IBM Spectrum Scale

AFM: Active File Management

Madhav Ponamgi
mzp@us.ibm.com

Agenda

AFM Organization
- Filesets
- Home / Cache cluster
- AFM Cache Cluster Semantics

AFM Deployment Architectures
- Multi-site data sharing
- Migration to new storage
- Disaster Recovery

AFM Caching Details
- Cache operations and semantics
- Cache failure scenarios

AFM Tuning & Setup

AFM Implementation Details

Store everywhere. Run anywhere.

Enable Global Collaboration

Challenge

- Multiple sites working on same data
  - Remote access is slower than local
  - Consistent metadata & data locking
  - Support for mission critical transactional replication
  - Manage unreliable, remote sites

Advanced File Management, Routing & Caching

- Global namespace with fast, consistent metadata
- Latency aware
- Multi-writer and multi-reader
- Automatic failover and seamless file-system recovery
- Asynchronous

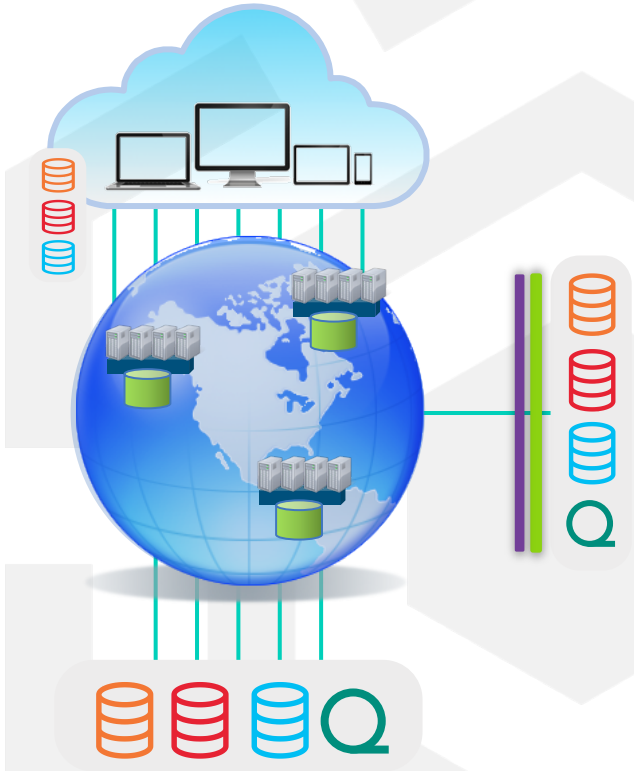

Evolution of the global namespace

Active file management (AFM)

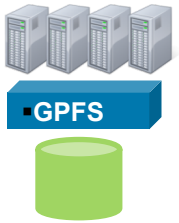

▪GPFS introduced concurrent file system access from multiple nodes.

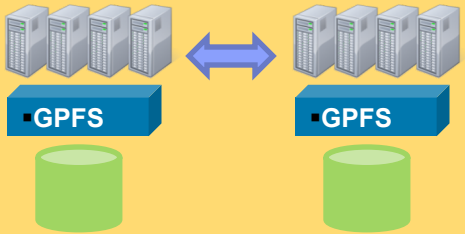

▪Multi-cluster expands the global namespace by connecting multiple sites

▪ AFM takes global namespace truly global by automatically managing asynchronous replication of data

▪1998 ▪2005 ▪2011

Spectrum Scale Advanced File Management (AFM)

Spans geographic distance and unreliable networks

- Caches local ‘copies’ of data distributed to one or more Spectrum Scale clusters
- Low latency ‘local’ read and write performance
- As data is written or modified at one location, all other locations see that same data
- Efficient data transfers over wide area network (WAN)

Speeds data access to collaborators and resources around the world

- Unifies heterogeneous remote storage

Asynchronous DR is a special case of AFM

- Bidirectional awareness for Fail-over & Fail-back with data integrity
- Recovery Point Objectives for volume & application consistency

Fileset

In most systems, file hierarchy is represented a series of directories that form a tree-like structure.

- Each directory contains other directories, files, etc.
- Every file system has a name associated with it

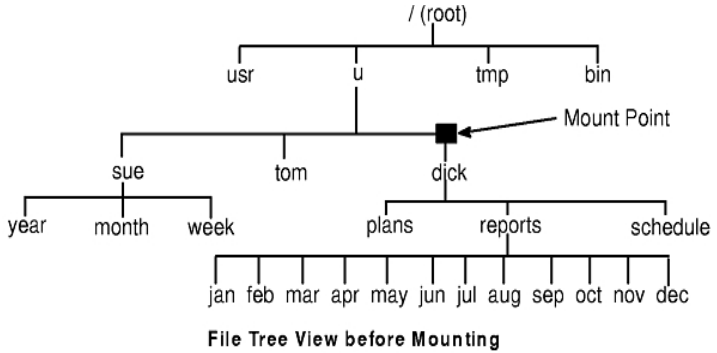

File Tree View before Mounting

A **fileset** is a subtree of a file system that in many respects like an independent file system.

- Filesets partition file system for finer granularity
- Filesets can have defined quotas for data and inodes
- The owning fileset is an attribute of each file for policies
- Fileset snapshots can be created
- Independent filesets have their inode space

Setting up AFM

On the home

- Create NFS export
- Set Home export configuration (`mmafmconfig`)

On the cache

- Define one or more Gateway nodes – one is designated a MDS
- Create cache GPFS fileset and associate a single home cluster with it
- Configure **maxfilestocache** to allow large number of pending requests at gateway nodes
- Mount the devices, create fileset and link filesets

AFM basics

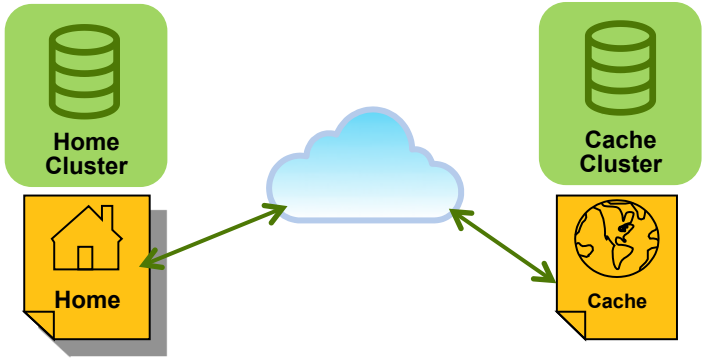

Cache basics
- Data updates are asynchronous
- Writes can continue when the WAN is unavailable

Two sides to a cache relationship
- Home
  - Where the information lives
  - Data is shared from to other cache sites
- Cache
  - Data is copied to the cache when requested
  - Data written to the cache is copied back to home as quickly as possible

Multiple cache relationships per file system
- Cache relationships are at a fileset level
- A cluster file system can contain multiple homes, caches and non-cached data (home, cache or both)

Multiple caching modes
- Read-Only
- Local Update
- Single Writer
- Independent Writer

AFM Structure – Standard Edition License

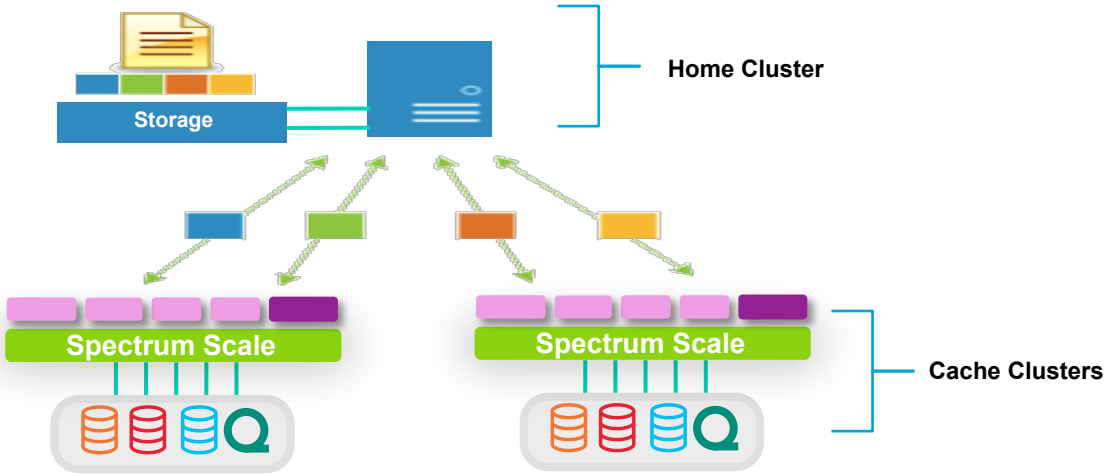

- Home is NFS export: local FS, GPFS, GPFS fileset. AFM uses proprietary protocol over NFS.
- Cache cluster comprised of Gateway (MDS) for fileset and Application servers
- Home can map to many cache clusters for a single fileset
- Each cache MDS owned fileset maps to one designated home

AFM Modes

Single Writer

- Only cache can write data. Home can’t change.
- Peer cache needs to be setup as read only

Read Only

- Cache can only read data, no data change allowed.

Local Update

- Data is cached from home and changes are allowed like SW mode but changes are not pushed to home.
- Once data is changed the relationship is broken i.e cache and home are no longer in sync for that file.

Independent Writer

- Data can change at home and any caches
- Different caches can change different files

AFM mode: Read-only caching

Read caching mode

- Data exists on the home fileset and one or more cache sites
- Data in the cache is read only – no modifications allowed in cache fileset

Data is moved to the cache on-demand.

- File Metadata caching: Listing the contents of a directory moves the file metadata information into the cache
- Data – Opening a file copies the data in the cache
- Getting data to the cache
  - On-demand when opened
  - Pre-fetch using a GPFS policy
  - Pre-fetch using a list of files

Caching behavior

- Many to one home
- Optional LRU cleaning of cache

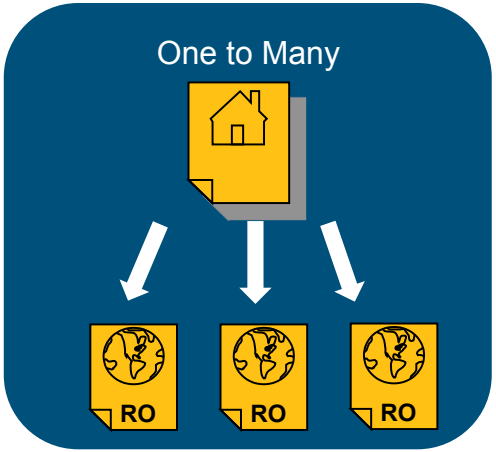

AFM mode: Local Update

Similar to Read caching mode

- Data exists on the home fileset and one or more cache sites
- Cached file set data can be modified but are considered local updates
- Local updates are never pushed to home
- Once a file is modified in the cache, file is no longer compared to home file
- Appending, truncating or writing to an uncached file causes the file to be fetched before making the change locally
- NOTE: any small change done in the LU fileset directory could cause the fileset to be marked local and lose context with the home.
- For example, running **chown** or **chmod** on root directory in LU fileset causes the entire fileset to lose context with home

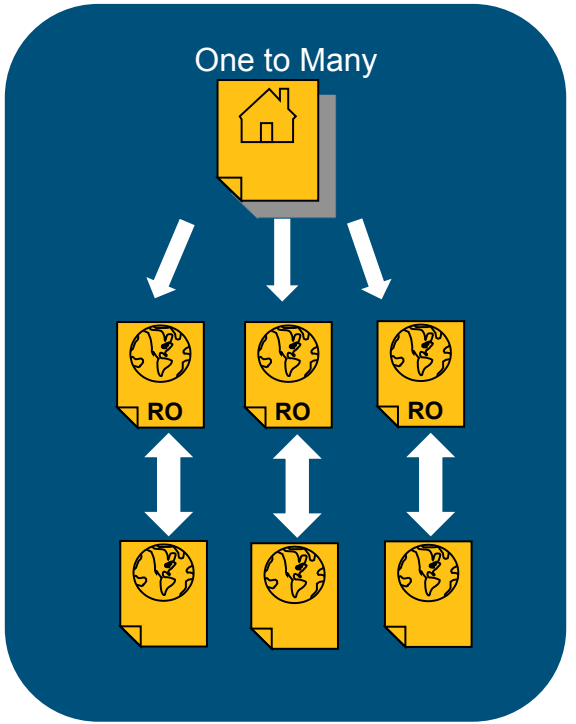

AFM mode: Single-writer

- While one cache does all writing, all others are in RW or LU mode (not enforced).
- In SW mode, home is not allowed to write.
- Asynchronous replication back to home
- Appending to or truncating does not fetch file, but queues it back home.
- The usage of peer snapshots via mmpsnap command to provide snapshot of home and cache

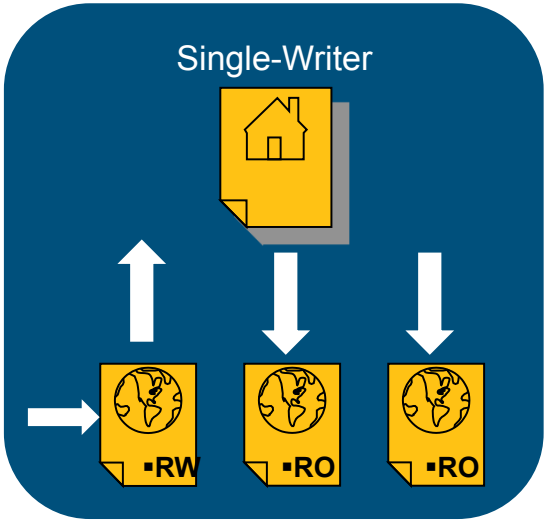

AFM mode: Independent Writer

Multiple cache nodes

All nodes can read/write data

Conflict resolution

- Default: The last writer wins (no locking enforcement)
- Partition data among writers to avoid conflicts

Appending or truncating a file in IW mode does not fetch the data into cache, but queues it to home

Factor in time-zones if home/caches are separated

Determine AFM mode:

mmlsfileset file_system fileset –L --afm

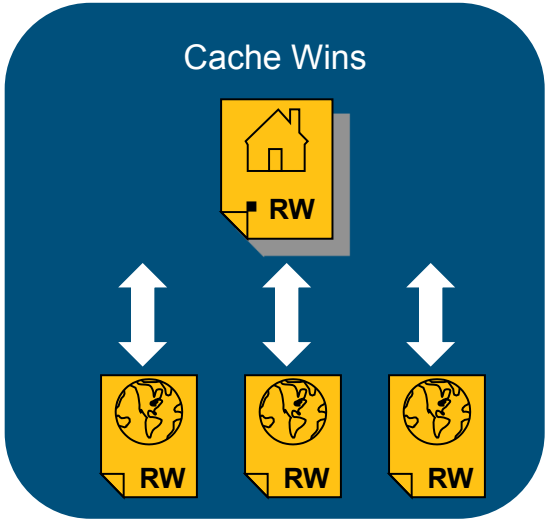

AFM Mode Conversion

- AFM Mode conversion can be done
  - Unlink for the file set
  - Issue a “**mmchfileset –p afmode”**
- A SW or IW with pending requests cannot be converted
- LU cannot be converted
- When going from IW to SW all other cache sites should be converted to RO or LU to avoid conflicts
- You cannot change mode or disalble AFM while fileset is linked
- Changing Modes
  - SW, IW & RO mode cache can be changed to any other mode
  - LU cache cannot be changed

AFM Architectures

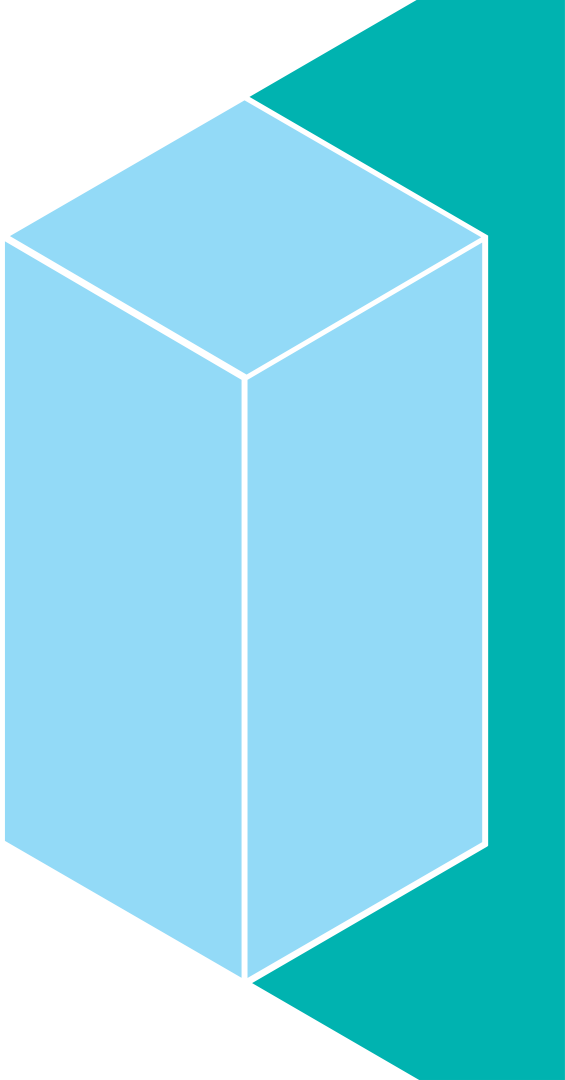

Global namespace

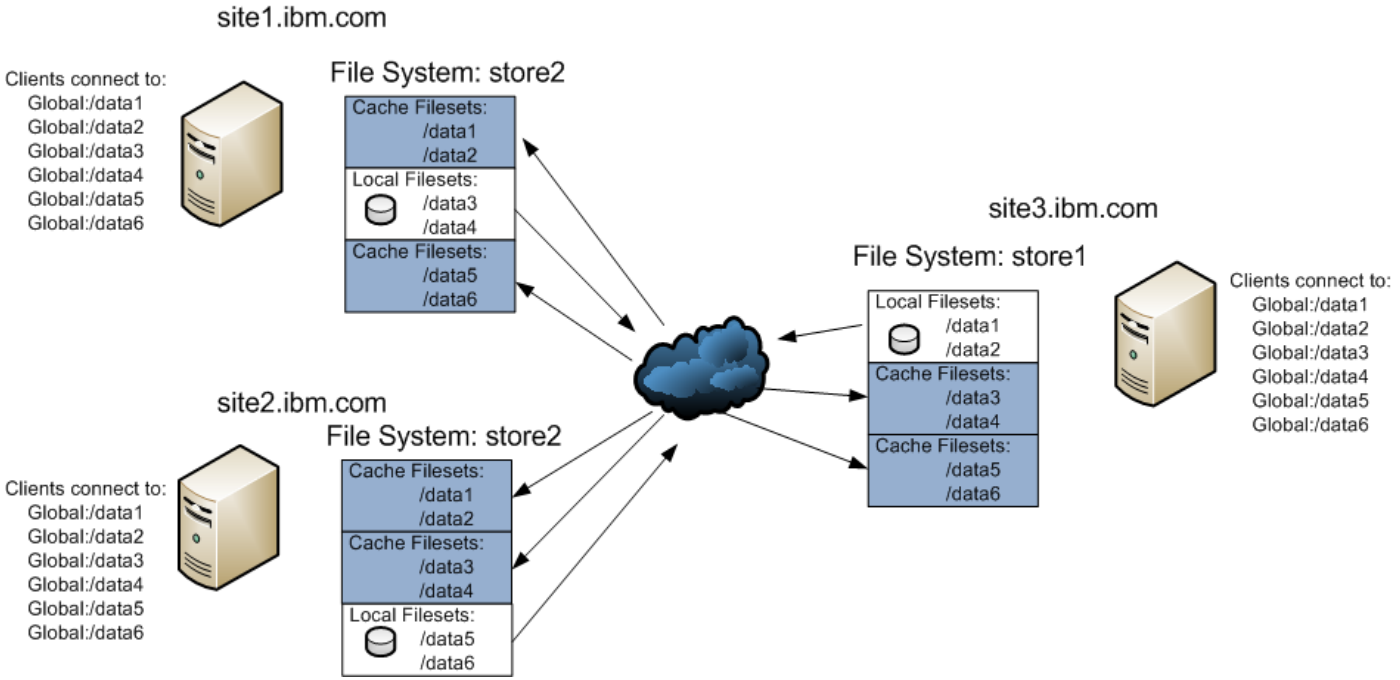

Multi-Site Pharmaceutical AFM Deployment

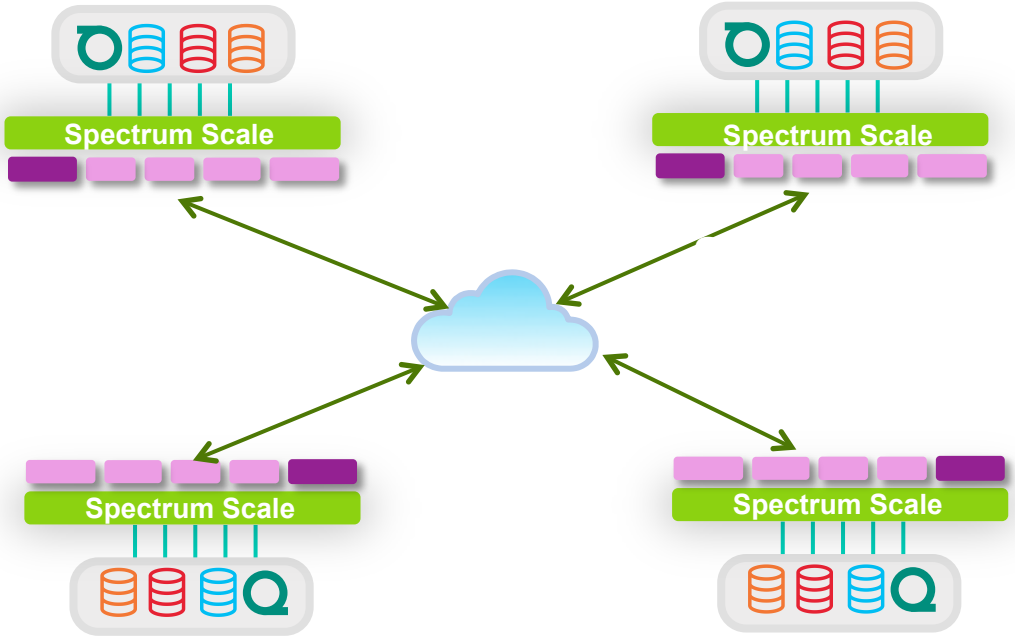

- 4 Sites with roving users
- IW Mode with no Pre-Fetch
- Each site consists of ESS storage with two CES nodes
- Users are in separate directories without common shared files
- Campus inter-connect is 1GigE

AFM: Asynchronous NFS Migration

Asynchronous Replication to new hardware
- Cache site does the writing
- Home site can have failback with IW

Cache fails
- New cache can be defined/ home continues

Home fails
- New Home can be defined or restarted

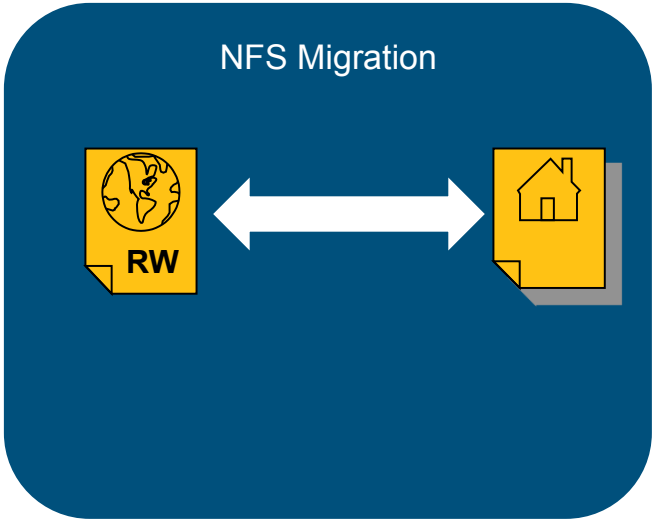

AFM-based NFS migration

NFS migration is a process of migrating data from any legacy storage appliance to a GPFS cluster via NFS protocol. This is useful in the event of upgrading hardware or buying a new system where the data from old hardware needs to be moved a new hardware.

- New hardware should be running at least GPFS 4.1 with data source NFS v3
- The export can be GPFS or non-GPFS (earlier than GPFS 3.4 is considered non-Gpfs)
- Migration comes in two type: Incremental or Progressive. Migration does not pull system details like quotas, snapshots, filesystem tuning parameters, policies, fileset definitions, encryption keys
- On GPFS data source, AFM moves Eas, ACLs, filesparseness.
- On non-GPFS data source POSIX permissions or ACLS are migrated (not NFS v4/CIFS ACLs)
- Incremental Migration (IW mode allows you to failback)
  - The metadata tree and critical data for migration for application are populated to AFM cache cluster
  - Applications continue to run on source, when cache is pre-populated downtime and transfer to cache
- Progressive Migration (LU mode does not allow you to failback)
  - Similar to incremental except applications are started the mmafmctl prefetch may still be running

Brown University AFM Migration

Brown University purchased new storage and used AFM NFS migration.

- Original DS3400 (circa 2009) with 1TB running 3.5.1
- New system composed of DCS 3700
- Network: FDR Infiniband using IP over IB no rdma
- New FS cache of old (progressive)
- Force it to populate the new cache

HOME TO THE LEGENDS OF TENNIS

Asynchronous Disaster Recovery (Advanced Server License Required)

Async DR is implemented one-to-one active-passive model.

- There is a Primary and Secondary pair
- The primary site is RW where applications are running with RW access to data
- The secondary is recommended to be RO and should not be used for direct writes
- The primary and secondary can be composed totally different hardware/networks
- After two sites are linked, primary is continuously available regardless of secondary status
  - Replicate all file user data, metadata, EAs (not atime / inode numbers), hard links, renames, clones
  - Filesystem attributes such as quotas, replications factors, dependent filesets, etc. are not replicated
  - File-based snapshots (psnaps) specify RPO point-in-time view of data.
  - Frequency of RPO are dependent on network between sites and size of data updates (15 min or >)

In the event of disaster of primary, the secondary can be promoted

- If required, the secondary's filesets can be restored using last RPO
- Applications can be moved or failed over to this acting primary
- Async DR offers capability to reconfigure old primary (or establish new) and failback

Hedge Fund AFM Disaster Recovery

- Primary Site
  - Geographically dispersed to different sites
  - Consists of mix of Native-Raid based disk storage (8 PB)
  - Also contains SSD tiers for meta-data and Tier 1 storage
  - Connected to Archive/Backup
  - Several thousand servers
  - Hundreds of users running dozens of applications
  - 10 GigE
- Secondary Site
  - Uses primarily DCS 3700 storage
  - Smaller amount of SSD tiers
  - No tape or archive backup
  - 10GigE

Components of Running AFM System

- Global Name Space / Migration / DR
- Storage needed at Home and Cache cluster
- Define Caching Mode
- Communication network (NFS / GPFS protocol)
- Define Gateways / Application nodes
- Define disconnect operation
- Create cache cluster policies
  - Eviction
  - Synchronization / validity intervals
  - Policies
  - Cache Operations

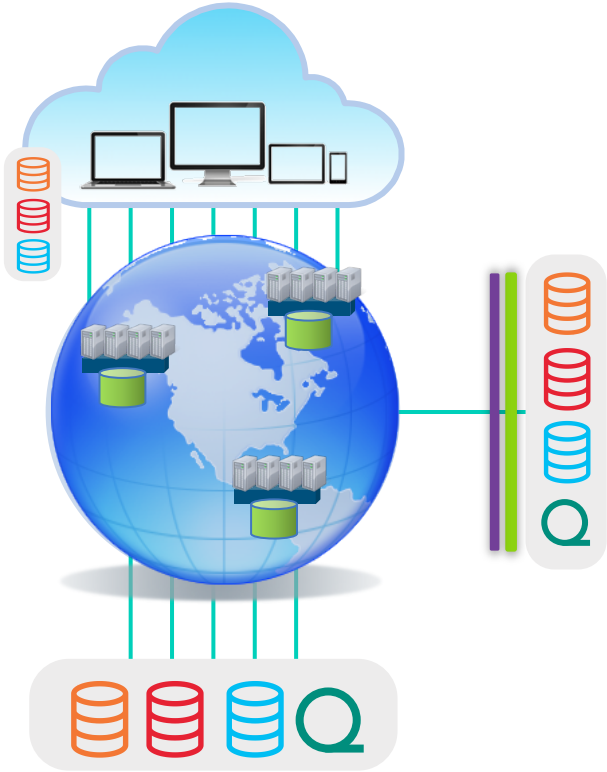

Cache Operations

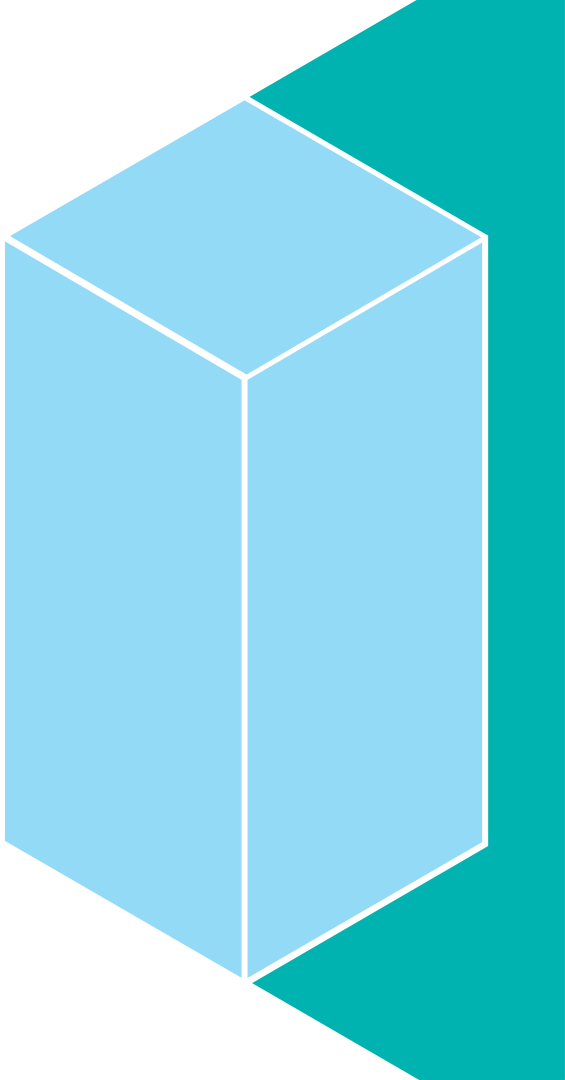

File System Caching and Synchronization

- **Asynchronous**
  - Application can proceed as soon as request is queued on gateway node
  - An MDS of fileset can delayed by the synchronization lag:
  - Files and directories are refreshed from home based on validity lag and file system updates are synchronized to home based on the synchronization lag
  - Within a single cache cluster, application nodes will experience POSIX semantics
  - AFM sends data from cache to home as root since root is not limited by quotas
  - Some asynchronous commands are: write, chmod, create, mkdir, remove, rename, link and symlink
- **Synchronous**
  - Synchronous operations block until completed at the home
  - READ: AFM does whole-file caching (afmPrefetchThreshold)
  - An uncached file cannot be evicted, resynced with home or failed to a new home
  - File and directory attributes are cached on-demand and can continue to work after disconnect

AFM Cache States

AFM cache can have different states, depending on fileset, queue and connection properties. The command **mmafmctl getstate** command displays the current cache state.

- Active
- Dirty – There are pending changes in cache not relayed to home. Normal activity continues.
- Disconnected – This state occurs in NFS exported cache (not GPFS protocol) when the MDS cannot connect to home NFS. To go back to Active, admin must correct errant NFS server at home.
- Dropped – File state moves to dropped state for multiple reasons
- Expired – When a RO cache is disconnected, a time period can be set to expire contents.
- FailbackCompleted/FailbackInProgress – These are only applicable for IW cache.
- FailoverInProgress – This state shows that the cache is in the middle of a failover and cache moves to Active when complete.
- FlushOnly – This state indicates operations are queued but have not been flushed. When a write is initiated should move to Active.
- NeedsResync – This is applicable for SW when accidental corruption at home is detected
- Umounted --- Cache using NFS/GPFS protocol go into this state if home is not available.

Failure and Recovery

Failures can occur in AFM if either an application node or gateway fails.

- The **mmafmconfig** command must run at home for recovery (GPFS based)
- Failures are not catastrophic and do not result in the loss of data or loss of AFM's ability to update the home cluster with local modifications.
- AFM stores on disk state information to replay updates to the cache at the home cluster
- Recovery process recovers any object that does not belong to a snapshot
- Application node failure is a GPFS node failure and does require special AFM processing
- Gateway node failure has special processing:
  - If no updates to AFM fileset, failure is harmless and application nodes experience no delays
  - If in recovery mode, all application node requests to AFM-filesets are temporarily blocked
  - If pending updates exists for AFM-fileset (write), cache cluster put in recovery mode
  - In recovery mode, cache requests not sent home are queued as determined by policy scan
  - Recovery is triggered for SW / IW and can run parallel across filesets (only 1 per fileset)

Failover of cache filesets

- Filesets continue to function independent of the home in the event of home failures. The AFM filesets serve the cache applications with cached data.
- The admin runs the "**mmafmctl** Device **failover –j** Fileset" command to point to new home
- An interrupted failover (gateway failure/quorum loss) restarts on next access to home
- Caches in SW and IW can be failed to new home (modes RO and LU cannot)
  - Failover fills the new home with contents from cache (data and metadata)
  - Extended attributes get synced when new home runs the **mmafmconfig** command before failover
  - Failover does not work if the new target is an NFS mapping (no such restriction if target is GPFS)
  - When there are multiple IW caches, admin must choose a primary IW and fail this over to the new empty home. It is recommended that all other IW caches be deleted and recreated.

NOTE: Admin must disable automatic eviction when home fails before starting failover to make sure eviction does not free cached data on the cache.

Prefetching / Peer snapshots

- Prefetch files selectively from home to cache
- Run asynchronously in the background (schedule for off peak times)
- Parallel multi-gateway prefetch (starting in 4.1) – one job per fileset
- Metadata-only without fetching files (starting in 4.1)
- User exit when completed
- Choose to files to prefetch based on policy
- Display status of prefetch
- Explicit prefetch **mmafmctl** Devicee **prefetch –j** Fileset
- Cache fileset peer snapshots are pushed home, and corresponding home snapshot is created
  - This results in a pair of peer snapshots
  - Use the **mmpsnp** command to create/delete snapshots (not mmdelsnapshot) on SW only
  - If cache is disconnected from home when cache snapshot is created, the cache notes that the peer snapshot on home is not created and attempts to complete after re-connect

AFM is on disk managed data

Data is managed like a cache but stored on disk in a GPFS file system.

Duration of data in a cache is dependent on configuration

- No cache cleaning (`afmAllowEviction`)
- Set duration of data in cache as good (`afmExpirationTimeout`).

Eviction can be triggered manually

```
mmafmctl Device evict -j FilesetName
```

- Use when
  - Cache smaller than home
  - Data fills up in cache faster than it can be pushed to home.
  - Need to create space for caching other files or space for incoming writes.
  - Eviction is linked with fileset quotas

AFM Tuning & Setup

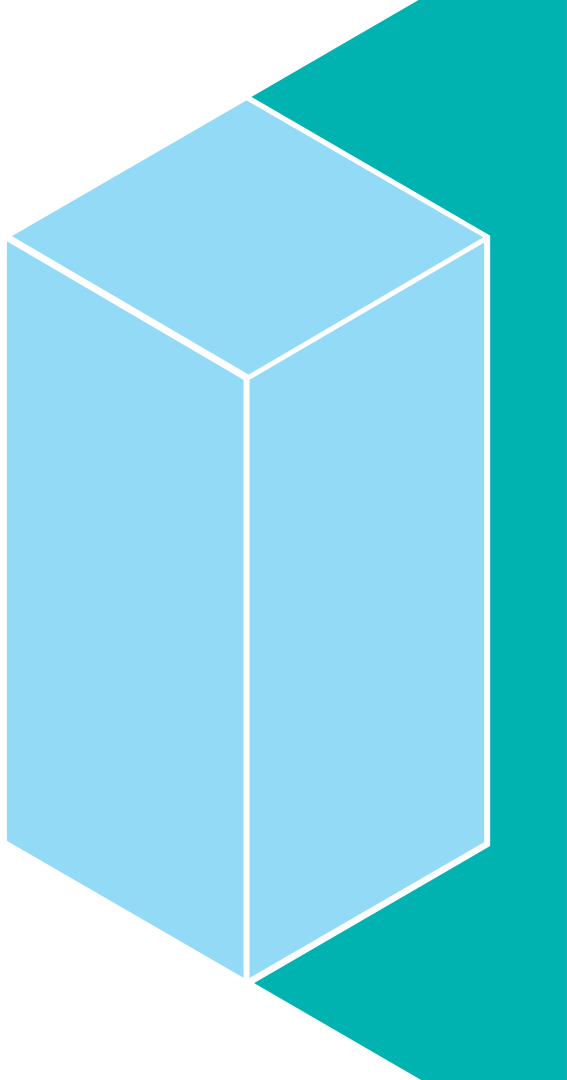

AFM and parallel I/O

The gateway server acting as the metadata server is the channel for communication with the home cluster. GPFS 4.1 now allows a cache cluster to be set up to perform parallel I/O and leverage all gateway servers in a cluster.

Multiple NFS servers are required at the home cluster and each gateway node in the cache cluster gets mapped to a specific NFS server at home—allowing for I/O load distribution

- One or more gateway nodes can be mapped to an NFS server, but each gateway server can only map to one NFS server
- Mapping is configured via the **mmafmconfig** command

If native GPFS protocol support is in place, gateway nodes can be mapped to any other node in the same cache cluster

- In the absence of a mapping definition, all gateway nodes will be used for the I/O

AFM Tuning

Tuning for the NFS client and NFS server

- Set the TCP values appropriate for the delay (buffer size = bandwidth * RTT)
  - For example, ping time is 50ms, and the end network consists of 100bt Ethernet and OC3 (155 Mps)
  - Buffer size = 0.05 sec * 10MB/sec = 500KB
  - If you are connected via T1 (1Mps) or less, default values are fine (faster need tuning)
  - The **sysctl** values do not activate until TCP connection is re-created via remounting (GPFS restart is required on the cache cluster)
  - On RedHat 6.1 and later, the server and client auto-tune
  - You can experiment with various net.ipv4/core settings described p. 231 of Adv manual
  - NFS server key settings to tune are **max_block_size, threads, nfsPrefetchStrategy**

AFM native GPFS protocol support

GPFS 4.1 enables native GPFS protocol support in place of NFS when using AFM

Native GPFS protocol utilizes the remote file system mount over multi-cluster to function as the AFM target.

This requires a multi-cluster setup to exist between the home and cache before AFM can use the home cluster's file system mount on the remote cluster

AFM will work with any file system on the home cluster, but ACL's, extended attributes, and sparse files are only supported when the home file system is GPFS

- Note: This is true whether using NFS or GPFS

The mmafmconfig command is used to enable native GPFS protocol support

AFM parameters

Set using `mmchconfig`, `mmcrfileset`, `mmchfileset`

- `mmchconfig` parameters are global defaults
- Fileset level setting override defaults

AFM Tuning

- Options are dynamic

`mmchfileset` afm options (`-p afmAttribute=Value`)

- `afmAllowEviction`
- `afmAsyncDelay`
- `afmDirLookupRefreshInterval`
- `afmDirOpenRefreshInterval`
- `afmExpirationTimeout`
- `afmFileLookupRefreshInterval`
- `afmFileOpenRefreshInterval`
- `afmMode`
- `afmShowHomeSnapshot`

Creating a cache

Cache is defined at the fileset level

`mmcrfileset` command

Usage:

```
mmcrfileset Device FilesetName
  [--inode-space=new [--inode-
  limit=MaxNumInodes[:NumInodesToPreallocate]]
   | --inode-space=ExistingFileset]
  [-p Attr=Value[,Attr=Value...]...]
  [-t Comment]
```

Example:

```
mmcrfileset cache2 master_t1 -p
  afmTarget=nfsnode:/gpfs/m1/m_t1 -p
  afmMode=cw --inode-space=new
```

Controlling AFM

Usage:

```
mmafmctl Device {resync | cleanup | expire | unexpire} -j FilesetName
```

or

```
mmafmctl Device {getstate | flushPending | resumeRequeued}
              [-j FilesetName]
```

or

```
mmafmctl Device failover -j FilesetName
               --new-target NewAfmTarget [-s LocalWorkDirectory]
```

or

```
mmafmctl Device prefetch -j FilesetName
              [[--inode-file PolicyListFile] | [--list-file ListFile]]
               [-s LocalWorkDirectory]
```

or

```
mmafmctl Device evict -j FilesetName
              [--safe-limit SafeLimit] [--order {LRU | SIZE}]
              [--log-file LogFile] [--filter Attribute=Value ...]
```

Setting up AFM

On the home

- Create NFS export
- Set Home export configuration (`mmafmconfig`)

On the cache

- Define one or more Gateway nodes – one is designated a MDS
- Create cache GPFS fileset and associate a single home cluster with it
- Configure **maxfilestocache** to allow large number of pending requests at gateway nodes
- Mount the devices, create fileset and link filesets

Thank you.

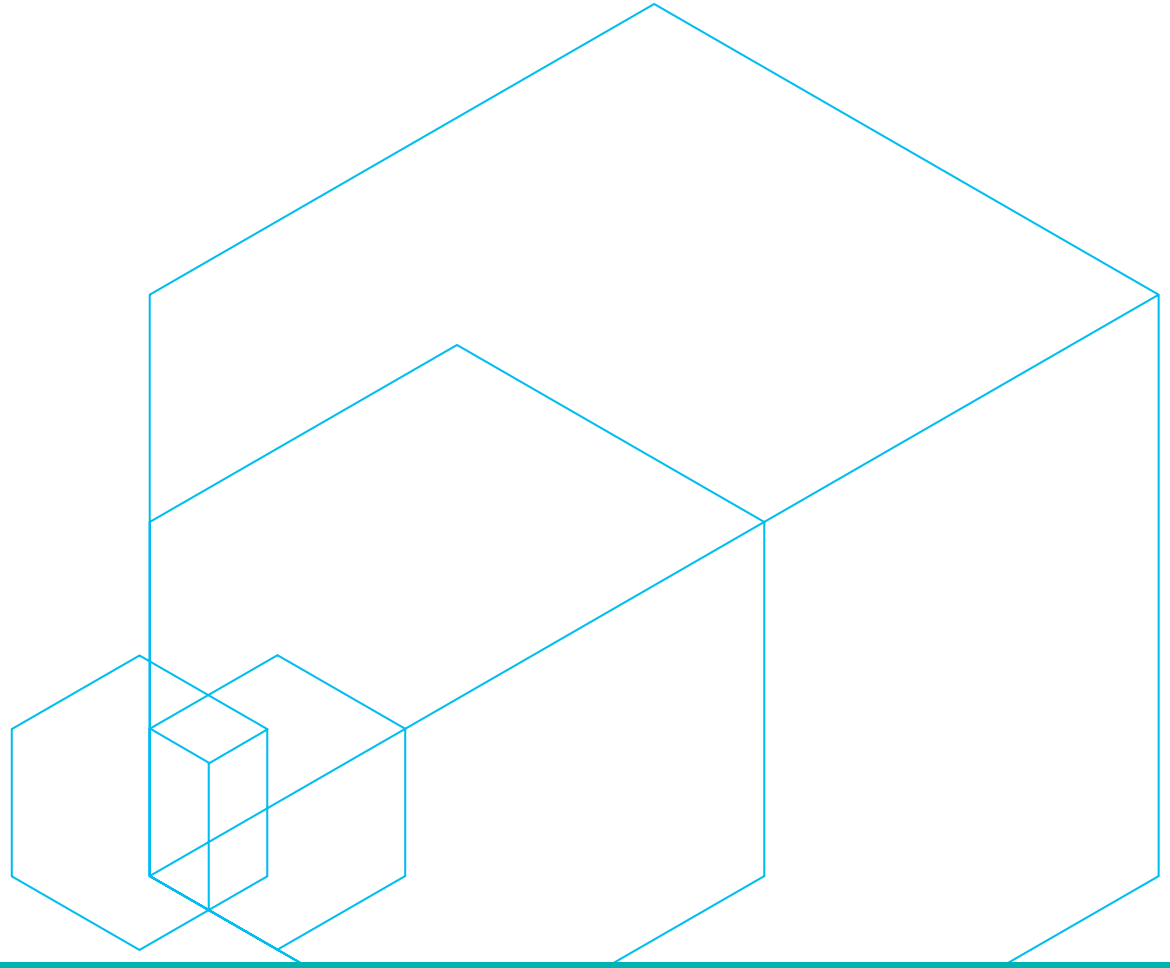

ibm.com/systems

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.

DCP03213USEN-06