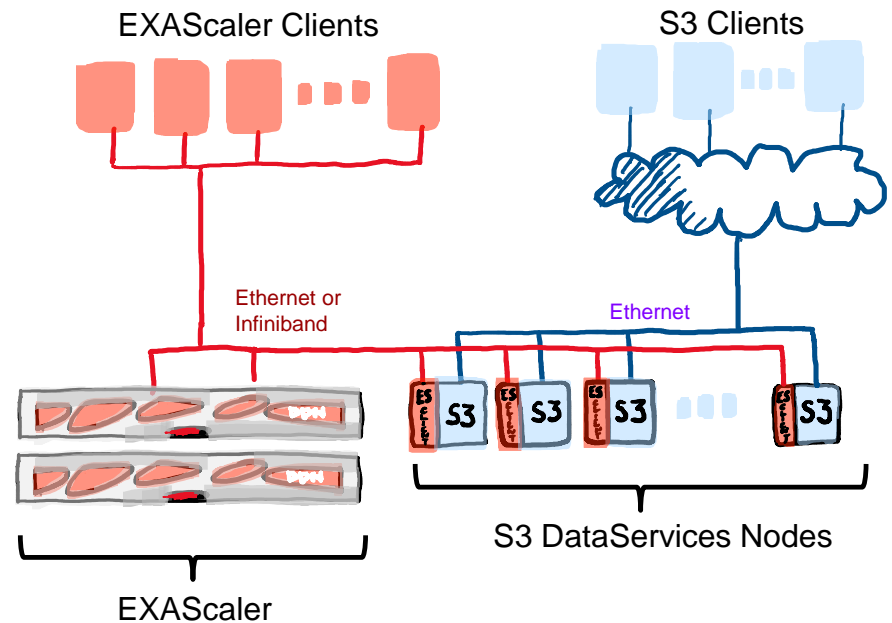
- ▶ Hot Pools V1 actively migrates old, least accessed files from flash tier to HDD tier
- ▶ API provided for use cases with externally driven data movement (e.g. scheduler directives)
- ▶ Hot Pools V2 adds capability for also promoting hot objects in the HDD pool into SSD pool
- ▶ Hot Pools uses Stratagem File Scanning & File Level Replication





# EXA5 S3 Data Services

- ▶ Highly Available Scale-Out S3 Data Services With Single S3/File Namespace
- ▶ Initial v1 Release Requires at least 2 External Data Services Nodes (min 32GB RAM, 8 Core (TBD))
- ▶ Data stored via the S3 API shall be stored on EXAScaler as files in the file system.
  - **V1:** PUT/GET through S3 and read-only through EXAScaler File interface.
  - **V2:** adds write capability through EXAScaler as Files so that files are accessible as S3 objects
- ▶ Files located in EXA at  
"/s3\_content/<bucket>/<filename>..."
- ▶ S3 object metadata shall be written into EXAScaler as extended attributes so that they are searchable through filesystem tools.
- ▶ Management REST API for system admin





# DDN A3I

**Accelerated, Any-Scale AI**

# DDN A<sup>3</sup>I X-APPLIANCES

## NEXT GENERATION AI DATA PLATFORMS



### AI200X

**24r/20w**  
**800K IOPS**  
**32/64/128/256 TB**



### AI400X

**48r/32w**  
**1.5M IOPS**  
**32/64/128/256 TB**




### AI7990X

**24r/20w**  
**800K IOPS**  
**1/2/4 PB**

**HDR100/EDR IB OR 100GbE • FULLY INTEGRATED AND OPTIMIZED FOR AI AND DL • VALIDATED UP TO 96 NVIDIA DGX-2s!**

# NVIDIA IS THE AI400X LAUNCH CUSTOMER



# **NVIDIA IS THE AI400X LAUNCH CUSTOMER**

**10 APPLIANCES**

**4 HOURS TO DEPLOY**

**400 GB/s RD, 310 GB/s WR**

**40 AI400X IN 2020 – 15 PB ALL**

**NVME**

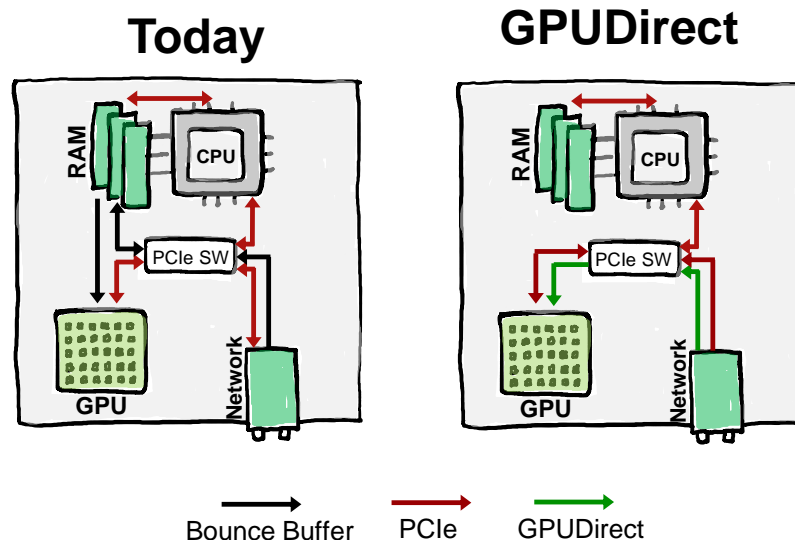
**JOINT IO500 SUBMISSIONS**



# DDN A<sup>3</sup>I SOLUTIONS

## Magnum IO (GPUDIRECT TO STORAGE)

- EXAScaler Integration with GPUDirect\* enables Direct Memory Access from EXAScaler to GPU memory
- 2x-8x higher BW data transfers between Storage and GPU.
- 3.8x lower latency with no faulting and bounce buffers
- Stable and flat latencies as GPU concurrency increases.
- Lower consumption of host CPU or memory subsystem
- The GPU is the computing element with the highest IO bandwidth, e.g. 215 GB/s vs. the CPU's 50 GB/s.
- Very fast access to petabytes of remote storage faster than even the page cache in CPU memory.



# A3I WITH GPUDIRECT

## DOUBLE NVIDIA DGX-2 THROUGHPUT!

**80 GB/s per client,  
20 X performance gains**


Native client integration with EXA5, fully-transparent to users and applications

Enables a direct path to transfer data between GPU memory and data storage

Eliminates unnecessary memory copies, lowers CPU overhead, reduces latency, bypasses hardware architecture limitations

Improves AI, DL, HPC application performance

### GPU READ THROUGHPUT WITH AI400



**5 DGX-2**  
**10 AI400**  
**400 GB/s**


EXA5 GPUDirect Integration delivers  
up to 80 GB/s of throughput per DGX-2

Enables a direct path to transfer data  
between GPU memory and data storage

Performance scales linearly and provides  
maximum at-scale application acceleration



### GPU READ THROUGHPUT SCALING WITH TEN AI400s



**DDN DELIVERS LINEAR PERFORMANCE SCALING**



# DDN A<sup>3</sup>I SOLUTIONS – NVIDIA SUPERPOD REFERENCE ARCHITECTURE

- **AI400 + DGX-2 SuperPOD at-scale testing and validation published by NVIDIA:**
  - ▶ The AI400 All-flash appliance delivers incredible sequential and random read performance, as required by the heaviest DL workloads.
  - ▶ Metadata performance scales well from 1 to 96 nodes, with no degradation as the number of nodes and threads increases.
  - ▶ The AI400 is a fully-integrated platform that's easy to deploy. DDN provides excellent technical deployment and support services.
- **RA document available from NVIDIA website**
  - ▶ <https://www.nvidia.com/content/dam/en-zz/solutions/data-center/documents/nvpod-superpod-ddn-ra09734001.pdf>



**NVIDIA DGX-2 SUPERPOD  
REFERENCE ARCHITECTURE**

# Thank You!

Keep in touch with us



[Team-JPSales@ddn.com](mailto:Team-JPSales@ddn.com)



Tokyu Bancho Bldg. 8F  
6-2 Yonbancho Chiyoda-ku,  
Tokyo 102-0081



[@ddn\\_limitless](https://twitter.com/ddn_limitless)



+81-3-3261-9101  
+81-3-3261-9140



[company/datadirect-networks](https://www.linkedin.com/company/datadirect-networks)