



**IBM Spectrum Scale**

**The next step...**


Spectrum Scale

# Please Note

IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion.

Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision.

The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.



# The Cloud Directive

## Let's do it!

- On demand pricing
  - Someone else runs it
  - Perceived lower cost
- 

## Oh, maybe not...

- I need high IO performance
- There is a High Cost of moving data
- How to get data to and from the cloud
- How to get your applications using object data.

## A two pronged approach



To the cloud

Internal Cloudification

## So what about Spectrum Scale ?



Public Cloud Platform

Private Cloudification Platform

# Spectrum Scale cloud platform features

## Existing Features

- Supports standard Power and x86 servers
- Supports any storage
- Extreme scalability
- High performance IO
- OpenStack Cinder, Swift and Manilla
- Hadoop Support for any storage architecture

## New Features

- Transparent Cloud Tiering
- Remote Boot Support (iSCSI)
- RESTful API
- OmniPath

# Public Cloud Platform

## On the cloud

- Use Cloud storage infrastructure
- HPC in the cloud

## Being the cloud

- Be cloud storage infrastructure
- Support VM's
- High performance HPC in the cloud



# Internal Cloudification

## What does it mean?

- Reduce costs
- Simplify storage management
- Provide high performance IO
  
- Use SMB and NFS for backwards compatibility
- Heavy hitting nodes running Spectrum Scale software
  
- Shared disk analytics
- Large capacity content store
  
- The more data the more this makes sense

## Why?

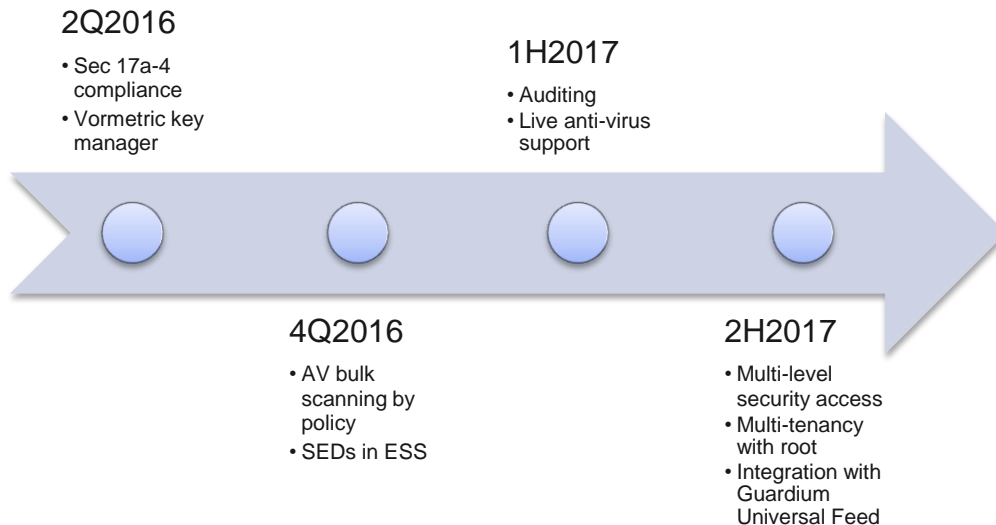
- NAS appliances not hitting the mark
  - Capacity
  - Performance

# Market Domain: Big Data and Analytics

- 2Q2016
  - Experience Documents (spread over time)
- 4Q2016
  - Support noSQL databases (MongoDB support plus others)
- 1H2017:
  - Scale-out infrastructure optimization
- 2H2017
  - Alluxia (formerly Tachyon)
  - Apache Kudu integration
  - Additional experience documents

# Technology Domain: Security

- 2Q2016
  - Sec 17a-4 compliance, certified by KPMG
  - Vormetric encryption key manager in addition to ISKLM
  - Documentation: how to configure AV bulk scan
- 4Q2016
  - AV bulk scan by policy
  - SED support in ESS
- 2017
  - Auditing
  - Live anti-virus support
  - Multi-level security access: reduce need for remote root
  - Multi-tenancy with root access permitted
  - Protocol node support for more than one Active Directory



# Technology Domain: Client Experience

- Phase 1 2Q2016:
  - Extensive monitoring of nodes: understand and resolve performance issues
  - Improved / streamlined gathering of diagnostics
  - Core performance metrics for cloud/object storage tier
- Phase 2 4Q2016:
  - Comprehensive monitoring across more areas of Scale include AFM, system health, and cloud/object tier
  - Capacity, performance and event information always in view
  - Navigation between related assets to simplify workflows
- Beyond:
  - Simplified, streamlined workflows for faster time to value
  - Unified management views across file, object, archive (including cloud) and backup (Spectrum Protect)

# Transparent Cloud Tiering

- Use the cloud or local object store as a secure, reliable, transparent storage tier fully integrated with Scale
- Peer-to-peer: Share data between Scale and cloud openly with read/write on all sides
- Cleversafe, Softlayer Swift, AWS S3, and other Swift / S3 compliant stores
- 2Q2016: Cloud / Object Storage Tier
  - RHEL 7 on x86 only for sync nodes
  - Scalability up to a single group of four transport nodes with failover
  - One filesystem per cluster synced to one object store
- 4Q2016:
  - RHEL 7 on Power for transport nodes
  - Multiple filesystems per cluster, multiple target object stores
  - Import / export on demand between Scale and object stores
  - AFM support
- 2017
  - SigV4 Amazon S3 support for US Government requirements
  - Multi-reader / multi-writer with eventual consistency
  - Additional stores including Amazon Glacier, IBM Ready Archive

# Technology Domain: Technology and Common Currency

- iSCSI support
  - Scale supports iSCSI Target: i.e. Scale presents its storage as iSCSI LUNs to other consumers.
  - This also allows a diskless server to remote boot from a Scale server's filesystem over iSCSI
  - Platforms: We intend to provide this functionality on RHEL operating system only, on both x86 and Power platforms.
- Omni-Path
  - RDMA support

# Technology Domain: Integrations







- Lightweight Events (LWE)
  - Scale can run policy rules as a “real time” response to events such as file open or file changed
  - Scale customers can use the policy engine to drive arbitrary custom events including AV scan, external indexing, audit logging...
- RESTful API
  - Programmatic access to Scale information and configuration
  - First release 4Q2016:
    - Query and manage filesets
    - Query cluster configuration
    - Authentication integrated with AD and LDAP
  - Later releases 1H2017:
    - Manage protocol (CES), performance monitoring, AFM, and other resources


# Technology domain: Linux on z Systems

- Top priorities by z Systems specific integrations
  - Support for emulated Fixed Block Access (FBA) devices as NSDs
    - cheap, fast, efficient
  - Software / Hardware accelerated encryption
    - Using CPACF on the processor
  - Support for heterogeneous Clusters (mixed NSD servers from x-,p- and z)
    - Customer requests to use ESS or AIX NSD server with z clients
  - Hardware accelerated compression using internal PCI compression card
    - HW support is expected first for SUSE operating system
    - Independent of CPU usage
  - Remote cluster mount
  - KVM for IBM z Systems support for Spectrum Scale
    - Spectrum Scale running in a KVM host and guest




# IBM Spectrum Storage Family


 <b>IBM Spectrum Control</b>	Analytics-driven data management to reduce costs by up to 50 percent
 <b>IBM Spectrum Protect</b>	Optimized data protection to reduce backup costs by up to 53 percent
 <b>IBM Spectrum Archive</b>	Fast data retention that reduces TCO for active archive data by up to 90%
 <b>IBM Spectrum Virtualize</b>	Virtualization of mixed environments stores up to 5x more data
 <b>IBM Spectrum Accelerate</b>	Enterprise storage for cloud deployed in minutes instead of months
 <b>IBM Spectrum Scale</b>	High-performance, highly scalable storage for unstructured data



**Any Storage**



**FlashSystem**



**Private, Public or Hybrid Cloud**

