

SUSE SES 5.5 Real life deployment

SUSECon19 - Nashville
Florian Rommel, Datalounges Oy

 @datalounges
 <https://www.datalounges.com>

Welcome!



About Us

Who we are not:

- The traditional run-off the mill IT company

Who we are:

- Cloud Gurus with a level of passion that is not very common
- Excited about new things and extremely good at helping customers learn and embrace new tech
- We work on things like Openstack, Ceph , Kubernetes, Nextcloud etc. (see the nice pictures in the footer??) and make them work for normal companies.
- We work on one of the worlds largest Openstack deployments and own our own cloud
- We have a lot of fun while working extreme hours to make the customers happy
- If you cannot approach us with a challenge or a project and we cannot help you right away, it makes us try harder and come up with a solution that will make you and us happy



Datalounges



nextcloud

kubernetes

SUSE

ceph

openstack.

2 for 1

- We will go over 2 customer case studies
- Deployment decisions were pretty much the same
- Challenges were different and workloads were different



Why SUSE SES?

- Ceph with management and promise of easy deployment
- Licensing is flexible
- Professional support
- Local partners available for extreme cases



Case 1: Cinia Networks

- Once upon a time in a hotel room, far far away..



Datalounges



Design and Decision making

- Design was fairly simple and based on best practises with additional twists
- Decision making was based on cost and offering as well as support and local specialist availability
- 2 Competing Solutions (vendors)
- SUSE SES won out after we showed a real life deployment and helped with a misconfigured cluster they had

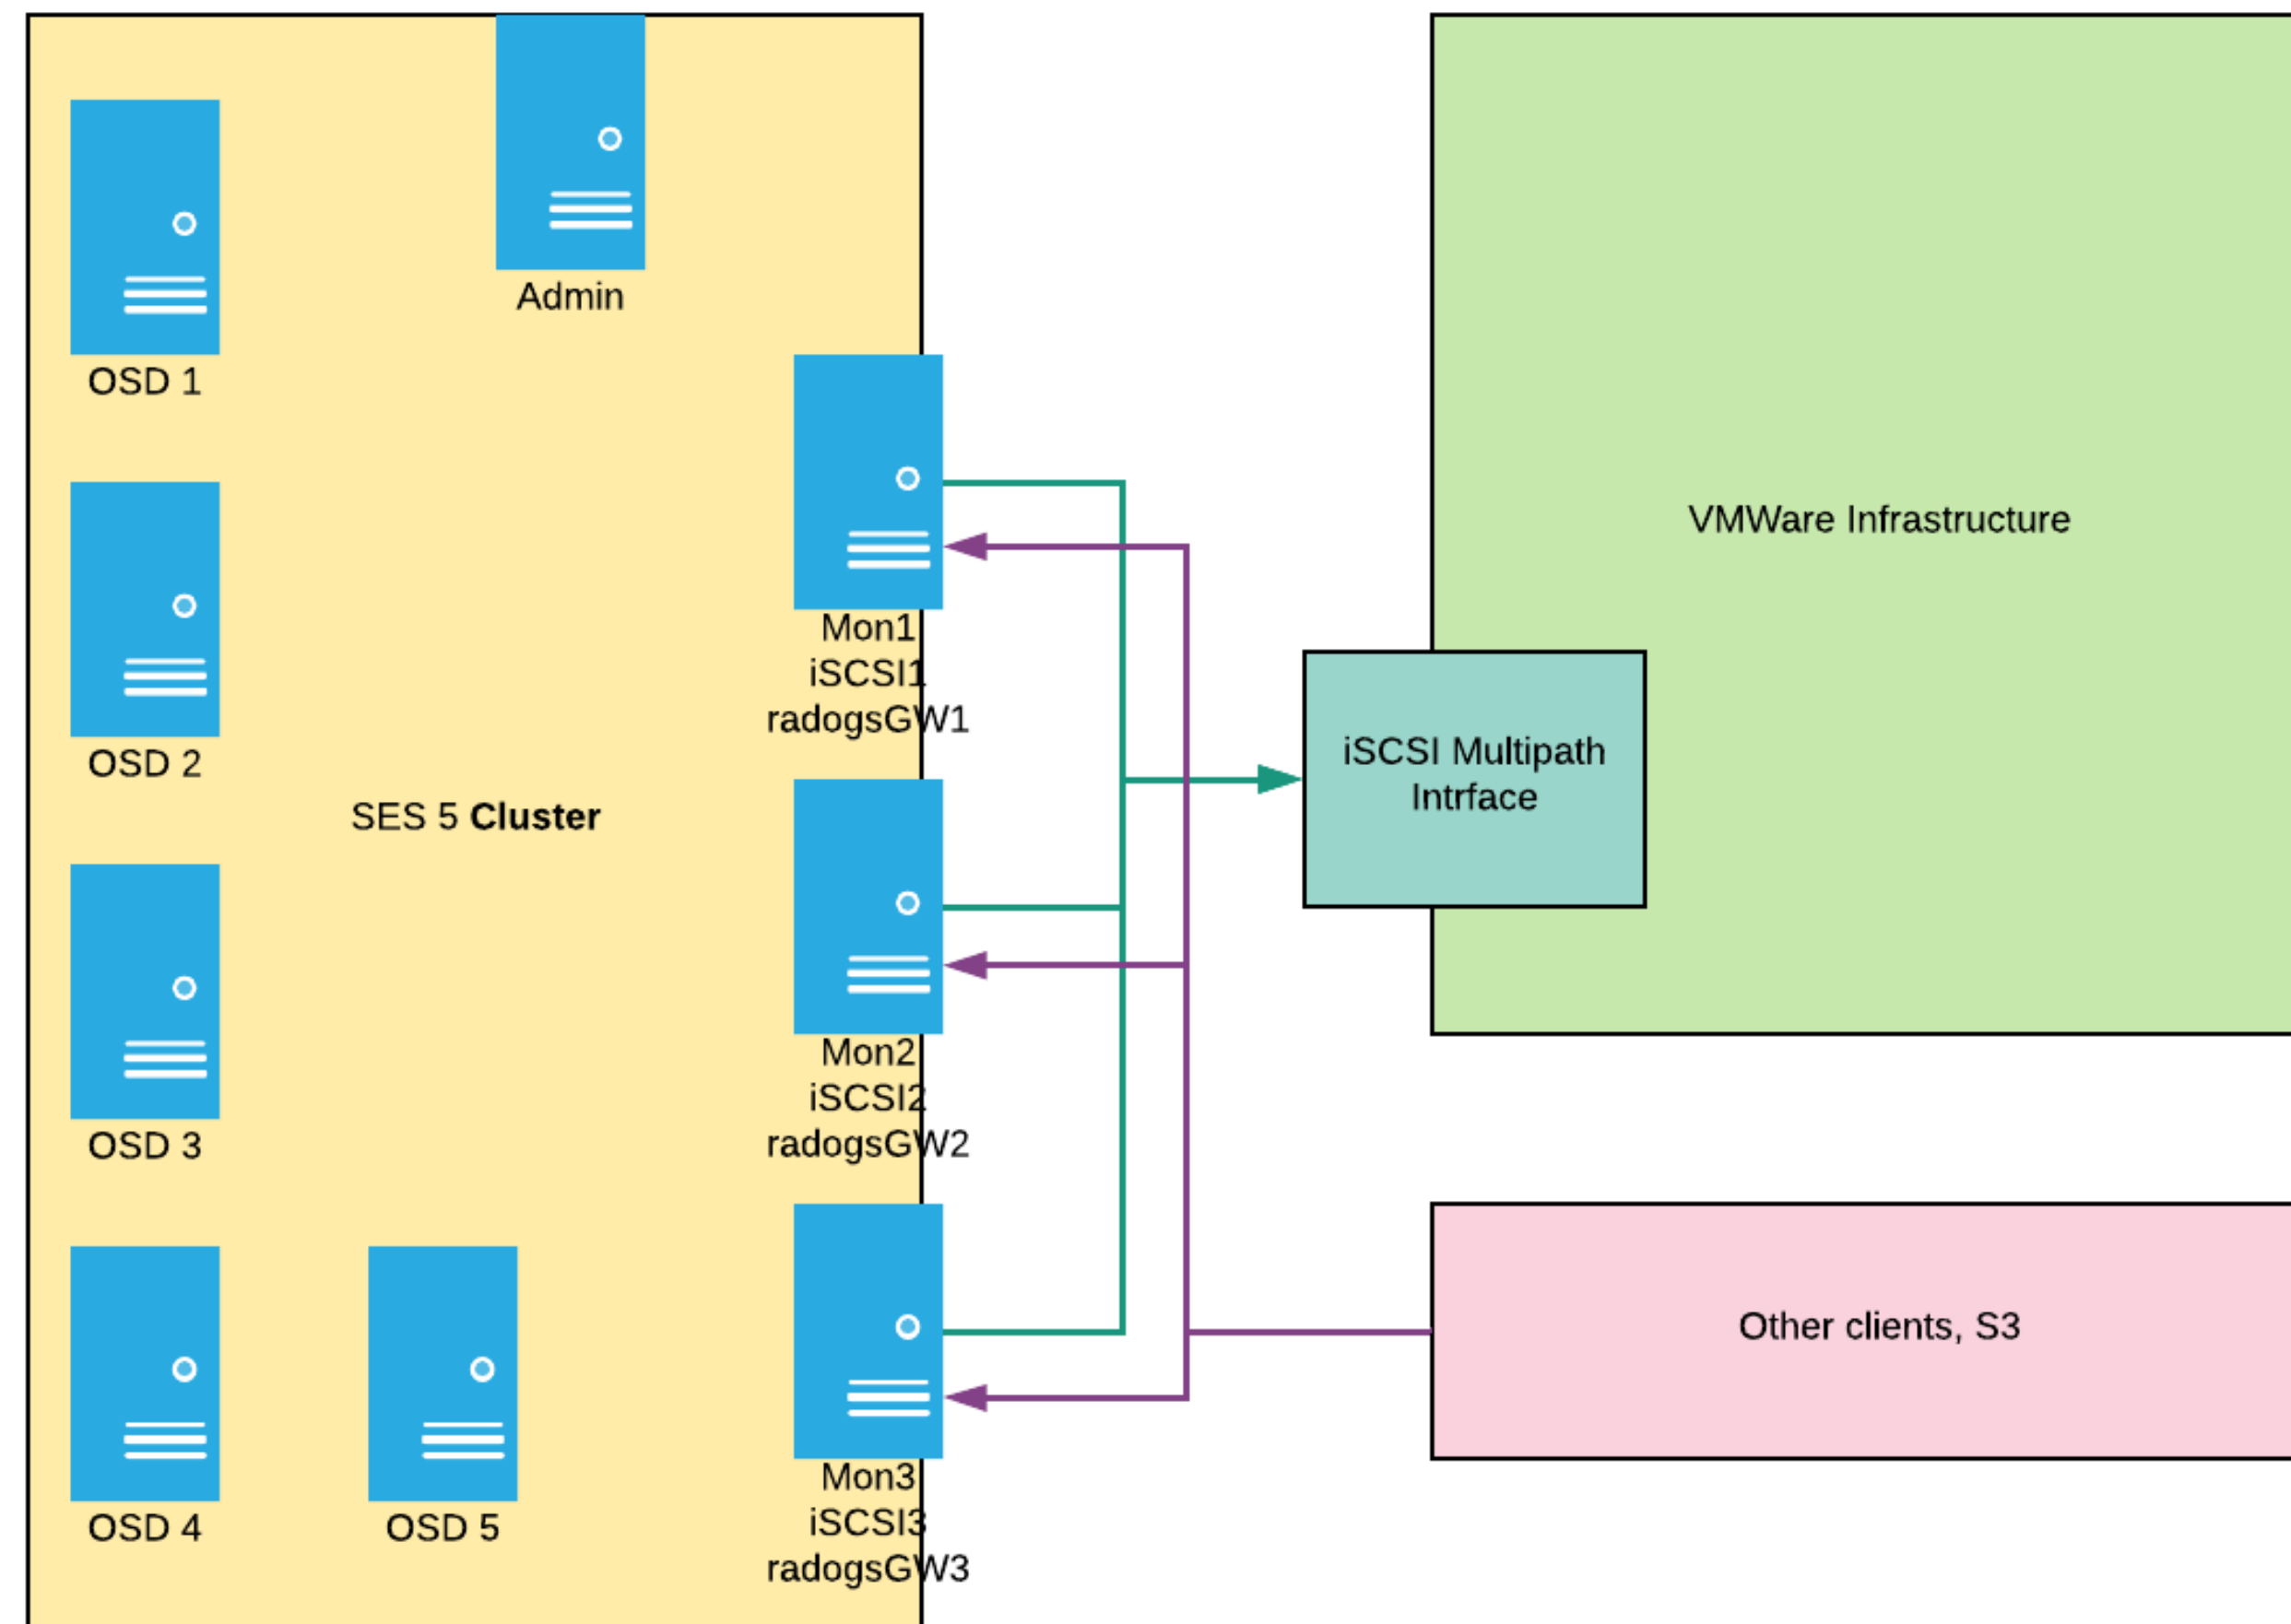


Deployment

- Initially deployed with SES 5 on vanilla hardware with bluestore
- Design was initially all spindles, (2 NVMEs were available because erroneous ordering)
- Pre-work was 0.5 man days
- Deployment prep and discovery run took less than 30 minutes
- Actual Deployment took less than 3 hours for all nodes, OSDs and Monitors
- Service Availability for customer testing, 4 hours after initial start.



Architecture

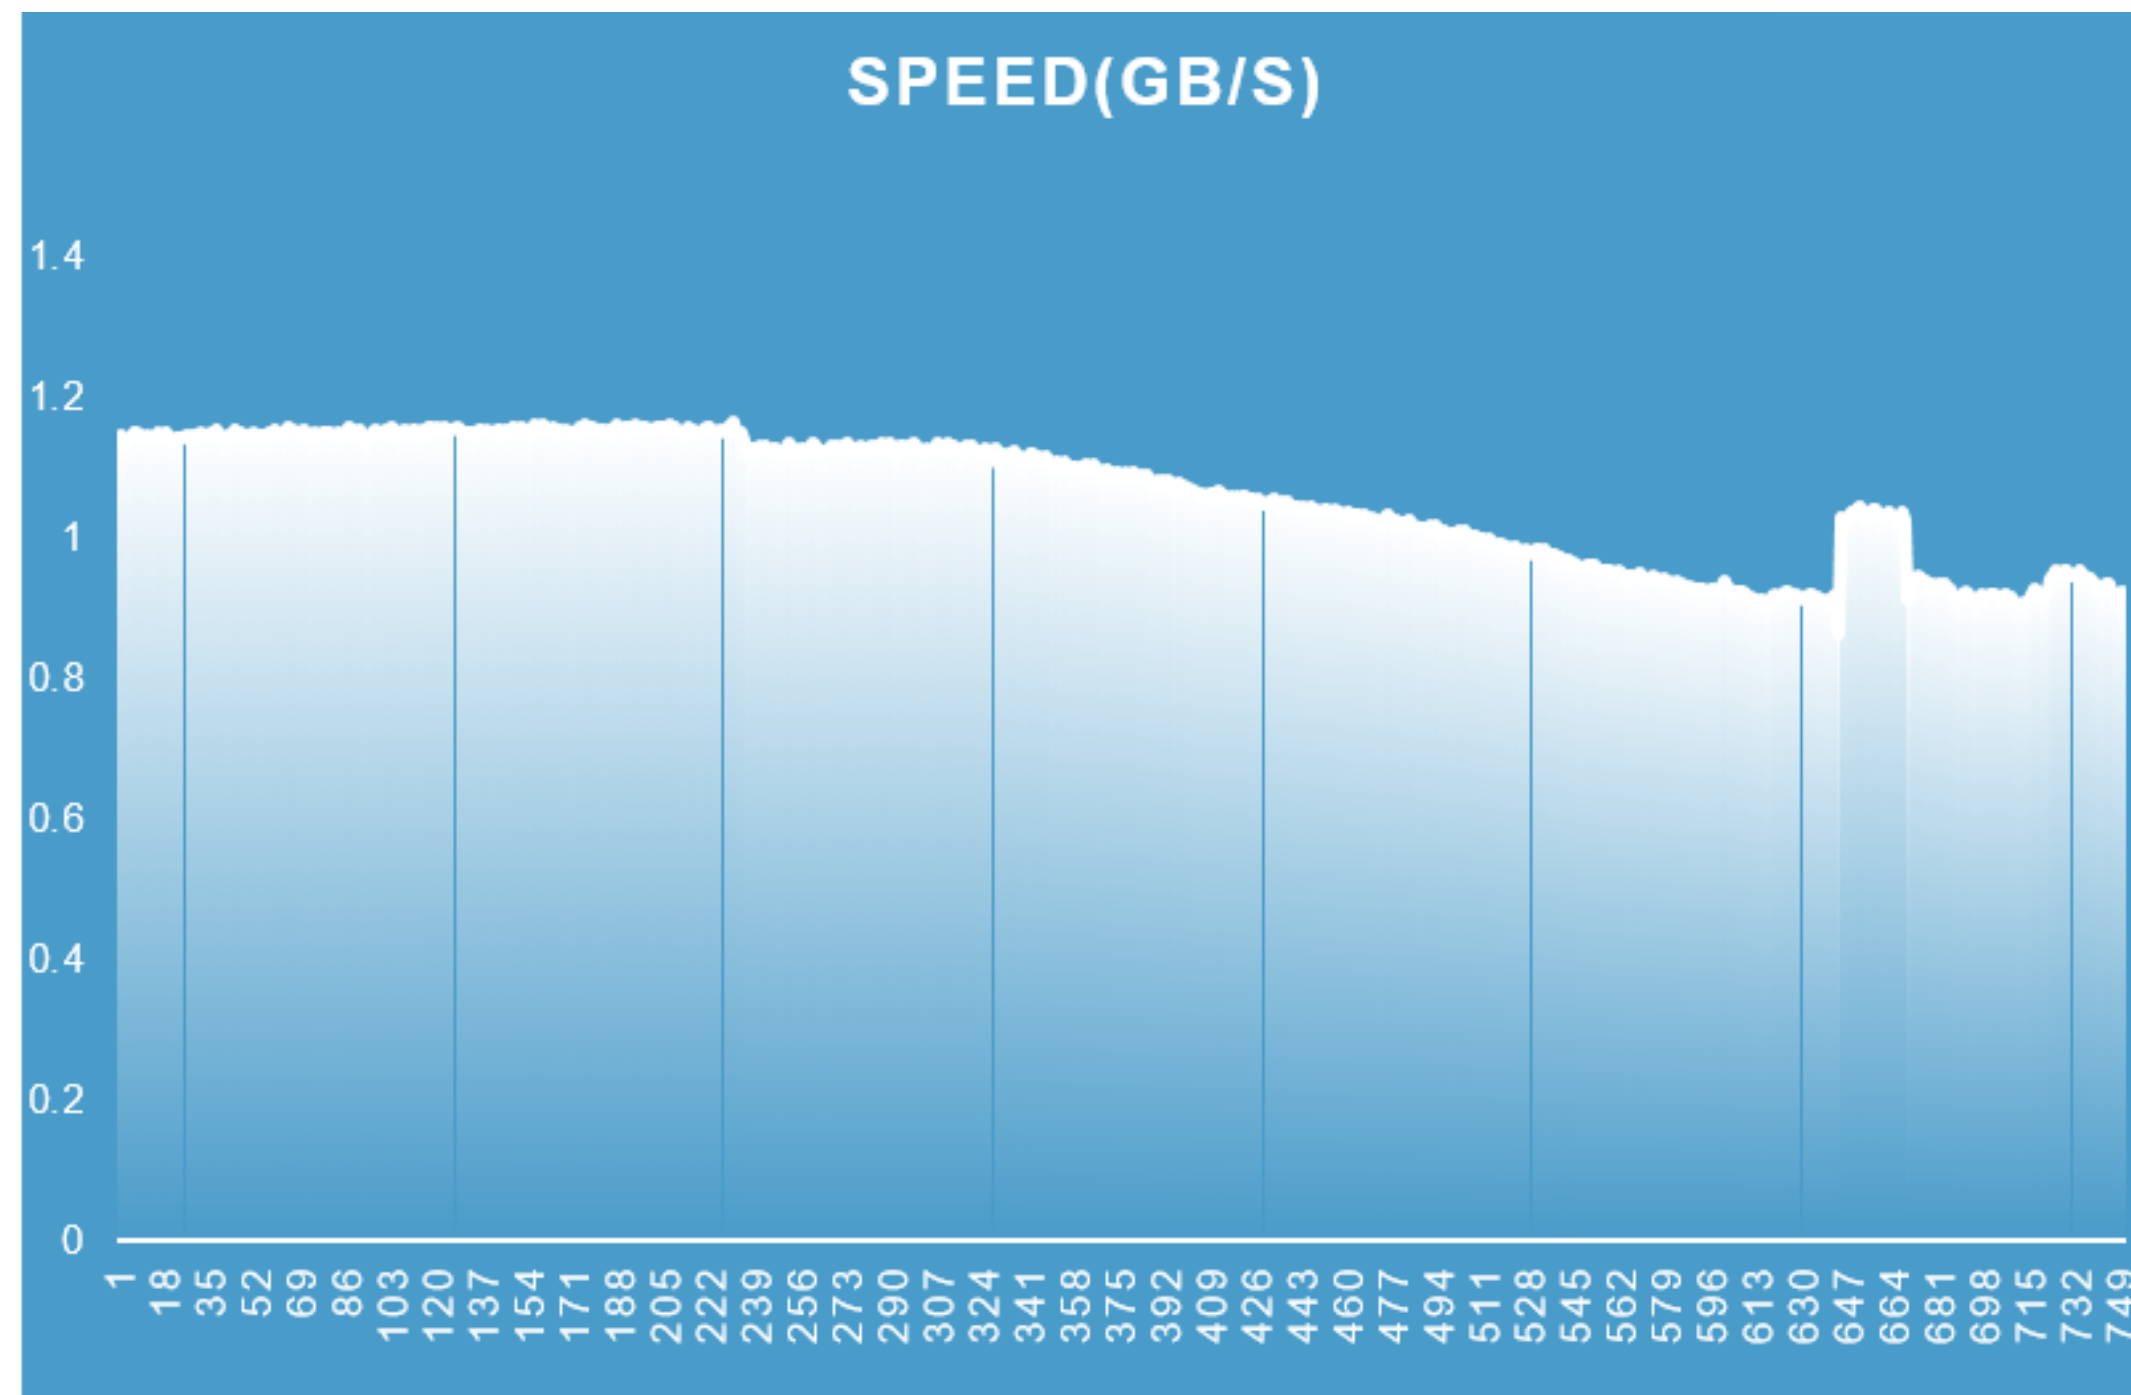


Pitfalls

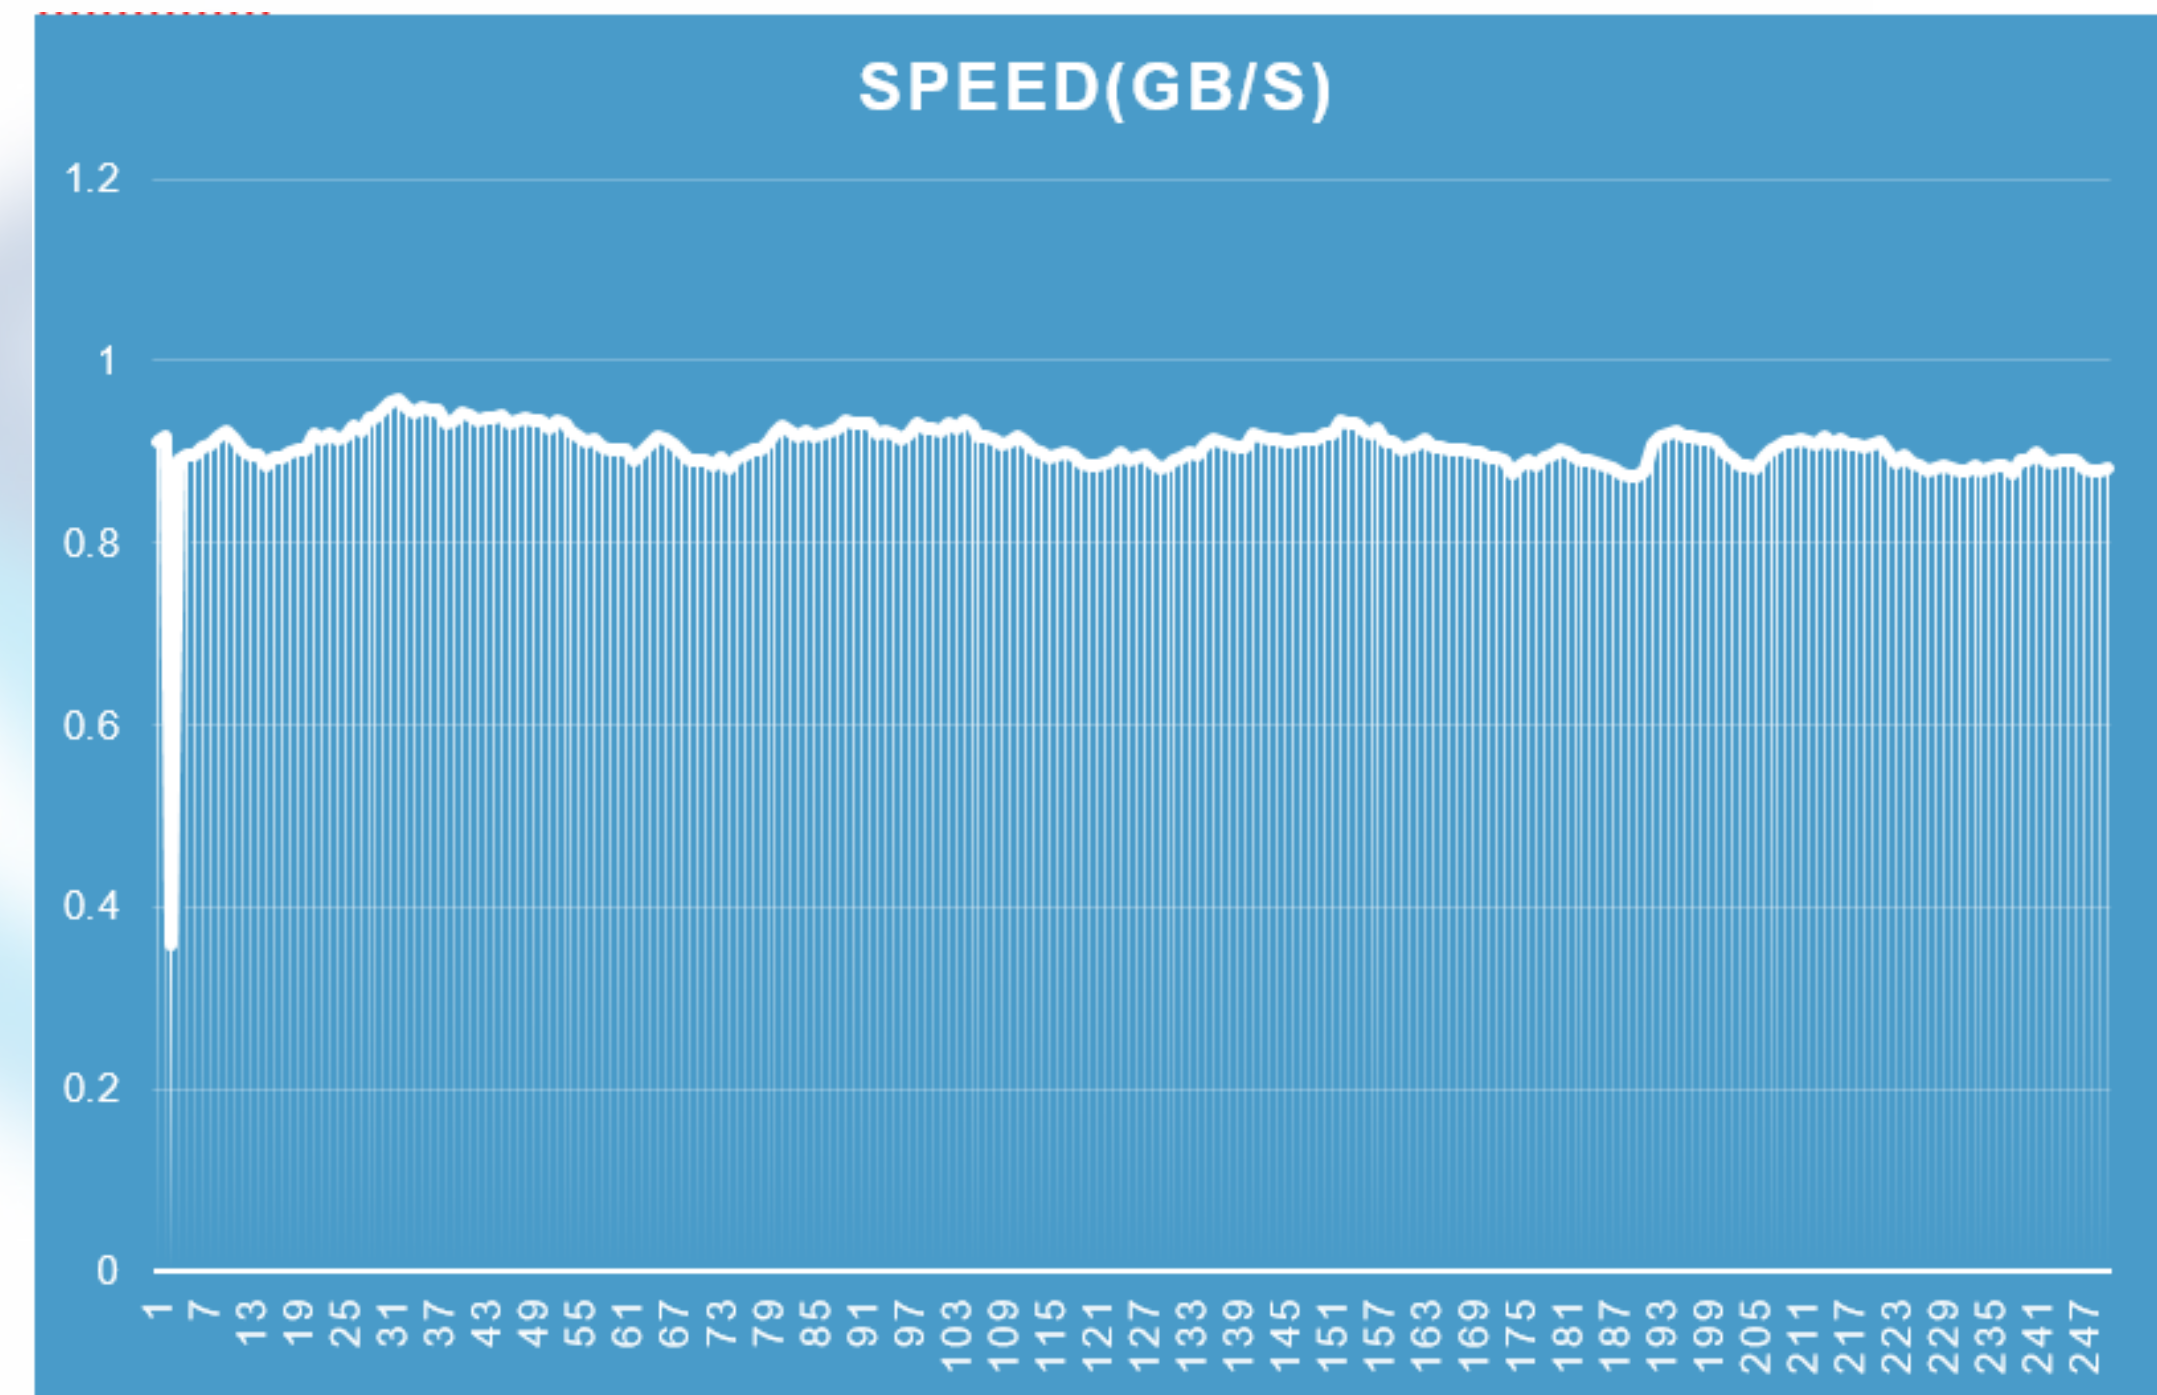
- SES 5.5 was released 3 days after deployment...
- iSCSI Gateway was a "requirement" and a stumbling block
- S3 Gateway HA was misconfigured
- NVME WAL/DB mishap



When it all worked...



72 hours, 4MB Blocks, 200GB Sets



40 hours, 4K Blocks, 200GB write sets



Case 2: Finnish Meteorological Institute (FMI)



Datalounges



nextcloud

kubernetes

SUSE

ceph

openstack.

Design and Decision making

- Design was also relatively simple but became complex
- Ceph cluster replication was required
- Initial CEPH cluster was already present
- Licensing was a big issue
- Local Expertise was needed (and still is)



Deployment: 1

- Initially deployed with SES 4 on vanilla hardware with filestore
- Upgrade to SES 5 went without a hitch but with service interruption and complexity
- SES 5 was then migrated to Bluestore with NO service interruption
- Upgrade to SES 5.5 was performed as a rolling upgrade
- Cluster Expansion went without a hitch and only 25% of performance drop. Per OSD node replication
- Each New node was brought in 1 by 1 due to workloads on the cluster
- Total workload speed improvement was almost 3 times of SES 4.
- RadosGW deployment was new with multi homed gateways running on 2 different networks on the same node



Deployment: 2

- Due to nature of workload, Cluster had to be replicated
- Feature not available at the time of design so it went async
- Replication of data is a complex script that runs every 10 minutes to sync the data off to another location
- Monitoring was an issue, especially Logs and error detection for both clusters in a single location



Pitfalls

- SES 4 was not SALT based installation
- Hardware failure during upgrade
- Replication script needed a lot of work
- Monitoring requirements needed a special solution



When it all worked...

- Total space went from 400TB to 800TB on each location
- Throughput went up by about 40%
- Access to RadosGW was available through https to partners
- Ganesha NFS was available for internal users (scientists) as well as internal S3
- Log analysis works in realtime with statistics and error alerting right away based on location
- Things still ongoing especially with the replication and management



Thank you for watching the show, questions?



 <https://www.datalounges.com>

 @datalounges



Datalounges

