



Lotusphere2011

IBM Software

BP118 Proactive Server Management: Learn How to Maximize Your Server Uptime

Daniel Nashed | CTO | Nash!Com

Harvey Pope | Solution Architect | Hewlett Packard





About the Speakers

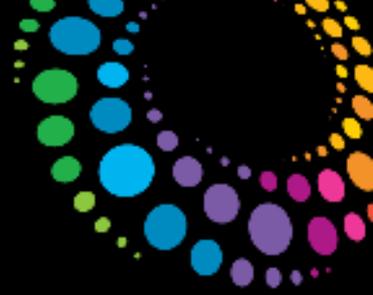
- Daniel Nashed
 - Nash!Com – German IBM/Lotus Advanced Business Partner/ISV
 - Member of The Penumbra group -- an international consortium of selected Business Partners pooling their talent and resources
 - focused on Cross-Platform C-API, Domino® Infrastructure, Administration, Integration, Troubleshooting and Lotus® Traveler
 - Platform Focus: Windows® 32/64, Linux®, AIX® and Solaris®
 - Author of the Domino on Linux/UNIX® Start Script
- Harvey Pope
 - Hewlett Packard Germany
 - Domino Consultant and Solution Architect
 - Problem Management and Specializing in Domino Troubleshooting





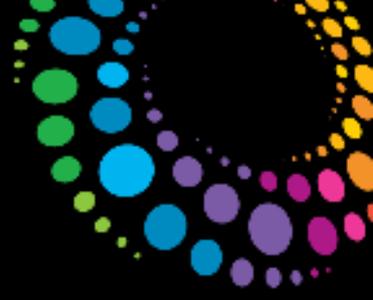
Agenda

- Introduction, General Considerations
- Right Architecture
- Maintaining your servers
- Proactive Monitoring
- Recovering from crashes / Crash Analysis
- Q & A



Aspects of Proactive Server Management

- **Right Architecture for your Domino Environment**
 - Right sizing and Platform (64bit OS instead of 32bit OS)
run Domino "Services" instead of Domino "Servers" → Clustering
 - Use Translog, Domino aware Backup Solution
- **Maintaining your Domino Servers**
 - Ways to make the server run more smoothly (with better performance etc)
 - Configuration Best Practices (DCT)
 - Daily / weekly operations (like compact)
 - Keep your servers up to date (fixpacks, etc)
- **Proactive Monitoring**
 - Ways to "notice" and fix problems before it reaches end users
 - DDM, monitoring free space, server availability, check statistics
- **Settings that help recover from crashes and pin point problems**
 - NSD, Fault Recovery, Debug Settings



Right Architecture



Domino Clustering

- Why Clustering?
 - Email and Domino Applications are mission critical for most customers
 - Domino needs to be available during business hours or even 24x7
 - It's not just about the TCO of your servers.
It's more about the cost of not having a service available when you need it
- High availability for Domino Servers
 - Better defined as High Availability of Domino Service instead of individual server
 - Total availability depends on
 - a.) Planned Down-Time and
 - b.) Unplanned Downtime for your Domino Service
 - You cannot reduce both to zero
- You should avoid all kind of Single Point of Failures (SPOF) for mission critical services



Difference between Domino and other High Availability Solutions

- Most solutions only use a stand-by servers
 - SAN is mirrored to a different data center and stand-by server is started with the same name and IP address when the other server or data center fails
- Domino works completely different
 - Separate servers with own name, own IP address and own storage
 - Data is almost real-time replicated between servers
 - Both servers are online and can be used
- Domino Clustering provides protection against
 - Disaster Recovery (when a whole data center goes down)
 - Server failure
 - Failure of a single database
 - Server crashes
 - Database Corruptions (usually corruptions do not replicate)
 - Shutdown servers for maintenance

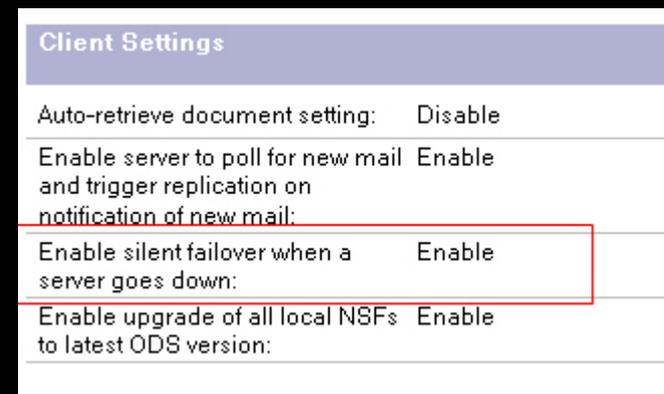
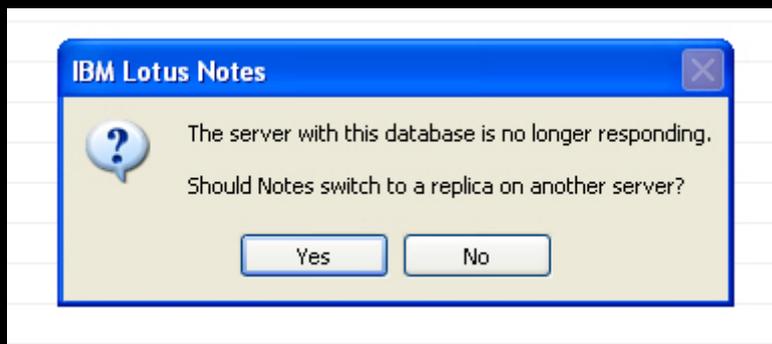


Client Failover Improvements

- Notes 8.x provides enhancements for failover
 - You are prompted to select another replica when the server goes down
 - Notes.ini Settings allow to automatically failover in some areas
 - HidePromptFailoverInc=1
 - Replaced by new functionality in 8.5.2
- New Notes 8.5.2 silent failover option to completely disable failover messages
 - Needs a desktop policy setting enabled
 - Sets **notes.ini FailoverSilent=1**
 - Also switches back to original server when it is back online! -- Finally



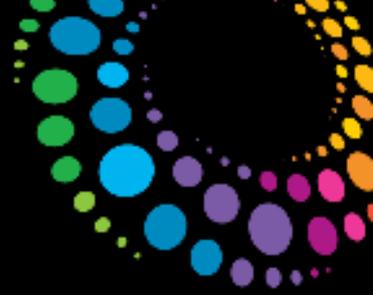
Tip





Transaction Logging

- Recommended for all Domino server types !!!
- Changes the way databases are locked for concurrent access
 - Lockmanager optimizes performance
 - Changes are written sequentially into translog
 - Asynchronous Log Manager writes data into databases afterwards
 - Process can continue to run meanwhile
- Without Transaction Logging, databases open at crash time are **inconsistent**
 - Needs fixup for **all open databases** which causes load on server and delays the time until Domino Server is completely online after a crash
 - Data Loss possible without Transaction Log in crash situations
- Recommend: Domino Backup API aware Backup Solution
 - Or shutdown your Domino Server for backup at night
 - All other backup operations are completely unsupported
 - Unless you need incremental backups (backup space saving & point in time recovery) circular translog is the right choice

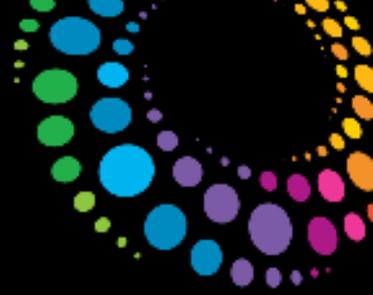


64bit OS Environments Recommended

- Windows 32 limits the total address space for applications to 2 GB!!!
 - You should also run Domino 32bit Servers on 64bit to have a dedicated address space with 4 GB (not only but particularly important for Windows)
- Add-On Applications need to be recompiled and linked with Visual Studio .Net 2003/2005 with link flag **/LARGEADDRESSAWARE**
 - Else any process (servertask) exceeding 2 GB limit (local + shared memory) it would crash!
- Check via dumpbin (available from Microsoft®)
 - Example: `dumpbin /headers nnsbdbcat.exe`
 - Result:
 - FILE HEADER VALUES
 - Application can handle large (>2GB) addresses
- Default Max Shared Memory is set to 2GB on Domino 32bit
 - notes.ini **ConstrainedSHMSizeMB=3072** gives you a maximum of 3GB



Heads Up!



Translog Settings

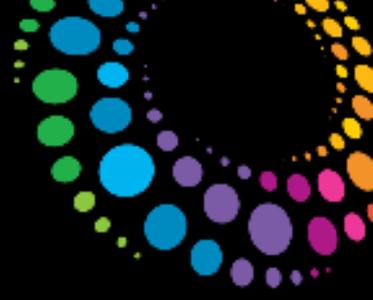
- Strongly recommended dedicated disk or SAN LUN
 - Should be the fastest available storage
- notes.ini parameter Create_R85_log=1
 - Ensures disc alignment for translog
 - You need to recreate Translog
 - On Windows you should use diskpart to create partitions with right alignment
 - Avoids overhead on disk operations (sector alignment)
 - <http://technet.microsoft.com/en-us/library/aa995867%28EXCHG.65%29.aspx>
 - Quite old article but describes it in detail
 - Same best practice applies also for all other file-systems!





Maintaining your Servers

“Prevention is better the cure”



Server Silent Install



Server Fix-Packs and Hot-Fixes (Windows)

- Manual Effort from customer experience
 - In the last 24 months we received 38 Hotfixes based on various ND8.5x releases
 - 2 of our customers have together over 200 Domino Servers distributed worldwide
 - $200 \times 38 = 7600$ installs
- In March 2010 IBM announced in Technote 1424074 a “Silent Installer for Fix Packs”
- e.g. 851FP1.exe /**NOUSER**
 - NOCHECKSUM - skip Application Integrity Verification
 - NOUSER - automatic upgrade mode (must be used with REVERT option for reverting hotfixes only)
 - REVERT Label - (with nouser) in revert mode will try to revert to the specified label
 - LISTFIX - lists installed hotfixes and fixpacks
 - FIXDETAILS - lists files belonging to hotfixes and their status
 - CONTENTS - lists files belonging to this hotfix



Current Limits of Silent Installer

- You have to make sure server and NSD service are not running
- No return code is set by the installer
 - Workaround **upgrade.log** in binary directory is created
 - search for text "**Notes has been successfully upgraded.**"
- You should delete/move the old upgrade.log before you install a new version
- In some error cases (like server is still running) a dialog box pops up
 - Needs very low-level work-arounds to check windows by name and close it
 - Works in most cases if you shutdown the server cleanly
 - Should be fixed in near future → SPR #JMIL8DNJQT
- Can deinstall previous fixpacks/hotfixes
 - But cleaner to use the installed fixpacks/hotfixes to de-install before upgrade



Fixpack/Hotfix Installer for Linux/UNIX

- Same functionality but leverages the “install script”
- Different type of script which works a lot faster than the Window based installer
- Example: `./install -script script.dat`
- You need to untar the installation tar file
- Current issue for normal installs
 - IBM Technote 1451499 published an issue with session window size
 - Happens because licence agreement contains a long URL
- If not updating a standard location “/opt/ibm/lotus” you have to set an environment variable
 - `export NUI_NOTESDIR=/opt/domino/lotus`



Server Major/Maintenance Releases

- Silent install also works for major and maintenance releases
- Different installer but can be automated using a response file
 - Sample response file included in installation kit
 - Unpack installation files (exe, tar file on Unix)
- Change settings in response file and invoke setup / install routine
 - Windows Example: `setup.exe -options "sample_response.txt" -silent`
 - UNIX Example: `./install -options "unix_response.dat" -silent`
- Current Limits
 - No Return-Code from installer
 - Workaround: DominoInstall.log is written in Domino Binary directory
 - but does not contain a good text line to search
 - check if it exists and has a certain size
 - Move/delete existing DominoInstall.log before install
 - You still need to make sure server and NSD service is not running
 - Recommended: Boot Windows Server after installation



Check for new notes.ini Parameters in each release and fixpack

- “search” the Fix Pack “Whats New” tab for any new notes.ini settings.
<http://www.lotus.com/ldd/r5fixlist.nsf/WhatsNew/>
- Important setting in 8.5.2 FP1
 - SPR# KBRN899NF6 - Fixed a problem where setting **Disable_Random_RW_File_ATTR=1** in the notes.ini does not disable setting RANDOM RW file flag
 - You need this notes.ini setting if you are running Windows 2008 64bit with more than 300 GB of Domino data
 - Cache Manager on Windows 2008 does not work well with the FILE_FLAG_RANDOM_ACCESS used by Domino since Release 7
 - Long story, check details in Daniel's blog

<http://blog.nashcom.de/nashcomblog.nsf/dx/new-performance-problem-with-domino-on-windows-2008-64bit.htm>



Heads Up!



Server MR, FP and HF Caveats

- Shortfall in IBM Domino “silent install”
 - no automatic software distribution to local and remote servers
 - no automatic Server termination or restart capability
 - no error reporting, or alarms by failure
 - no built-in Automatic roll-back (de-install)
 - no Central Repository update history overview

Food for Thought

- Consider creating an Application to distribute and install MR's/FP's/HF's.

U/D	Log	Friendly Name	Ref. Number	Target Release	MR	FP	HF	Dist Date	Dist Time	Created by	Modified
t09-03		Ref-09-12-2010-14-45-52	851FP1HF134	851	851FP1		851FP1HF134	09-12-2010	14:55:00	Template Development	09-12-2010 16:26:48
t09-22		Ref-09-12-2010-13-29-32	851FP1HF288		851FP1		851FP1HF288	09-12-2010	13:39:00	Template Development	09-12-2010 14:45:47
t09-04		Ref-09-12-2010-10-39-29	851		851			09-12-2010	10:48:00	Template Development	09-12-2010 11:46:58
t09-03		Ref-09-12-2010-10-38-30	851FP1HF288		851FP1		851FP1HF288	09-12-2010	10:44:00	Template Development	09-12-2010 13:29:27
t09-2		Ref-09-12-2010-09-17-41	851		851			09-12-2010	09:24:00	Template Development	09-12-2010 10:38:05
t09-01		Ref-09-12-2010-09-11-15	852HF300				852HF300	09-12-2010	09:17:00	Template Development	09-12-2010 10:38:09
t08-22		Ref-08-12-2010-22-36-16	852HF137				852HF137	08-12-2010	22:40:00	Template Development	09-12-2010 09:11:12
t08-01		Ref-08-12-2010-20-00-02	852		852			08-12-2010	20:04:00	Template Development	08-12-2010 22:38:03
t08-01		Ref-08-12-2010-11-58-15	851FP1HF288				851FP1HF288	08-12-2010	12:07:00	Template Development	08-12-2010 21:56:22
t07-01		Ref-07-12-2010-15-43-18	851FP1HF134		851	851FP1	851FP1HF134	07-12-2010	15:50:00	Template Development	08-12-2010 11:57:40
t02-04		Ref-02-12-2010-17-32-15	852HF300				852HF300	02-12-2010	17:36:00	Template Development	07-12-2010 16:42:10
t02-03		Ref-02-12-2010-16-29-18	852HF300				852HF300	02-12-2010	16:34:00	Template Development	02-12-2010 17:31:54
t02-02		Ref-02-12-2010-13-24-55	MR852		852			02-12-2010	13:32:00	Template Development	02-12-2010 15:02:10
T-02-01		Ref-02-12-2010-10-06-48	851FP1		851	851FP1		02-12-2010	10:16:00	Template Development	02-12-2010 13:24:25
t01-01		Ref-01-12-2010-14-52-12	852HF137		852		852HF137	01-12-2010	15:01:00	Template Development	02-12-2010 10:06:42
t30-03		Ref-30-11-2010-15-41-38	851		851			30-11-2010	15:47:00	Template Development	01-12-2010 14:52:07
t30-02		Ref-30-11-2010-11-33-04	852		852			30-11-2010	11:40:00	Template Development	30-11-2010 15:30:31
t30-01		Ref-30-11-2010-09-37-06	851FP1HF134		851	851FP1	851FP1HF134	30-11-2010	09:54:06	Template Development	30-11-2010 11:31:59
t29-11		Ref-29-11-2010-20-44-05	851FP3HF250		851FP3		851FP3HF250	29-11-2010	21:55:00	Template Development	30-11-2010 11:08:02
t29-10		Ref-29-11-2010-19-11-10	851FP2HF170		851FP2		851FP2HF170	29-11-2010	19:29:00	Template Development	29-11-2010 20:43:59
t29-01		Ref-29-11-2010-16-12-48	851FP1		851	851FP1		29-11-2010	16:29:48	Template Development	29-11-2010 21:22:38
t29-01		Ref-29-11-2010-15-12-28	852		852			29-11-2010	15:17:00	Template Development	29-11-2010 21:22:43
t28-10		Ref-28-11-2010-17-03-00	851FP1HF288				851FP1HF288	28-11-2010	17:14:00	Template Development	28-11-2010 18:34:07
t28-09		Ref-28-11-2010-16-13-42	851FP1HF134				851FP1HF134	28-11-2010	16:19:00	Template Development	28-11-2010 17:02:57
t28-04		Ref-28-11-2010-13-27-57	851FP1HF288				851FP1HF288	28-11-2010	14:23:00	Template Development	28-11-2010 16:13:37
t28-03		Ref-28-11-2010-13-12-34	851FP1HF288				851FP1HF288	28-11-2010	13:15:00	Template Development	28-11-2010 13:27:28
t28-02		Ref-28-11-2010-00-37-27	851FP1HF288				851FP1HF288	28-11-2010	12:54:00	Template Development	28-11-2010 13:12:28
t28-01		Ref-28-11-2010-00-02-41	851FP1HF288				851FP1HF288	28-11-2010	00:07:00	Template Development	28-11-2010 00:37:23
t27-07		Ref-27-11-2010-21-55-05	851FP1HF134		851FP1		851FP1HF134	27-11-2010	22:04:00	Template Development	28-11-2010 00:02:38
t27-06		Ref-27-11-2010-20-48-51	851FP1		851FP1			27-11-2010	21:00:00	Template Development	27-11-2010 21:54:57
t27-05		Ref-27-11-2010-19-27-48	851		851			27-11-2010	19:34:00	Template Development	27-11-2010 20:48:48
t27-04		Ref-27-11-2010-18-01-40	852HF137		852		852HF137	27-11-2010	18:10:00	Template Development	27-11-2010 19:27:43
t27-03		Ref-27-11-2010-16-54-34	851FP1		851	851FP1		27-11-2010	17:00:00	Template Development	27-11-2010 18:00:05
t27-02		Ref-27-11-2010-15-24-45	852HF137		852		852HF137	27-11-2010	15:32:00	Template Development	27-11-2010 16:53:54
t24-01		Ref-24-11-2010-09-40-16	852HF137		852		852HF137	27-11-2010	14:40:00	Template Development	27-11-2010 14:53:32
t23-01		Ref-23-11-2010-14-54-02	851FP1HF134		851FP1		851FP1HF134	23-11-2010	15:02:00	Template Development	24-11-2010 09:40:13
t22-02		Ref-22-11-2010-12-29-16	851FP1		851FP1			22-11-2010	12:46:16	Template Development	23-11-2010 14:53:57
t22-01		Ref-22-11-2010-09-55-10	851		851			22-11-2010	10:19:00	Template Development	22-11-2010 12:29:12
t19-07		Ref-19-11-2010-16-25-42	851FP2		851FP2			19-11-2010	18:42:00	Template Development	19-11-2010 17:21:39
t19-04		Ref-19-11-2010-14-24-29	851FP1HF288				851FP1HF288	19-11-2010	14:33:00	Template Development	19-11-2010 16:25:37
t19-03		Ref-19-11-2010-13-17-40	851FP1		851FP1			19-11-2010	13:34:40	Template Development	19-11-2010 14:24:26
t19-02		Ref-19-11-2010-08-40-55	851		851			19-11-2010	11:40:00	Template Development	19-11-2010 13:11:32
t18-02		Ref-18-11-2010-14-03-17	851FP3HF250		851FP3		851FP3HF250	18-11-2010	14:30:00	Template Development	19-11-2010 08:40:51
t18-01		Ref-18-11-2010-14-02-20	851FP3HF250		851FP3		851FP3HF250	18-11-2010	14:19:21	Template Development	19-11-2010 08:40:46
t17-11		Ref-17-11-2010-15-11-04	852HF137		852		852HF137	17-11-2010	15:45:00	Template Development	18-11-2010 14:02:16
t16-02		Ref-17-11-2010-14-26-58	851FP3HF250		851FP3		851FP3HF250	17-11-2010	14:43:58	Template Development	17-11-2010 15:22:00
t16-01		Ref-17-11-2010-12-30-56	851FP2HF170		851FP2		851FP2HF170	17-11-2010	12:47:56	Template Development	17-11-2010 14:26:25
t16-01		Ref-16-11-2010-08-34-05	851FP1HF134		851FP1		851FP1HF134	16-11-2010	11:16:00	Template Development	17-11-2010 12:26:35
t15-06		Ref-15-11-2010-21-21-16	852		852			15-11-2010	21:48:00	Template Development	15-11-2010 22:09:08
t15-01		Ref-15-11-2010-11-46-18	851		851			15-11-2010	13:03:00	Template Development	15-11-2010 15:41:23

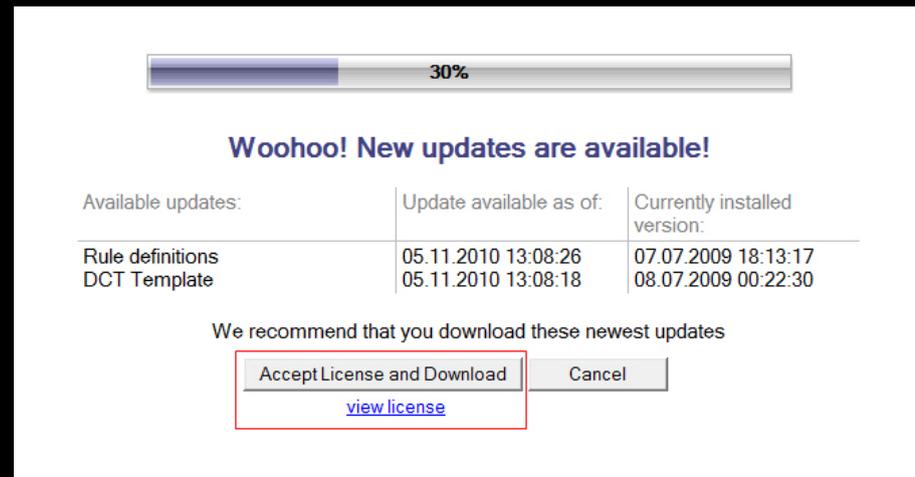
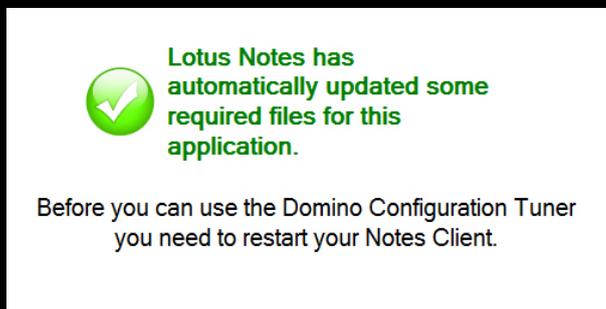


Domino Configuration Best Practices



Domino Configuration Tuner (DCT)

- DCT.nsf as started as a separate project and is now included in Domino
- Local Client application
 - Contains predefined Rules from IBM
 - checks notes.ini, Database Settings and Server Configuration
- Updated directly via Internet
 - Updates Rules and also Template
- Gives recommendations with cross reference
 - You should still double check if parameters match your environment





Run Scan on one or more Servers

- Button to run the scan is a bit hidden

Run New Scan View Results

Servers Update list | Select all | Deselect All **Run Scan**

Select servers to scan from the list.

last updated 18.06.2010 13:38:18

- domino.nashcom.de/Srv/NashCom-Net
- newton.nashcom.de/Srv/NashCom-Net
- notes.nashcom.de/Srv/NashCom-Net



DCT Results

newton.nashcom.de/Srv/NashCom-Net	
Critical	notes.ini CONSOLE_LOG_ENABLED causes the Domino Server to write console data to console.log in the IBM_TECHNICAL_SUPPORT directory.
	notes.ini CONSOLE_LOG_MAX_KBYTES sets the maximum size of the console.log file in Kbytes.
	notes.ini DAOSMinObjSize indicates the size (in bytes) of the smallest attachment that will use DAOS. (all operating systems except iSeries)
	notes.ini DEBUG_CAPTURE_TIMEOUT captures semaphore time stamps and logs them to semdebug.txt
	notes.ini DEBUG_SHOW_TIMEOUT instructs the Domino server to log semaphore information to semdebug.txt
	notes.ini DEBUG_THREADID instructs the Domino server to output verbose logging with thread specific information.
Warning (High)	database ODS version should be as high as possible
	notes.ini CREATE_R85_DATABASES=1 enables the creation and upgrade of database versions to ODS 51
	notes.ini FTG_USE_SYS_MEMORY reduces memory fragmentation related to full text index searching (post 6.5.4 inclusive)
	notes.ini GROUP_CACHE_SIZE specifies the size of the GroupCache (in KBytes.)
	notes.ini LOG_MAILROUTING specifies how much logging detail the mail router sends to the server console and log file
	notes.ini NLCACHE_SIZE specifies the size (in bytes) of the NameLookup cache.
	notes.ini NSF_BACKUP_MEMORY_CONSTRAINED limits the amount of memory used by backup applications and may prevent a server crash
	notes.ini SERVER_MAX_CONCURRENT_TRANS can be used in conjunction with SERVER_POOL_TASKS to control the number of NRPC threads that can be executing server transactions. These settings must be used together.
	notes.ini SERVER_NAME_LOOKUP_NO_UPDATE can prevent unnecessary Domino Directory views updates (post 8.0.1 inclusive)
	notes.ini SERVER_POOL_TASKS can be used in conjunction with SERVER_MAX_CONCURRENT_TRANS to control the number of NRPC threads that can be executing server transactions. These settings must be used together.
	notes.ini TNEFEnableConversion enables Domino conversion of MS Transport Neutral Encapsulation Format (TNEF) data (post 7.0.3 inclusive)
	notes.ini UPDATE_FULLTEXT_THREAD enables using a separate thread to do full text indexing which can improve server performance.
	Server Document 'Store Admin Process log entries when status of no change is recorded' allows the administrator to indicate whether or not 'no change' log documents are saved in admin4.nsf.
	Server Document setting 'Compare public keys' allows an administrator to indicate whether or not Notes and Domino public keys should be compared and at what level.
	ServerDoc 'Transactional Logging' Enabled/Disabled
Warning (Low)	database property 'Compress database design' can reduce I/O demands
	database property 'Compress document data' can reduce I/O demands
	database property 'Don't overwrite free space' can reduce I/O demands
	database property 'Use LZ1 compression for attachments' can reduce I/O demands
	generic notes.ini rule: diagnostic settings should usually remain disabled

- Severity**
- Critical
 - Warning (High)
 - Warning (Low)
 - Normal
 - Exception

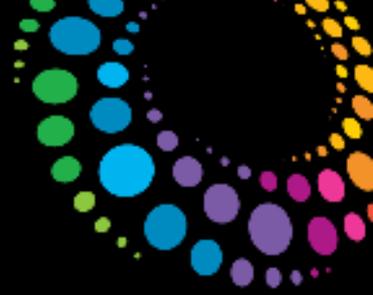
Servers Scanned:

newton.nashcom.de/Srv/NashCom-Net

Rules Evaluated:

- Activity Trends: enable Activity Logging
 - Adminp: the administration process should...
 - database ODS version should be as high as possible
 - database property 'Compress database design' can reduce I/O demands
 - database property 'Compress document data' can reduce I/O demands
 - database property 'Don't overwrite free space' can reduce I/O demands
 - database property 'Use LZ1 compression for attachments' can reduce I/O demands
 - generic configuration document rule: The...
 - generic notes.ini rule: diagnostic settings should usually remain disabled
 - generic notes.ini rule: setting value size of...
- ▶ 138 more...





Additional Refined tips

- NSF_Backup_Memory_Constrained=1 sets backup memory to 20 MB
 - Resonable size e.g. 200 MB NSF_Backup_Memory_Limit=200000000 (value in bytes!)
 - See TN #1211241 for details
- Remove DISABLE_DIR_DEL_IF_EMPTY
 - Old work-around for an unmount issue with junctions on windows
 - Causes FT index directories are not removed when deleting FT index!
- FT_FLY_INDEX_OFF=1
 - Disables on-the-fly FT index creation in agents when simple search queries are used
 - “Extremely inefficient” message in log
- NSF_DbCache_Maxentries=n
 - Increase number of cache entries for larger servers
 - Default 3 times the buffer pool size (new standard 512 MB for Domino 32bit)
 - Too low for large servers. Increase it to the number of users on your system
 - See reference slides for details



Tip



Server Maintenance

COMPACT



Compact -B (online)

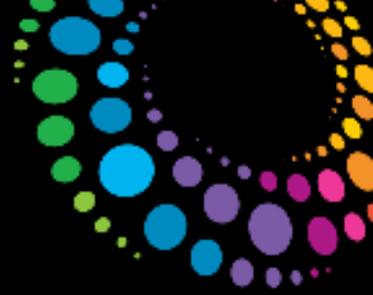
- Recommend running once a week on all databases via a Program Document.
 - To reduce database sizes
 - In case of Archive Style Translog/Backup run compact before backup because it generates a new DBIID – new instance that needs to be backed up afterwards
 - There is a bug in 8.5 where database size is not reduced (fixed in 8.5.2FP1)
- No scheduled fixup needed
 - Use DDM to repair databases automatically (DDM is covered later)
- Check your full-backup runtime to prevent an overlap
- Periodically check semdebug.txt to if there are problems with particular database(s) (more about semdebug later).





Database Quotas – Compact needed?

- If you are using quotas users care if you do a daily compact
 - With Translog enabled they cannot compact their mailfile on the server any more
- But what they really care if their free space in the database counts for quota or not
- You could consider a daily compacts for DBs with large free space
 - E. g. compact with more than 100 MB free space → compact -B -S 100M
 - Take care that with Archive Style Translog you need a new backup (new DBIID)
- Different Methods how quotas a calculated
 - By default the quota checks the total physical file size instead of the used space
 - Setting in Translog Tab to allow Domino calculating the quota based on used space
 - Change “**Quota enforcement**” to “**Check space used in file when adding a note**”
 - This will calculate and display the size based on space used
- This reduces the need for daily compacts :-)



Compact – Multiple Instances

- For large Servers consider using multiple simultaneous compact tasks
 - Notes.ini parameter (DEBUG_ENABLE_COMPACT_8_5=1)
 - **Disclaimer: It's currently not official supported**
- # nn Execute up to nn compactions in parallel using different threads
- -W nn Only compact databases which haven't been compacted successfully in last nn days.
- -w Exclude system databases (e.g. log.nsf, names.nsf).
- -X nn Limit each database compaction to no more than nn minutes.
(applies to in-place compaction only)
- -x nn Limit total compaction time to no more than nn minutes.



Tip



Compact -D (offline)

- Recommend running at least once a year (use a quiet period e.g. Public Holiday).
 - will repair corruption if not seriously damaged (similar to compact -c)
 - Removes view indexes which will reduce the database size significantly
 - !! Caution – first access maybe slow as the view index is rebuilt.
- Use IND files to allow multiple simultaneous compact tasks.
- Or use the multi-Thread compact option # nn
- **Caveat Multi Compact instances are not visible on the Domino Console. But we verified via an NSD that multiple threads are used ;-)**



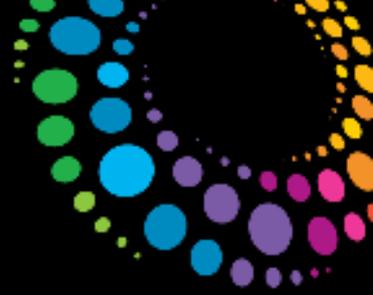
Server Maintenance

DATABASE INDEXERS



Understanding Update Task and its' Queue

- Little control over update, other than starting or stopping the task.
- Loaded at server startup by default and runs continually from it's own queue. Compact immediate requests are also added to the same message queue.
- Update queue is checked every 5 seconds - first-come first-served basis, then batched for the Suppression Time (default 5minutes), before updating.
 - Can be changed with `Update_Suppression_Time=minutes`
- After updating view indexes, it then updates all databases that have full-text search indexes set for immediate or hourly updates.
- The directory indexer runs at one-minute intervals and is dedicated to keeping Domino Directory view indexes up-to-date.



Do multiple Update Task improve performance?

- Updaters= n
- Update was initially designed to only run one process.
- Update task may be in contention for the same database semaphore.
- Check at intervals your semaphore timeouts (topic covered later).
- Another indicator that multiple updaters are contending for the same view is seen in the console.log
 - “Btree context signature is NULL”update
- Consider removing Updaters= and monitor the queue (more later)



Full-Text Index Update Strategies

Reference Technote #1298543

Update frequency	Description
Daily	Performed by the nightly "updall" task. If this nightly task is not run, the daily updating is not performed.
Scheduled	Performed by a Program Document that runs "updall". You must set the frequency to Scheduled and create the proper Program Document. You can also use this method to update different databases at different times.
Hourly	Triggered by the Chronos task and performed by the update task (if the update task is running). If not running, Chronos performs the update. If the Chronos task is not running, the update is not performed.
Immediate	Performed by the Update task. If Update is not running, the update is not performed. All immediate requests are processed as they are received.
Deferred	Deferred requests are held for 15 minutes before they are processed. Requests to update the same database that occur in that time are ignored as duplicate requests.



Updall

- Refreshes views and/or updating full text indexes
- Contrary to Update:
 - It does NOT work from a queue
 - Can use options (e.g. Updall -R) on a single, or multiple databases
 - by Default in notes.ini (ServerTasksat2=Updall)
- Caveats:
 - ServerTasksat2=Updall means it runs “**every day**” and may conflict with other operations (e.g. weekly backup)
 - Without options only refreshes outdated views and full text indexes and discards view indexes.
 - Looses visibility in the notes.ini
- Recommendation:
 - Remove from notes.ini & create a program document avoiding full backup runtimes

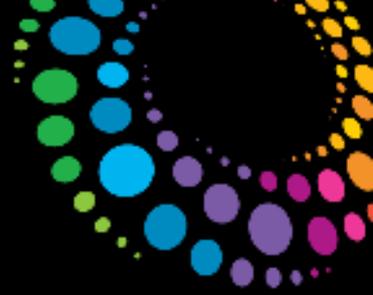


Tip



Server Maintenance

FULL TEXT INDEXES



Full Text Index sizes

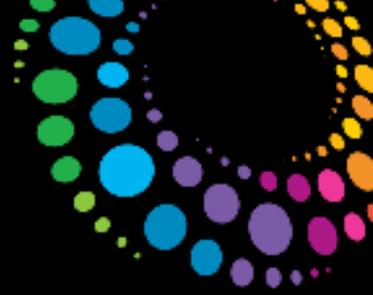
- If database sizes are reduced (i.e. mail archiving) the corresponding FT Index size is NOT automatically reduced.
 - Seen examples of 2GB mail databases and FT Index is 1GB
 - After Archiving the mail database 200MB but the FT Index remains 1GB.
- Recommend - monthly deletion of all FT Indexes and rebuild
 - Via a program document Updall -X
 - Will help improve FT search times as well as reducing disk space used
- Caveat – Minor obstacle, until IBM improve the Program documents scheduler forced to make this a weekly operation
 - Enhancement Request to for a monthly and quarterly option SPR #ETHU5GHR3Q
 - If you have the same requirement you might want to open a PMR to add weight to the SPR



Full Text Indexing - Dedicated Thread

- Full Text indexing is often inhibited by the Update task
 - err_invalid_btree ERROR: Btree context signature is NULL [F:\Lotus\Domino\data\app\ABC.nsf]
 - Error full text indexing app\ABC.nsf: Database is currently being indexed by another process.
 - Set server notes.ini parameters (in DCT catalog of best practices):
 - Update_Fulltext_Thread=1
 - FTUPDATE_IDLE_TIME=4
 - Restart the update task for the changes to take effect.





Full Text Indexing - Update_Fulltext_Thread=1

- Enter “sh stat update” on the Admin console
- Stats explanation <http://www.ibm.com/support/docview.wss?rs=203&uid=swg21258366>

without Update_Fulltext_Thread=1	With Update_Fulltext_Thread=1
Update.DeferredList = 197	Update.DeferredList = 8
Update.DeferredList.Duplicates = 114719	Update.DeferredList.Duplicates = 4869
Update.DeferredList.Max = 459	Update.DeferredList.Max = 24
Update.DeferredList.NextDueTime = 12/20/2010 15:02:46 CET	Update.DeferredList.NextDueTime = 12/20/2010 15:28:11 CET
Update.DeferredList.Processed.AllViews = 66265	Update.DeferredList.Processed.AllViews = 2778
Update.NAB.Updates = 604	Update.DeferredList.Processed.FullText = 4593
Update.PendingList = 141	Update.FullTextList = 0
Update.PendingList.Duplicates = 20995	Update.FullTextList.Max = 30
Update.PendingList.Max = 281	Update.FullTextList.Processed = 25340
Update.PendingList.Processed.AllViews = 55699	Update.NAB.Updates = 695
Update.PendingList.Processed.Compactions = 34	Update.PendingList = 0
Update.PendingList.Processed.SingleView = 481	Update.PendingList.Duplicates = 295
	Update.PendingList.Max = 2
	Update.PendingList.Processed.AllViews = 47
	Update.PendingList.Processed.FullText = 48
	Update.PendingList.Processed.SingleView = 291



Full Text Indexing – Improving Merging

- How to fix and/or prevent:
 - R8.5.2 and earlier:
 - Optimize merging operations by full text index reorg regularly via “updall -x”
 - R8.5.3 and later, 2 notes.ini variable controls:
 - FTG_MERGE_THRESHOLD = x - the number of database updates before FT merge (default is 100)
 - FTG_MERGE_HOURS = low, high - the hours between which merges are allowed (e.g. FT_MERGE_HOURS=23,2 will allow from 11:00PM to 2:59:59AM)



View_Rebuild_Dir

- Changes the way view rebuild is working
 - Optimized for speed instead of RAM used
 - Make sure the directory is in the Operating System Anti-Virus exclusions list
 - Domino should delete all temporary index files (DTF Files) by a Server restart

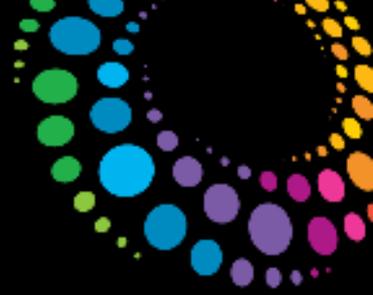
- Use fast disks preferably local – but fast SAN disk would also work
 - Don't use the same disk as Transaction Logging
 - You can share **view_rebuild_dir** with physical OS disk but should use a separate file-system e.g. logical disk (fragmentation, temporary files)
 - You might want to use a temp file-system on Linux/Unix



Keep your Disk defragmented

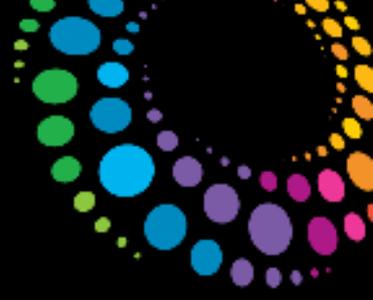
- Domino uses smaller incremental allocations for new documents which can cause fragmentation
 - Fragmentation can be a performance issue specially for backup operations on large sets of data
 - Best Practice is to defragment once per week at off-peak times
- Several tools on the market for Windows. Some of them are free
 - Contig V 1.55 – Free Tool from Microsoft
<http://technet.microsoft.com/en-us/sysinternals/bb897428>
 - Domino Defrag 3.2 – OpenNTF Project
 - <http://www.openntf.org/internal/home.nsf/project.xsp?action=openDocument&name=DominoDefrag>
 - Servertask uses Contig internally



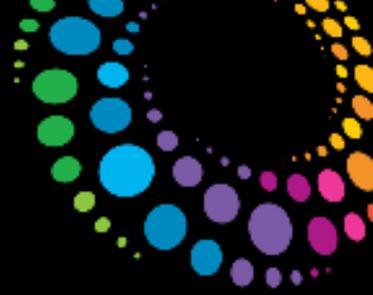


General Recommendations on File-Systems

- It's a general best practices to have at least 30% free space available to allow the file-system to optimize where to write data and avoid fragmentation
- Keep file-systems below 1 GB on all platforms
 - Best practice from performance and also disaster recovery point of view
 -
 - Tip: Use mounts on Linux/Unix and junctions on Windows
 - If you use DAOS the file-system used for NSF is up to 70% smaller
 - DAOS File-system needs less resources and is written once per file sequentially
 - Less fragmentation and defragmentation effort
 - Diskspace requirement for NSF does grow linear instead of exponential



Server Operation



Server Operation – General

- Many of Dominos activities are controlled by the names.nsf
 - Use the dircat task and create a synchronized Version of the names.nsf
 - Use the sample script provided in this session to copy across those documents dircat omits.
 - Create as many none-standard views as need without affecting Directory lookups.
 - Ever deleted critical person/server/configuration documents from the names? Then use the synchronized Version without having to scurry for Backup tapes.

Server Document where dircat should run

Server: ACMEHUB01/Server/ACME ACME.com

Basics | Security | Ports... | Server Tasks... | Internet Protocols... | MTAs... | Miscellaneous

Administration Process | Agent Manager | Domain Catalog | **Directory Cataloger** | Internet

Basics

Directory Cataloger filenames:

Schedule

Schedule:

Run Directory Cataloger task at:

Repeat interval of:

Days of week:

Target names → In view Directory

Extended Directory Catalog

Basics | Advanced | Administration

Basics

Directories to include:

Additional fields to include:

Note: No fields means ALL

Remove duplicate users:

Group types:

Include Mail-In Databases:

Include Servers:

Restrict aggregation to server:

Send Aggregation reports to:

Comments:

- Sample modified names.nsf with None-Standard views available



Sample modified names to Sync your NAB

- Brings what's at document level to the view level

NABSync

Active On	On Server	Run Program	Command Line	Schedule	Active On
▼ Friday					
▼ TestMailCluster01/USA					
		Updall	-R	03:00	So,Mo,Tu,We,Th,Fr,Sa
		Updall	-X	04:00	So,Mo,Tu,We,Th,Fr,Sa
		Updall	-Mistake	04:00	So,Mo,Tu,We,Th,Fr,Sa
		Compact	-B	20:00	Fr
▼ TestMailCluster02/USA					
		Updall	-R	04:00	So,Mo,Tu,We,Th,Fr,Sa
		Compact	-B	20:00	Fr
▶ Monday					
▶ Samstag					
▶ Sunday					
▶ Thursday					
▶ Tuesday					
▶ Wednesday					

Domino Directory

On Server	Run Program	Command Line	Schedule	Repeat
▼ TestMailCluster01/USA				
	Updall	-R	03:00	
	Updall	-X	04:00	
	Updall	-Mistake	04:00	
	Compact	-B	20:00	
▼ TestMailCluster02/USA				
	Updall	-R	04:00	
	Compact	-B	20:00	



Proactive Monitoring



Lotus Domino Statistics

- Valuable resource of information
 - Combines Domino Statistics and Platform statistics
 - Platform statistics depend on the OS platform but are sort of unified between platforms
 - Check events4.nsf for a description of each platform stat per platform
- You should collect Server stats at least every 15 minutes (default is 90 minutes)
 - Enable collect task, configure settings in events4.nsf
 - Configure statistic events for important stats with the right thresholds
 - Keep long term data to compare current and historic data
 - You can also leverage SNMP to query stats
 - Limitation: Only works for 1 partition per OS instance on all platforms
- Very detailed session about how to get most out of the stats collected in statrep.nsf
 - SHOW102 Statistics and Events Base Camp: Proactive Monitoring of Your IBM Lotus Domino Servers



Server Availability Index (SAI)

- Can give you and “idea” / “indication” for current server response times
- SAI was broken until Lotus Domino 8.5
 - Finally fixed D8.5 :-)
- SAI calculation on fast servers still might not work for you out of the box
 - Internally LOADMON routine uses micro seconds
 - On a fast server at off-peak times transactions can take a couple of micro seconds
 - Compared to normal performance (e.g. 1 ms) this can result in very high XF
 - Causes a low SAI for normal performing servers
- “Tuning”: D8.5 Set range of minimum and maximum values
 - notes.ini: **Server_MinPossibleTransTime=1500**
 - notes.ini: **Server_MaxPossibleTransTime=20000000**
 - Important: Delete **loadmon.ncf** when server is down to delete old minimum values
 - Tip: Nash!Com Domino on Linux/Unix start script can delete loadmon.ncf automatically



**Best
Practice**



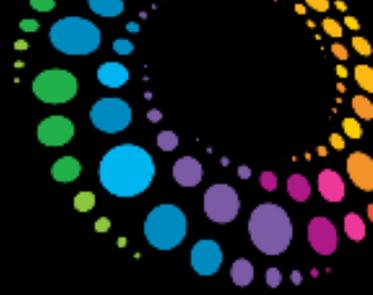
Domino Domain Monitoring (DDM)

- Comprehensive Monitoring
 - ddm.nsf contains focused monitoring results
 - Detailed error messages including names of resources
 - Suggestions for problem solution including actions!
- Based on the foundation build by event monitoring
 - Event categorization and severity defined in events4.nsf
- Additional build in probes into the code
 - Replication (detailed reporting for failing replication)
 - Agent Manager (long running agents, high memory/CPU usage, ...)
 - Compact / Fixup reports detailed information about corrupted documents etc
- You can also leverage statistics and platform statistics



The Five C's of DDM

- Consolidate
 - Track multiple, related errors in a single event report
- Check
 - Assess health of functional areas
- Collect
 - Access all domain wide reports from a single database
- Correlate
 - Locate related reports from other servers
- Correct
 - Assess knowledge base of explanations and recommendations.
 - Click to resolve issues.



DDM Probe Configuration

- Three flavors of probes (some probes have more than one):
 - **Scheduled probes** - run according to a configurable schedule and defaults are supplied
 - **Embedded probes** - run inside feature area and catch problems/issues as they occur
 - **Listening probes** - run when particular error codes are logged
- Probe configuration is quick and flexible:
 - Defaults probe configuration documents supplied for “out-of-box” values
 - “Special target servers” concept allows out-of-the-box probing without having to specify named servers
 - Thresholds and result content are highly customizable what the probe will actually check
 - probe sensitivity (when will they generate an event)
 - what severity event the probe will generate
 - Schedule is highly customizable for scheduled probes
- Probes can be enabled/disabled per server/server group



DDM Example

- Very detailed information about corrupted database including actions for fixup

Assign... Change State Add Comments... Common Actions

Open Event
(This event is not assigned to anyone)

Generated by: [redacted]
in Domain: [redacted]
Event class: Enhanced

For server: [redacted]
Database: [redacted]
Agent: [redacted]
User: [redacted]

Most Recent Event | Correlated Events | Prior Events

Fatal

Friday, 10. December 2010 - 16:32:07
Database is corrupt – Cannot allocate space

Explanation

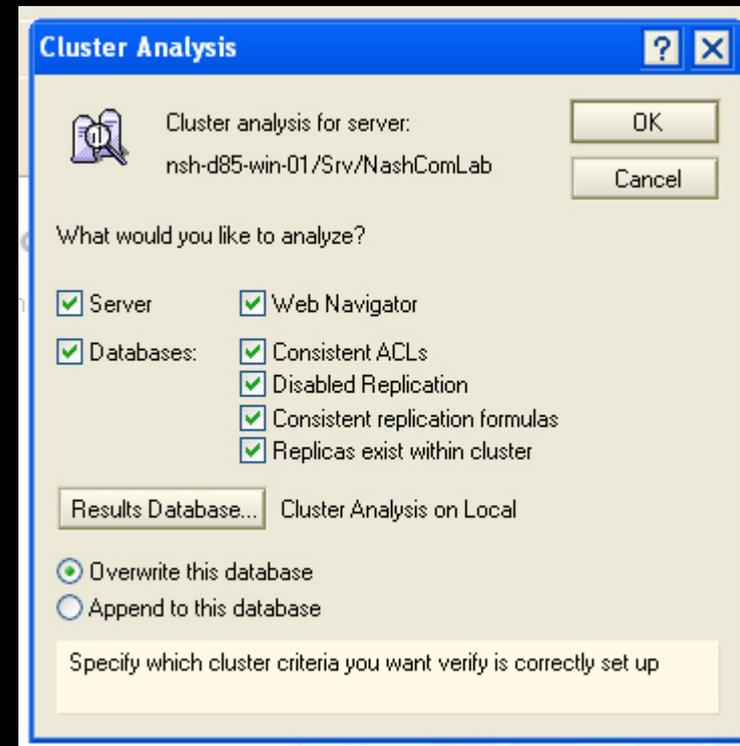
- Reported by: Router
- Severity and type: **Fatal** in Database
- Probable cause: The most common causes of database corruption are improper shutdown of the operating system due to system reboot, system crash, power failure, or failure to follow proper shutdown procedure. An API program may be improperly accessing the database.
- Possible solution: The database should be fixed-up automatically the next time it is opened. Run the FIXUP task from the Domino server console.
- Corrective action:
 - Issue a command to the server [redacted] to run FIXUP on the target database - "Load Fixup -F -J [redacted]"
 - Inspect and manipulate databases on [redacted] using tools from the Domino Administrator 'Files' tab.

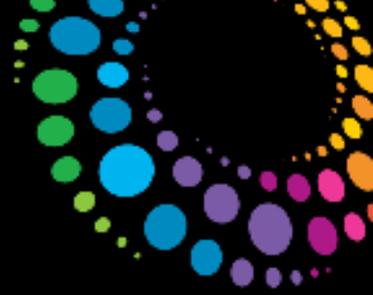
Event Change History:
10.12.2010 01:46 PM [redacted] - changed state to Open



Analysis Tools

- Domino Admin Client has build-in analysis Tools
 - Located in Server/Analysis Tab
 - Cluster Analysis
 - Log Analysis
- You should regularly run a Cluster Analysis to check that all databases can replicate correctly
- Activity logging can also help for troubleshooting
 - Needs to be enabled in Server Config Document





Recovering from crashes / Crash Analysis



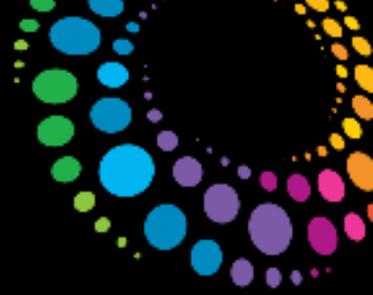
Take benefit of Domino RAS Functionality

- RAS = Reliability Availability Serviceability
- RAS is the effort to improve the Domino Product suite so that:
 - Reliability - Client/Server doesn't crash or hang as often
 - Availability - Client/Server performs well, Server is available to clients
 - Serviceability - The ability to quickly pin-point and fix problems
- Ongoing effort in each incremental release
 - Some features are even back-ported from Domino 8.5.x to Domino 7
- It's not just about NSD & Memcheck but all parts of Domino
 - Logging, Debug Options, etc
- Great help for Admins, Developers and Troubleshooters



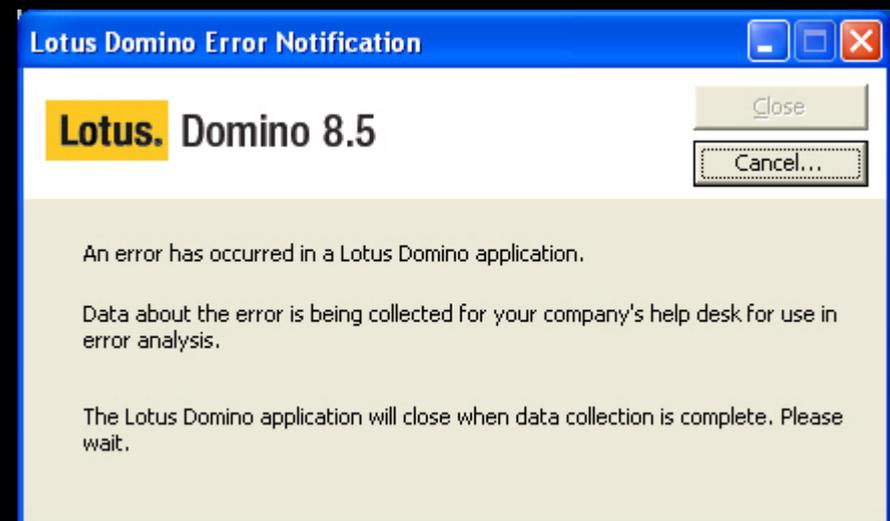
Fault Recovery

- Domino Server detects crash and restarts automatically
 - Panic routine calls fault recovery code
- How Fault Recovery works
 - Run NSD if configured
 - Cleans up resources
 - Restarts Server
- Enabled in Server document
 - Run NSD To Collect Diagnostic Information: **Enabled**
 - Automatically Restart Server After Fault/Crash: **Enabled**
 - Mail Fault Notification to: **LocalDomainAdmins**
 - D8.x: Improved crash info contained within e-mail notification
- Results
 - Server is restarted quickly without admin intervention
 - You have all logs and diagnostics from each server crash for analysis



Automated Diagnostic Collection (ADC)

- Enables you to set up a mail-in database to collect the diagnostic information generated from the ND Client/Server crashes in one central repository.
 - **Senddiag** server task runs on startup to collect information like NSDs
- Configured in Server Configuration Doc / Diagnostics Tab
 - Fault-Report Database (Indfr.nsf) as Mail-in Database
 - Size for diagnostic data, retention days, ...
 - Filter pattern to add to data collection (file-patterns!)
- Tip: ADC and Fault Recovery is also available for Notes Clients
 - Configured in desktop policy
- Caution Mail Size Limit in Config Document also applies for diagnostic mails (Client+Server)





Fault Analyzer Servertask

- Fault Report Database
 - Typically located on admin server
 - Mail-In Database with Fault Report Template
 - Diagnostic TAB of Configuration Doc
- Fault Analyzer Servertask
 - Used to annotate, categorize NSDs
 - Similar call-stacks, Same Domino releases, Client or Server



**Best
Practice**

- You should have separate databases for clients / servers
 - But template already has different views for servers and clients

Diagnostic Collection Options	
Mail-in Database for diagnostic reports:	<input type="text" value="FaultReports"/>
Maximum size of diagnostic message including attachments (in MB):	<input type="text" value="50"/>
Maximum size of NSD output to attach (in MB):	<input type="text" value="40"/>
Maximum amount of console output file to attach (in KB):	<input type="text" value="10240"/>
Diagnostic file patterns:	<input type="text" value="semdebug*.txt"/>
Remove diagnostic files after a specified number of days:	<input type="text" value="Yes"/>
Number of days to keep diagnostic files:	<input type="text" value="120"/>
Fault Analyzer	
Run FaultAnalyzer on Fault DBs on this server:	<input type="text" value="Yes"/>
Run Fault Analyzer on:	<input type="text" value="All mail-in databases on this server"/>
Remove attachments from duplicate faults:	<input type="text" value="No"/>



Fault Reports – Free Modified NTF as Sample

Use the Views to see what your Servers are telling you

The screenshot shows the 'Fault Reports' window in the 'Fault Analyzer' application. The interface includes a toolbar with 'Open Report', 'Edit Report', 'Forward', and 'New Response' buttons. A tree view on the left shows various filters like 'By Notes/Domino Version', 'By Date', 'By OS Version', 'By Uptime', 'Clients', and 'Servers'. The main table displays the following data:

User Name	Occurrences	Unique IDs
Release 8.5 HF447 May 14, 2009	100	
Release 8.5FP1 HF147 August 18, 2009	360	
Release 8.5FP1 HF242 September 16, 2009	6	
Release 8.5.1FP1 HF25 February 03, 2010	10	
	476	



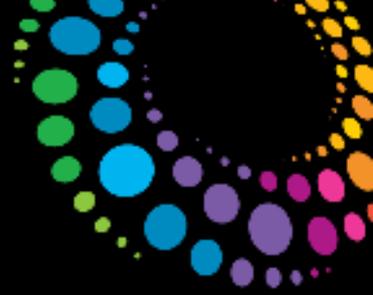
Fault Reports – Database

Are other Servers experiencing identical crashes?

Tip – Open an IBM Support PMR and stipulate “data to follow”. Then use the forward button. (Important enter PMR number in the Subject field)

The screenshot shows the 'Fault Analyzer' application window. On the left is a navigation pane with categories like 'By Notes/Domino Version', 'By Date', 'By OS Version', 'By Uptime', 'Clients', and 'Servers'. The main area displays a table of fault reports. A red box highlights a group of entries under the 'APPL03/Server/ACME' folder, including several 'Additional Occurrence' entries.

Notes/Domino Version	Occurrences	Unique IDs	Process
APPL03/Server/ACME			
Release 8.5 HF447 May 14, :	7	2	nAMgr
Release 8.5 HF447 May 14, :	1	1	nAMgr
Release 8.5 HF447 May 14, :	5	1	nAMgr
Release 8.5 HF447 May 14, :	1	1	nAMgr
Release 8.5 HF447 May 14, :	1	1	nAMgr
Release 8.5 HF447 May 14, :	4	1	nAMgr
Release 8.5 HF447 May 14, :	2	1	nAMgr
Additional Occurrence (Exact fault for APPL02/Server/ACME on			
Release 8.5 HF447 May 14, :	1	1	nAMg
Release 8.5 HF447 May 14, :	7	2	nAMg
Additional Occurrence (Exact fault for APPL02/Server/ACME on			
Additional Occurrence (Exact fault for APPL02/Server/ACME on			
Additional Occurrence (Exact fault for APPL02/Server/ACME on			
Additional Occurrence (Exact fault for APPL01/Server/ACME on			



Fault Reports – Sample database zips data

The Free Sample

- Zips the attachments and zip file-name is the PMR number + Server name .zip
- Also a simple Action button to open a live admin console within the form

Fault Report

Diagnostic Data

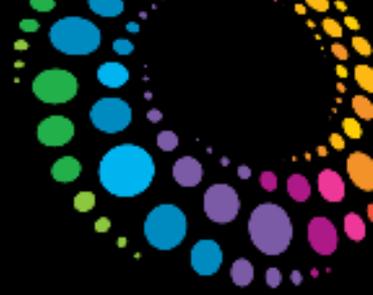
Name:	APPL01/Server/ACME
Notes/Domino Version:	Release 8.5.1FP3 HF250 July 07, 2010
Machine Name:	APPL01
OS Versions:	
Start Time:	08.12.2010 04:10:26
Crash Time:	09.12.2010 02:47:39
Uptime:	0 day(s) 22:37:13
Outage Length:	0 day(s) 00:16:12
Error message:	PANIC: DBCONTWritePage: Untwiddle failed (i:\zo01sn10\la_tkalkbre.nsf Dbid: 0x5D57465E 0x...
Processes:	
Process List:	C:\Lotus\Domino\Inservice.exe C:\Lotus\Domino\InSERVER.EXE



Fault Reports – Use it for Problem Management

The Free Sample database contains enhanced views to Manage your issues

Fault Reports		<input type="button" value="Open Report"/> <input type="button" value="Edit Report"/> <input type="button" value="Forward"/> <input type="button" value="New Response"/>				Resolved	Status
Step	Action Date	Notes/Domino Version	PMR Ref	SPR#			
APPL01/Server/ACME							
<input type="checkbox"/> Release 8.5.1FP3 HF250 July 07, 2010 1234123,999,999 XX1234123 ✔ Waiting for IBM							
1	27.01.11	✔ Notes.ini settiing - Update_Fulltext_Thread=1					
2	28.01.11	⚠ Talk to Vendor					
3	30.01.11	✖ Run database maintenance					
4	31.01.11	✔ PMR History					
<input type="checkbox"/> Release 8.5FP1 HF250 September 18, 2009 ✖ Waiting for Vendor							
1	27.01.11	⚠ Notes.ini settiing - Create_R85_Log=1					
		Release 8.5.1FP1 HF25 February 03, 2010					No Status
		Release 8.5 HF447 May 14, 2009					No Status
		Release 8.5.1FP1 HF25 February 03, 2010					No Status
		Release 8.5FP1 HF147 August 18, 2009				✖	No Status
		Release 8.5FP1 HF147 August 18, 2009				✖	No Status
		Release 8.5FP1 HF147 August 18, 2009				✖	No Status



Fault Reports – Use it for Problem Management

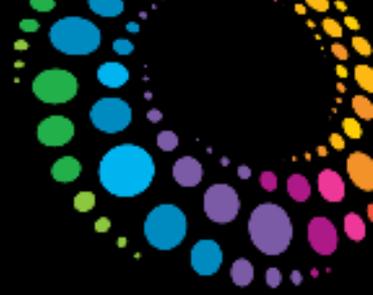
The Free Sample – Can you remember when and why you implemented a particular notes.ini parameter?

Fault Reports

Open Report Edit Report Forward New Response

INI	Step	Action Date	Notes/Domino Version	Status
▼ Create_R85_Log=1				
APPL01/Server/ACME				
1	27.01.11	Release 8.5.1FP3 HF250 July 07, 2010		Open
▼ Update_Fulltext_Thread=1				
APPL01/Server/ACME				
1	27.01.11	Release 8.5.1FP3 HF250 July 07, 2010		Successful

By Notes/Domino Version
None Standard - By Version
None Standard - By Actions
None Standard - By notes INI
By Date
Standard
Fault Analyzed
By OS Version
By Uptime



Fault Reports – Make the most of the database

- Open the Fault Report and run standard database maintenance on the open databases listed (Fixup, Compact and Updall)
- Use the Administrative Section of the Fault Report. Enter your own reference if there is no IBM PMR number
- Create Responses
 - Documents for all actions taken a form of simple step-by-step action plan
 - Separate ones for IBM recommendations
 - Separate ones for any parameters set (e.g. special debug “Settings”)
 - Final Response document should be the entire PMR history
 - Create own views sorted by “PMR ID” and “Settings”
- Caveats
 - The Original Server attachments (nsd's console.log etc) are not zipped
 - Maximum mail size to IBM is 20MB. Larger than 20MB must be sent to IBM ftp
 - No link to the IBM Knowledgebase (LND tool has this capability)



NSD - Notes System Diagnostics

- Has been around for years in Domino on Unix
 - Fully available since Domino 6.0 for Win32
 - Replaced RIP in Domino 6 for Win32 long time ago
 - Not a „Just in Time“ (JIT) Debugger
- It's invoked automatically if Server/Client crashes
 - Or you can manually invoke it for troubleshooting
- NSD provides a huge collection of system diagnostics information on Domino and Operating System level
 - Used by Admins, Developers and Support for Troubleshooting
- NSD and Crash Analysis would be a complete session on it's own



SYM File Support for Add-On Products

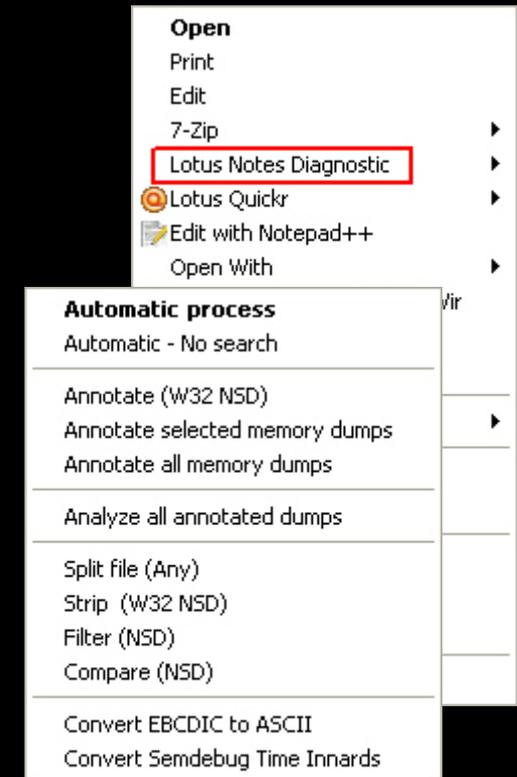
- On Windows add-on products need SYM files to annotate call-stacks in NSD
 - Needed to quickly pin point problems
 - Without SYM files there is no way to annotate those call-stacks quickly
 - Every C-API based application should ship with SYM files
- Domino uses a special SYM file format integrated into one large SYM file
 - Microsoft mapsym cannot be used to generate sym files for Notes/Domino
 - Since D6.5.1 Domino is able to read SYM files for individual binaries
- Lotus Development (Iris) Tool Map2iSym is part of the Lotus C-API Toolkit
 - Ability for NSD to integrate 3rd party "Domino family products"
 - Starting with D6.5.1 NSD it works also extended Domino products
 - Has not been shipped with some versions of the C-API toolkit
 - If you cannot find a current version we are happy to send it to you
 - Known Issue: For Domino 64bit you need to set the following link option
 - /BASE:0x400000





Lotus Notes Diagnostics (LND)

- Tool to annotate NSDs, semdebug files, memory dumps etc
 - Current Version 2.9 downloadable from IBM
 - <http://www.ibm.com/support/docview.wss?rs=899&uid=swg24019151>
 - Also works with Win7 64bit!
- Can be used to annotate crash NSDs
 - Ships with notes database, plugs into Explorer
- Very helpful tool
 - Helps you to find crashing call-stack and categorizes
 - the various NSD sections
 - Also matches the data section of the thread in memcheck
 - But you still have to know much about the background to interpret the results
 - Can be used to annotate SEMDEBUG.TXT files
 - Timestamps, match against NSDs





Enable Semaphore Debugging

- Domino uses Semaphores to protect resources for concurrent access
 - Databases, Views, Thread Pool, etc
 - Without Semaphores (or Mutexes) no multi-threaded/multi-tasking environment would work
- But if semaphores are locked for too long other parts of the servers might need to wait too long to continue processing
 - AKA hang situation
 - In rare conditions you can run into dead-lock situations
- To be prepared you should keep semaphore debugging permanently enabled
- Semaphores are logged when held for longer than 30 seconds
 - `debug_threadid=1` (dynamic - on by default from 8.5.3)
 - `debug_capture_timeout=1` (restart required - on by default from 8.5.3)
 - `debug_show_timeout=1` (restart required - on by default from 8.5.3)
 - Generates SEMDEBUG.TXT in IBM_TECHNICAL_SUPPORT



Client Clocking

- Can be used to track Notes Client/Server Transactions (NRPC)
 - Logs
 - transaction name
 - transaction data
 - response time (ms)
 - bytes send, received
- Example:
 - (15-78 [15]) OPEN_NOTE(REPC1256B16:0072BCBE-NT00000E3E,00400020):
0 ms. [52+1454=1506]
- Enable on Client via
 - client_clock=1
 - debug_console=1
 - Enables a debug text window -- never close this manual, causes a crash
 - debug_outfile=c:\debug_notes.log
 - Writes a debug log file



Client_Clock Use-Cases

- Test number of transactions and data send over the network before your deploy a new application as a best practice
 - Performance might be good in a local test environment but could be painful in WAN
 - Might work well for small number of users but many transactions could cause a lot of load on server
- Can be used to trace performance bottlenecks on client side
 - Enable in background and let it run for a while
- Use free annotation tool OpenNTF Client_Clock Parser
 - <http://www.openntf.org/Projects/pmt.nsf/ProjectLookup/Notes%20RPC%20Parser>
- Caution: Can be only annotated against the same server because Note-ID instead of UNIDs are logged



Client_Clock Reference: Selected Transaction Types

START_SERVER	Start User Session
OPEN_DB/CLOSE_DB	Open/Close a database
OPEN_NOTE/NIF_OPEN_NOTE	Open a Note
UPDATE_NOTE	Update a Note – there is no close transaction
OPEN_COLLECTION/ CLOSE_COLLECTION	Open/Close a view/folder collection
READ_ENTRIES	Reads data from a view/folder
UPDATE_COLLECTION	Updates a view/folder collection
FIND_BY_KEY	Finds notes in a view/folder collection
FINDDESIGN_NOTES	Finds design notes
SEARCH	Search operation with formula
GET_MODIFIED_NOTES	Find table of modified notes
GET_ALLFOLDERCHANGES_RQST	Get changes in all folders Unread count in mail folders



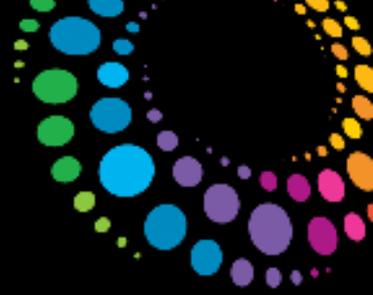
Client_Clock Reference: Selected Transaction Types

NAME_LOOKUP	Lookup information in Domino Directory
ALLOC_UPDATE_OBJECT	Create or update an object
READ_OBJECT	Read data from an object
WRITE_OBJECT	Write data to an object
READ REPLICATION HISTORY	Read Replication History
DB_INFO_GET	Get database info buffer
GET_NOTE_INFO	Get Note information
DB_MODIFIED_TIME	Get the modified date of DB
DB_REPLINFO_GET	Get Replication info
POLL_DEL_SEQNUM	Get Delivery Sequence Number
DB_GETSET_DEL_SEQNUM	Get or set Delivery Sequence Number
GET_SPECIAL_NOTE_ID	Get a special Note-ID by number



Server_Clock & Show Trans

- The server keeps track of all transactions
 - Also used for LOADMON (part of SAI calculation)
 - You can display transaction summary via “show trans”
 - And reset the summary counters via “show trans reset”
 - Or you can display transactions via console log via `server_clock=1`
- `Server_clock=1` has some limitations
 - Only shows transaction information but no user or database information
 - 38965515 ms 'OPEN_DB' 0 ms (0 ms NETIO) TCPIP 000403B1 Rcvd 0 Sent 216
- New `server_clock` options have been introduced in Domino 8.5.1

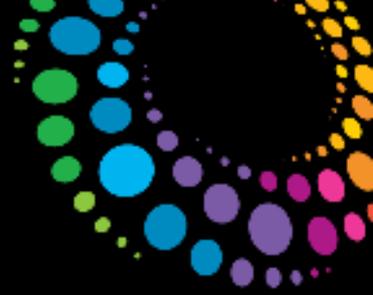


Additional Server_Clock in Domino 8.5.1

- Server_Clock=2
- Will dump more information
- Username, Database, IP Address, and if transaction is used for LOADMON (Lm 1)

```
39255671 ms 'OPEN_DB' 0 ms (0 ms NETIO) TCPIP 000403B1 Rcvd 0 Sent 254  
User 'Daniel Nashed/NashCom/DE' Db 'acl.nsf' Ip '192.168.100.3' Lm 1
```

- Server_Clock=3, DEBUG_TRANSACTION_TIME=n
 - Dumps only transaction taking longer than the specified time
 - Can help to reduce the number of transactions dumped and only lists “slower” transactions
 - For example: 5000 ms
 - Take care: But some transaction like open view collections might take longer than 5 seconds without indicating a problem



Check out the following presentations

- BP111 "CSI Domino": Investigating Server Crashes and Deciphering NSD's
 - Speakers: Francie Tanner, Susan Bulloch
 - SW 7-10 - Wednesday 4:15pm - 5:15pm

 - SHOW102 Statistics and Events Base Camp: Proactive Monitoring of Your IBM Lotus Domino Servers
 - Speakers: Andy Pedisich, Rob Axelrod
 - SW Osprey - Sunday 08:00am - 10:00am
- Created using the



Q&A

- Thanks for your attention!
 - Questions?
 - Now , after the session or via email

- Harvey Pope
 - harvey.pope@hp.com

- Daniel Nashed
 - nsh@nashcom.de
 - <http://blog.nashcom.de>
 - <http://www.nashcom.de>





Additional Material

- The following slides are intended as additional material



General File-System/Disk Considerations

- Mis-configured I/O Subsystem leads to performance bottlenecks!
 - Use RAID1/RAID10 instead of RAID5 at for Transaction Log
 - Have separate disks for translog, view-rebuild, data
 - Put different file-systems on different controllers/ multi channel controllers for large servers
- Have always at least 20-30% free space on file-systems for data
- SAN (Storage Area Network) is faster than local disks in most cases
 - It is recommended to have different file-systems in different SAN volumes
- Use LVM (Logical Volume Manager)
 - Acts as a layer between the physical disk and file-systems
 - Enables you to resize file-systems and add more disks/SAN space
 - Depending on SAN use multiple virtual disk per Volume Group
- Don't use UNIX symbolic links or dir-links
 - Use mounted file-systems instead – also on Windows



File-System Recommendations - “Notesdata”

- Recommended: RAID5 or better RAID10 if you can
 - Allocation Size: 16K
- Access pattern:
 - Random Access I/O – 16K blocks
 - Files Opened Fully Buffered
 - Domino does not do simple predictable sequential reads
- Read Ahead Caching should be set to a low value
 - Unrequired read-aheads can increase the I/O traffic, as well as filling a SAN’s I/O cache buffer with unwanted blocks
 - But you need read ahead caching for backup and some read ahead helps Domino performance too
- Faster drives perform disk I/O in less time
 - Choose drives with a combination of low seek time & high RPM



File-System Recommendation “Translog”

- Size: 5 GB for Circular Translog
 - For archive-style translog size highly depends on your data
- Recommended: RAID1
- Allocation Size: 4K
 - Use larger disk block size and matching Stripe size
- Access Pattern:
 - TX logging writes fixed sequential 4k blocks
 - Transaction log files are opened in a synchronous mode
 - OS file system cache is not used
 - NVRAM cache in the disk subsystem helps - make the size of the largest write to the transaction log files.
- On SAN use 4Gbit Fibre Channel rather than 1 GB.
- Have a dedicated channel



File-System Recommendations - “DAOS”

- Recommended: RAID5 to save costs
 - Allocation Size: 256K up to the size of an attachment – several MB
- Access pattern:
 - Sequential I/O Read – 16-32K blocks
 - Sequential I/O Write – 16-32K blocks
 - Only occurs with new object – static afterwards
- Read Ahead Caching depends on the performance you need
 - You need read ahead caching for backup and normal operations
- Entry Level Storage is sufficient from performance point of view
 - Low I/O rate and specially low IOPS



Conclusion Storage

- You can use SAN or Local Storage
 - Both work great and makes sense depending on your environment
 - What you use depends on many factors including corporate strategy
- In case of SAN performance really depends on the right best practices
 - Tuning
 - Latest Versions, Drivers etc
- Domino has high demand on storage performance
 - It's getting better in Domino 8.x / 8.5 with storage optimization
 - DAOS Store has much lower I/O pattern -> only needs entry level storage
 - Take benefit on Local Drives already existing in the machine even if you plan to use SAN
 - e. g. local translog



Top Statistics – NSF Buffer Pool

- Used for Buffering Database I/O
- Check Server Stats
 - Database.Database.BufferPool.Maximum.Megabytes
 - Database.Database.BufferPool.PercentReadsInBuffer
- Interpretation
 - Bad < 90% < PercentReadsInBuffer < 98% < Perfect
- Tune: notes.ini NSF_Buffer_Pool_Size_MB=n (in MB)
 - Default: 512 MB



Top Statistics – NSF Cache

- Used for Caching Open Databases
- Check Server Stats
 - Database.DbCache.HighWaterMark
 - Database.DbCache.CurrentEntries
 - Database.DbCache.MaxEntries
 - Database.DbCache.OvercrowdingRejections
- Interpretation
 - Good = HighWaterMark < MaxEntries
 - Good = 0 OvercrowdingRejections
- Tune: notes.ini NSF_DbCache_MaxEntries = n
 - Default: NSF_BUFFER Pool size multiplied by 3



Top Statistics – (Cluster) Replication

- Use to check Cluster Replicator Performance
- Check Server Stats
 - Replica.Cluster.Failed
 - Replica.Cluster.SecondsOnQueue
 - Replica.Cluster.WorkQueueDepth
- Interpretation
 - Perfect < 10 < SecondsOnQueue > 15 > Bad
 - Perfect < 10 < WorkQueueDepth > 15 > Bad
- Tune:
 - Add more cluster replicators
 - optimize cluster server usage (e.g. Split active users between cluster mates)



Top Statistics – Transactions

- Use for Indication of Server Load
- Check Server Stats
 - Server.Trans.PerMinute
- Interpretation:
 - Heavy < 30 < Trans.PerMinute (per User) > 10 > Light
- Tune: Analyze Heavy users and try to avoid load



Top Statistics – Concurrent Tasks

- Use to check Simultaneous Active Database Connections
- Check Server Stats
 - Server.ConcurrentTasks
 - Server.ConcurrentTasks.Waiting
- Interpretation
 - Waiting should be ZERO
- Tune:
 - Server_Pool_Tasks = n (e.g. 80)
 - Server_Max_Concurrent_Trans = m (e. g. Server_Pool_Tasks * Number of Ports)



Top Statistics – Platform Memory

- Used to check Allocated using memory pools and sub-allocations
- Check Server Stats
 - Mem.Allocated
 - Mem.Allocated.Process
 - Mem.Allocated.Shared
- Interpretation
 - Memory Leaks when increasing over days / weeks
- Tune
 - By several parameters (bufferpool, cache, namelookup...
- Note Be careful interpreting this statistic... Not all memory might be included



Top Statistics – Platform CPU

- Used to check CPU Utilization on Server
- Check Server Stats
 - Platform.System.PctCombinedCpuUtil
 - Platform.System.PctTotalPrivilegedCpuUtil
 - Platform.System.PctTotalUserCpuUtil
- Interpretation:
 - OK < 90% CombinedCpuUtil > 90% > TOO HIGH
- Tune
 - Many Root Causes Possible



Top Statistics – Paging File

- Use to check Server Memory Swapping to Disk
- Check Server Stats
 - Platform.PagingFile.Total.PctUtil
- Interpretation
 - OK < 0% < PctUtil.Avg > 10% > BAD
- Tune
 - OS Level tuning, Check Memory



Top Statistics – Platform Disk

- Used to check Disk Performance
- Check Server Stats
 - Platform.LogicalDisk.1.AvgQueueLen
 - Platform.LogicalDisk.1.PctUtil
- Interpretation
 - Good < 2% < AvgQueueLen > 5% > BAD
 - Good = PctUtil < 80%
- Tune
 - By several parameters (bufferpool, cache, namelookup), and OS / Disk Tuning
- Note
 - Platform.LogicalDisk.1.AssignedName=C points to the disk



Legal Disclaimer

© IBM Corporation 2011. 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.

IBM, the IBM logo, Lotus, Lotus Notes, Notes, Domino and Lotusphere are trademarks of International Business Machines Corporation in the United States, other countries, or both.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel is trademark or registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

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.

All references to Acme refer to a fictitious company and are used for illustration purposes only.