NetApp E/EF i FAS storage sistemi

Nebojša Ilić
COMING – Computer Engineering

COMING COMPUTER ENGINEERING NetApp

NetApp storage portfolio

Data ONTAP-v™
for public cloud

NPS for AWS
for hybrid cloud

AltaVault®
Backup to the Cloud

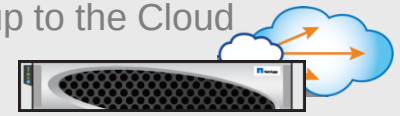

All-Flash Arrays
for highest performance

All Flash FAS

All-Flash SolidFire

All-Flash EF

NetApp® E-Series Arrays
purpose-built for performance, simplicity and reliability

FlexPod®, FlexPod Select
converged infrastructure

StorageGRID®
object storage

StorageGRID

FlexArray™
third-party array virtualization

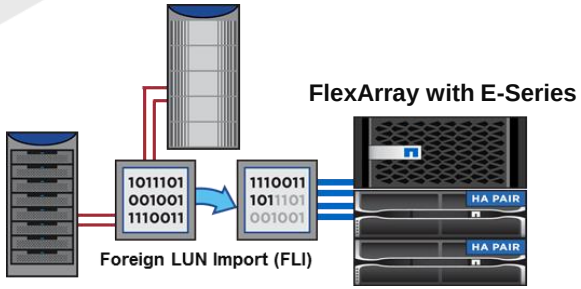

OnCommand®
management software

Hybrid Scale-out Arrays
for shared infrastructures

Professional and Support Services

COMING COMPUTER ENGINEERING

NetApp

NetApp E/EF serija

System highlights

Entry and enterprise-class block storage designed for superior price/performance, storage density, and throughput

Versatile reliable storage

- E2700
- E5600
- EF560

- Easy installation, low maintenance, and affordable
- Modular design: choice of shelves, drives, I/O
- Tunable performance: both IOPS, bandwidth, and latency
- Flash integration: performance and efficiency
- Sustained performance and automated management

E-Series focus on performance

Second Platform Enterprise

Workloads and Applications

- Classic database and OLTP
- Data warehouses/in-memory
- At-scale backup

Third Platform Enterprise

Workloads and Applications

- Hadoop
- NoSQL databases
- Analytics

Dedicated Workloads

Workloads and Applications

- HPC
- Cybersecurity
- Video analytics

NetApp E-Series at a glance

1M

Systems shipped (nearing)

20+

Years of OS hardening

200+

Patents

85,000

Items in current interoperability matrix

#1

SANtricity® is the #1 SAN OS

COMING COMPUTER ENGINEERING NetApp

NetApp E-Series storage systems

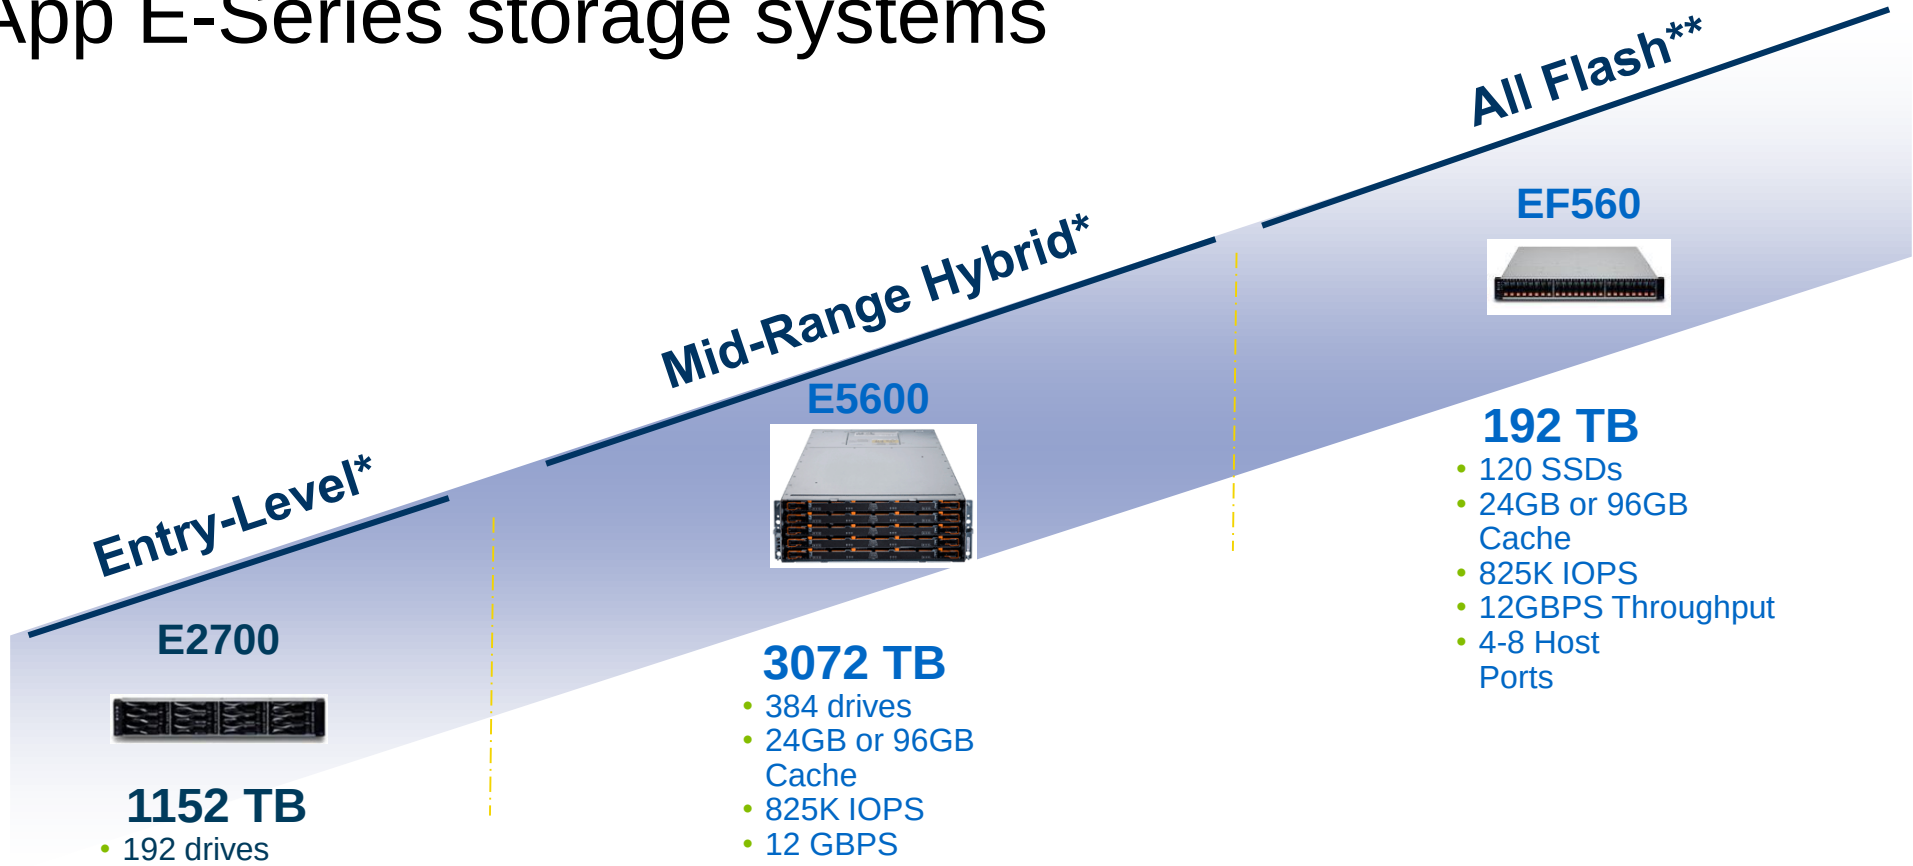

COMING COMPUTER ENGINEERING NetApp

Flexible modular system design

Controllers + Disk Shelves

Controllers	E2700 E5600	E2700 E5600 EF560	E2700 E5600
Disk Shelves	DE6600 (4U/60) 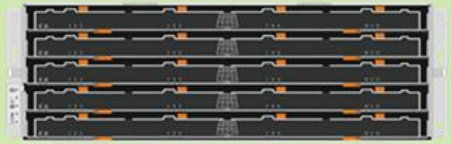 ▪ (60) 3.5” or 2.5” drives ▪ Highest throughput ▪ Largest capacity/density ▪ NL-SAS, 10K SAS & SSD drives	DE5600 (2U/24)  ▪ (24) 2.5” drives ▪ Highest throughput/RU ▪ Great performance/watt ▪ 10K SAS & SSD drives	DE1600 (2U/12)  ▪ (12) 3.5” drives ▪ Lowest entry price ▪ NL-SAS drives
	E5660/E2760	EF560/E5624/E2724	E5612/E2712

SANtricity disk structure

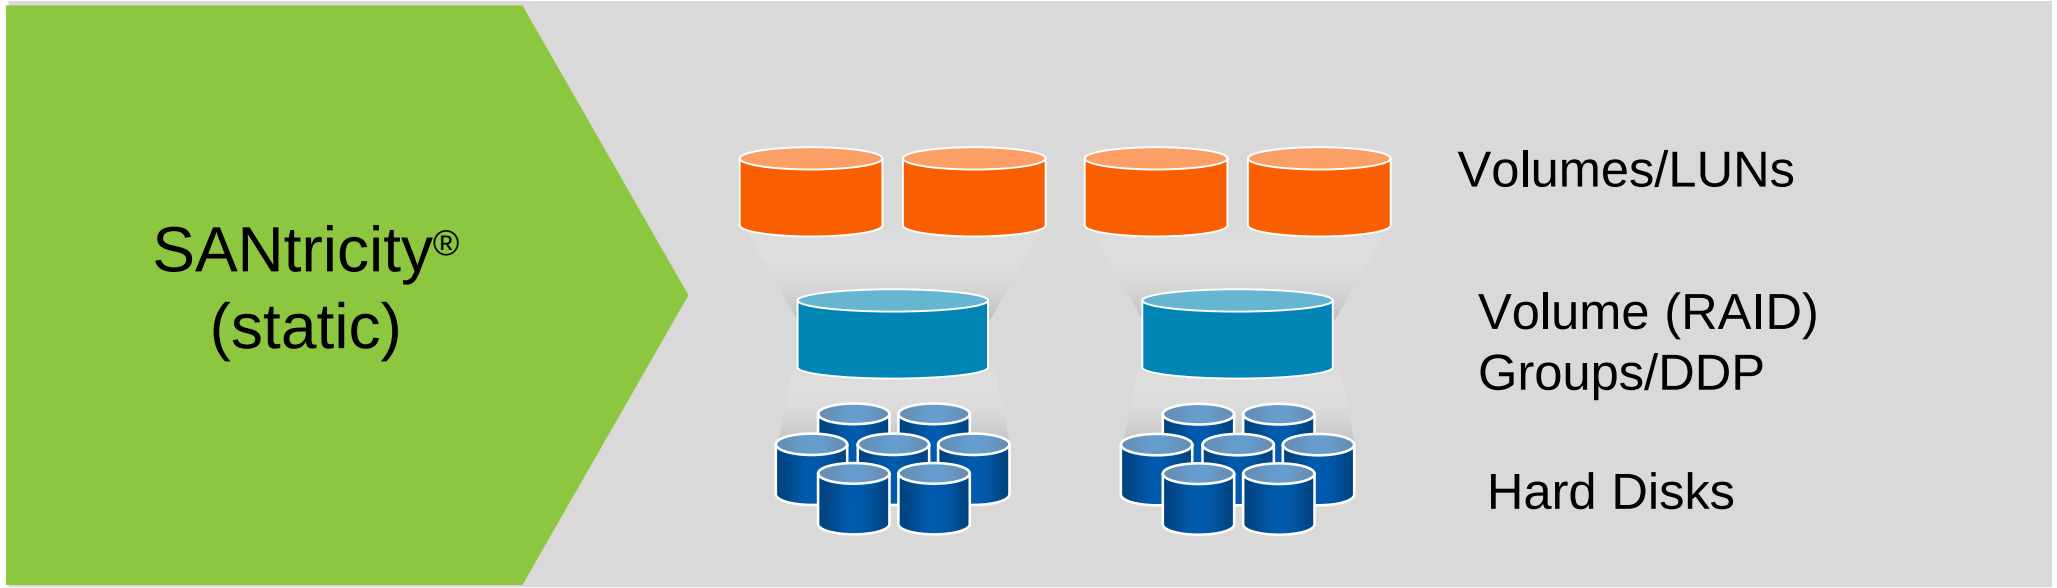

- Volumes are preallocated and statically mapped
- Based on RAID groups or disk pools
- Benefit: architected for performance and simplicity
- Volume groups can include up to 30 drives, except RAID 10 can be all drives; in practice, most groups should not exceed 12 drives
- DDP can use all drives in the system, best performance with 30 to 60 drive pools

SANtricity RAID levels

Block-level striping with a distributed parity

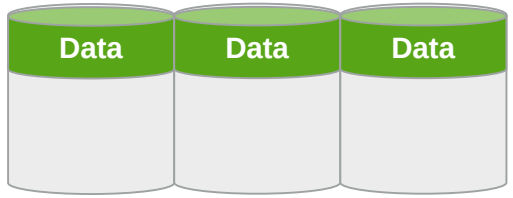

- RAID 0: striped

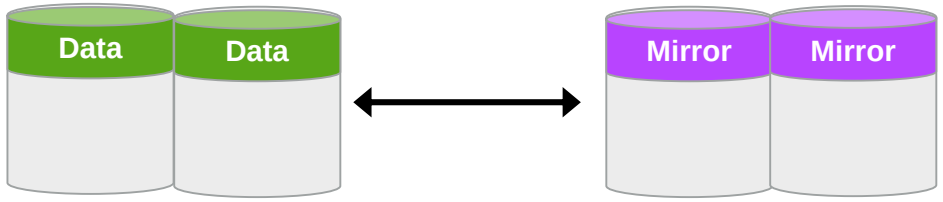

- RAID 1 (10): mirrored and striped

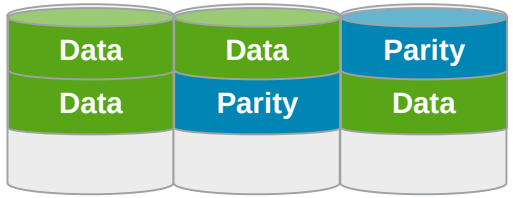

- RAID 5: data disks and parity

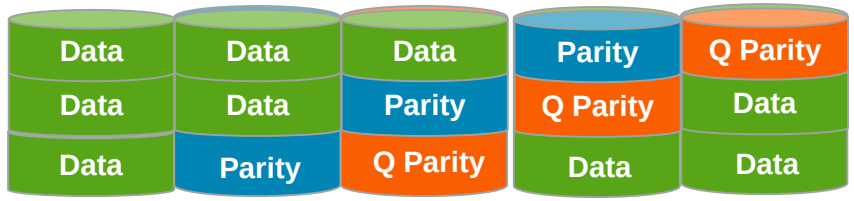

- RAID 6 (P+Q): data disks and dual parity

Traditional RAID volumes

- Disk drives organized into volume groups
- Volumes reside across the drives in a volume group
  - Performance is dictated by the number of spindles
- Hot spares sit idle until a drive fails
  - Spare capacity is “stranded”

24-drive system with 2 10-drive groups (8+2) and 4 hot spares

Traditional RAID: drive failure

- Data is reconstructed onto hot spare
  - Single drive responsible for all writes (bottleneck)
  - Reconstruction happens linearly (one stripe at a time)
- Performance for all volumes in that group are significantly affected

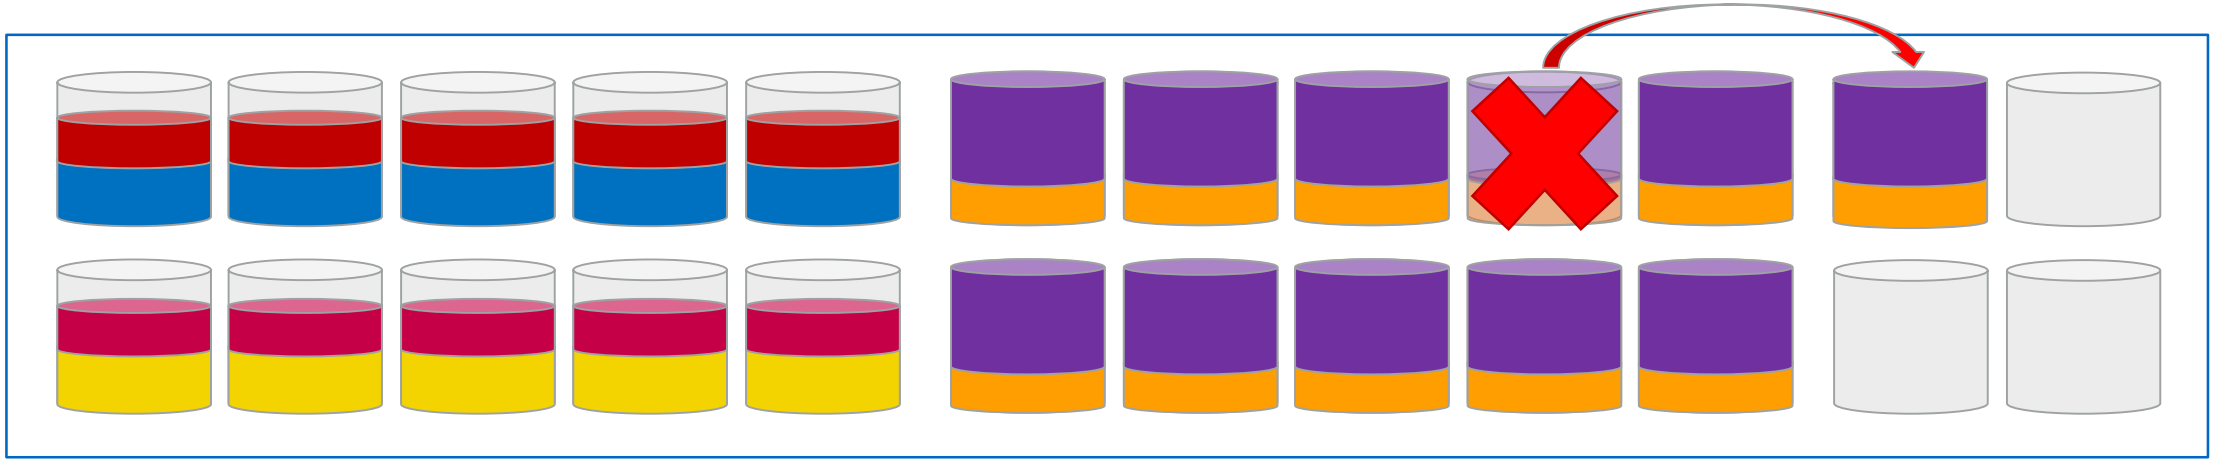

24-drive system with 2 10-drive groups (8+2) and 4 hot spares

Dynamic Disk Pools (DDP) stripes

- Each stripe resides on 10 drives within the pool
  - Always 10 drives, regardless of pool size
- Intelligent algorithm defines which drives are used
  - A different set of 10 drives is used for each stripe

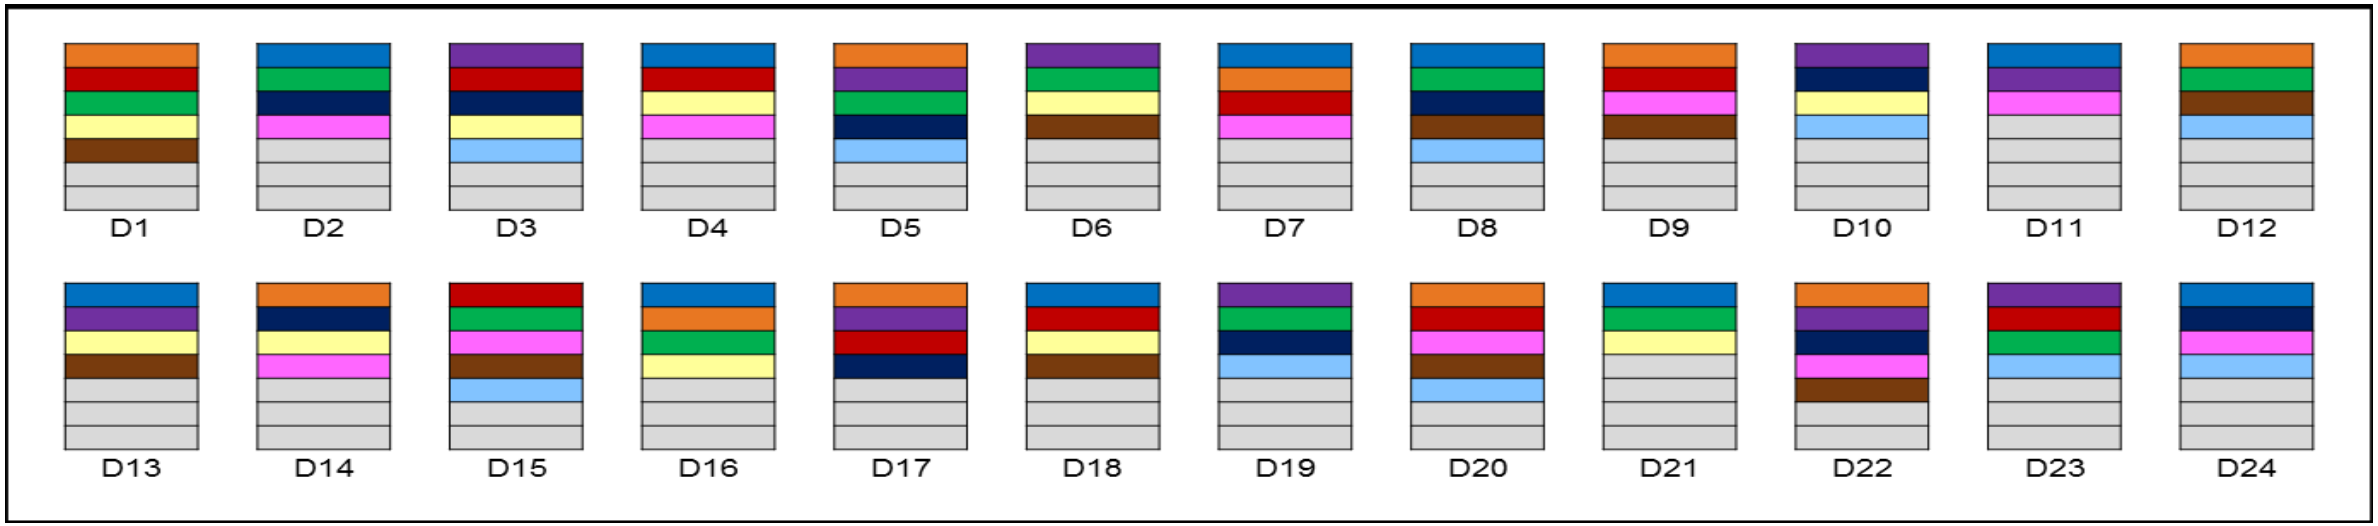

24-Drive Pool

Dynamic Disk Pools: disk failure

- For each stripe that has data on a failed disk:
  - Segments on other disks are read to rebalance data
  - A new disk is chosen to rebalance segments from failed disk
- Rebalance operations run in parallel across all drives

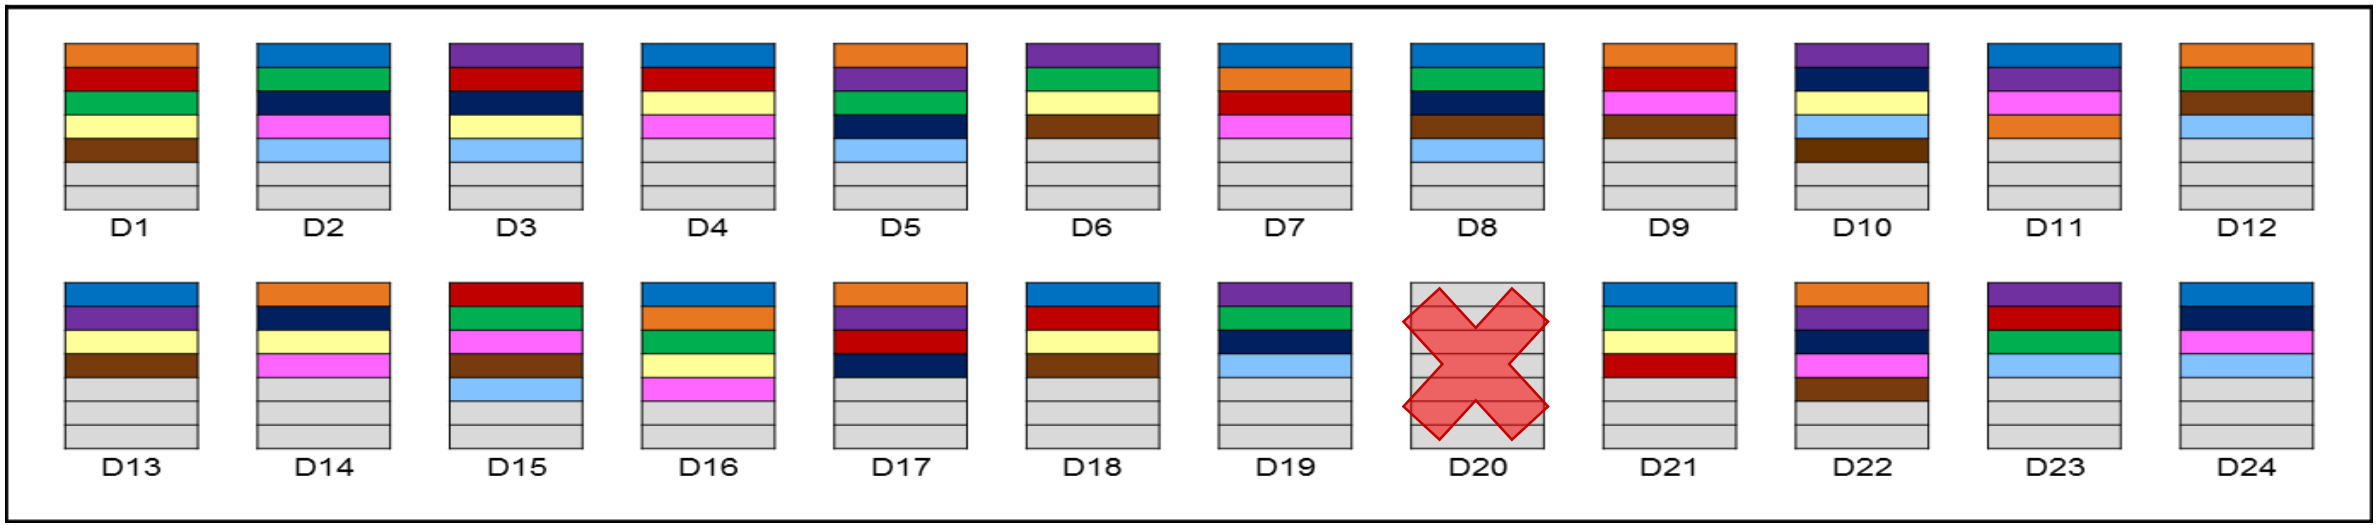

23 ~~24~~-Drive Pool

Data rebalancing in minutes vs. days

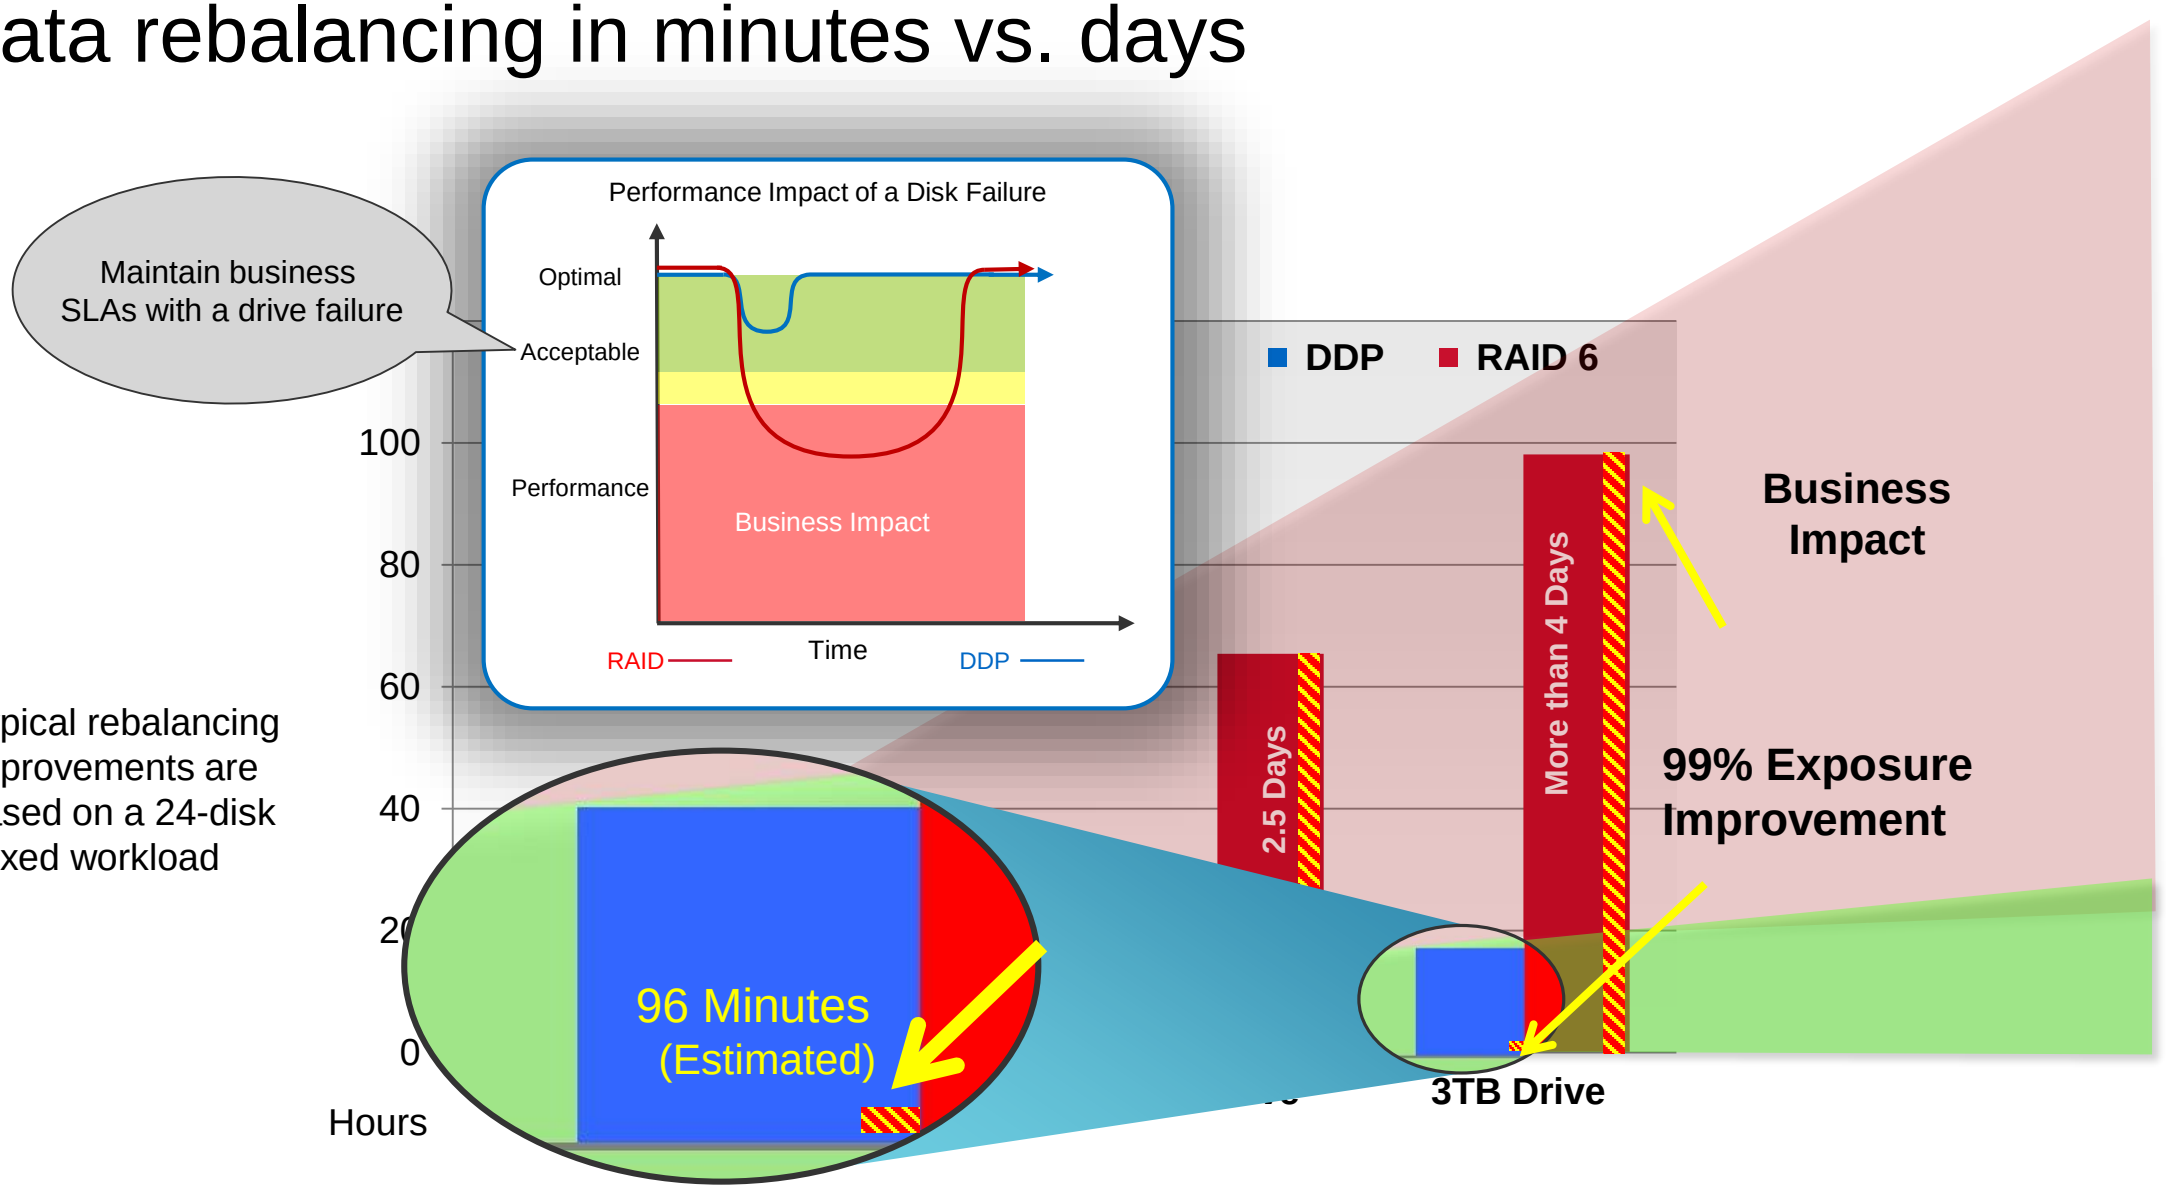

Typical rebalancing improvements are based on a 24-disk mixed workload

COMING COMPUTER ENGINEERING NetApp

SANtricity OS features

Encryption requires additional feature key; all others bundled with SANtricity

Dynamic Disk Pools

8x faster rebuild times (minutes, not days) and continuous high performance during drive rebuild

SSD Cache and Hybrid Storage

Expedited access to “hot” data through automated, real-time caching to SSD; mix and match SSDs and HDDs

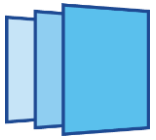

NetApp® Snapshot® and Volume Copies

More precise recovery point objectives and faster recovery

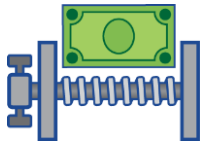

Thin Provisioning

Improve storage utilization by up to 35% and eliminate overprovisioning

Enterprise Mirroring (Sync and Async)

Cost-effective enterprise-class disaster recovery of data with FC and IP replication

Encrypted Drive Support

Extended security enhancements for compliance and regulations

COMING COMPUTER ENGINEERING NetApp

NetApp FAS – Clustered Data ONTAP

HA pair of FAS controllers

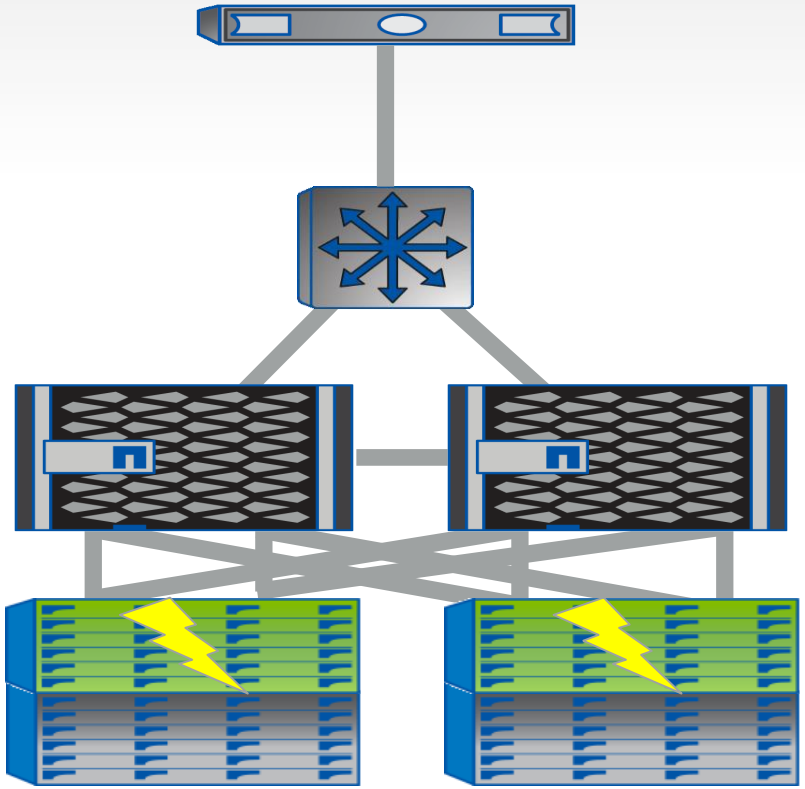

HA pairs may reside in the same chassis

- Provides redundancy
- Has an interconnect to mirror writes to partner's nonvolatile RAM (NVRAM)
- Each controller has connections to its partner's disks
- Provides redundant network connections

- = 7-mode Data ONTAP

Aggregates and RAID groups

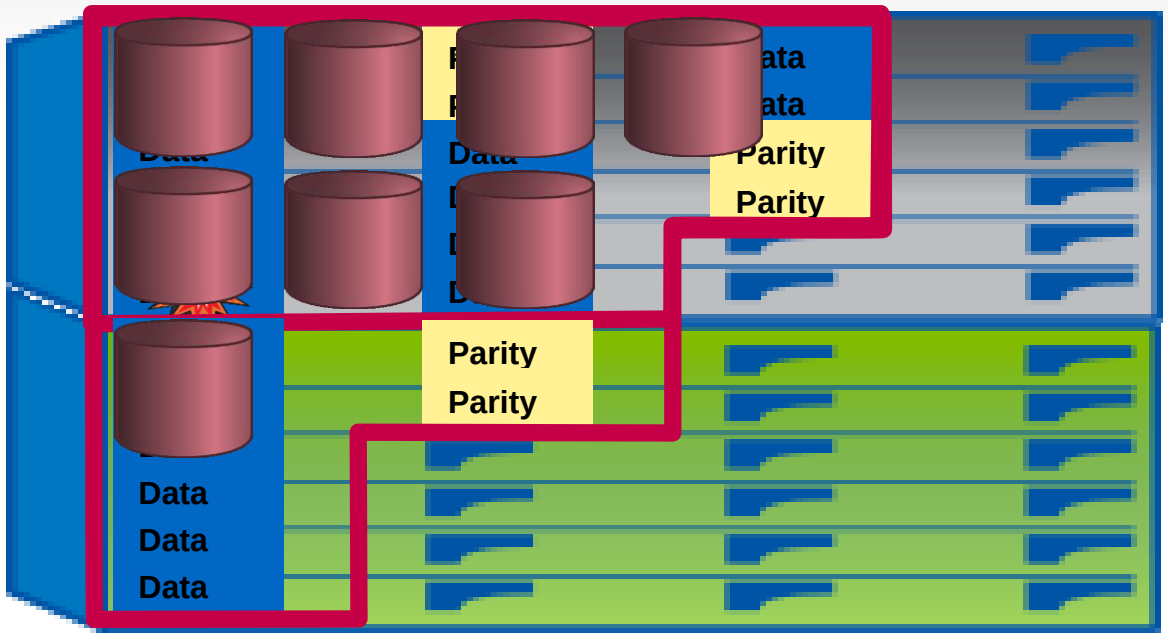

- Data is protected using NetApp® RAID DP® or RAID-4
- RAID DP can survive two simultaneous disk failures per RAID group
- One or more RAID groups form an aggregate
- An aggregate can consist of SSDs, spinning media, or both
- Aggregates provide physical storage for one or more flexible volumes

Logical interfaces (LIFs)

- Provide client and host connectivity to storage
  - NAS LIFs have an IP address
  - SAN LIFs have a WWPN
- LIFs use a network port for physical connectivity
- NAS LIFs can nondisruptively move to another network port
  - Client I/O continues uninterrupted
- SAN LIFs do not need to move
  - Paths to a LUN can be added or removed
  - The host's ALUA MPIO stack will switch to the most direct path nondisruptively

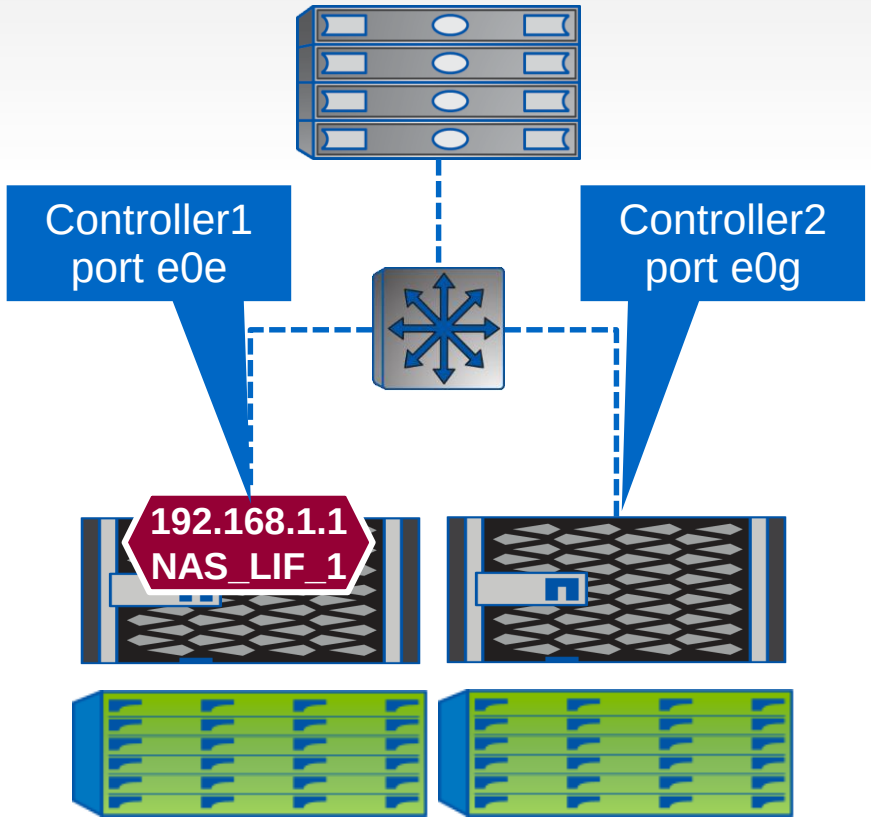

Cluster overview

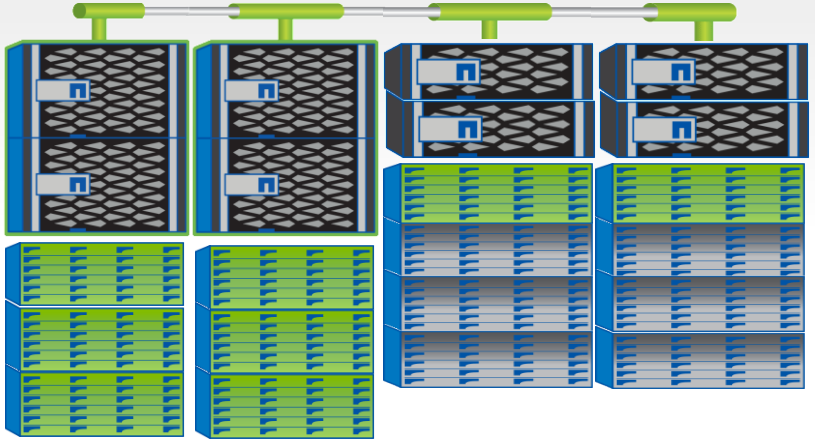

- Nodes are paired for high availability
- Consists of 1 to 24 nodes
- Up to 103PB in a single system
- Is connected with a dedicated, dual-fabric 10GbE interconnect

Cluster overview

- Set of connected storage nodes that work together as a single system
- Infrastructure changes do not affect user data access

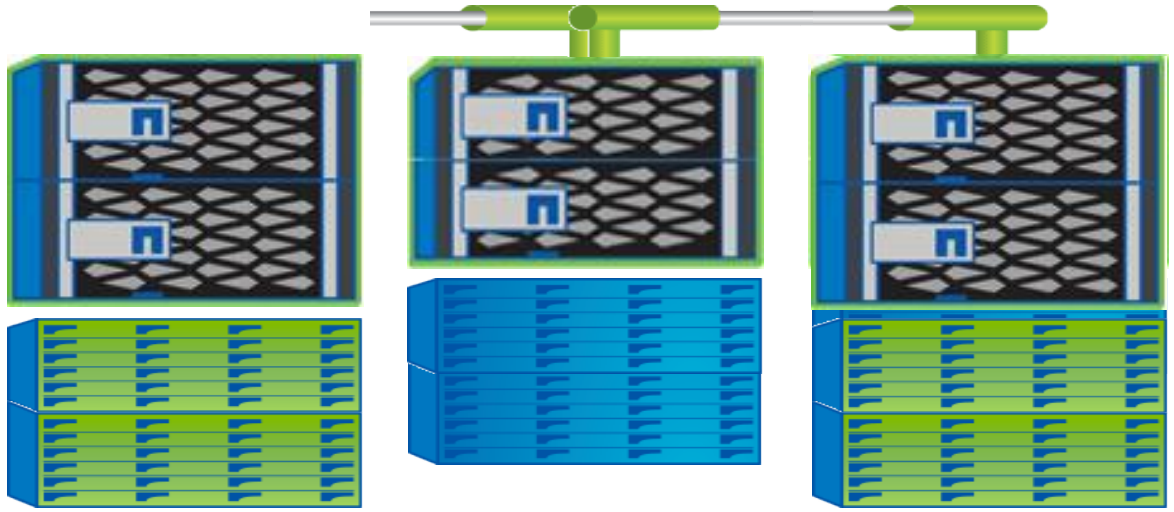

The host and client view of a cluster

Storage Virtual Machines (SVMs)

- Virtual storage systems
- Define storage services available to tenants or applications
- Required for data access
- Serve SAN, NAS, or both
- Include FlexVol® volumes and LUNs
  - Can use physical storage on any cluster node
- Include logical interfaces (LIFs); a LIF may have an IP address or a WWPN
  - Can have LIFs on any cluster node

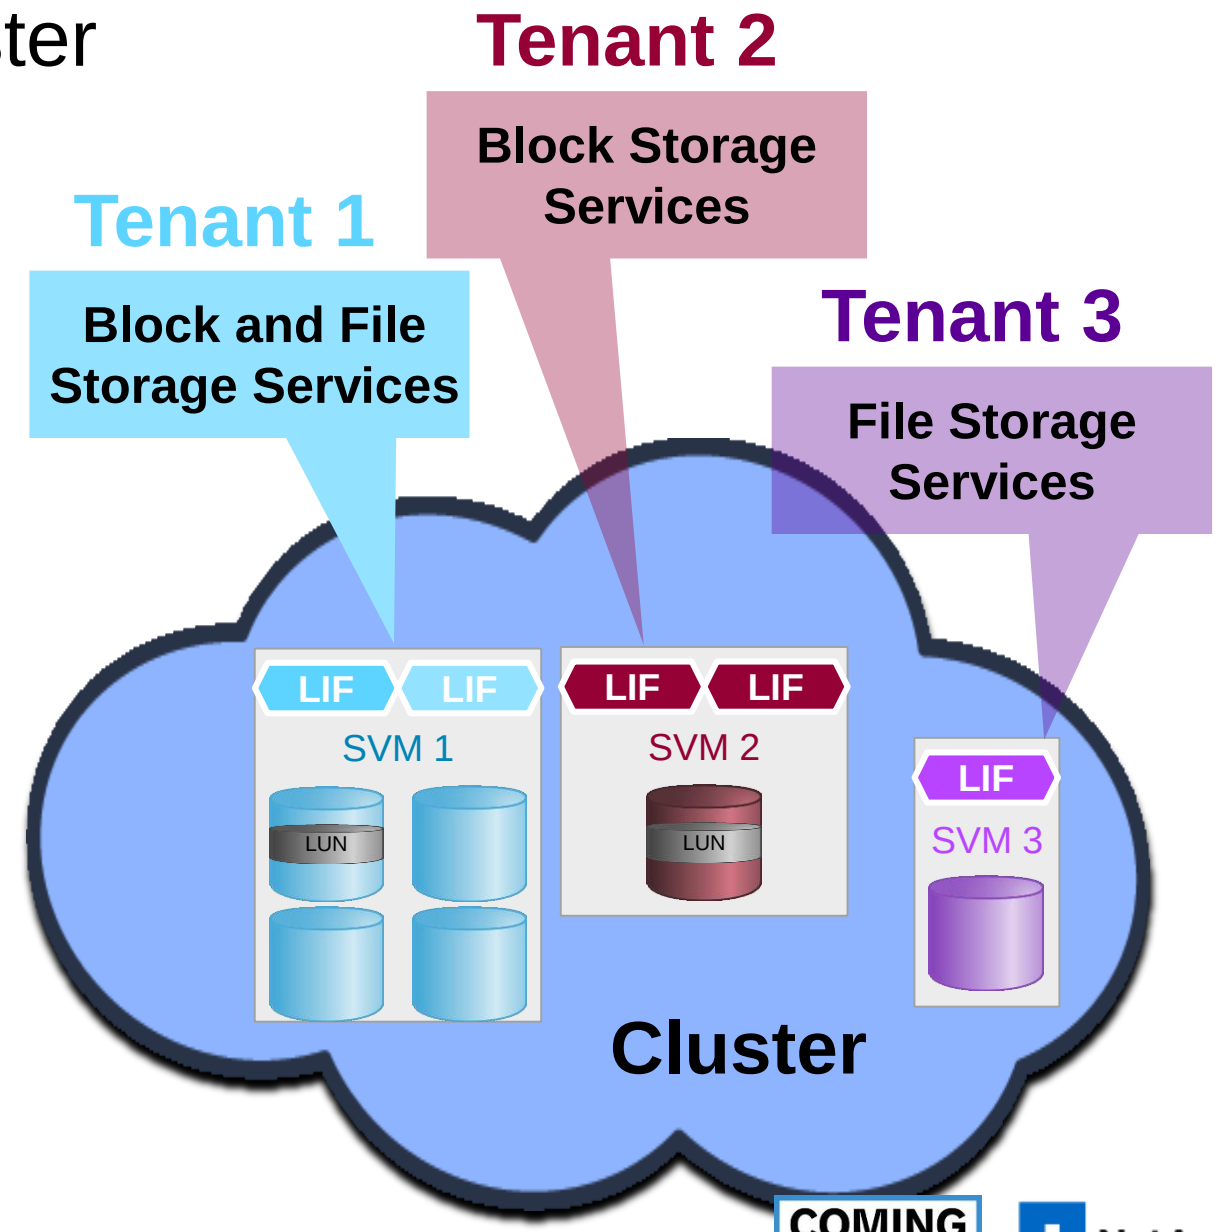

The host and client view of a cluster

Storage virtual machines decouple storage services from physical hardware

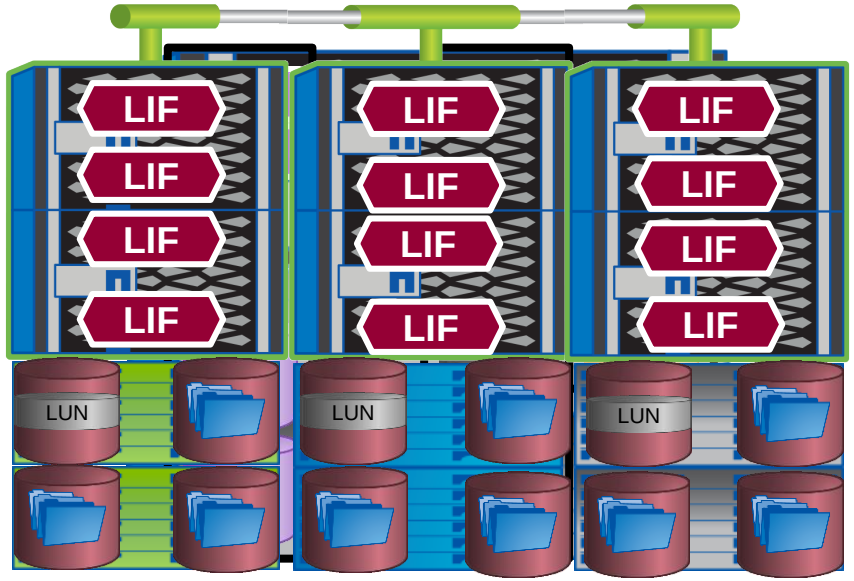

- A cluster can act as many virtual storage systems
  - Enabling multi-tenancy
- A cluster can act as one virtual storage system
  - Enables scale-out

Nondisruptive operations

- Move flexible volumes, LUNs, and LIFs while data is accessed
- The client or host view remains the same as data and network access changes

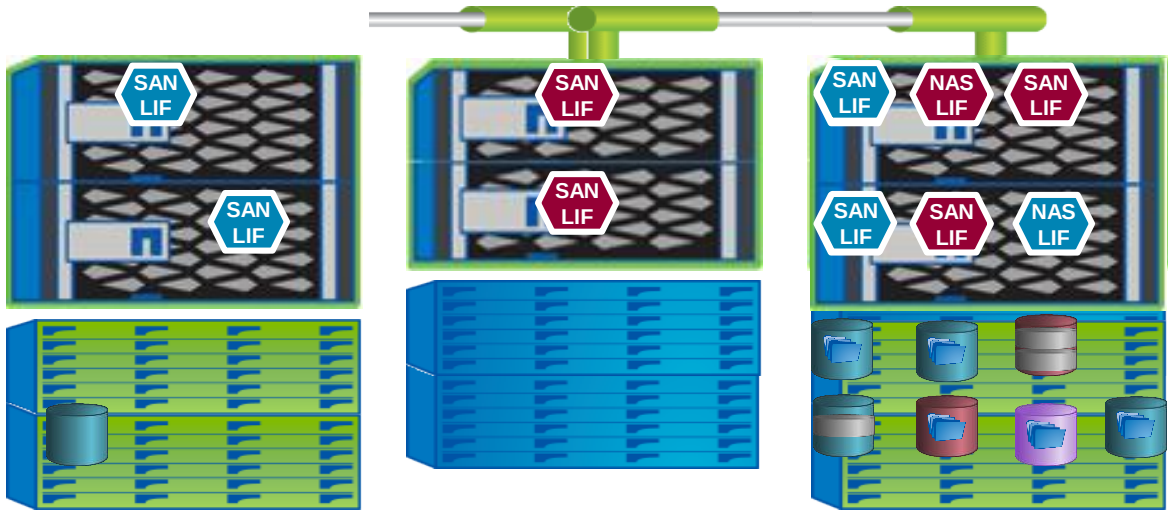

Storage Virtual Machine (SVM) 1

Storage Virtual Machine (SVM) 2

Storage Virtual Machine (SVM) 3

Backup, recovery, and business continuity

Data protection and DR between sites

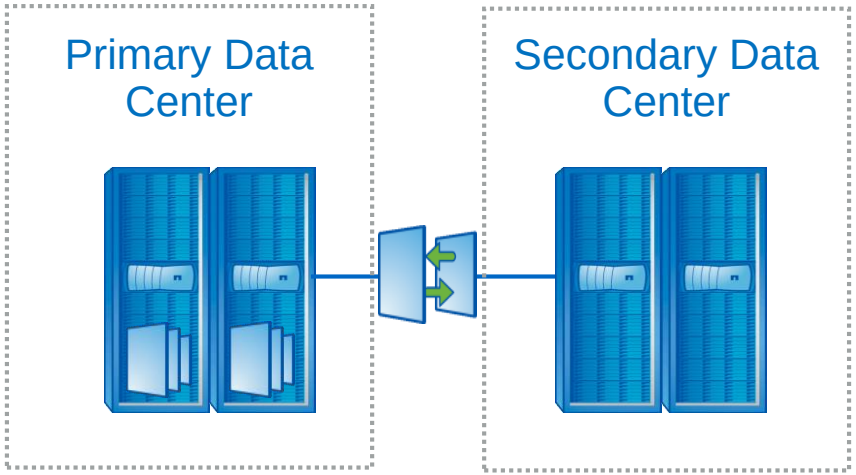

- NetApp® SnapMirror® software
  - Asynchronous mirroring to the same or a different cluster
  - Can become the writable primary copy during a DR event
  - Can resync with the original source when primary site comes online
- NetApp SnapVault® software
  - Long-term NetApp Snapshot® archive in the same or a different cluster
- NetApp MetroCluster™ software
  - Synchronous mirroring between two 2-node clusters up to 200km apart
  - Simple failover procedure

Quality of Service (QoS)

- Measure and control I/O
- Policy-based
- Limit I/O operations per second (IOPS) or raw throughput
- Use on SVMs, flexible volumes, data files, compute virtual machine disks, or LUNs
- Control bully workloads
- Manage tenant expectations

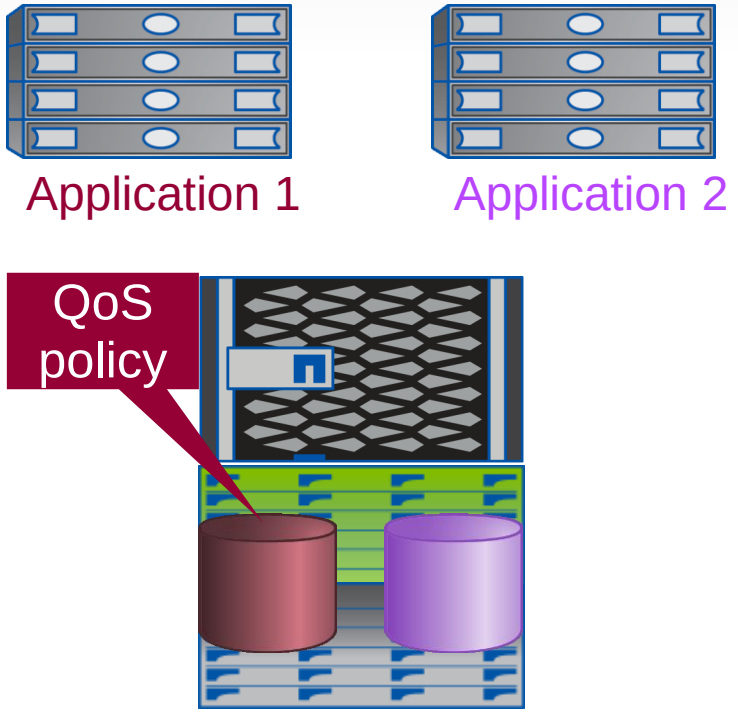

Proven efficiency

- Deduplication
- Compression
- NetApp® FlexVol® cloning
- LUN cloning
- File cloning
- Thin provisioning
- Virtual storage tiering
- Use individually or in combination

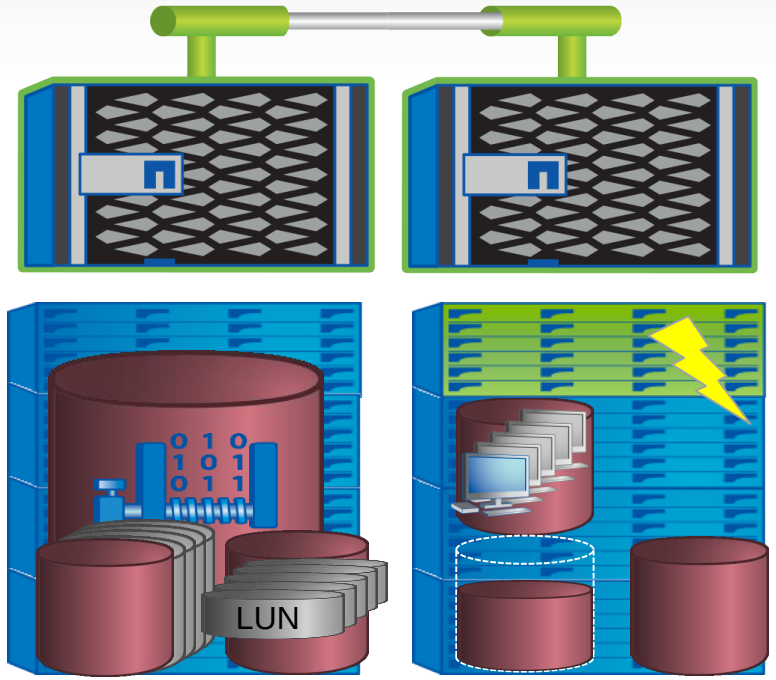

Proven efficiency

- A single resource pool
- Unified SAN and NAS
- On-demand performance and capacity optimization
- NetApp® AutoSupport™ system
- Single system to manage
- Leverage third-party storage arrays with NetApp FlexArray™ software
- Import third-party LUNs with Foreign LUN Import

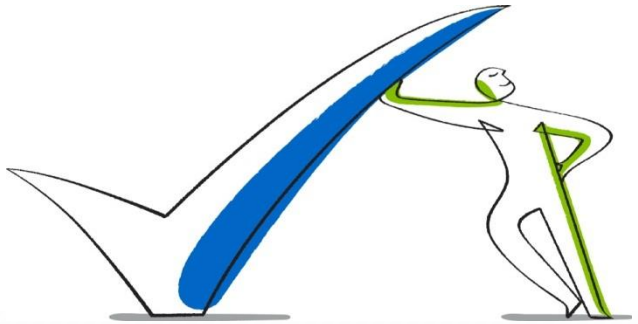

COMING COMPUTER ENGINEERING NetApp

Seamless scalability

Scale out to meet tenancy and application requirements

- Scales to 24 NAS-enabled nodes
  - Clients can access NAS volumes in a single namespace
  - Up to 103PB maximum capacity
- Scales to 8 SAN-enabled nodes
  - 96,000 LUNs per cluster

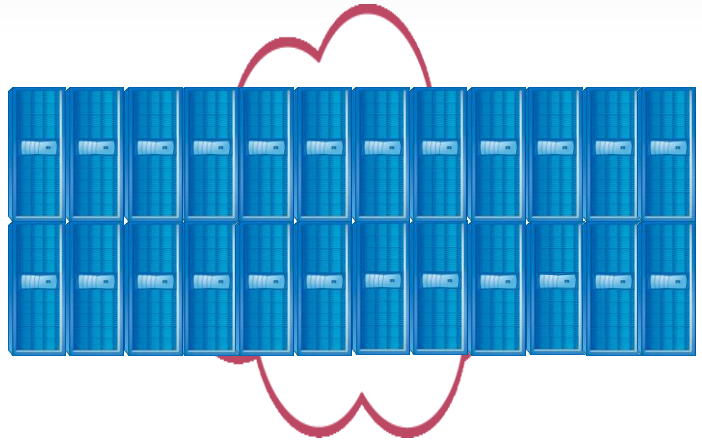

Scalable storage solution portfolio on a single architecture

Targeted solutions from FAS storage to third-party storage to the cloud

50TB
EBS Storage

Cloud ONTAP™
for AWS

504TB
84 Drives
16TB Hybrid Flash

FAS2520

749TB
144 Drives
16TB Hybrid Flash

FAS2552

864TB
144 Drives
16TB Hybrid Flash

FAS2554

2880TB
480 Drives
24TB Hybrid Flash

FAS8020

4320TB
720 Drives
48TB Hybrid Flash

FAS8040

7200TB
1200 Drives
72TB Hybrid Flash

FAS8060

8640TB
1440 Drives
144TB Hybrid Flash

FAS8080EX

Enterprise-Grade FAS	Extensive portfolio of hardware platforms from entry to enterprise
FlexArray™ Software	Reuse third-party disk arrays using FlexArray software
Storage Shelves	SATA, SAS, and SSD shelves with mixed deployments
NPS for Cloud	Scale compute in the cloud while maintaining control of your data

COMING COMPUTER ENGINEERING

NetApp

NetApp All Flash FAS

The next big wave in flash

Flash is better than high-performance disk in every way—including cost and capacity

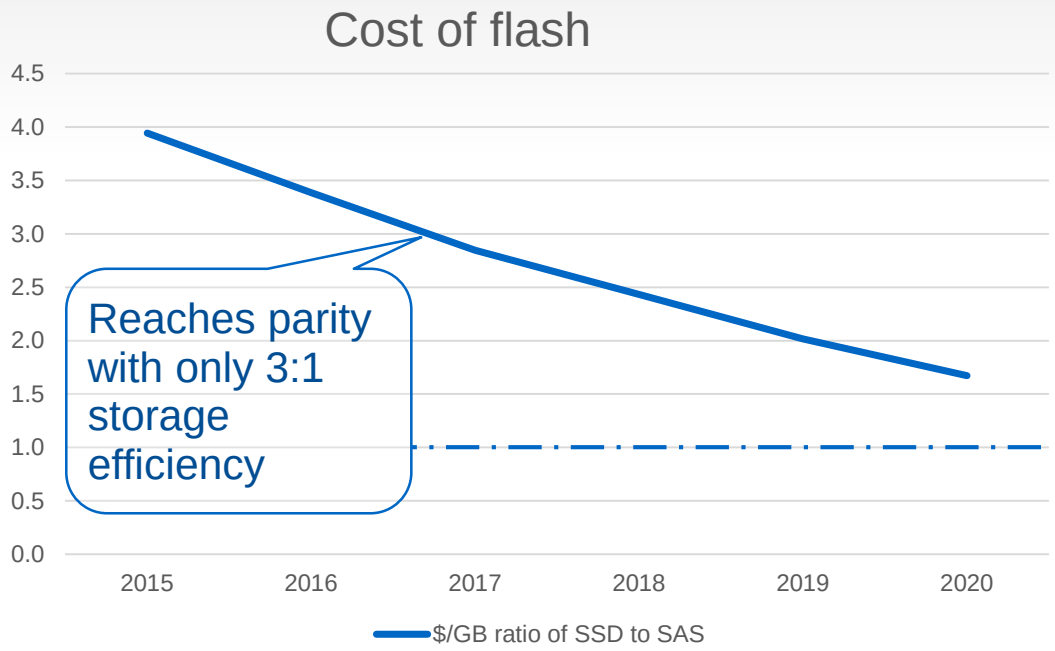

Flash is almost the same cost as SAS

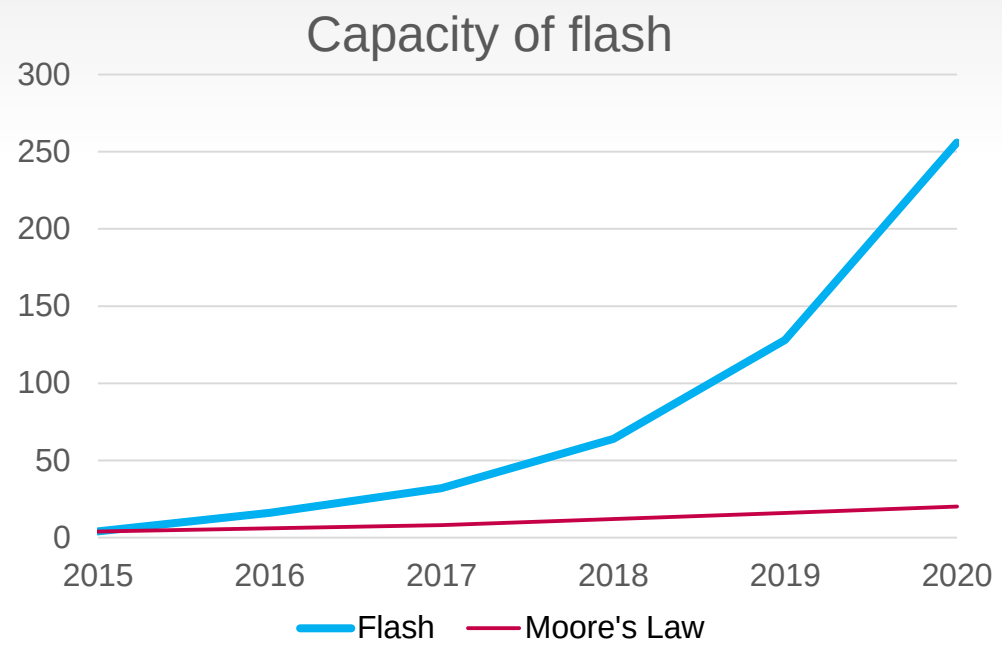

SSD capacities are growing faster than Moore’s Law; they will exceed HDD capacity this year

How is AFF Different From FAS?

- Same Controller
  - AFF8000 only
- Same DataONTAP
  - 8.3.x required for AFF
- Same SSDs
  - FAS support HDD as well

- AFF Specific Features
  - Storage efficiencies are enabled by default
  - DOT loads Flash Optimized code paths when booting on an AFF system.
    - Optimized SCSI drivers
    - Read Path Optimizations
    - Thresholds lowered to take advantage of SSD
  - Advanced Drive Partitioning
    - Increases usable capacity by up to 17%

All Flash FAS runs Data ONTAP 8.3.x

Protocols

- FC
- FCoE
- iSCSI
- NFS/pNFS
- CIFS/SMB

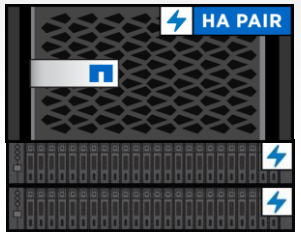

All Flash

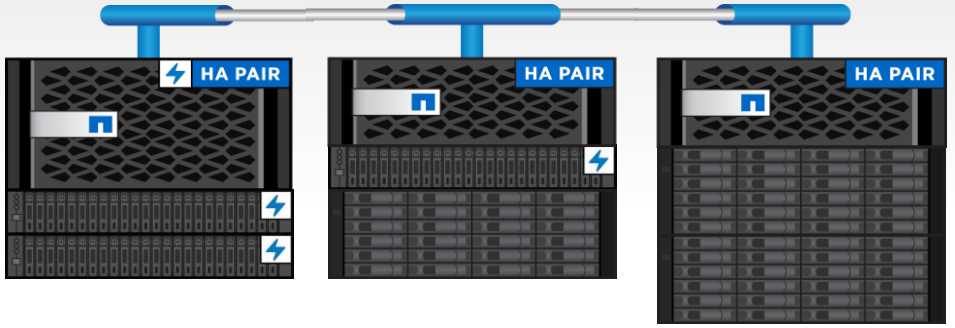

All Flash and Hybrid Flash/ HDD

Unified Management

- Scale-out and non-disruptive operations
- Data Mobility within a cluster
- Integrated data protection (Snapshots, SnapMirror, SnapVault)
- Storage efficiencies (RAID-DP, Thin Provisioning, FlexClone, Dedupe, Inline Compression)
- Advanced application integration
- Secure multi-tenancy, QOS, add without re-architecting

Multiple Deployment Options

All Flash FAS HA Pair

- Performance for dedicated workload

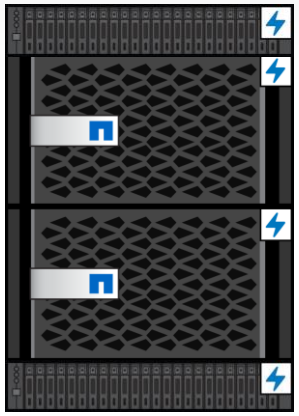

All-Flash FAS cluster

- Scale-out performance and capacity
- Up to 24 nodes and nearly 5PB (NAS)

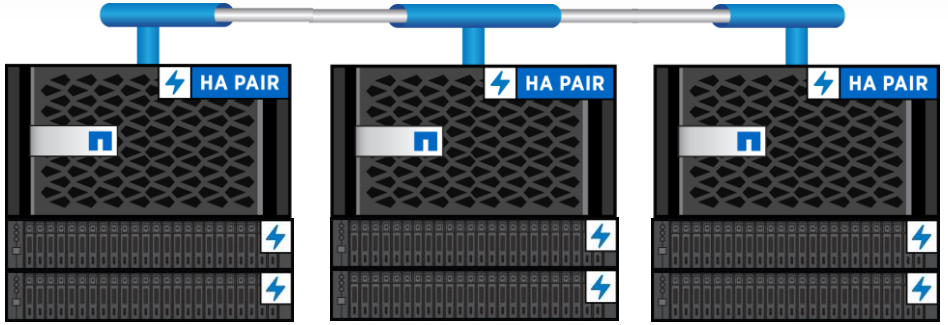

AFF w/Hybrid and HDD only FAS nodes in DataONTAP cluster

- Leverage AFF nodes for high-performance apps
- Transparently balance workloads between all three tiers within the cluster

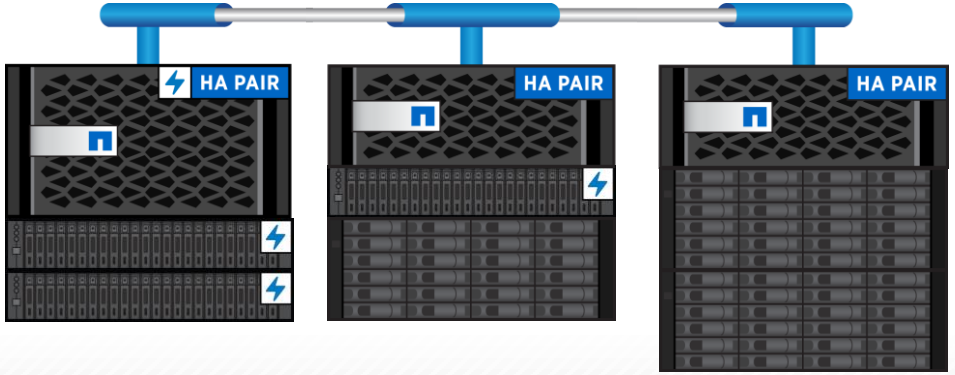

COMING COMPUTER ENGINEERING

NetApp

Data ONTAP 8.3.1 AFF Read Optimizations

Data ONTAP 8.2.x vs 8.3 vs 8.3.1

- Read Fast Path to Storage optimizations streamline read responses and requests
- Cumulative random read latency reduction of ~1 ms with Data ONTAP 8.3.1

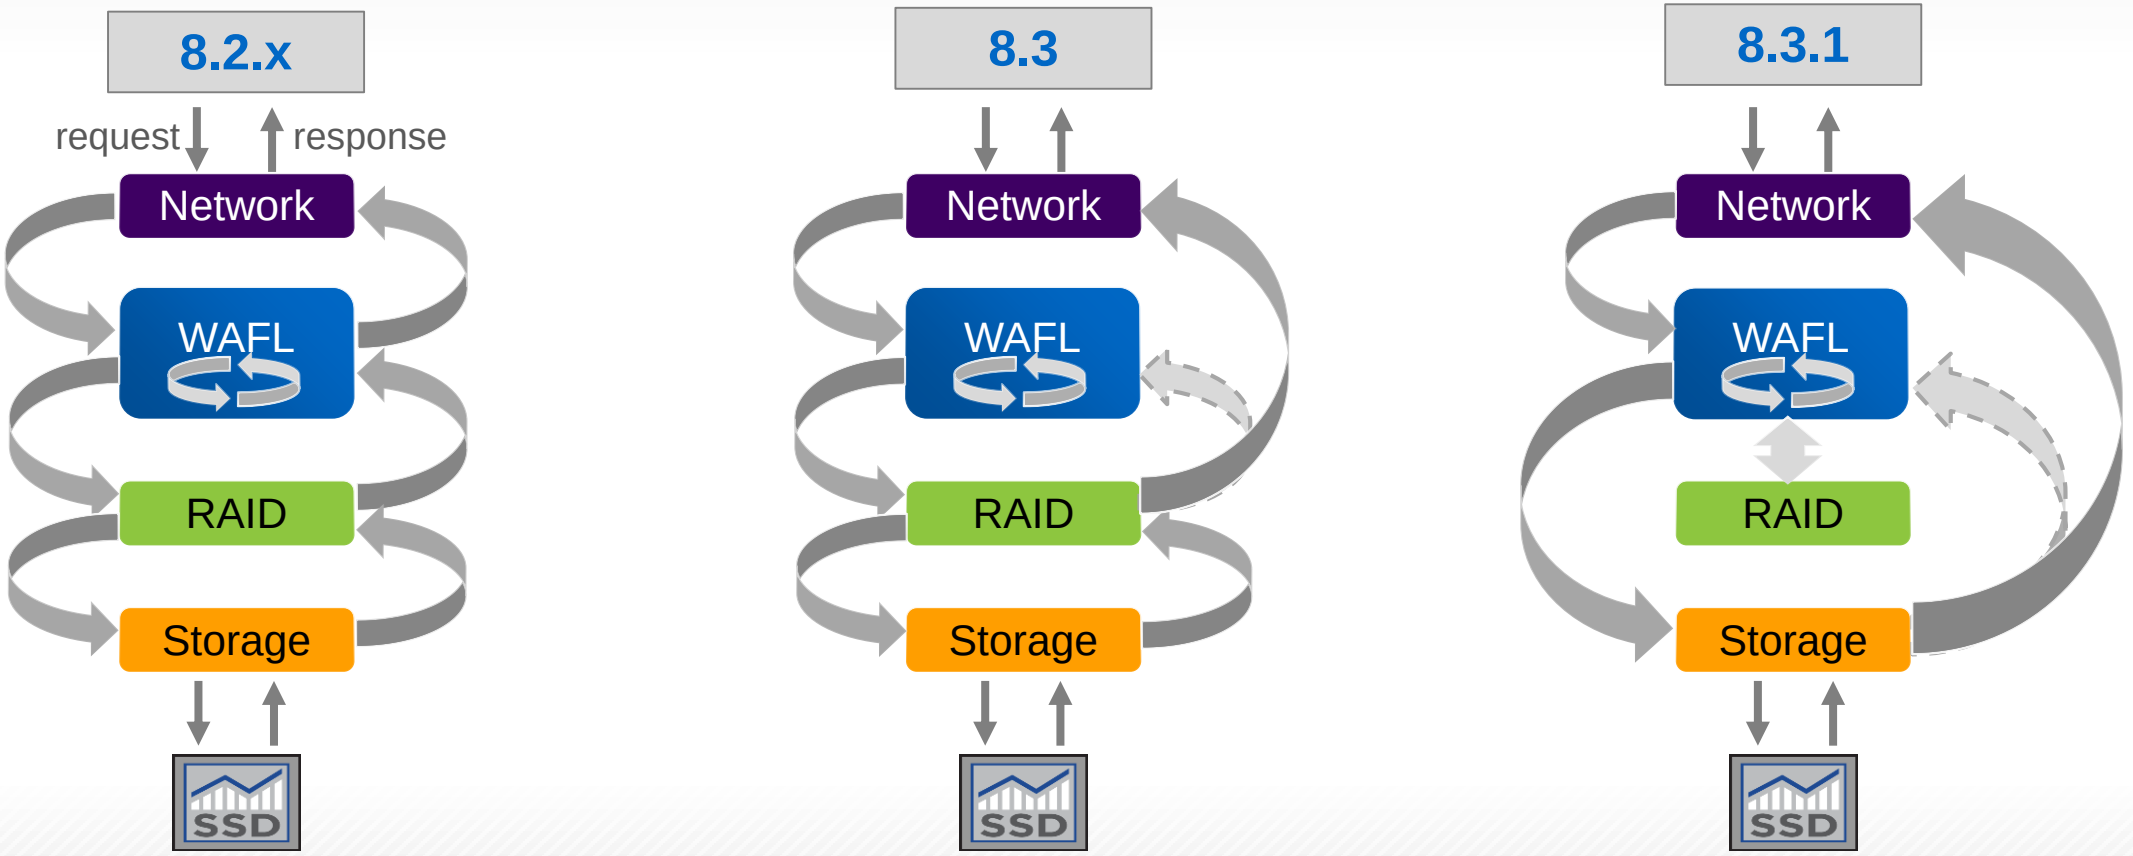

COMING COMPUTER ENGINEERING NetApp

Inline (Zero) Deduplication

- All data entering the system is hashed and compared
- Data that is similar is marked for deeper bit for bit inspection
- Zeros are immediately detected and eliminated with Inline Deduplication

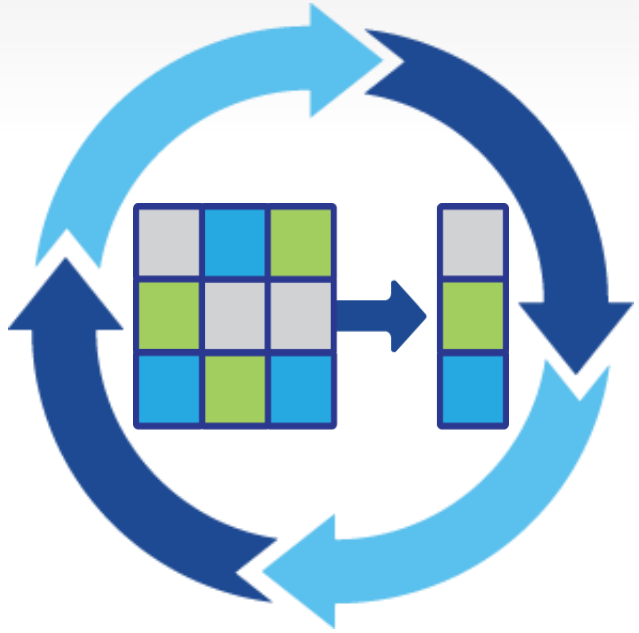

Inline 8KB Compression

- Compresses 8KB aligned I/O ‘segments’ into a 4KB WAFL block
  - Optimized for database applications
  - Maximum space savings is 50%, if all 8KB segments are compressible
  - Larger I/Os require multiple compression operations – e.g. 4 ops for 32KB
- On by default for new AFF volumes created with 8.3.1
  - Can be enabled on pre-existing volumes – only new writes are compressed
  - Optional policy for Flash Pool aggregate volumes
- Protocol and application independent algorithm
- Quick assessment reduces CPU consumed on incompressible I/O
- Predetermine savings with NetApp Space Savings Estimation Tool

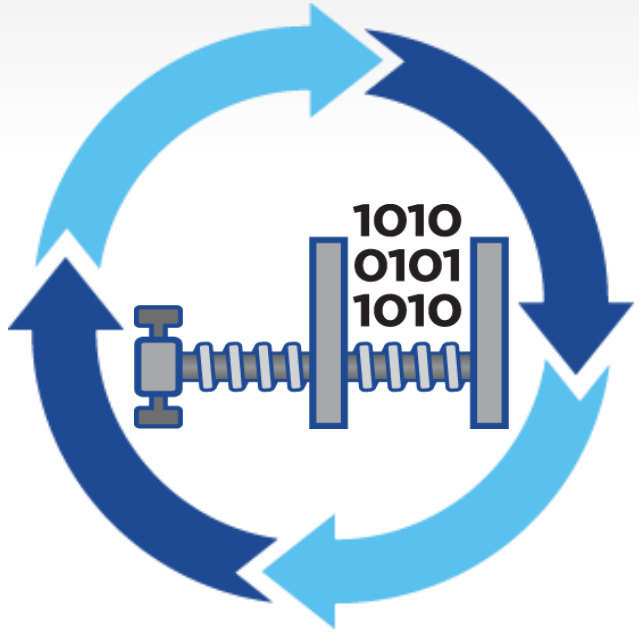

Always-On Deduplication

Storage efficiency with consistent latency!

- Aggressive deduplication schedule running every minute!
- Provides consistent latencies
- Client work takes precedence over background operations!
- All Flash FAS only
- Virtual Desktop environments ONLY!

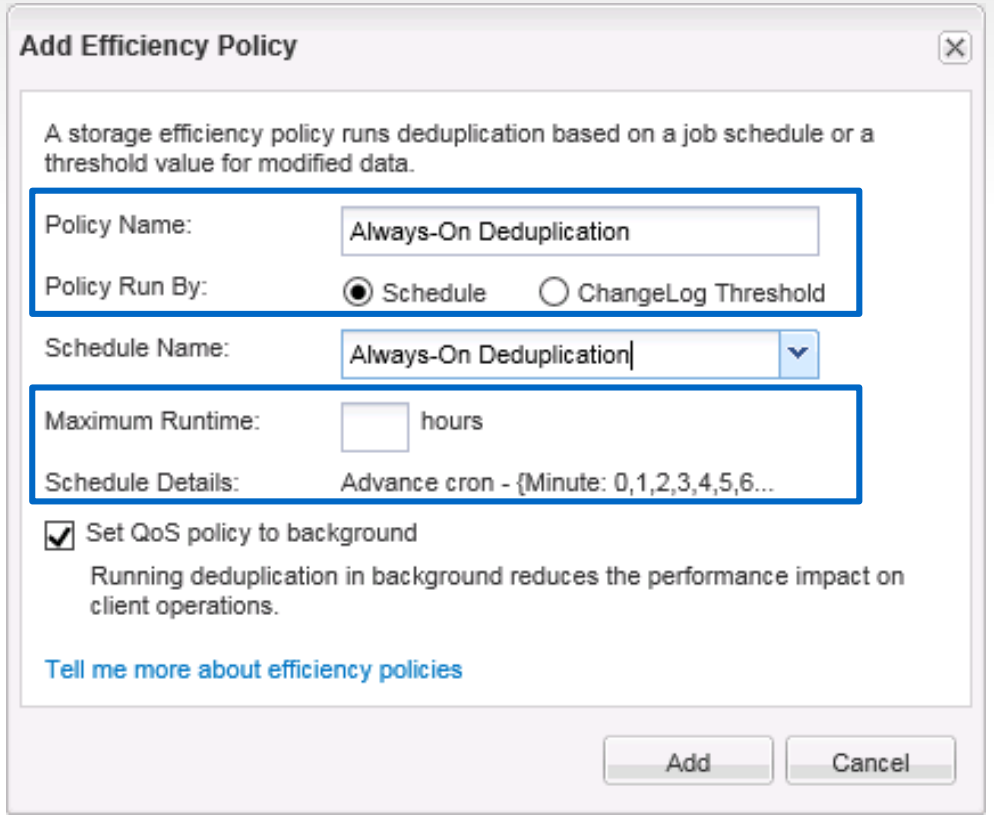

COMING COMPUTER ENGINEERING NetApp

Storage Efficiency – Reducing Data Footprint, Lowering Costs

- AFF efficiency features
  - Inline 8KB Compression
    - Default in 8.3.1 for AFF
  - Inline Zero Deduplication
  - Always-On Deduplication
    - Optimized for VDI
  - Advanced Drive Partitioning
- Data ONTAP efficiency features
  - Snapshots
  - FlexClones
  - Thin provisioning
    - FlexVol and LUN
  - Space efficient replication
    - Only consumed space replicated

6.4:1 observed savings among install base (ASUP data)

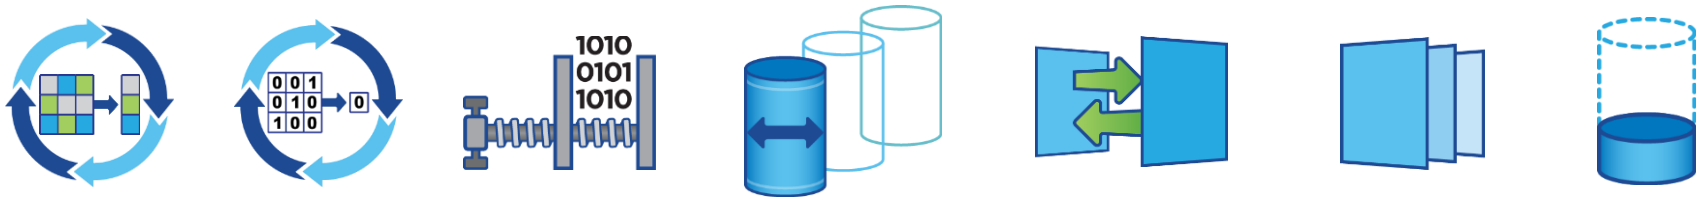

NetApp All Flash portfolio

NetApp offers the broadest all-flash portfolio in the industry

	EF-Series	SolidFire	All Flash FAS
Lead Differentiation:	▪ Speed/lowest latency	▪ Scale	▪ Data services
Buyer/Influencer:	▪ Infrastructure buyer, application owner	▪ Cloud architect	▪ Infrastructure architect
Deployment Model:	▪ Standalone apps/database acceleration	▪ Cloudlike infrastructure	▪ Virtual infrastructure

COMING COMPUTER ENGINEERING NetApp

NetApp Flash Leadership

Metrics through January 2015 (Q4 FY15)

180+ PB of flash

~200 flash-related patents

100% year-over-year growth (PBs)

70,000+ hybrid controllers

4,000+ all-flash controllers

Positioning AFF and EF

All Flash FAS | Data ONTAP®

Performance and versatility

EF560 | SANtricity®

Fast and single purpose

COMING COMPUTER ENGINEERING NetApp®

When to deploy

	FAS	E-Series	EF-Series
Protocol	▪ Unified and file	▪ Block	▪ Block
Data Management	▪ Needed in the storage layer	▪ Advanced management provided up-stack (app, OS, or file system)[1]	▪ Advanced management provided up-stack (app, OS, or file system)[1]
Workload Characteristics	▪ Shared, concurrent ▪ Transactional	▪ Dedicated, sequential ▪ IOP-intensive random IO	▪ Extreme low-latency
Performance Requirements	▪ Mixed I/O ▪ Scale-out IOPS	▪ Bandwidth ▪ Scale-up IOPS	▪ Consistent ultralow latency
Capacity	▪ Scale-out to 65PB	▪ Scale-up to 3PB	▪ Up to 192TB raw
Data Protection	▪ Integrated	▪ Application-driven	▪ Application-driven
Efficiencies	▪ Dedupe ▪ Compression	▪ Capacity/U ▪ Performance/U	▪ Performance/U
Flash	▪ Hybrid cache tier (read or read/write) ▪ Persistent storage	▪ Hybrid cache tier (read only) ▪ Persistent storage	▪ Persistent storage

1. SANtricity® provides value-add on-box management features and application integration plug-ins for enhanced application-based management.

Hvala na pažnji!