




IBM Spectrum Scale

Transparent Cloud Tiering Deep Dive



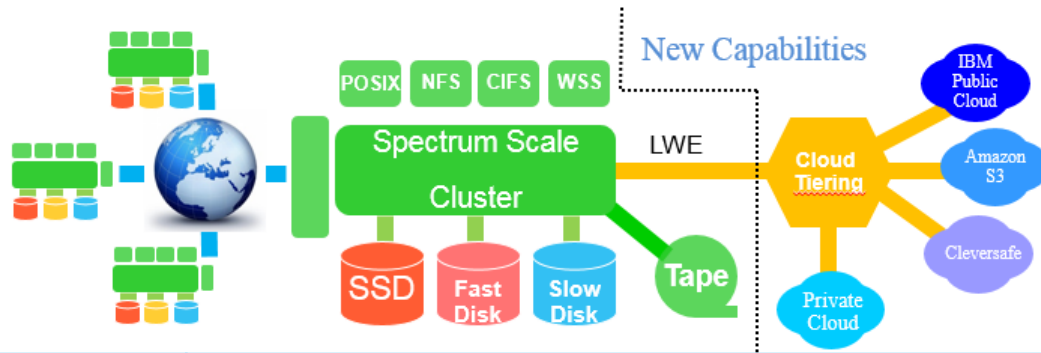
Disclaimer

- 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. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.
- Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the users 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 results similar to those stated here.

Agenda

1. Transparent cloud tiering overview
2. Multi-node support deep dive
3. Futures: Import / Export discussion
4. Questions

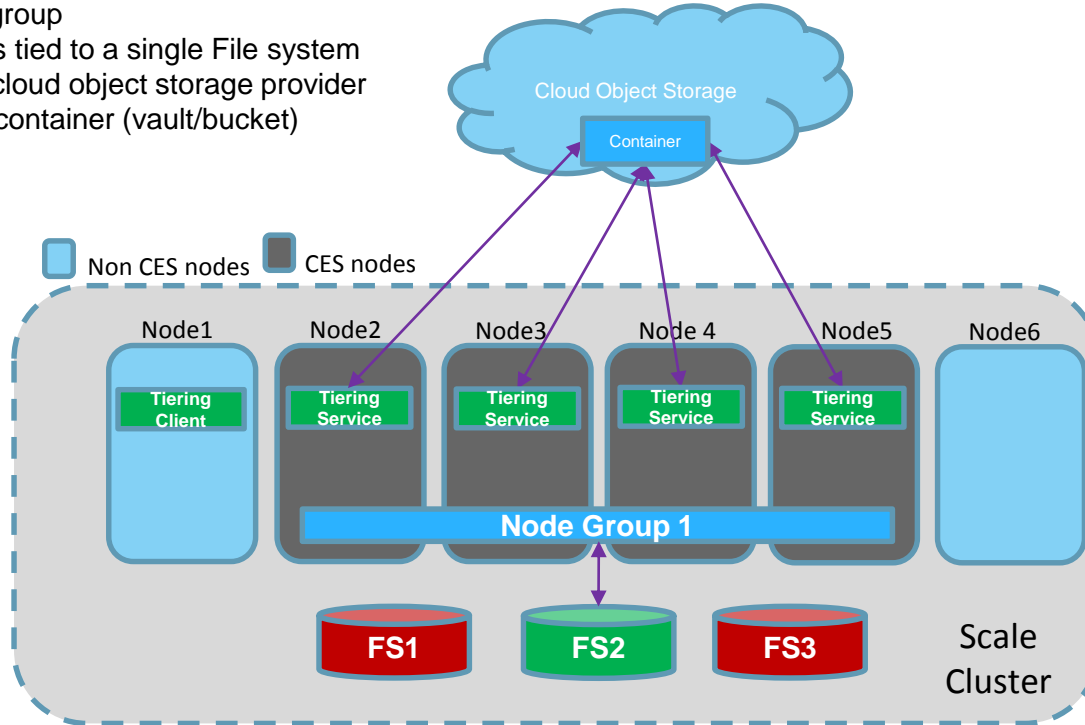
Transparent Cloud Tiering Overview



Key Client Use Cases	<ul style="list-style-type: none"> Enable a secure, reliable, transparent cloud storage tier in Spectrum Scale for cool or cold data with single namespace <ul style="list-style-type: none"> Leverage Spectrum Scale Information LifeCycle Management (ILM) policies Hybrid Cloud support: Integrate your Spectrum Scale storage with cloud storage of your choice Migrate and recall files between Spectrum Scale and cloud storage thereby saving storage costs and freeing up capacity on Spectrum Scale for your more active data
Platforms supported	Spectrum Scale 4.2 or later cluster, X86 Linux only. No mixed-node cluster support
Clouds supported	AWS S3, Swift, Swift S3 Emulation, and Cleversafe, HTTPS support for load balancer connections, validation tool checks compatibility of this service storage services
Scalability	Up to four nodes per cluster on RHEL7 only. One file system may be integrated to the cloud tier.
User Interface	CLI only – some limited monitoring available
Data Integrity	eTag integrity checking on all cloud data and meta-data
Security	Local key store, ISKLM Support
Delivery model	To be fully integrated as a part of Spectrum Scale Advanced Edition

Initial Release Multi-Node Support

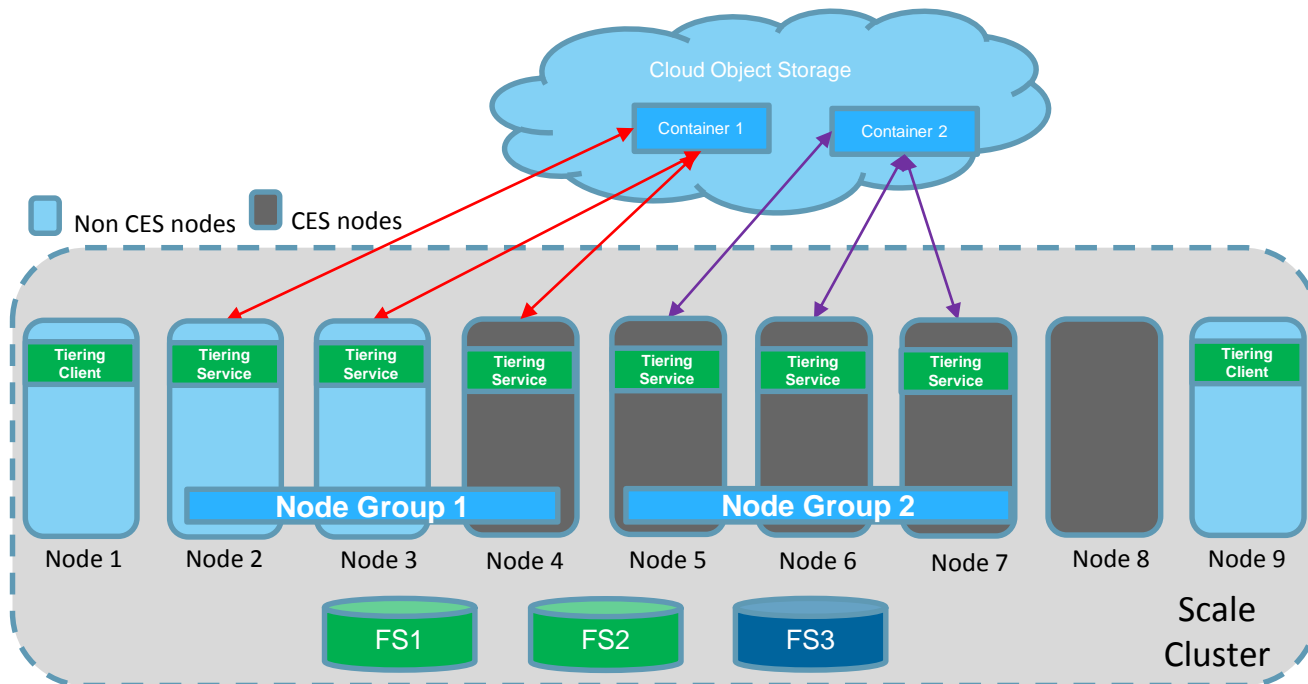
- Up to 4 cloud tiering service nodes
- Single node group
- Node group is tied to a single File system and a single cloud object storage provider with a single container (vault/bucket)



- Cloud tiering service nodes can only be deployed on CES nodes
- Parallel processing for performance - Default 32 threads (JVM) per cloud tiering service node
- Load balancing across the nodes for performance – ILM Policy distributes workload across all active node.
- Cloud tiering service nodegroup' Concept – node 2, node 3, node 4, node 5
- Single nodegroup for initial release

First Refresh Proposed Plan

- Multiple cloud object storage provider with a unique container set(vault/bucket) per file system (node group)
- 8-16 Cloud tiering service nodes per node group.
- Support multiple file systems in a node group (same file system to multiple node groups not supported)
- 4 node groups
- Cloud tiering service can be deployed on CES node or non-CES nodes
- Each node can exclusively belong to one node group only



▪ Load Balancing, HA, Fail Over, Security Design topics

Category	Initial Release	Near-term enhancements
Load Balancing	<p>Migration - ILM Policy to distribute files during migrate</p> <p>Recall – ILM Policy to distribute files during recall</p> <p>Transparent Recall – Load distribution via random node selection</p>	<p>Allow each node in a node group its own target address instead of all nodes using a shared target address</p>
HA	<ul style="list-style-type: none"> • No single point of failure • Cloud tiering service node failure – Rest of the nodes work in degraded mode • Idempotent on failures (recommended action is to rerun policy until successful) • Some failure mechanisms have no automated fail back when cloud tiering service is back up -- requires manual intervention. 	<p>Improved failback based on heart-beat detection and service restart when restart is successful</p>
Monitoring	<p>Monitor and retry for GPFS shutdown, node shutdown/reboot events and other failures.</p>	<p>Heart-beat based monitoring and integration into Scale eventing infrastructure</p>
Security	<p>External Key management with ISKLM</p> <p>Local key store</p> <p>AES 256-bit</p>	
Reconcile	<p>Manages removal from the cloud of deleted files.</p> <p>Older versions of files must be deleted manually</p>	<p>Support for clean-up of older versions of files.</p>

Performance Considerations

- Properly configured a single transparent cloud tiering node will mostly fill a 10 Gigabit network connection
 - Performance with WANs will be slower due to latency issues, packet drop rates, etc.
 - No attempts have been made on anything faster than a 10 Gigabit network
 - Implication is if you are doing multi-node and you want to fully utilize performance a load balancer is going to be needed
- Performance is slow for small files
 - Object storage generally has a lower IOPs rate than file or disk
 - Recommendation: Retain smaller files in other tiers and keep median file size over one megabyte

Import / Export Requirements Gathering Discussion

Use Cloud Storage as data distributor

Single Writer / Multi-Reader using cloud storage for distribution

Example: Data produced On-premise on Spectrum Scale

- Data is Exported to the cloud using Swift-On-File for straightforward file to object mapping
- Example only, data can flow other way into Scale via Import
- Controlled, serialized multi-writer is possible
- Flow is controlled manually by customer

Consistency

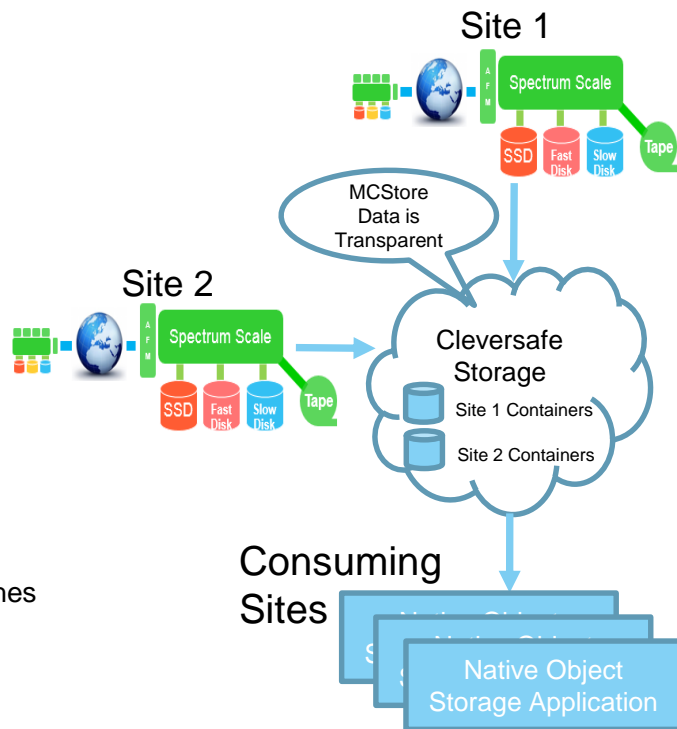
- Single-writer access on initial configuration

Full Transparency

- Meta-Data and Data are transferred to the cloud with no changes for direct consumption
- Some corner case exceptions on metadata size and path length naming mismatches

Communicating what data is available for consumption

- Object storage listing services can be challenging (Cleversafe claims significant degradation if listing services function is enabled)
- Readers poll (do we need asynch notification?) a well-known location for a journal of what is written to avoid the list services dependency



Back-up

Legal notices

Copyright © 2015 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 intellectually 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 inquires, 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.