

# Spectrum Scale Cloud Deployments - Solutions

Christopher D. Maestas

Ashutosh Mate

Global Technical Architects  
IBM Systems Storage

# Deploying and running Spectrum Scale on Open Compute Project (OCP) hardware

# What is OCP?

A rapidly growing, global community whose mission is to design, use, and enable mainstream delivery of the most efficient designs for scalable computing.

As a founding member, Facebook partnered with IBM to evaluate IBM Spectrum Scale, the industry's leading global and scalable, clustered file-system.



**OPEN**  
Compute Project



**OCP U.S. SUMMIT 2017**

Santa Clara, CA

# Spectrum Scale runs flawlessly on OCP platform

IBM Storage & SDI

OCP  
Server 1

4 OCP servers each having:

2 sockets

28 Intel cores

256GB memory

1 SAS raid controller using RAID 6

1 10Gbe

2 OCP KNOX storage arrays each  
having:

20 X 8TB drives

OCP  
Server 2

OCP  
Server 3

OCP  
Server 4

KNOX Storage Array 1

KNOX Storage Array 2



Facebook deploying IBM Spectrum Scale on OCP platform

## Hardware is what?

- Server, storage, network switches and cards
- Firmware updates and are they consistent?
  - Found raid and NIC firmware out of date
- Hardware Raid or no RAID
  - Use Spectrum Scale Replication only?

## Consistent OS installation?

- RedHat or CentOS kickstart!
- Verify running kernel and kernel headers are the same
  - `uname -r && rpm -qa | grep kernel`
- Verify gcc binaries are installed
  - *python-lxml if you want object!*

# Episode II - The software can strike back!

IBM Storage & SDI

## Configure ssh keys

```
ssh localhost uptime  
ssh $(hostname) uptime
```

## Verify time is accurate on all nodes

date or see uptime output

## Update /etc/hosts with hostnames/ IP's

You can try to rely on DNS, but better safe than sorry

**Did you know?** /etc/hosts: IP, FQDN, shortname  
DEFAULT for localhosts is incorrect on some dists!

## Verify storage on each node to be used as NSD disks

```
fdisk -l /dev/DEVICENAME (e.g. sdb)
```

## SELINUX and Firewalls?

Use or disable them – *search Spectrum Scale docs:*  
getenforce, setenforce, systemctl disable|stop firewalld

## Repositories or local CD?

```
yum check-update
```

# Spectrum Scale runs flawlessly on Open Power OCP platform!!!

OCP Server 1

OCP Server 2

2 OCP Barreleye servers each having:

2 Sockets

20 Cores

256 GB memory

1 SAS raid controller using Raid6

1 10GbE

1 OCP KNOX storage array having:

15 X 4TB drives

KNOX Storage Array 1



- Facebook deploying IBM Spectrum Scale on OCP platform



# Open Power Notes

Install look a bit more  
but if we remember  
our episodes

In the beginning –  
use consistency

The software will  
strike back

So it took an hour  
after resolution



# What we accomplished and where to next?

## ✓ 30 mins to install

- ✓ Facebook driven
- ✓ SNC cluster type

## ✓ Chef based Spectrum Scale Installation Toolkit

- ✓ Install GUI and CLI based
  - ✓ Rpm or deb based install still possible

## ✓ Zero changes to scale running on OCP

- ✓ Software defined solution

## ✓ Functional and Performance testing on 10 Gb Ethernet

- ✓ Tested I/O bandwidth of 1+ GB/s

## ✓ Monitoring Framework - Zimon

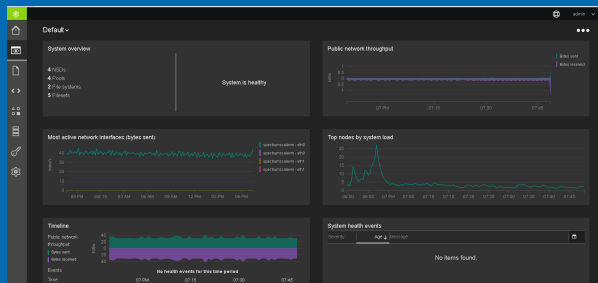
- ✓ Integrates with grafana

## ❑ REST API released and updated

## ❑ More Storage Options and other OCP hardware configurations

## ❑ **Excelero and others – but why?**

## ✓ Open Power testing with Spectrum Scale



# Other opportunities

## Excelero Opportunity

Software defined block storage – 20 TB DME to client

## Other OCP partners

Yahoo Japan

Rackspace

Applying same concepts Cisco Versastack V7000 Unified replacement

Spectrum Scale Cloud Gateway project!

# Spectrum Scale and the weather



Long, long ago  
GPFS 3.5 and AFM  
3 site scenario  
data source to site A,  
push to site B and then push  
to site C



# Spectrum Scale Deployment

Install toolkit

GUI to monitor failed NSD drives

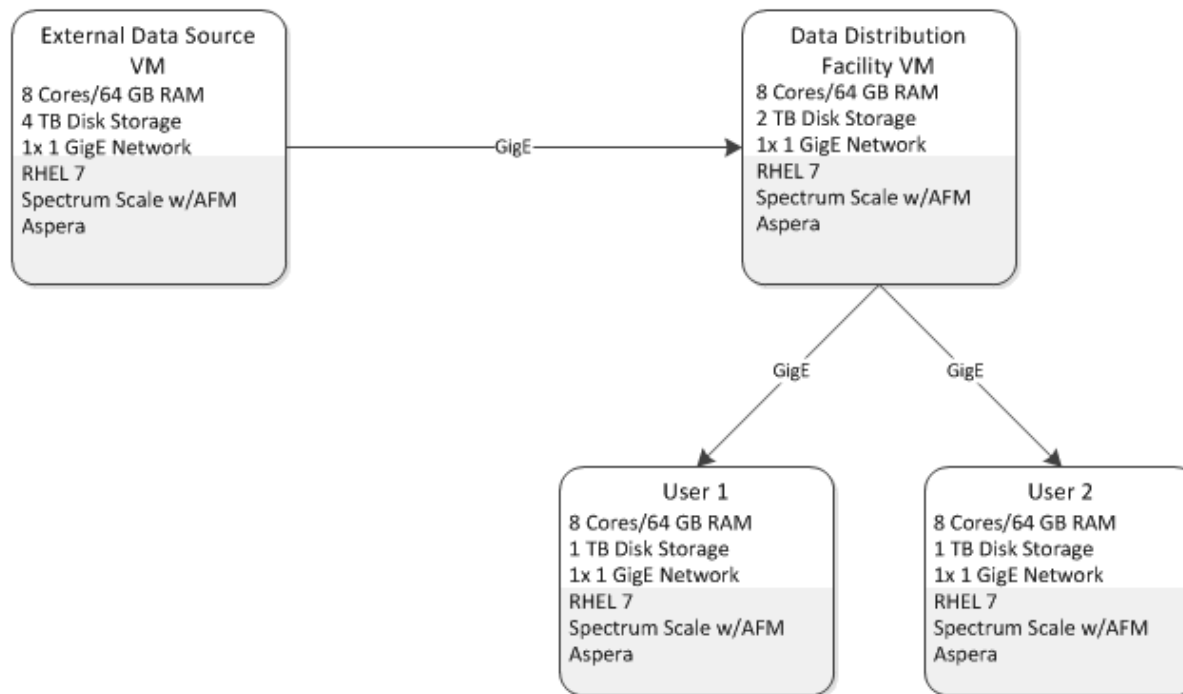
Easier to let new operators experience the install and maintenance of the system

Repeatable ... regardless of version

4.2.1 to 4.2.2.3

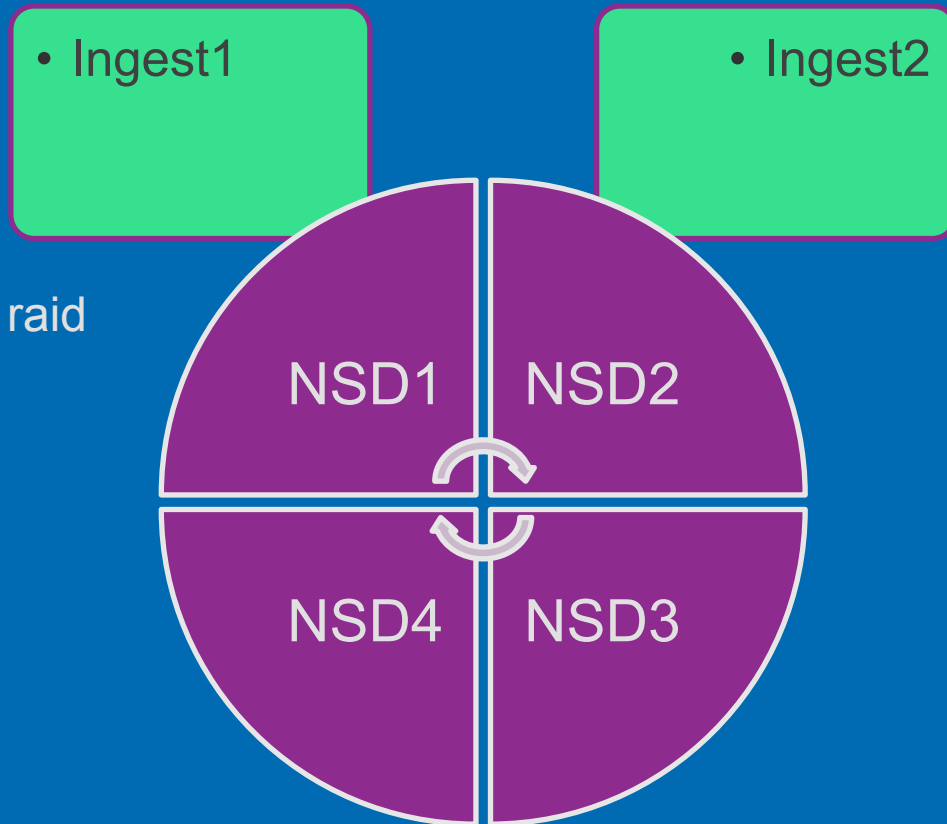
# Initial Bluemix (Softlayer) Architecture

## SoftLayer



# Bluemix on bare metal testing

2 ingest servers talking Aspera  
4 NSD nodes with 24 drives  
use replication instead of hardware raid  
controller – **but why?**  
10 Gige, can fill the pipe





# Bluemix current prototype

- 2 ingest servers talking Aspera
- 2 NSD nodes with 12 drives each
  - use replication instead of raid controller
  - 10 Gige, can fill the pipe
- 2 Clusters, 2 sites
  - Initially write to both sites at once
  - Let Scale do AFM transfers within Bluemix
  - No cost to move data in Bluemix

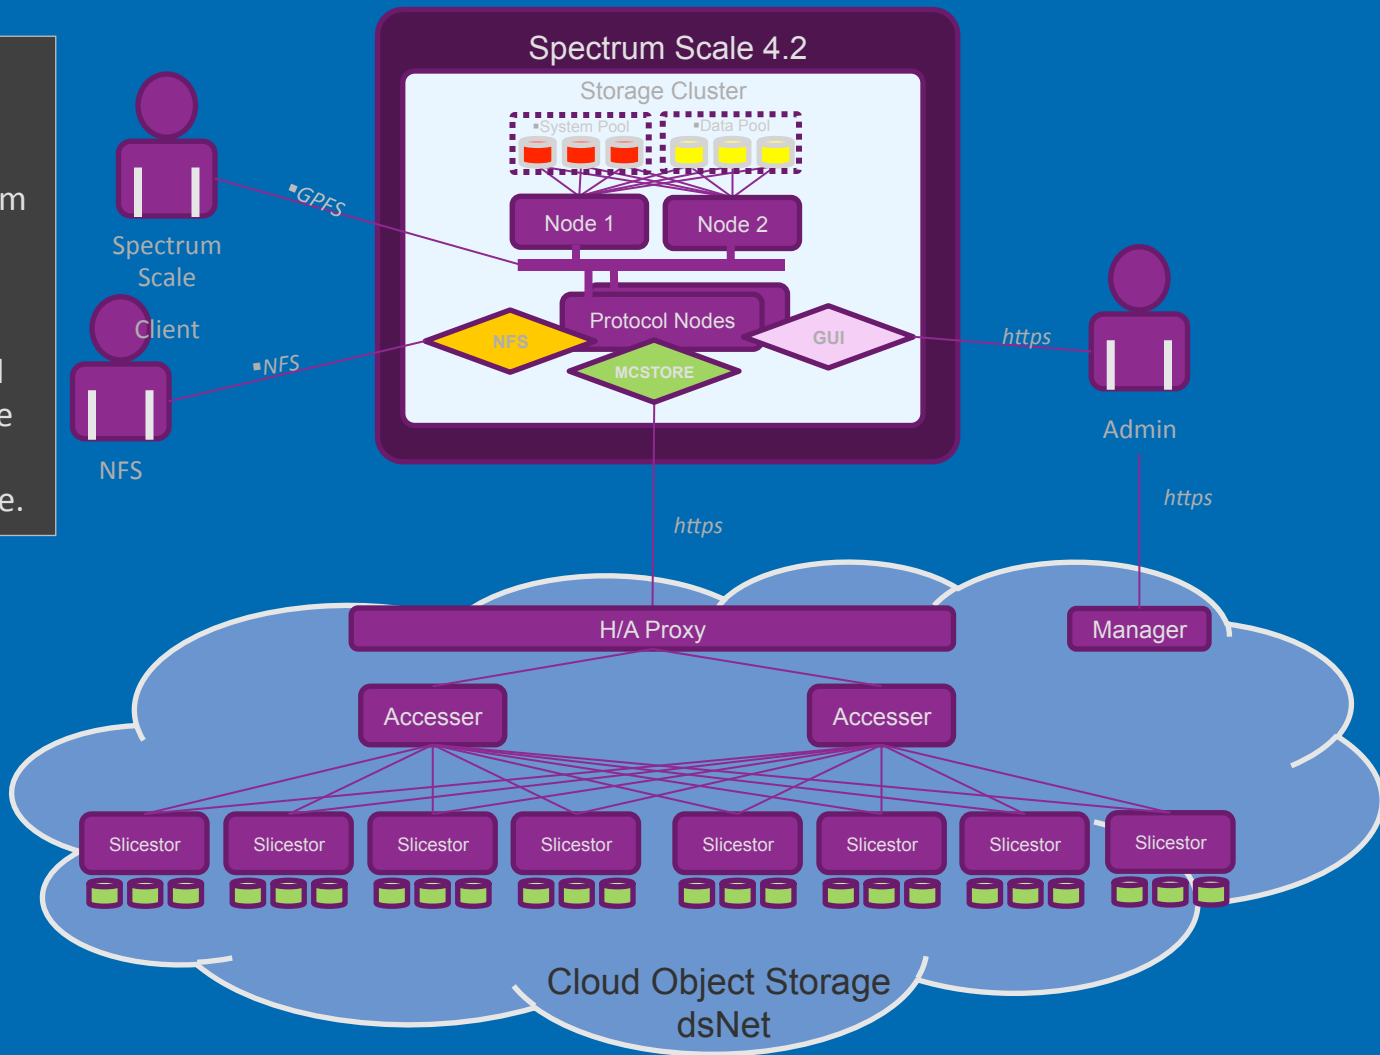


# City of Hope – Spectrum Scale Cloud Gateway v1.0

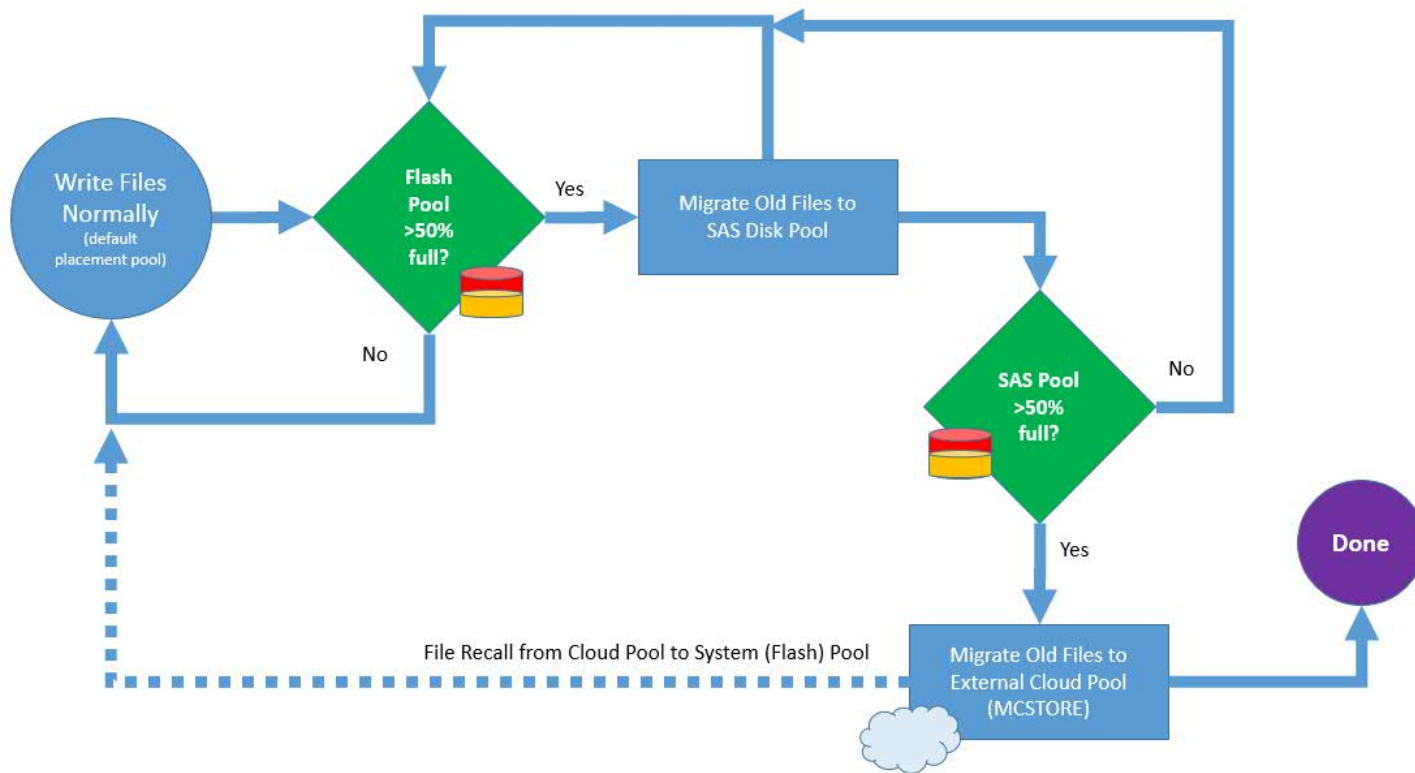
## Functional POC

Using Development Lab System  
in Tuscon

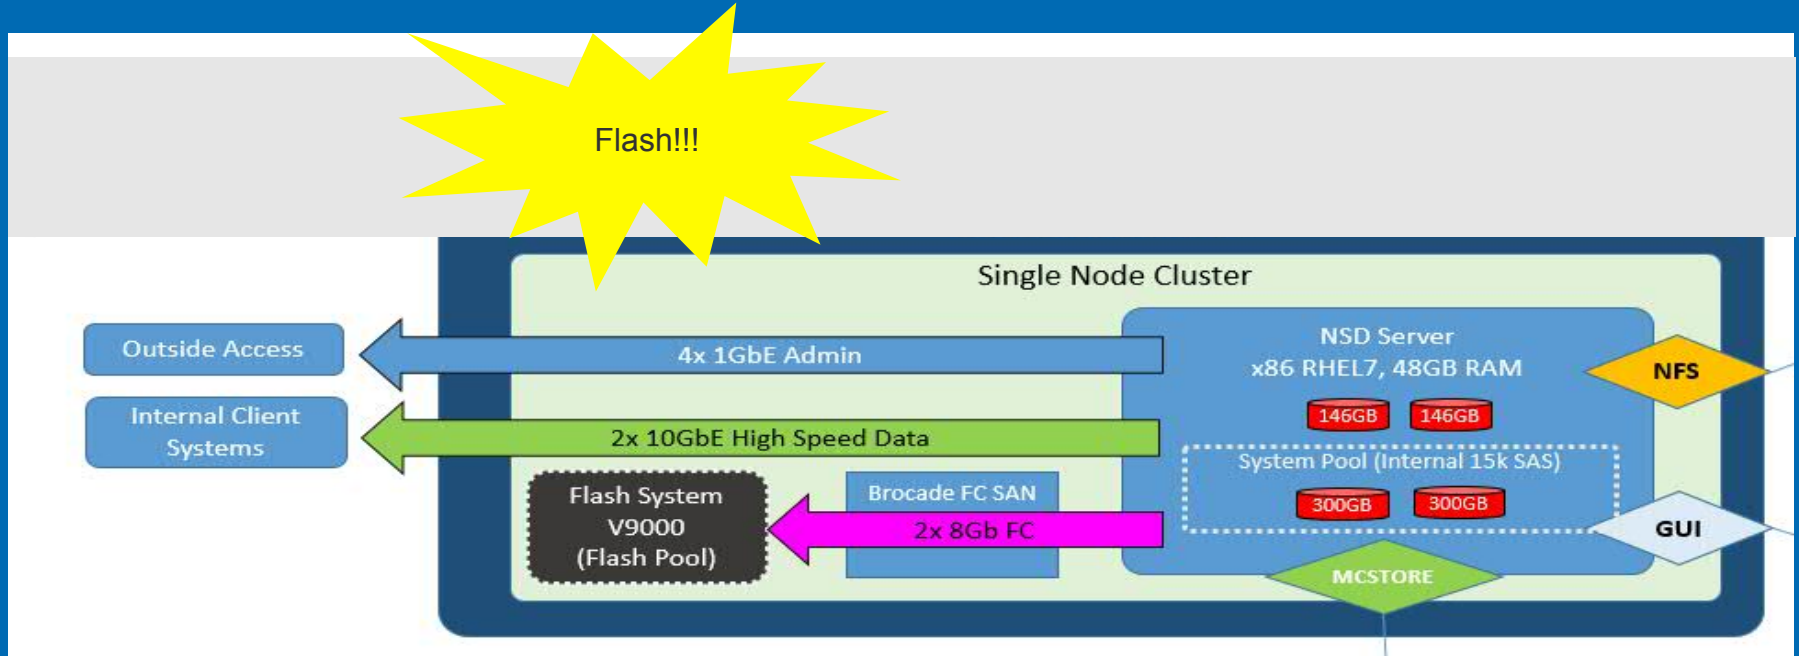
Note: All systems are vitalized  
(VMs) and share a single, large  
disk system. This is for  
functionality, not performance.



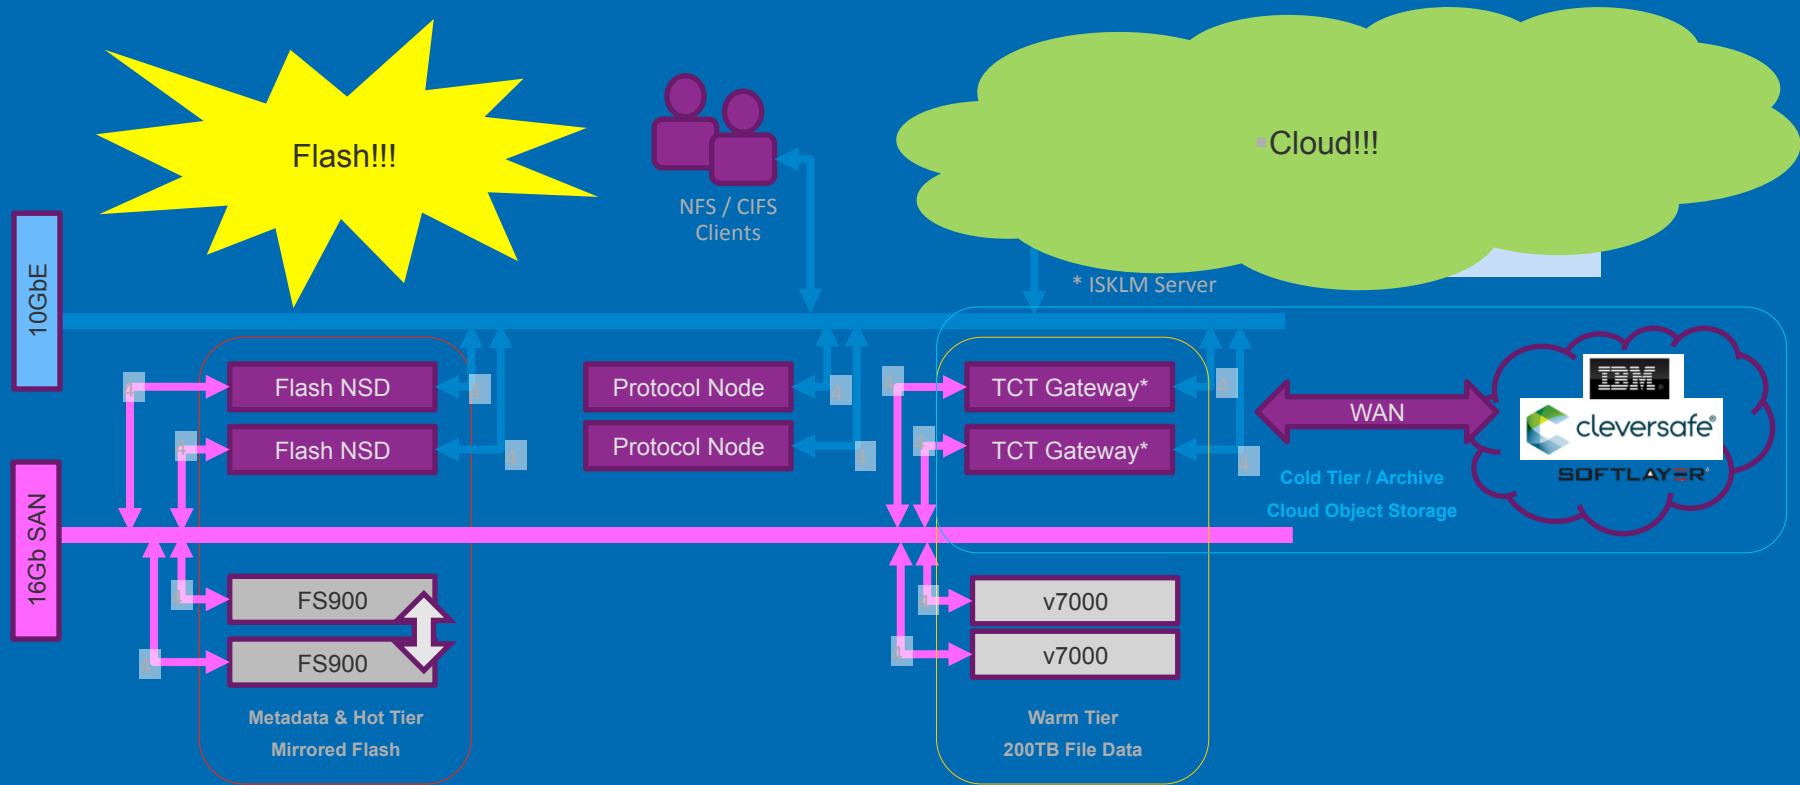
# Demo Workflow – Migrate Files based on Threshold and Age



# Proof of Concept onsite



# Production Environment: FLASH + CLOUD == FLOUD!



Note: Each Flash NSD will have 2x 16Gb SAN Connections to each FS900 (criss-cross). The same goes for TCT Gateways and v7000 Controllers

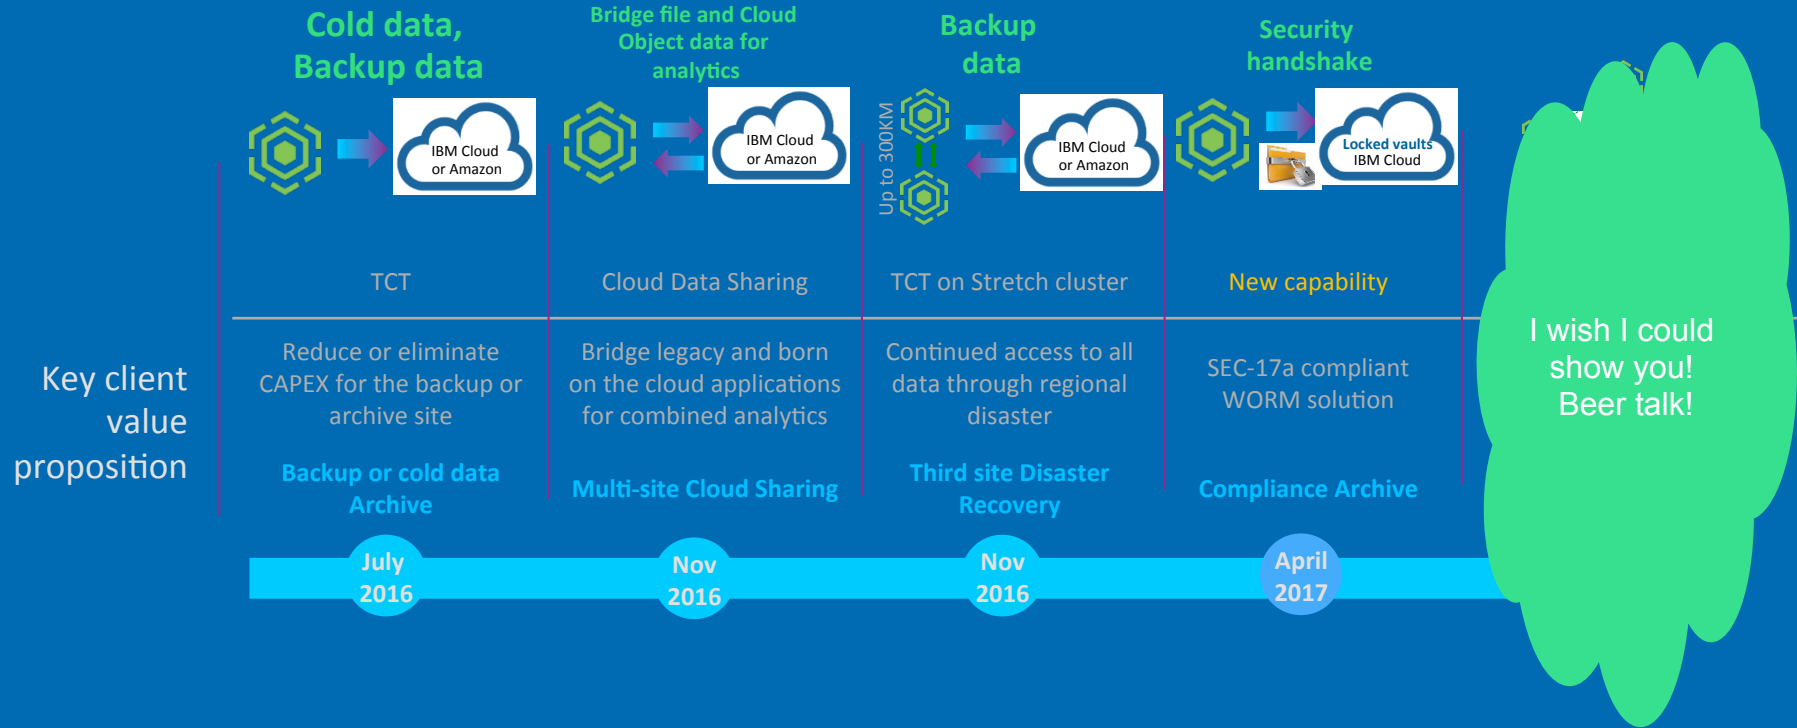
# Feedback - Spectrum Scale Cloud Gateway

Ashutosh Mate

Christopher D. Maestas

# Building out Hybrid-Cloud offerings with Spectrum Scale

IBM Storage & SDI





# Objective

Create a Spectrum Scale Cloud Gateway (SSCG) as an on-prem appliance offering to move data into IBM Cloud Object Storage (ICOS) on IBM Cloud:

- ✓ Use Spectrum Scale with NFS, SMB, Swift, S3, AFM & TCT along with the single name space
- ✓ Could bundle the SSCG on a standard hardware config
- ✓ GUI for quick configuration & management
- ✓ It's on the truck!

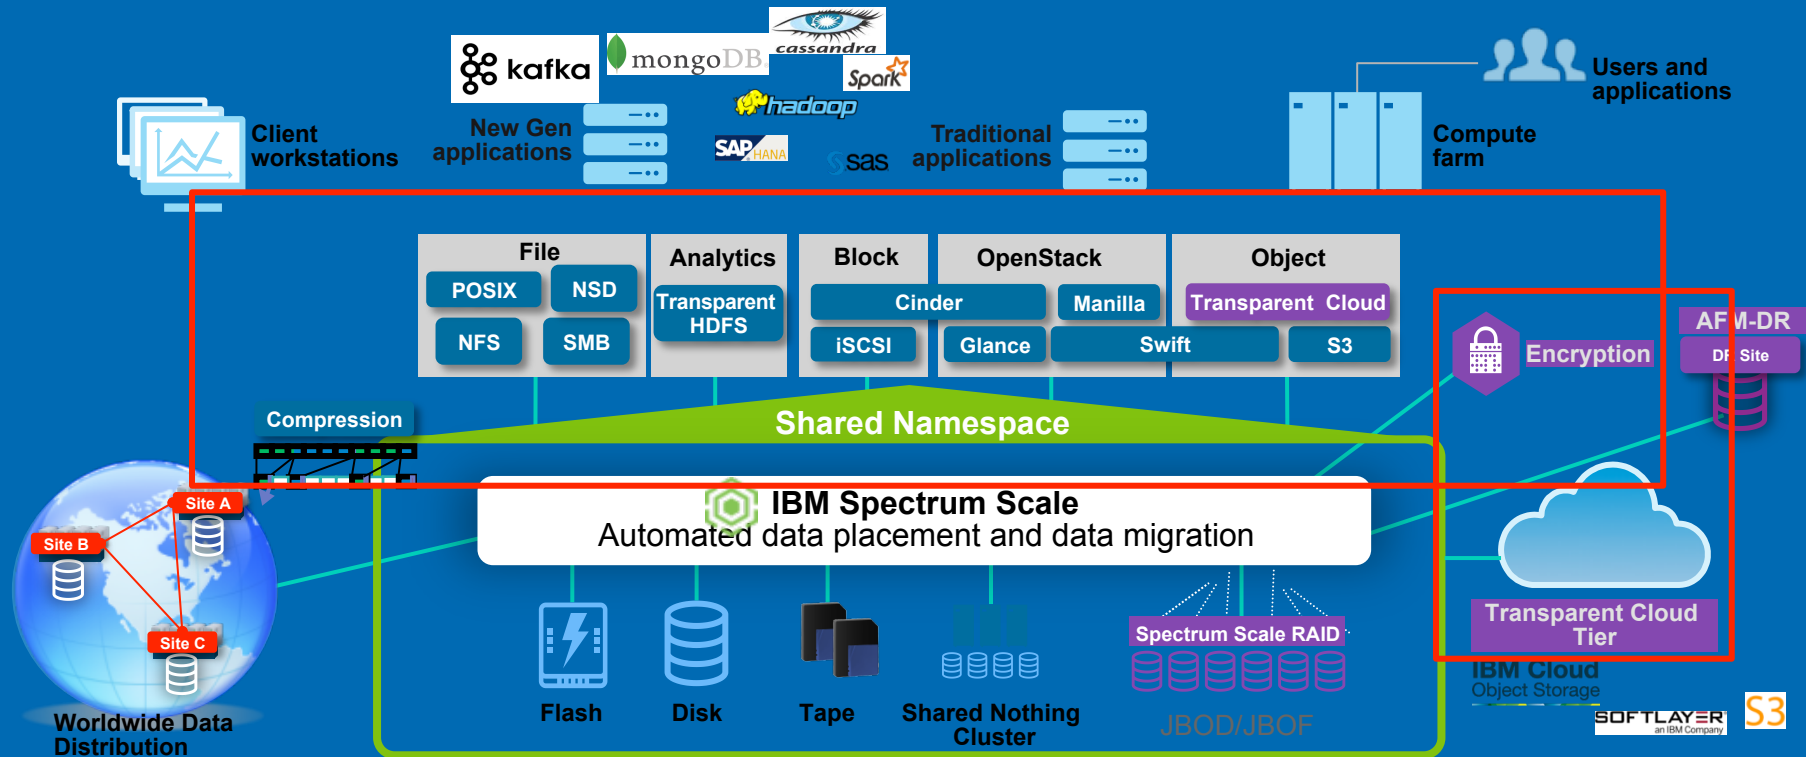
# Use Cases

- ✓ Use case is to discover old data & pull it into the SSCG via AFM & migrate it to ICOS
- ✓ Another use case is to allow data from USB / Flash drives to be copied into the SSCG to be uploaded to ICOS
- ✓ Could also be used to pull data from other public clouds
- ✓ SSCG would have fixed SSD storage for initial staging – FLOUD (Flash + Cloud) appliance
- ✓ Configurable ILM policy will automatically move cold data to ICOS through TCT Cloud Services

# Imperatives

- ✓ Easier to deploy, configure & manage than other solutions!
- ✓ Comparable price-point to Panzura / Ctera / Nasuni gateways
- ✓ Hardened software quality & high performance data movement (Optional: Aspera FASP)
- ✓ Scheduled Data Migration & Continuous Throttled Migration
  - ✓ QOS integration
- ✓ Transparent Cloud Tiering & Cloud Data Sharing support
- ✓ Data availability during migration

# Spectrum Scale: Unleash new storage economics on a global scale



Consolidate all your unstructured data storage on spectrum scale with unlimited and painless scaling of capacity and performance

# Spectrum Scale Cloud Gateway

IBM Storage & SDI

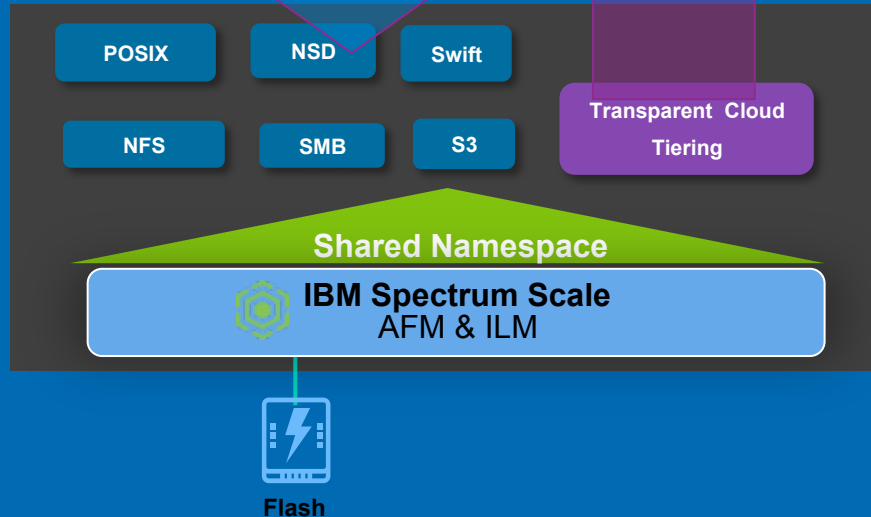
From NAS via AFM/NFS/SMB

From Mobile devices via Swift / S3

From other Public Clouds via S3

REST interface

To & From Private or Public ICOS via TCT  
To Public ICOS via TCT Export

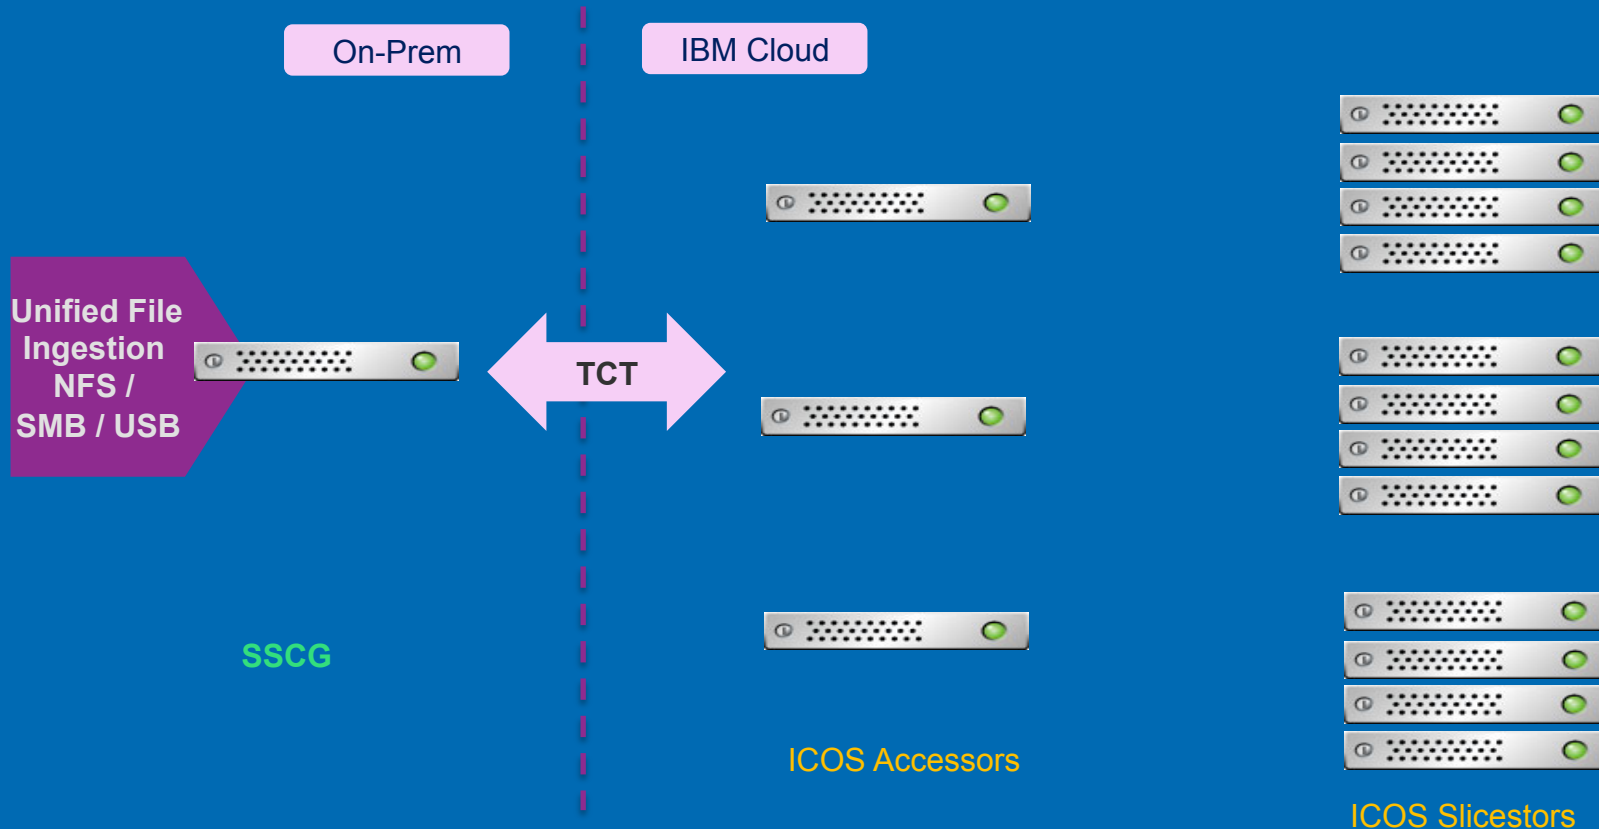


IBM Cloud  
Object Storage

SOFTLAYER  
an IBM Company

S3

# SSCG Deployment



# Deployment Models

City of hope is about 50 TB of FS900

SMALL - 6TB – Intel based

- Reuse existing COS appliance F5100

MEDIUM - 12 TB – 812LC and

LARGE - 32TB – 822LC

- depending on 4.2.3 PTF with TCT Power LE support

Does this make sense?

... and Spectrum Scale rocks!



Thank You.  
IBM Storage & SDI

# Legal notices

Copyright © 2016 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 1 0504- 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 inquiries, 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.

# JUJU charms!



- Spectrum Scale deployment with Juju charms
  - Juju is an open source service deployment tool
  - Charms is a software component that contains instructions for deployment
  - Spectrum scale Juju charms now available for deployment and configuration
  - Integrated to work with base openstack charms
  - See: <https://jujucharms.com/q/spectrum/scale>