Enabling the journey to the Hybrid Multi Cloud with Storage Technology

Andy Walls

IBM Fellow CTO Flash Storage Old Geezer with lots of experience



Critical areas of Storage Design Focus for IBM

Hybrid Multi Cloud and Al

- Virtualize in the Public Cloud
- Spectrum Scale
- Containerize



Differentiated Flash

- High and consistent performing Data Reduction
- QLC everywhere
- Accelerators for software stacks
- SCM Exploitation
- NVMe End to End and fabric leadership



Simple, Simple, Simple

- Determine the right options instead of asking
- Simple to manage
- Simple everything

AI Empowered Storage

- IBM Storage Insights
- Best Practice Insights
- Troubleshoot tough network issues



It is a Hybrid Multicloud World 85% of companies operate a Hybrid multicloud environment	Companies average	Migrate from public cloud Security, Performance, Cost, Cost, Control
91% of public cloud adopters use internal private cloud	80% of companies moved applications or data from Public Cloud in 2018	98% of companies will use Hybrid Multiclouds in three years

Hybrid and multicloud are the new normal



Key customer value drivers that are driving cloud adoption in the market



1. BUSINESS TRANSFORMATION: SHIFT CLOUD STRATEGY FROM 'UTILIZATION & COST IMPROVEMENTS' TO 'DELIVERING STRATEGIC BUSINESS

VALUE'



4. BIG DATA: SHIFT & SCALE BIG DATA FROM ON-PREMISE TO CLOUD TO EXPLOIT HIGH VALUE TECHNOLOGIES LIKE AI, IOT & BLOCKCHAIN



2. Application modernization: ADD CLOUD SERVICES TO EXISTING APPS



5. SECURITY: THE MODERN DAY BUSINESS WORKPLACE CREATES A NEW SET OF SECURITY CONSTRAINTS FOR THE HERITAGE DATA CENTER



3. IT OPTIMIZATION: OPTIMIZE THE IT THAT CUSTOMERS OWN & RUN

၀, ၇ ၀, စ စ ၀ ၀

6. ECOSYSTEM: GAIN ACCESS TO ECOSYSTEMS WHERE DATA, APPS, SERVICES, AND IP

CAN BE EXPANDED, LEVERAGED, AND SHARED

Storage challenges

Old

65% Cut infrastructure costs

59% Avoid vendor lock-in...

...Avoid data stored with different vendors' technologies in silos across the enterprise.

59% Improve latency

Source: IBM Institute for Business Value, 2018

New

Storage connectivity in the

hybrid multicloud era matters. Storage must connect to a hybrid multicloud world

Storage location in hybrid

multicloud matters depending on data security, regulatory environments, performance, and availability requirements.

Speed of adoption matters.

NVMe, storage class memory (SCM), software-defined storage (SDS)

© Copyright IBM Corporation 2019











smart is open.

We've now combined the industry expertise of IBM with the open source leadership of Red Hat. Let's unlock the world's potential. Let's put smart to work.[™]

ibm.com/redhat

open is smart.

We've now brought together the hybrid cloud portfolio of Red Hat with the scale and security of IBM. Let's unlock the world's potential. Let's put smart to work."

ibm.com/redhat

With red hat, IBM can now deliver the next chapter of digital transformation, innovation, and growth with a consistent stack for open hybrid multicloud with containers as its underpinning

Containers are a fundamental technology for clients moving to the cloud

Containers are used to create new cloudnative apps and to modernize existing ones Ease container management and container orchestration with kubernetes

It's fundamental to the experience

Differentiated Flash

- NVMe Based
- Right mix of software and hardware accelerators
- Sophisticated flash controllers
- SCM Enablement
- Connect with the cloud
- High Availability





A Simple, Robust NVMe Platform

- 2 Rack Units
- Up to 378TB of Useable Storage
- Up to 865TB of Effective storage with hardware compression
- Up to ~1.8 PB (or so) after software DR
- Full suite of storage services
- SCM Enabled
- Cloud connections provided

Custom Differentiated Flash Controller





9150 NVMe based dual controller Storage System



The FS9100



- Ths software stack is Spectrum Virtualize
- But nothing is virtualized!
- Each Controller connected
 over NVMe to the Flash
- The controllers are connected over PCIe to each other
- The Flash does the Compression in band
- This a fully self contained All Flash Array

Building a World Class Flash Controller – Starts with the Flash!

- Develop characterization platforms
- Test tens of thousands of flash blocks
- Work across the industry and partners to stay ahead of and help establish trends
- Invent techniques to extend flash endurance
- Use flash "knobs" to extend performance
- Validate with application use models to reduce write amplification in <u>real life use cases.</u>
- Teams of researchers, semiconductor and flash
 experts across the Industry



IBM's Write Amplification Reduction

- Comparing write amplification reduction of
- Circular Buffer GC
- Greedy Window GC (i.e., always selecting the block with the most invalid pages for GC)
- FlashSystem N-bin GC, no heat segregation
- FlashSystem N-bin GC, with heat segregation



Inline, At Speed, Hardware Compression

The FlashSystem data compression/decompression algorithm is a Modified Dynamic GZIP algorithm

- Takes advantage of already existing LSA mapping
- There is less data to transfer in backend making up for small added latency (7uS)
- Decompression done in line with hardly any extra latency
- Data is decompressed along side compression and checked bit by bit. Never any corruption!
- Data protection (ECC) is implemented on top of compressed data.
- Allows garbage collection and other background data transactions to operate on compressed data

Compression and Decompression completely transparent above the Flash module except for management of space.





Other Key Technologies in Flash Core

- ECC that never has to reread (Hard decision)
- Calibration algorithms that adjust voltage thresholds according to Characterization findings.
- Health Binning
- Write Heat Segregation. (Read Heat coming)
- Smart Data Placement
- Synergies and accelerators for the Software stack

Use Cases for IBM Spectrum Virtualize for Public Cloud





Optimize Public Cloud Block Storage Extend On Premises to Hybrid MultiCloud

Business Continuity on Public Cloud

Protection of Data in Public Cloud

- Lower Cost, Improve Performance of native Public Cloud IaaS
- Thin-Provisioned Volumes, Space Efficient Snapshots, AI based Auto Tiering
- Increase Scalability of Public Cloud Storage for Enterprise Apps

- Add cloud capabilities to existing Storage on Prem
- Temporary or permanent data migration to/from Public Clouds, and between Public Cloud providers.
- Move data to cloud resources, such as containers, VMs.
- Consistent Management

- Create a DR site in the public cloud
- Synchronize local storage data with sync or async storage replication
- Protect on prem data for virtualized, containerized, or bare metal applications

- For workloads moved to cloud
- Use Sync or Async Mirror to protect cloud data center deployments
- Supported within and between disparate Cloud Provider data centers

IBM Spectrum Virtualize for Public Cloud on AWS



####