

Spectrum Scale auf Amazon AWS

—

Harald Seipp
STSM, CoE for Cloud Storage
IBM Systems EMEA Storage Competence Center

EMEA Storage Competence Center

The EMEA Storage Competence Center (ESCC) provides **Pre- to Post-Sales Storage services to assist you across IBM's complete Storage portfolio (Hardware / Software / Solutions)** via:

Advanced Technical Skills (ATS)

Skill Transfer, New Product Intro., Solution Enablement, Architectural Guidance

Lab Services (LS)

Skill Enablement, Hands-On-Training, Implementation Assistance, Data Migration & Performance Tuning

Development & Product Test Systems Lab

Proof of Concept (PoC), Hands-on Product Training ...

Product Field Engineering (PFE)

Last level defect support

Client Care

Pro-active Customer Care Management



ESCC@IBM Intranet

http://escs.mainz.de.ibm.com/e_dir/e4uweb.nsf/site.xsp



ESCC@IBM Partnerworld

https://www.ibm.com/partnerworld/page/pw_com_bp_ecampus

Spectrum Scale on Cloud Background

Workloads are moving to the cloud.

Private, Public, Hybrid, Multi-Cloud.

We want to offer support for Scale in **cloud environments** that **our customers are moving to**.

This is part of our hybrid strategy.

There is need to have a **cloud offering** in order to participate in new customer opportunities which are **not starting on-premises**.

Key goal is to deliver **Spectrum Scale's value proposition** on existing Cloud Environments and **expand the flexibility** of our offerings to **better support hybrid environments**.

Spectrum Scale on AWS value proposition

Easily use a HA scalable filesystem on AWS with POSIX semantics.

Deployment is available through **Amazon Marketplace**

On-demand provisioning **rapid** (6-60 minutes) deployment

No Spectrum Scale admin skills required!

Targeted for HPC use on AWS to move workloads to public clouds.

Global namespace across AWS instances.

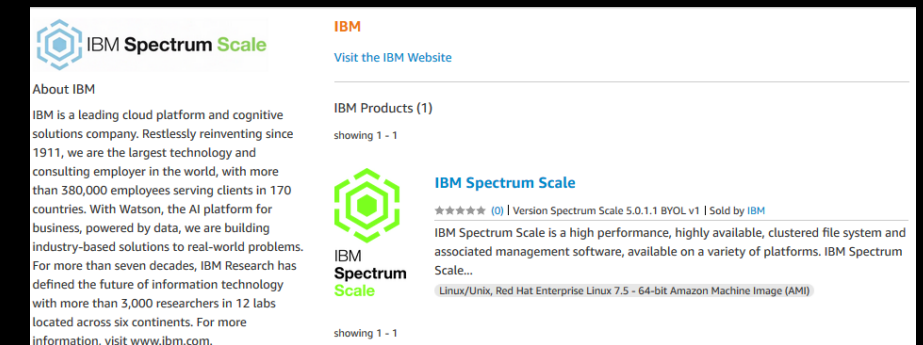
Create cluster → ingest data → use by application → export data → destroy cluster

Try things without infrastructure investments (POCs, code evaluation, etc.)

Provides operational expense model focus (over capital expense model)

Lower costs for some (e.g. 'bursty') workloads

Agile deployment configuration adoption for transient or temporary use cases



IBM Spectrum Scale Offerings on AWS

Marketplace Offering With BYOL

- Available since Sept 2018
- Provides an AMI (boot image) with Spectrum Scale DME (currently 5.0.2.1) on RHEL
- Automated deployment through AWS CloudFormation orchestration service
- Targeted for HPC use on AWS
- Only Bring Your Own License (BYOL)
- You have to pay Amazon for AWS resources used plus RHEL and other software you consume
- Marketplace listing:
<https://aws.amazon.com/marketplace/pp/B07DRLMG2W>

AWS Quickstart

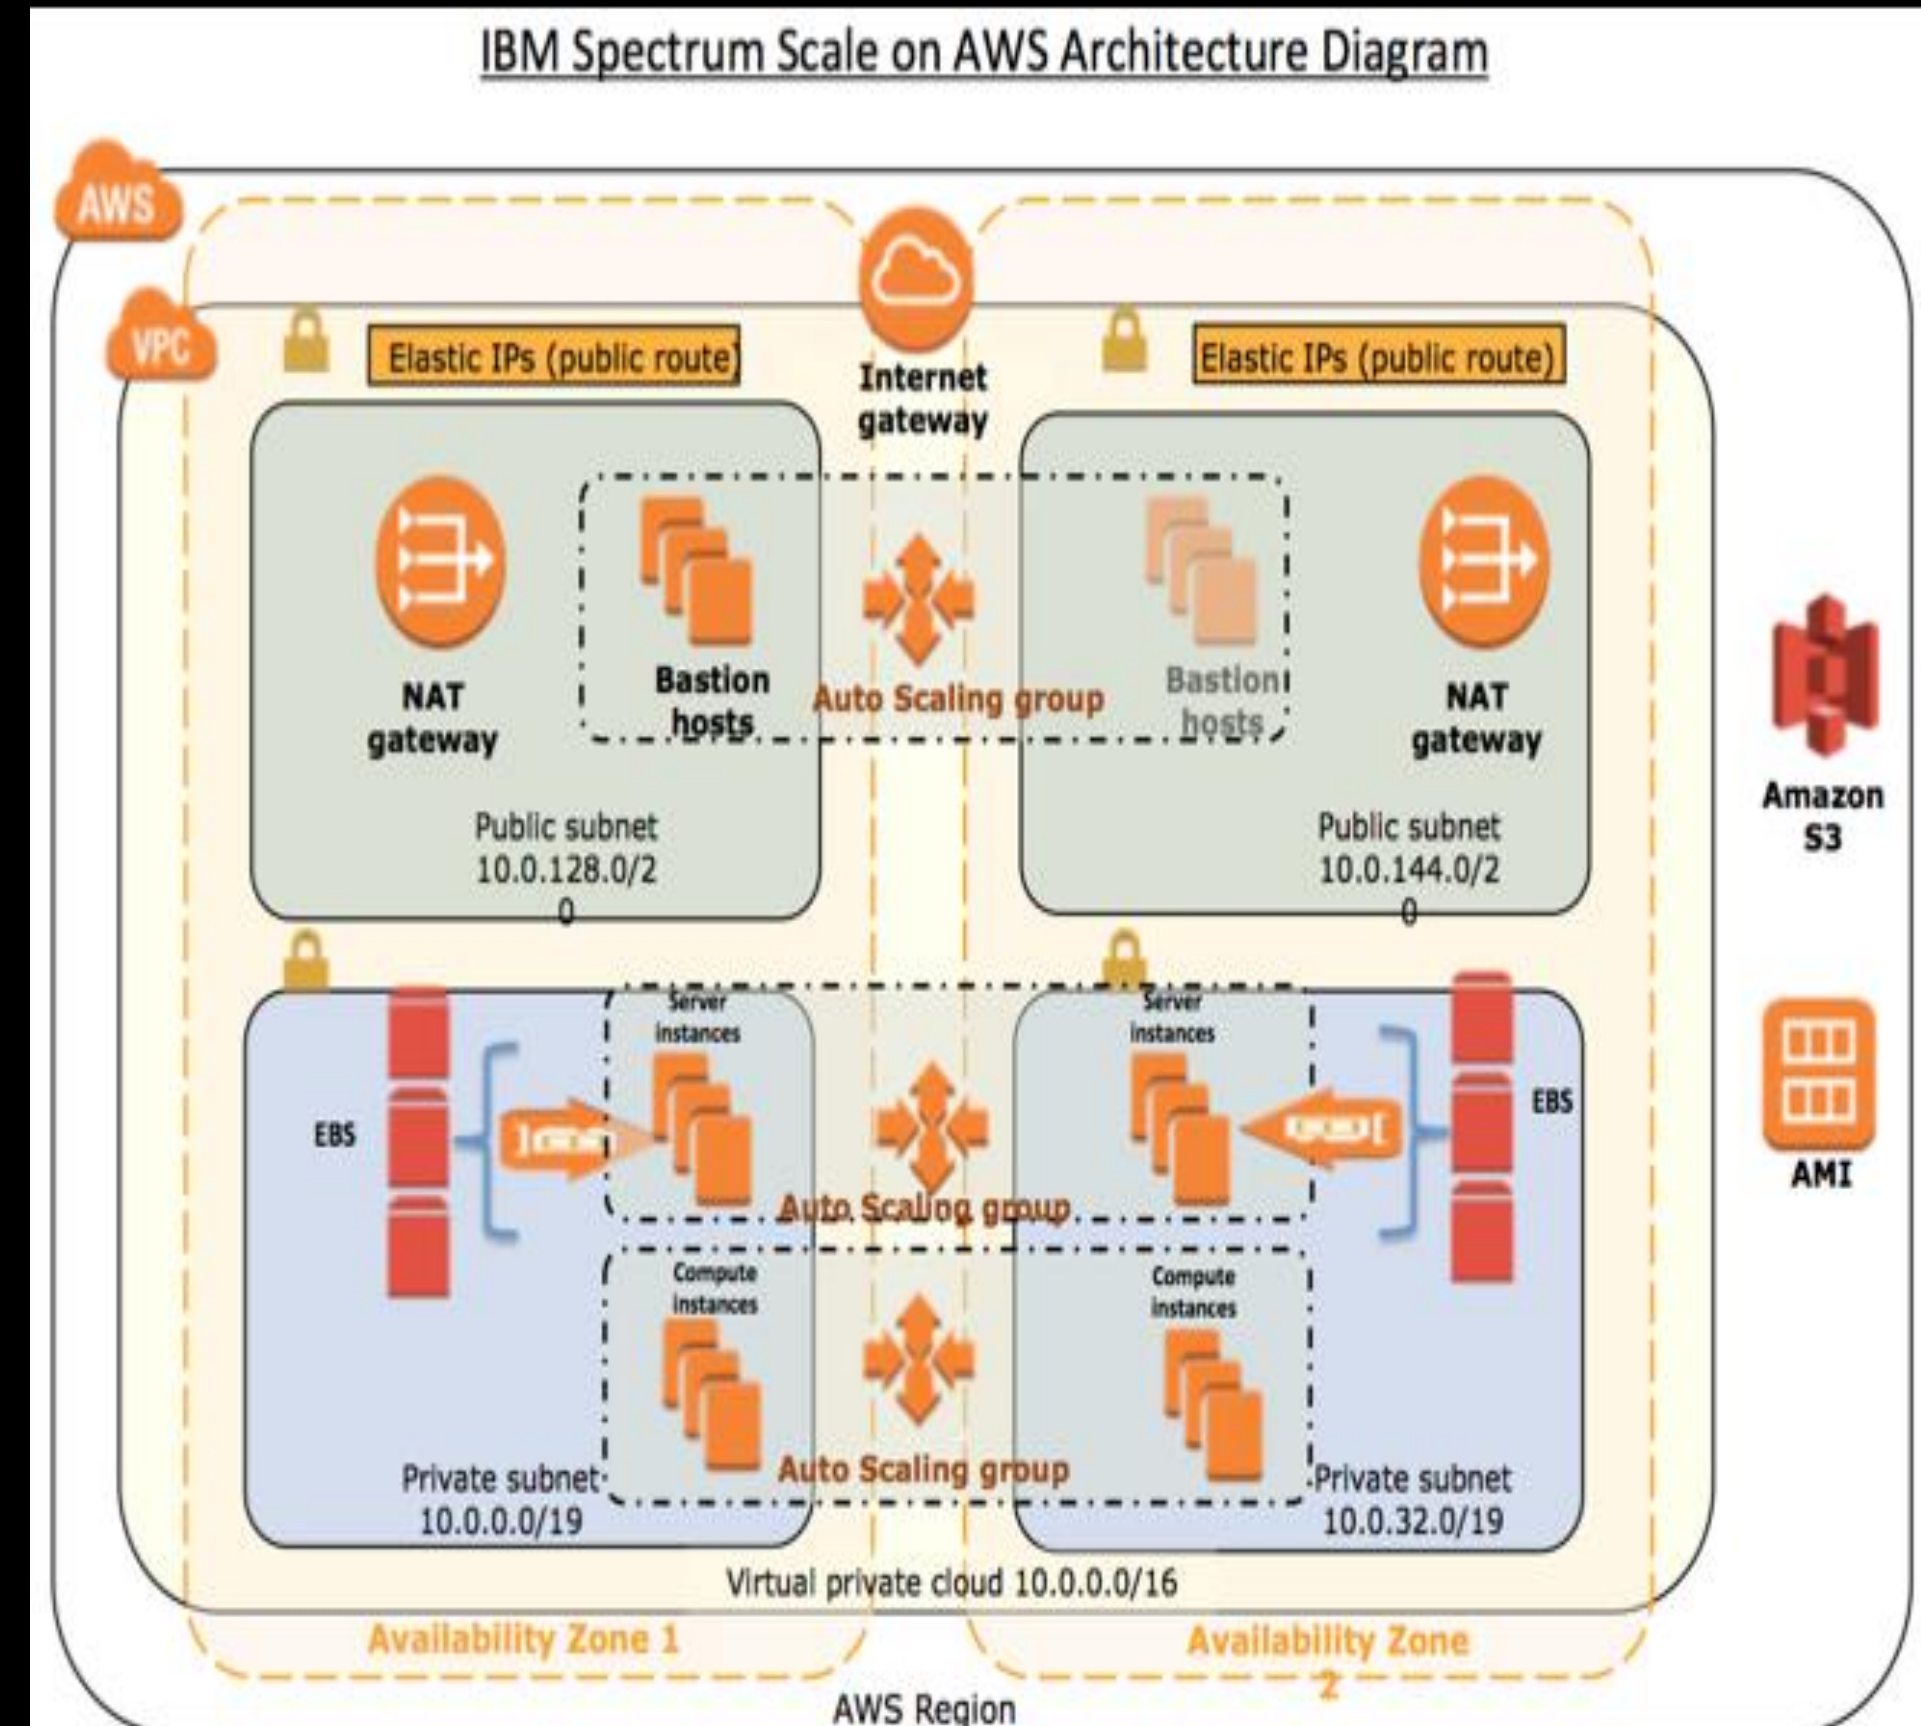
- 90 day trial version
- Available since Sept 2017
- Quickstart:
<https://aws.amazon.com/quickstart/architecture/ibm-spectrum-scale/>

**A cluster of 16 EC2 instances
can be launched and
configured with a shared
filesystem mounted on all
nodes in less than an hour**

(& does not require any Spectrum Scale Admin Skills)

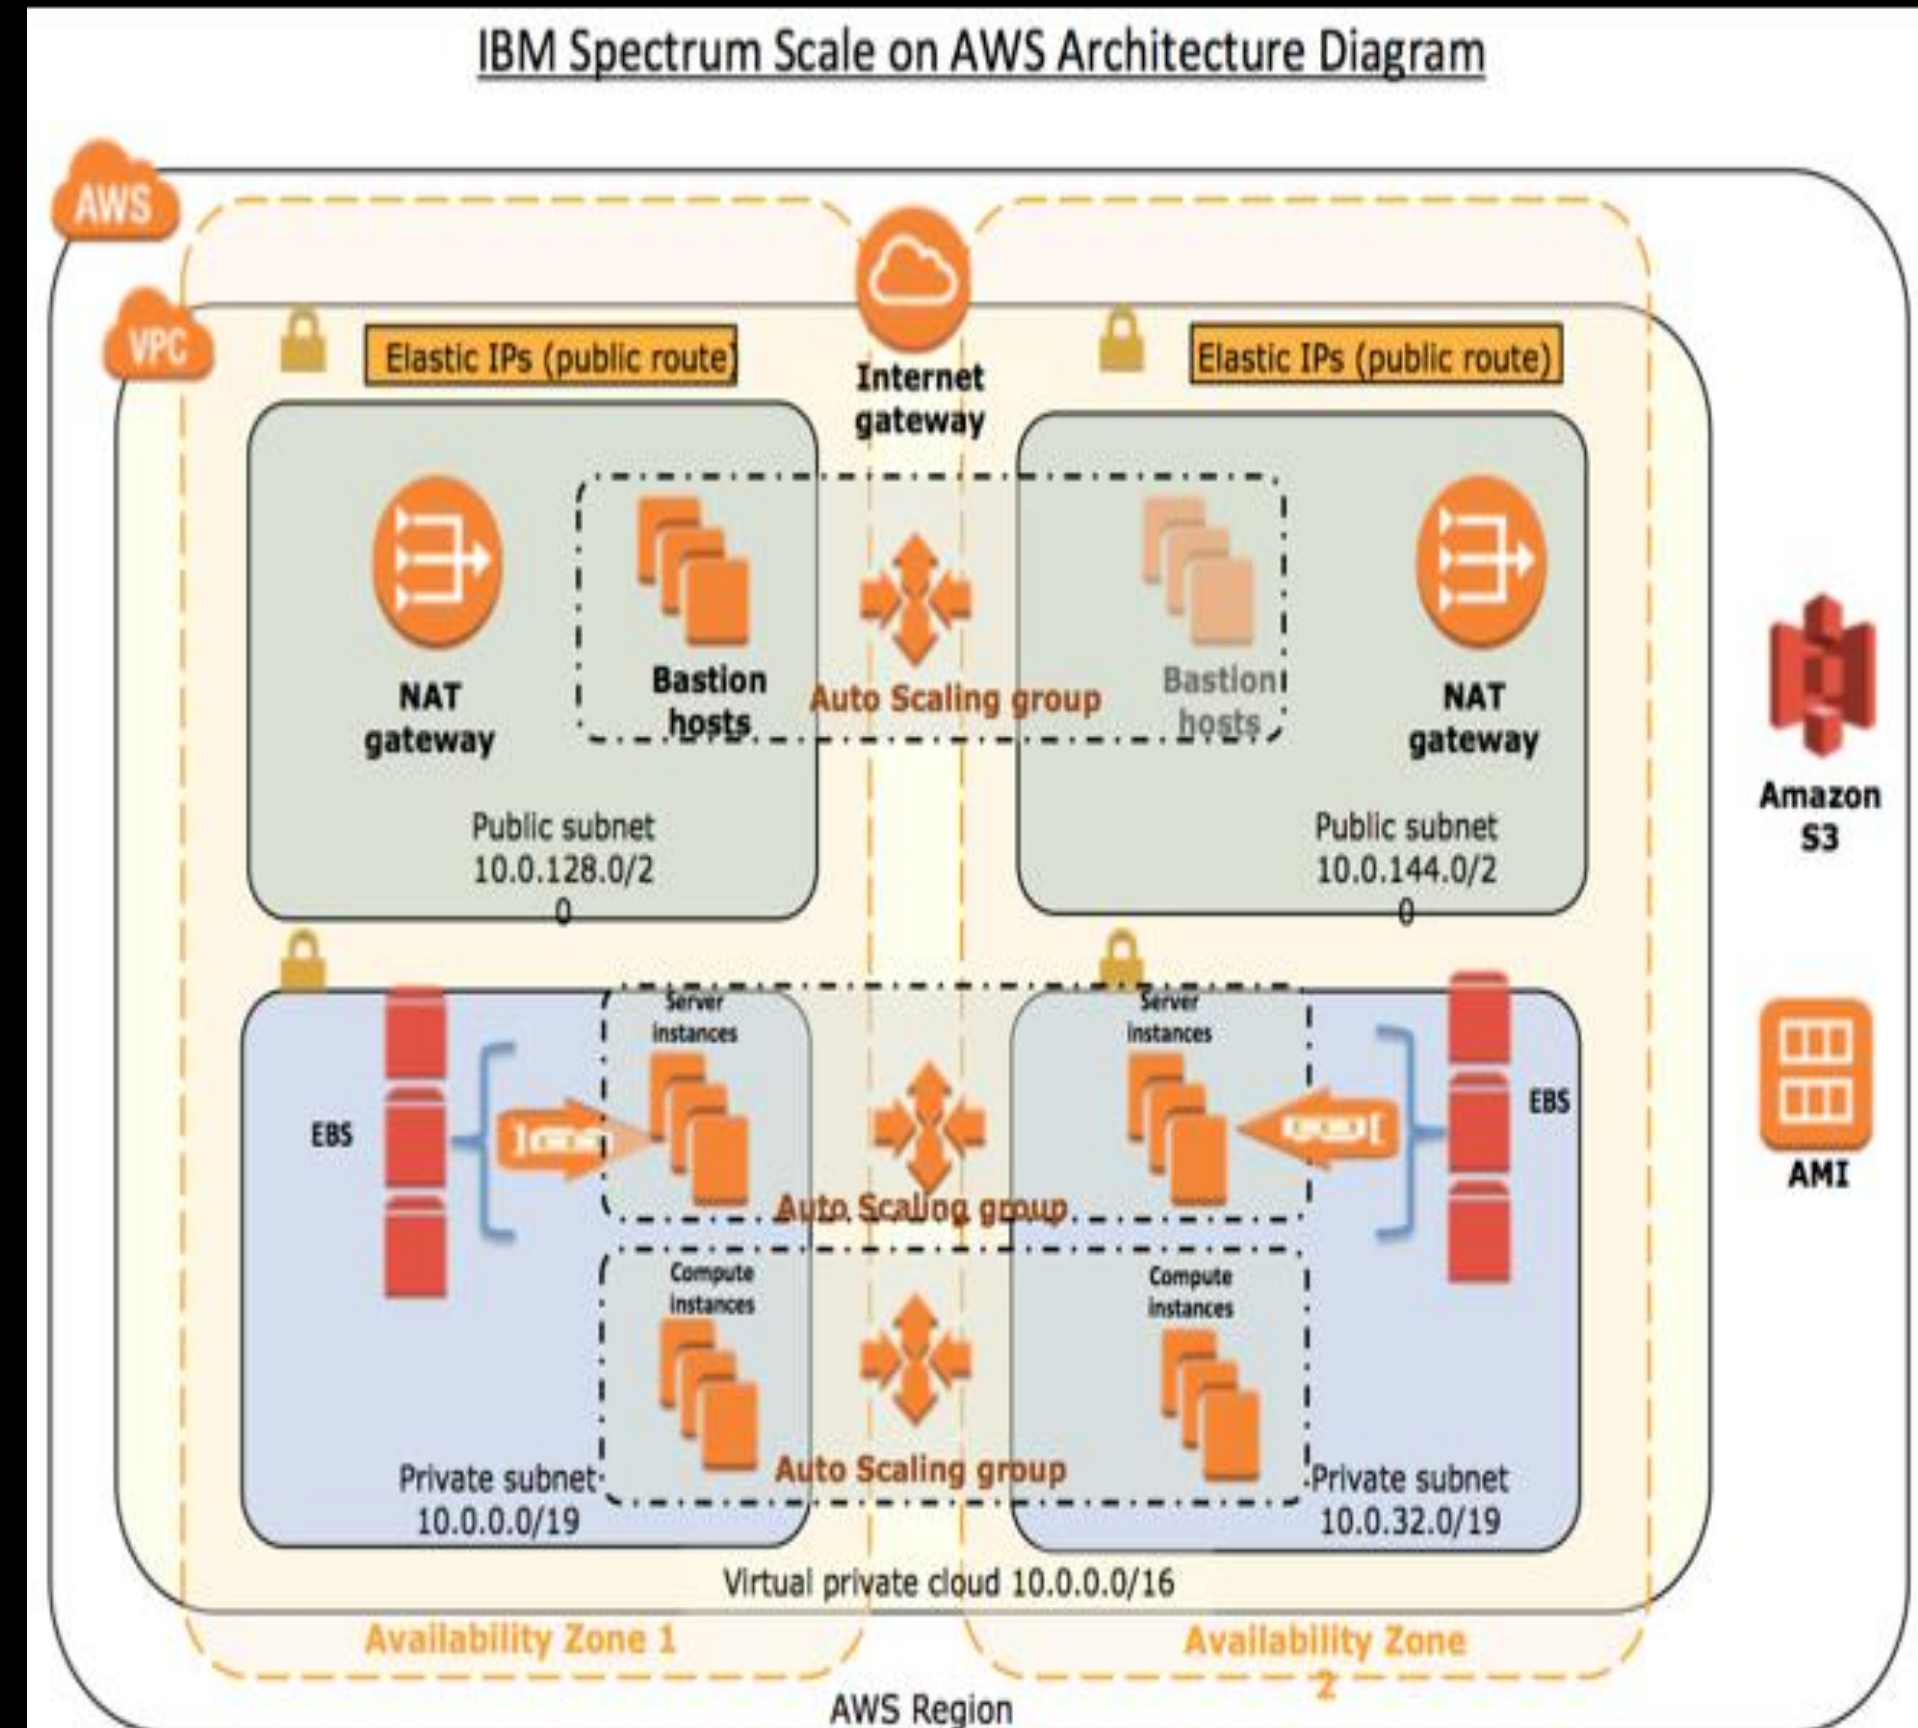
Spectrum Scale on AWS – Architecture

- Solution based on **EC2 instances** and **EBS volumes**
- **NSD servers** using AWS EBS storage are **separate** from “**Compute**” nodes that mount the shared filesystem
- Data is replicated across **multiple availability zones**



Spectrum Scale on AWS – Architecture (cont.)

- Support for creation of **new VPCs** or launching instances into **existing VPCs**.
- Connection to these instances only **via a Bastion Host** (as per AWS best practices)
- Focuses on simplicity/usability of deployment (e.g. reduced config. options) and leverages Amazon features such as **Cloud Formation templates** and **AMIs** (Amazon Machine Image)



Spectrum Scale Running on EC2 Instances Backed by EBS

Managing AWS Workflows + Spectrum Scale cluster

```
mmaws add_nodes
```

```
mmaws remove_nodes
```

```
mmaws list_instances
```

```
mmaws start_nodes
```

```
mmaws stop_nodes
```

```
mmaws create_lambda_functions
```

```
mmaws collect_debug_data
```

Add compute/server nodes

Remove compute/server nodes

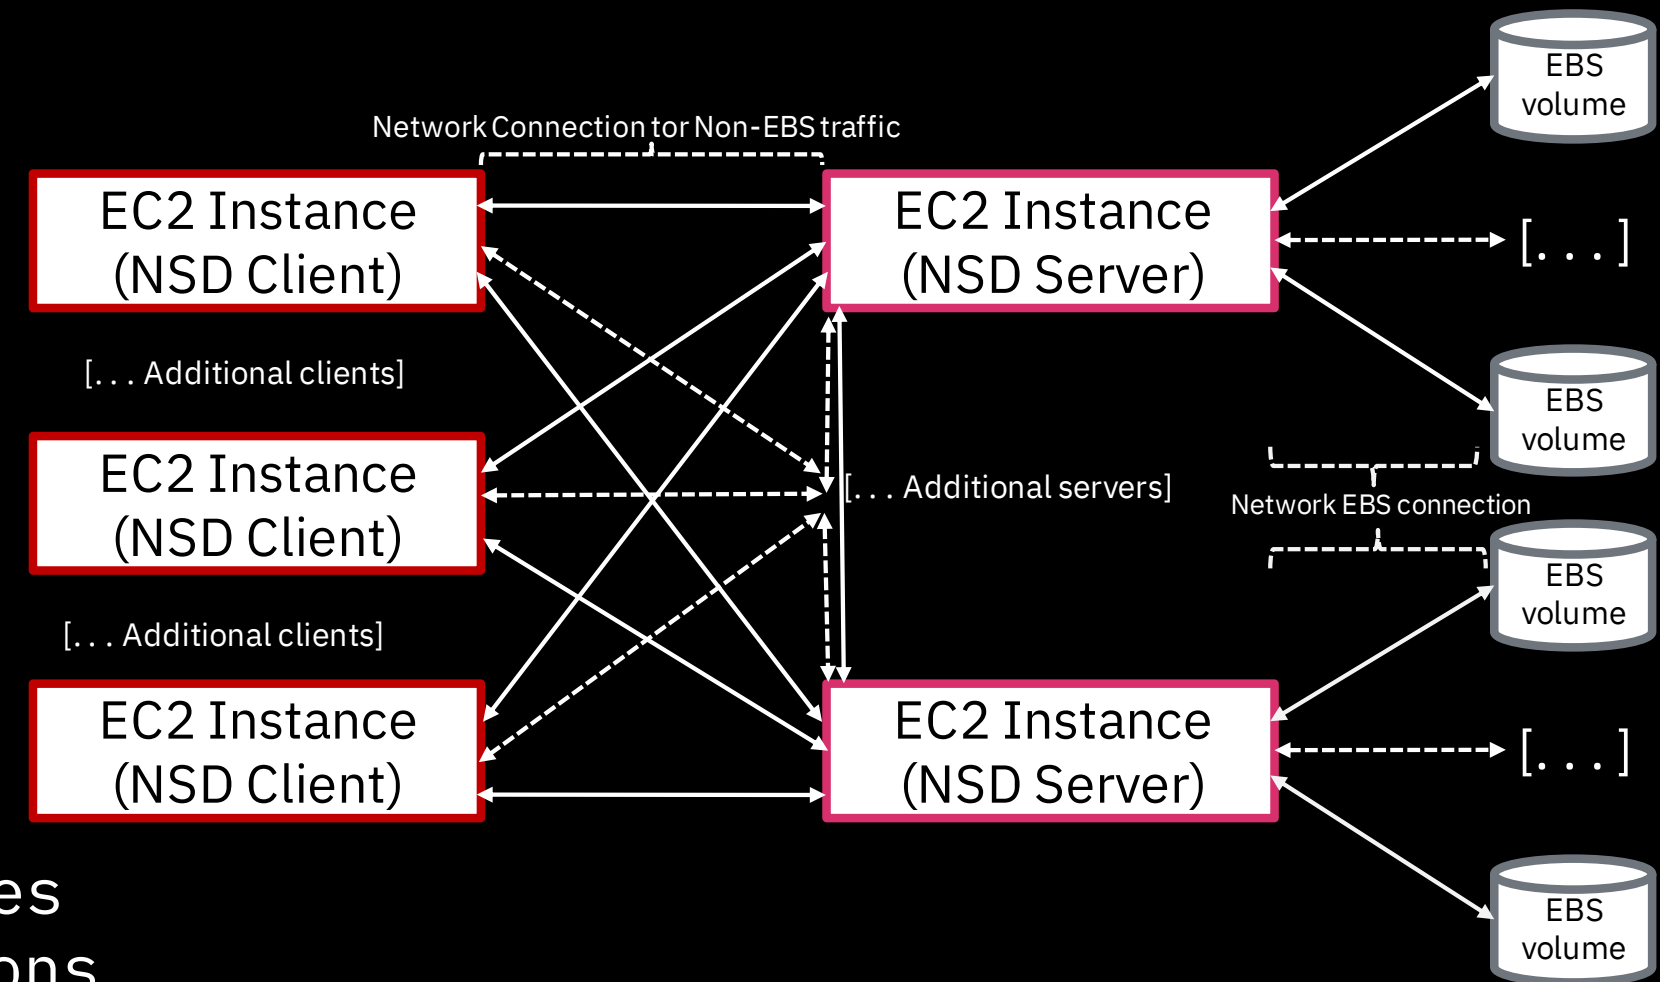
List instances in the vpc

Start nodes

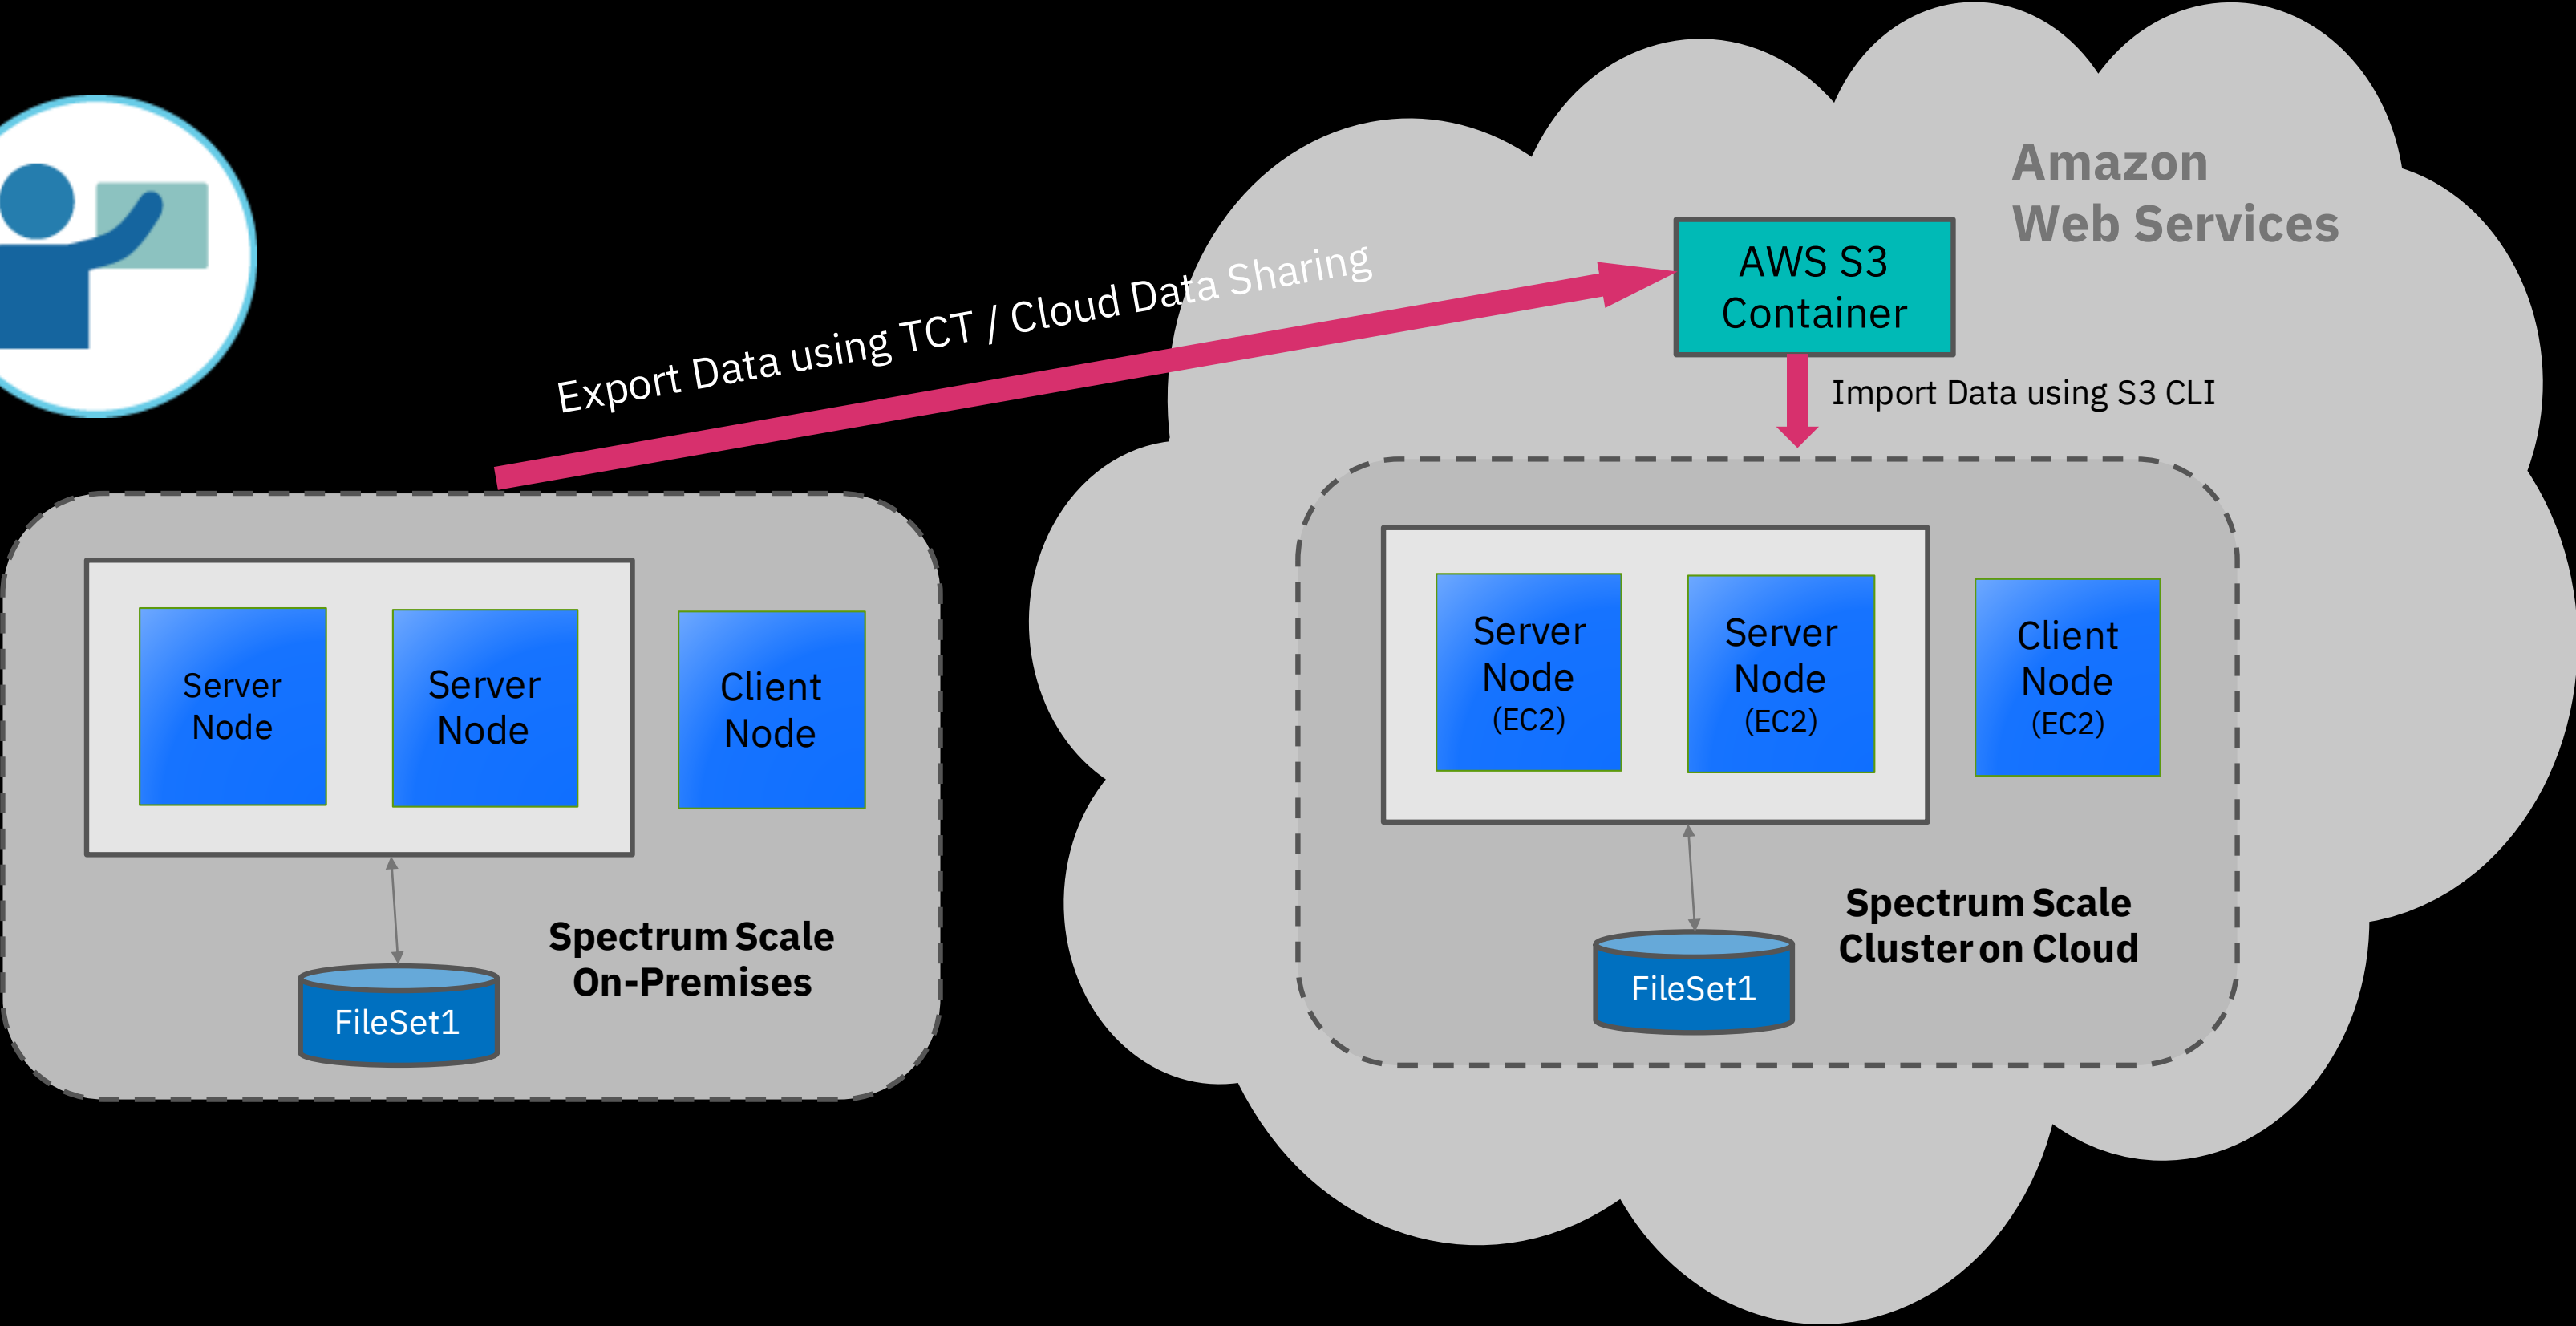
Stop compute/all nodes

Create Lambda functions

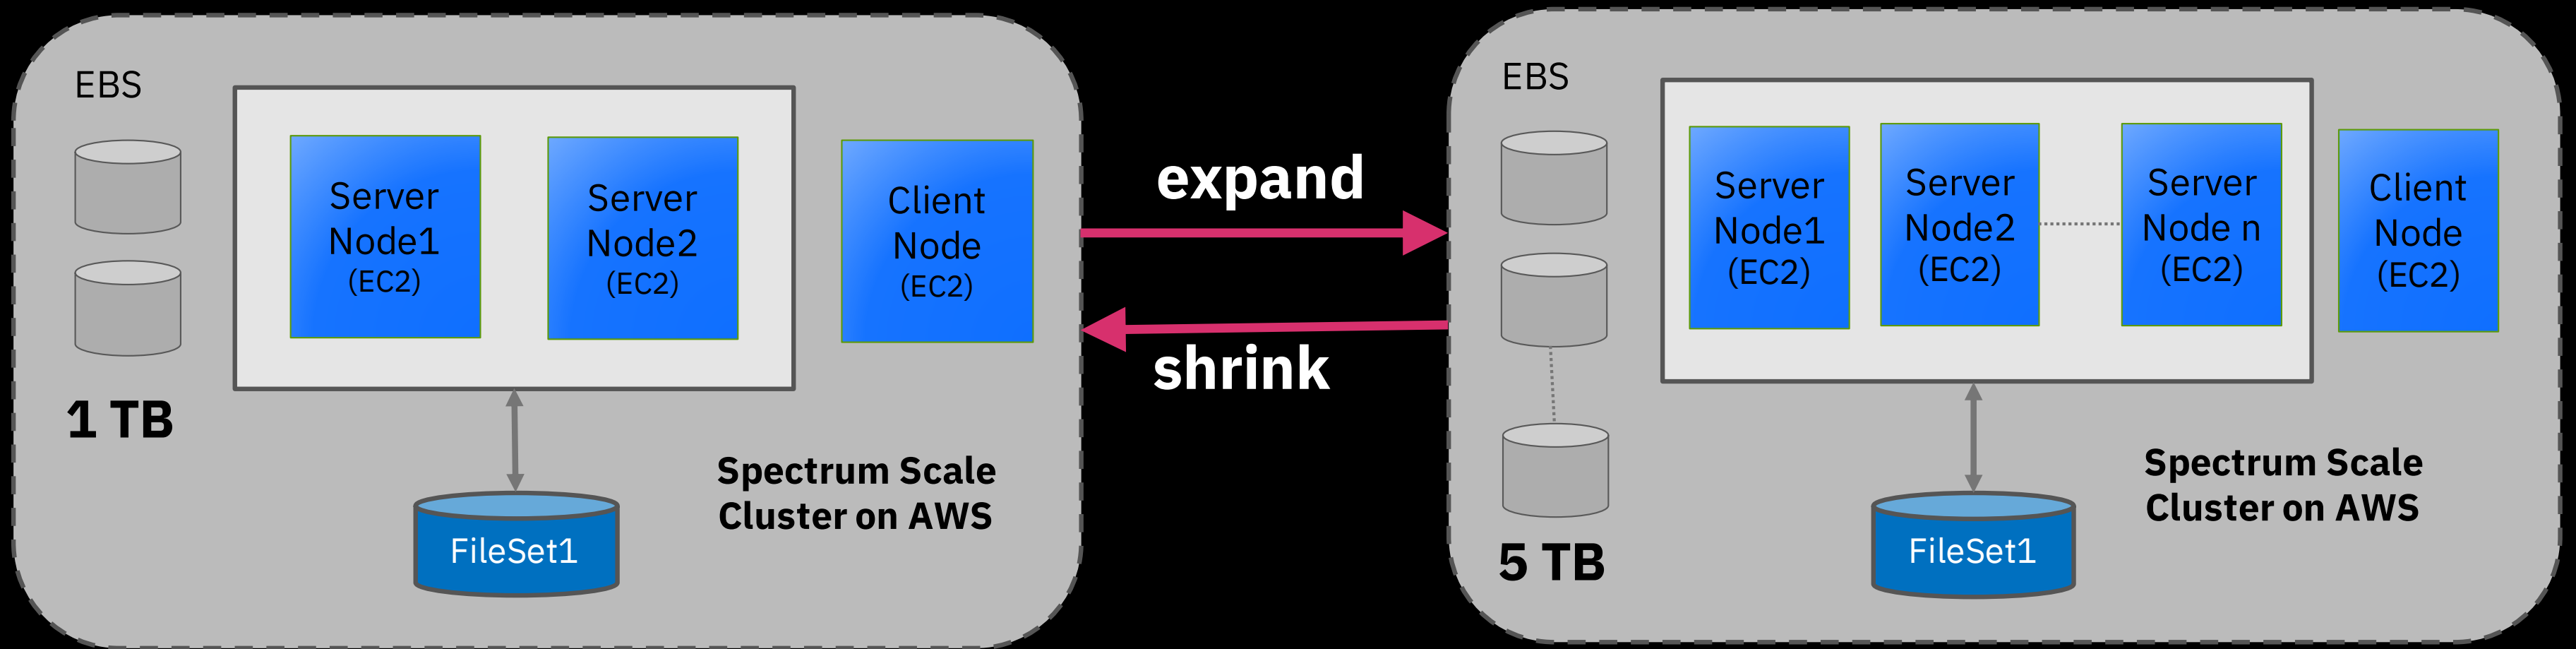
Collect AWS debug data



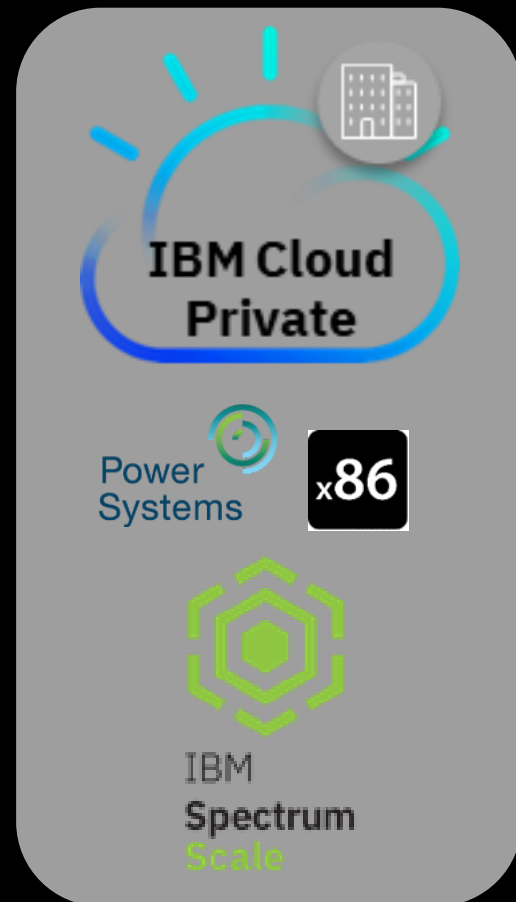
Demo Use Case: CloudBurst



Demo Use Case: Scaling in the Cloud



Hybrid Multicloud Data Agility and Collaboration with IBM Spectrum Scale

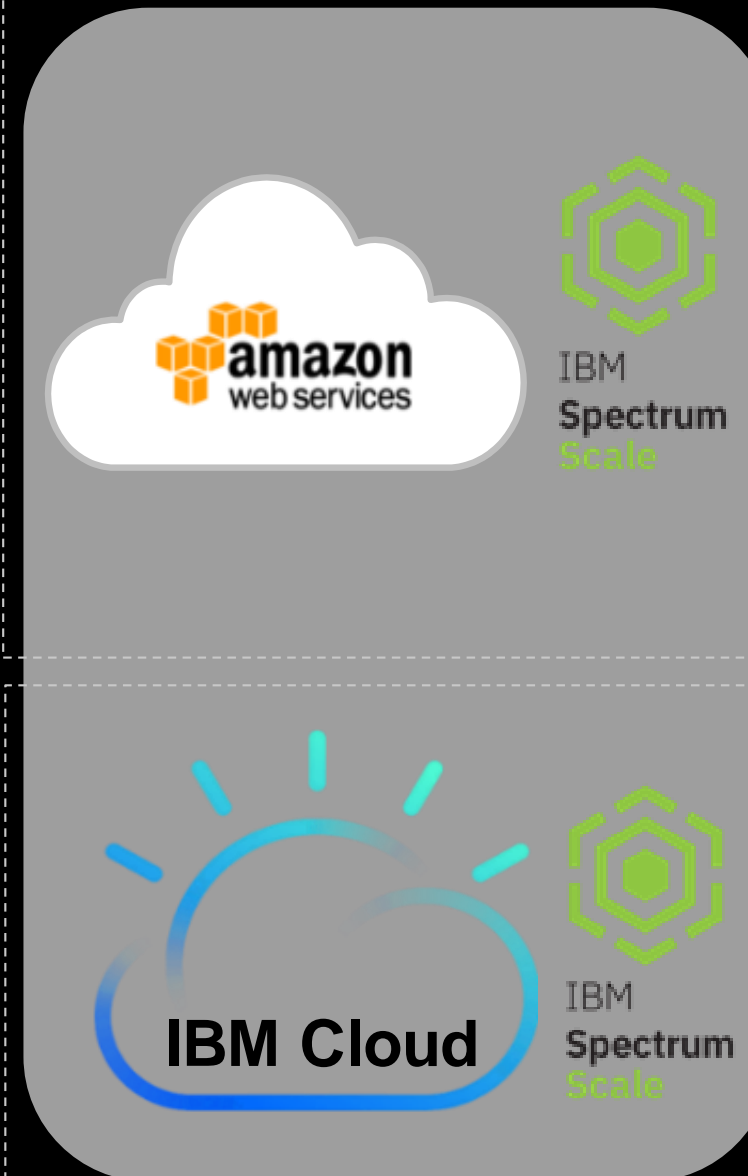


1. On-Premises Private Cloud Using:

- IBM Cloud Private
- IBM Cloud Automation Manager for self-service multicloud
- IBM Spectrum Scale

Connect via VPN over WAN

2. Data Mobility via **IBM Spectrum Scale Advanced File Management** (tech preview)



3. Automated deployment using **IBM Cloud Automation Manager** leveraging EC2 instances and EBS volumes. Deployed in 3-node cluster.

4. **Archive to S3** within AWS

5. Deployed leveraging BareMetal servers and block storage. Deployed in 2-node cluster.

6. Archive to **IBM COS**

Thank you!

Harald Seipp
Senior Technical Staff Member, Center of Excellence for Cloud Storage
IBM Systems EMEA Storage Competence Center
seipp@de.ibm.com



Legal notices

Copyright © 2019 by International Business Machines Corporation. All rights reserved.

No part of this document may be reproduced or transmitted in any form without written permission from IBM Corporation.

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements and/or changes in the product(s) and/or program(s) described herein at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead.

THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER OR IMPLIED. IBM LY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, according to the terms and conditions of the agreements (e.g., IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. IBM makes no representations or warranties, ed or implied, regarding non-IBM products and services.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-785
U.S.A.

Information and Trademarks

IBM, the IBM logo, ibm.com, IBM System Storage, IBM Spectrum Storage, IBM Spectrum Control, IBM Spectrum Protect, IBM Spectrum Archive, IBM Spectrum Virtualize, IBM Spectrum Scale, IBM Spectrum Accelerate, Softlayer, and XIV are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.shtml>

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

IT Infrastructure Library is a Registered Trade Mark of AXELOS Limited.

Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

ITIL is a Registered Trade Mark of AXELOS Limited.

UNIX is a registered trademark of The Open Group in the United States and other countries.

* All other products may be trademarks or registered trademarks of their respective companies.

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

This presentation and the claims outlined in it were reviewed for compliance with US law. Adaptations of these claims for use in other geographies must be reviewed by the local country counsel for compliance with local laws.

Special notices

This document was developed for IBM offerings in the United States as of the date of publication. IBM may not make these offerings available in other countries, and the information is subject to change without notice. Consult your local IBM business contact for information on the IBM offerings available in your area.

Information in this document concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquires, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

The information contained in this document has not been submitted to any formal IBM test and is provided "AS IS" with no warranties or guarantees either expressed or implied.

All examples cited or described in this document are presented as illustrations of the manner in which some IBM products can be used and the results that may be achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions.

IBM Global Financing offerings are provided through IBM Credit Corporation in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government clients. Rates are based on a client's credit rating, financing terms, offering type, equipment type and options, and may vary by country. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice.

IBM is not responsible for printing errors in this document that result in pricing or information inaccuracies.

All prices shown are IBM's United States suggested list prices and are subject to change without notice; reseller prices may vary.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Some measurements quoted in this document may have been estimated through extrapolation. Users of this document should verify the applicable data for their specific environment.

Backup

AWS Services / Components

VPC (Virtual Private Cloud): Provision a logically isolated section of the AWS Cloud.

AMI (Amazon Machine Image): It provides the information required to launch an instance.

S3 (object Storage) : Object storage built to store and retrieve data from anywhere on the web

CloudFormation: Allows user to use a simple text file to model and provision, in an automated and secure manner, all the resources needed for the applications across all regions.

AutoScaling: Automatically launch or terminate instances based on user-defined policies, health status checks etc.

Auto Recovery: Automatically recovers the instance if it becomes impaired due to an underlying hardware failure.

SNS: Pub/Sub messaging and mobile notifications.

IAM Policies, Roles: Identity-based policies are permission policies that can be attached to a principal (or identity), such as an IAM user, role, or group.

CloudWatch: Monitoring service for AWS resources. Collect, track metrics and react immediately.

Lambda: Compute micro service and runs code in response to events such as image uploads, in-app activity, website clicks, or outputs from connected devices.

Amazon EFS vs IBM Spectrum Scale

	EFS	Spectrum Scale
Access Type	NFSv4.1	Posix (NFS/SMB to come in future)
Regional Availability	Not available on all regions like London, Paris, Mumbai etc.	Available on all regions
File system management	Fully automatically managed	Automated management with more control over disk types, instance types to be used, tuning based on workload type, fine grained tuning
Max file size	47.9 TiB	5PB
Feature support	Encryption, IAM based access control, replication	IAM based access control, replication, quotas, ILM (AFM, CES to come in future)
Offering	Pay-as-you-go	BYOL (Pay-as-you-go offering coming in 2019)
Performance (dd tests for 2GB file)	With fsync: 109 MB/s No fsync: 81.1 MB/s	With fsync: 166 MB/s No fsync: 135 MB/s