

STORAGE SOLUTIONS FOR AI WORKLOADS




Dr. Didier Gava

EMEA & LATAM Senior Solution architect

September 2024

Every enterprise is different, but four challenges appear across the board




#1

Complexity


- Uncertainty and difficulty of DIY infrastructure
- Fast evolving software stack

Scaling AI projects

- Scaling across multiple AI projects
- Handling multiple data sources especially across clouds



#2




#3

Data Integration Trials


- AI Environment troubleshooting, ROI, visibility-monitoring
- Support for multivendor technology stack

Data Security

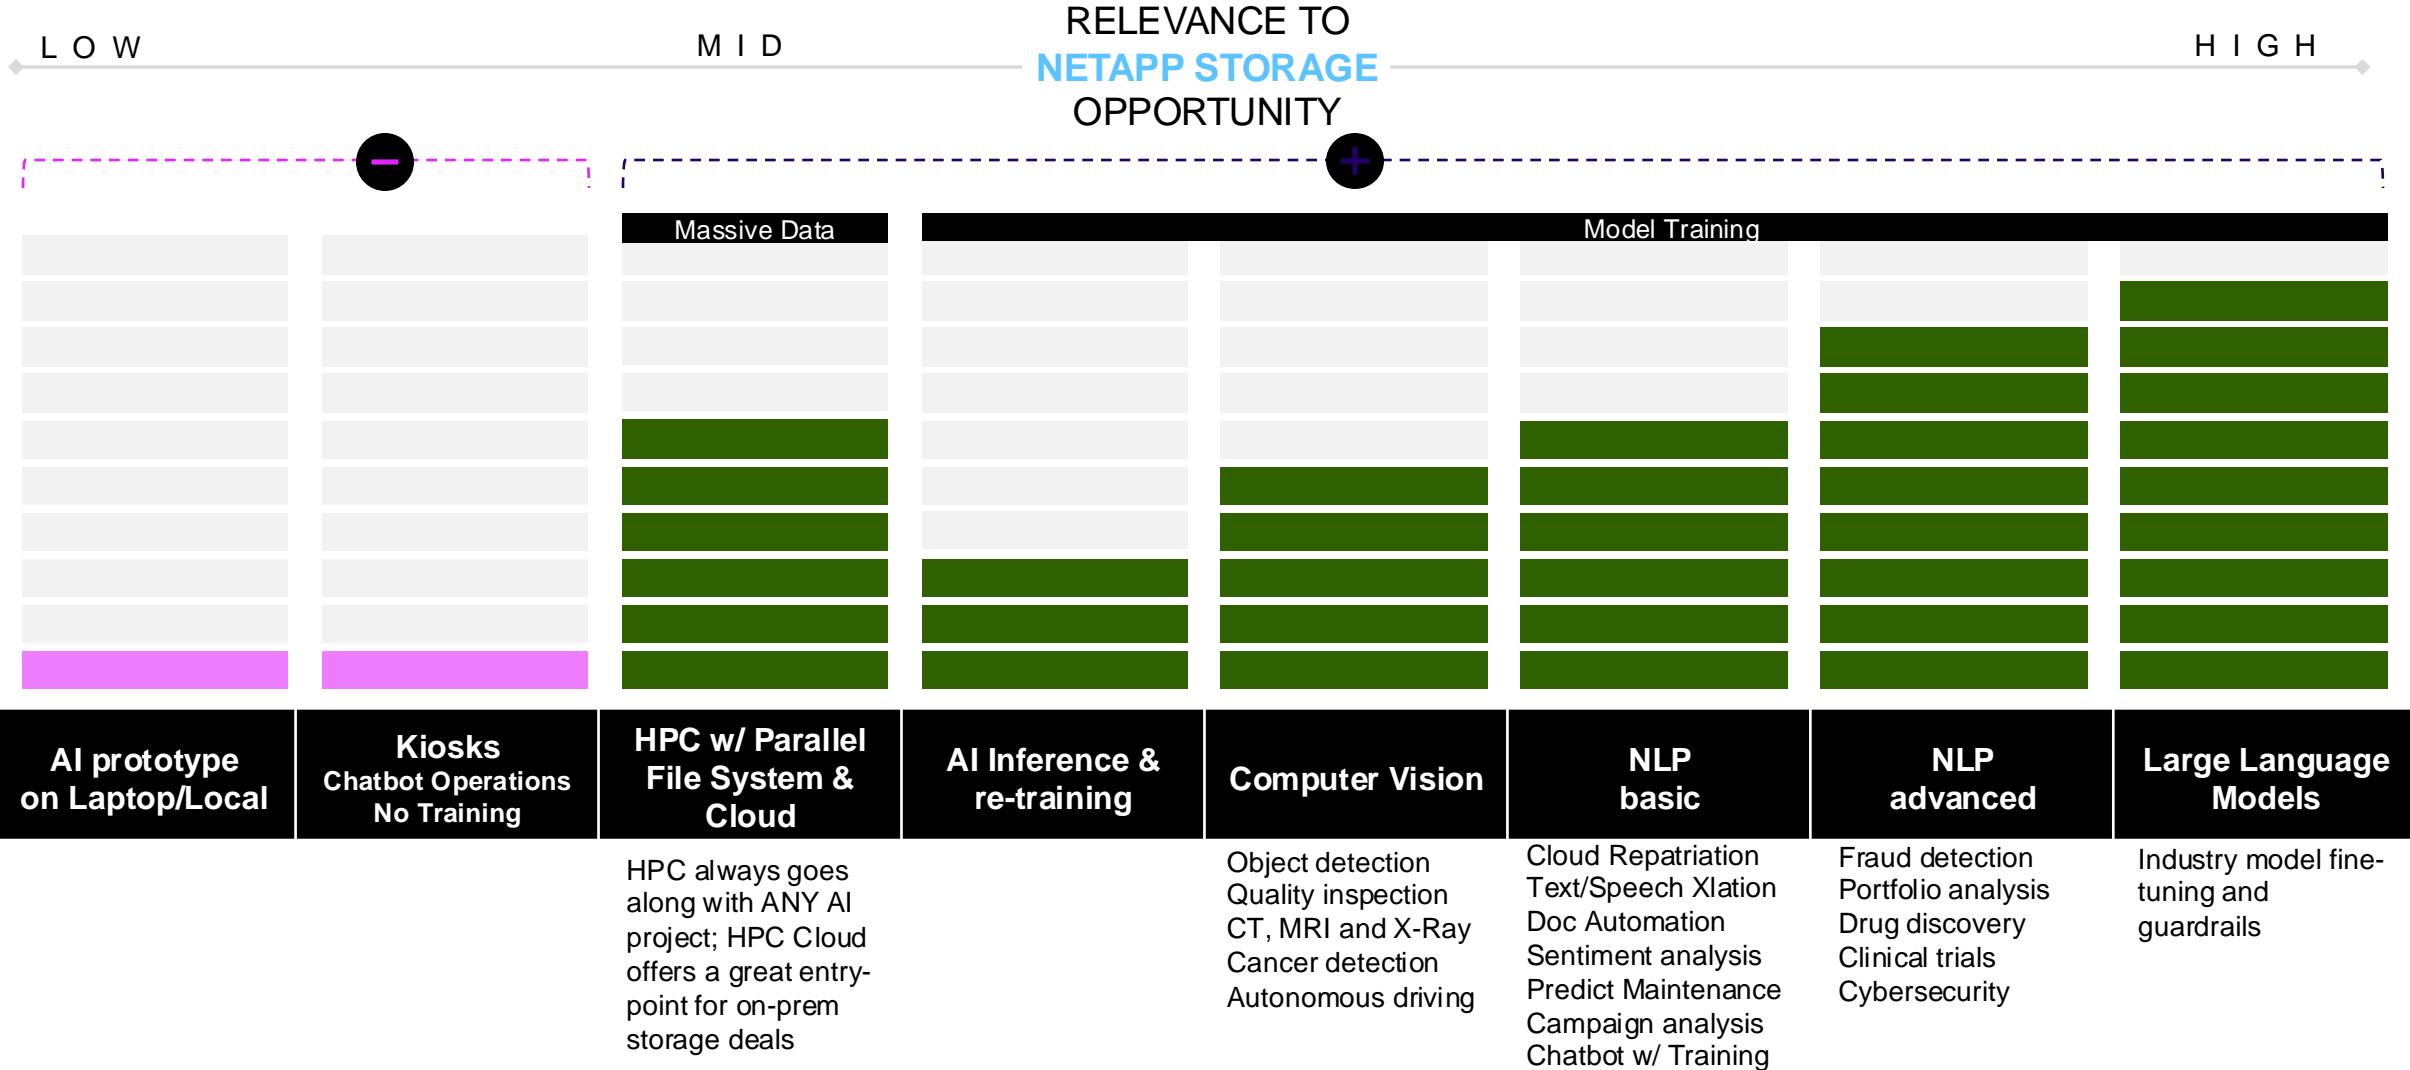
- Outdated policies and frameworks
- Tools that find attacks rather than prevent them



#4

 **FASTEST GROWING CY2022**


Key Target AI Use Cases



NetApp's Portfolio spans the data pipeline

EDGE


Ingest



- Data collection
- Edge-level AI

CORE


Data Prep Training Cluster Deployment



- Aggregation
- Normalization
- Exploration
- Training
- Deployment
- Model serving

CLOUD




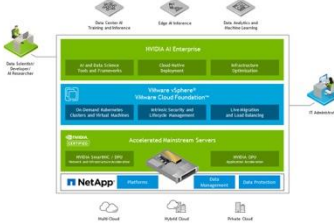
Analysis / Tiering



- Cloud AI (GPU instances)
- Cloud tiering / bursting

NetApp and NVIDIA portfolio

Deploy the industry's leading AI infrastructure, your way

NetApp AI	DGX Cloud NetApp / Equinix US	NVIDIA DGX SuperPOD with NetApp	NVIDIA AI Enterprise with NetApp
	 nvidia NetApp		
Reference architecture Certified for success with maximum flexibility <small>(Others: FlexPOD AI, E-Series AI, ONTAP with Lenovo, ONTAP with Fujitsu)</small>	Integrated solution Easy to procure and installed onsite with premium support	Fully managed Subscription-based service for AI development infrastructure	AI for any organization Mainstream AI for a wide variety of enterprise workloads
Turnkey solution for rapid deployment		Massive scale supercomputing	AI for enterprise operations
Field-proven design formula			

Storage “Building Blocks” for SuperPOD

BeeGFS on NetApp Verified Architecture

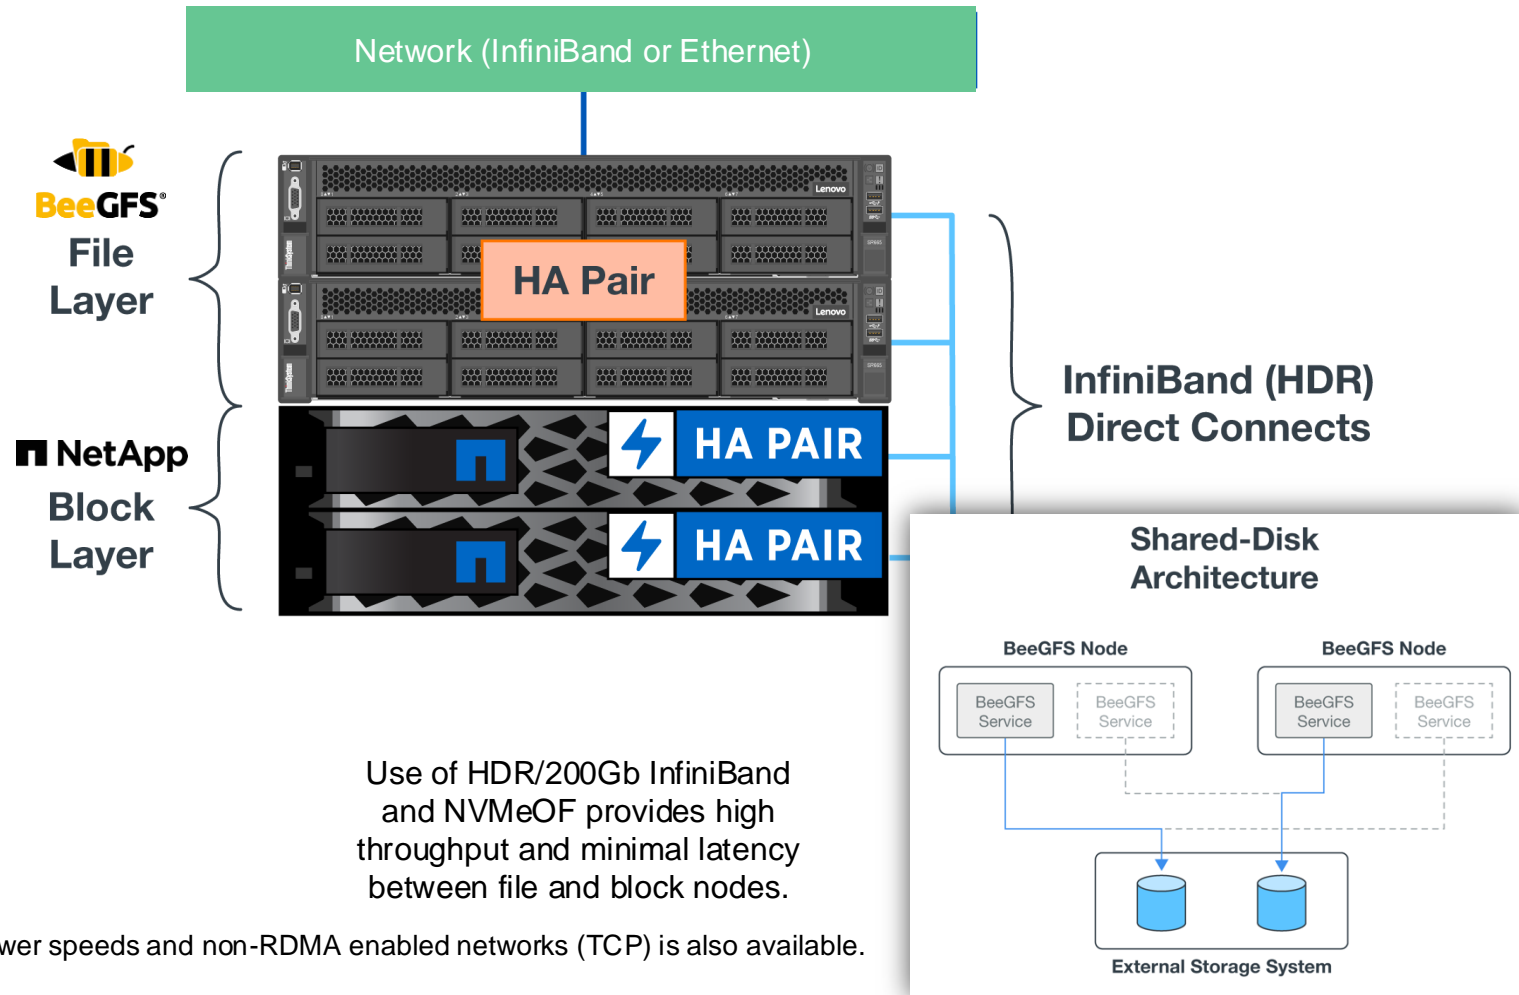
Highlights

Building blocks provide high availability in a two-tier design that separates fault domains for file and block layers.

BeeGFS services fail over between two file nodes directly connected to two block nodes.

BeeGFS metadata and data protected by NetApp E-Series technology with up to 99.9999% reliability.

Interfaces for the client network support InfiniBand (EDR/HDR) and RDMA over Converged Ethernet (RoCE -100Gb/200Gb).*



*Support for slower speeds and non-RDMA enabled networks (TCP) is also available.

Typical Deployment Steps

Step-by-step details or automate end-to-end deployment using Ansible.

E-Series/Server Preparations:

- Create E-Series volumes and map to servers.
- Setup frontend/backend networking including server/E-Series connections (e.g., NVMe-IB).
- Format and mount E-Series block devices used as BeeGFS metadata/storage targets.

BeeGFS Installation:

- Add the BeeGFS package repository file.
- Install services using a package manager.
- Configure BeeGFS services starting with the management service and ending with clients:
- Tune additional parameters `/etc/beegfs` (optional).
- Start/enable BeeGFS services with `systemd`.

Beegfs Installation

- Add the BeeGFS package repository file.
 - x86-64 packages provided for RHEL 8.x. and 9.3, SLES 15.1, Debian 10- 12 and Ubuntu 20.04 or 22.04 (or build BeeGFS from source).
Eg yum install beegfs-xx where xx are services
- Install services using a package manager.
 - yum install beegfs-mgmtmd
 - yum install beegfs-meta libbeegfs-ib
 - yum install beegfs-storage libbeegfs-ib
 - yum install beegfs-client beegfs-helperd beegfs-utils
- Configure BeeGFS services starting with the management service and ending with clients:
 - Setup scripts in /opt/beegfs/sbin/ to set management IP on other services and initialize storage/metadata targets.
 - /opt/beegfs/sbin/beegfs-setup-mgmtmd -p /data/beegfs/beegfs_mgmtmd
 - /opt/beegfs/sbin/beegfs-setup-meta -p /mnt/meta -s 2 -m node01
 - /opt/beegfs/sbin/beegfs-setup-storage -p /mnt/myraid1 -s 3 -i 301 -m node01
 - /opt/beegfs/sbin/beegfs-setup-storage -p /mnt/myraid2/ -s 3 -i 302
 - /opt/beegfs/sbin/beegfs-setup-client -m node01

- Tune additional parameters /etc/beegfs (optional).
 - If using RDMA update beegfs-client-autobuild.conf.
- Start/enable BeeGFS services with systemd.
 - systemctl start beegfs-mgmtmd
 - systemctl start beegfs-meta
 - systemctl start beegfs-storage
 - systemctl start beegfs-helperd
 - systemctl start beegfs-client
- Verification with system status

```
root@admin1-Express5800-R120a-2-N8100-1501E:~# systemctl status beegfs-storage
● beegfs-storage.service - BeeGFS Storage Server
   Loaded: loaded (/lib/systemd/system/beegfs-storage.service; enabled; vendor preset: enabled)
   Active: active (running) since Sat 2023-04-01 17:58:12 CEST; 1min 35s ago
     Docs: http://www.beegfs.com/content/documentation/
    Main PID: 85274 (beegfs-storage/)
      Tasks: 2 (limit: 14266)
     Memory: 1.0M
        CPU: 341ms
    CGroup: /system.slice/beegfs-storage.service
            └─85274 /opt/beegfs/sbin/beegfs-storage cfgFile=/etc/beegfs/beegfs-storage.conf runDaemonized=false

abr 01 17:58:12 admin1-Express5800-R120a-2-N8100-1501E systemd[1]: Started BeeGFS Storage Server.
root@admin1-Express5800-R120a-2-N8100-1501E:~# systemctl status beegfs-meta
● beegfs-meta.service - BeeGFS Metadata Server
   Loaded: loaded (/lib/systemd/system/beegfs-meta.service; enabled; vendor preset: enabled)
   Active: active (running) since Sat 2023-04-01 17:58:00 CEST; 1min 59s ago
     Docs: http://www.beegfs.com/content/documentation/
    Main PID: 85205 (beegfs-meta/Mai)
      Tasks: 2 (limit: 14266)
     Memory: 372.0M
        CPU: 3.714s
    CGroup: /system.slice/beegfs-meta.service
            └─85205 /opt/beegfs/sbin/beegfs-meta cfgFile=/etc/beegfs/beegfs-meta.conf runDaemonized=false

abr 01 17:58:00 admin1-Express5800-R120a-2-N8100-1501E systemd[1]: Started BeeGFS Metadata Server.
root@admin1-Express5800-R120a-2-N8100-1501E:~#
```

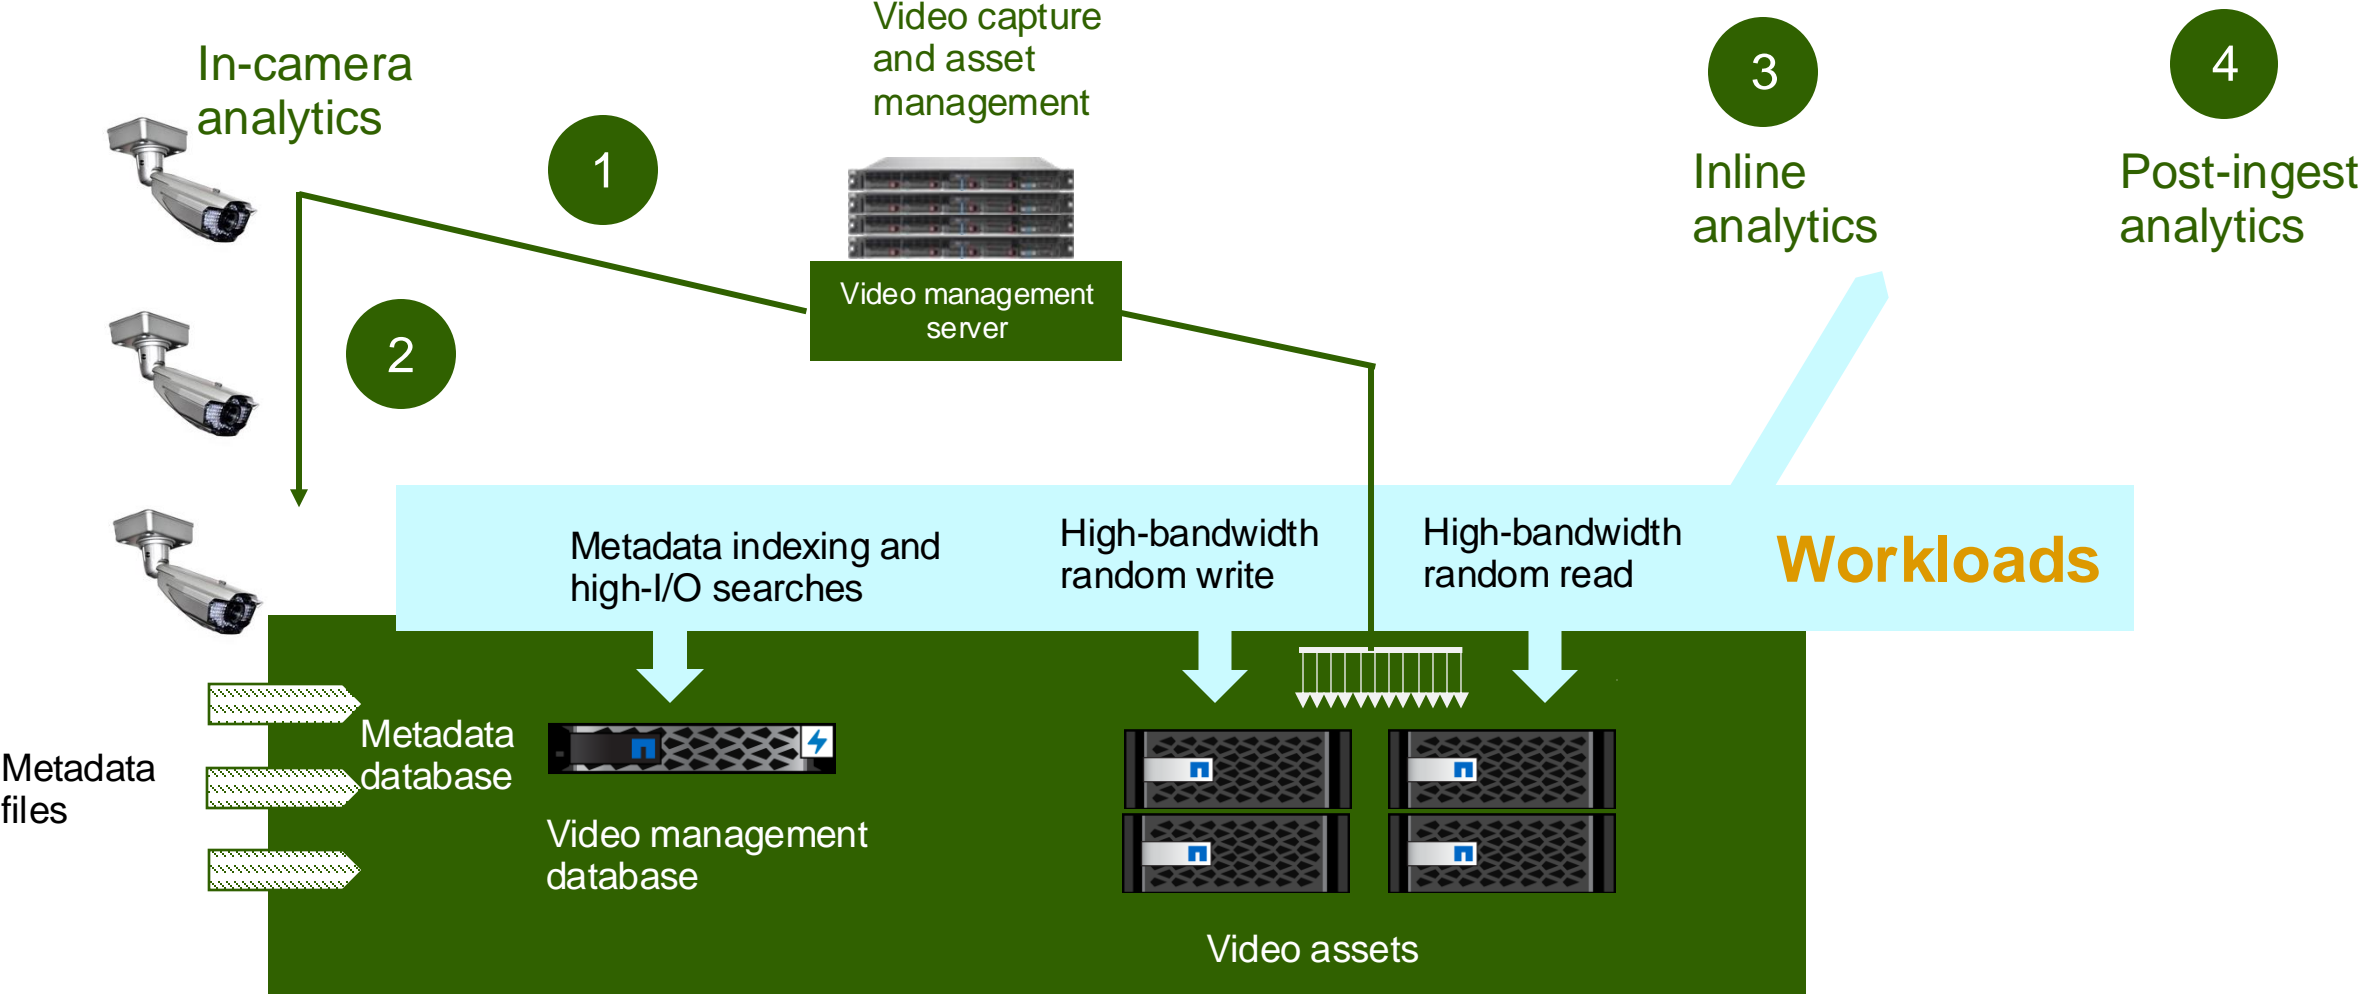

ANSIBLE scripts using for deployment

- GitHub / [ANSIBLE GALAXY](#)
- santricity: [1.4.0](#)
 - Santricity OS: RCB_11.80GA_6000_64cc0ee3.dlp
- beegfs: [3.2.0](#)
 - BeeGFS 7.4.3
- host: [1.3.3](#)
 - Lenovo Server SR665v3
 - RHEL 9.3
 - Connectx-7 dual-port NICs
 - Min. MOFED v23.10 LTS RHEL 9.3 drivers for the IB NICs
 - OpenSM MOFED v23.10 LTS based
- Install collections and dependencies on ansible control node:
- ansible-galaxy collection install netapp_eseries.beegfs
- Repository
- on the ONTAP NAS S3 via http/https

Ansible : HW installation/ configuration Automatization

```
[root@ictm1625h1a hsl_automation]#  
[root@ictm1625h1a hsl_automation]# ./ansible-playbook.sh -i inventories/beegfs_rhel_e5760.yaml playbooks/golden_configs/beegfs.yaml
```

Example 1: AI in Video Surveillance



Video Analytics and an Example of Deep Learning Application

NVIDIA's DeepStream SDK

Three Steps of Video Analytics





Image Classification




Object Detection



Object Tracking

NVIDIA DeepStream Workflow (Pipeline)




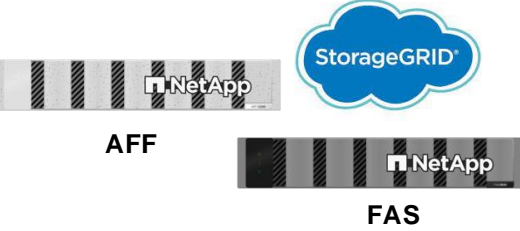



Source text placeholder

Sample Object Detection Pipeline - Output

Comes with DeepStream



Summary : NetApp AI/HPC workloads

		Enterprise AI (Building Center of Excellences)			
		High Performance Model Training	Data Prep/Data Lake Modernization	Model Training & Fine Tuning	RAG & Inferencing
Use Case	<ul style="list-style-type: none"> Build and Train Large Foundational models Parallel file system and/or InfiniBand requirements 	<ul style="list-style-type: none"> Modernizing Data Lakes Build Data Staging platform for AI workflow 	<ul style="list-style-type: none"> Build Models and Train them on enterprise data in an iterative process Produce models ready for inferencing Fine Tune pre-trained models with proprietary data 	<ul style="list-style-type: none"> Augment open-source models with proprietary data Storing Gen AI output data 	
Solutions	<p>EF-Series with BeeGFS for NVIDIA DGX SuperPOD</p> 	 <p>AFF</p> <p>FAS</p> 	<p>AFF A-Series or C-Series for NVIDIA DGX BasePOD and OEM Servers with GPUs</p> 	 <p>FSX</p> <p>InstaClustr</p> <p>NetApp</p> <p>Lenovo</p> <p>FlexPod</p>	
Key Points	<ul style="list-style-type: none"> Enterprise level support Competitive price/performance High availability 	<p>One Platform Cloud Connected NVIDIA Certified Ecosystem integrated Secure & Efficient</p>			



THANK YOU

