Deep Dive: Deploy & Configure IBM Storage Scale Anywhere - The Open-Source Way

IBM Storage Scale Days 2025 DE

March 19th – 20th, 2025 | Heidelberg, Germany

Sasikanth Eda (sasikanth.eda@in.ibm.com)

Muthu Muthiah (mutmuthi@in.ibm.com)

Disclaimer

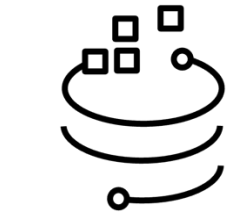

- IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion. Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code, or functionality. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.
- IBM reserves the right to change product specifications and offerings at any time without notice. This publication could include technical inaccuracies or typographical errors. References herein to IBM products and services do not imply that IBM intends to make them available in all countries.

Hybrid cloud storage platforms are increasingly used to leverage cloud agility and elasticity. These solutions empower I&O leaders to update storage and data services across on-premises, edge and public cloud environments, providing mobility, resiliency and operational efficiency.

Gartner Research - Market Guide for Hybrid Cloud Storage

Source: https://www.gartner.com/en/documents/4019387

Hybrid Cloud Use cases

1. Backup or Disaster recovery
2. Global Data Orchestration
3. Burst for Processing or Compute in cloud
4. Burst for Capacity or Tiering

Source: https://www.gartner.com/en/documents/4019387

Sample Scenario: Cloud Bursting

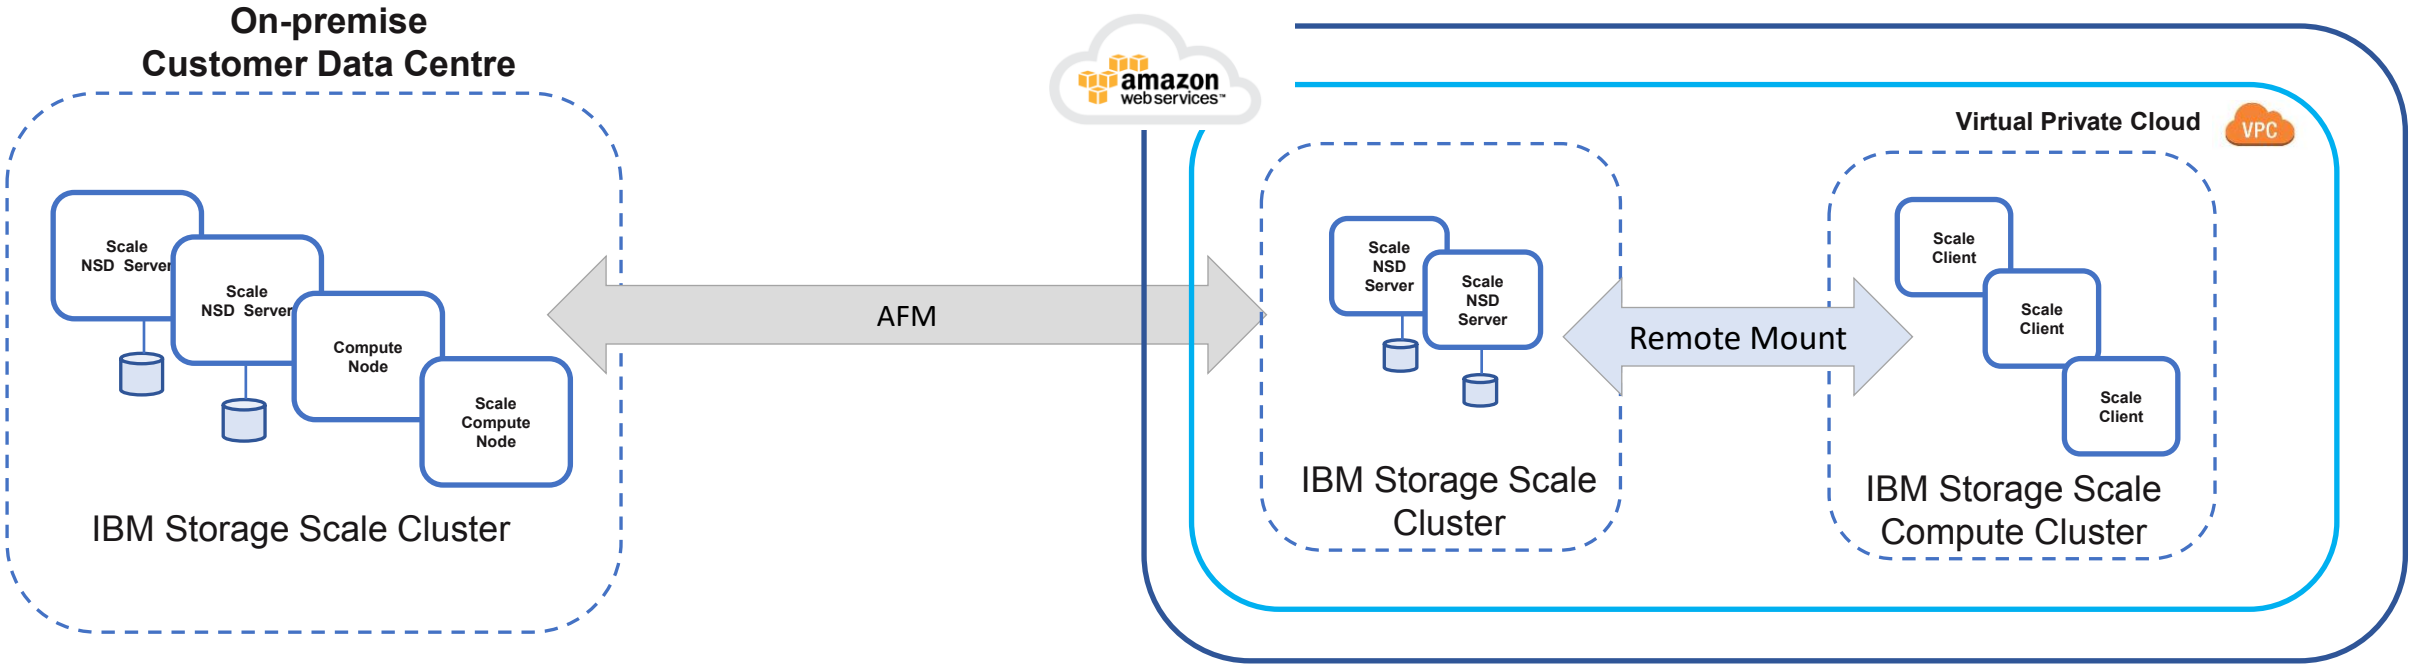

- Burst to cloud for excess capacity requirements
- Permanent Storage Cluster on Cloud (Storage capacity as much as data to be cached on cloud)
- Ephemeral Compute Cluster on Cloud
  (Rapid Deployment / Expansion / Contraction / Destruction)

Sample Scenario: High Performance Tier

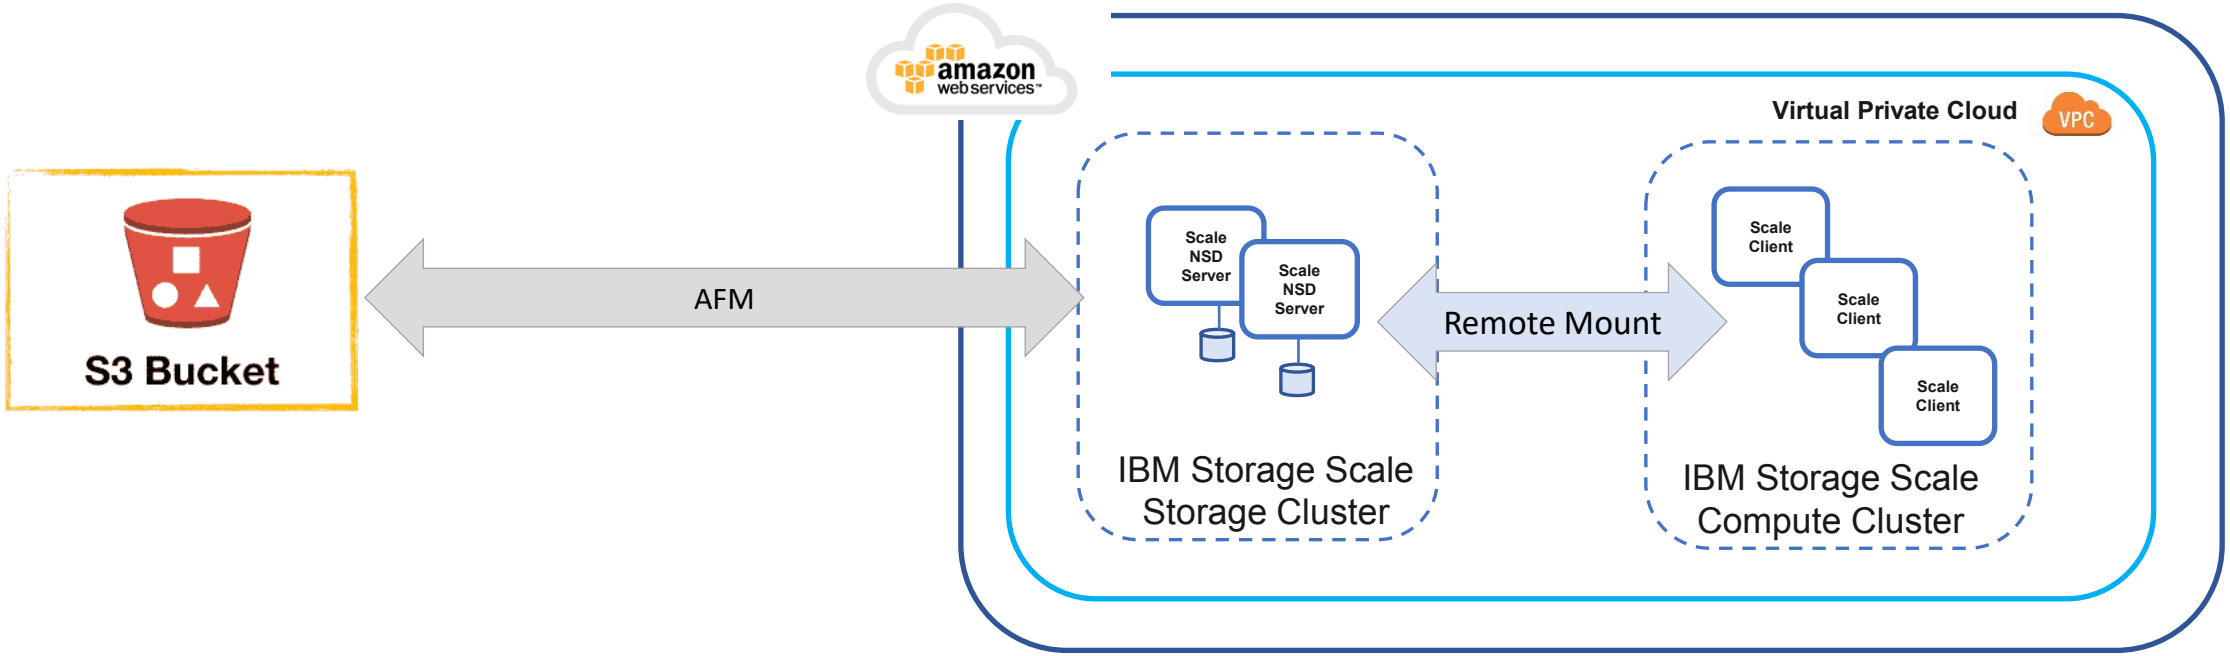

- Data available in Cloud Object Storage (e.g. AWS S3)
- Permanent Storage Cluster on Cloud (Storage capacity as much as data to be cached on cloud)
- Ephemeral Compute Cluster on Cloud
  (Rapid Deployment / Expansion / Contraction / Destruction)

Sample Scenario: Persistent Storage for Containerized workloads on Cloud

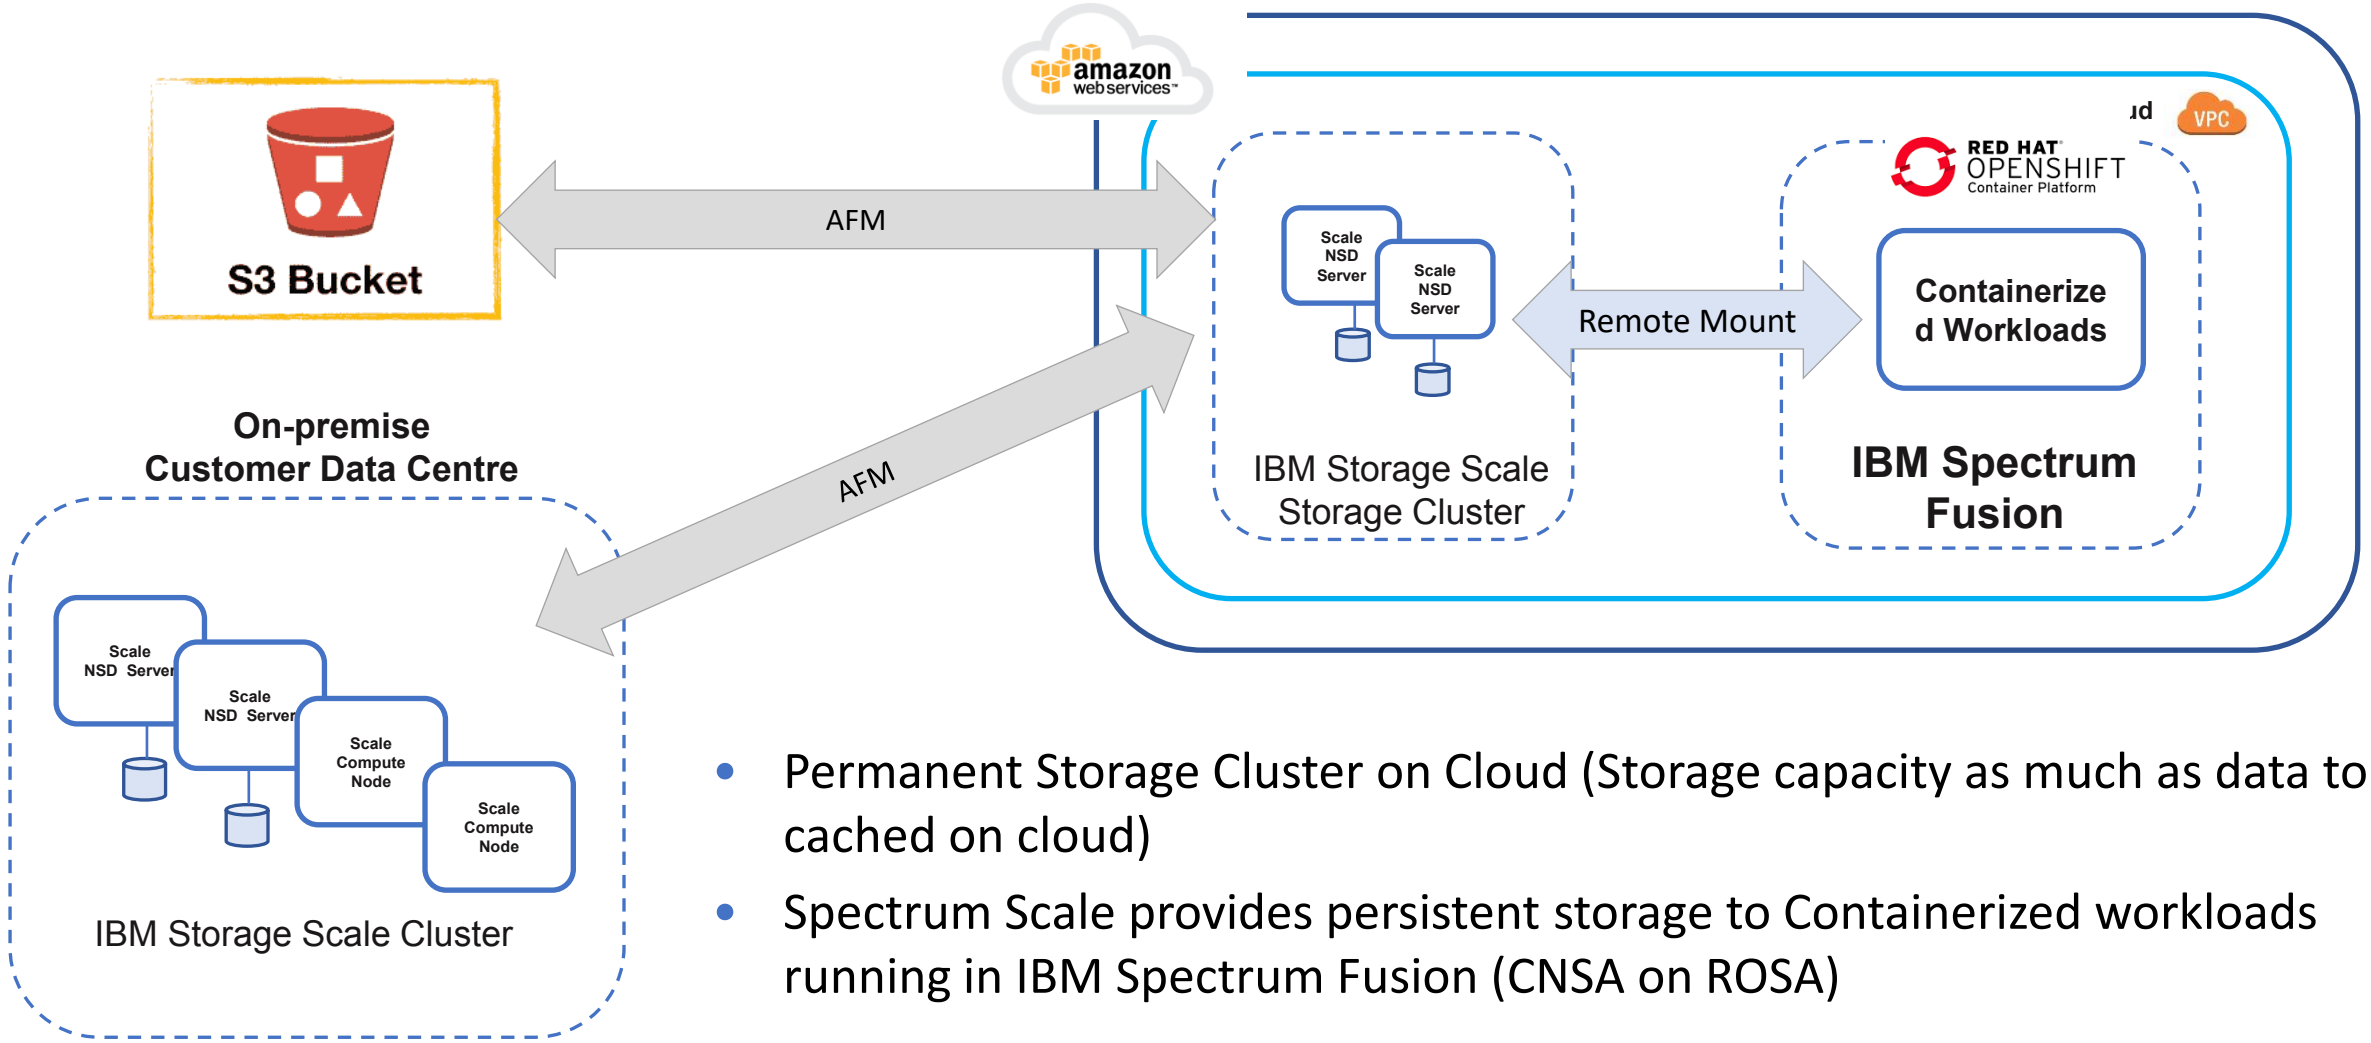

- Permanent Storage Cluster on Cloud (Storage capacity as much as data to be cached on cloud)
- Spectrum Scale provides persistent storage to Containerized workloads running in IBM Spectrum Fusion (CNSA on ROSA)

Cloud Deployment Modes

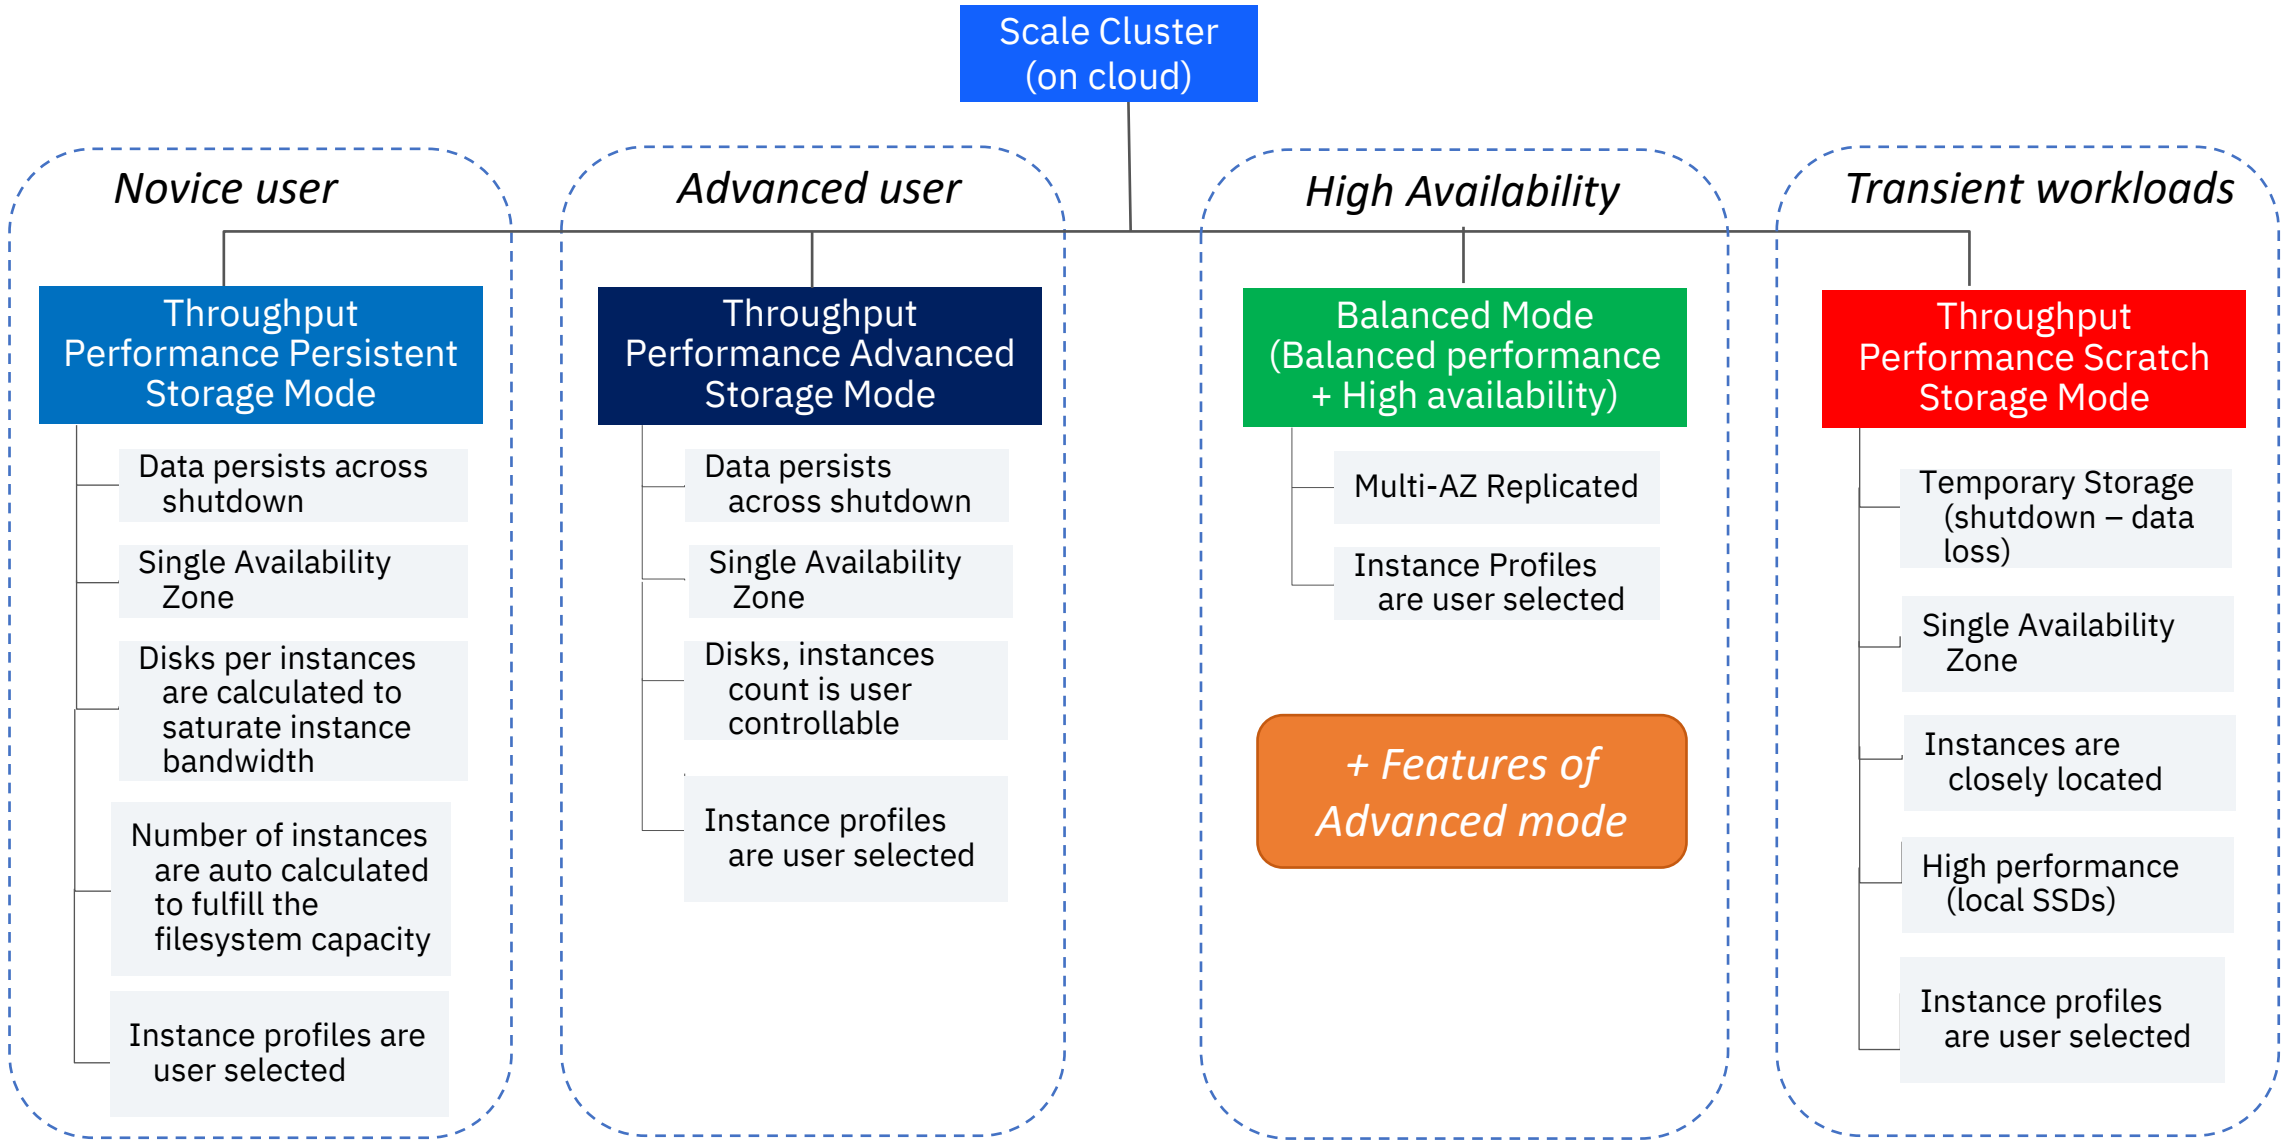

Configuration Management (Ansible)

Ansible Terminology

- Control node
- Managed host (target node)
- Playbook
- Ansible Roles
- Inventory

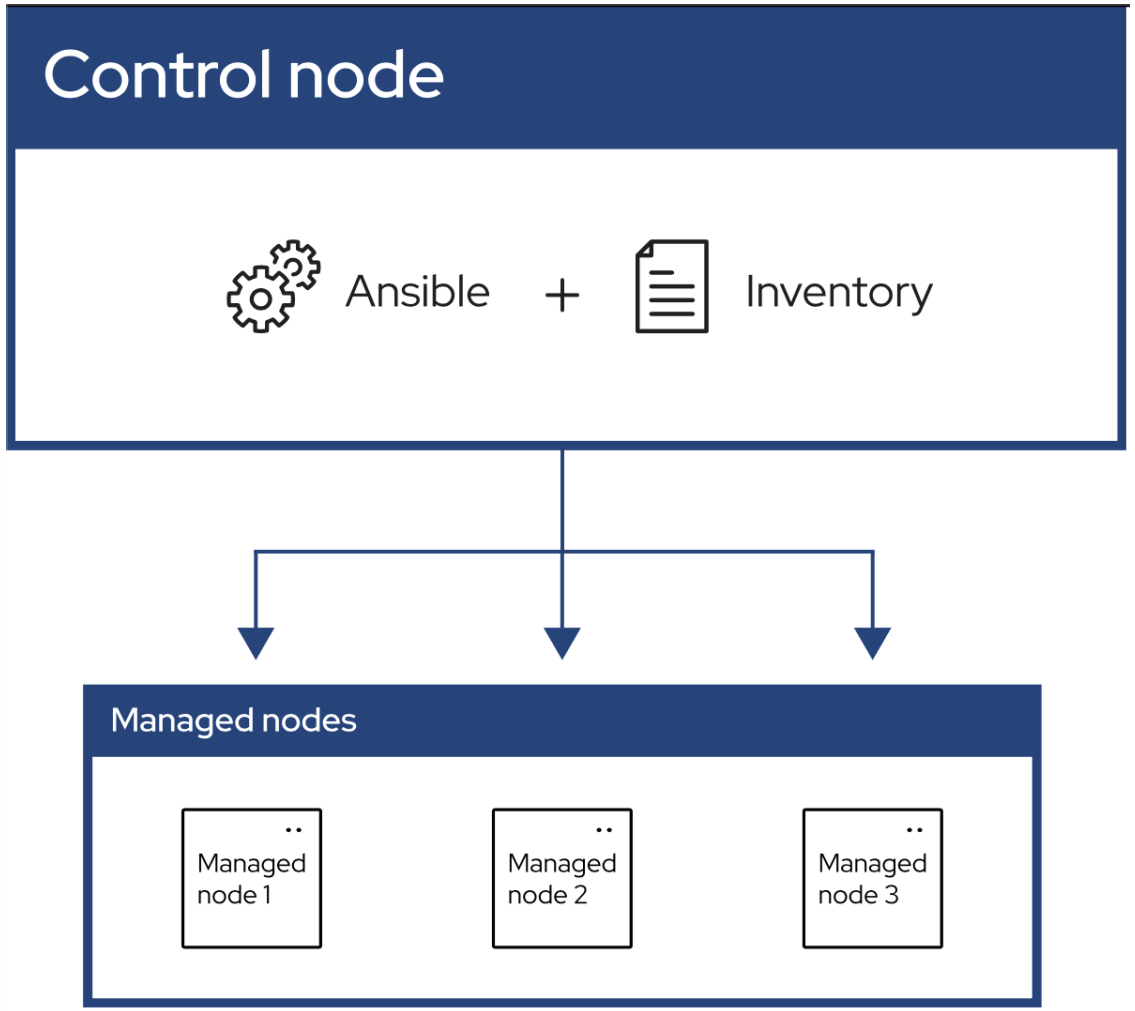

Source: https://docs.ansible.com/ansible/latest/getting_started/index.html

Configuration Management (Ansible)

```
usage: ansible-playbook [-h] [--version] [-v] [--private-key PRIVATE_KEY_FILE]
                        [-u REMOTE_USER] [-c CONNECTION] [-T TIMEOUT]
                        [--ssh-common-args SSH_COMMON_ARGS]
                        [--sftp-extra-args SFTP_EXTRA_ARGS]
                        [--scp-extra-args SCP_EXTRA_ARGS]
                        [--ssh-extra-args SSH_EXTRA_ARGS]
                        [-k | --connection-password-file CONNECTION_PASSWORD_FILE]
                        [--force-handlers] [--flush-cache] [-b]
                        [--become-method BECOME_METHOD]
                        [--become-user BECOME_USER]
                        [-K | --become-password-file BECOME_PASSWORD_FILE]
                        [-t TAGS] [--skip-tags SKIP_TAGS] [-C] [-D]
                        [-i INVENTORY] [--list-hosts] [-l SUBSET]
                        [-e EXTRA_VARS] [--vault-id VAULT_IDS]
                        [-J | --vault-password-file VAULT_PASSWORD_FILES]
                        [-f FORKS] [-M MODULE_PATH] [--syntax-check]
                        [--list-tasks] [--list-tags] [--step]
                        [--start-at-task START_AT_TASK]
                        playbook [playbook ...]
```

Source: https://docs.ansible.com/ansible/latest/getting_started/index.html

Configuration Management (Ansible)

Getting Started with IBM Storage Scale Ansible Roles

$ mkdir my_project
$ cd my_project
$ git clone -b main https://github.com/IBM/ibm-spectrum-scale-install-infra.git collections/ansible_collections/ibm/spectrum_scale

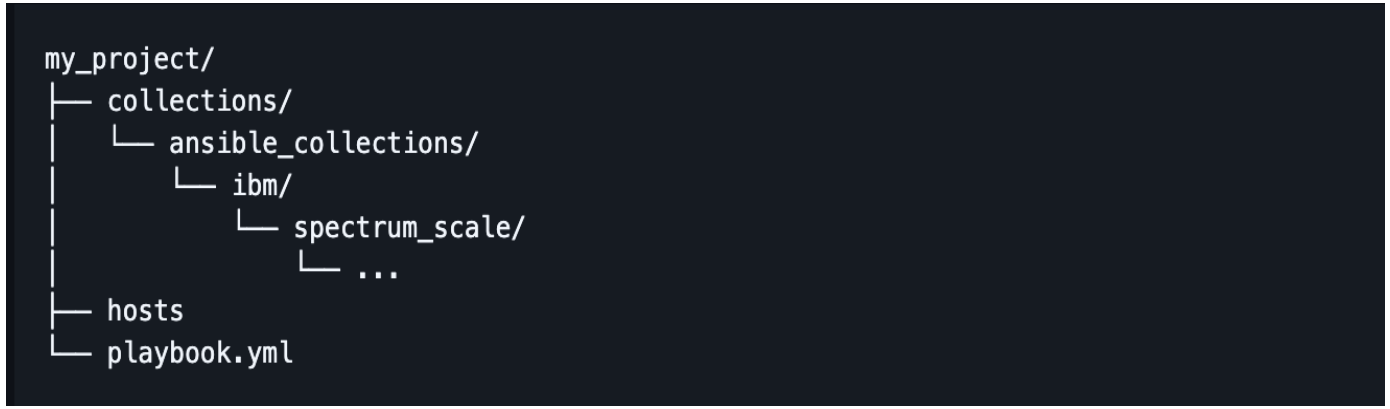

Project Directory Structure

Configuration Management (Ansible)

Create Ansible inventory

- Define IBM Storage Scale nodes in the Ansible inventory (e.g. hosts) in the following format:

```
# hosts:
[cluster01]
scale01 scale_cluster_quorum=true scale_cluster_manager=true
scale02 scale_cluster_quorum=true scale_cluster_manager=true
scale03 scale_cluster_quorum=true scale_cluster_manager=false
scale04 scale_cluster_quorum=false scale_cluster_manager=false
scale05 scale_cluster_quorum=false scale_cluster_manager=false
```

The above is just a minimal example. It defines Ansible variables directly in the inventory. There are other ways to define variables, such as host variables and group variables.

- Numerous variables are available which can be defined in either way to customize the behavior of the roles. Refer to VARIABLES.md for a full list of all supported configuration options.

Configuration Management (Ansible)

Create Ansible playbook

- The basic Ansible playbook (e.g. playbook.yml) looks as follows:

```
# playbook.yml:
---
- hosts: cluster01
  collections:
    - ibm.spectrum_scale
  vars:
    - scale_install_localpkg_path: /path/Spectrum_Scale_Standard-5.2.2.0-x86_64-Linux-install
  roles:
    - core_prepare
    - core_install
    - core_configure
    - core_verify
```

- **Run the playbook to install and configure the IBM Storage Scale cluster**

  Using the ansible-playbook command:
  $ ansible-playbook -i hosts playbook.yml

Best Practices for IaaC Management

1. Version Control (Pin)
2. Use Vault for sensitive inventory storage
3. Implement audit for roles execution
4. Don’t commit your inventory or store in forks
5. Upstream can break your pipeline
6. Code for Idempotency
7. Security/IAM management

Thank you

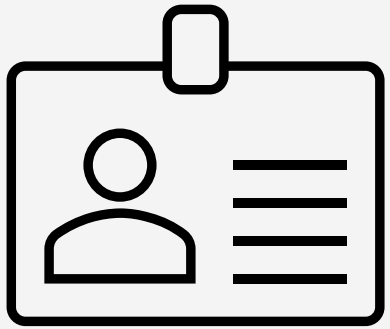

Sasikanth Eda (sasikanth.eda@in.ibm.com)
Muthu Muthiah (mutmuthi@in.ibm.com)

Thank you for using

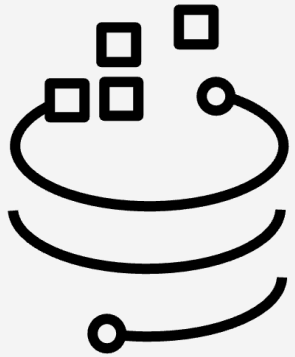

Storage Scale

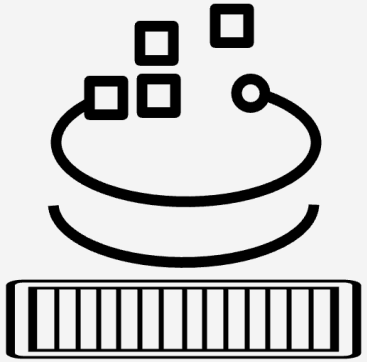

Storage Scale System