IBM **Spectrum Scale**

OpenStack Integration

Update on Swift, Cinder, Manila, IceTier

Harald Seipp, Leader, CoC for OpenStack Storage

With input from Bill Owen, Gaurang Tapase and many more...

IBM

Spectrum Scale 4.2 Object Highlights

- Swift Storage Policy Integration
- Unified File (POSIX, NFS, SMB) and Object (Swift, S3) access
- Improved S3 emulation
- Multi-region active-active object store
- Redpaper on Active Archive with Spectrum Scale Object and Spectrum Archive EE

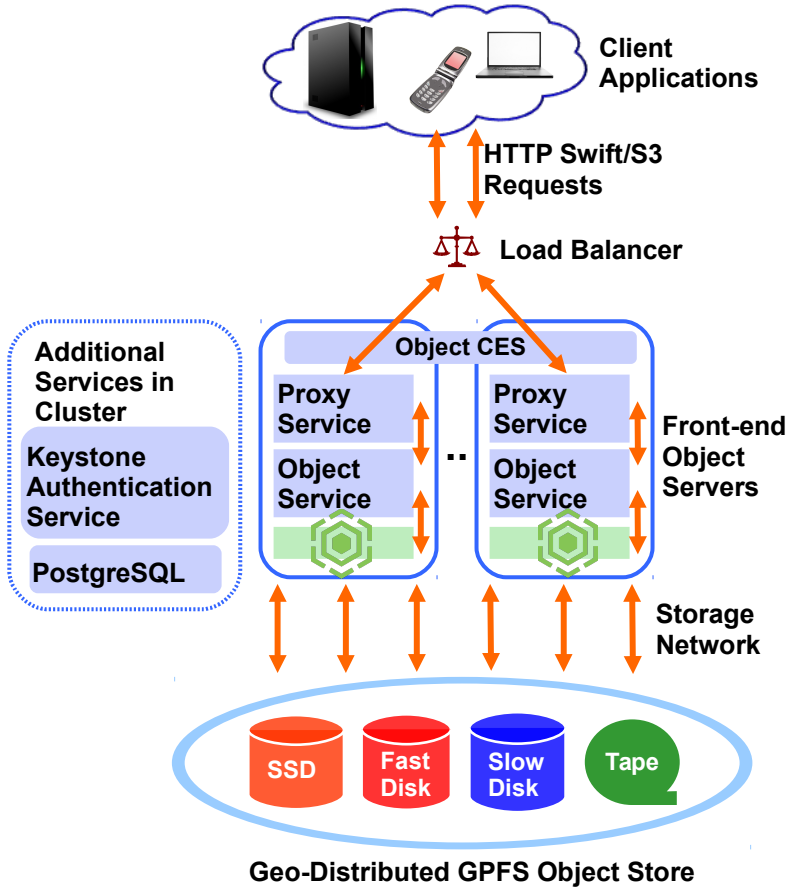

Swift Storage Policies

- 1:1 relation between Containers and Storage Policies
- Storage Policy is associated to Object Ring that defines
  - Redundancy (# of replicas)
  - Where to store (nodes)
  - How to store (device)

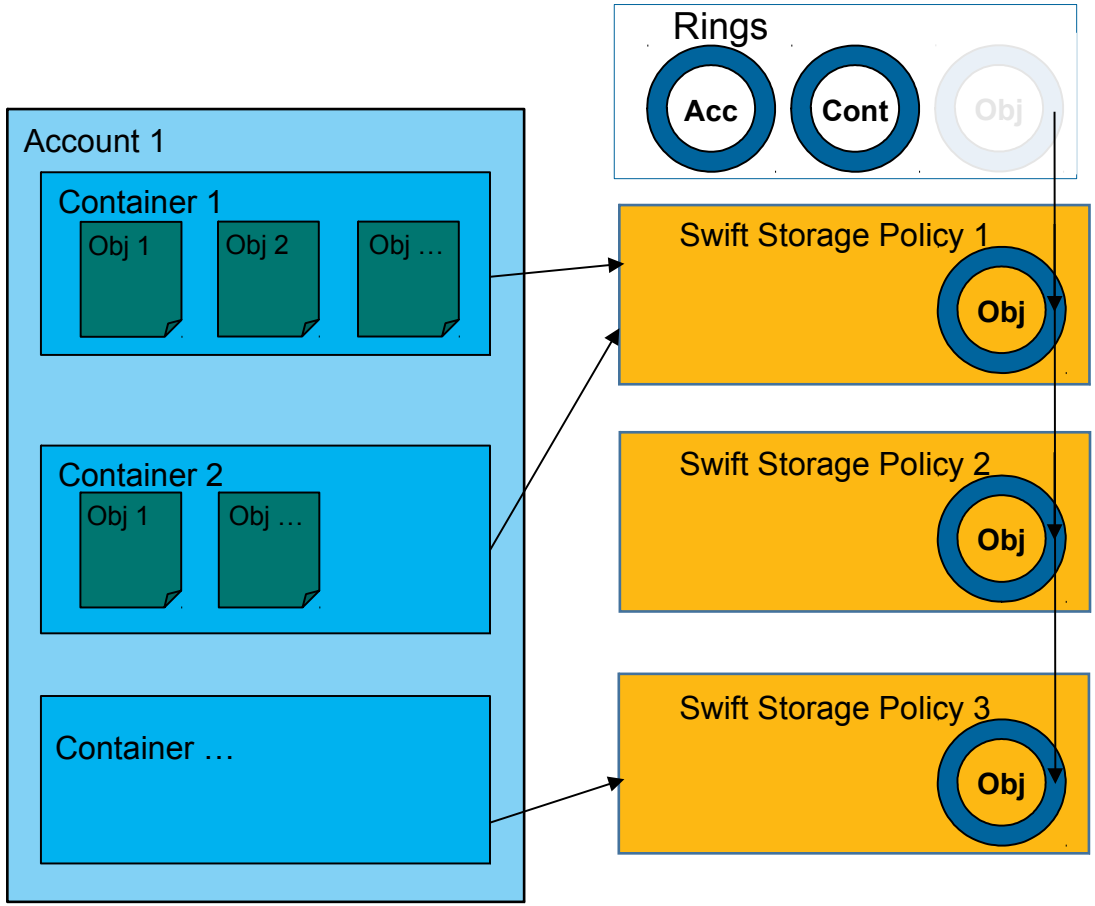

Spectrum Scale Swift Storage Policy Integration

- Spectrum Scale handles Swift “devices” as directories
- Each Swift storage policy will be mapped to an independent Fileset
- Create Storage Policy with a single command:

```
# mmobj policy create CompressionTest –-enable-compression\
--compression-schedule "0:*:*:*"
```

(run automated compression on Containers with “CompressionTest” policy every top of the hour, every day, every weekday)

- ...then assign Storage Policy as Swift metadata to a container:

```
# swift post MyCompressedContainer -H "X-Storage-Policy: CompressionTest"
```

Unified File and Object Access – Benefits

- Seamlessly integrate File (POSIX/NFS/SMB) and Object (REST) worlds
- Access cloud data through applications that need file access
- File and object data on same namespace, no data duplication
  - Enables “Data Oceans” of different types with multiple access options
- Placement policies for files can be leveraged for objects
- Plays well with the Spectrum Scale Hadoop connectors for Object data analytics

Unified File and Object Access – Filesystem layout

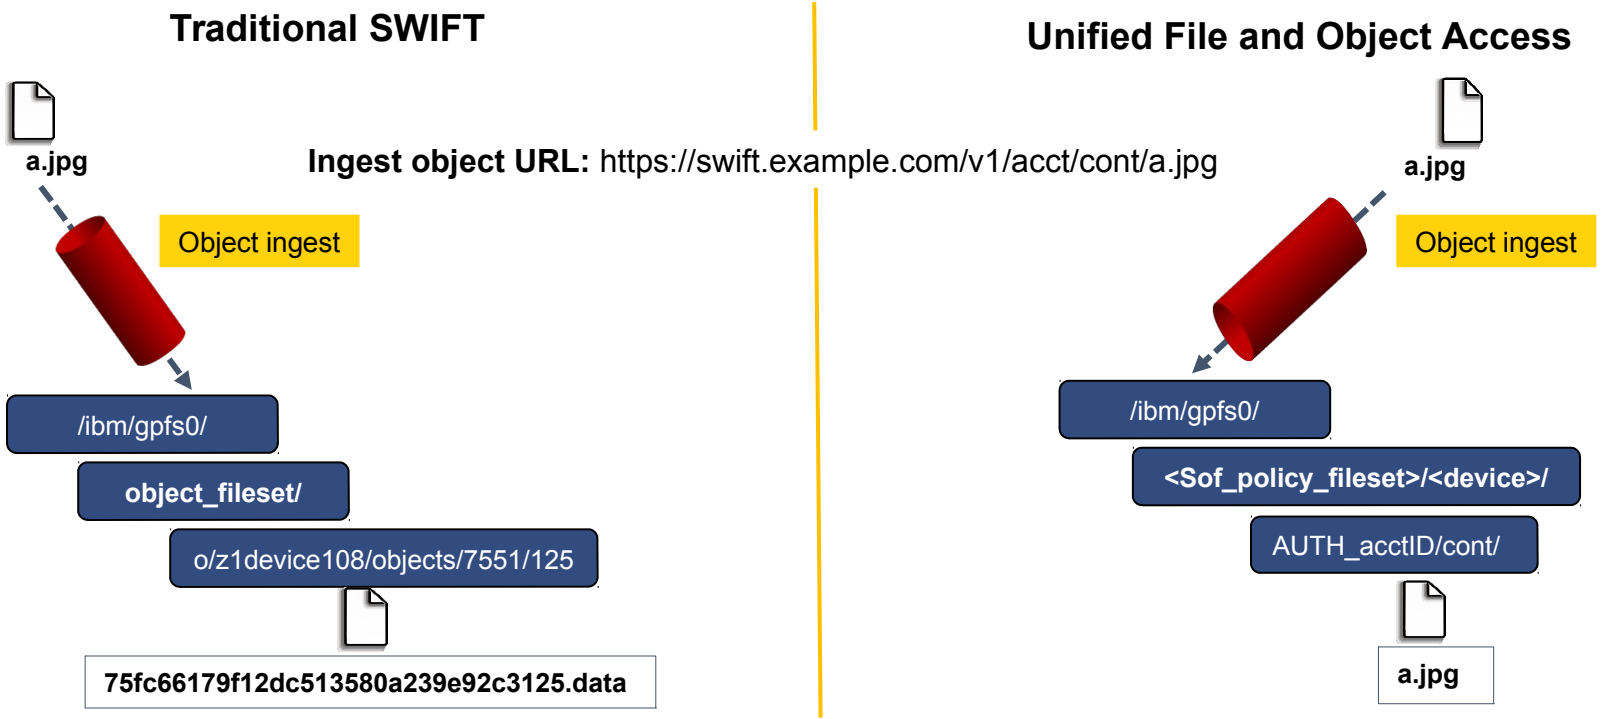

Unified File and Object Access – Ingest through Object

- Objects ingested are available immediately for File access
- Flexible ID management modes (explained later)
- Compatibility of Object and File Applications
  - Object authorization semantics are used during object access
  - File authorization semantics are used during file access of the same data

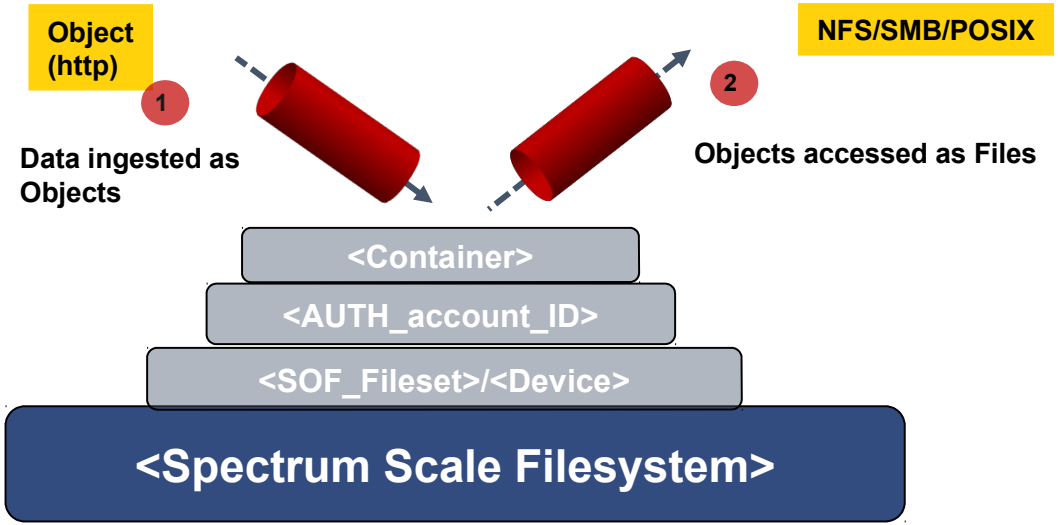

Unified File and Object Access – Ingest through File

- ibmobjectizer system service is responsible for objectization
- Objectization converts files to be available from the object interface
  - Run every 30 minutes by default, configurable
  - Immediate objectization through `# mmobj file-access ...`
- New files need to be visible to the Swift database to show correct container listing and container or account statistics

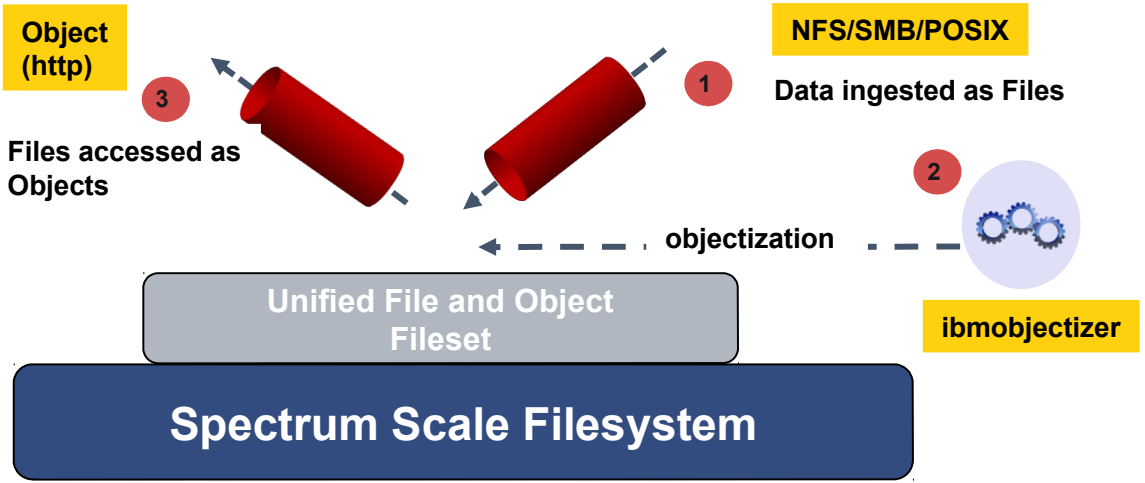

Unified File and Object Access – Flexible identity management

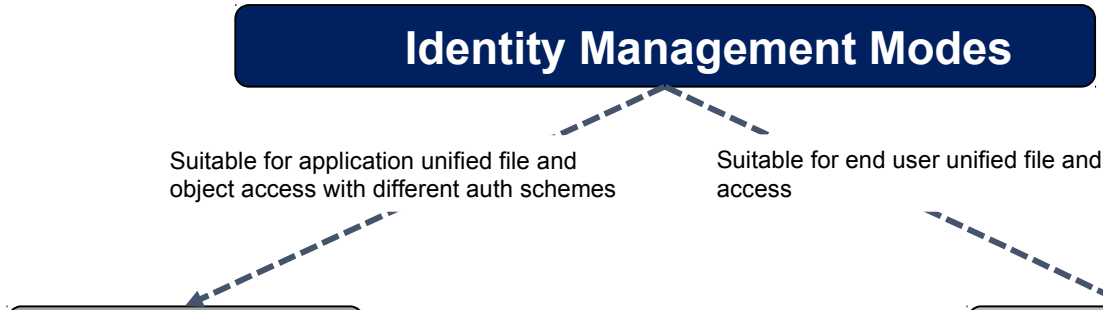


```
#mmobj config change
--ccrfile object-server-sof.conf
--section DEFAULT --property id_mgmt
--value unified_mode | local_mode
```

- Object created by **Object interface will be owned by internal "swift" user**
- Application processing the object data from file interface will need the required file ACL to access the data.
- **Object authentication setup is independent of File Authentication** setup

- Object created from Object interface should be **owned by the user doing the Object PUT** (i.e FILE will be owned by UID/GID of the user)
- **Owner of the object will own and have access to the data** from file interface.
- Users from Object and File are expected to be **common auth and coming from same directory service** (only AD+RFC 2307 or LDAP)

Unified File and Object Access – Unified identity management

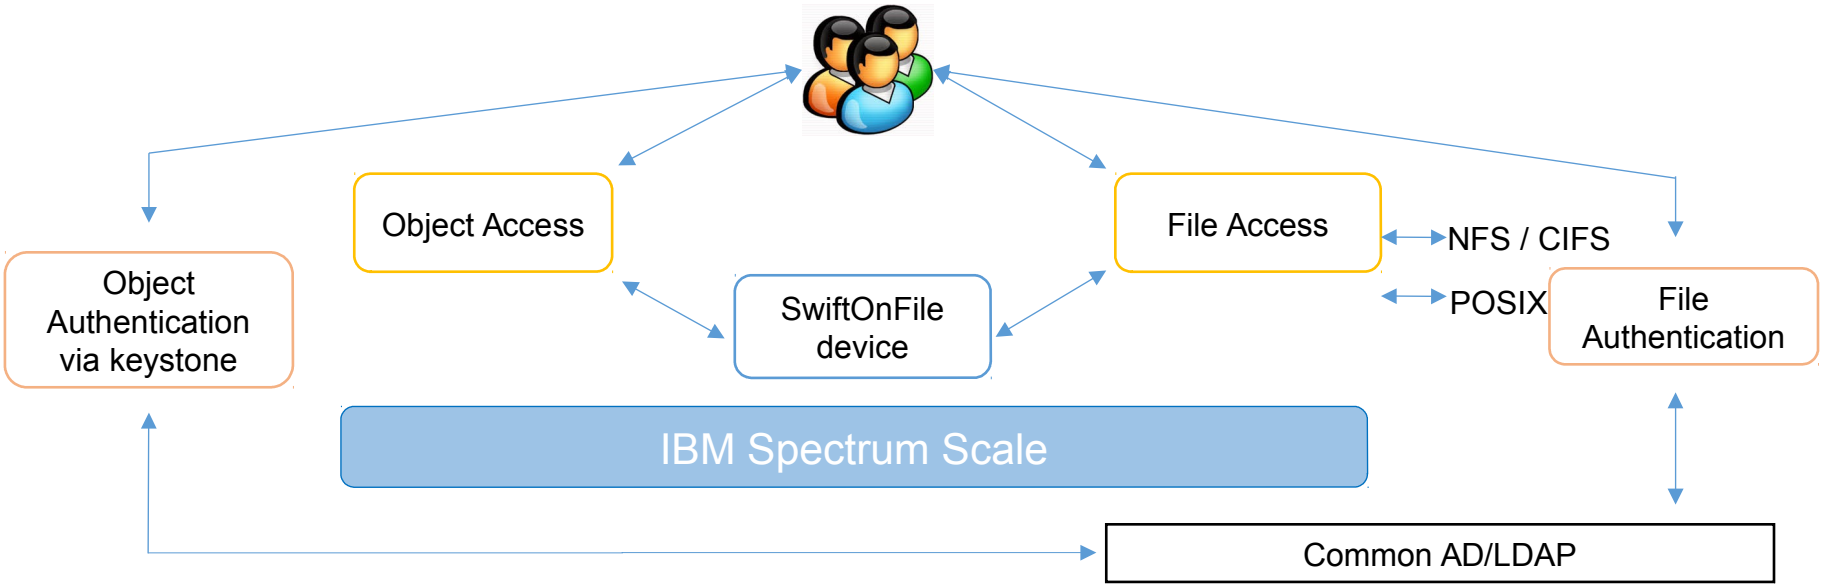

Improved S3 emulation

- What is S3 Emulation?
  - Object interface allows a client to communicate with either Swift API or Amazon S3 API (both are HTTP/REST based)
- Compatibility matrix: https://wiki.openstack.org/wiki/Swift/APIFeatureComparison
- How to Enable S3 Emulation (no change from 4.1.1):
  - On install using s3 flag:
    - `# spectrumscale config object -s3 on`
    - `# mmobj swift base … -- enable-s3`
  - After install:
    - Steps defined here: http://www.ibm.com/support/knowledgecenter/STXKQY_4.2.0/com.ibm.spectrum.scale.v4r2.adm.doc/bl1adm_ChangeconfigurationenableS3.htm
- New Features in 4.2:
  - S3 ACL Support:
    - define access control lists on bucket (container) and object level
    - ACLs stored in object/container file xattrs
  - Full support for Multi-part upload
  - Automated compliance tests

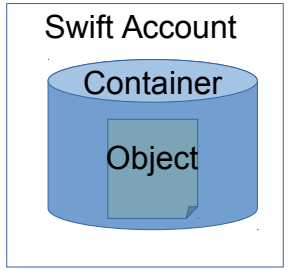

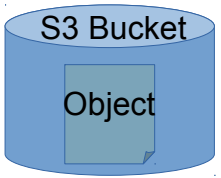

Multi-Region Active-Active Multi-Site Storage Cloud

Global Distribution

Ingest and Access from Any Data Center

Multi-Site Availability

Objects Replicated Across 2 or more Sites

Flexible

Async or Sync Replication

Multi-Region Architecture Details

- Provides Disaster Recovery of data center failures in a Active-Active storage cloud
- Binds separate Spectrum Scale clusters into a practically limitless capacity storage cloud
- Objects are stored in one or more regions depending on
  - Required performance (1, 2 or 3 copies)
  - Required number of supported data center failures (up to 3 sites)
- Objects are accessible from ANY site
  - Non-local objects are retrieved from remote region
- Supports asynchronous or synchronous replication
- Always returns latest copy across all sites

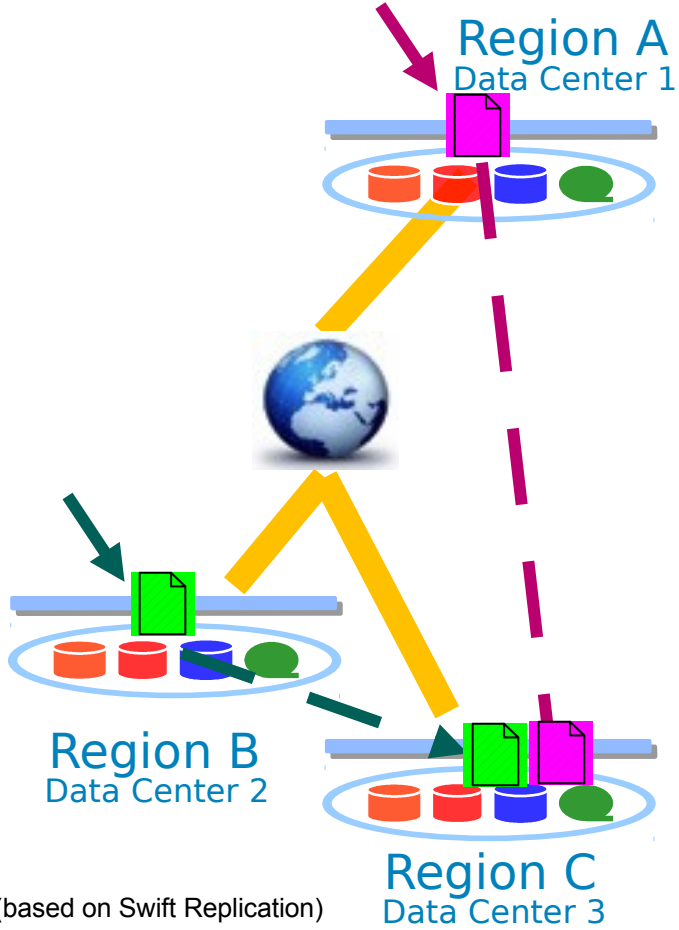

Note: currently only supported for Object access (based on Swift Replication)

Redpaper on Spectrum Scale Object + Spectrum Archive

2 main methods to leverage the Spectrum Archive tape tier in the Spectrum Scale object store:

1) Specific S3 buckets and containers with immediate migration to tape
   - Advantage: fine-grain control of migration
   - Disadvantage: requires end-user or application awareness
2) Single namespace to contain warm and cold data
   - Advantage: no data copy or movement by user/application required
   - Disadvantage: no control of movement to tape tier → application timeouts

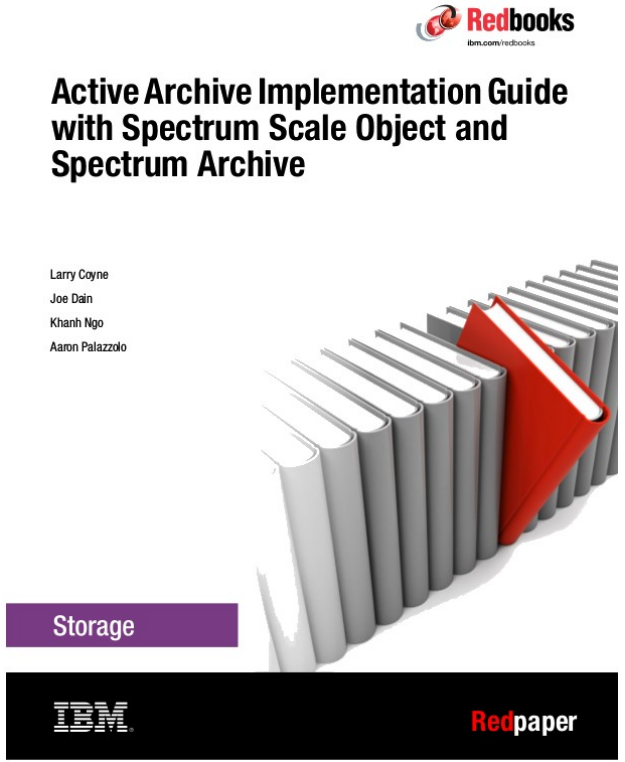

IceTier Update – OpenStack Swift and High-latency Media

- New Open Source SwiftHLM middleware for explicit migration/recall control though Swift API
- Works with custom HLM backends to enable controlled movement to HLM
  - Prototype backend available for Spectrum Scale Object and Spectrum Archive

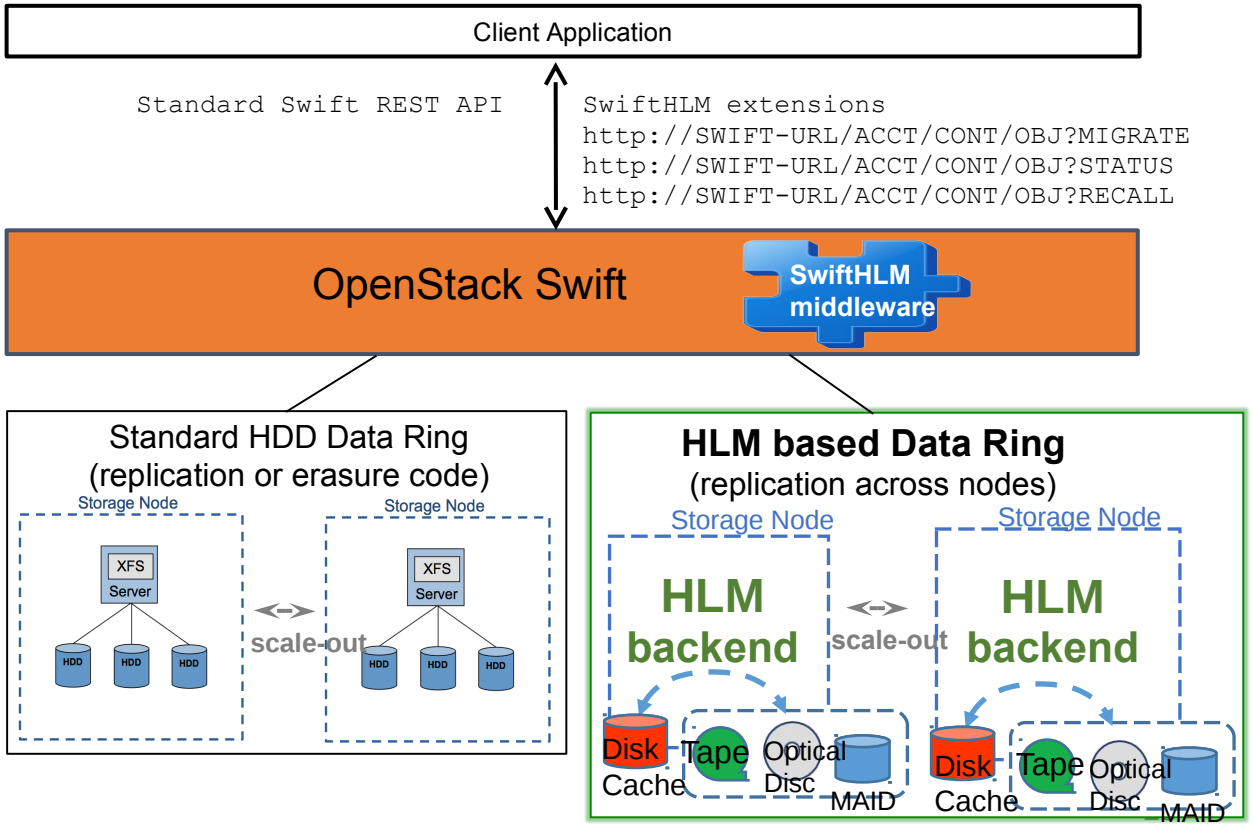

Cinder Spectrum Scale Driver News

- Volume Consistency Group support (Kilo)
- NFS support (Liberty)
  - Enables GPFS driver functionality without GPFS client code on OpenStack nodes
- GPFS encryption support (Liberty)
  - Creation of volumes in encrypted fileset through ExtraSpecs / Volume Type
- Currently working on
  - Volume replication
  - Consistency Group enhancements

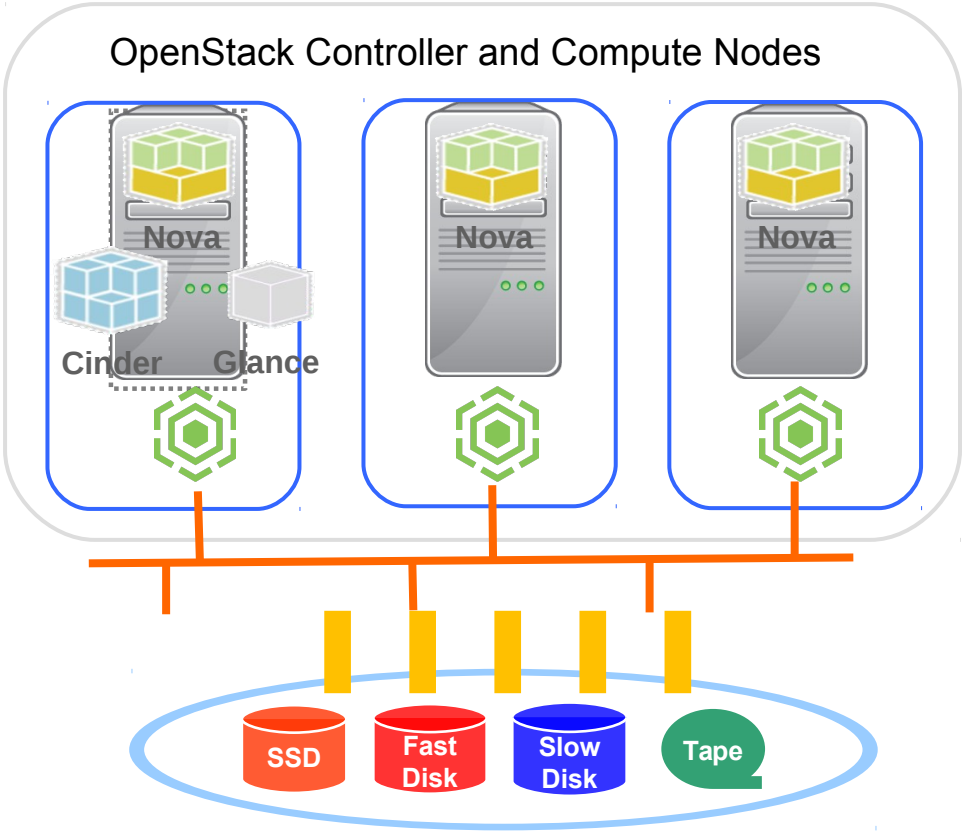

Manila Spectrum Scale Driver Kilo/Liberty News

- Initial official Manila driver release (Kilo)
- Adoption to 3rd party CI system (Liberty)
  - Required by community to stay upstream
- Added capability to extend share (Liberty)
  - Extend the size of the share by changing the quota on the backing GPFS fileset
- Currently working on
  - Making driver compatible with CES

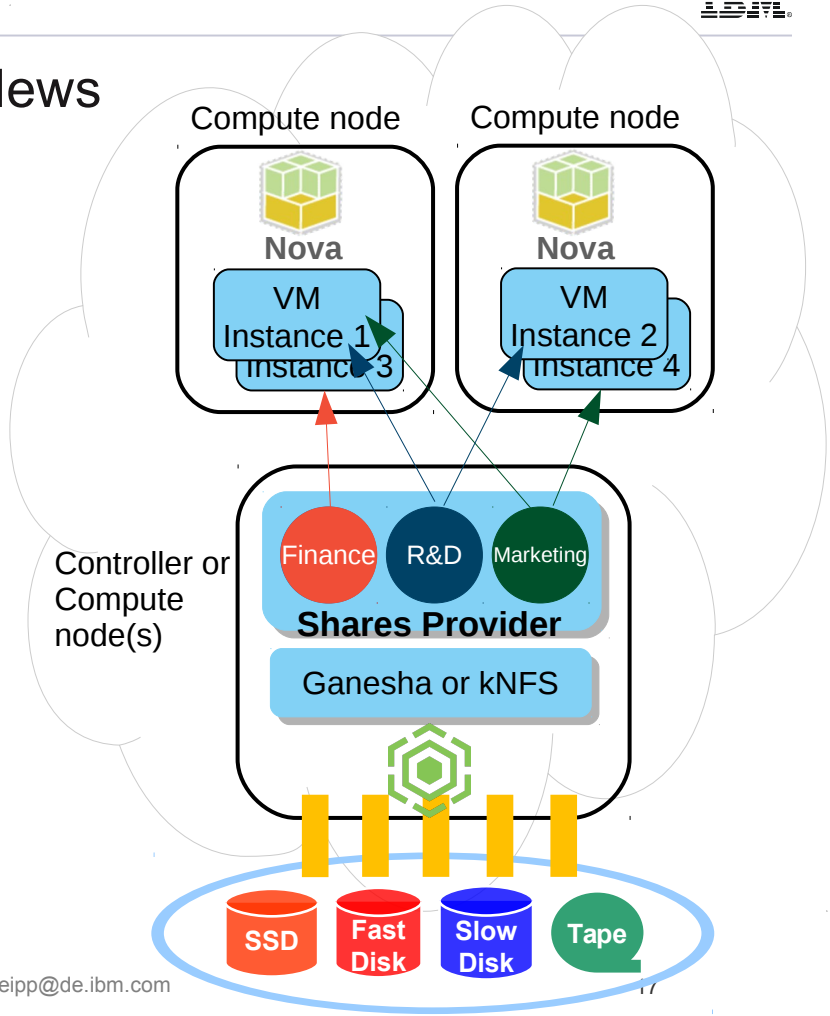

Integration with IBM Cloud Manager with OpenStack

- Requires IBM Cloud Manager with OpenStack (CMWO) 4.3 FP 3 or later
- Integrated installation experience, CMWO will also install & configure Spectrum Scale (with Object)
  - Common firewall configuration
  - Spectrum Scale optimizations for Nova, Glance and Cinder exposed through CMWO Chef parameters

```
    "ibm-spectrum-scale": {
      "object": {
        "nodes": [
          "cmwo_scale11-gpfs.mainz.de.ibm.com",
          "cmwo_scale12-gpfs.mainz.de.ibm.com",
...
        ],
        "export_ip_list": [
          "10.0.1.10",
...         ],
        "load_balance_ip": "10.0.1.2"
      },
      "package": {
        "repository": "http://repo-srv.mainz.de.ibm.com/repo/IBM/GPFS-4.2.0.1",
        "file": "Spectrum_Scale_Protocols_Standard-4.2.0.1-x86_64-Linux.install"
      },
      "filesystem": {
        "metadata_replicas": 3,
```

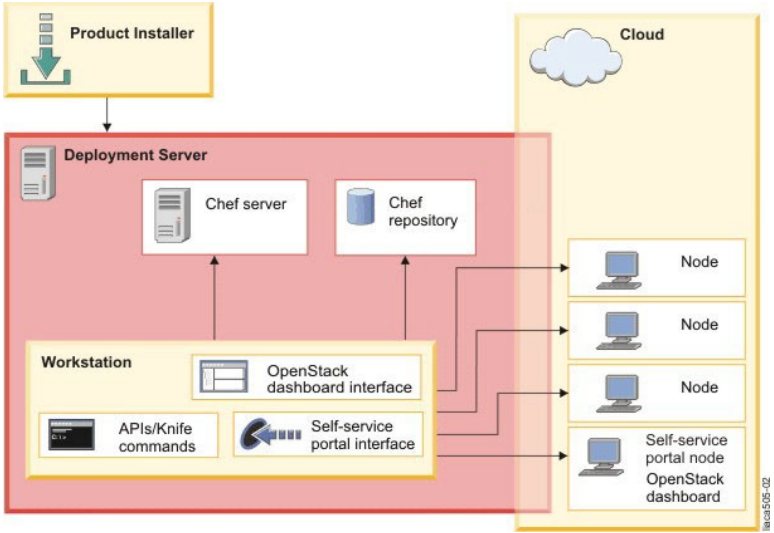

IBM Cloud Manager with OpenStack

The Vision: One solution for all your data needs

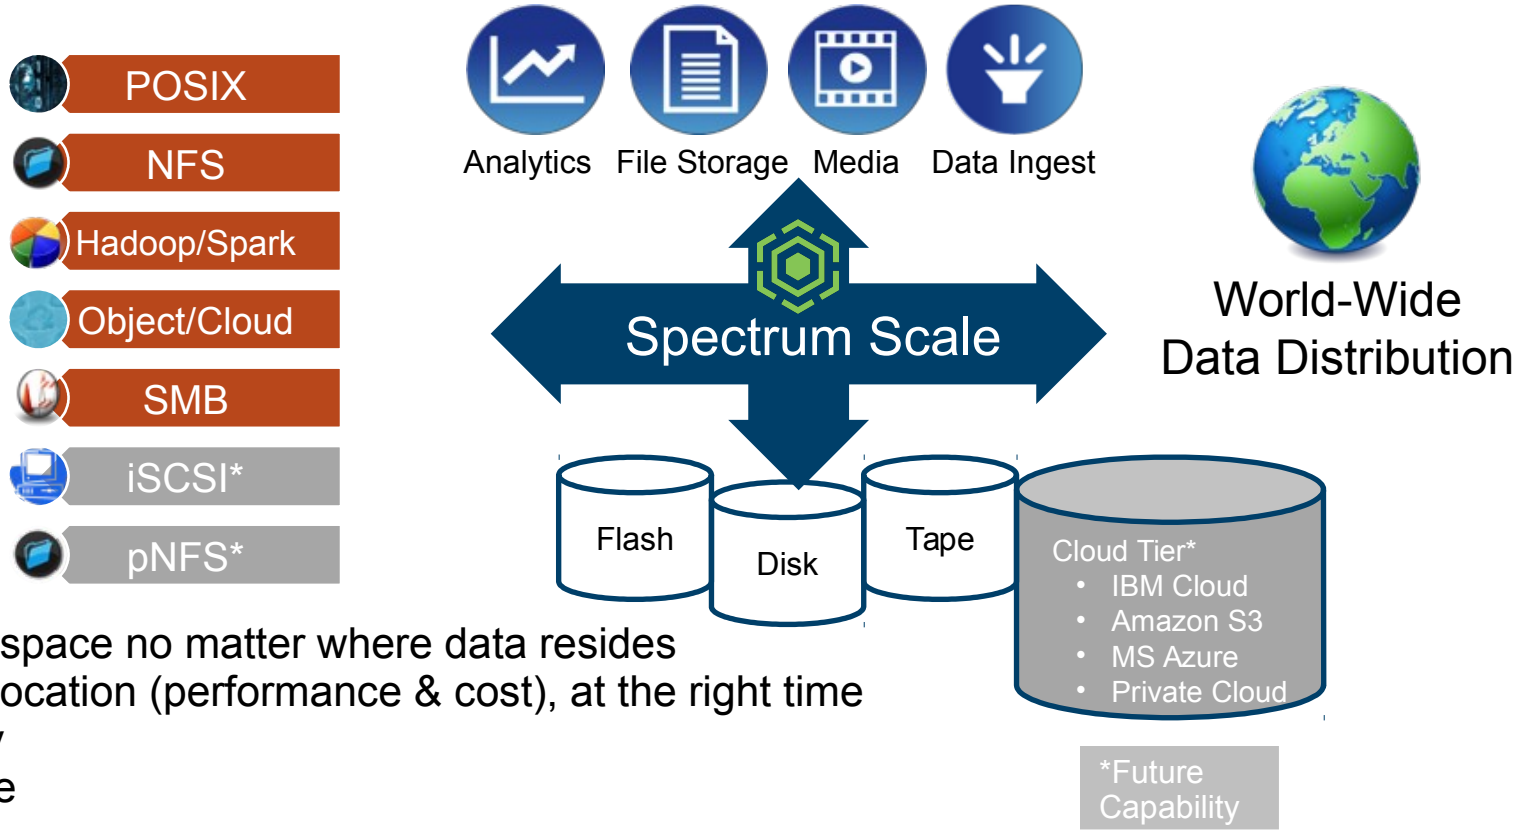

- Single name space no matter where data resides
- Data in best location (performance & cost), at the right time
- Multi-tenancy
- All in software

Reference Information

Openstack governance (“Big Tent” model with project maturity indicator tags): http://git.openstack.org/cgit/openstack/governance/tree/reference/projects.yaml

Paper with information regarding Icehouse news: https://ibm.biz/BdRG3U

Cinder Volume re-type details: https://blueprints.launchpad.net/cinder/+spec/volume-retype

Cinder Volume replication details: https://blueprints.launchpad.net/cinder/+spec/volume-mirroring

Cinder Backup/recovery API details: https://blueprints.launchpad.net/cinder/+spec/cinder-backup-recover-api

IBMNAS GPFS support details: https://review.openstack.org/#/c/102154/6/specs/juno/support-GPFS-nas-ibmnas-driver.rst

Consistency Groups: https://blueprints.launchpad.net/cinder/+spec/consistency-groups
https://review.openstack.org/#/c/96665/8/specs/juno/consistency-groups.rst

GPFS Cinder driver support for Consistency Groups: https://bugs.launchpad.net/cinder/+bug/1381877

GPFS Cinder driver volume attribute support: https://github.com/openstack/cinder/commit/8417b9ac0f6dc5e1f684591dffa3908c2a5427d0

Spectrum Control Cinder driver: http://www.ibm.com/developerworks/servicemanagement/sm/tpc/downloads.html

OpenStack Manila overview: https://wiki.openstack.org/wiki/Manila

Supporting NFS Ganesha in Manila shared file systems: https://www.youtube.com/watch?v=1zUSFzbhThM

GPFS Manila integration review: https://review.openstack.org/#/c/114311/4/

Blog covering Spectrum Scale Manila driver: https://www.ibm.com/developerworks/community/blogs/3aea761a-302b-4d05-be70-6fd652f7636a/entry/using_openstack_manila_shared_file_system_with_ibm_spectrum_scale?lang=en

Try out Unified File and Object with the Spectrum Scale Trial VM: http://www.ibm.com/developerworks/servicemanagement/tc/gpfs/evaluate.html

Redpaper on Spectrum Scale Object and Spectrum Archive: https://www.redbooks.ibm.com/redbooks.nsf/RedpieceAbstracts/redp5237.html?Open

IBM Cloud Manager with OpenStack trial downloads: https://www.ibm.com/developerworks/servicemanagement/cvm/sce/downloads.html

IceTier and SwiftHLM : http://www.research.ibm.com/labs/zurich/sto/tier_icetier.html
https://wiki.openstack.org/wiki/Swift/HighLatencyMedia
https://github.com/ibm-research/SwiftHLM

Trademark and disclaimers

Trademarks and disclaimers (cont.)

NOTES:

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.

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. Prices are suggested US list prices and are subject to change without notice. Starting price may not include a hard drive, operating system or other features. Contact your IBM representative or Business Partner for the most current pricing in your geography. Any proposed use of claims in this presentation outside of the United States must be reviewed by local IBM country counsel prior to such use.

Information is provided “AS IS” without warranty of any kind. The information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

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. Users of this document should verify the applicable data for their specific environment.

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.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM makes no representation or warranty regarding third-party products or services.