IBM Storage: Announcement Q4 2016

Ralf Colbus 11/2016
IBM Storage DACH
Member WW CTO Team
colbus@de.ibm.com

Storage Industry Evolves by Layers

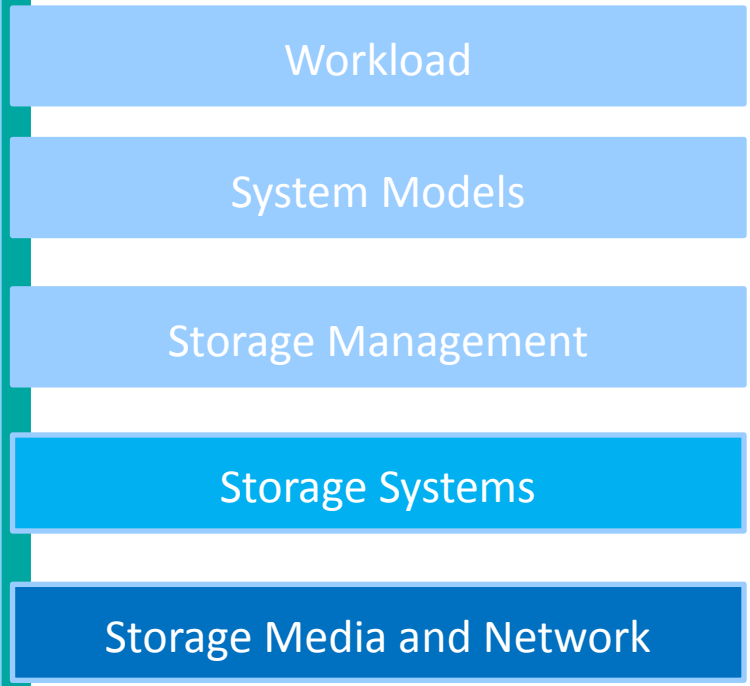

- **Cognitive and Analytic, Containerized Workloads**
- **Highly automated, Cloud/hybrid Cloud**
- **Shift to SDS and management of services through APIs**
- **Flash optimized, Efficiency, Cold Storage, Object as „3rd Storage Platform“**
- **IO Performance, Cost, Power and Space, Low latency Network,**

IBM Storage- The Next Decade:

Flash

IBM Spectrum Storage SDS

IBM Spectrum Computing, Optimized Storage Solutions

IBM Software Defined

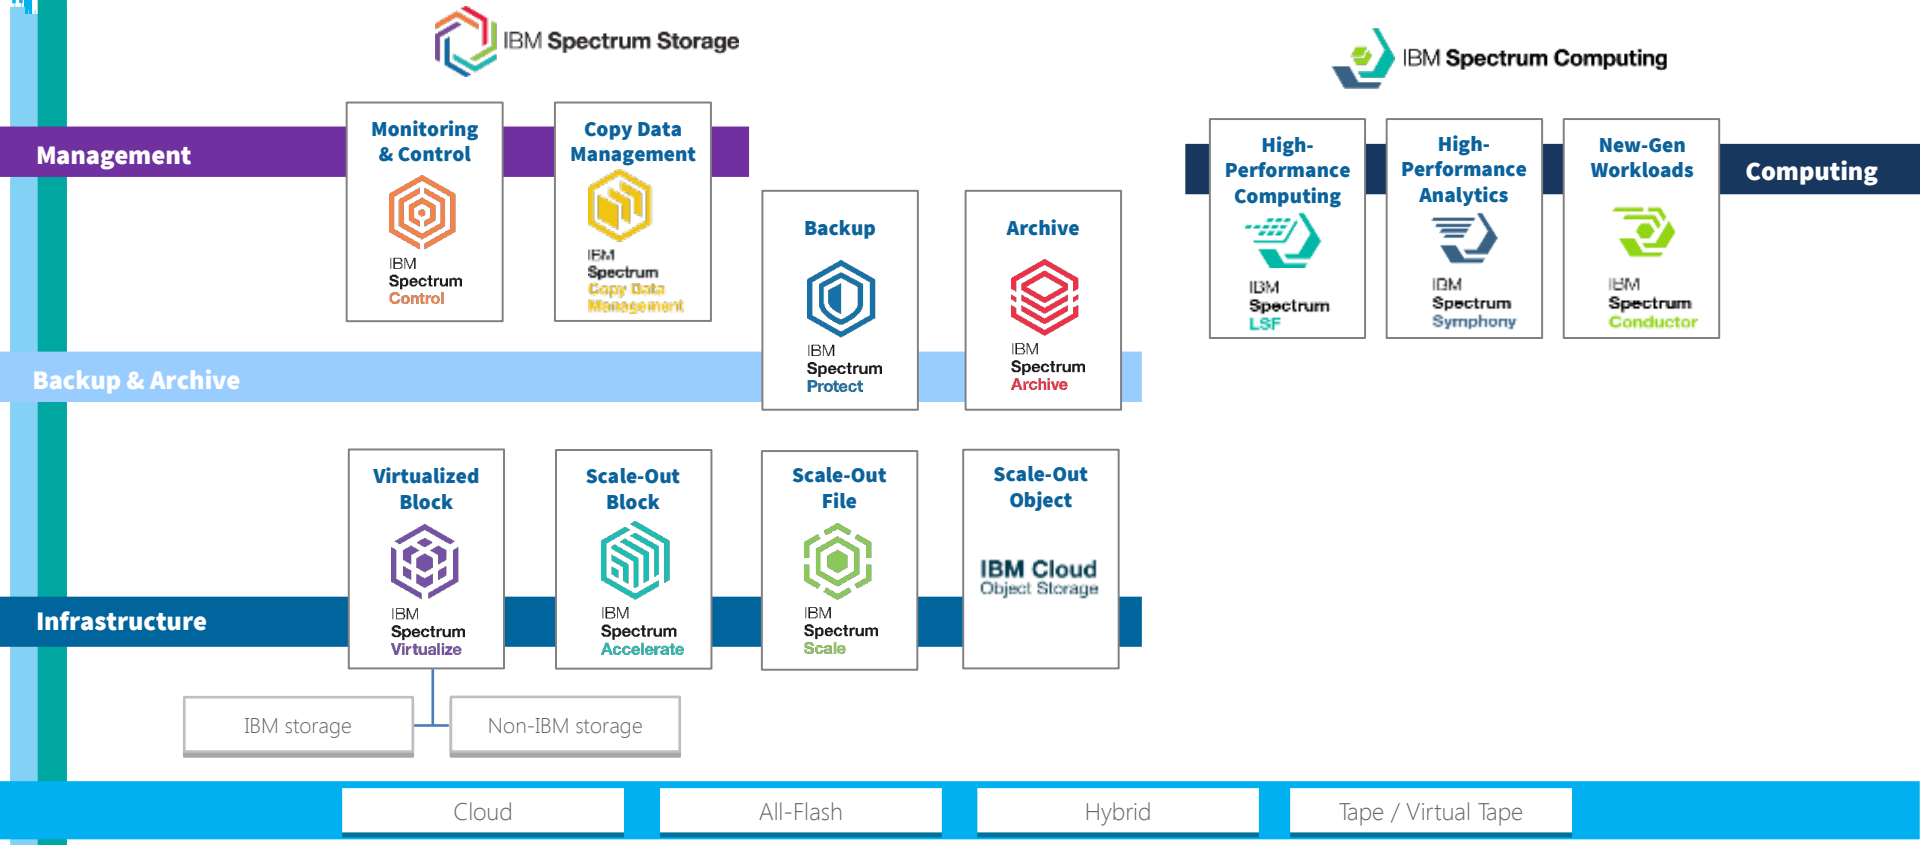

IBM Storage 2016

Storwize „F"

FS900

V9000

A9000 R

A9000

DS8888

DeepFlash 150
DeepFlash ESS

Flash-Storage

IBM **Spectrum** Accelerate

IBM **Spectrum** Archive

IBM **Spectrum** Control

IBM **Spectrum** Protect

IBM **Spectrum** Scale

IBM **Spectrum** Virtualize

IBM **Spectrum** Copy Data Management

IBM Cloud Object Storage

IBM Spectrum Storage

Storwize V5000-V7000

SVC

Storwize

Tape

VTL-Systems

DPR

DS8000

XIV

ESS

High End

Optimized Storage

IBM Storage Announcement Q4:

- IBM DeepFlash Elastic Storage Server
- IBM Spectrum Scale 4.2.2 with cloud tiering
- Spectrum Virtualize v 7.8
  - Cloud Tiering
  - New HD Drawer „Atlas“
  - Flash-drive and SKLM support

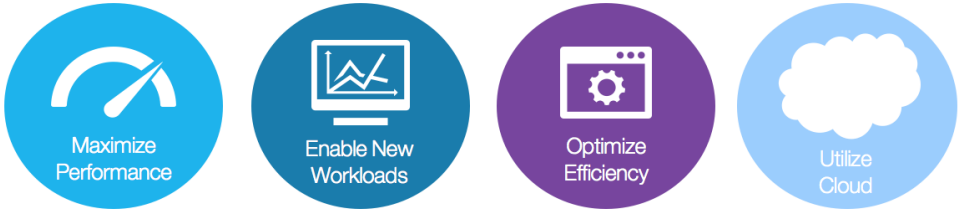

- DS8880 AFA with High-Performance Enclosure Gen2
- Spectrum Copy Data Management (CDM)

IBM All-Flash Across the Portfolio

Flash for every workload

V5030F
V7000F

FlashCore

FS900

V9000

A9000 R

A9000

DS8888

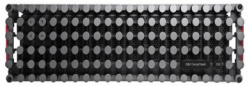

DeepFlash 150
DeepFlash ESS

Flash-Storage

IBM DeepFlash 150

Q3

Higher Reliability

- Hot-swappable architecture - easy FRU of fans, SAS expander boards, power supplies, flash cards

Game changing performance

- 2 M- IOPs
- <1ms latency
- 20GB/s throughput per enclosure

Higher Density

- 3U chassis with 128TB -256TB- 512TB
- >170TB per Rack Unit
- SAS attached

Solution with Spectrum Scale

- Unified data access including, file, object, HDFS and OpenStack
- Scales to exabytes of capacity and hundreds of TB per seconds of throughput under a single name space
- Low overhead encryption
- IBM lab services implementation

Lower Energy Consumption

- 150W(idle), 750W(abs -max), typical workload 450W
- 30% to 50% lower power and cooling power requirement of a equivalent HDD array

New

IBM DeepFlash Elastic Storage Server for File & Object

Configurations

- **Model GF1**
  - One DeepFlash 150 JBOF: 180TB usable
  - Dual IBM Spectrum Scale Servers
  - 7U
- **Model GF2**
  - Two DeepFlash 150 JBOF: 360TB usable
  - Dual IBM Spectrum Scale Servers
  - 10U

Performance

- **8 times faster response time***
  - 8 times lower latency compared to HDD based building blocks
- Extraordinary throughput
  - **GF1:** 13.6GB/s read, 9.2GB/s write
  - **GF2:** 26.5GB/s read, 16.5GB/s write

Advanced Data Protection

- IBM Spectrum Scale RAID erasure coding
- End-to-end checksum
- "Drive hospital" failure management

Solution with IBM Spectrum Scale

- Big data/analytics with transparent HDFS and Spark
- Low-latency tier for compute cluster
- Content management and media serving e-commerce back-end

* Based on preliminary SPEC SFS measurements

IBM DeepFlash™ Elastic Storage Server

New

ESS GF1

1 Flash Enclosures, 7U
180TB of usable Flash;
Max Read **13.6 GB/sec;**
Max Write **9.3 GB/sec**

DeepFlash JBOF

Spectrum Scale I/O server

DeepFlash JBOF

ESS GF2

2 Enclosures, 10U
360 TB of usable Flash
Max Read **26.6 GB/sec;**
Max Write **16.6 GB/sec**

*based on SPEC SFS results

New Big Data alternative: instead of HDD, use Big Data Flash

Move from Big Data **HDD** configuration

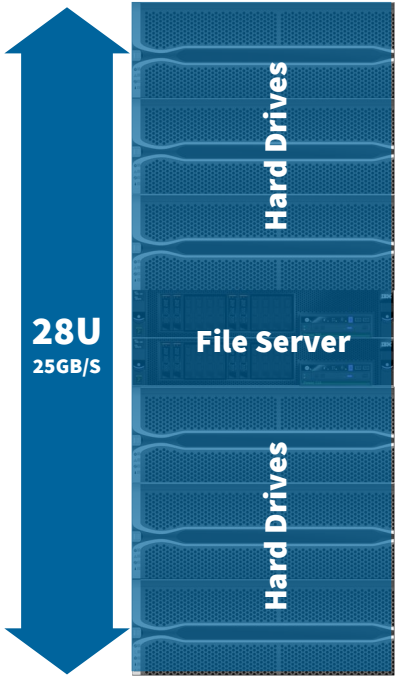

Improve application response time by 8X

Increase throughput/rack unit by 2.8X

Reduce MTBF

Reduce power & cooling costs by 30%-50%

To this **Big Data Flash** configuration

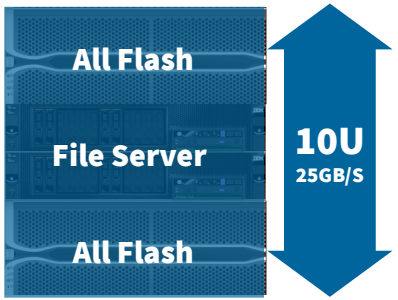

8X faster response time and same throughput as the HDD version

IBM Spectrum Scale 4.2.2

Big Data - Analytics - Unstructured Data – In one System

IBM Spectrum Scale

New

Encryption and security

- **Sec17a-4 certification by KPMG**
- **Bulk AV scanning added to policy**
- **Third party audit logging with lightweight events**

Usability and GUI update

- **Enhanced troubleshooting & comprehensive monitoring**
- **Streamline daily tasks with guided interfaces**
- **Self-tuning installation parameters with system query**

Enterprise flexibility

- **iSCSI client support, including diskless boot**
- **RESFTful API for management**

Big data analytics

- **HDFS connector updated for compatibility with latest apps**
- **Federated HDFS capabilities, inclusion of existing HDFS storage**
- **Data ocean: multiple HDFS domains in single storage pool**

HDFS Transparency Gen2

New

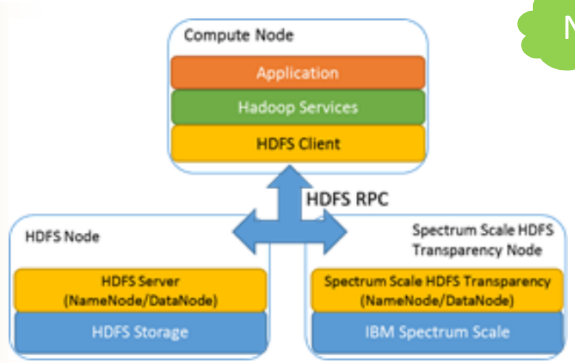

- No Spectrum Scale Client is needed on every Hadoop node
- Kerberos support
- Leverage cache of HDFS client
- **Now with Transparency Gen2:**
  - - Auto HA: automatically fail over to standby NameNode
  - - Distcp: easy migration from one cluster to S.Scale
  - - Federation: HDFS Transparency NameNode has no scaling issue (data lake to data ocean)

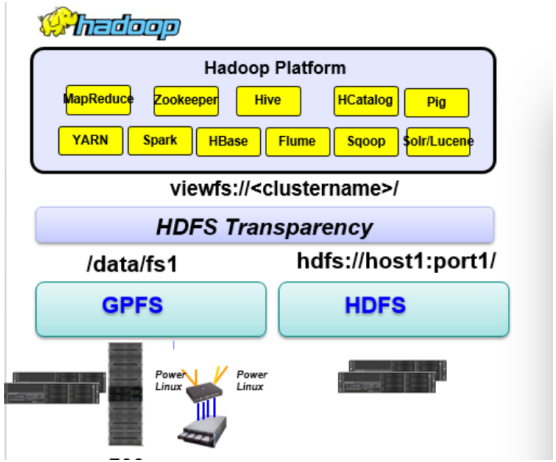

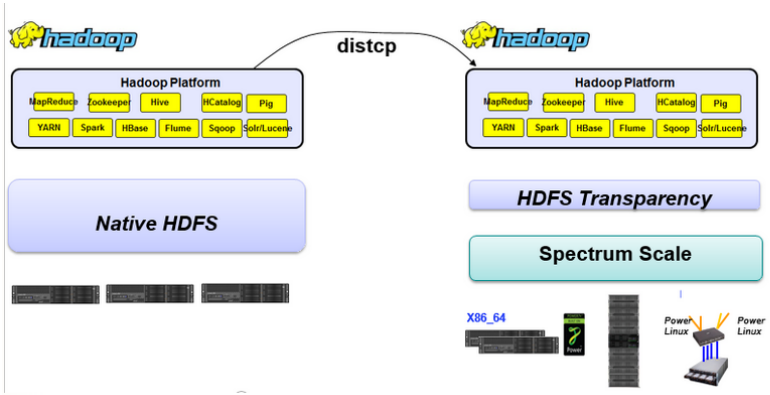

IBM Spectrum Virtualize v7.8

New Entry All-Flash - IBM V5030F

Q3

- Two Nodes: each with 2x 6-Core CPU and 32 GB Cache (total 64GB)
- Cluster
- SSD: 400GB-800GB-1.6TB-3.2TB
- RI-SSD: 1.92 TB und 3.84 TB RI-SSD (7TB und 15TB SoD)
- External Virtualisierung, RtC, HyperSwap, Encryption intern/extern, Tiering…

IBM V5030F

Midrange All-Flash: IBM V7000F (2076-AF6)

Q3

- Two Nodes- each with 10-Core 2.2 GHz Broadwell and 64GB Cache (total 128GB)
- HW assisted Compression Acceleration
- SW like V7000:
  RtC, Externe Virtualisierung, Tiering, Mirroring, Hyperswap, Snapshots, Encryption etc.

- SSD: 400GB-800GB-1.6TB-3.2TB
- RI-SSD: 1.92 TB and 3.84 TB RI-SSD (7TB und 15TB SoD)

V7000 AFA
2PB

IBM Spectrum Virtualize as Software

Q3

- **R1 als Bare Metal only**

- Rel 1. support for Lenovo x3650 M5
- **Single or dual socket Intel CPU with 32-64GB RAM**
- **2 Boot Drives**
- **Emulex 16 Gbs FC HBA and Intel 10 Gbps iSCSI HBA**

- **SoD für Cisco, Supermicro, HP**
- **SoD für Hypervisor oder Container support**
- **SoD für iSCSI Networking**

High Density Expansion for Spectrum Virtualize

New

- **92 x drives in 5U**
- 2.5” and 3.5” **SSD-HDDs** supported
  - **920TB with 10TB NL-SAS HDDs**
  - **1380TB with 15TB SSD**
- 12Gb SAS connection
- Storwize V5000 Gen2, V7000 Gen2/Gen2+, V9000, SVC (DH8 and SV1) and VersaStack w/ v7.8
- Intermix with 2U enclosures in same system
- GA: 9.12

IBM Storwize V5000 Scalability Improvements

New

System	V5010	V5020	V5030
SAS chains (per controller)	1	1	2
Dense expansions	4	4	8
Drives/controller	392	392	760
Drives/cluster	392	392	1,056
Max cluster capacity, NL HDDs (raw)	3.8 PB	3.8 PB	10.3 PB
Max cluster capacity, RI SSDs (raw)	5.7 PB	5.7 PB	15.4 PB

IBM V5010
with 5.7PB SSD

v7.8: Fully integrated Management of RI SSDs and and Easy Tier

New

- **New terminology recognizes different drive types**
  - **Tier0 Flash** (formerly SAS SSD)
  - **Tier1 Flash** (new)
  - **Tier2 HDD** (formerly SAS HDD)
  - **Tier3 NearLine** (formerly SAS NearLine HDD)
  - **Easy Tier new warm promotion feature.**

- **Easy Tier and drive configuration now recognize different drive types**
- **Easy Tier supports up to 3 tiers**

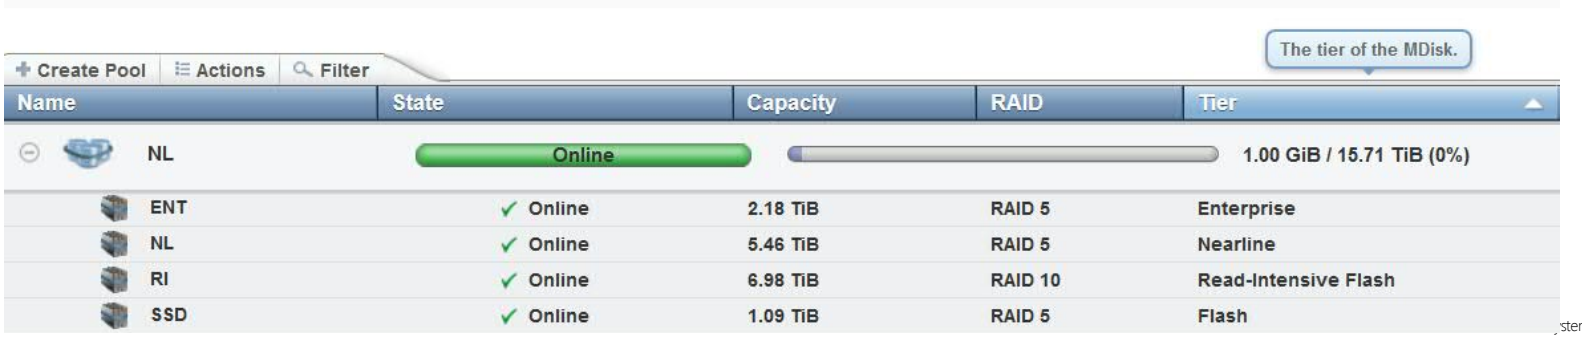

IBM Spectrum Virtualize Transparent Cloud Tiering (TCT)

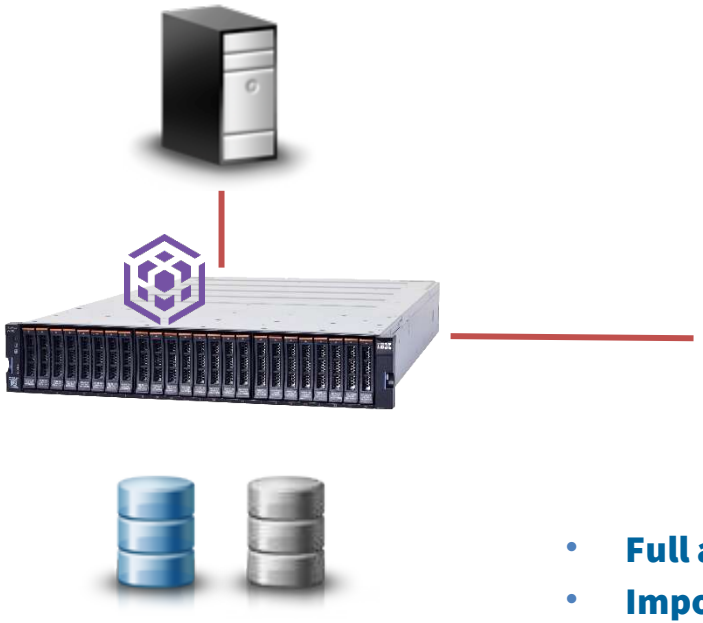

Almost 400 storage systems from IBM and others

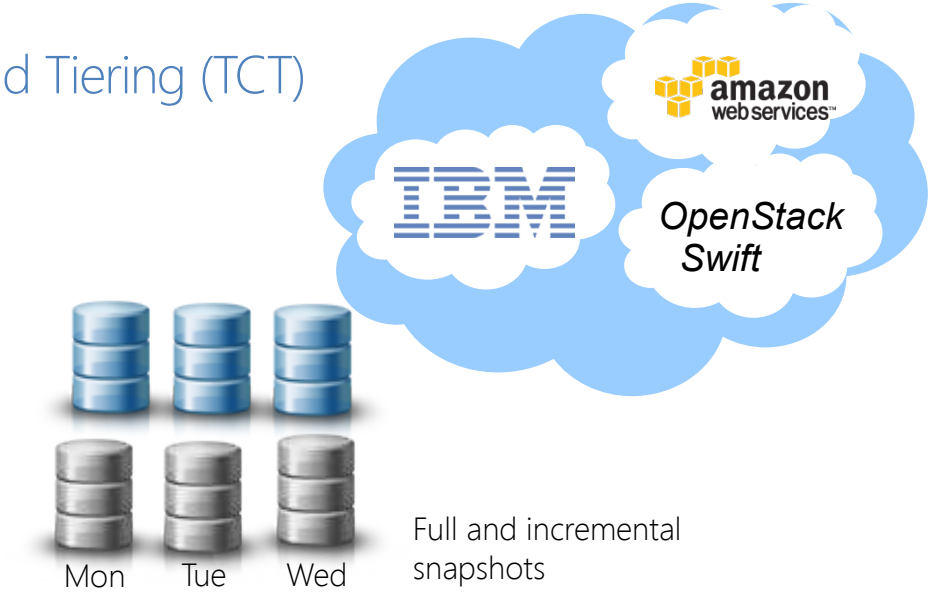

Full and incremental snapshots

- **Full and incremental** restore from cloud snapshots
- **Import and export** of volumes between systems
- **Supported cloud providers:** IBM Softlayer, Openstack Swift, Amazon S3
- V7000 G2/G2+, V9000, SVC, VersaStack, S.Virtualize SW
- GA: 9.12

Your Choice: IBM Cloud based on IBM COS

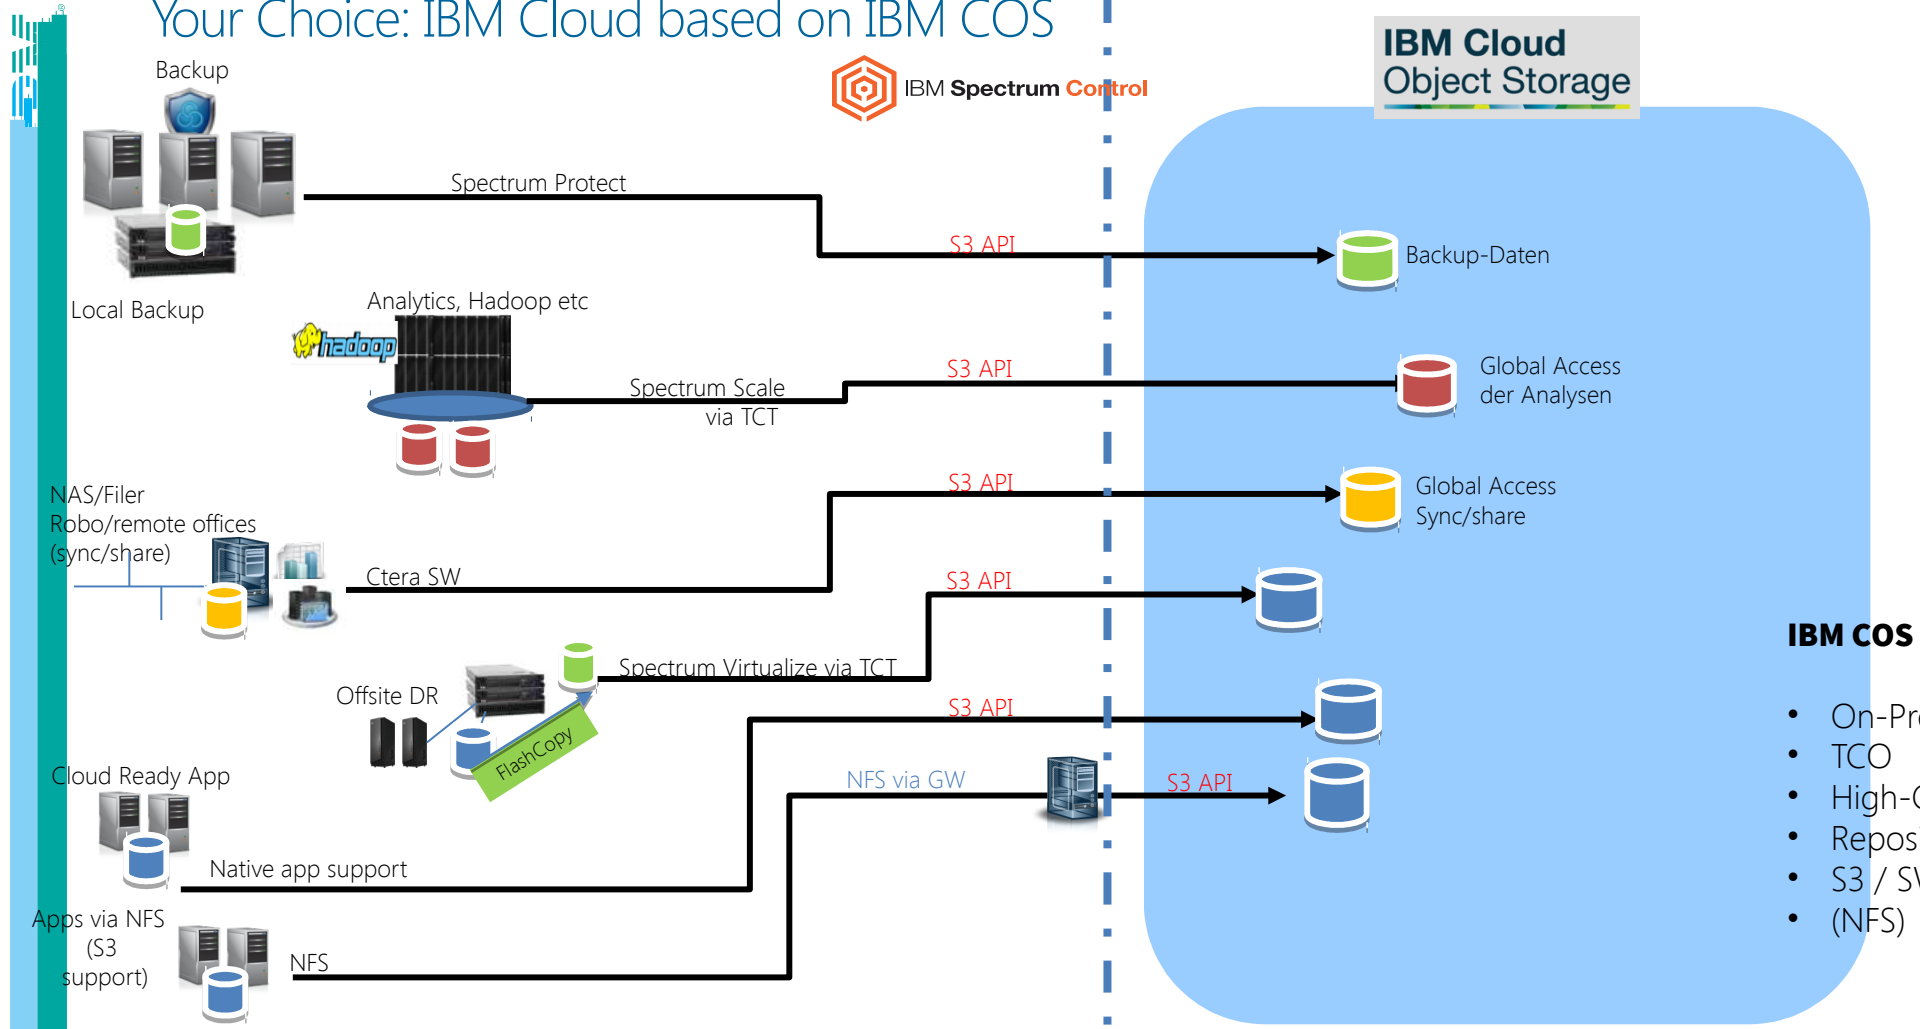

IBM DS8880 AFA
New Flash- Enclosure

IBM DS8888 All-Flash for System-z:

New

High-Performance Flash Enclosure Gen2

- ✓ up to 2x performance improvement per enclosure[1]
- ✓ 6 times more capacity per enclosure for storage
- ✓ using RAID 6 as default

DS8880 with new AFA Enclosure

Thin Provisioning space release for z System volumes

- ✓ space release functionality for Count Key Data (CKD) volumes
- ✓ Improves storage efficiency by releasing free extents and making them available making them available for reuse

IBM DS8880 High Performance Flash Enclosure Gen2

New

		HPFE		HPFE Gen2	
Flash card options		400 / 800 GB		400 / 800 / 1600 / 3200 GB	
Maximum raw capacity per enclosure		24 TB		153.6 TB **Over 6x**	
RAID protection		RAID 5 and 10		RAID 5, 6 and 10	
Performance		IOPS	Throughput	IOPS	Throughput
	Read	340,000	3.8 GB/s	650,000 **+90%**	8 GB/s **+110%**
	Write	200,000	2.7 GB/s	300,000 **+50%**	5 GB/s **+85%**

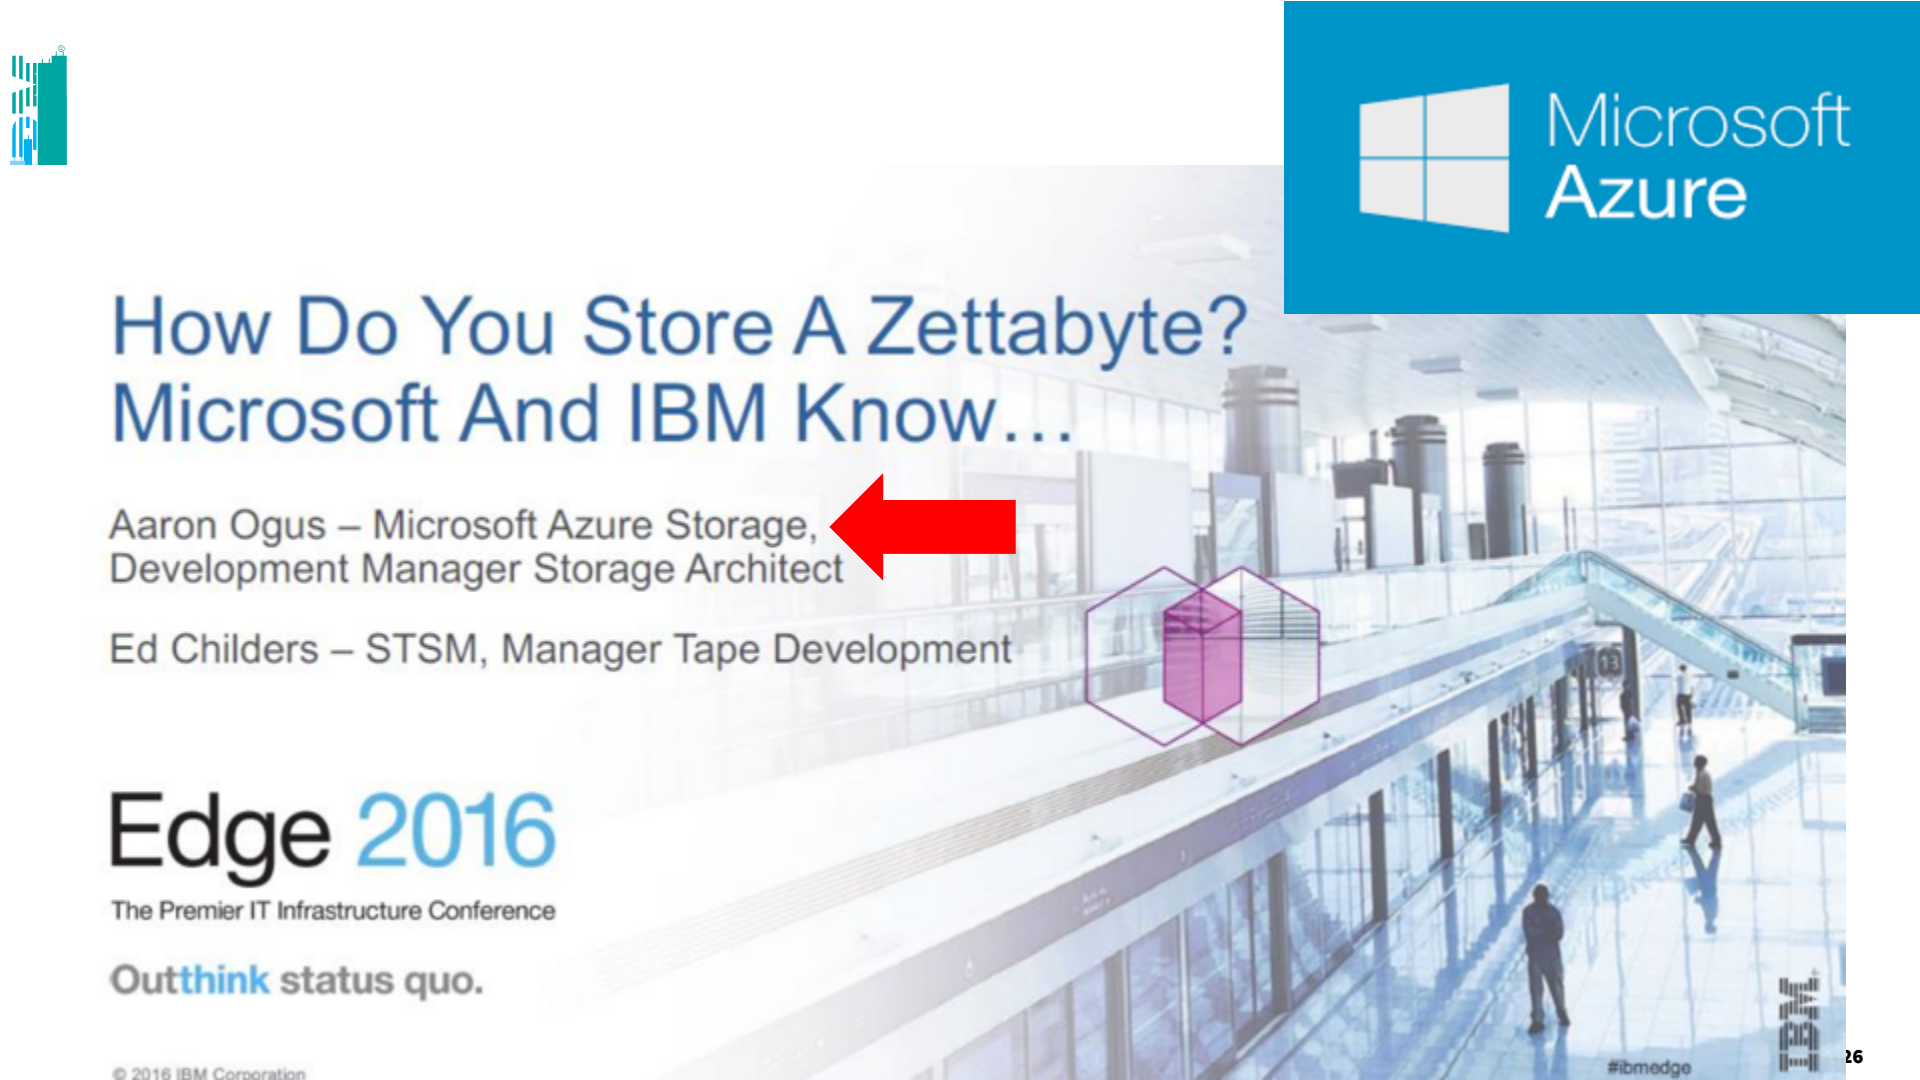
Microsoft
Azure
How Do You Store A Zettabyte?
Microsoft And IBM Know…
Aaron Ogus – Microsoft Azure Storage,
Development Manager Storage Architect
Ed Childers – STSM, Manager Tape Development
Edge 2016
The Premier IT Infrastructure Conference
Outthink status quo.
© 2016 IBM Corporation
#ibmedge

IBM Spectrum Copy Data Management

IBM
Spectrum
Copy Data
Management

IBM CDM: Complete Copy Automation

Create Copies

Snapshots, Replicas, Clones

Track Copies

Catalog

Refresh Copies

Automated DR and Dev-Test refresh

Use Copies

Map LUNs, Spin Up Systems

Transition Copies

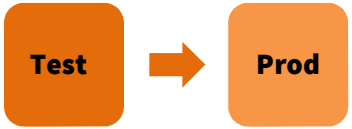

Test to Production

Delete Copies

Clean Up

IT Modernization through "In Place" Copy Data Management

Your Infrastructure

IBM FlashSystem A9000

IBM FlashSystem A9000R

IBM FlashSystem V9000

IBM Storwize V7000F

IBM Storwize V5000F

Also supports:

- **Storwize HDD systems**
- **SAN Volume Controller**
- **Spectrum Virtualize**
- **Spectrum Accelerate XIV**
- **Storage Arrays**
- **VersaStack**
- **EMC VNX and Unity**
- **NetAPP**

IBM Spectrum Copy Data Management

Software-Defined

Copy Data Management Platform

Catalog
- Discover
- Search

Automate
- SLA compliance
- Policy-based

Transform
- Cloud integrated
- DevOps enabled

LEVERAGE

Use Cases

- **Protection and Disaster Recovery**
- **DevOps, Test/Dev**
- **Automated Copy Management**
- **Hybrid Cloud**

Applications

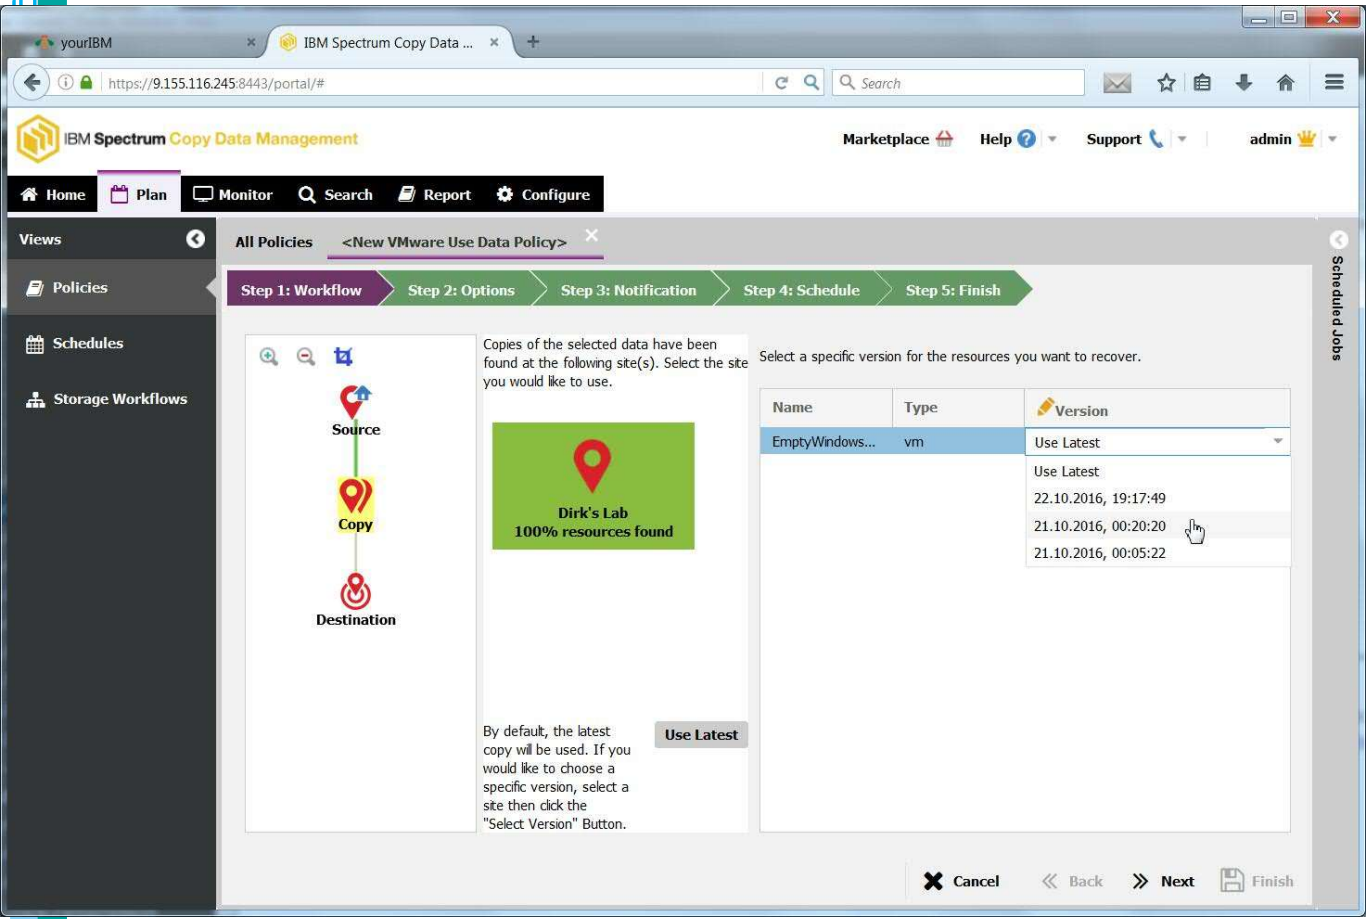

- VM admin connect to Spectrum CDM portal
- VM admin can restore and mount snapshots without the storage admin.
- Mounting is possible even to a new source/destination.

High-Performance Computing

IBM Spectrum LSF

High-Performance Analytics

IBM Spectrum Symphony

New-Gen Workloads

IBM Spectrum Conductor

IBM Spectrum Conductor

(SDI: BigData, Hadoop, Spark, Container)

Compute and Data Silos to Analyze Data Is Costly and Inefficient

New modern scale-out applications and frameworks …

…leads to costly, complex, siloed, under-utilized infrastructure and replicated data

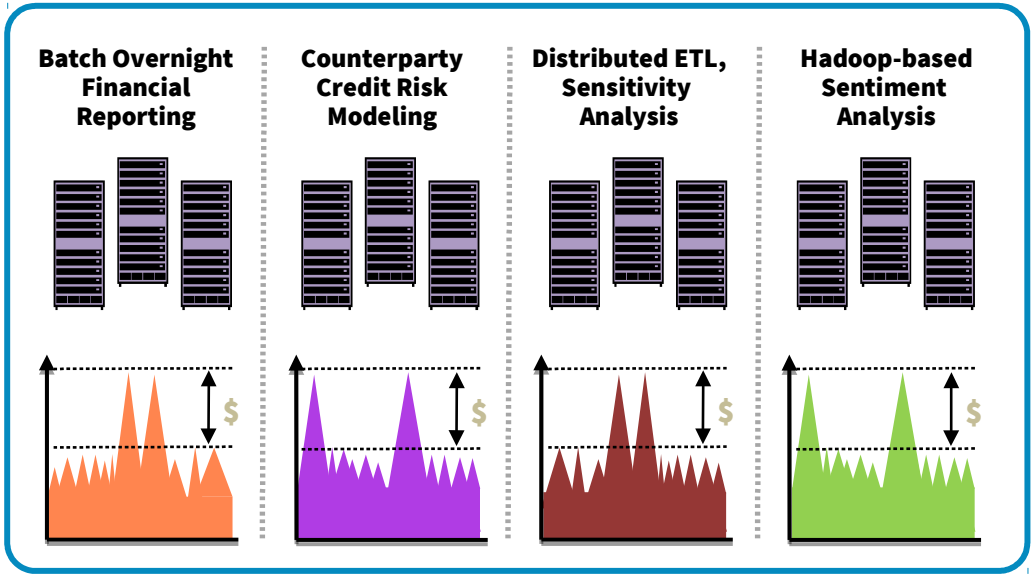

Low Utilization + Poor Performance = Higher cost

Software Defined Infrastructure for Next Generation Workloads

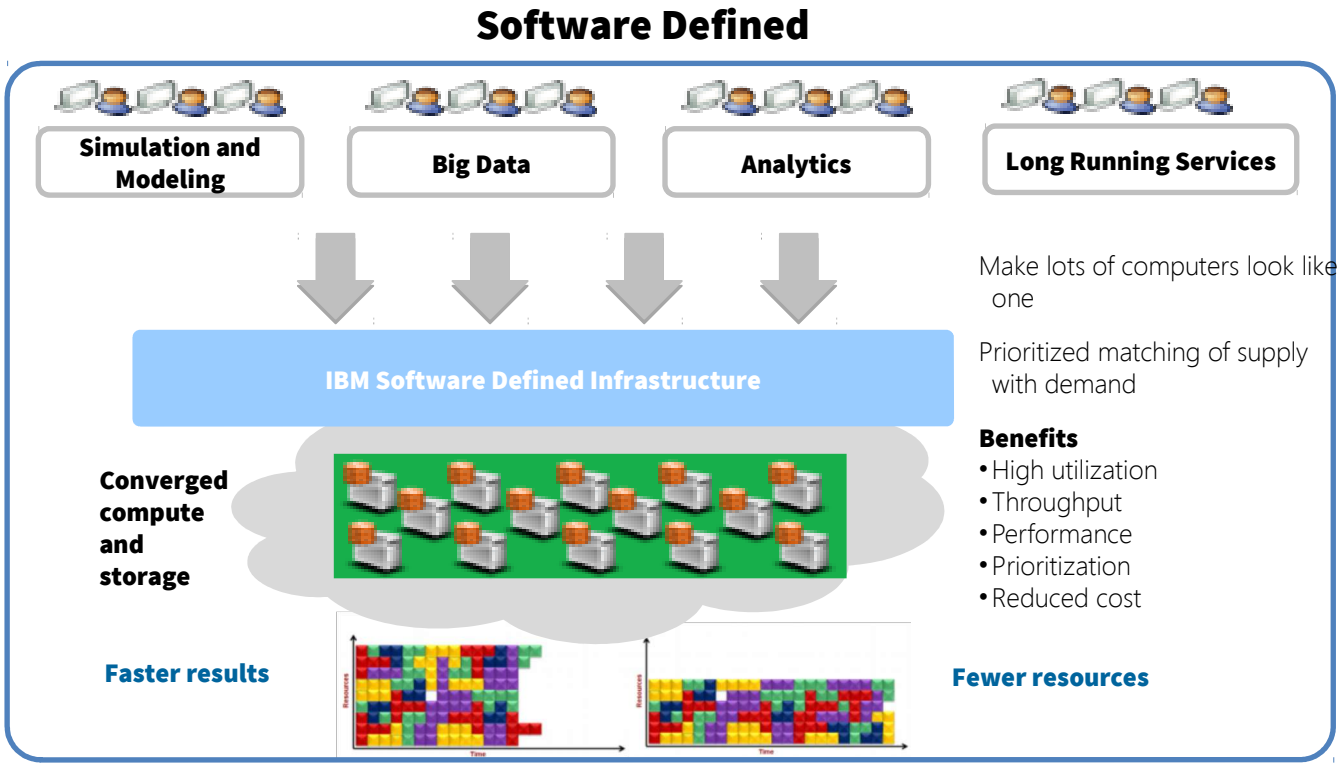

Make lots of computers look like one

Prioritized matching of supply with demand

Benefits

- High utilization
- Throughput
- Performance
- Prioritization
- Reduced cost

IBM Spectrum Conductor

IBM Spectrum Conductor for Containers

- Is a generalized **workload manager** for complex long-running applications on top of EGO (Enterprise Grid Orchestrator)
- Different applications have different **computing requirements, that can change dynamically**
- Some of the applications may be **containerized, others not containerized**
- Different applications are able to **share resources** in an efficient way to increase utilization
- Automated and **intelligent placement of application** instances
- Support for **scaling, high availability,** and dependencies of

IBM Spectrum Computing

IBM Software Defined Infrastructure- Spectrum Conductor

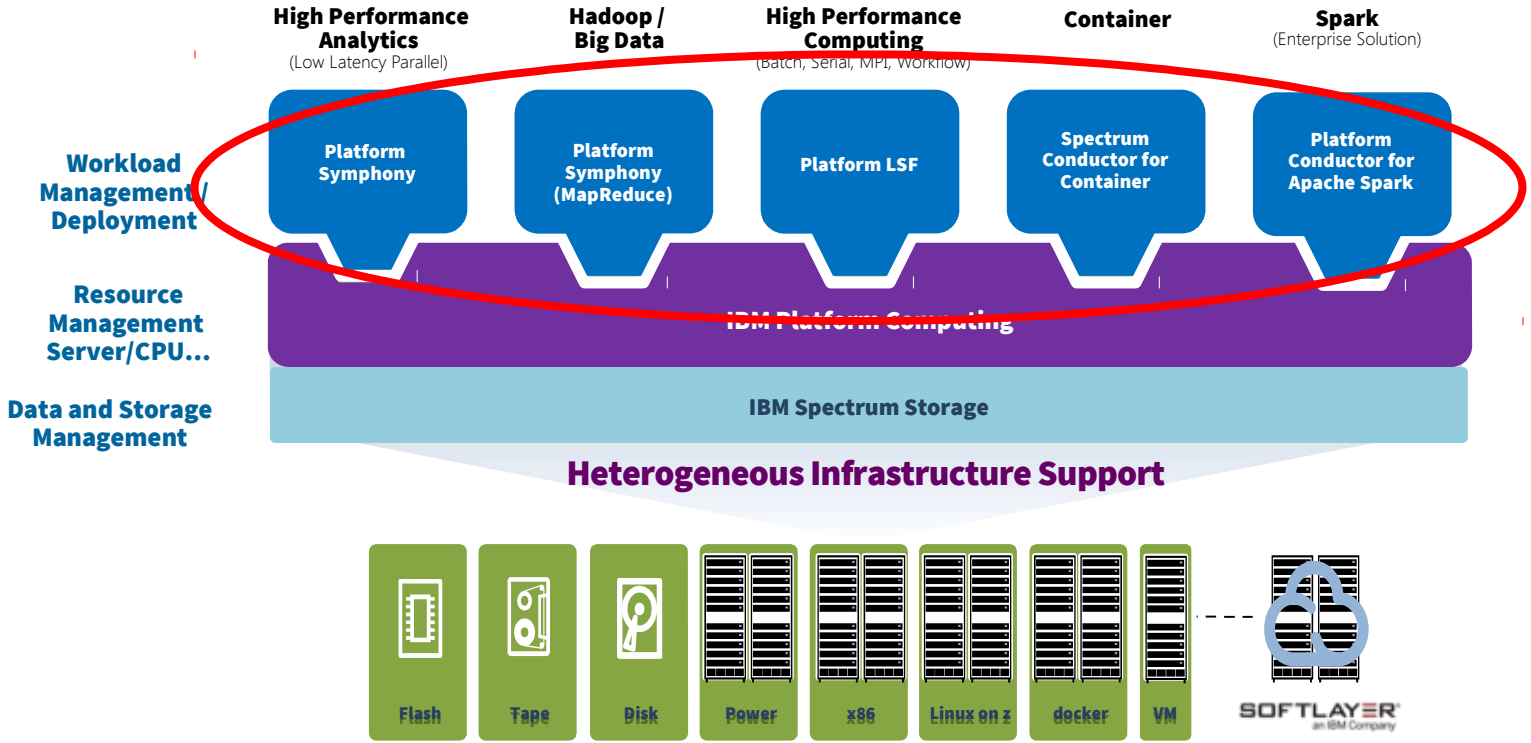

On-premise, On-cloud, Hybrid Infrastructure
(heterogeneous distributed computing and storage environment)

Spectrum Conductor for Container

IBM Spectrum Conductor for Containers

- **Full Lifecycle Management for Container**
- **Container orchestration**
  - Resource management
  - Application life-cycle management/schedule/deployment
  - Scaling, Rolling upgrade, Discovery
- **Container infrastructure**
  - Load balancer
  - Multi-host networking
  - Distributed storage management
  - Image/Software repository management
  - Configure management
  - Logs/Meters/Alerts
  - User/Account management (multiple tenancy & RBAC)
- **Ops management**
  - Installation/upgrade
  - Health check

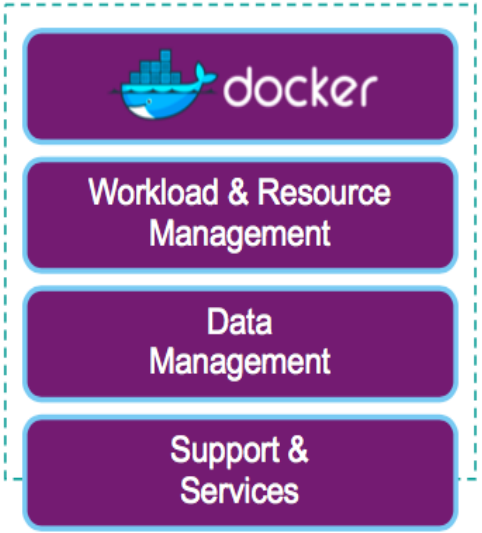

What can Users do?

User can: **Share and Manage the Data over many Applications**

Data volumes are created on demand and owned by LOBs/Individuals

Data is shared between users applications. Seen as a shared filesystem

The desired data volumes are connected to Docker containers

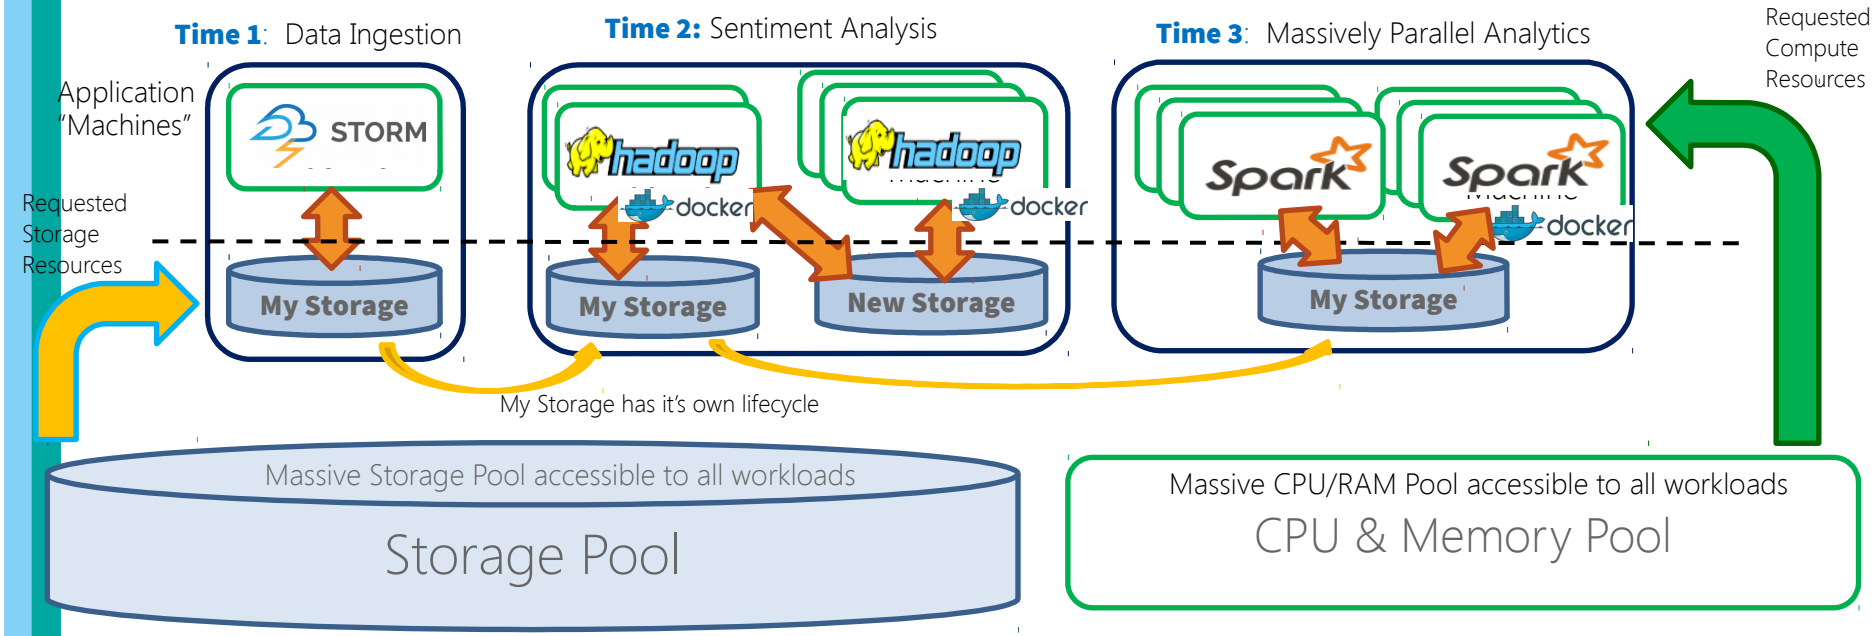


IBM Spectrum Conductor for Containers
admin
Master Host: 192.168.122.149
Dashboard
Applications
Policies
Images
Nodes
App Center
Storage
System
System Overview
Nodes 2
100%
Active
Active 2
Inactive 0
Shared Storage 0 B
0%
Available
Available 0 B
Used 0 B
Released 0 B
Failed 0 B
Applications 2
100%
Healthy
Healthy 2
Unhealthy 0
Resource Allocation
CPU
13%
1 of 8 Shares
Memory
6%
796 MiB of 14 GiB
Local Disk
0%
0 B of 171 GiB
100%
50%
0%
-60s
-30s
0

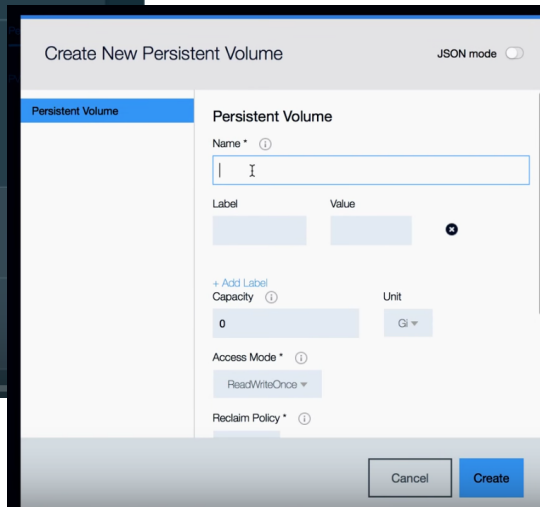
Create New Persistent Volume
JSON mode
Persistent Volume
Persistent Volume
Name *
Label
Value
+ Add Label
Capacity
Unit
0
Gi
Access Mode *
ReadWriteOnce
Reclaim Policy *
Cancel
Create