IBM Spectrum Scale on Google Cloud

Google Cloud

Agenda

1. Introduction
2. GCP resources and terminology
3. Scale Deployment on GCP
4. Performance
5. Maintenance
6. Q&A

Presenters

1. Jeff Ceason (IBM Business Dev.)

Disclaimer

- This presentation does not represent Google or IBM Spectrum Scale
- There is no committed date to put Spectrum Scale on the Google Cloud Marketplace. This is an investigation by IBM business development.
- All performance and other information regarding Google Cloud Platform is based on testing in us-central region and may not reflect current or other region performance

Regions and Zones

- A region is a specific geographical location where you can run your resources. Each region has one or more zones; For example, the us-central1 region denotes a region in the Central United States that has zones us-central1-a, us-central1-b, us-central1-c, and us-central1-f.
- Resources that live in a zone, such as instances or persistent disks, are referred to as zonal resources. Other resources, like static external IP addresses, are regional.
- Regional resources can be used by any resources in that region, regardless of zone, while zonal resources can only be used by other resources in the same zone

Google Cloud

Machine Types

- Machine types combine vCPUs, RAM, and disk
- Predefined
  - Standard, High memory, High CPU
- Customized

Machine name	Description	vCPUs[1]	Memory (GB)	Max number of persistent disks (PDs)[2]	Max total PD size (TB)
n1-standard-1	Standard machine type with 1 vCPU and 3.75 GB of memory.	1	3.75	128	64
n1-standard-2	Standard machine type with 2 vCPUs and 7.5 GB of memory.	2	7.50	128	64
n1-highcpu-32	High-CPU machine type with 32 vCPUs and 28.8 GB of memory.	32	28.8	128	64

Types of Storage

- **Persistent Storage** (Persistent Disk)
  - **Two types:** pd-standard and pd-ssd
  - Moderate latency & throughput - **dependent on capacity and # vCPUs**
  - Multi-mount support (RO PD)
  - Large, persistent through reboots
- **Local Storage** (Local SSD)
  - Lowest latency, high throughput
  - Small, **non-persistent**
- **Object Storage** (Google Cloud Storage)
  - Object Interface (GET, PUT, DELETE)

Images

- **Public**
  - Free – Centos-6,7, Core OS, Debian, Suse ,Ubuntu
  - Premium - RHEL, SLES, Windows Server
- **Custom**
  - Pre-saved images stored in project
  - Scale on GCP uses Centos7-byol
    - Centos-7 + rpms, root ssh-keys, /scaleGCPinstall scripts, .bashrc, .hushlogin, SELINUX=disabled

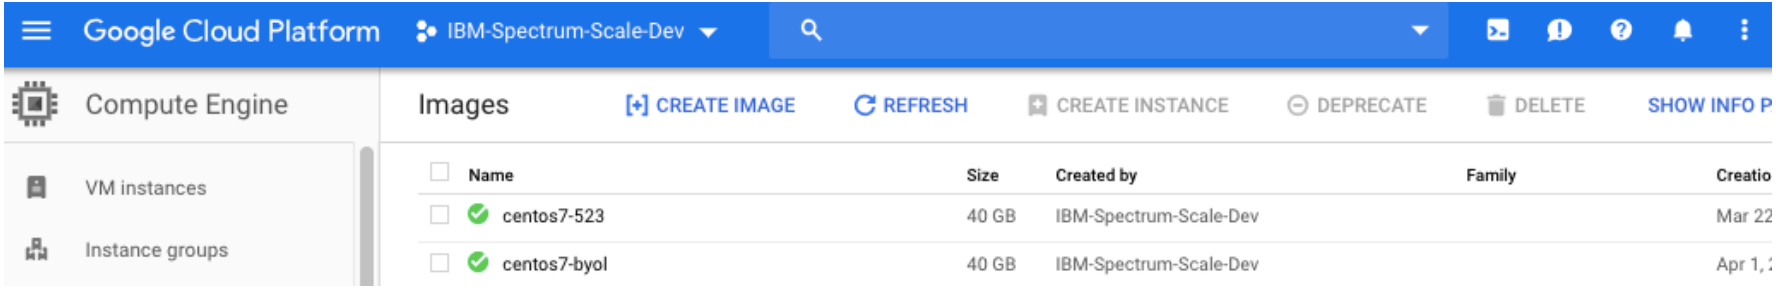

Networking

- **GCP has Virtual Private Cloud (VPC)**
  - Virtual version of a physical network with WW scope
  - Contains regional subnets , each instance belongs to a subnet
  - IPv4 only
  - Can be shared between project/deployments within an account

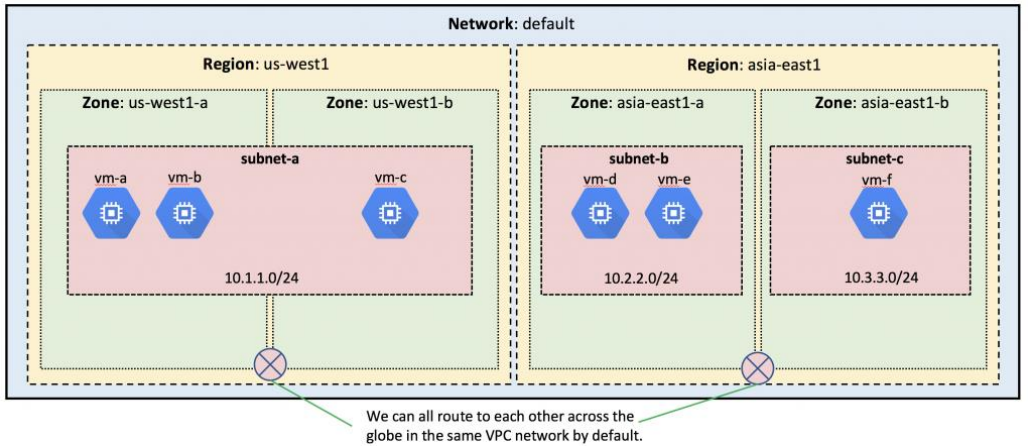

Firewalls

- **Firewalls**
  - Virtual version of a physical network with WW scope
  - Contains regional subnets , each instance belongs to a subnet
  - IPv4 only
  - Can be shared between project/deployments within an account

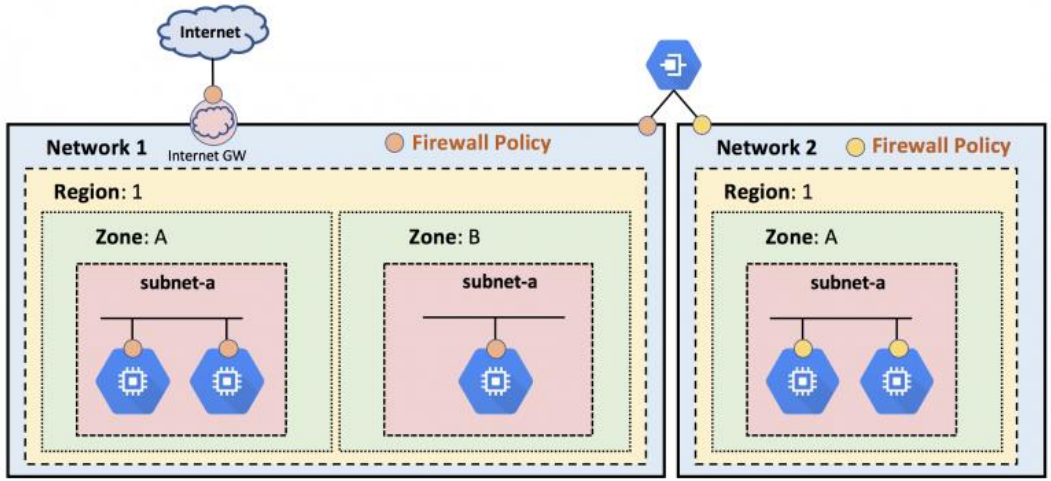

Instances

- Instances - deployed combination of resources and properties into a VM or as containers
  - Combines resources
    - Machine Type
    - Image (Centos7-BYOL)
    - Disks
    - NICS
    - Belongs to a subnet
  - Properties
    - Deletion Protection
    - Metadata

Gcloud CLI

- **gcloud** is a tool that provides the primary CLI to Google Cloud Platform
  - Available windows, Linux, OSX
  - Query
  - Start, stop, delete
  - Update / edit
  - SSH,SCP

- Examples:
  - gcloud config list
  - gcloud compute --project "ibm-spectrum-scale-dev" ssh --zone "us-central1-b" mn
  - gcloud compute scp scalesetup.py mn:/scaleGCPinstall
  - gcloud compute instances stop nsdsn2

Scale on GCP Marketplace - Demo

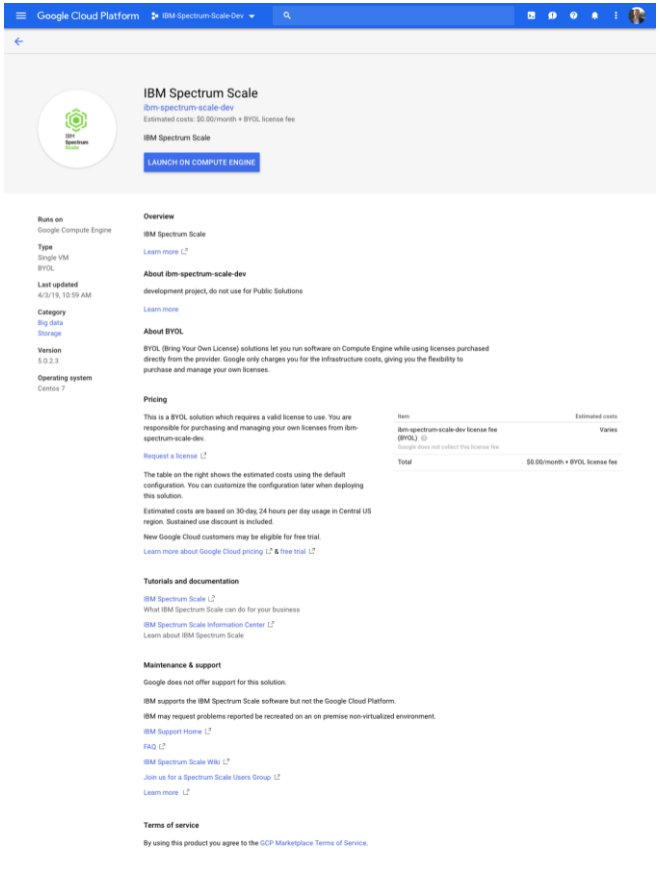

Google Cloud

A deployed cluster

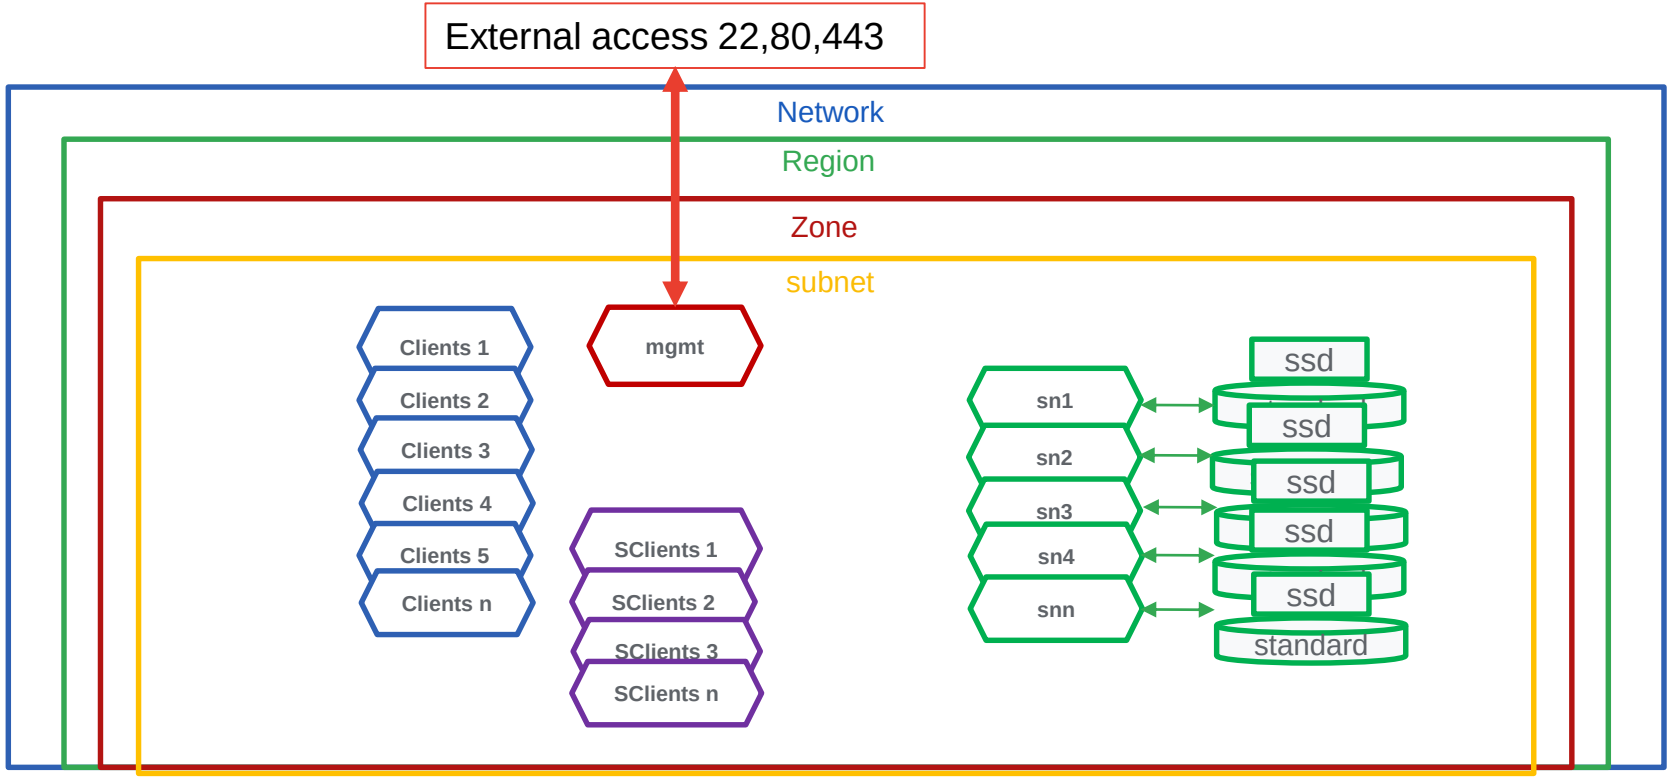

Google Cloud

Disk Devices

```
[root@sn1 ~]# cd /dev/disk/by-id/
[root@sn1 by-id]# ls -alt
drwxr-xr-x 2 root root 200 Apr 16 17:13 .
lrwxrwxrwx 1 root root  10 Apr 16 17:13 google-boot-part1 -> ../../sda1
lrwxrwxrwx 1 root root  10 Apr 16 17:13 scsi-0Google_PersistentDisk_boot-part1 ->
../../sda1
lrwxrwxrwx 1 root root   9 Apr 16 17:13 google-pd-ssdnsd1-1 -> ../../sdb
lrwxrwxrwx 1 root root   9 Apr 16 17:13 scsi-0Google_PersistentDisk_pd-ssdnsd1-1 ->
../../sdb
lrwxrwxrwx 1 root root   9 Apr 16 17:13 google-pd-standardnsd1-1 -> ../../sdc
lrwxrwxrwx 1 root root   9 Apr 16 17:13 scsi-0Google_PersistentDisk_pd-standardnsd1-1 -
> ../../sdc
lrwxrwxrwx 1 root root   9 Apr 16 17:13 google-boot -> ../../sda
lrwxrwxrwx 1 root root   9 Apr 16 17:13 scsi-0Google_PersistentDisk_boot -> ../../sda
drwxr-xr-x 6 root root 120 Apr 16 17:13 ..
```

Performance

- **Bottlenecks**
  - **Block limited on**
    - Number of VCPU , need 32 to get max throughput and IOPs
    - Size of block, pd-ssd 900 GB
    - **Type**
      - pd-ssd 400 MB/s
      - pd-standard 280MB/s
    - Per instance , multiple blocks do not add
  - **Network**
    - Number of VCPU
    - 2Gbs per VCPU
    - Max 16 Gbs ~ 1.86 GB/s (Think 20 Gb/s ethernet)

Network Performance

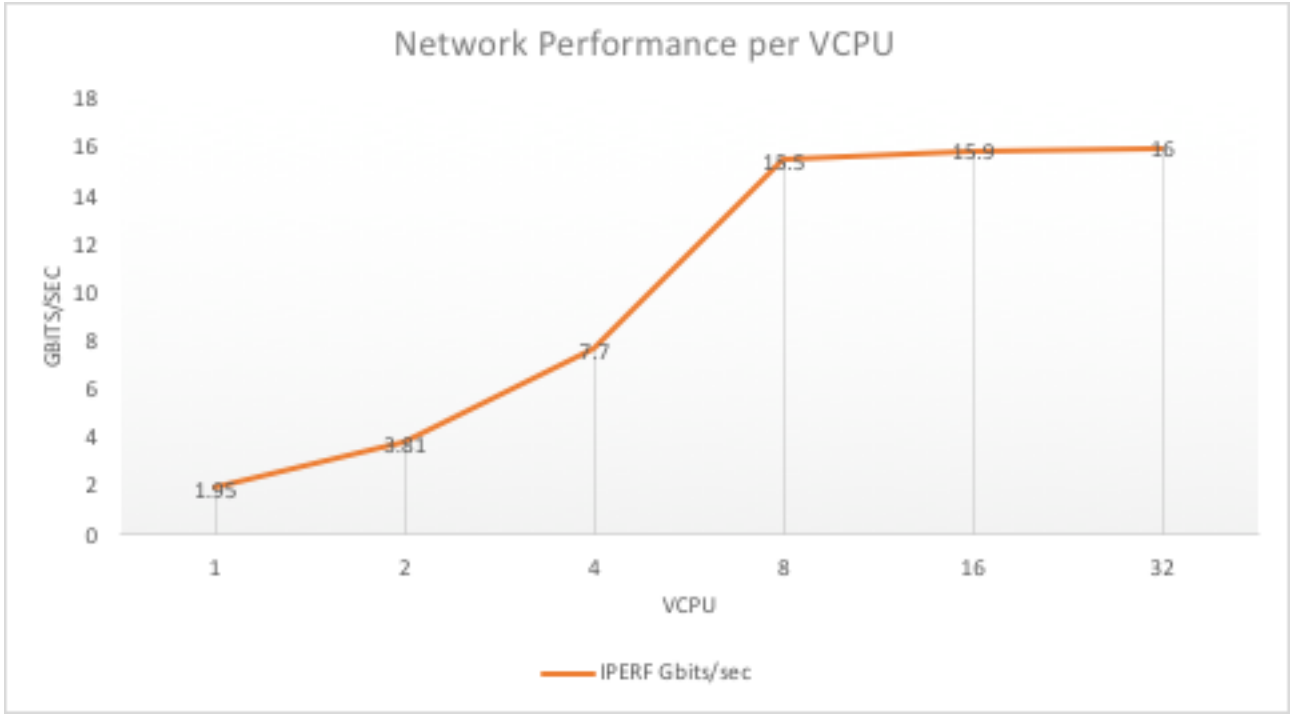

Google Cloud

PD-SSD Performance

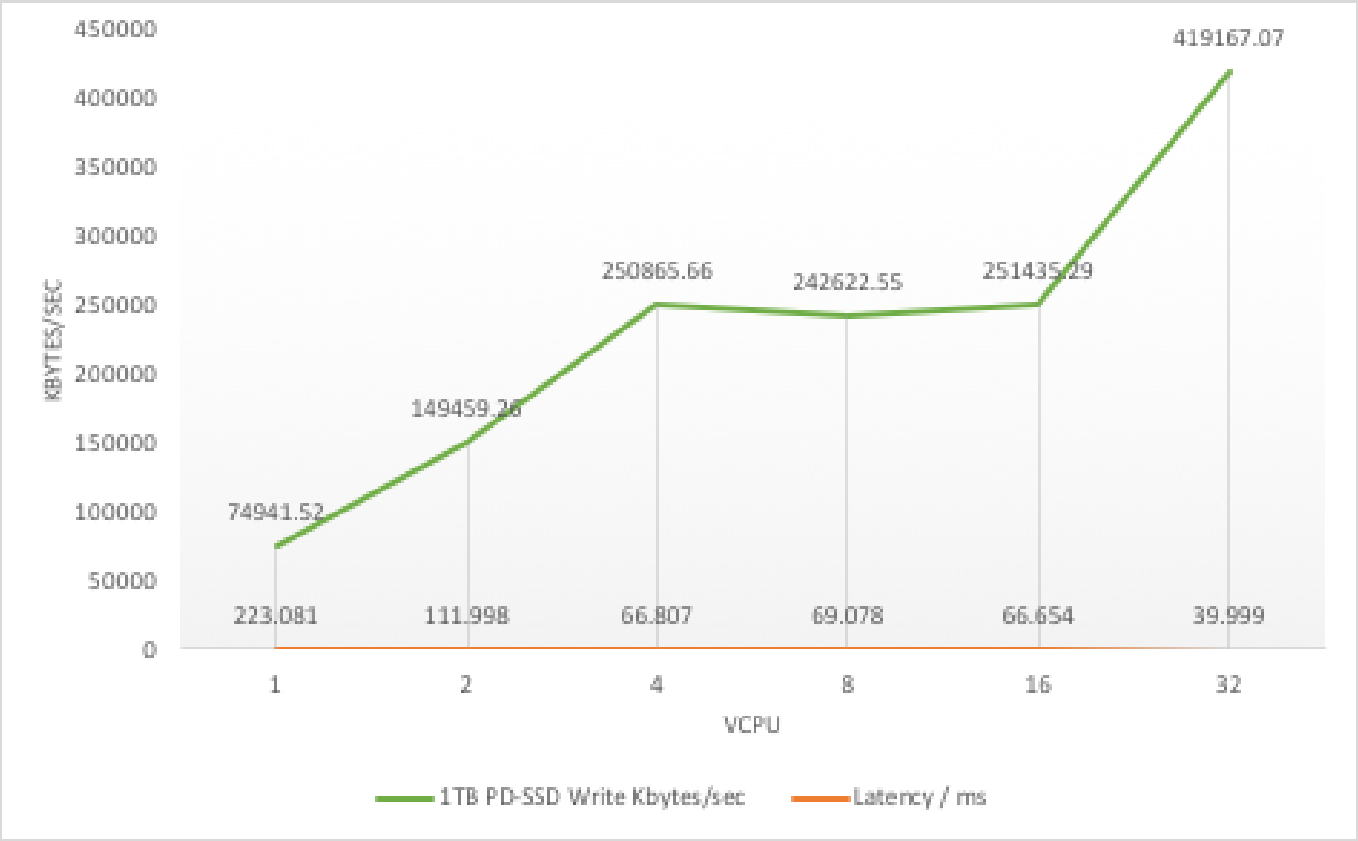

Google

PD-Standard Performance

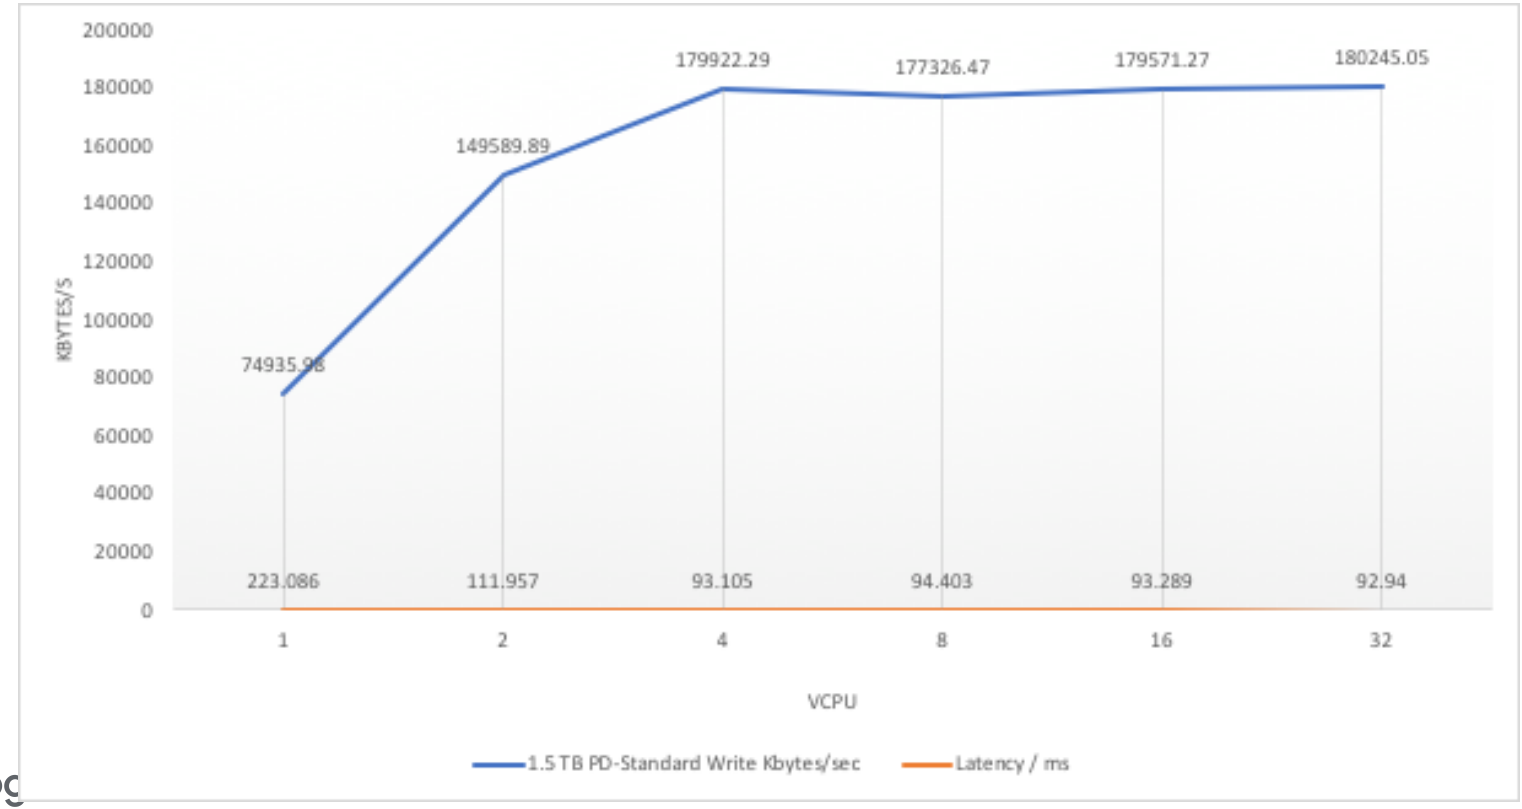

Upgrades

Software

- Updated Spectrum Scale
- Update image software
- Replace image

Hardware

- Adding storage.
- Can attached additional disks (depends on VCPU for how many are allowed)
- **Do not resize disks**. Could result in corruption
- Add nodes, be sure to put in same VPC Network and Subnetwork.

Use Cases

Keep in mind

- **Only 20Gbe network**
- **Cloud may not be cheaper**

1. Test new Scale releases
2. Isolate projects from on-premise Scale solution
3. Provide infrastructure to remote sites
4. Start small and grow to on-premise size
5. Evaluate Scale vs other file systems

Next Steps and Needs

What advanced functionality do users want ?
Example: AFM , TCT, CES

What Linux Distros ?

What configuration options?

Who would like to collaboration?

Would you run compute on NSD servers?

Thank You!

IBM **Spectrum Scale**

Google Cloud