



IBM Spectrum Storage

**Spectrum Scale Cloud Enablement
SPXXL User Group Meeting Sept 28 2017**

John Lewars and Gautam Shah – Spectrum Scale Development
jlewers@us.ibm.com and gautam@us.ibm.com

- The information presented on these charts does not represent any official commitment in terms of Spectrum Scale future direction or deliverables.
- 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.

Agenda

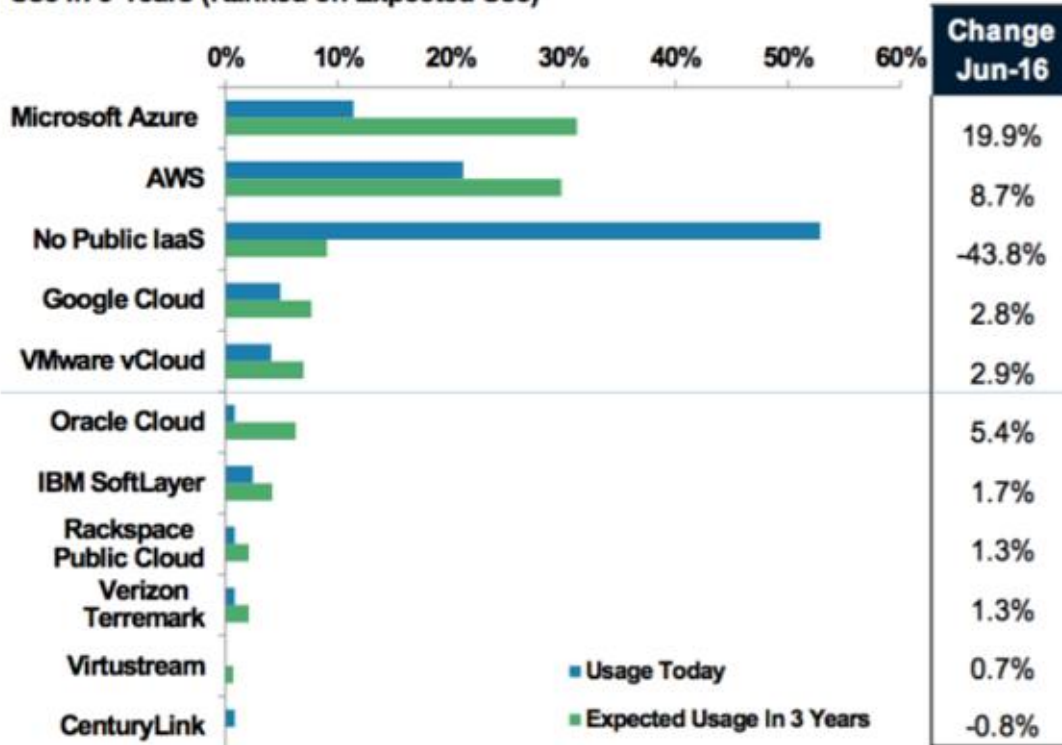
- Public and Private Cloud Enablement in Spectrum Scale
- Spectrum Scale Amazon Quick Start Demo
- Container support
-

Public and Private Cloud Enablement

Discussion on Where Market Data Suggests Cloud Dollar Are Heading

Exhibit 15: Cloud-based IaaS Use Today and Expected Use in Three Years

% of Respondents Using Cloud-based IaaS Today and Expected Use In 3-Years (Ranked on Expected Use)



http://www.networkworld.com/article/3091083/cloud-computing/microsoft-could-overtake-amazon-in-the-cloud-morgan-stanley-survey-finds.html#tk.rss_cloudcomputing

Drivers for Spectrum Scale in Cloud Environments

- **Some benefits potentially offered by Public Cloud Deployments**
 - An increasing number of customers are seeing pressure to move workloads to public clouds (e.g. government mandates requiring public cloud consideration)
 - Offers a way to try things without infrastructure investments (POCs, code evaluation, etc.)
 - Provides operational expense model focus (over capital expense model)
 - On demand provisioning and rapid deployment
 - In an agile manner, the ability to adapt deployment configuration to meet the specific needs of the transient/temporary use cases
 - Lower cost for some workloads (particularly 'bursty' workloads)
- **Other observations:**
 - Several workloads seem to be moving to the cloud (private and public)
 - Some (new) companies have only deployments on the cloud

Spectrum Scale Cloud Strategy

Primary Work Items for Cloud Scrum

1. Enable Spectrum Scale on following Cloud Environments and expand the flexibility of our offerings to better support hybrid environments.
 - a) Support available Spectrum Scale offerings on SoftLayer and support/collaborate regarding new offerings based on Scale
 - b) We've delivered a Trial Offering support of Spectrum Scale on Amazon (AWS) and we intend to continue to enhance this offering, based on customer feedback
 - c) We've prototyped Spectrum Scale support on Azure. We are seeking customer input on Azure demand and evaluating next steps
 - d) Support OpenStack Cloud environments via deliverables such as Cinder & Manila driver support
2. Support and enhance our OpenStack Swift based Object Storage solution which expands the methods by which Spectrum Scale data can be accessed. For example, we've done work to prototype moving our latest Object solution, based on Mitaka, to Occata, and we are evaluating moving to Pike instead of Occata.
3. Investigate better integration, and more flexible support of storage for, virtualized environments (focused on containers)
4. Working on providing automation for application and verification of performance related tunables/configuration options

Questions:

Would you considering running Spectrum Scale on AWS as part of a trial (if the software was made available for free)?

Would you pay for Spectrum Scale software on AWS?

Would you considering running Spectrum Scale on Azure as part of a trial (if the software was made available for free)?

Would you pay for Spectrum Scale software on Azure?

What about any other public cloud offerings?

- IBM has collaborated with Amazon to release a Spectrum Scale Quick Start (launched Sept. 17/17)
<https://aws.amazon.com/quickstart/architecture/ibm-spectrum-scale/>
- This Quick Start follows Best Practices for AWS deployments, e.g., security and high availability
- Focus on simplicity/usability of deployment (e.g. reduce config options) and leverage Amazon features such as Cloud Formation templates and AMIs (Amazon Machine Image)
- Deploy using auto-scaling (for efficiency) then disable auto-scaling in favor of auto-recovery (to allow failed instances to be relaunched with same root device image, same IP address, etc)
- Clusters are created with two Availability Zones (AZ's) in a single region and 2X Spectrum Scale replication is used (one replica per AZ)
- Limited number of configuration choices (e.g. EBS volume types, EC2 instances types, networking options, etc) allows for better customization, more expansive testing, performance measurements, etc
- Amazon provides a rich ecosystem and we pursuing integration of features (e.g. CloudWatch, S3 data movement, data lifecycle management options available)

Documentation regarding use of Spectrum Scale in Cloud environments:

- Spectrum Scale Amazon AWS Quick Start and Deployment Guide:
<https://aws.amazon.com/quickstart/architecture/ibm-spectrum-scale/>
<https://s3.amazonaws.com/quickstart-reference/ibm/spectrum/scale/latest/doc/ibm-spectrum-scale-on-the-aws-cloud.pdf>
- Team has provided Guide for running Spectrum Scale on SoftLayer systems:
<http://www.redbooks.ibm.com/redpapers/pdfs/redp5410.pdf>
- Team has provided OpenStack Redpaper:
<http://www.redbooks.ibm.com/redpapers/pdfs/redp5331.pdf>
- Team has produced Redpaper for Spectrum Scale Unified Object Access :
<http://www.redbooks.ibm.com/abstracts/redp5113.html>
- Knowledge center OpenStack documentation for Spectrum Scale:
https://www.ibm.com/support/knowledgecenter/en/STXKQY_4.2.1/com.ibm.spectrum.scale.v4r21.doc/bl1ins_openstackusecase.htm

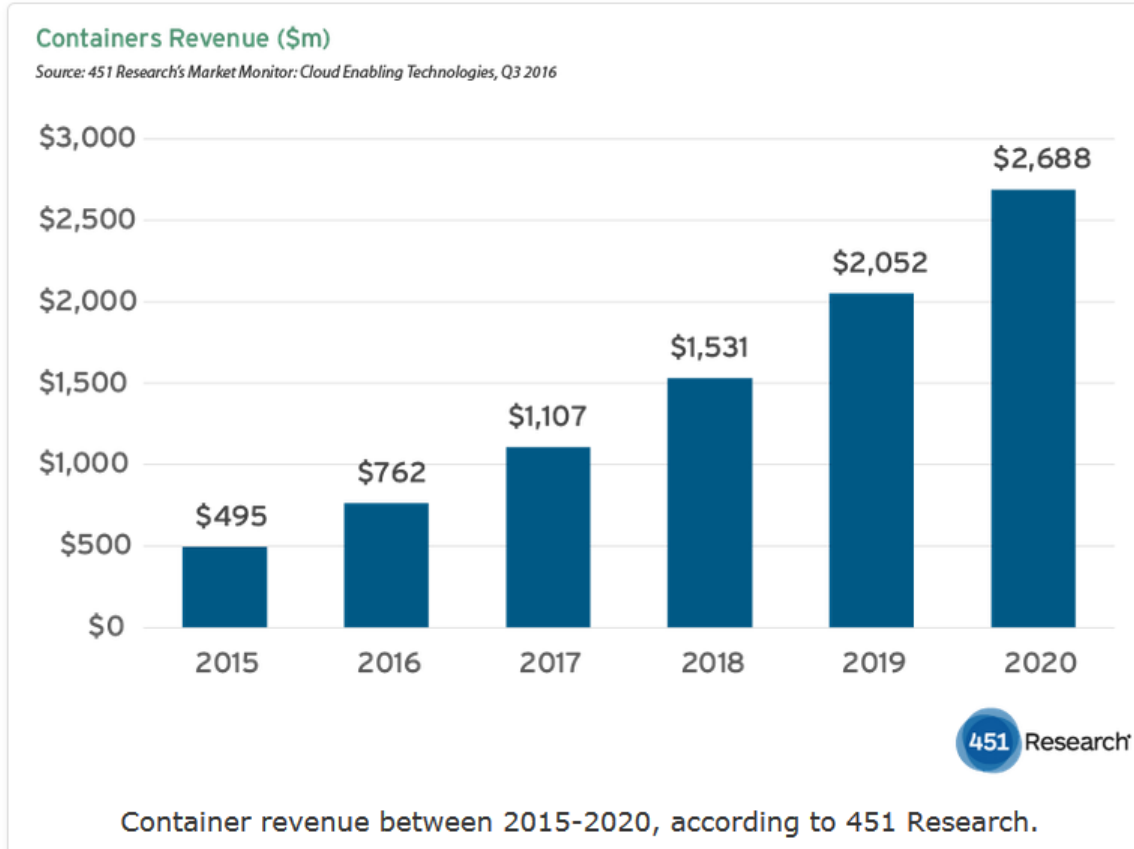
AWS Spectrum Scale QuickStart Demo

Container Support

(Cloud Technology Enablement)

- Containers growth predicted to be at high rate:

<http://www.itworldcanada.com/article/application-container-market-small-but-mighty-451-research-says/389916>



Containers

A large yellow shipping container is being lifted by a yellow tractor in an outdoor industrial setting. The container is suspended in the air, and the tractor is positioned below it. The background shows a cloudy sky and some industrial structures.

Self Contained (Portable)

Data Access Isolation (Multi-Tenancy)

Data Management

High Performance Data Access

Fast and lightweight

Resource Utilization

Open-Source

Global Repositories

High Level Strategy for Containers


1. Team is has done investigations into using GPFS and NFS (exported GPFS file systems) in a variety of environments, including Kubernetes and Docker Swarm. These investigations have included some testing of running/deploying workloads in container.
2. We are focusing on meeting the requirements of Spectrum Compute's IBM Cloud Private product (which uses Kubernetes), and investigating options for new integration.
3. We're interested in customer feedback on the use case for Spectrum Scale in containers.

Questions regarding container usage:

- Are you using containers and, if so, what kinds of workloads are you running with them?
- What orchestration method(s) (e.g. Kubernetes, Docker Swarm) are you looking at?
- How do you see containers and cloud coming together?
For example, how do you see containers fitting into use cases you may be looking at on Amazon or Azure?

What is Possible with Spectrum Scale and Containers?

- Can we use Spectrum Scale Storage within Docker containers?
 - Yes - Any Spectrum Scale storage (be it native protocol or a Spectrum Scale filesystem that has been exported via NFS) can be made available inside Docker containers
 - We can manually provision, mount, and attach (via the docker '-v' run flag) Spectrum Scale volumes to Docker containers, e.g. '-v /gpfs/projects/project1:/gpfs' or '-v /gpfs/projects/project2:/gpfs'



- Can we automate this process via a plug-in Docker driver support?
 - Yes the hostpath driver supports automating bindmount function for Docker/Kubernetes
 - The Spectrum Scale Research team developed Docker/Kubernetes driver (Ubiquity) with additional function but it is not part of the Spectrum Scale product
- Can we currently containerize Spectrum Scale (i.e., deliver Scale via a container)?
 - No – not with the current kernel extension design of Spectrum Scale. Dockerizing all of Scale would be a significant effort but we could potentially exploit containers in the delivery, e.g. we *could* protocol support or the elements needed for GPL builds

Backup

OpenStack

High Level Strategy for OpenStack

Support OpenStack Cloud environments with Cinder & Manila drivers and investigation of best practices (e.g. multi-tenancy recommendations and accessing Spectrum Scale data from VMs)

Questions:

- 1) Who is using OpenStack now?
- 2) Why are you using it or why are you not using it?
- 3) Are you thinking about using it in the future?
- 4) What else would you like to see from Spectrum Scale in terms of OpenStack support?
- 5) Do you agree that OpenStack has become a private cloud play?

Options for Spectrum Scale in the Cloud

	1a.) Spectrum Scale as a backend file store for IBM cloud services	1b.) Spectrum Scale in managed public cloud deployments "On Cloud"	2.) Spectrum Scale in OpenStack private clouds	3.) Cluster as a Service with Spectrum Scale	4.) File Storage as a Service with Spectrum Scale
	Custom Solution Stack in Public Cloud		Private Cloud	*-as-a-Service	
What is MVP	Storage rich server based recommended architecture, Workload specific tuning guidance		<ul style="list-style-type: none"> - Development and validation of required drivers and integration points - Certification with key distribution (Bluebox/ Mirantis/ RDO) 	<ul style="list-style-type: none"> - Provision a Single tenant, dedicated Spectrum Scale cluster on a public cloud <u>without manual intervention</u> - Ability to upload data using NFS/Object - Run compute jobs on cluster 	<ul style="list-style-type: none"> - Provide file storage as service based on some SLAs in Softlayer/Bluemix - Solve multitenancy problem
Examples/ Details	Scale as a backend for bluemix dashdb/ Spark/ Watson Health services	Prana usecase, or any managed cloud deployment in Softlayer/ Amazon requiring clustered file system	CLIMB, eMedLab, Uvictoria, internal XCloud in CDL	Spectrum Scale Amazon AMI Cluster as a Service on Bluemix/Amazon	<ul style="list-style-type: none"> - File service on bluemix - Architecturally possible to implement Offprem as well as Onprem - Similar to Amazon EFS?
What do we have today	Recommendation on architecture choices (redpaper)		Redpaper for deployment guidance, Cinder driver, Manila driver	Working prototype of Amazon AMI	-

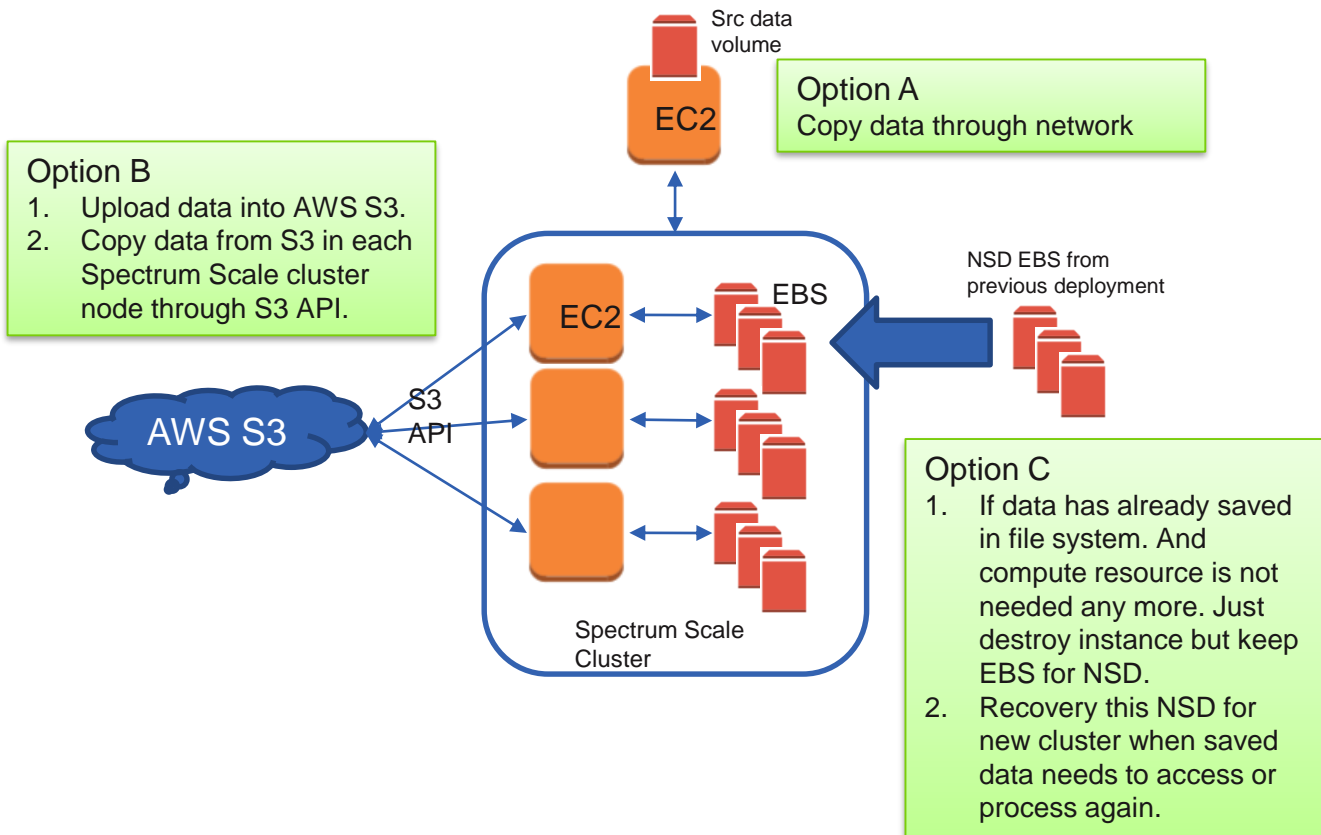
- Team has engaged with Spectrum Scale customers using SoftLayer to support these strategic use cases. Some examples:
 - Supported Prana Animation Studio to do 'cloudbursting' for temporary workload resource needs using approximately 900 bare metal servers running on SoftLayer:
<https://www-03.ibm.com/press/in/en/pressrelease/50966.wss>
 - We have enabled IBM teams to use Spectrum Scale as part of their solutions offered on SoftLayer, e.g.:
 - DashDB database solution needs clustered file system and also leverages Spectrum monitoring functions (e.g. Spectrum Scale detects node failures and facilitates running of callbacks in response to system events):
https://www.ibm.com/support/knowledgecenter/SS6NHC/com.ibm.swg.im.dashdb.doc/admin/local_file_system.html
 - The IBM Platform Computing Cloud service solution is a managed HPC computing service that uses Spectrum Scale:
<http://www.redbooks.ibm.com/redpapers/pdfs/redp5214.pdf>
 - Supported Cloud Data Spark Service's use of Spectrum Scale.
 - We are now exploring options for providing a NFS based storage service (backed by Spectrum Scale) for a variety of workloads (e.g. Bluemix services running in containers).

1. Spectrum Scale on AWS

- a. Testing of Spectrum Scale on AWS instances using EBS volumes
- b. Developing CloudFormation deployment methods for Spectrum Scale clusters
- c. Investigating providing Spectrum Scale support on AWS initially as a trial offering delivered via an AMI, not for production, that we can share with customers via Amazon's AMI sharing facility
- d. Developing guidance for troubleshooting issues with the AWS infrastructure and worked with L2 team to ensure we're ready to support the AWS environment
- e. Collecting baseline scaling performance data for Spectrum Scale running on EC2 instances using EBS volumes

2. Develop a plan for Spectrum Scale support on Azure.

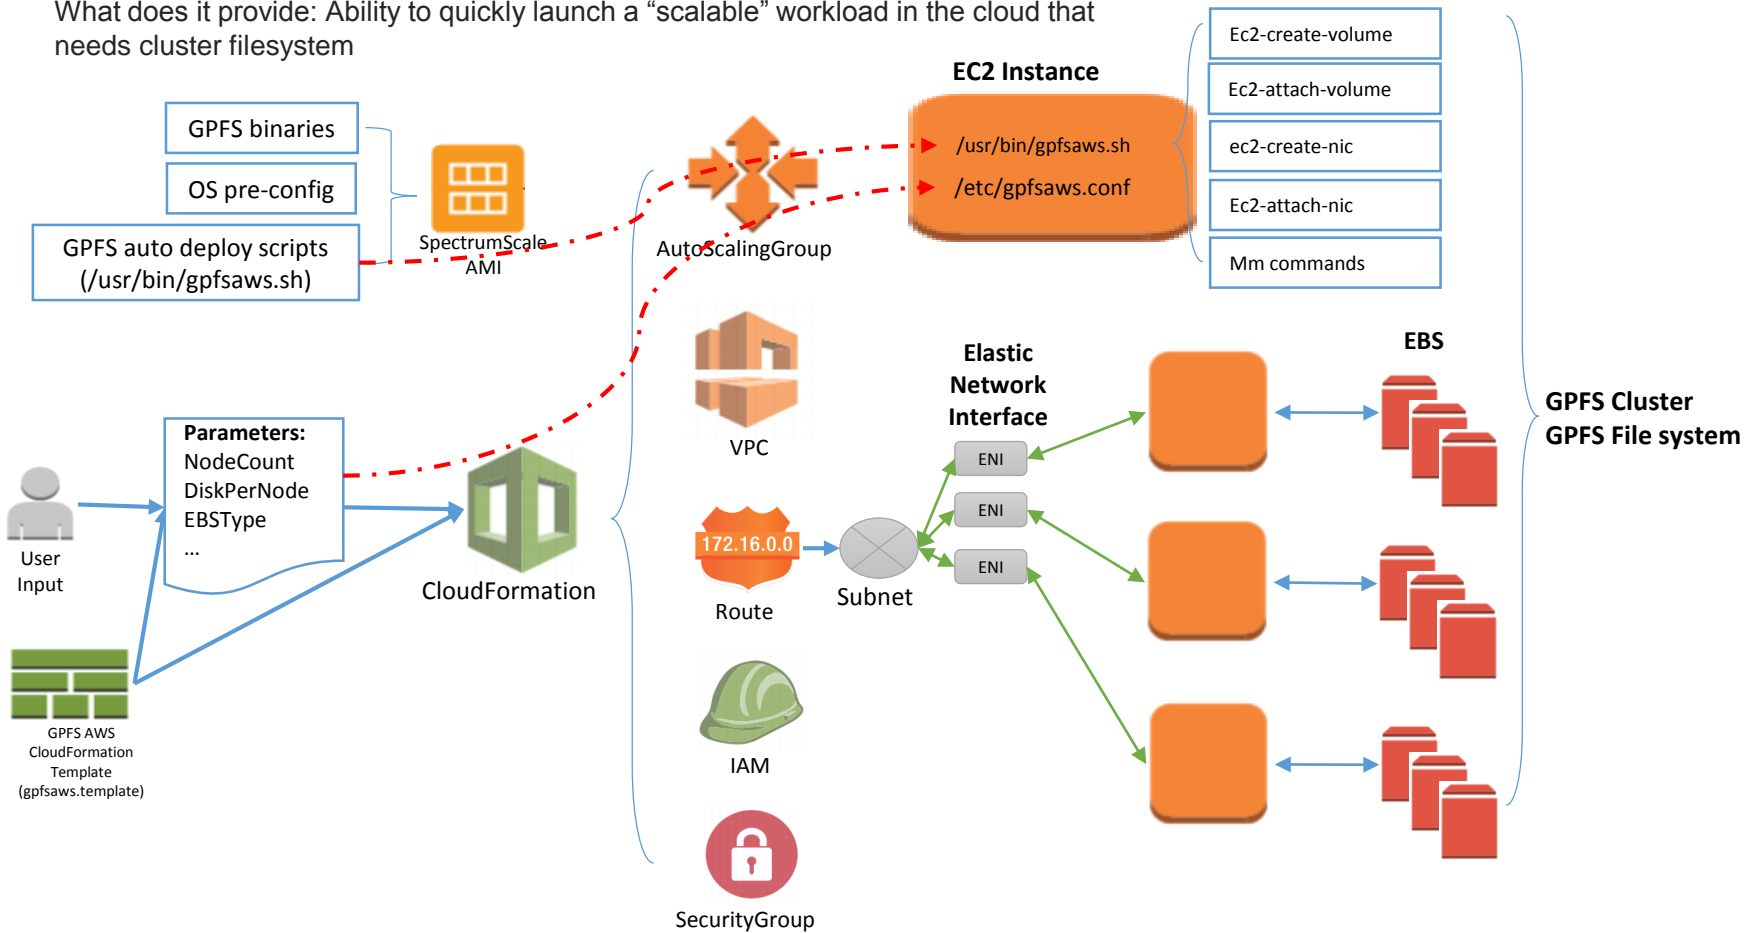
- a. Learning about storage and compute services provided by Azure
- b. Prototyping the use of Spectrum on Azure

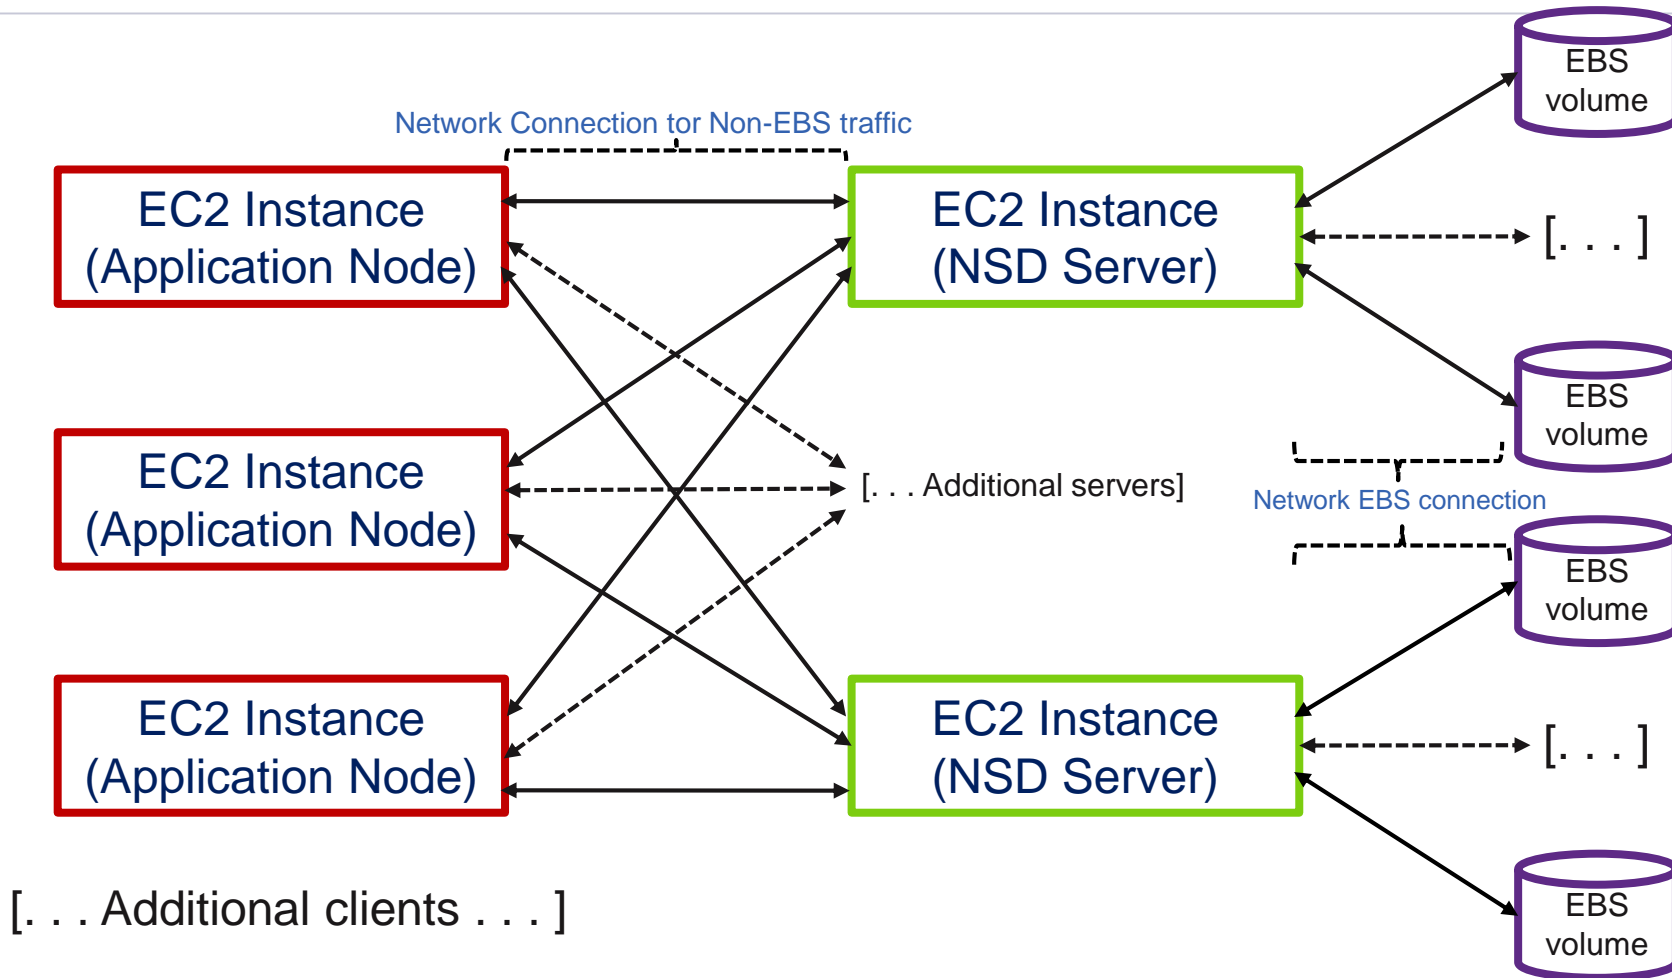


Delivery on Amazon Use CloudFormation

Goals: Deploy a cluster filesystem on Amazon AWS/EC2 using EBS


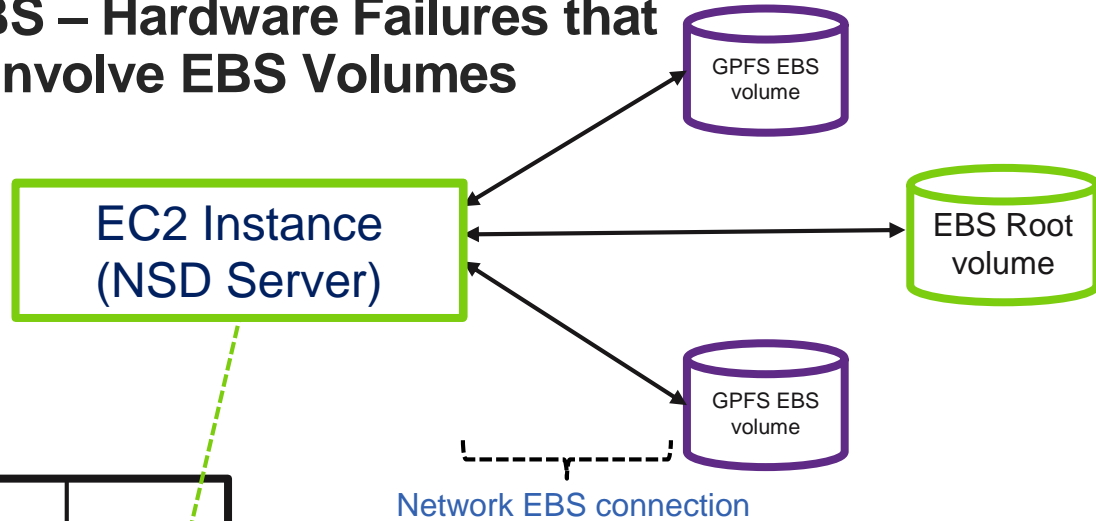
What does it provide: Ability to quickly launch a “scalable” workload in the cloud that needs cluster filesystem





AWS Spectrum Scale Using EC2 Instances

Backed by EBS – Hardware Failures that Do Not Involve EBS Volumes



In the case of a hardware failure leading to termination of an EC2 instance, if the hardware failure is not tied to an EBS volume used by the instance, simply restarting the instance should allow it to be launched on new hardware, using the same EBS root volume and GPFS data disks.

Options for Handling Failures on EBS Volumes Used by Spectrum Scale

1. Run without any Spectrum Scale replication but this assumes a willingness to accept potential data loss if an EBS volume used for GPFS data/metadata fails. (This approach assumes that snapshots/backups of the Scale data are taken to allow recovery from such failures). The SLA Amazon provides for EBS volumes should be carefully reviewed before choosing this option.

From <https://aws.amazon.com/ebs/details/#elasticvolumes>:

Amazon EBS volumes are designed for an annual failure rate (AFR) of between 0.1% - 0.2%, where failure refers to a complete or partial loss of the volume, depending on the size and performance of the volume. This makes EBS volumes 20 times more reliable than typical commodity disk drives, which fail with an AFR of around 4%. For example, if you have 1,000 EBS volumes running for 1 year, you should expect 1 to 2 will have a failure.

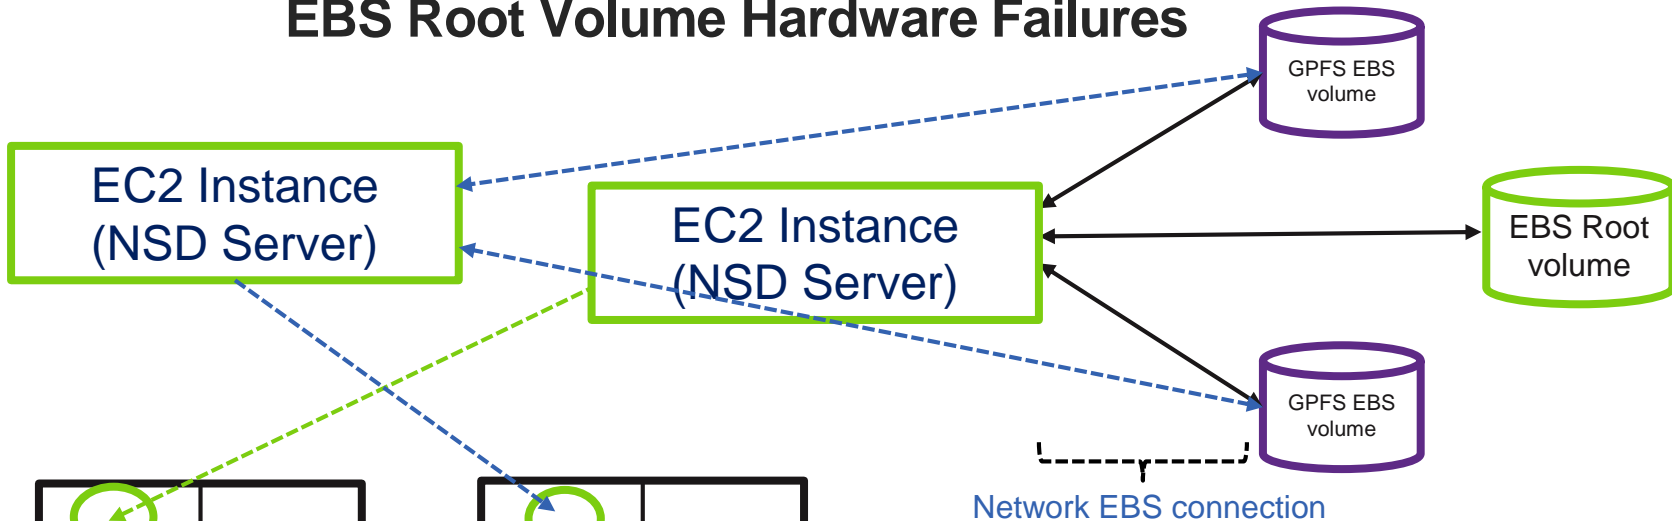
See [Amazon EC2 and EBS Service Level Agreement](#) for full details

2. Use Spectrum Scale replication to replicate all GPFS EBS volumes used by GPFS. Two-way data and metadata replication is suggested, but other options could be considered, depending on data protection requirements (e.g. 2 way data replication and 3-way metadata replication will allow the filesystem to stay mounted, should two separate NSD servers experience failures related to GPFS EBS volumes).

In the case of an EBS volume used by GPFS failing, the 'mmrestripefs -r' command can be used to re-replicate data across the disks and restore the cluster replication levels.

AWS Spectrum Scale Using EC2 Instances Backed by EBS

EBS Root Volume Hardware Failures



Network EBS connection

A hardware failure causing the loss of an EBS root volume will result in the termination of an EC2 instance, thereby preventing that instance from being reactivated. The loss of any GPFS data volumes in such a failure case can be handled as per the previous slide's recommendations. Alternatively, (but this requires unmounting the filesystem), any EBS volumes assigned to the failing instance can be reassigned via `mmchnsd` to existing instances in the cluster or a new instance can be activated to replace the failed instance.

Storage for Virtual Machines	Disks "Persistent block storage for Azure IaaS VMs"	Files "Fully Managed File Shares in the Cloud"	
Unstructured Data	Blobs "Highly scalable, REST based cloud object store"	Data Lake Store "HDFS as a service"	
Structured Data	Tables "Key Value, high scale, auto-scaling NoSQL store"	DocumentDB "NoSQL document database service"	Azure SQL DB "Fully managed database-as-a-service built on SQL"

- ARM = Azure Resource Manager Template
 - Deploy solution as a group, repeatedly along with dependencies, Programmed using JSON
- Resource Manager Group
 - Logical separation of resources, identifiable using Unique name.
- General Purpose Storage Account
 - Namespace to store Blobs and Azure virtual machine disks.
- SKU (Stock Keeping Unit) = Virtual Machine Image (in our case)
 - OS VHD file.
- `copy()`, `copyIndex()`
 - Duplicates resources, Iterative loop function.

All Azure virtual machines have at least two disks;

- A Linux operating system disk
- A temporary disk.

The operating system disk is created from an image, and both the operating system disk and the image are actually virtual hard disks (VHD's) stored in an Azure storage account.

Virtual machines also can have one or more data disks, that are also stored as VHD's.

Operating system disk:

Every virtual machine has one attached operating system disk. It's registered as a SATA drive and is labeled `"/dev/sda"` by default. This disk has a maximum capacity of 1023 gigabytes (GB).

Temporary disk:

The temporary disk is automatically created. On Linux virtual machines, the disk is typically `"/dev/sdb"` and is formatted and mounted to `/mnt/resource` by the Azure Linux Agent.

Data disk:

A data disk is a VHD that's attached to a virtual machine to store application data, or other data you need to keep. Data disks are registered as SCSI drives and are labeled with a letter that you choose. Each data disk has a maximum capacity of 1023 GB. The size of the virtual machine determines how many data disks you can attach to it and the type of storage you can use to host the disks.

About VHD's:

The VHD's used in Azure are ".vhd" files stored as **page blobs** in a standard or premium storage account in Azure.

Block Blobs: (Media Files, Log files)

Block blobs let you upload large blobs efficiently. Block blobs are comprised of blocks, each of which is identified by a block ID. You create or modify a block blob by writing a set of blocks and committing them by their block IDs. Each block can be a different size, up to a maximum of 100 MB (4 MB for requests using REST versions before 2016-05-31), and a block blob can include up to 50,000 blocks.

Ubiquity Github (Released by Research Team)

Not part of the Spectrum Scale product but available via:

Ubiquity Service <https://github.com/ibm/ubiquity>

Ubiquity Docker Plugin <https://github.com/IBM/ubiquity-docker-plugin>

Ubiquity K8s DP and FV <https://github.com/IBM/ubiquity-k8s>

Supported on a best effort basis by the research team

Thanks!