

NetApp INSIGHT 2023

Always On

数据存储与管理前沿技术系列课程

第三讲：重新定义 AI 的可能性

Redefining what is possible with AI

2023年11月28日



目录

Contents



- **The AI & Analytics Opportunity**
人工智能和分析存储项目机会

AI在液晶面板行业—AI之眼



准确率：90%+；

识别速度：从2秒缩短到了500-600毫秒；

人工替代：140+人；

成本优化：降本超过1000万人民币；

华星光电导入了国内面板业缺陷判别的首例人工智能项目——ADC项目（Auto Defect Classification 自动缺陷分类），实现对面板质检环节的智能化；

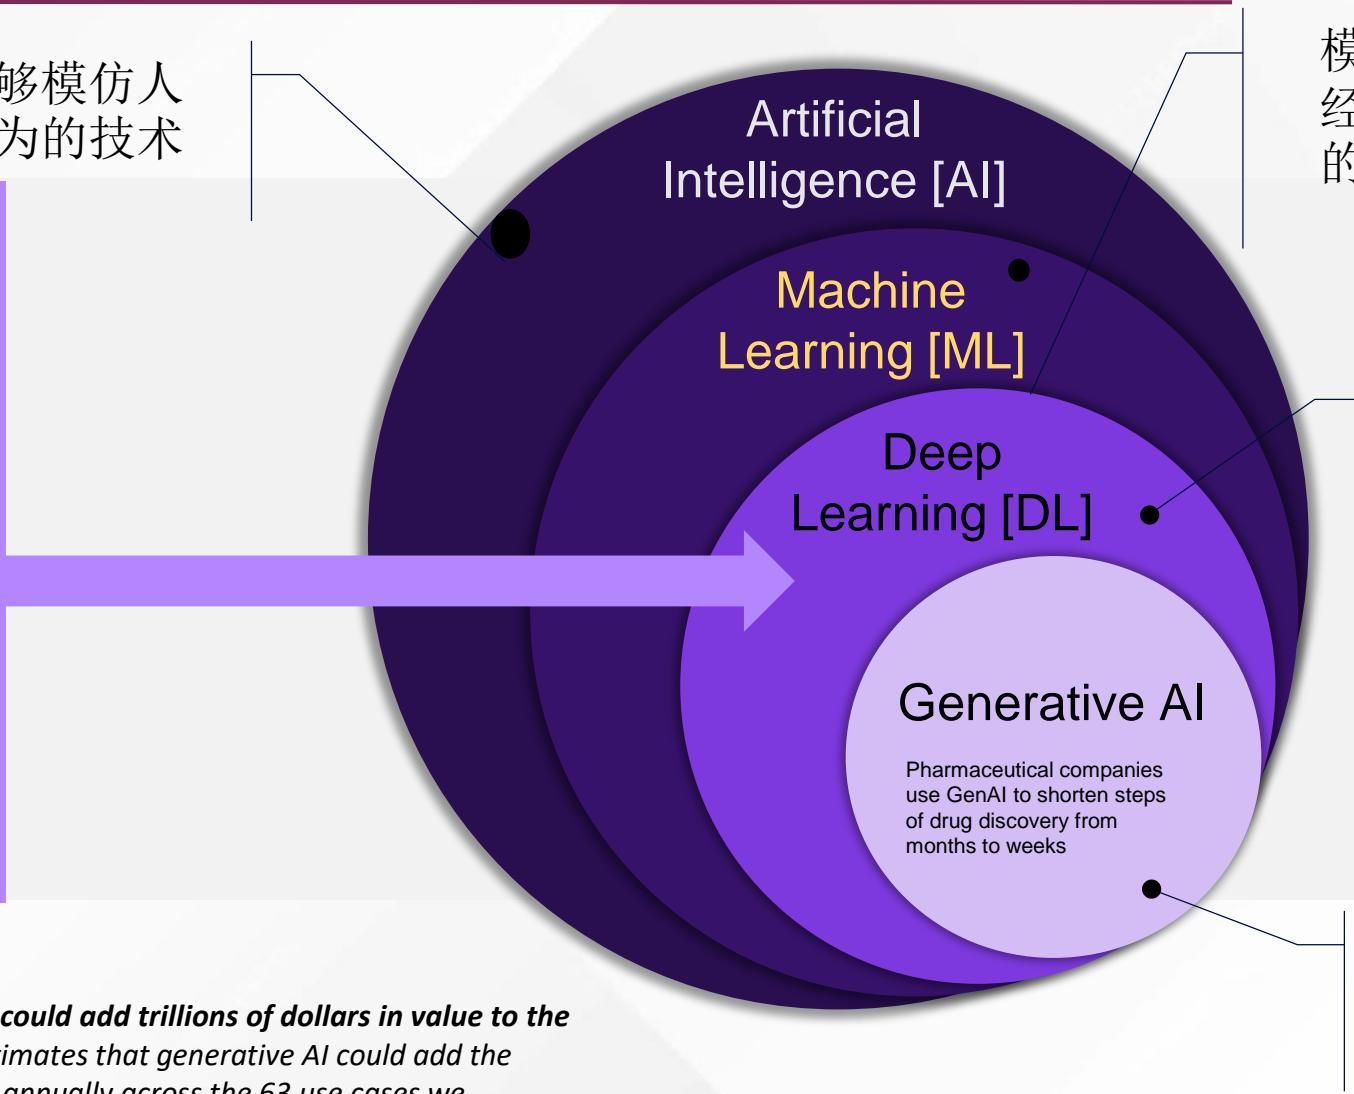
在华星光电工厂的Monitor Room中，每名图片质检员每天需要面对1万张产品图片进行检测，一名成熟的质检员，平均要花费2秒的时间完成对每张产品图片的缺陷分类工作，同时人员需要确认缺陷是否落在线路上、对产品的影响严重程度是否需要开立对应的生产异常单。



正在不断刷新对AI认知：生成式AI

任何使计算机能够模仿人
类行为的技术

NetApp's 5+ Years AI
Business with 100's of
Customers



模式识别和统计技术，随积累
经验逐步提升完成决策和任务
的能力。

使用多层神经网络和海
量数据算法，经过训练
和学习数据中的潜在特
征，使软件完成决策和
更复杂任务

通过从现有数据中学习
和训练，生成新数据

McKinsey, June 2023

Generative AI's impact on productivity could add trillions of dollars in value to the global economy. Our latest research estimates that generative AI could add the equivalent of \$2.6 trillion to \$4.4 trillion annually across the 63 use cases we analyzed—by comparison, the United Kingdom's entire GDP in 2021 was \$3.1 trillion.

利用 AI 最大限度地提高生产力

The Future of Work: Maximizing Productivity with AI

AI is about uncovering new inputs and rethinking operations to reach new levels of productivity.

While few organizations are realizing the productivity benefits of AI today, many are making the right moves.

One of the largest challenges organizations face is **harnessing the latent potential of a company's data estate.**

To do so, they must consider transforming their data environment **and this is where NetApp can help.**



Across 63 uses cases, generative AI has the potential to create

\$2.6 to \$4.4 trillion

in value through increased productivity across industries

"The economic potential of generative AI"
McKinsey June 2023

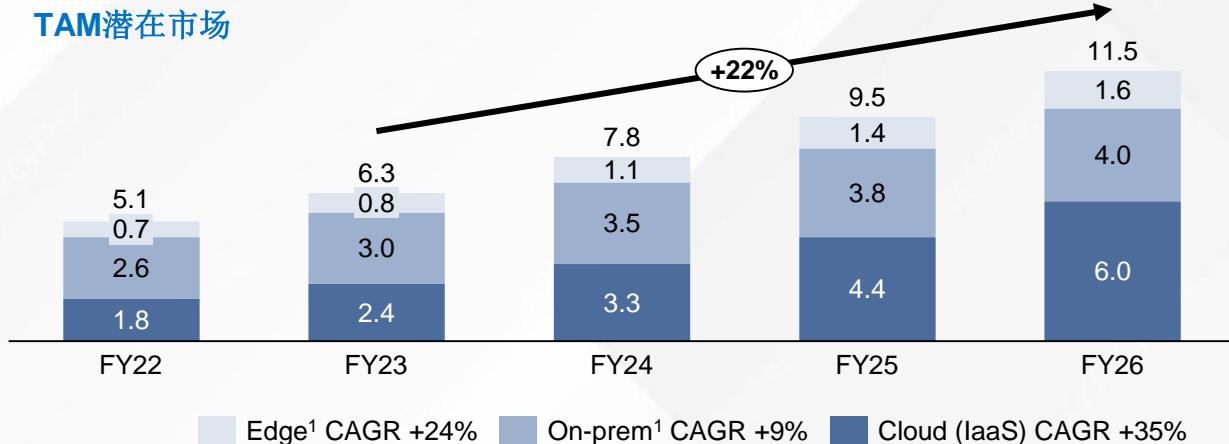
AI对存储需求的预测

The AI market can be ~30% of NetApp revenue in 3-5 years

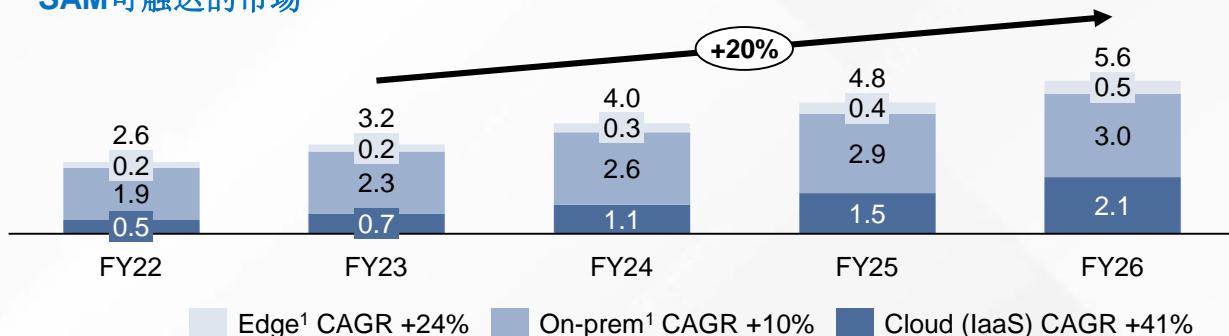
AI storage market forecast

\$B, CAGR (FY23-26)

TAM潜在市场



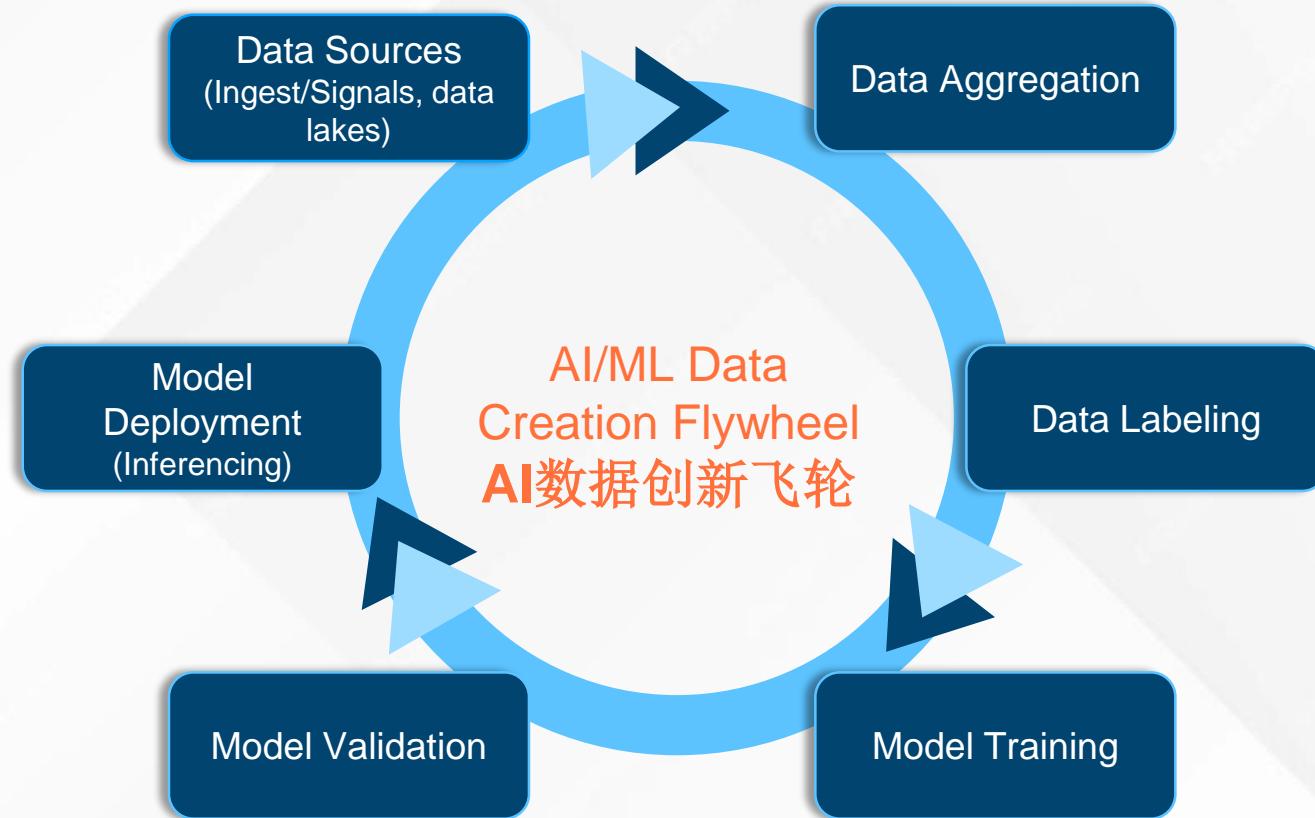
SAM可触达的市场



Note: [1] Includes storage and support revenue

Sources: (a) NetApp CSO, June 2023; (b) 2022H1 IDC AI tracker; WW IaaS/PaaS workloads overview and trends, 2023; (c) WW Public Cloud IaaS and PaaS workloads forecast, 2021

Data fuels AI - 数据推动人工智能 (Update)



数据架构不够灵活，无法跟上人工智能的快速发展

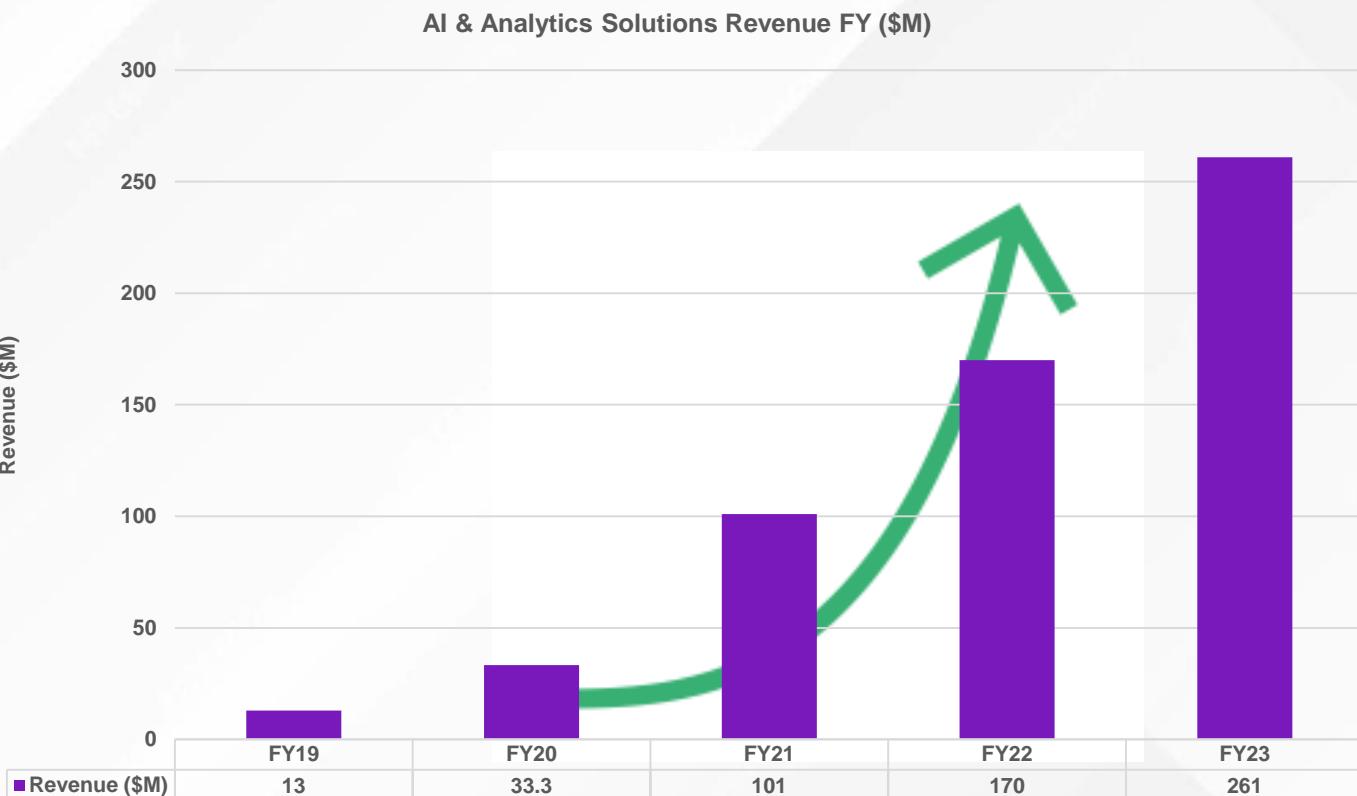


 Complexity	 Scaling AI projects
<ul style="list-style-type: none">▪ Uncertainty and difficulty of DIY infrastructure▪ Fast evolving software stack	<ul style="list-style-type: none">▪ Scaling across multiple AI projects▪ Handling multiple data sources especially across clouds
 Data Integration	 Data Security
<ul style="list-style-type: none">▪ AI Environment troubleshooting, ROI, visibility-monitoring▪ Support for multivendor technology stack	<ul style="list-style-type: none">▪ Outdated policies and frameworks▪ Tools that find attacks rather than prevent them

The root of the problem: **overly streamlined and simplified data environments**

人工智能和分析：增长最快的市场之一

Driving Flash and Object



102% NetApp AI Sales Growth

CAGR last 3 FYs



Data fuels AI. We fuel data (Update)

5x

Run 5 times more data through your pipeline.

<60 seconds

Copy datasets in seconds rather than in hours or days.

~20 minutes

Quickly configure your AI infrastructure with Ansible integration.

NetApp AI has won eight AI industry awards since 2019

A Leader in Gartner Magic Quadrant for Primary Storage Arrays

A Leader in Cloud File Systems

A Leader in Object Storage

A Leader in File-Based Storage



4 years of AI experience with hundreds of published documents and 30+ dedicated AI team members

Top NVIDIA partner, 4 years of continual development with NVIDIA

Over 400 customers worldwide

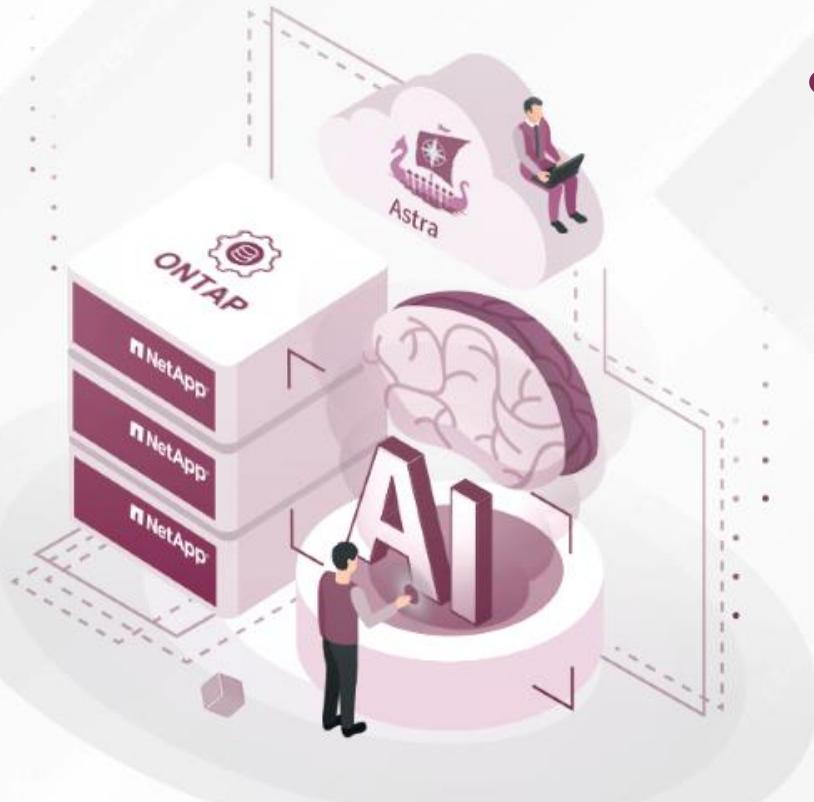
>50% year-over-year revenue growth

数据为人工智能提供动力。我们NetApp为数据提供动力。

目录

Contents

- NetApp AI 解决方案

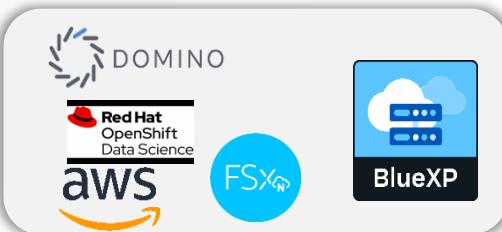


NetApp AI 在2023年关键更新



AI Training & Hybrid MLOps

Hybrid MLOps platform provides consistent experience for data scientist
Domino Data Labs and Red Hat OpenShift AI



Solution Update

NVIDIA BasePOD with NetApp AI
Reference Architecture,
DGX A100 & DGX H100



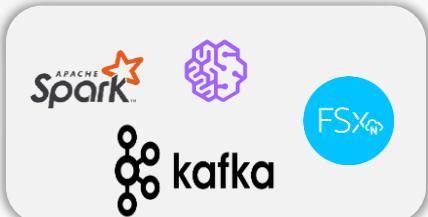
AI Training & Hybrid MLOps

Reduced cost Flash
NFS over RDMA with GDS



SageMaker, Kafka, Spark FSxN S3

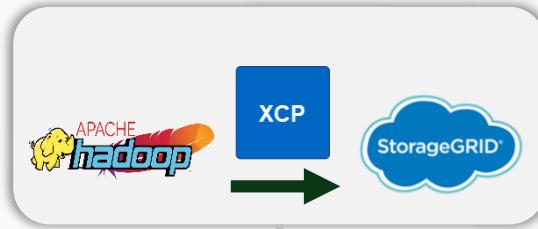
Optimized workloads on Object



C-Series QLC

Modernize Hadoop

Unlock PBs of storage for
AI and Analytics



ML/DL Training Performance Improvement

NFS over RDMA with GDS
Keeps scarce GPUs processing



INNOVATION STREAM

2023年新发布AFF-C系列增强

BlueXP 登录

NetApp TV

 NetApp

您想搜索什么内容?



解决方案和产品 ▾

走近 NetApp ▾

如何购买 ▾

支持和培训 ▾

AFF C 系列为基于 NVIDIA DGX BASEPOD 的可扩展 AI 创建了一个新的入口点

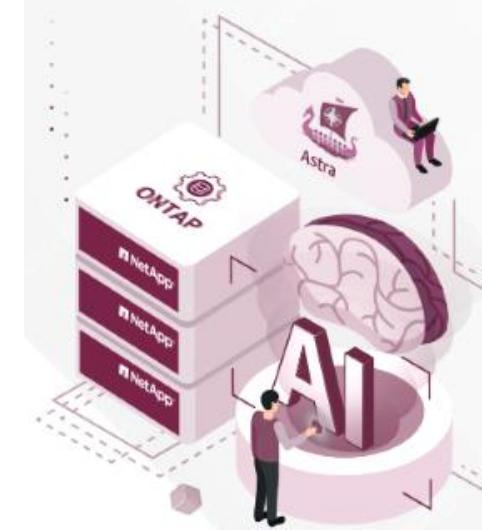


目录

- 简化 AI 解决方案
- 降低成本壁垒
- 从小规模起步，逐渐扩展

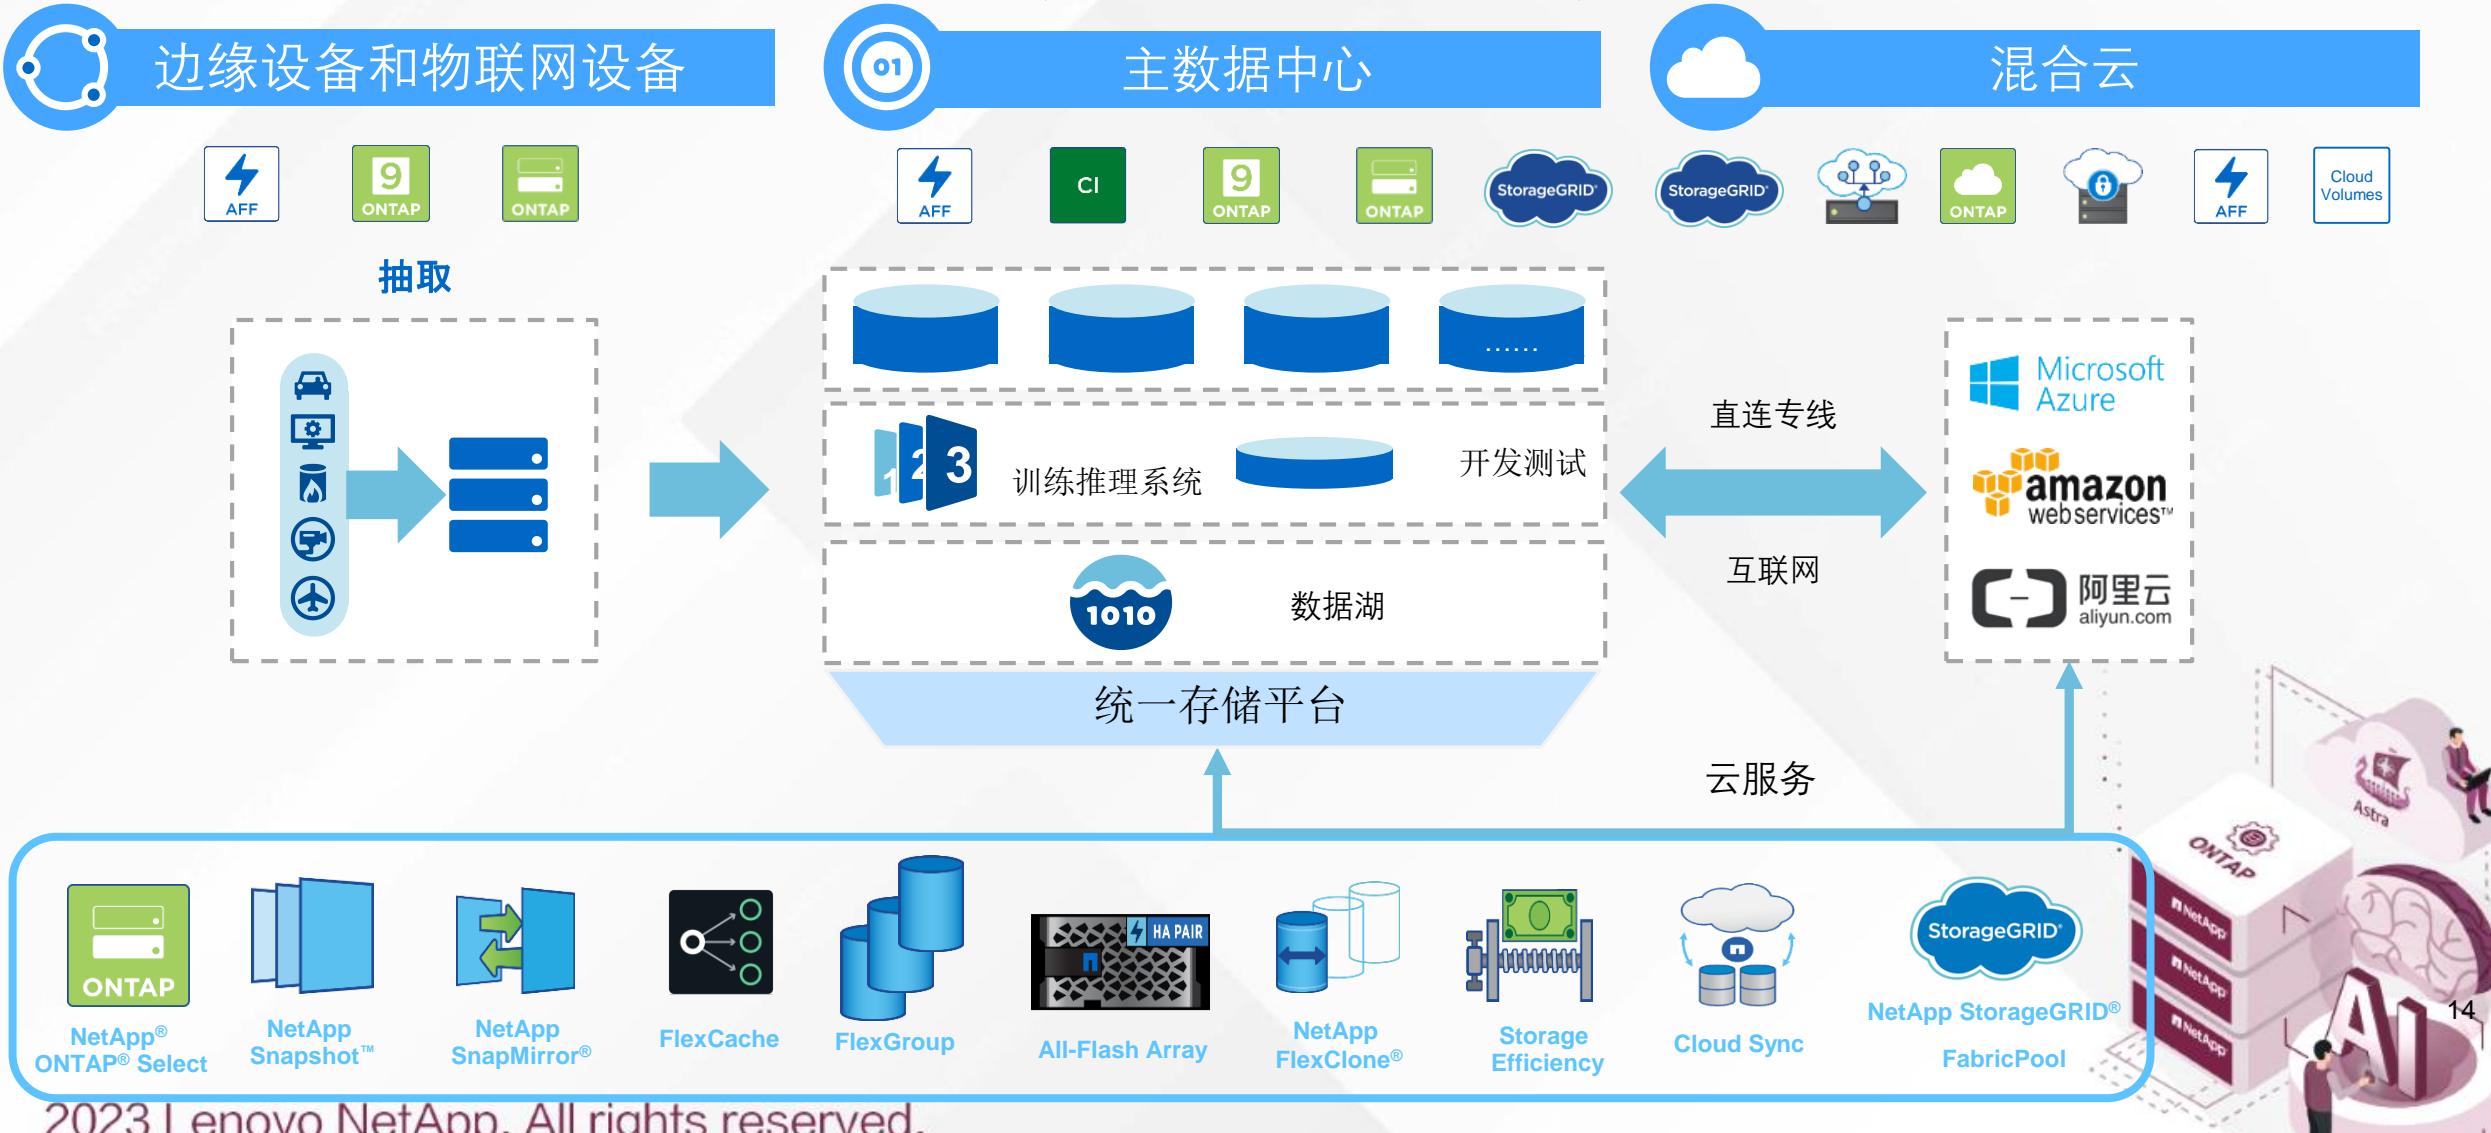
**Mackinnon Giddings**
2023-10-25

在当今快速发展的技术环境下，对可扩展人工智能 (AI) 解决方案的需求正在迅猛增长。但是，构建和管理 AI 基础架构可能既复杂又昂贵，这给技术



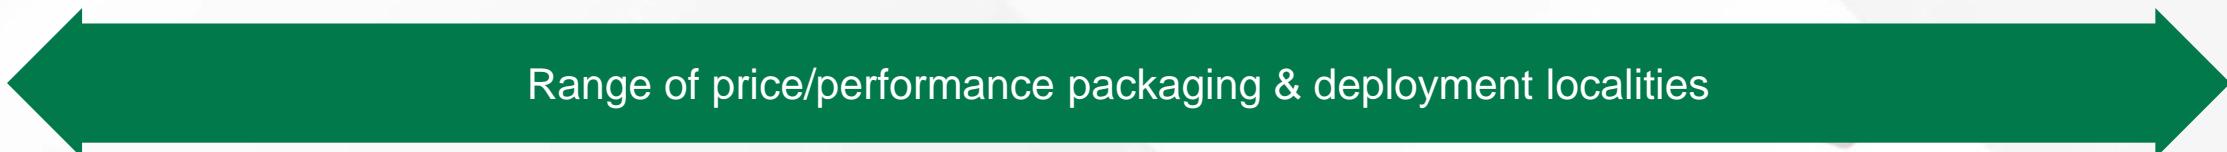
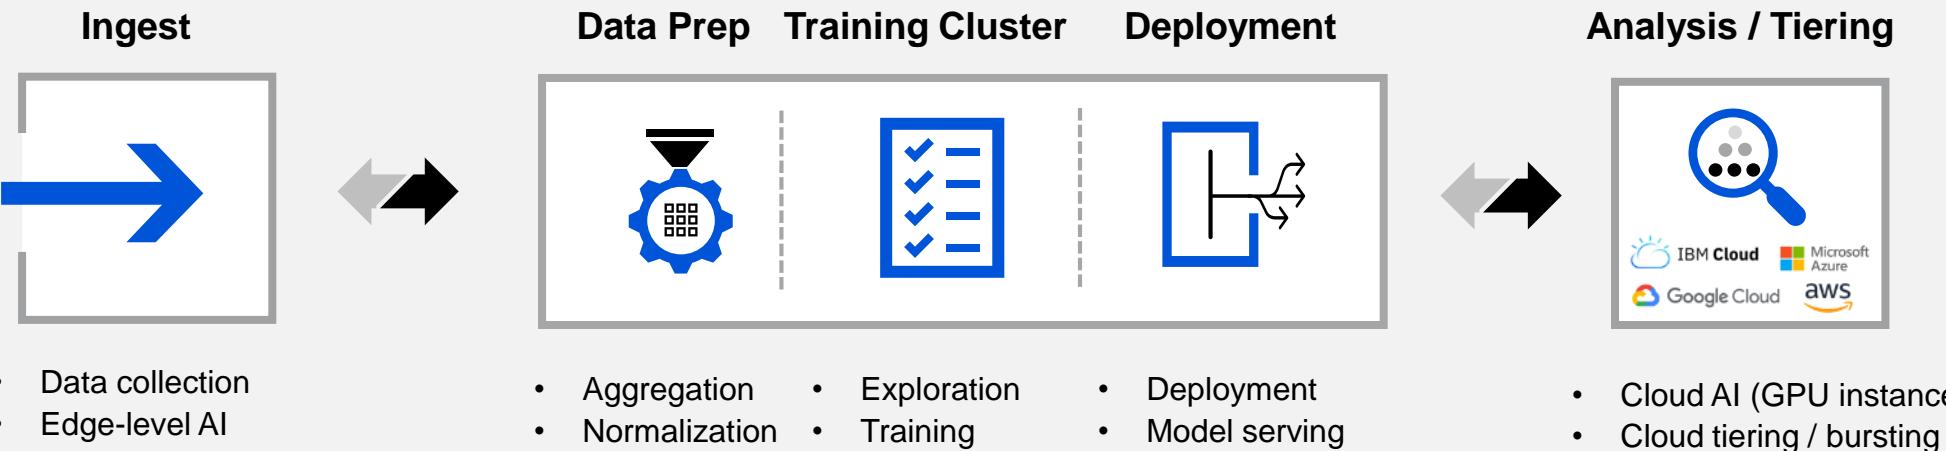
边缘-核心-云无缝的Data Fabric数据流动管理

在满足可靠性、性能、敏捷性要求的同时，减少了数据副本数量，降低了数据流动成本



NetApp 平台与 AI 数据管道的一致性(更新)

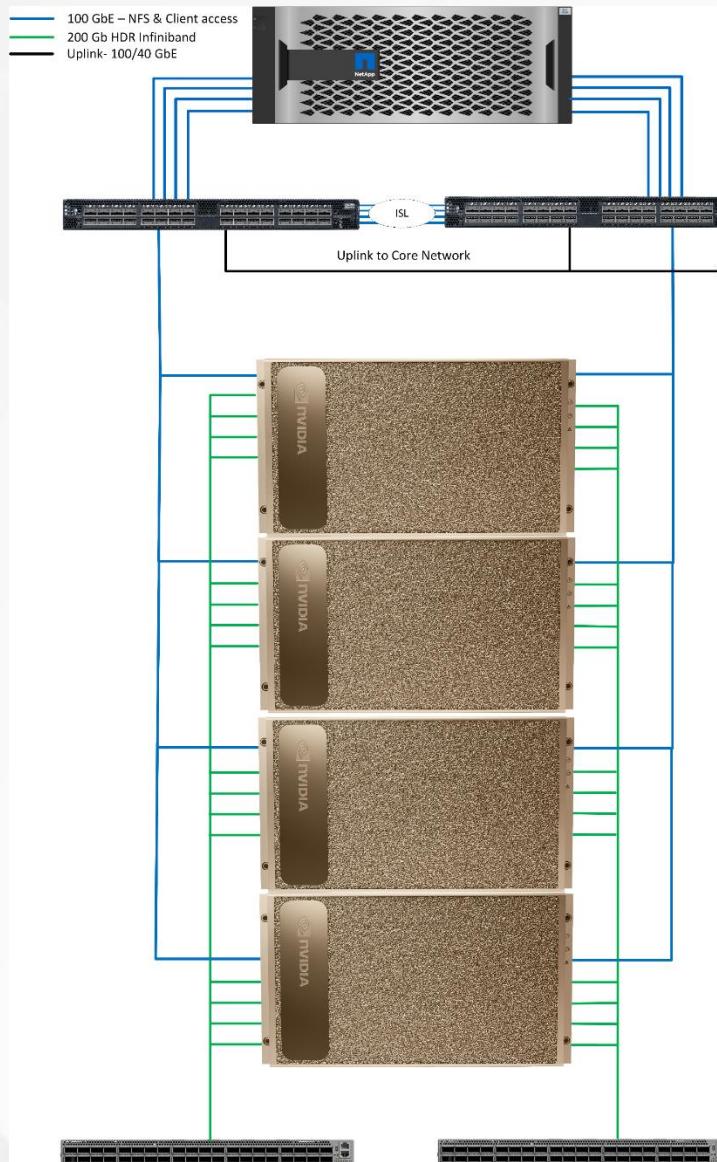
AI Data Pipeline (simplified)



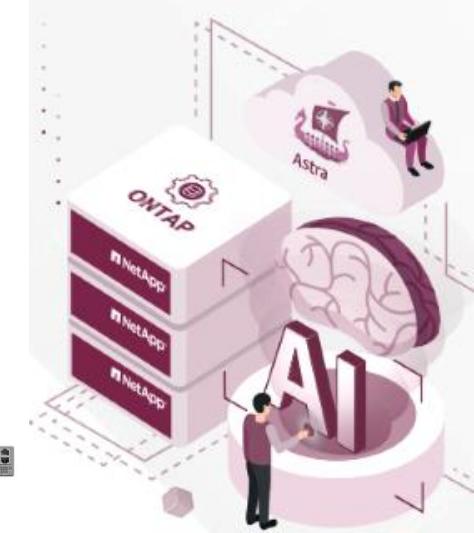
NetApp ONTAP AI with DGX A100/A800 参考架构

4节点DGX配置为例

- DGX A100/A800 服务器
 - 8x A10/A800 0 SXM2 GPUs –
 - 55k CUDA cores, 3456 TensorCores – 高达5 PFLOPs FP16
 - 600GB/s NVLink
 - 8x HDR IB端口, 2x 100Gb 以太网端口
- AFF A800 / C800 存储系统
 - 1x HA Pair / 4 A100 节点
 - 48x NVMe SSD
 - 100GbE 网络
- 独立网络设计
 - 2x Mellanox QM8700 计算网络
 - 200GB HDR InfiniBand
 - 2x Mellanox SN3700 存储和客户端访问网络
 - 100GB Ethernet (NFS网络)



POWERED BY

C系列让参考架构可灵活配置



	Starter	Small	Medium	Large	
GPU Compute	1x DGX A100	2x DGX A100	4x DGX A100	8x DGX A100	
File Storage	1x NetApp C400	1x NetApp C400	1x NetApp C800	2x NetApp C800	
Networking	1x SN3700, 1x AS4610	2x Mellanox SN3700, 2x Mellanox QM8700, 1x Mellanox AS4610			
Infrastructure	Mgt Svrs, 1x Rack, PDUs	Management servers, racks, PDUs			

In addition to the 4 standard configurations above, customers can now specify any number of DGXs 1-8, and their choice of an A400, A700, or A800. Arrow will quote pricing on this custom configuration plus the additional BOM elements necessary for the ONTAP AI integrated solution.



AI客户首选NFS数据协议

基于用户选择的调研数据

- 数据类型（随机小文件和顺序大文件）
 - 传感器数据
 - 日志
 - 图片（激光雷达、高清摄像头）
 - 激光点云、地图、决策
 - 数据库(RDBMS/NoSQL)

Question: What is your choice of file system for the AI/DL Data Tier		Answers:
NFS		32.38%
HDFS		20.00%
S3		15.24%
GPFS		10.48%
Ext4		7.62%
Lustre		6.67%
ZFS		4.76%
XFS		2.86%

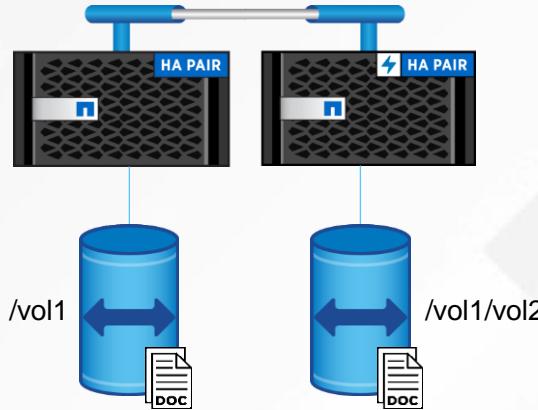
应用场景	充足的读缓存	DGX缓存能力	推荐网络类型	网络文件系统选择
数据分析	NA		10Gbe	对象存储, NFS, 或其他并发读及小文件性能优的存储
HPC	NA		10/40/100Gbe IB	NFS, 或其他支持大量客户端, 单点存储性能优的HPC并行存储系统
DL 256x256 图片	Yes	63 million images	10Gbe	NFS或小文件读写效率高的存储
DL 1080p 图片	Yes	13 million images	10/40Gbe IB	高性能NFS, HPC存储系统, 高并发
DL 4K 图片	Yes	5 million images	40Gbe IB	高性能NFS, HPC存储系统, 高并发, 单节点3GB/s+
DL 无压缩图片	Yes	1 million images	IB 40/100Gbe	高性能NFS, HPC存储系统, 高并发, 单节点3GB/s+
DL 不缓存数据集	no	NA	IB 10/40/100Gbe	性能同上, 总的性能需要满足所有应用并发使用的需求

https://docs.nvidia.com/dgx/bp-dgx/index.html#storage_scaling



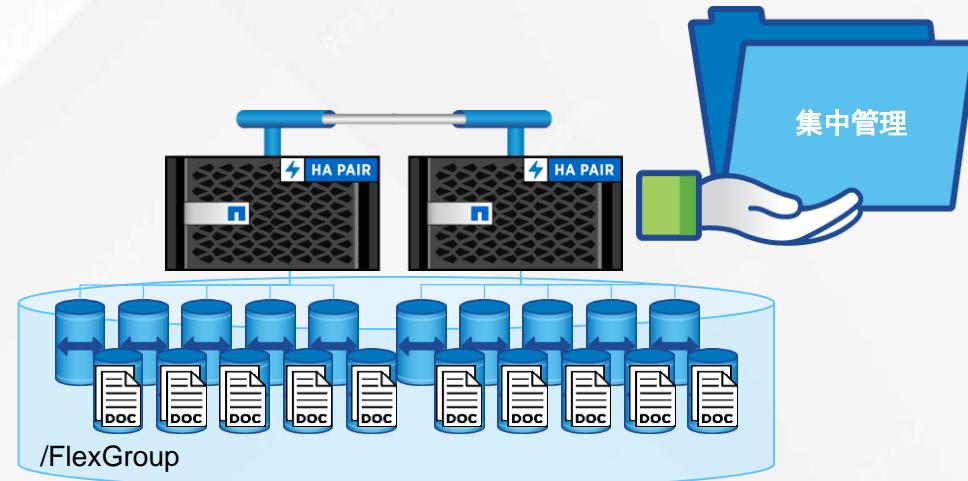
适合海量存储的FlexGroup卷组

技术原理对比



FlexVol volumes 常规卷

- 隶属单一节点，单一Aggr
- 最大支持100TB、20亿个文件
- 基于单个文件系统的管理方式

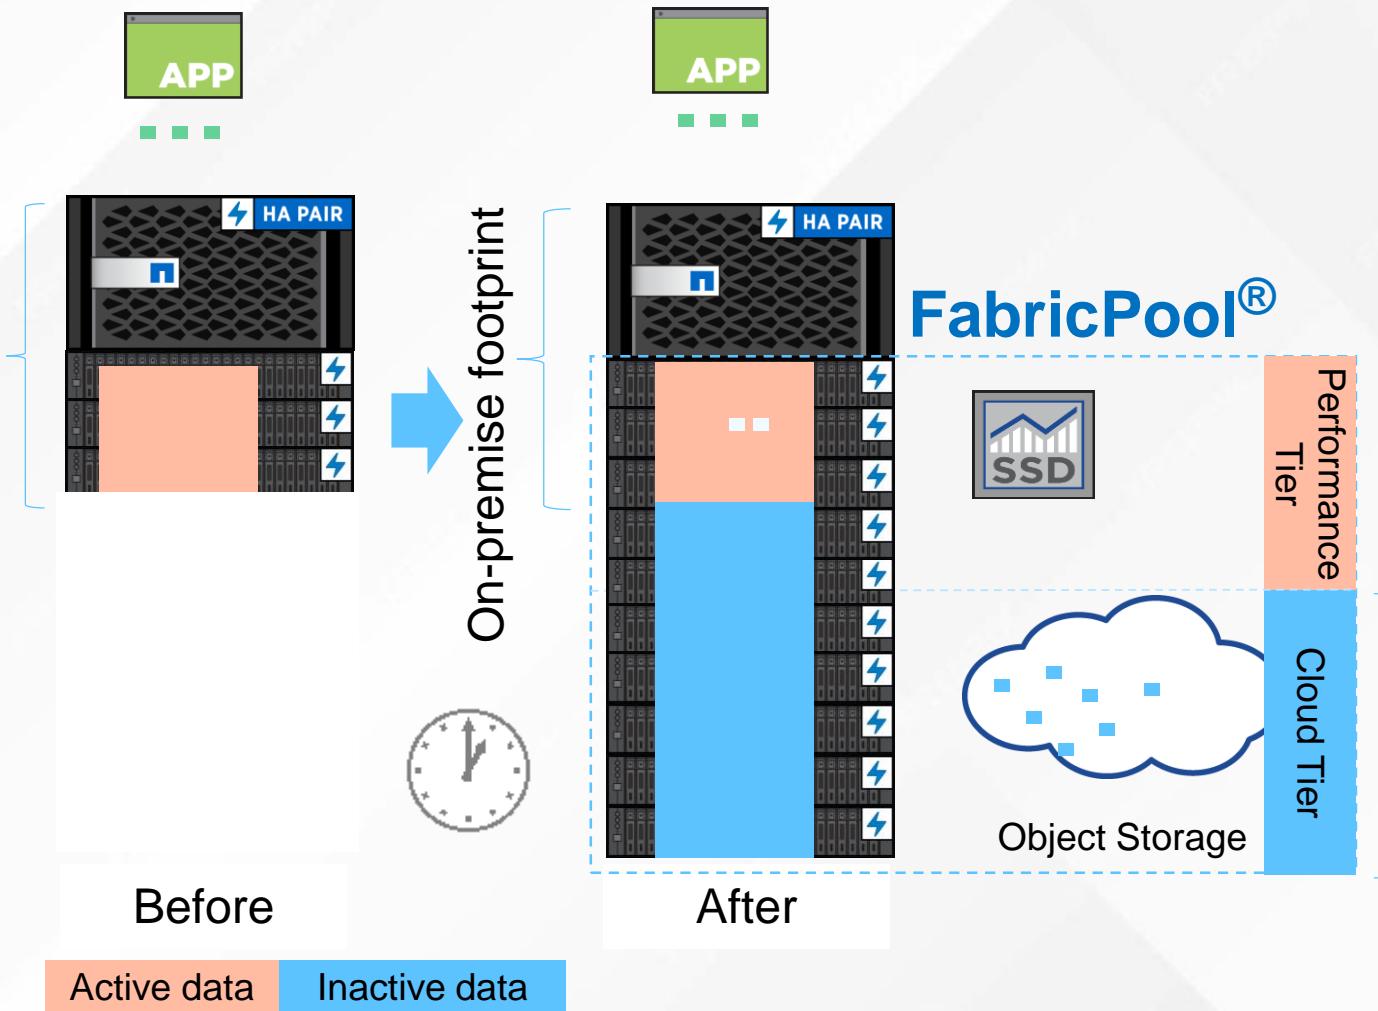


FlexGroup volumes 分布式卷组

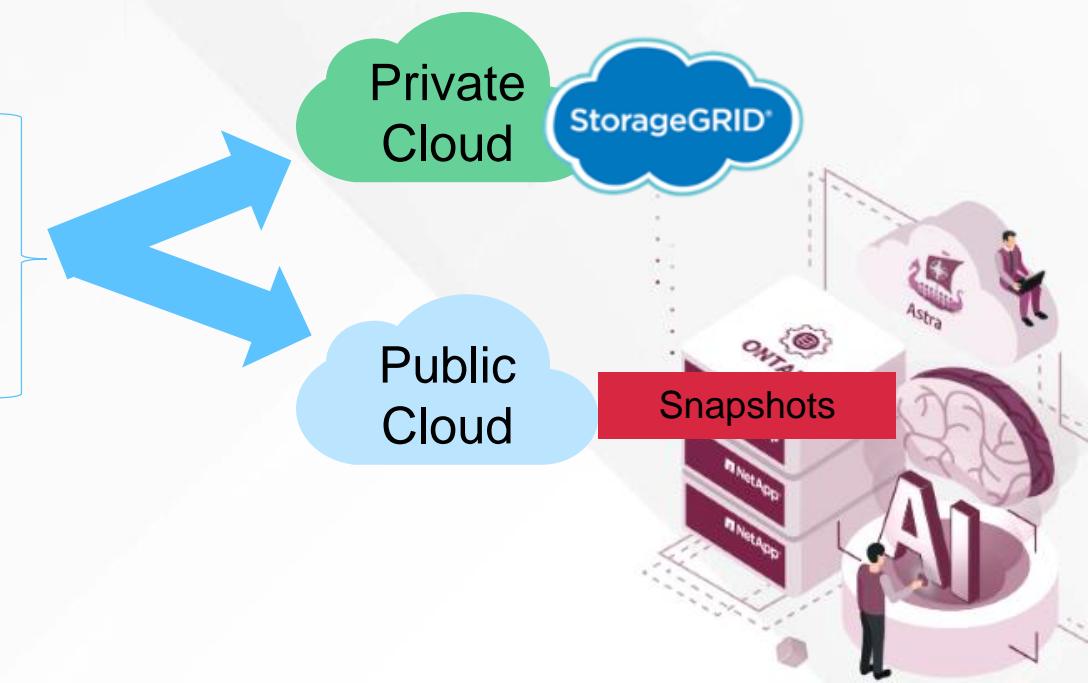
- 单一命名空间，统一管理
- 跨越整个集群所有节点
- 最大支持20PB、4000亿个文件
- 充分利用集群内部硬件资源，实现负载均衡
- 4KB文件块、分布式metadata存储，性能线性增长



全闪高性能AI数据分层到大容量S3对象存储

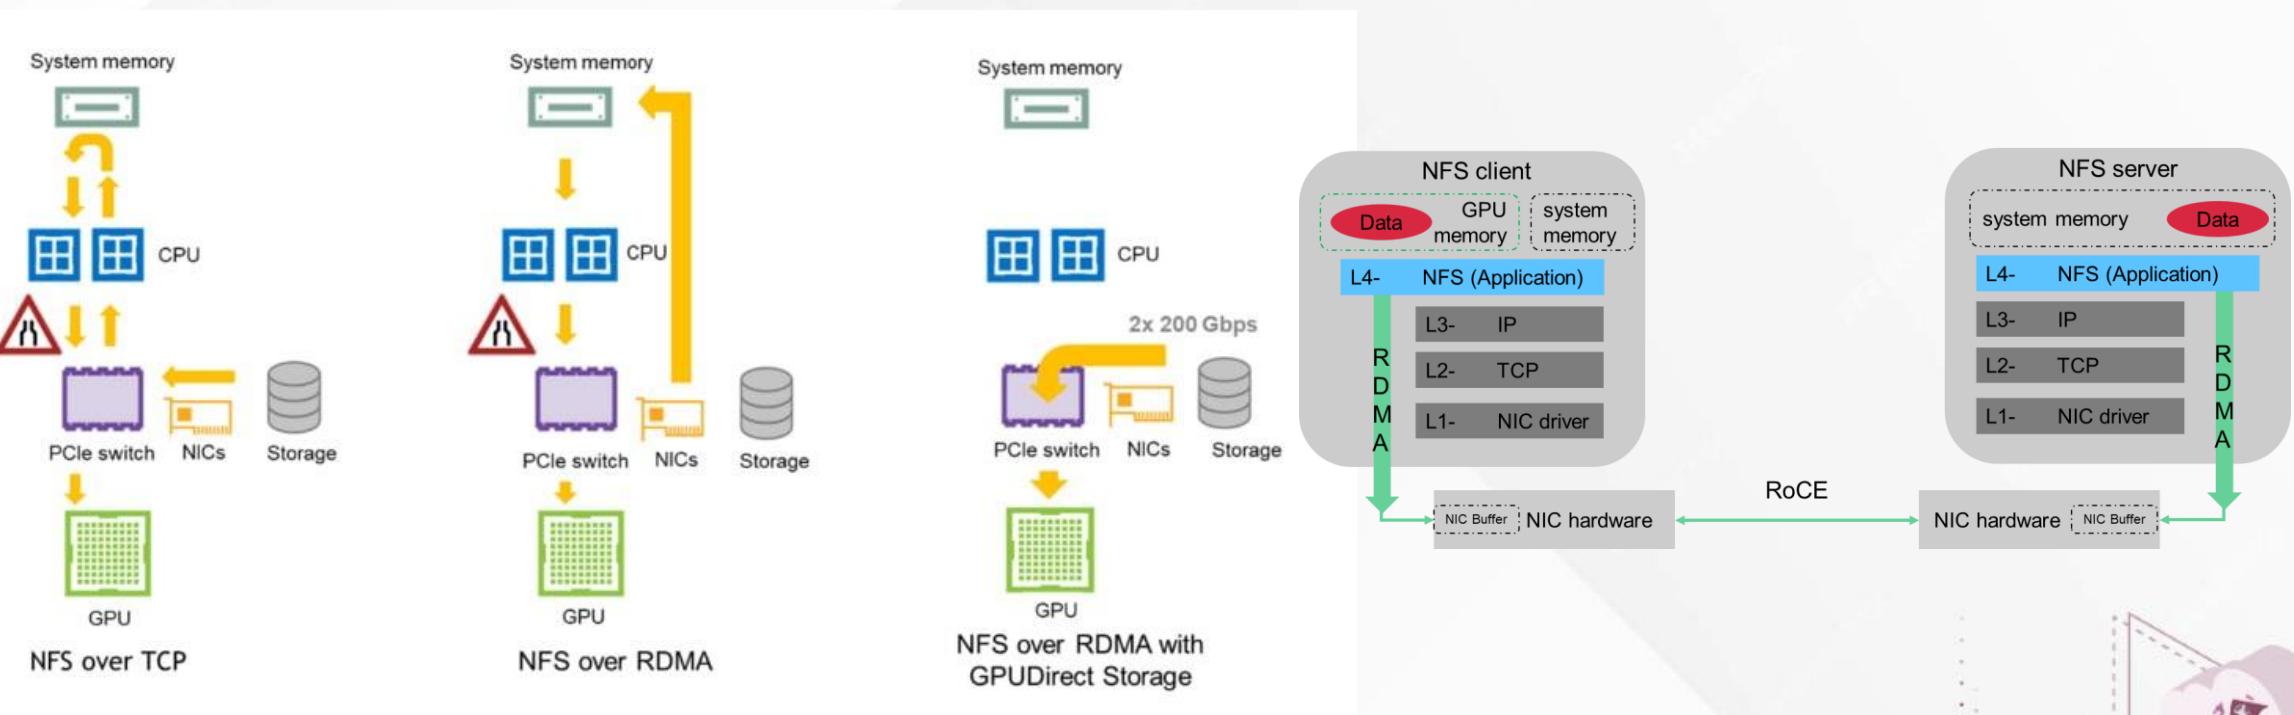


- 自动分层
- Zero-touch管理
- 更低的拥有成本

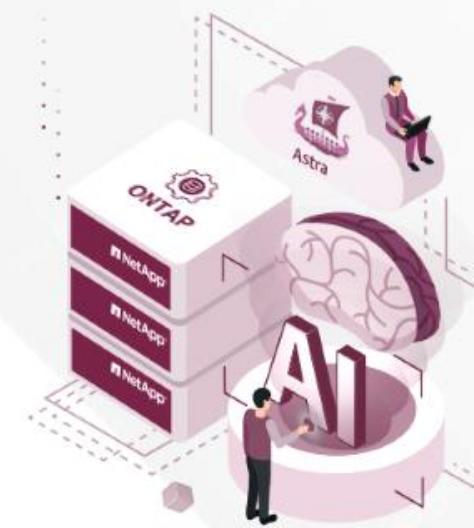


全面支持NVIDIA GPUDirect Storage – GDS显卡直接访问存储

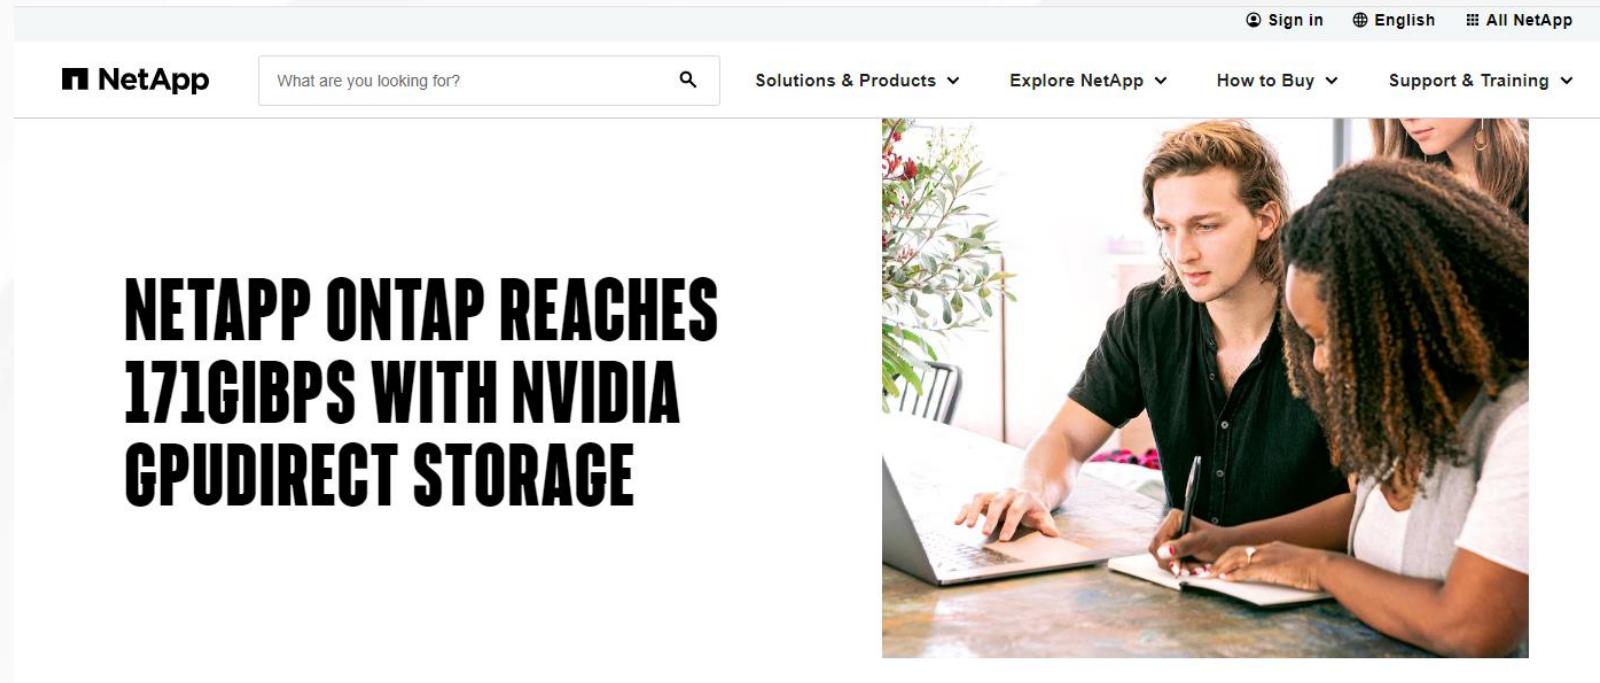
解放CPU转发压力，大幅降低GPU和存储之间的访问延时



GDS – GPU Direct Storage直接存储技术，使本地或远程存储和 GPU 内存之间实现了直接数据路径。GPU 直接 RDMA 和 GPU 直接存储器，避免了通过 CPU 内存中缓冲区的额外拷贝，并使 NIC 或存储器的直接内存访问（DMA）引擎，能够在直接路径上将数据移入或移出 GPU 内存，因此，所有这些都不会给 CPU 或 GPU 带来负担。



ONTAP+GDS实现卓越的性能表现



The screenshot shows the NetApp website's header with navigation links for Sign in, English, and All NetApp. Below the header is a search bar with placeholder text "What are you looking for?". The main content area features a large image of two people working together at a desk, one using a laptop and the other writing in a notebook. To the left of the image is a bold headline: "NETAPP ONTAP REACHES 171GIBPS WITH NVIDIA GPUDIRECT STORAGE". Below the headline is a table of contents and a blog post by David Arnette.

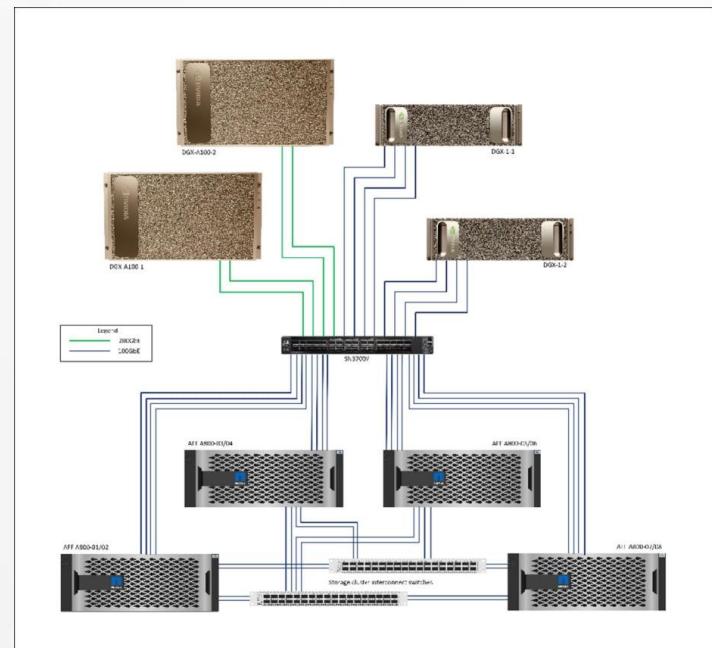
NETAPP ONTAP REACHES 171GIBPS WITH NVIDIA GPUDIRECT STORAGE

Table Of Contents

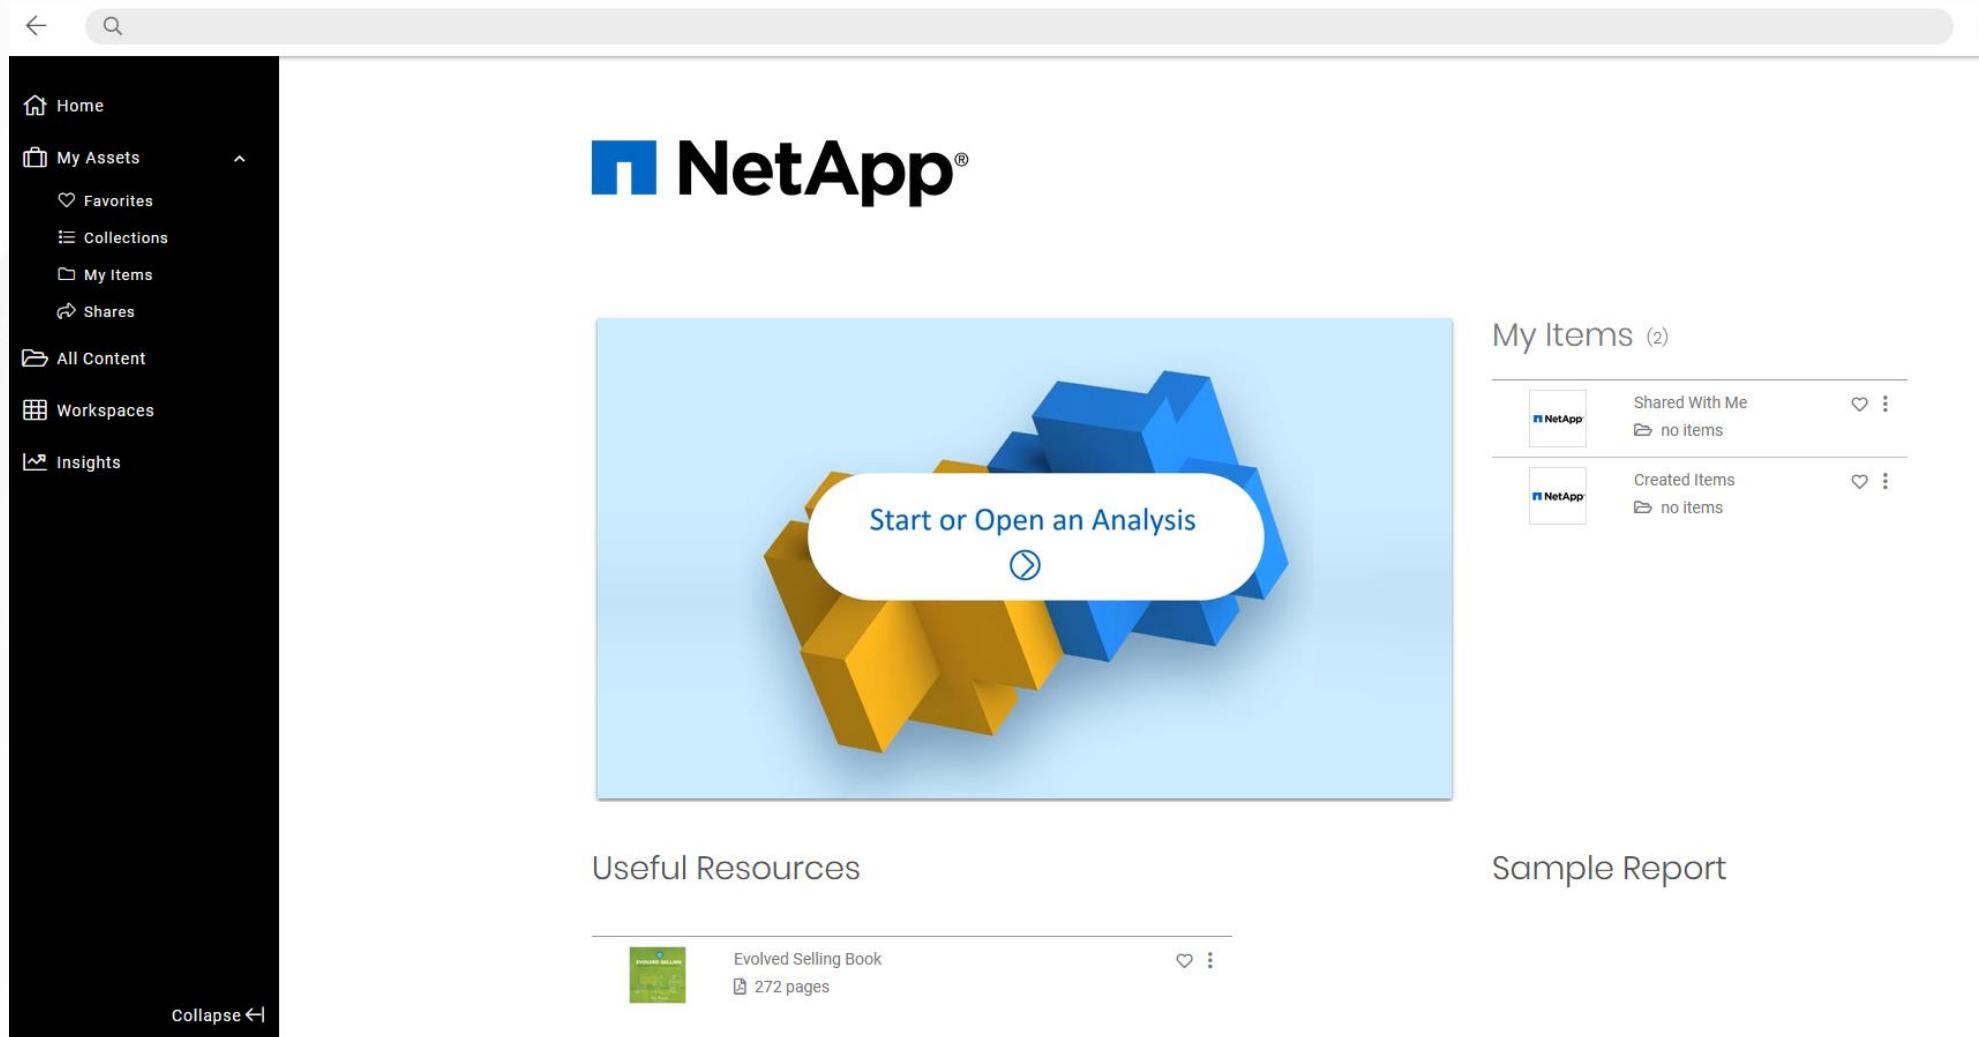
- A reminder about how NVIDIA GDS works
- How ONTAP and NFS improvements help GDS optimize resources for AI/ML
- Testing to prove optimized performance

David Arnette
April 28, 2023

A year ago, I wrote a blog post about [NetApp® ONTAP® support for NVIDIA GPUDirect Storage™ \(GDS\)](#) technology at a time when GDS and the surrounding ecosystem were still very new. As with many things that are new, there's always a lot of room for improvement, and our initial release of NFS over RDMA was no exception. I'm happy to say that, as



NetApp AI TCO评估工具



The screenshot shows the user interface of the NetApp AI TCO assessment tool. On the left is a dark sidebar with navigation links: Home, My Assets (Favorites, Collections, My Items, Shares), All Content, Workspaces, and Insights. A 'Collapse' button is at the bottom of the sidebar. The main area features the NetApp logo and a central graphic of stacked 3D bars (blue and yellow) with a call-to-action bubble that says 'Start or Open an Analysis'. Below this is a section titled 'Useful Resources' with a link to the 'Evolved Selling Book' (272 pages). To the right is a 'My Items' section showing two items: 'Shared With Me' (no items) and 'Created Items' (no items). At the bottom right is a decorative graphic illustrating AI integration with NetApp ONTAP storage.

Home

My Assets

- Favorites
- Collections
- My Items
- Shares

All Content

Workspaces

Insights

Collapse

NetApp®

Start or Open an Analysis

Useful Resources

Evolved Selling Book

272 pages

My Items (2)

Shared With Me

no items

Created Items

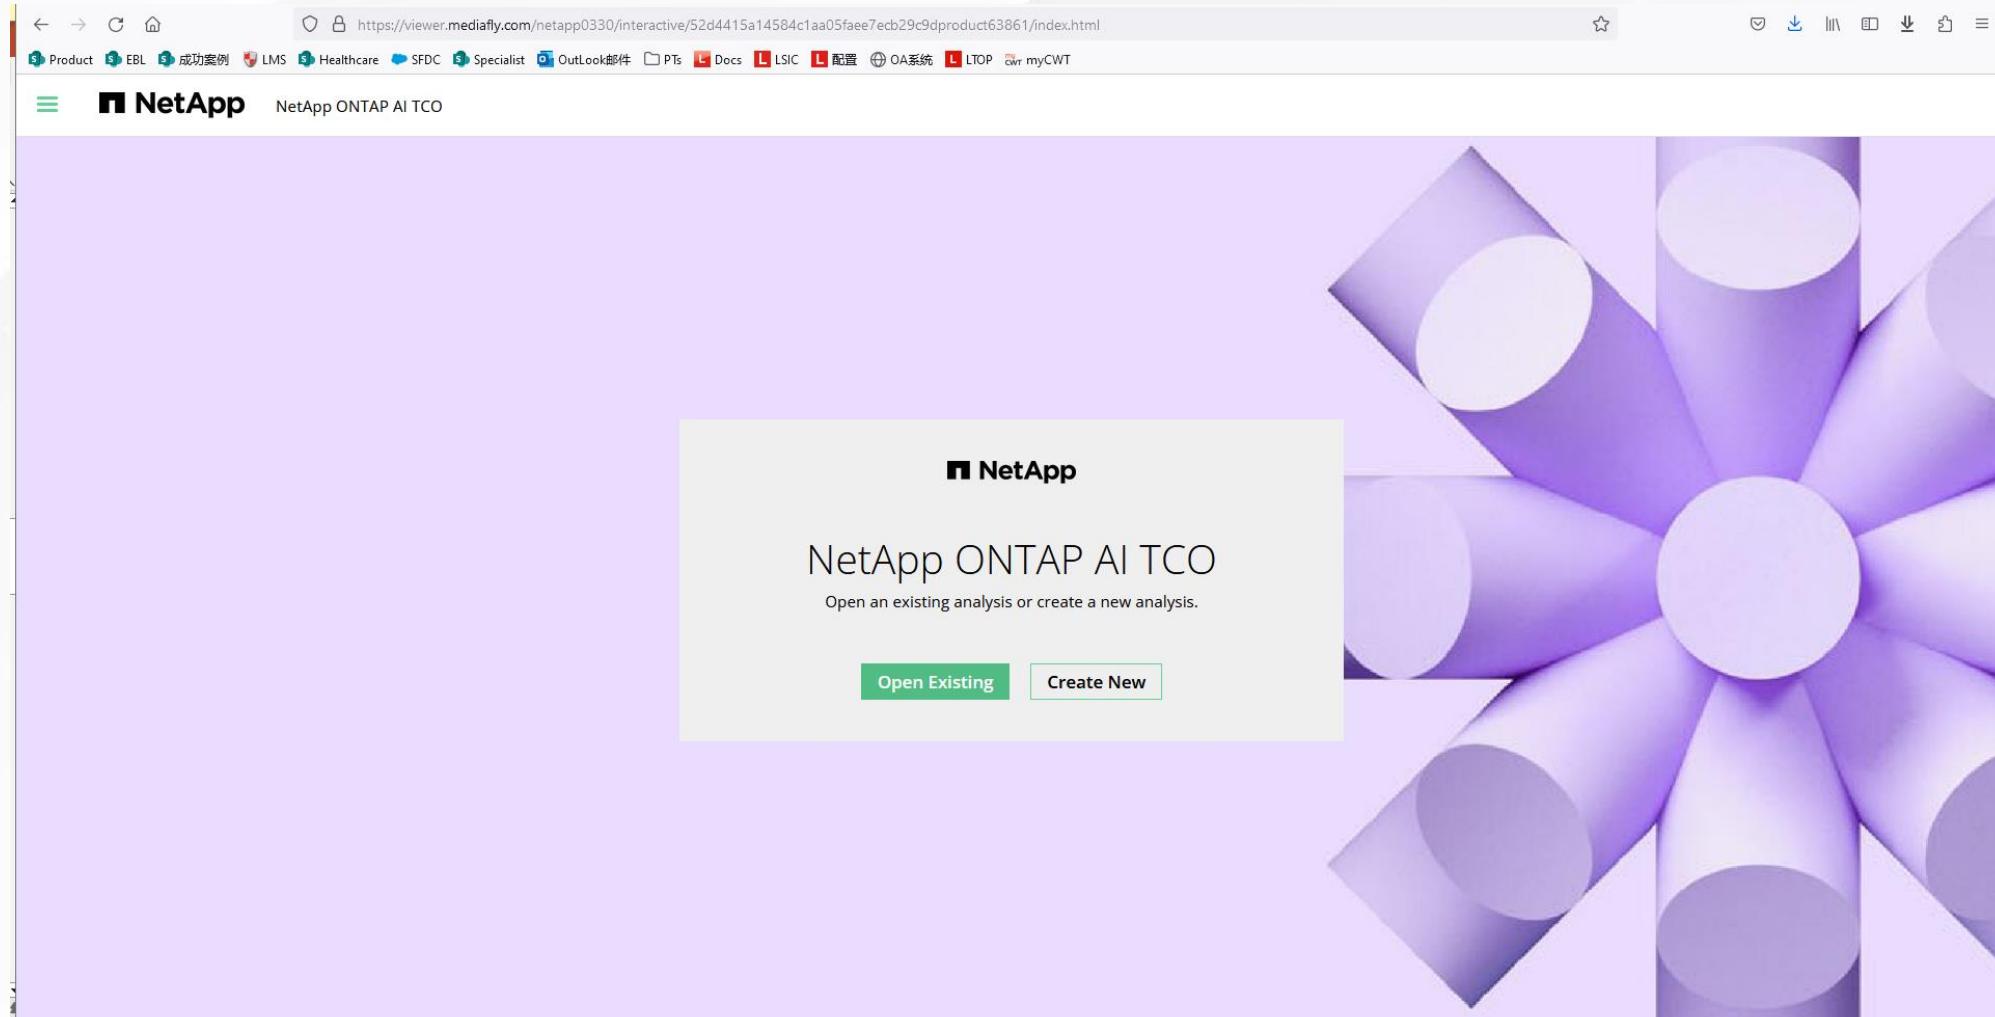
no items

ONTAP

Astra

AI

NetApp AI TCO评估工具



The screenshot shows a web browser displaying the NetApp ONTAP AI TCO tool. The URL in the address bar is <https://viewer.mediafly.com/netapp0330/interactive/52d4415a14584c1aa05faee7ecb29c9dproduct63861/index.html>. The page features a large, abstract purple geometric background. In the center, there is a white rectangular box containing the NetApp logo and the text "NetApp ONTAP AI TCO". Below this, a smaller text says "Open an existing analysis or create a new analysis." with two buttons: "Open Existing" and "Create New". To the right of the main content area, there is a decorative graphic of several overlapping purple cylinders.



The illustration at the bottom right depicts a stack of storage units labeled "ONTAP" and "NetApp". To the right of the stack, a person is standing next to a large brain icon, with the letters "AI" written vertically nearby. Above the stack, another person is working on a laptop inside a cloud-like shape labeled "Astra".

2023 Lenovo NetApp. All rights reserved.

NetApp AI TCO评估工具

NetApp ONTAP AI TCO Analysis Profile Results Next Steps

Analysis Results

Based on your profile, research on costs and business impacts, and your specific assumptions, a comparison of the total cost of ownership and incremental business value potential of NetApp ONTAP API versus cloud.

Total Cost of Ownership (TCO)

Cloud

	Monthly	1 Year	3 Years
Cloud Computing Explore	\$1,384,430	\$16,613,165	\$49,839,494
Cloud Storage Explore	\$13,850,000	\$166,200,000	\$498,600,000
Cloud Administration Explore	\$13,411	\$160,938	\$502,383
Total Cloud Cost	\$15,247,842	\$182,974,102	\$548,941,877

On-Prem

	Monthly	1 Year	3 Years
On-Prem Compute / Network Hardware Explore	\$454,984	\$5,459,808	\$16,379,423
On-Prem Storage Infrastructure Explore	\$2,023,200	\$24,278,400	\$72,835,200
On-Prem Operations Explore	\$144,461	\$1,733,532	\$5,343,502
Total On-Prem Cost	\$2,622,645	\$31,471,740	\$94,558,125

Savings of On-Prem

[Review Assumptions](#)

Select Review Assumptions and click on each TCO and business value element to review the metrics used and calculations. Every element can be reviewed and personalized for a precise analysis for your specific opportunity.

By Month By Year Over 3 Years

\$182,974,102

Total Cloud Costs

\$31,471,740

Total On-Prem Costs (Leased)

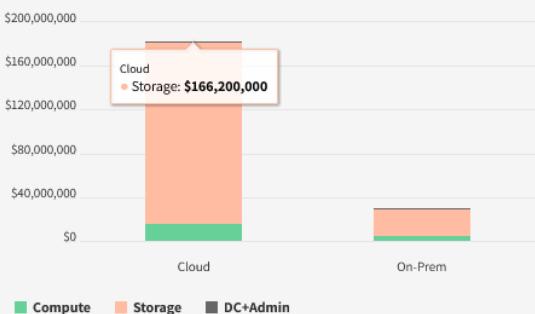
\$151,502,363

Savings of On-Prem

\$153,470,184

Total Savings and Business Value Benefits

TCO - Yearly Cloud vs. On-Prem Cost



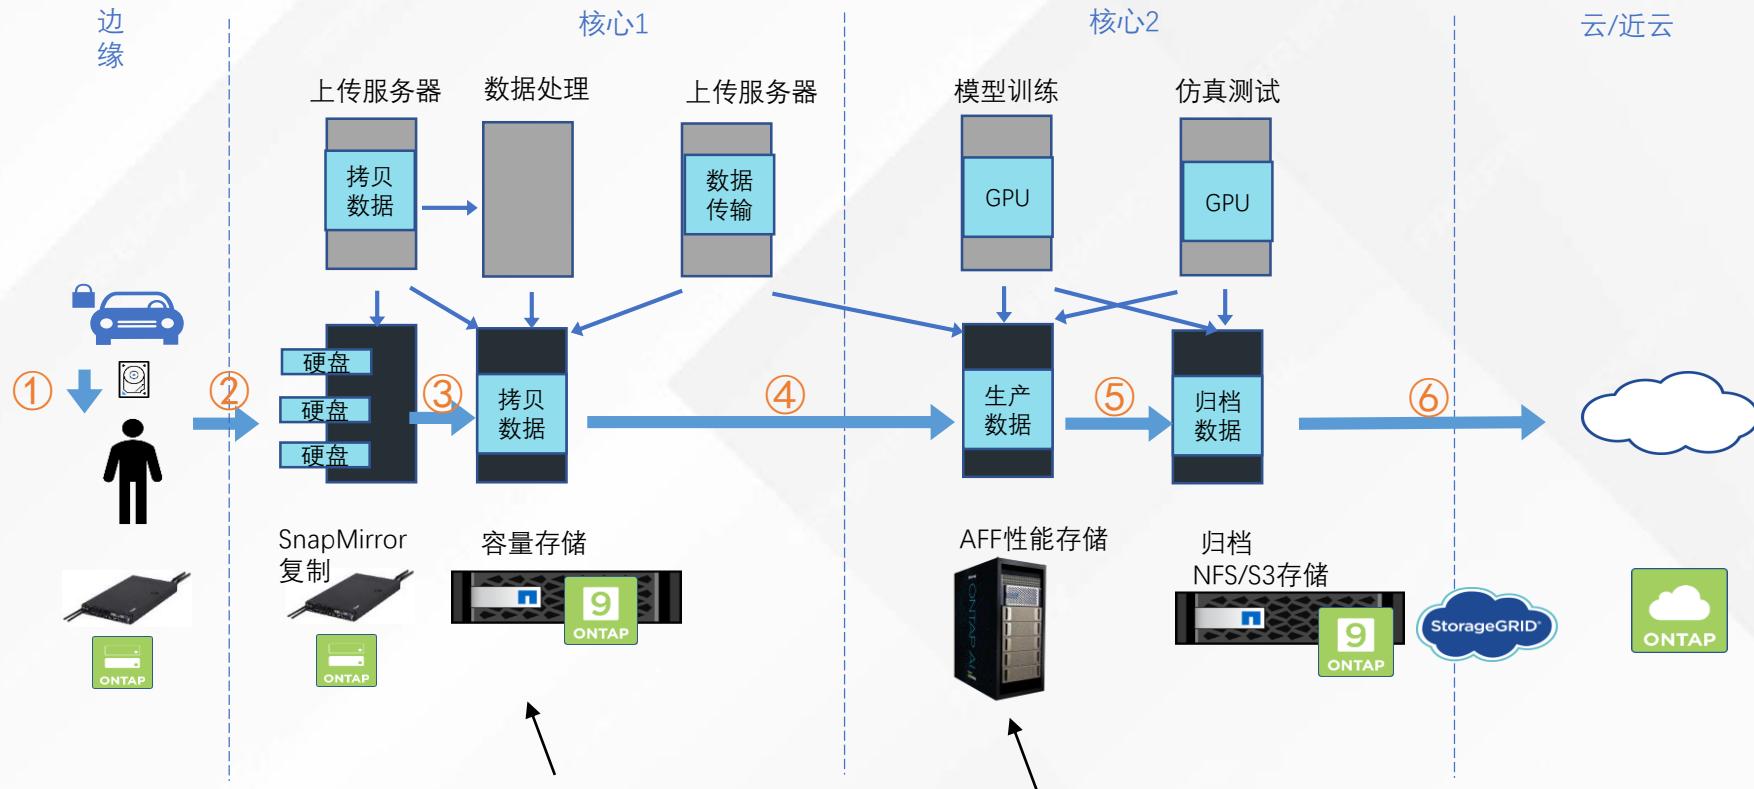
Incremental Business Value Potential of NetApp ONTAP AI vs. Cloud

\$1,500,000

Next Steps

[Profile](#)

国内头部基于激光雷达自动驾驶用户



现有FAS8700，用于存放原始数据和清洗后数据

本次项目推荐A900全闪存储，极致性能满足仿真训练需求
参考架构：Nvidia basepod
认证架构ONTAP AI

客户仿真训练存储平台性能要求高带宽低延时，**带宽需求最终确认为19GB/s；我们方案可以提供33GB/s**，满足现在及未来需求。

我们现有一台FAS8700，从整体数据流数据流转方便性的角度出发，我们推荐集中式全闪存储架构（Volume Move），并且该架构符合ONTAP AI 的官方认证。

同时，GPU 服务器之间采用高速 RoCE 网络（100Gb）。

方案设计同时帮客户考虑未来边端、云端的数据打通路径！

<https://www.nvidia.cn/data-center/dgx-basepod/>



总结：为什么选用NetApp ONTAP-AI？

79PB

> 25x Raw storage capacity
versus competition

6x~12x

Reduction in time to develop and
deploy new AI services

#1 NFS

DGX systems usage
NFS by default

Edge solutions

NetApp® ONTAP®
Select Lenovo

Cold data tiering

Unified data lake for AI with
NetApp StorageGRID®

Storage efficiency

Deduplication
Compression
Compaction

50%

reduction in operating costs

Cloud connected

AWS
Microsoft Azure
Google Cloud

16x

storage capacity
reduction

Leader validated architectures

FlexPod® 10 years and going



目录

Contents



- NVIDIA DGX SuperPOD
with NetApp

经过认证的SuperPOD参考架构

Storage Building Block Overview: BeeGFS + EF600

NVIDIA DGX A100 Systems

20-140 systems for SuperPOD
Each system contains 8x A100 GPUs



NVIDIA Quantum Switches

400 / 200Gb HDR InfiniBand



1 Storage Building Block Contains:

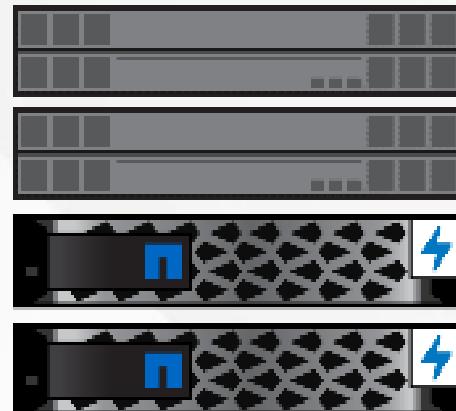
2x 2U x86 BeeGFS Servers (400Gb in Gen2)
2x 2U NetApp EF600 all flash storage



NetApp BeeGFS Pod Building Blocks

Per Building Block	Performance / Capacity
Metadata + Storage Option (Converged)	
Bandwidth R/W GB/s	66 / 22
IOPS R/W	1.3M / 900K
MLPerf Training img/sec	155K
Est. Max Usable Capacity	431 TB
Storage Option (Standalone)	
Bandwidth R/W GB/s	74 / 24
IOPs R/W	-
MLPerf Training img/sec	-
Est. Max Usable Capacity	538.6 TB

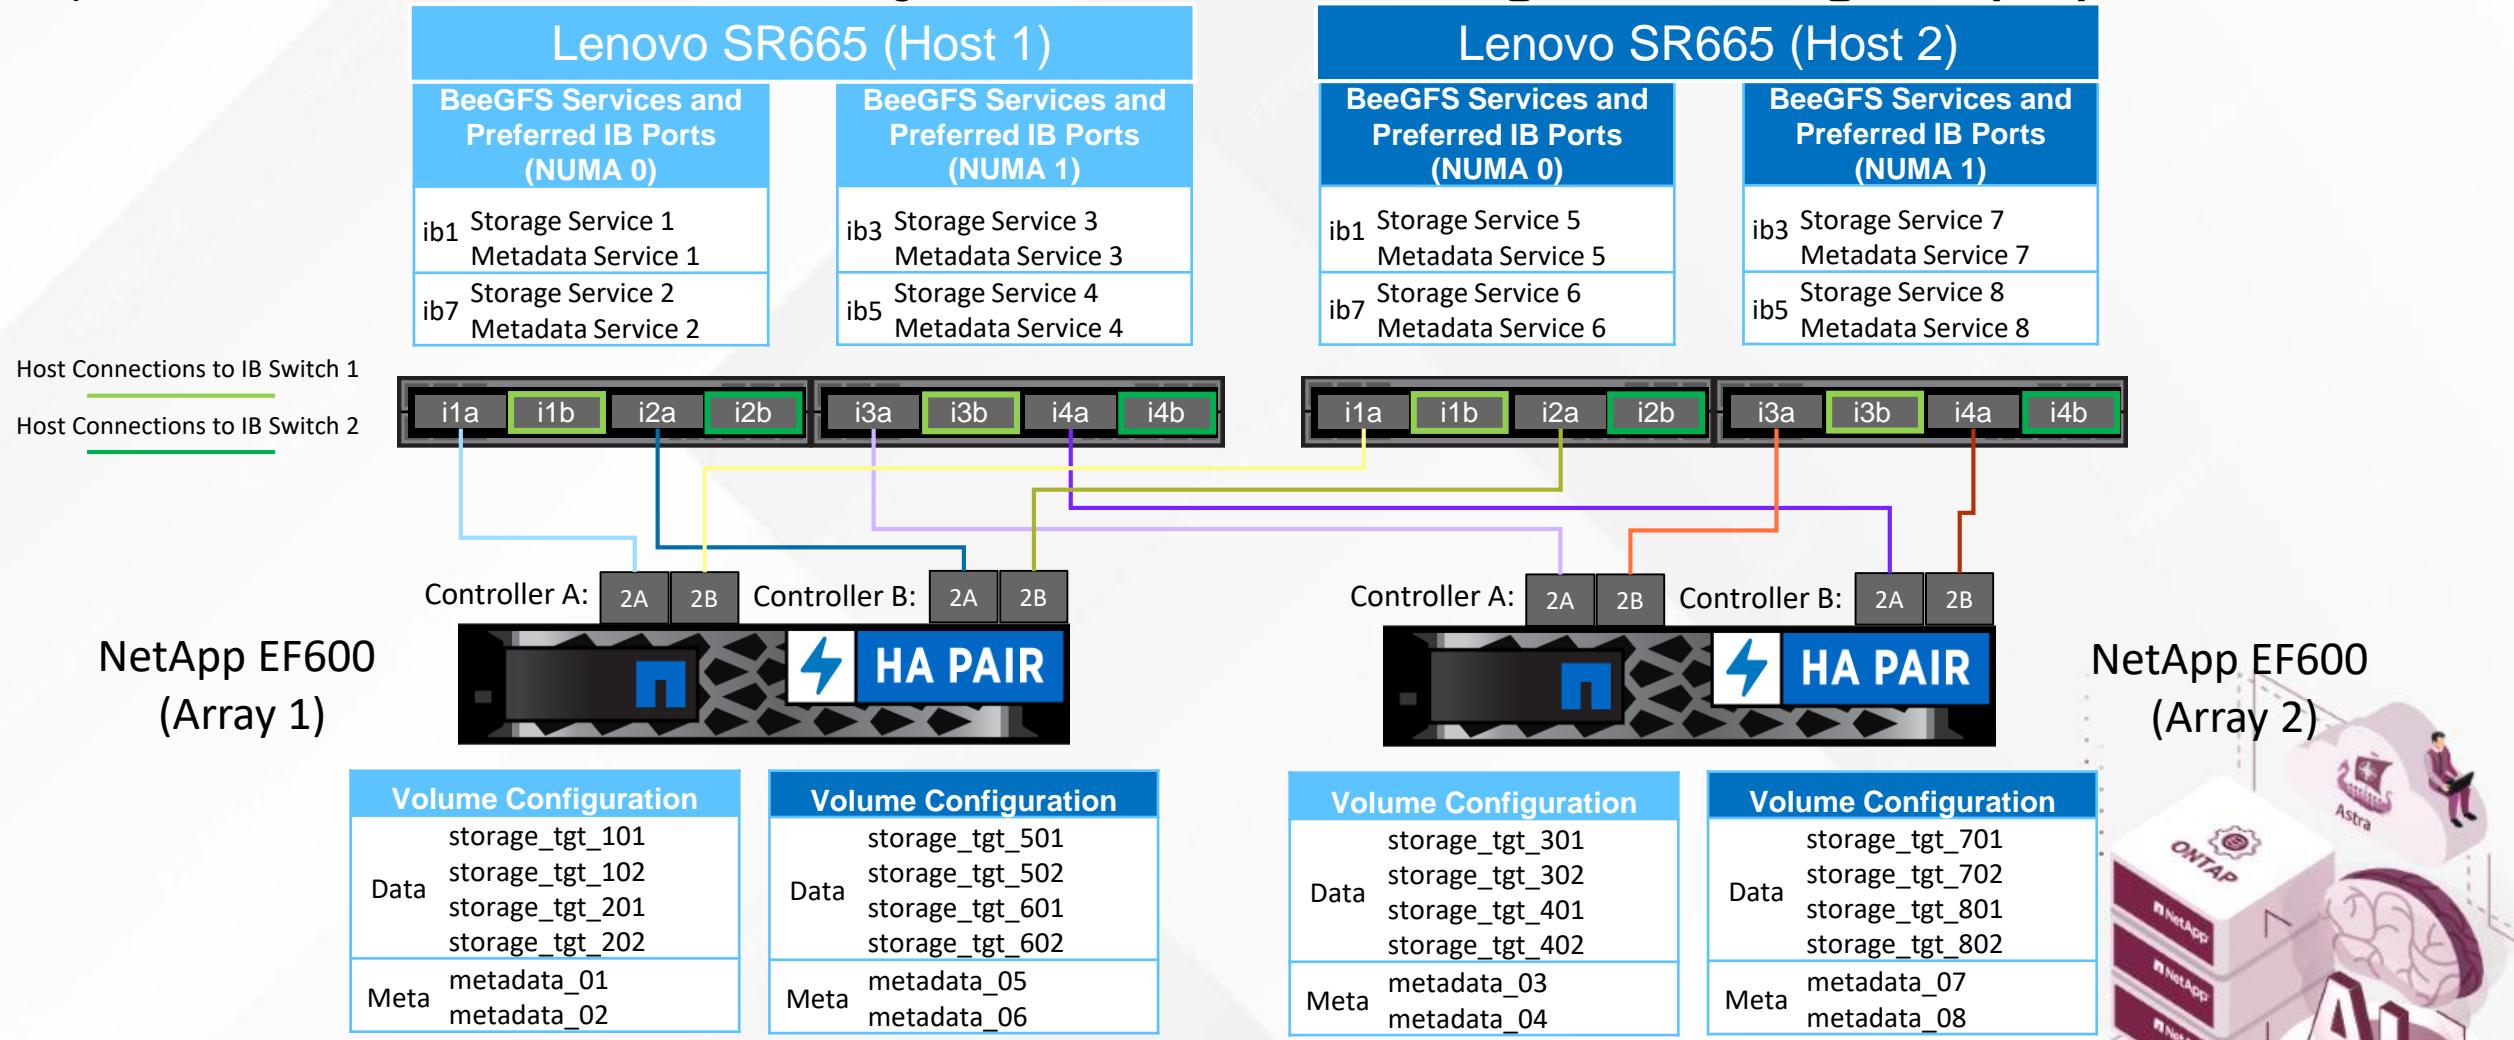
Storage Building Block:
2-socket x86 servers
EF600 storage arrays
HDR (200Gb) IB
8U rack density



NetApp BeeGFS Pod Building Blocks(Gen2)

Replicated for each EF600/SR665 building block

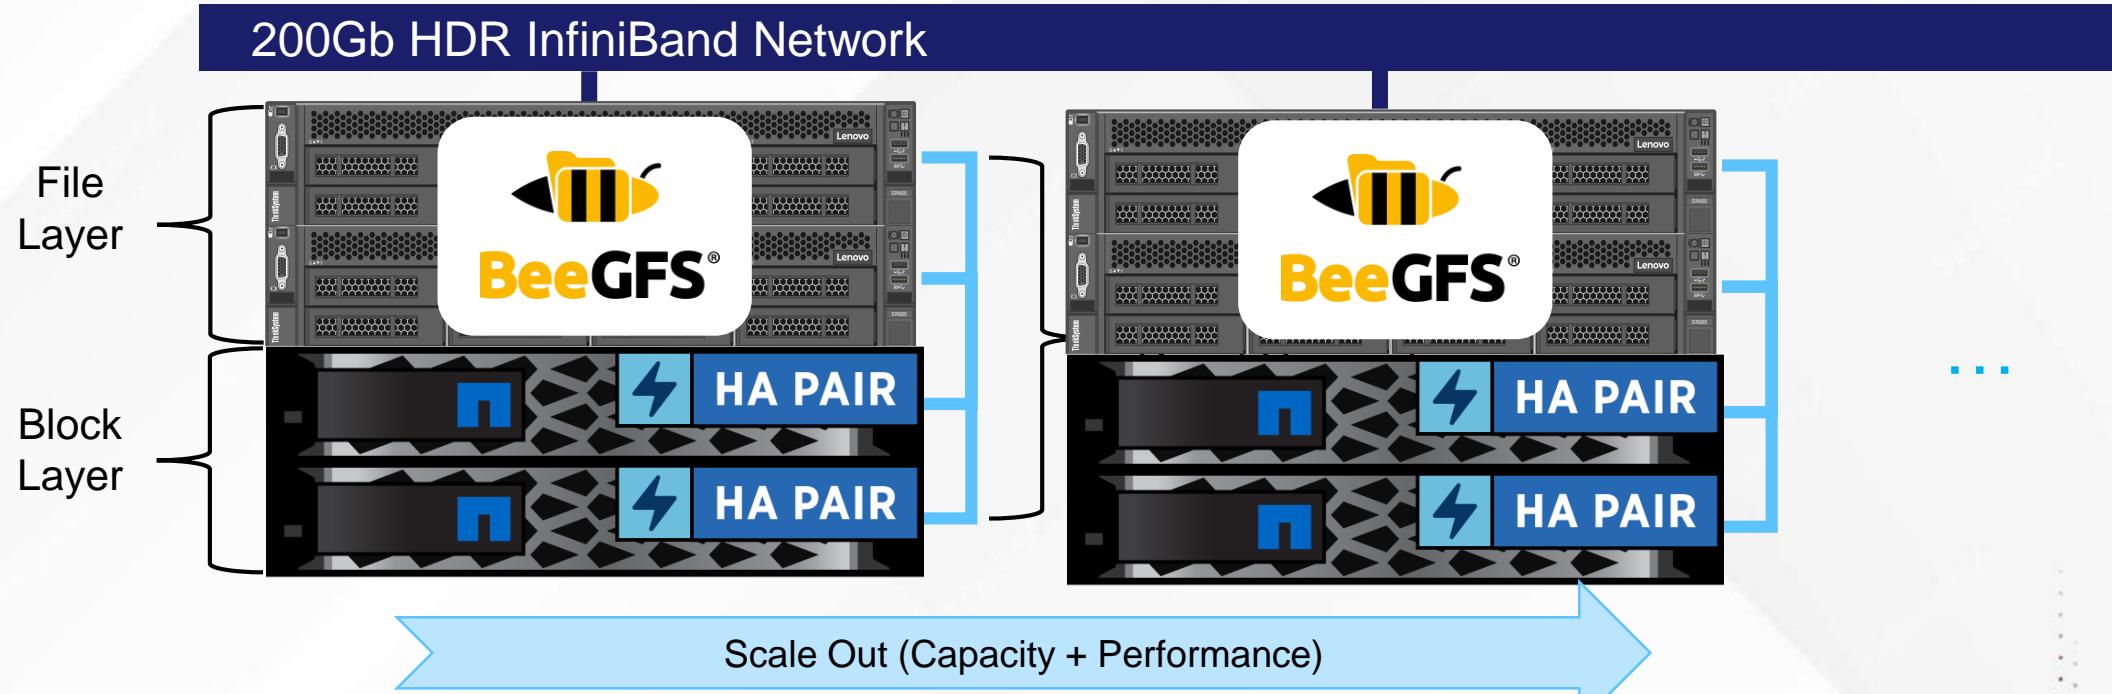
Gen2 Building Block Diagram (8U)



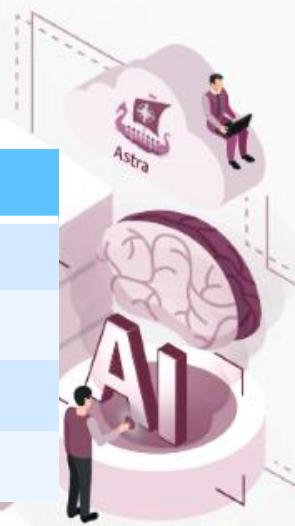
Note: All volumes are mapped to both hosts allowing for optional creation of a shared-disk high availability cluster.

NetApp BeeGFS Pod Building Blocks

高可用架构兼顾线性的性能和容量扩展



Performance / Capacity	1 Block	2 Blocks	3 Blocks	...
Bandwidth R/W GB/s	66 / 22	132 / 44	198 / 66	...
IOPS R/W	1.3M / 900K	2.6M / 1.8M	3.9M / 2.7M	...
MLPerf Training img/sec	155K	310K	465K	...
Est. Max Usable Capacity	431 TB	862 TB	1293 TB	...



EF600提供超低延时下的超高带宽

EF600 HPC 超算架构下 NVMe 极低延时下性能输出.

One EF600 with 24 SSDs Delivers

250 μ s
latency

2M
IOPS*

44GB/s
throughput

367TB
raw capacity

25
EF600
Systems

50M IOPS* 250us avg latency

1,100GB Delivering the Data Bandwidth



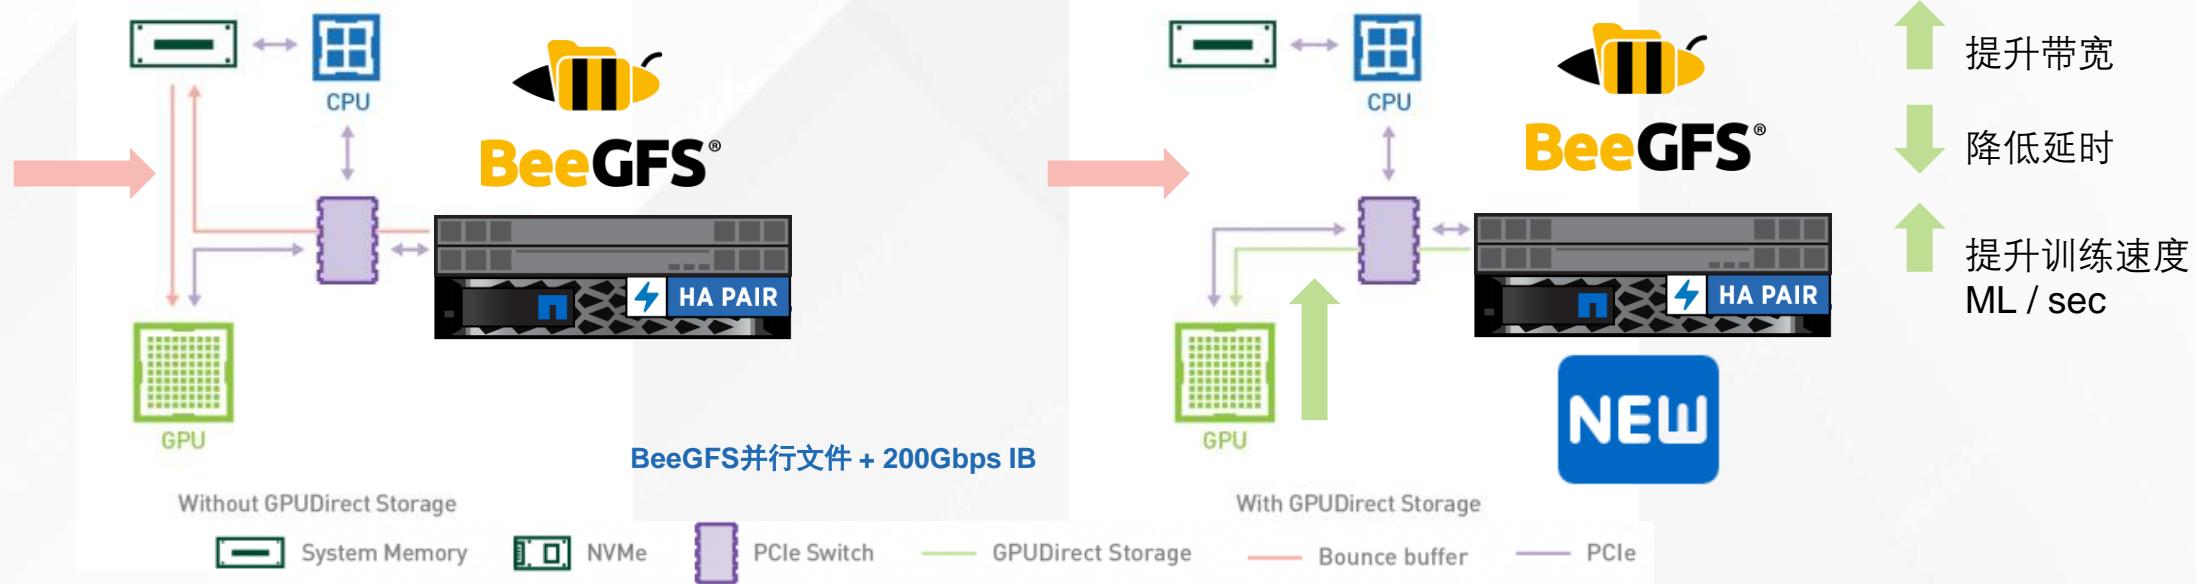
* 4K block random reads dual parity

2023 Lenovo NetApp. All rights reserved.



全面支持NVIDIA GPUDirect Storage – GDS显卡直接访问存储

解放CPU转发压力，大幅降低GPU和存储之间的访问延时



GPUs vs Throughput

— GDS (NVMe ->GPU) — mmap (Read from Active page cache) + cudaMemcpy + cudaMalloc Memory
 — mmap (Read from Disk) + cudaMemcpy + cudaMalloc Memory — mmap + cudaMallocManaged (Stock Bits)

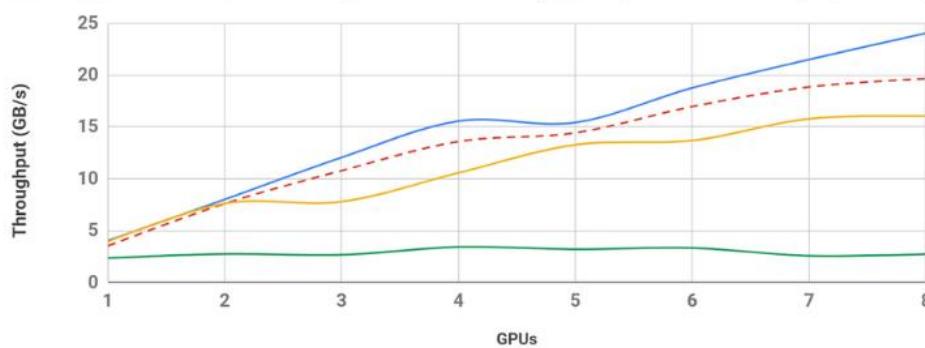


Figure 4. GPUs vs. throughput

Average Worker Latency

— GPUDirect Storage — cuDF optimized - bounce buffer but no faulting
 — Original cuDF with faults, bounce buffer, unpinned

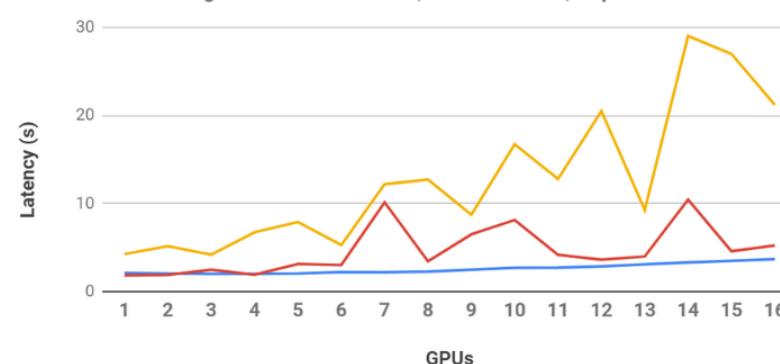


Figure 5. Latency comparison for cuDF read_csv

Comparison of Transfer Methods: Bandwidth/CPU Utilization

— GDS — CPU bounce buffer — Page cache

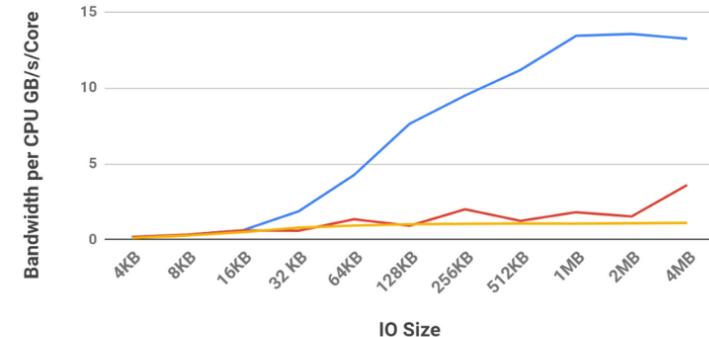
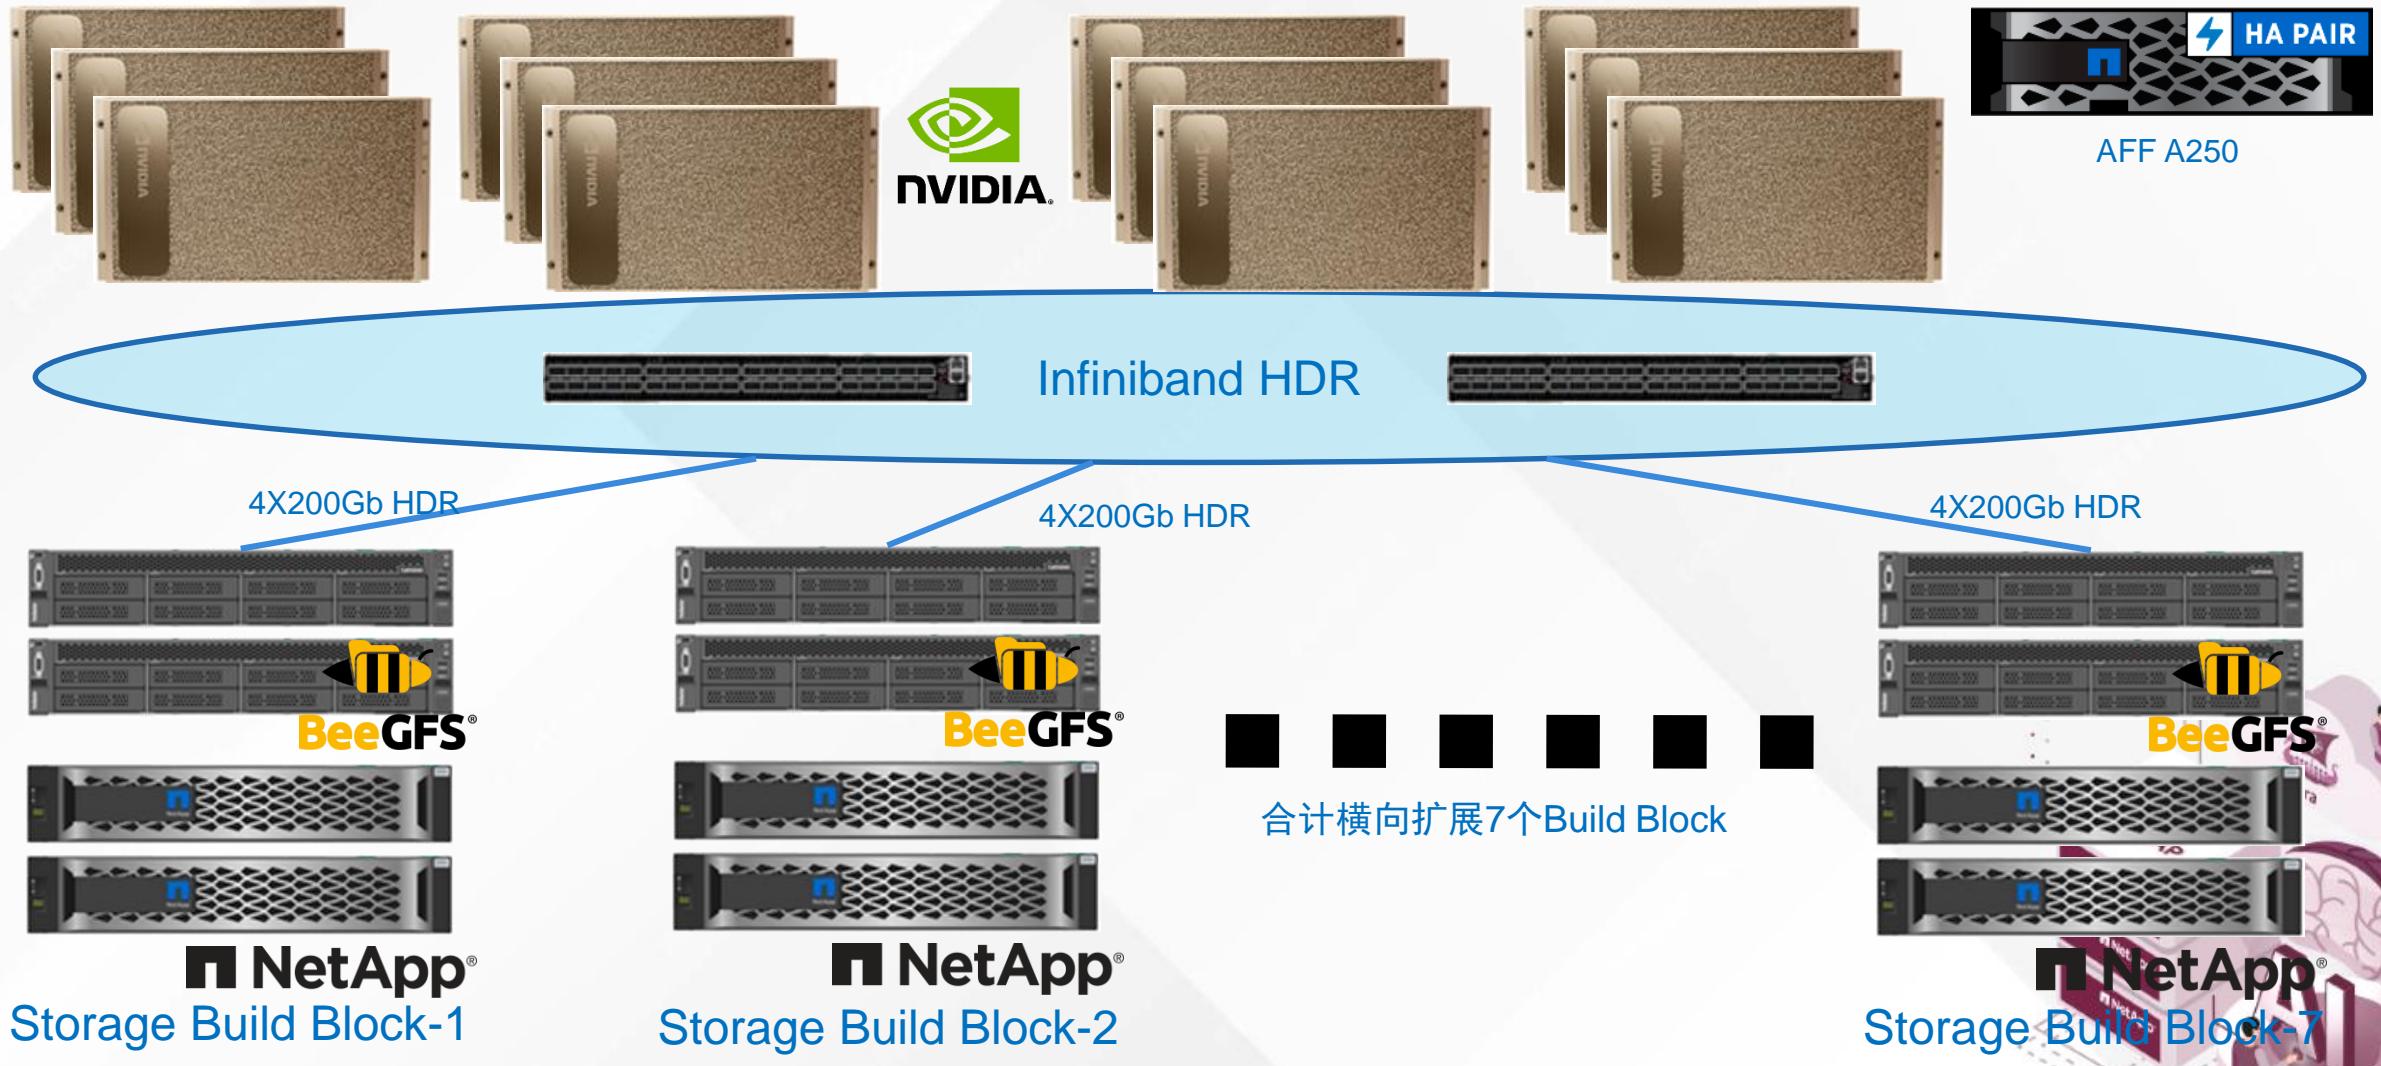


Figure 7. Bandwidth divided by the fractional utilization of the CPU core. GPUDirect Storage incurs far less burden on the CPU at larger transfer sizes.

国内首个SupperPOD成功案例



谢谢！

联想凌拓官方联络方式

官方网站: www.lenovonetapp.com

服务热线: 400-828-3001 (呼叫中心)

400-116-0099 (销售热线)

官方社交平台账号:



联想凌拓
官方微信



Bilibili联想凌拓
空中沙龙



联想凌拓
渠道微信

