

# **DAOS Technical Overview for NELOTUS**

**9/15/2009**

**Patrick Mancuso**  
**Senior Software Engineer**  
**Domino Server Development**

**patrick\_mancuso@us.ibm.com**

**Collin Murray**  
**Program Director**  
**Domino Product Management**

**[murrayc@us.ibm.com](mailto:murrayc@us.ibm.com)**

# Agenda

- Storage Savings & Messaging Economics
- What is DAOS?
- How DAOS Works
- DAOS Deployment
- DAOS Backup and Restore

# Industry Leading Messaging Economics

- Lotus Domino addresses key components of TCO
- Storage can be a significant piece of IT spend
- Recent IBM/Lotus storage savings initiatives
  - Design Compression
  - Document Compression
  - View Index savings
  - Attachment Compression
  - Attachment De-Duplication (DAOS)

# Storage Savings

- 40 – 60% storage savings
- Reduced administrative costs
- Reduced network traffic
  
- IBM saves \$\$ millions yearly – compression alone!

*"Upgrading to Lotus Domino 8.5 provided an immediate ROI and significant cost savings. Implementing DAOS resulted in storage savings of 77%, avoiding the need to purchase additional hardware for storage and the administrative costs associated with installation, etc."*

- Chief Hobbs, Forest Park Police Department

## Why DAOS?

- Continues storage saving capabilities
- Addresses customer concerns:
  - How can I reduce mail file and application size?
  - How can I reduce redundant storage?
  - How can I reduce I/O usage to improve performance?
  - How can I reduce storage associated costs?
    - e.g.: disks, backup, administration

# What is DAOS?

- DAOS = Domino Attachment and Object Service
- Attachment data stored externally in NLO (Notes Large Object) files
- Reference to NLO stored in NSF
- Multiple documents in multiple NSF files on same server can reference same NLO
- DAOS shares NO code with SCOS
- Transparent - works with all client versions
- Available on 8.5+ server only

# What is DAOS?

- -->NOT LIMITED TO MAIL<--
- Reduces size of NSF files in data directory
- Segregates static and dynamic data
- Reduces backup footprint
- Lower traffic to DAOS repository allows cheaper storage alternatives
- Reduces I/O, especially on multiple-recipient deliveries
- Oh, and...reduces overall storage footprint

# How DAOS Works

- NLO Files
  - The filename represents a checksum of the contents
  - Same filename...same contents...regardless of metadata
  - Domino compression and encryption done BEFORE checksum
  - DAOS encryption done AFTER checksum using server key
- NLO Storage
  - NLO files stored in numbered subdirectories
  - Files balanced based on max per directory (40K default)
  - NLO file never modified after creation - except to delete it

# How DAOS Works

- Write Operation
  - Attachment data written to temp file (xxxP.NLO)
  - File renamed to checksum name on write completion
  - Entry added to daoscat.nsf with refcount=1
  - Ticket written to NSF

# How DAOS Works

- Write Operation – duplicate
  - Same process for spooling and computing checksum
  - When checksum known, daoscat.nsf searched for existing entry
  - Refcount incremented, temp file deleted
  - Ticket written to NSF

# How DAOS Works

- Read Operation
  - Ticket opened at server
  - DAOS Base Path + Hint + name = NLO name
  - If file not at hint location, daoscat.nsf consulted
  - Contents of NLO file returned to client
  - Client sees no difference in operation

# How DAOS Works

- Document Deletion
  - Ticket read
  - daoscat.nsf entry retrieved, refcount decremented
  - If resulting refcount=0, start deferred deletion clock
- NSF Deletion
  - NSF scanned for list of tickets
  - NSF deleted
  - List of tickets processed in background

# How DAOS Works

- Pruning Unreferenced NLO Files
  - Prune done nightly at 2am
  - Scans daoscat.nsf for NLO entries
  - Deletes NLO files with 0 references older than deletion interval
  - Default interval is 30 days

# How DAOS Works

- Quotas and sizing
  - Logical size = NSF size + size of all attachments in DAOS
  - Physical size = NSF size
  - Quota enforcement based on logical size
  - DB properties box displays logical size
  - Domino Administrator files view shows sizes
  - User/script/API view of DB size shows no indication of DAOS

# How DAOS Works

- daoscat.nsf
  - Index of NLO files and reference counts
  - Index of NSF files using DAOS
  - No views
  - Rebuilt from scratch by resync
  - Do NOT back up or restore

# How DAOS Works

- daos.cfg
  - Contains counts of files in each subdir of DAOS repository
  - Created from scratch on startup if missing
  - Do NOT back up or restore

# How DAOS Works

- Replication with non-DAOS copy
  - Notes client, pre-8.5 server, non-DAOS 8.5x server replica
  - Attachment data sent with document contents
  - Client sees no difference

# How DAOS Works

- Synchronization
  - daoscat.nsf maintains master counts
  - Each NSF header contains count and timestamp
  - 'Needs Resync' if any discrepancy is detected
    - NSF deleted
    - Reference to NLO with 0 refcount
    - Reference by NSF not listed in daoscat.nsf
    - Difference in daoscat.nsf info and NSF header info
  - NR only affects prune operations
  - Run resync to get back to 'Synchronized' state
  - Resync usually not immediately necessary

# How DAOS Works

- New 8.51 Features
  - Smart Replication
    - If destination server already has NLO, reference sent instead
  - Smart Reply
    - Reply With Attachments – won't re-send attachment content
  - Smart features work for ANY source (including non-DAOS)
    - Destination must be 8.51 DAOS enabled
    - Source must be 8.51 (client or server) w/ODS 51

# DAOS Deployment

- Roll out 8.5 server
  - Upgrade to 8.5 Administrator Client
  - Upgrade Domino Directory design to 8.5 – hub first
  - Convert to ODS 51 `CREATE_R85_DATABASES=1`
  - Enable data + design compression
  - Enable LZ1 attachment compression
  - Enable transaction logging
  - Establish stable environment

# DAOS Deployment

- Location for DAOS directory
  - Logical Volume suggested to allow for expansion
  - Less I/O intensive than data directory
  - Opportunity for cheaper storage solution
- Minimum participation size
  - DAOS Estimator can help with overall & minimum participation size
  - 8.50 Default value too small, suggest 64000 instead (8.51 default)
- Deferred Deletion Interval
  - Coordinate with backup interval
- Encryption of NLO files – in addition to normal encryption
  - DAOS encryption on by default, disable with `DAOS_ENCRYPTION=0`

# DAOS Deployment

- Update backup/restore procedures
- Disable A/V Scans for NLO files
- Enable mail.box for DAOS
- Enable mailfiles for DAOS
  - `compact -c -daos ON somefile.nsf`
- Verify operation in Domino Administrator
  - Files view

# DAOS Deployment

- Some Actual Results – YOURS WILL VARY

System	NSF size before	NSF size after	NLO size after	Total size after	NSF size reduction	Total size reduction	Average I/O volume	CPU usage	Relative I/O rate
Mail server, AIX - 5.3.7.7 64-bit, 112 mail files	65.9 Gb	23.9 Gb	25.3 Gb	49.2 Gb	63%	25%	8% reduction	No change	DAOS directory disk has 1.3% of the busy rate of the data directory disk
Application server, Windows Server 2003 32-bit	17.3 Gb	9.8 Gb	5.3 Gb	15.1 Gb	43%	12%	N/A	N/A	N/A
Mail server, AIX - 5.3.7.7 64-bit, 389 mail files	70.9 Gb	40.3 Gb	15.4 Gb	55.7 Gb	43%	21.5%	12.9% reduction	10.5% reduction in average CPU busy	N/A
Mail archive server, OS400 - V5R4M0, 89 mail files	87.3 Gb	16.9 Gb	34.6 Gb	51.5 Gb	80%	41%	N/A	N/A	N/A

# DAOS Deployment

- Best DAOS Candidates:
  - Lots of Attachments
  - Lots of Duplication
  - Located on same server

# DAOS Backup

- 8.5x contains no changes to B/R API
  - Existing utilities will work
  - Only procedural changes need to be made
    - Back up NSF files first
    - Back up NLO files second
    - Do NOT back up daoscat.nsf, daos.cfg
  - NLO files treated as flat files
  - Interval should be less than Deferred Deletion Interval
  - Encrypted NLO files can NOT be restored to other servers

# DAOS Backup

- NLO files can be backed up incrementally
  - Not modified after initial creation
  - Significant savings in backup processing
  - Backup volume savings tracks closely to NSF footprint
- NSF files are smaller
  - Less NSF data to back up

# DAOS Restore

- Restore NSF file
- Check for missing NLO files as needed
  - Console: "tell daosmgr listnlo missing somefile.nsf"
  - Only needed if restore is outside Deferred Delete Interval
- Resync DAOS catalog
  - Required on any restore operation
- NLO files are encrypted with server key by default
  - Encrypted NLO files not portable between servers

# Q&A

- Additional resources
  - DAOS Wiki
    - <http://www-10.lotus.com/ldd/dominowiki.nsf/archive?openview&title=DAOS&type=cat&cat=null&tag=DAOS>
  - DAOS Backup and Restore
    - <http://www-10.lotus.com/ldd/dominowiki.nsf/dx/daos-backup-and-restore>
  - DAOS Estimator
    - <http://www.ibm.com/support/docview.wss?rs=463&uid=swg24021920>

© IBM Corporation 2009. All Rights Reserved.

The information contained in this publication is provided for informational purposes only. While efforts were made to verify the completeness and accuracy of the information contained in this publication, it is provided AS IS without warranty of any kind, express or implied. In addition, this information is based on IBM's current product plans and strategy, which are subject to change by IBM without notice. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this publication or any other materials. Nothing contained in this publication is intended to, nor shall have the effect of, creating any warranties or representations from IBM or its suppliers or licensors, or altering the terms and conditions of the applicable license agreement governing the use of IBM software.

References in this presentation to IBM products, programs, or services do not imply that they will be available in all countries in which IBM operates. Product release dates and/or capabilities referenced in this presentation may change at any time at IBM's sole discretion based on market opportunities or other factors, and are not intended to be a commitment to future product or feature availability in any way. Nothing contained in these materials is intended to, nor shall have the effect of, stating or implying that any activities undertaken by you will result in any specific sales, revenue growth or other results.

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 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 results similar to those stated here.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

IBM, the IBM logo, Lotus, Lotus Notes, Notes, Domino, Quickr, Sametime, WebSphere, UC<sup>2</sup>, PartnerWorld and Lotusphere are trademarks of International Business Machines Corporation in the United States, other countries, or both. Unyte is a trademark of WebDialogs, Inc., in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.