HCL Domino RunFaster=1

RNUG Forum, 10-11 October 2019, Moscow
Daniel Nashed (Nash!Com)

Nash!Com
Communication Systems

Rnug

About “Daniel Nashed”

- Nash!Com – HCL Business Partner from Germany
- Focus
  - Cross-Platform C-API, Domino® Infrastructure,
  - Administration, Integration, Performance, Security,
  - Troubleshooting and Traveler
- Author Domino on Linux®/UNIX® Start Script
- Co-Author Domino Docker Script
- Contact:
  - Mail: nsh@nashcom.de
  - Blog: http://blog.nashcom.de

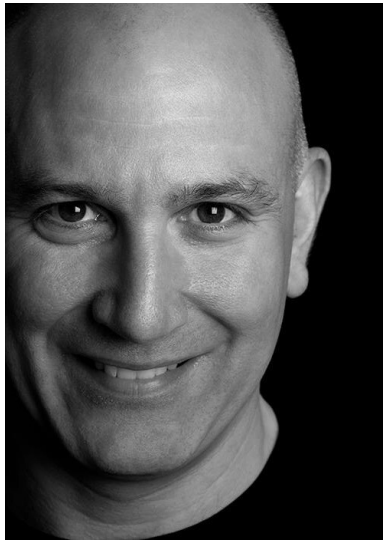

IBM CHAMPION

HCL Grandmaster

Nash!Com
Communication Systems

Rnug

Agenda

- Right Technology / Architecture
- Best Practices / Implementation

- Performance Troubleshooting

A Holistic Approach to Performance Management

- You need to take care about all layers
  - **Similar to security, the weakest component will define your performance!**
  - **Also how you run and maintain your environment matters!**

  - **Hardware**
    - Disk Drives, Memory, CPU
  - **Operating Systems**
    - Parameters, Subsystems, …
  - **Network**
  - **Virtualization Layer**

  - **Notes Infrastructure**
    - Server, Server settings, Client, Client settings
  - **Databases & Database Settings**
  - **Applications**
    - Use the right techniques designing & coding your applications

Classical Deployment

- Physical machines with single or partitioned servers used to be the standard deployment model
- Single Server usually cannot take full benefit of the full hardware resources
- Partitioned servers can better use the resources of a single machine with just one OS level instance
- Still often used with local disk or SAN storage
- Today most customers are moving to “virtual servers” for standalone and clustered servers

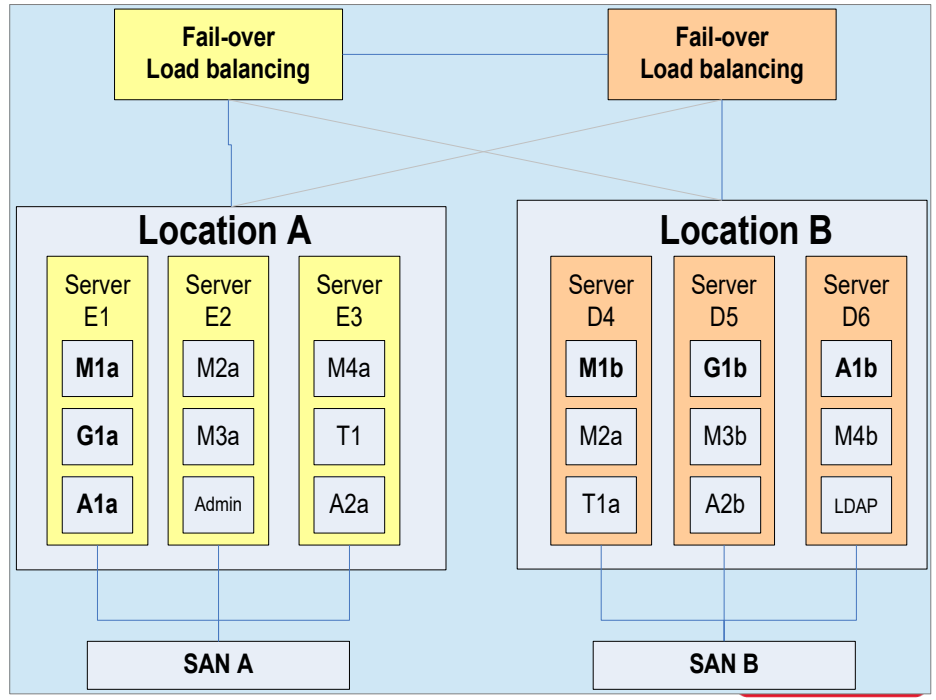

Nash!Com
Communication Systems

Benefits of Virtual Server Deployment Model

- Faster OS deployment and virtual hardware deployment
- Better use of hardware resources and more granular control of resources assigned
  - **Specially for smaller servers → No idle resources or server oversizing**
  - Load-balancing can move machines to other hosts, etc
  - Ease of adding resources like RAM and CPU
- Standardized machine layout and installation
  - Virtual hardware has the same drivers on OS level for every machine etc
- High availability by moving machines to different virtualization hosts
  - **That also means → Switching hardware without downtime**
  - **But you should still use Domino clustering**
- New backup options depending on used storage and configuration

REF - Domino 9.0.1 Virtualization Platform Support

- **VMware ESX and ESXi 4.0, 5.0, 5.1, 5.5 & 6.0 (including future Fix Packs) wherever VMWare ESX and ESXi are supported**
  - Most common used hypervisor for the Intel platform
  - Very stable large enterprise deployments with Domino and other IBM products
- KVM: KVM support is based on KVM/Kernel level. Support for level 2.6.32 x86_64 as baseline on either RHEL or SLES.
  - Automatically support patch releases beyond this level as well (but not feature releases).
- Microsoft Hyper-V Server 2008 R2,2012, 2012 R2 & new 2016

Best Practices Virtualization

- Don't overcommit CPU and RAM resources!
- Avoid high CPU ready time
  - It' the time needed the ESX host assigns CPU resources to a virtual machine
  - Good article: http://vmtoday.com/2010/08/high-cpu-ready-poor-performance/
    KB Article: https://kb.vmware.com/s/article/2002181
- Ensure you have sufficient I/O resources for the ESX host
  - Domino needs fast disks for NSF with a higher IOPS rate
- Don't allow RAM “ballooning”
  - Ensure you plan sufficient RAM resources for your Domino servers
  - RAM used for file-system cache dramatically reduces I/O read resources
- Network today should be 10 Gbit/second for your ESX hosts

Platform Focus / Architecture

- **Windows 64bit** and **RHEL Linux 64bit** are the two strategic platforms for Domino and other collaboration products

- Widely used by customers with physical and virtualized servers
- Mostly used **Domino 64bit** on Windows or Linux on the **ESX Server** Platform
  - We see far less customers on other virtualization products
  - ESX was the first “Intel” platform supported by Domino and is a “Tier 1” virtualization platform
  - Works best with **64bit** applications. E.g. Register optimization when application is native 64bit

Server Sizing

- Too large servers might not have best resource usage
  - Serve size does not only depend on number of **registered users**
  - But can also depend on
    - **Storage used**
    - **Concurrent users**
    - **Number of heavy users!**
  - Note: Large mail files **document count** and specially docs in **inbox** has impact on scalability!
- Usually a good size of a single Domino mail server instance is 1500 users
  - In a cluster environment 1500 per home-server and 3000 mailfiles per server (split between servers)

Network Considerations for Consolidated Environments

- Different Networks have different latency
  - LAN = around 1 ms
    - Take care: Wireless LAN connections can have different performance already!
  - WAN in same country depending on network around 40 ms
    - Today with current technology I have seen 6-8 ms latency!
  - WAN in Asia or other remote locations up to 400 ms
- Speed alone does not matter if you have many small transactions
  - Network Accelerators can help to some extend specially for re-transmissions, data flow optimization, ..
    - But don't expect miracles!
  - Notes client serializes transactions
    - Web client can request multiple resources at the same time

IBM SoftLayer Data Center & Network PoP an an Example

- https://www.ibm.com/cloud-computing/bluemix/data-centers
  - Datacenters spread around the world to allow low latency / Optimized between Data centers
  - Leveraging https://www.akamai.com technology to optimize

Nash!Com
Communication Systems

Rnug

Latency Example from the real world

- Avira Phantom VPN shows latency for different virtual locations
  - Tested from hotel in Switzerland
- Internet Latency in same county ~ 30 ms if you are lucky
- Latency is also important for cloud services

```
ping apps.na.collabserv.com

Pinging e10012.cd.akamaiedge.net [2.20.19.137] with 32 bytes of data:
Reply from 2.20.19.137: bytes=32 time=10ms TTL=58
Reply from 2.20.19.137: bytes=32 time=9ms TTL=58


ping outlook.office365.com

Pinging outlook.ms-acdc.office.com [40.101.88.194] with 32 bytes of data:
Reply from 40.101.88.194: bytes=32 time=29ms TTL=239
Reply from 40.101.88.194: bytes=32 time=31ms TTL=239
Reply from 40.101.88.194: bytes=32 time=65ms TTL=239
```

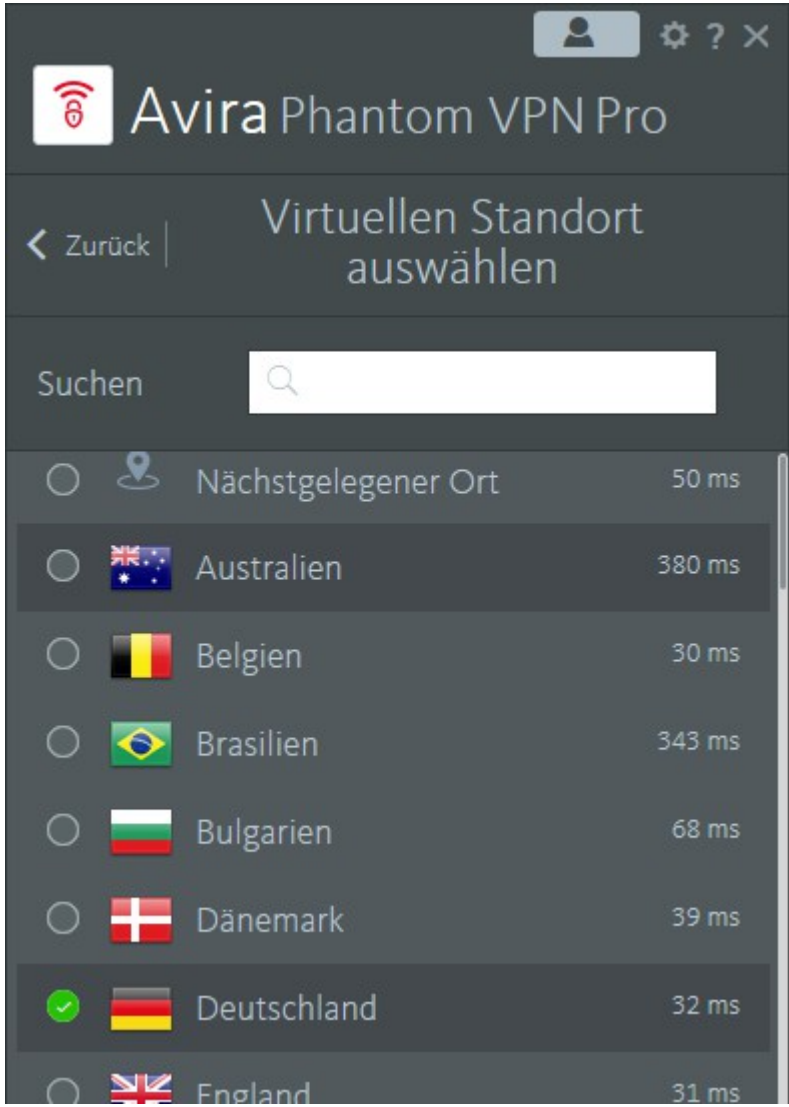

Nash!Com
Communication Systems

Rnug

Network Optimization Tips

- Network Compression
  - a.) Either Notes port compression
  - b.) Preferred: Compression on router level / Network appliance
    - Don't use both usually network level works best, if available
      - If you use Notes port encryption, enable Notes port compression as well because encrypted data cannot be compressed
- Take care about MTU size and auto MTU discovery
  - Specially when using VPN connections and mobile networks
  - MTU size needs to be aligned with smallest MTU of the connection
  - MTU Auto discovery might not always work, **Notes/Domino uses the full MTU size**
  - **You should check with your network administrator**

Applications and Latency

- For complex applications already 20 ms make it more difficult to work
  - Simple Example: 100 NoteOpen Transactions
    - in LAN = 100 ms
    - in WAN is already 2 seconds
    - Notes NRCP Operations are serialized!
- Multiple ways to optimize performance
  - Optimize Network → not always possible
  - Optimize Application
  - User different access method like **Web** or **Citrix**
  - Offload operations to the server (e.g. agent run on server)

Issue: Attachment Performance with Windows 2008 TCP/IP

- With 1 Gbit and 6 ms latency we got 2,5 MB/sec
  - After long research we found the reason
  - Default DefaultSendWindow is pretty small which has impact larger amount data send in chunks.
- Win2012 and higher does not have this issue and does not need tuning
- Workaround Registry Settings:
  - By default it was around 12 KB of data is send until the IP stack waits for an ACK.
    - After the change we had around **35 MB/sec** compared to **2,5 MB/sec**.
  - [HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\services\AFD\Parameters]

    "DefaultSendWindow"=dword:00080000

    "DefaultReceiveWindow"=dword:00080000
- Object Write Chunk-Size is 256 KB
  - There is a technote on attachment send chunk size which states that by default 64 KB is used.
  - But the mentioned Notes.ini SERVER_SEND_OBJECT_CHUNK_SIZE never made it in to the code and the current default is 256 KB

Windows 7 / Win2008 Network Tuning

- New Networking Stack in Windows 7 (also Visa) and Win2008 Server
- Some optimization can cause issues with Notes/Dominog
- Depending on your infrastructure you might need to change the autotuning level from “normal” to “restricted”
  - Set → **netsh interface tcp set global autotuninglevel=restricted**
  - Check → **netsh interface tcp show global**

```
netsh interface tcp show global

TCP Global Parameters
----------------------------------------------
Receive-Side Scaling State          : enabled
Chimney Offload State               : automatic
NetDMA State                        : enabled
Direct Cache Acess (DCA)            : disabled
Receive Window Auto-Tuning Level    : restricted
Add-On Congestion Control Provider  : none
ECN Capability                      : disabled
RFC 1323 Timestamps                 : disabled
```

Nash!Com
Communication Systems

Rnug

Linux Process Scheduler Optimization

- Doesn't impact all Linux versions but settings found do not impact newer Linux versions
- Timing issue with Domino server pool threads
  - We ran into it when moving mail-files via Adminp to a new server
  - Changed our performance from 25 MB/sec to 100 MB/sec attachment write performance
- Very specific Linux Tuning – Could change depending on Linux version used!
  - Reduce Kernel scheduling latency from 12 ms to 4 ms (value is in nano seconds)
  - **echo 4000000 > /proc/sys/kernel/sched_latency_ns**
  - OR you can permanently set it in **/etc/sysctl.conf**
  - **kernel.sched_latency_ns = 4000000**
  - Automatically set after boot or run once sysctl -p
- For more details check my blog post
  - http://blog.nashcom.de/nashcomblog.nsf/dx/domino-performance-issue-on-some-linux-versions.htm

Tips for Clustering

- Notes Clients are “cluster aware” and can silently failover to another cluster mate!
  - Still not enabled by default!
  - Client silent failover: notes.ini **FailoverSilent=1** or via Desktop Policy
    - Ensures client failover without prompting → also works when saving open documents!
- Usually you don't need to mirror your SAN data to another data-center when clustering on Domino level
  - Domino “cluster-mates” should be located on other ESX infrastructure – e.g. different data center in the same country
- Port Stealth Mode on Windows should be disabled!
  - Will cause delays failing over to a cluster partner – also kind of performance issue

REF - Port Stealth Mode Windows 2008 and higher

- Port-Stealth Mode is a firewall independent behavior to protect server ports
  - If an application does not respond, packets are dropped
  - Leads to long client wait times until the client fails over to another server
  - Can be disabled → Needs reboot after registry is changed!

- References
  - https://www.ibm.com/support/docview.wss?uid=swg21498755
  - http://msdn.microsoft.com/en-us/library/ff720058%28v=prot.10%29.aspx

- Registery Settings

[HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\WindowsFirewall\DomainProfile]

"DisableStealthMode"=dword:00000001

[HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\WindowsFirewall\PrivateProfile]

"DisableStealthMode"=dword:00000001 [HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\WindowsFirewall\PublicProfile]
"DisableStealthMode"=dword:00000001 [HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\WindowsFirewall\StandardProfile]
"DisableStealthMode"=dword:00000001

Additional RAM helps to reduce Read I/O

- A normal Domino itself usually does not need more than 4 GB RAM for all processes
  - Large Application Web servers and large Traveler servers need more RAM based in the application and load
- Additional RAM will be used as file-system cache on a 64bit OS!
- Dramatically reduces I/O read operations even with Domino random I/O pattern!
- Additional 16 – 64 GB RAM can makes sense – or even more!
  - Really depends the type of server and how much RAM you have left in your ESX server
  - Usually 64 GB RAM greatly improves performance already

Customer Example Linux 64bit on ESX with a lot of RAM

- Customer configured 1 TB RAM per Domino instance and had almost zero read I/O during the day
  - Customer business case: High speed for Adminp Processing for a full organization rename
  - Most Operations have been improved 20 times! → Usually great benefit already with 64 GB RAM

```
top - 16:49:40 up 11 days,  2:36,  1 user,  load average: 0.20, 0.27, 0.33
Tasks: 415 total,   1 running, 414 sleeping,   0 stopped,   0 zombie
Cpu(s):  0.7%us,  0.2%sy,  0.0%ni, 99.1%id,  0.0%wa,  0.0%hi,  0.1%si,  0.0%st
Mem:  1032574064k total, 1030662104k used,  1911960k free,  5423572k buffers
Swap:  2097144k total,      628k used,  2096516k free, 1015683544k cached

  PID USER      PR  NI  VIRT  RES  SHR S %CPU %MEM    TIME+  COMMAND
15777 notesp2   18   0 3025m 1.3g 1.2g S 21.9  0.1 984:23.84 server
18050 notesp2   18   0 2906m 1.0g 1.0g S  4.0  0.1 217:10.31 update
18049 notesp2   15   0 2952m 1.1g 1.0g S  2.0  0.1  62:09.21 router
```

Nash!Com
Communication Systems

Rnug

Large File System Cache for 64 Bit Windows

- By default on Win64 there is a very high physical memory limit for the file-system cache
  - It will try to use all memory for file-system cache which can cause Domino Memory to be swapped out
  - Therefore Domino uses Windows 64bit call “SetSystemFileCacheSize()” to limit the cache
    - See TN #1391477 for details
    - cacheset.exe -g will show the settings
- Default is 30% of memory
  - Can be tuned via notes.ini **MEM_FSCachePercentMem=n**
  - Set depending on the RAM in your machine – also required for native Domino 64
  - **Example: 16 GB RAM, 6 GB reserved for Domino/OS = MEM_FSCachePercentMem=65**
  - You can check the current settings with “cacheset.exe -g”
- Recommendation: 4-6 GB for OS & Domino itself, configure remaining memory as cache

Linux File-System Tuning

- Disable write of meta information via mount option -noatime
  - Meta information like last accessed time
- A real Runfaster=1 Parameter
  - Change the default scheduler from **CFQ** (complete fair queuing) to **NOOP**
  - CFQ tries to optimize disk access by reordering requests
  - But it would be better to send it to a SAN, RAID controller directly
  - Tests have shown that this works better for almost all SAN or local disk configurations
- Disable per device
  - echo noop > /sys/block/hda/queue/scheduler
- Disable globally via kernel boot parameter
  - Edit /boot/grub/grub.conf and enter in kernel line elevator=noop.

View Rebuild and Notes_TEMPDIR

- View Rebuilt will optimize view update
  - Notes.ini: **view_rebuild_dir=/dev/shm/view_rebuild**
  - Should have at least 20 GB free space
- You can move the Notes tempdir to a separate disk
  - Notes.ini: Notes_tempdir=/dev/shm/notes_temp
- On Windows best location is either separate file-system/virtual disk
  - Or system disk if you have sufficient space free!
- On Linux leverage “tempfs”
  - Works “similar” to a RAM disk but only needs RAM for used space
  - Enabled by default with 50% of your RAM, mounted by default in **/dev/shm**
  - Does swap to disk if memory can be better used by processes or file-system cache → very optimized
    - Ensure you have sufficient swap space → At least 50% of your physical memory
  - Tip: Nash!Com Start Script has options to create sub-directories automatically at server-startup

Disk Requirements in the virtual World

- In virtual environments we usually don't have much influence on disk configuration
  - But you can explain Domino requirements and usage patterns to help optimize
- In earlier days virtual servers with large storage used direct SAN storage via iSCSI etc.
  - Today most customers use VMDKs with larger RAID10 NAS/SAN Volumes
- It's still a good idea to split data into virtual disks similar to physical machine installations
  - They might be on same RAID group but high write concurrency in the same VMDK causes contention!
  - Also helps troubleshooting performance issues ;-)
- Recommendation
  - Translog should be on separate VMDK or on system VMDK
  - NSF Data should be on separated VMDK
  - DAOS should be on separate VMDK
  - FT and NIF (new in FP8) should be on separate VMDK

REF - File System Recommendation “TRANSLOG”

- Size: 4.1 GB for Circular Translog
  - For archive-style translog size highly depends on your data
  - For linear style translog size should be in the range of 8-20 GB
- 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

REF - File System Recommendation “NOTESDATA”

- Recommended: Fastest Disk – Usually RAID 10 instead of RAID 5
- 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 I/O traffic, as well as filling OS/SAN etc cache buffer with unwanted blocks
    - You need read ahead caching for backup and some read ahead also improves Domino performance

- Faster drives perform disk I/O in less time
  - Choose drives with a combination of low seek time & high RPM

REF - File System Recommendation “DAOS”

- Recommended: RAID5 to save costs (if available)
- 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
  - In most cases more read than write cache makes sense
- Entry Level Storage is sufficient from performance point of view
  - Lower I/O rate and specially lower IOPs than what we see for NSF

REF - File System Recommendation: “FT” and “NIF”

- Same recommendations like “NSF”
- Similar access pattern to “NSF”
- Should be a separate file-system / disk
  - Depending on size you can have one file-system with different directories
    - For larger servers separate virtual disks can make sense

Fulltext Index on separate Disk

- Note new – available since 8.5.3 full text index can be move to separate directory / disk
  - Specially useful to separate storage in larger environments and to distribute I/O
  - Also helpful when using snapshot backup
- Fulltext index can cause disk fragmentation
  - **FT Index should be rebuild regularly (e.g. once per month) → See DBMT slides**
  - Separating FT index reduces disk fragmentation on “NSF” disk
- Notes.ini **FTBasePath=d:\full_text**
  - Restart Server after moving FT Index
  - Rebuild of FT index will move the FT index data

Best Practices: FT Index

- UPDATE_FULLTEXT_THREAD=1
  - Separate Thread for Full-Text indexing
- FT_FLY_INDEX_OFF=1
  - Disables on the fly FT indexing when agents use search queries on a not FT indexed DB
  - Avoids “extremely inefficient” temporary FT index → Agent will not run and code should be changed
- FTG_USE_SYS_MEMORY=1
  - Use system memory instead of Domino pooled memory
- Avoid using “Immediate FT index” if possible
  Don't index attachments if possible!
  - If you have to use exclude lists for attachment types (Example below)
    - Notes.ini FT_INDEX_IGNORE=*.asf,*.avi,*.bin,*.bmp,*.CSS,*.dat,*.DXL,*.EOT,*.EPS, *.FLV,*.EPS,*.HTC,*.ICS,*.iso,*.JAR,*.JS,*.mpeg,*.MP4,*.PNG,*.rar,*.SVG,*.SWF,*.TIF, *.TTF,*.WMV,*.XML,*.XSL,*.zip

Virus Scanners & NSF Files

- You should disable on access virus scanning for all Notes related files!
  - Huge impact on performance and stability!
  - Use Domino aware virus scanners instead!
    - Scan attachments in mail.box and have at least scheduled databases scans!
- Either exclude directories if possible or extensions!
  - Important extensions: *.NSF, *.NTF *.BOX, *.DTF, *.ID, *.FT
  - Translog: *.TXN
  - DAOS: *.NLO
  - New extension: *.NDX for view index!
  - Extension used for copy style compacts: *.ORIG, *.TMP

Database Optimization – NSF Storage

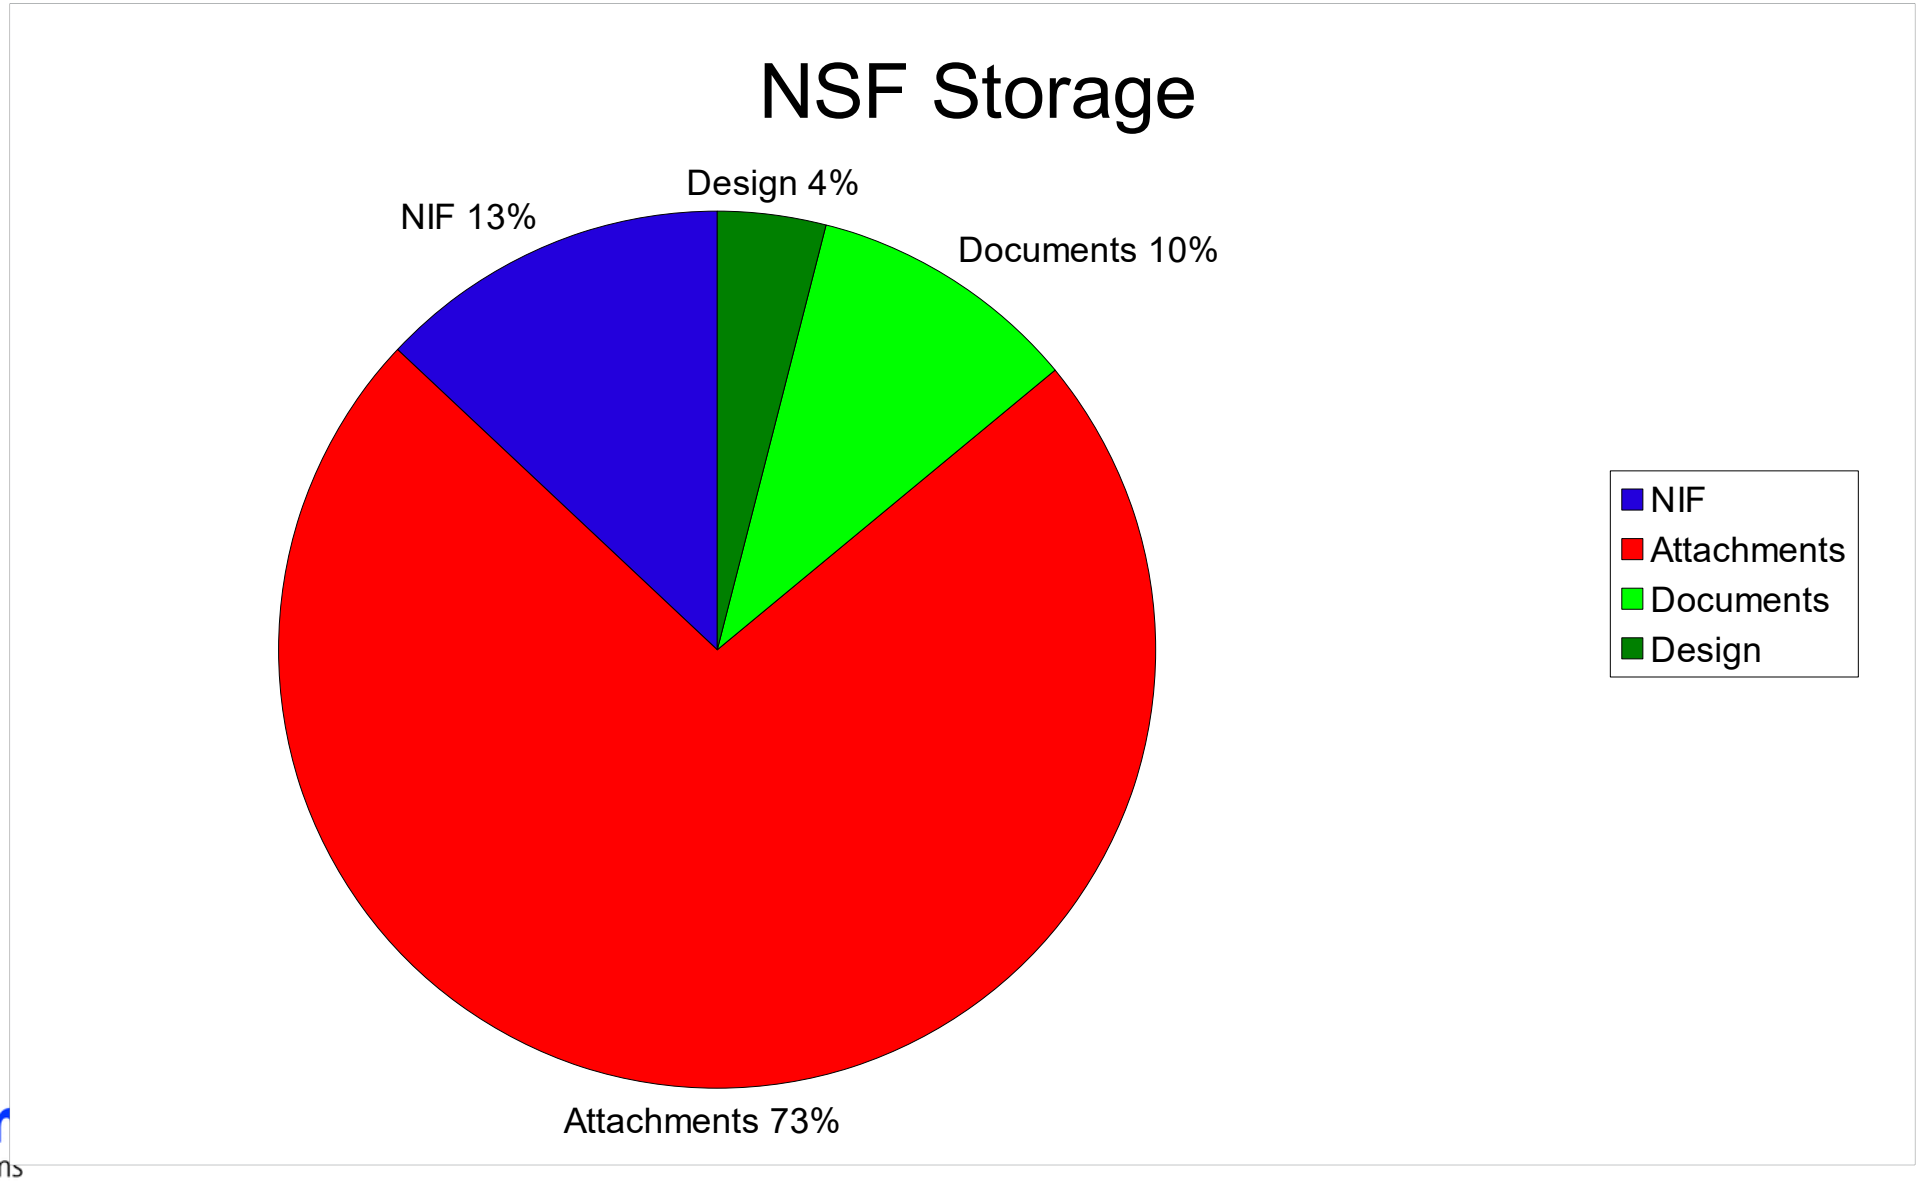

Database Optimization

- Database Design Compression
  - **50%** of Design → only **2%** total reduction, but it reduces the I/O!
- Document Compression
  - **50%** of Summary data → **5%** total and reduction of I/O
- Tip: Enable via compact, recover space with DBMT
  - Compact -c -n -v mail
  - In case you are using DBMT use without -c to set the options and have DBMT compact and compress data/design
- DAOS
  - **70%** NSF File size reduction, **25-35%** over all data reduction by deduplication
  - Also helps to improve performance at runtime and maintenance operations

DAOS Example Results

- 70% reduction of NSF
- 25-35% total reduction by deduplication

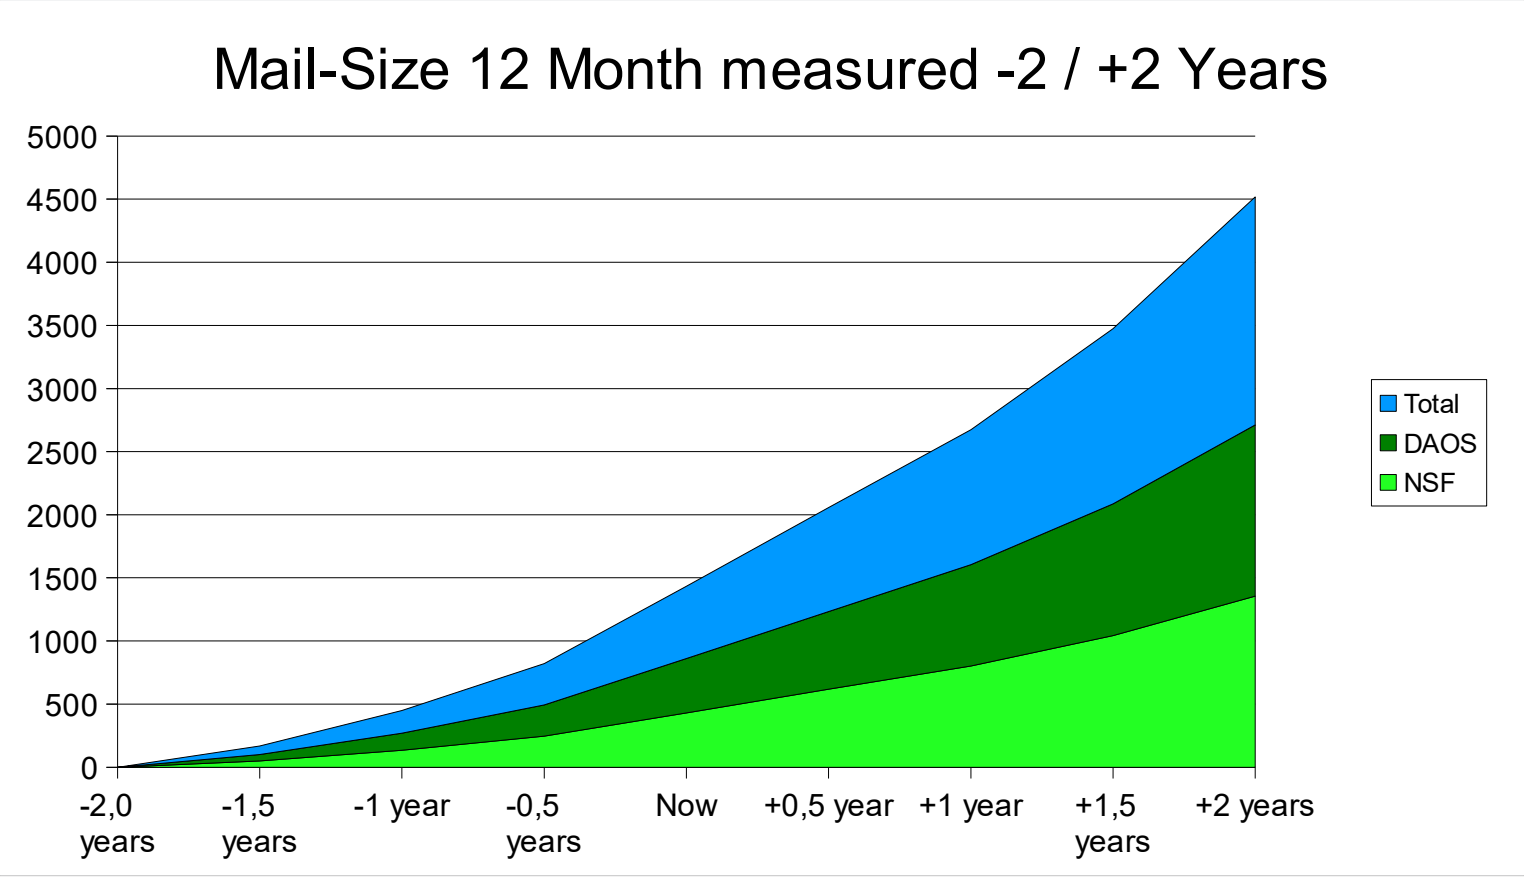

Nash!Com
Communication Systems

Rnug

DAOS Storage & Deduplication

- DAOS is supported on NAS storage even leveraging mounts

- Some NAS storage like NetApp supports storage deduplication in same storage container
  - Deduplication usually works on block level (e.g. NetApp uses 4 KB blocks)

  - If you have multiple mail-servers or archiver servers deduplication can reduce storage dramatically!
  - Only works if disabling NLO encryption → notes.ini **DAOS_ENCRYPT_NLO=0**

  - Also reduce the number of NLOs per sub-folders (default : 40000)
    - Notes.ini **DAOS_MAX_FILES_PER_SUBCONTAINER=20000** (or less in smaller environments)

- Tip: Some backup solutions like Tivoli support deduplication if configured
  - Speak with your backup team how to backup DAOS in the most efficient way!

“NIFNSF” – Move NIF outside NSF

- New since 9.0.1 Feature Pack 8
  - You should wait for FP10 IF2 to use it because of an important stability and performance fix
  - Needs a separate file-handle/cache entry. Might need to increase DbCache Entries
  - **NSF_DbCache_Maxentries=4000**
- Moving NIF (View Index data) outside the NSF
  - Reduces backup data and optimizes disk I/O and database locking
  - More parallel NSF and NIF operation possible
  - Allows further split of data into different file-systems
- NIF is stored in one separate “.NDX” file per database
  - New Notes.ini parameter can be used to move NDX files to separate directories/disks
    New Compact option to enable/disable NIFNSF per database
- Mail File NIF size: ~ 10% of the database size
  - With DAOS ~ 30% of the remaining physical file

Optimize Daily / Weekly Operations

- Newer Tools like “**DBMT**” allow optimized operations
- You should review your current strategy and operations!
- Often maintenance, backup and archiving are not aligned
  - Depending on your configuration this can make a big difference
- DBMT can have shorter runtime and has flexible and new options
  - Some ideas, tips, .. on the next slides with reference slides for more information

Database Management Tool - DBMT

- Multi-threaded Domino Server Add-In Task
- Performs automated coordinated administrative tasks on databases
- DBMT replaces the need for running Compact on Non-System databases
- Standard Maintenance
  - Copy-style Compacting (compact) → Only copy-style compact!!
  - Fixup of Corruption (fixup)
  - Full Text Indexing (updall)
  - Incremental View Indexing (update)
- Additional Maintenance
  - Folder Reorganizing
  - Deletion Stub Purging
  - Expiring of Soft Deleted Entries
  - Updating of Unread Lists

REF -DBMT Example

- dbmt -compactThreads 4 -updallThreads 4 -ftiThreads 4 -ftiNdays 30 -range 1:00AM 6:00AM -compactNdays 7 -force 1
  - Compact all databases not compacted for 7 days with 4 threads
  - Update all views with 4 threads
  - **Update FT Index with 4 threads, Rebuild if corrupt or older than 30 days**
  - Run between 1 AM and 6 AM every night and stay in memory during the day
  - Fixup database which did not compact successfully 5 times on Sunday

Additional Information

- See additional slides at the end of the presentation for more details about DBMT

D9.0.1 - DBMT – File Space Pre-Allocation

- A copy style compact generates a temp file and copies the database note by note into the database
  - Allocation of file-space is also done step by step
  - Depending on the file-system, free space database is already fragmented when the compact finishes
    - It is very likely that database has many fragments
- DBMT will pre-allocate the calculated space that the physical database needs
  - It's DAOS aware
  - Enabled by default in 9.0.1
  - Notes.ini **DBMT_PREFORMAT_PERCENT=n** allows you to increase or decrease the size
  - For example DBMT_PREFORMAT_PERCENT=120 will have 20% contiguous free space in the database after compact (if there is sufficient unfragmented space in the file-system)

Best Practices: Performance Tuning for larger Servers

- NSF_BUFFER_POOL_SIZE_MB
  - By default 1 GB for 64bit Domino → could be increased to e.g. 2048 MB if you have a lot of RAM
- NSF_DbCache_Maxentries=3000
  - Default value around 3 times the buffer pool size in MB
- EVENT_POOL_SIZE=41943040
  - Needed on all servers if higher number of log messages
- Increase internal pools for larger servers
  - CATALOG_POOL_SIZE_MB=100
  - dirman_poolsize_mb=100
  - nsf_monitor_pool_size_mb=200
  - MAX_NETPOOLSIZE to 1Gig (from 512 MB)
  - MAX_TASK_POOL _SIZE to 16 MB (from 4 MB)
  - NETSESSIONPOOL_SIZEMAXMB to 20000 MB (from 2000 MB)
  - MAX_GROUPCACHE_POOLSIZE to 32 MB (from 15)

Keep your Software up to date!

- Newer OS Versions and Domino versions usually have better performance!

- Domino Examples
  - Better I/O optimization
  - Internal optimization of transactions
  - Better database concurrent access
- Windows, Linux
  - Optimized network performance
  - Optimized process scheduling

- So it not only makes sense from security point of view to keep your servers and clients updated!

General Application Performance Tips

- Think about the amount of documents you might have per month/year!
  - How often are they created/updated/deleted?
  - Split databases and thin about document life-cycle including archive delete from day one! Complex databases with lot of changes should be below 100.000 documents!
  - Bigger more static databases can be much bigger
- A view is not always the best way to find data
  - Take a look at the NotesNoteCollection
    - Very fast und flexible way to search data!
- Avoid to complex views and split data in different views
  - Avoid too many reader fields and a single view
- Leverage @formula language!
- Disable unread marks if needed (system-databases, application that does not require it)
  - Can have larger impact special on Web applications invoking agents

REF - Some Application Pitfalls for WAN Connections

- Avoid Private on First Use Views/Folders
  - If you have more than 40 design elements with the same name the internal design cache breaks.
  - See Technote for details → #1322578, SPR #RSTN7K2EM4
    - Email me for a tool to test if the design cache is broken for a database
  - When the design cache is broken fall-back code is used to find design elements
  - Instead of **FINDDESIGN_NOTES** transactions the design collection is opened and searched sequentially until a match is found → This overhead breaks your performance with WAN connections!
- ComputeWithForm" does not use cache.ndk
  - It even does not leverage the newer finddesign transaction but does a lookup in the "**design collection**"
  - So every ComputeWithForm does read all forms, sub-forms, shared fields which are involved …
  - A lot of overhead. Maybe OK for a single ComputeWithForm but refreshing multiple documents can be painful
- **Idea/Workaround**: Use Run on Server Agents off-loading the work to the server
  - This can be a good approach for other operations that need many transactions

REF - Notes Client

- Data Directory should be always **local** not on a file-server
  - File-Server configurations are only supported for Citrix clients!
- Most important is cache.ndk which is used as a local cache for design elementes
  - Should not be deleted on client start!
- Ensure Client ODS is updated to the latest ODS!
  - Leverage policy settings or notes.ini settings directly
    Client will upgrade to the latest ODS automatically!
- Antivirus exclusions are important!
  - Similar settings than server (see previous slides)
- Avoid on access scan for Notes client binaries – specially for the thousands Eclipse files!
  - Instead use write scan and regular scheduled checks

Performance Troubleshooting

Performance Troubleshooting

- Can be quite time consuming and needs a structured approach
- Best would be to implement performance measurement directly when building a larger, complex application
- Embedded performance measurement into your application
  - Tip add thresholds which can be configured in production to report potential performance issues
  - Make performance measurement part of your application
- Performance is not just a Notes/Domino team effort
  - The bottleneck can be a combination of things including timing
  - Can be hardware, OS, Domino, Storage, Network, Application ..
- You need to understand first which native performance measurement we have available on all different layers

REF - Millisecond Timer for Applications

- Getthreadinfo(LSI_THREAD_TICKS)
  - LSI_THREAD_TICKS = 6
  - LSI_THREAD_TICKS_PER_SEC = 7 returns ticks per second (1000)
  - tick = millisecond (different than agent ticks!)

    Note: Can wrap around.

- Code Example
  - Dim lStart,lStop As Long

    lStart = Getthreadinfo(6)

    ... do some work ...

    lStop = Getthreadinfo(6)

    ms = lStop – lStart

Domino Platform Statistics

- Domino Platform Statistics can help
  - Simple long term monitoring in statrep.nsf
  - No so helpful for detailed troubleshooting because just updated once per minute from OS
  - Not all information like “await” on Linux available
- For detailed troubleshooting use OS level tools
  - Windows: Perfmon – “Average Disk Queue Len” < 1
  - TIP: New since Win2008: Resource Monitor shows I/O rates and response time! (can be live sorted)
  - Linux: iostat -x – “await” < 10 ms
- Long term OS monitoring can help to understand trends, differences during day /week
  - Also helps to identify issues occurring at certain times

The Perfomance Measurements You already have

- Backup Runtime
  - Useful to understand read performance, network performance
  - Depends on your backup infrastructure – you should at least have 60 MB/sec on a Gbit connection
- Compact Runtime
  - Gives an idea about Domino server performance
  - Idea: You can use compact -# to invoke multiple compacts to generate some load
- Catalog Runtime
  - Also gives an idea how good your server is performing
- Agent Runtime
  - See next slide
- If you measure all different runtimes and keep track, you notice performance improvements or slowdowns!

Agent Manager Runtime Analysis

- Increase log level - need for debug output
  - set config l**og_agentmanager=2**
  - Enable Debugging agent runtime analysis
    - **tell amgr debug p**
    - or set config debug_amgr=p and restart amgr servertask
  - Disable via "**tell amgr debug off**"
- Log Output of Amgr:
  - **11.11.2017 11:11:00 AMgr: Agent 'import data' in 'crm.nsf' ran successfully; elapsed time was '591' ticks**
    - 1 tick = 1/100 second
  - Extract and analyze this information into a log analysis DB or excel sheet
- Domino 6 & higher does not use a database for log analysis anymore
  - Search Results are now stored as single documents in log.nsf

Disk Statistics Windows/Linux

- Windows
  - AvgQueueLen should be below 1

- Linux
  - „Service Time“ just an indicator. “await” would be the better performance counter
  - Should be below 10 ms

```
Platform.LogicalDisk.1.AssignedName = HarddiskVolume1
Platform.LogicalDisk.1.AvgQueueLen = 0,2
Platform.LogicalDisk.1.AvgQueueLen.Avg = 0,8
Platform.LogicalDisk.1.BytesReadPerSec = 0
Platform.LogicalDisk.1.BytesWrittenPerSec = 0
Platform.LogicalDisk.1.PctUtil = 11,00
Platform.LogicalDisk.1.PctUtil.Avg = 50,00
Platform.LogicalDisk.1.ReadsPerSec = 0
Platform.LogicalDisk.1.WritesPerSec = 0
```


```
Platform.LogicalDisk.1.AssignedName = HarddiskVolume1
Platform.LogicalDisk.1.AvgQueueLen = 0,2
Platform.LogicalDisk.1.AvgQueueLen.Avg = 0,8
Platform.LogicalDisk.1.BytesReadPerSec = 0
Platform.LogicalDisk.1.BytesWrittenPerSec = 0
Platform.LogicalDisk.1.PctUtil = 11,00
Platform.LogicalDisk.1.PctUtil.Avg = 50,00
Platform.LogicalDisk.1.ReadsPerSec = 0
Platform.LogicalDisk.1.WritesPerSec = 0
```

“vmstat” provides detailed Linux Performance Information

- r = The number of processes waiting for run time.
- b = The number of processes in uninterruptible sleep e.g. waiting for I/O to complete!
- cs = Context switches – how often the CPU switches between threads
  - Should be below 10000 per Domino partition
- us/sy/id = User/System/Idle CPU
- wa = CPU time spend on I/O (should be below 30%)

```
 vmstat 1
procs -----------memory---------- ---swap-- -----io---- --system-- -----cpu-----
 r  b   swpd   free   buff  cache   si   so    bi    bo   in   cs us sy id wa st
 2  1      0 6831004      0 228364    0    0    12     9    0   100  10  0 80  10  0
 3  2      0 6831000      0 228364    0    0     0     0    0  6210   9  0 80  11 0
 5  2      0 6830964      0 228400    0    0    12    40    0  6820   5  0  80   15  0
```

Nash!Com
Communication Systems

Rnug

Linux I/O Stats “iostat”

- Linux iostat is your friend!
- TIP: Filter via grep – z.B. “iostat -x 2 |grep sda”

- %util =Disk Utilisation in % → 90% and higher are a hint for a bottleneck
- r/s = I/O Read per second
- w/s =I/O Write per second
- svctm = Service of disk itself (or virtual disk)
- await = Request time in total (application → OS queue → disk → and back)
  - Should be below 10 ms!

ESX Server Statistics

- Statistics per ESX Guest
- Can help to understand bottlenecks as well
- Layer between the actual hardware and the guest OS can have impact
- Take care of “CPU Ready” statistic mentioned earlier
  - Only place to see the statistic is guest statistics
- Different tools available but usually only have access to the graphical tools
  - There is also esxtop and scripting/add-on tools – That's more for the ESX server team

Network Performance Check

- Check latency first with “Ping”, “Traceroute”
  - Provides basic feedback about the application latency
- But we finally need application level latency information
  - C-API call NSPingServer(..) might not be a good choice
  - Better have a script which opens documents to benchmark client/server latency
    - I wrote a simple application which can be used by an end-user to do a quick test
  - Also test how long it takes to transfer larger amount of data – for example attachments
    - E.g. if you have central Traveler servers syncing data for users with decentralized servers.
  - If someone is interested, drop me a note. I can share some of my test tools.

Trace Notes Client Transactions

- Client_Clock=1
  - Can be used to track Notes Client/Server Transactions (NRPC)
  - Notes,ini debug_outfile=c:\debug_notes.log

- Client_Clock logs
  - Transaction sequence
  - Transaction name
  - Transaction data (ReplicaID, NoteID)
  - Response time (ms)
  - Bytes send, received

- Example:
  - (15-78 [15]) OPEN_NOTE(REPC1256B16:0072BCBE-NT00000E3E,00400020): 0 ms. [52+1454=1506]

- Free annotation tool OpenNTF Client_Clock Parser (Thanks Andrew ;-)
  - http://www.openntf.org/Projects/pmt.nsf/ProjectLookup/Notes%20RPC%20Parser

REF – Some important NRPC Transactions 1/2

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

REF – Some important NRPC Transactions 2/2

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 – Trace Transactions on Server Side

- "**Show trans**" shows transactions categorized
  - Tip: "**Show trans reset**" resets transaction counters
  - Can be used to get an idea about transactions and transaction types
- "**Server_Clock=2**" dumps all transactions to server console/ console.log
  - Nicely delimited output, that can be parsed and put into Excel or Notes database
  - Example:
    01.01.2018 12:37:57 39789504 ms '**OPEN_NOTE**' **1 ms** (0 ms NETIO) TCPIP 00010EDC **Rcvd 0** Sent **3026** User '**Full Admin/NashComLab**' Db '**mail\fadmin.nsf**' Ip '**192.168.100.101**' Lm 1

Server Hang Symptoms

- Server (or specific task) is still running, but client receives error messages "Server not Responding"
  - No error is produced on the console but an error may be written to log.nsf
- Console does not accept keyboard commands
- Servertask will not shutdown cleanly
- User report that other Domino server tasks have slowed down
- No NSD is generated and no Fault Recovery

How to troubleshoot Server Hangs?

- Check call-stacks for specific calls
  - e.g. a large number Semaphore Calls, SpinLock Calls
- Use Semaphore Debugging
  - DEBUG_SHOW_TIMEOUT=1
  - DEBUG_CAPTURE_TIMEOUT=10
  - DEBUG_THREADID=1 (default meanwhile)
  - Optional: DEBUG_SEM_TIMEOUT=X
    (in milliseconds, default 30000)
  - „Show stat Sem.Timeouts“ to check semaphores
- Run 3 nsd -nomemcheck in short sequence
  - plus one full NSD

Analyzing Semaphore logs

- Semdebug.txt in IBM_TECHNICAL_SUPPORT
  - Example line:
    ti="0025CA9C-C1257353" sq="00004CE8" THREAD [28208:00241-169659312] WAITING FOR SEM 0x0931 Task sync semaphore (@0F7711A4) (OWNER=28208:158743472) FOR 5000 ms
  -
  - Contains semaphores locked for more than e. g. 30 seconds
  - Information about process/thread, semaphore, time, ...
  - Also contains information who is currently holding the semaphore
    - But just the process/thread.id – You have to annotate on your own via NSD
    - Find the call-stack of the process requesting and olding the semaphoere
      - Can only be done thru NSD

Annotate Semaphore Logs

- Example:
  - ti="0025CA9C-C1257353" sq="00004CE8"
  - THREAD [28208:00241-169659312]
  - WAITING FOR SEM 0x0931 Task sync semaphore
  - (@0F7711A4) (OWNER=28208:158743472) FOR 5000 ms
- ti is the internal repesentation of the timedate
  - You can use LND to annotate the ti values
- Lotus Notes Diagnostics (LND) can be used to annotate semapore.debug files
- Current Version: LND 2.9
  - http://www.ibm.com/support/docview.wss?uid=swg24019151

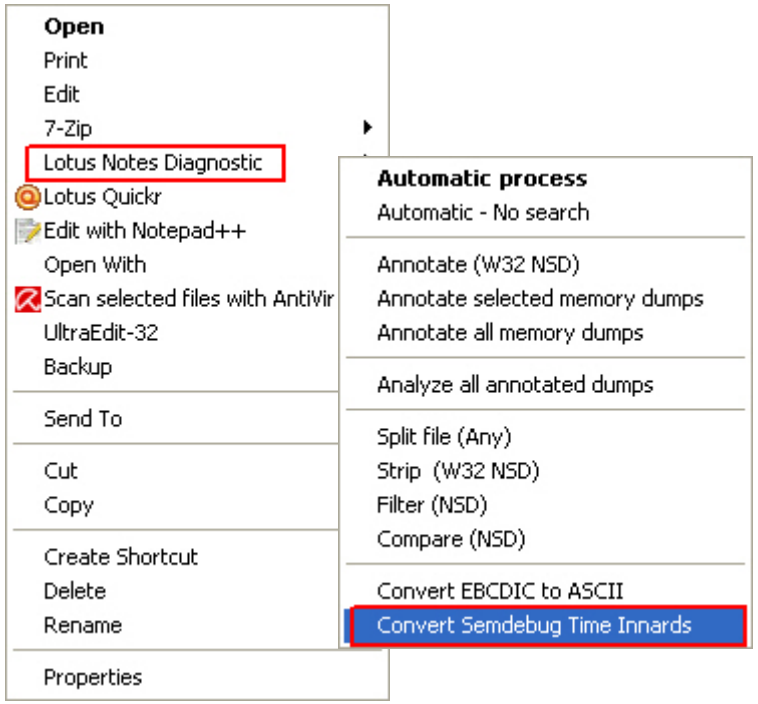

Long Held Locks - LkMgr

- LkMgr is used to lock databases for concurrent access
  - Is used on transaction logged databases
  - Locks can block other threads to continue with their work
  - It's OK when this happens for short time
  - But it could be a dead-lock (when different threads depend on resources that the other needs)

- How to figure out what is going on
  - Check how long a lock is waiting and match against lock “granted”
    - Req(**Status=Waiting** Mode=S Class=Manual Nest=0 Cnt=0.. Tran=0 Func=N/A dbopen.c:4279 [45854:00002-00000007740885l2(00007FF92E23A740)] **Delay=2min**)
    - Req(**Status=Granted** Mode=SIX Class=Manual Nest=0 Cnt=1 Tran=0 Func=N/A nsfsem1.c:2245 [45868:00031-00000019937707 52(00007F7776D68700)])

Thank you

- Thanks for your attention!
- Find me during the conference or contact me later
  - Mail: nsh@nashcom.de
  - Blog: http://blog.nashcom.de

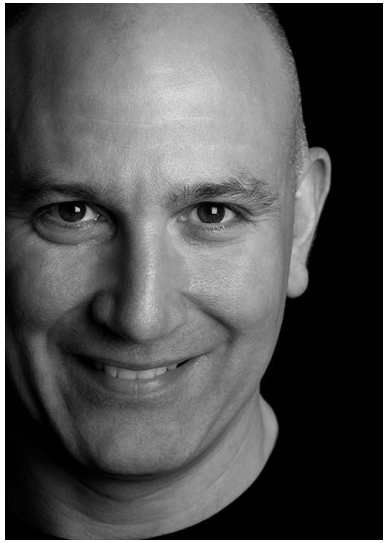

IBM CHAMPION

HCL Grandmaster

Nash!Com
Communication Systems

Rnug

Additional Material

DBMT Limitations & Tips

- For Compact Options you cannot specify how much Space should be free for compact to run (-S parameter)
  - But you can configure how often a database should be compacted ()
- Compact only runs in copy-style compact
  - Works fine in a cluster
  - Else you need to make sure compact only runs on weekends/ middle of the night
- You should only have one DMBT with -range Option
  - No second DMBT for the weekend with a bigger time window
  - Means you should keep compact operations separate and run them during the day on weekends
- In case you have archive style transaction logging you need to schedule compacts for the weekend anyway
- Ensure you have the right quota setting in the Translog tab:
  - “Check space used in file when adding a note”
- Command Line Options in Program Docs is limited to 128 chars!

Mail File Disable Compact Abort

- Since Domino 8.5.3FP3
- By default, compact operation will be aborted when mail is delivered
  - To make DBMT more efficient set notes.ini **MailFileDisableCompactAbort=1**
  - New mail in mail.box will sit undelivered in mail.box on the Domino Server
  - Mail router will periodically retry delivering new until the current compact completes
- Especially for large mail files, this may be an unacceptable extended period of time to go without receiving new email
  - But there is another new option in Domino 9.0 (see next slide)
- Best practice is to set
  - MailFileDisableCompactAbort=1

Mail File Enable Delivery Failover

- Since Domino 9.0,mail Router has ability to failover mail delivery to an Available mail file replica on a cluster mate when the local replica is unavailable because:
  - Copy-style compact of mail file is in progress
  - Fixup of mail file is in progress (Performing consistency check on...)
  - Mail file is missing (File does not exist)
  - To enable delivery failover, set **MailFileEnableDeliveryFailover=1** in notes.ini on Domino Server
  - **DBMT** uses the replica index in cldbdir.nsf to determine if it should compact the mail file on this Cluster Mate on this day to make sure all Cluster Mates are not compacting same replica on same day
    - **Only works in combination with -compactNdays**
  - **$MailClusterFailover** item will be appended to the note with a value of the Domino Server name of the server where the note was actually delivered on
  - Note: For Mail-Rules you need to allow your servers to use “monitors” (Security tab)

DBMT command line options

- **-compactThreads <n>**
  - Use <n> threads for compact where **n** is between **0 & 100** inclusive
  - Number of threads DB copy style compact operations. **Default is 1 thread**.
    If **0** is specified, no compact operations are performed.
- **-updallThreads <m>**
  - Use <m> threads for updall, where m is between **0 & 100** inclusive
  - Number of threads doing updall operations. Default is **1**.
- **-ftiThreads <g>**
  - Use <g> threads for the rebuild of the full text indices, where g is between **1 & 100** inclusive
  - Configures number of threads for rebuild of full text indices if corrupt or **-ftiNDays** specified. Default is **1**; 0 is not allowed.
- **-ftiNdays <f>**
  - **Rebuild** full text indices every <f> days from initial creation
  - Default is to rebuild full text indexes only when found corrupt.

DBMT Schedule Options

- **-range <s> <e>**
  - only run between times <s> and <e> daily
  - Minimum of 10 minutes between time <s> and time <e>
  - Times must be specified in 12hour format with a required AM/PM (e.g. 11:50PM)
  - DBMT process will remain loaded but dormant until -range comes around again the next day

- **-timeLimit <q>**
  - only run for <q> minutes
  - DBMT process will exit once **-timeLimit** expires and all active DBMT processing completes
  - -stoptime <e>
  - only run between now and time <e> once
  - Time must be specified in 12hour format with a required AM/PM (e.g. 6:00AM)
  - DBMT process will exit once -stoptime is reached and all active DBMT processing completes

DBMT Compact Options

- **-compactNdays <p>**
  - Compact databases that have not been compacted in last <p> days
- **-noCompactLimit**
  - Allow compact to run past the end time to finish, but no new compacts will be started
- **-force <n>**
  - Take DBs offline for fixup & compact if flagged for force
    - Fixup is run only when 5 or more consecutive compact operations fail (and the failure is not due to database in use).
  - <n> between 0 and 7 where 1= Sunday, 2 = Monday, and so on.
    - If n is 0 (zero), the fixup operation will run any day.

DBMT for System Databases

- DBMT does not compact System Dbs by default
  - but will perform -updall and -fti operations
  - You can explicitly specify databases on command-line or “ind” file

- System-Database List
  - names.nsf
  - events4.nsf , admin4.nsf, catalog.nsf, ddm.nsf
  - log.nsf, statrep.nsf, domlog.nsf, lndfr.nsf
  - busytime.nsf, cldbdir.nsf, clubusy.nsf
  - daoscat.nsf, dbdirman.nsf, dircat.nsf
  - mtdata\mtstore.nsf

DBMT Compact Filter indirect file

- dbmt_compact_filter.ind
  - Exclude list for Databases
  - DBMT will not compact database(s) **< ODS41** – Database(s) automatically added to dbmt_compact_filter.ind by DBMT
  - If -range is specified and **-noCompactLimit** is not specified
    - If DBMT compact of database takes longer than ~full -range specified and terminated because end of range reached , database automatically added to dbmt_compact_filter.ind by DBMT
    - Prevents DBMT from churning on same database every night
  - Administrators may manually add additional database(s) to **dbmt_compact_filter.ind** if e.g. they have an Application/database that needs to be Highly Available all the time

DBMT -updallThreads

- By default, specific views in mail databases are
  - brought up-to-date
  - **Created**
  - **marked non-discardable**
  - dramatically improving end user's fail over experience in a cluster

- Need Inheritance Enabled and Template Name one of StdR7Mail,StdR8Mail,StdR85Mail or StdR9Mail
  - ($Inbox), ($Drafts), ($Sent), ($All)
  - ($RepeatLookup), ($ToDo), ($Calendar), (Haiku_TOC)
  - ($Alarms), (iNotes), ($Users), (iNotes_Contacts), ($ThreadsEmbed)

- If you use customized mail template(s) use the following notes.ini
  - DBMT_MailTemplate=templatename1,templatename2,templatename3

DBMT -updallThreads

- Custom Database View Configuration
  - Required Inheritance Enabled and Template Name set on database(s)
  - DBMT_TemplateName=ViewNameOrAlias1;ViewNameOrAlias2;...ViewNameOrAliasN
  - Substitute TemplateName after the DBMT_ with inherited database Template Name
  - Substitute the ViewNameOrAlias1-N with database View Name or Alias and separate the ViewNameOrAlias1-N with either semicolons or commas

- Examples in notes.ini
  - DBMT_StdNotesLog=SecurityEvents, ReplicationEvents, MailRoutingEvents
  - DBMT_StdR4AdminRequests=All Requests by Time Initiated;All Request by Server

- Remember notes.ini setting(s) are restricted to 128 characters!

DBMT Statistics

- Stastistics for Compact and Updall when DMBT terminates
  - DBMT.Compact.Began 04.06.2016 21:27:44
  - DBMT.Compact.Finished 04.06.2016 21:27:59
  - DBMT.Compact.Successful 26
  - DBMT.Compact.Unsuccessful 1
  - DBMT.Compact.Unsuccessful.InUse 1
  - DBMT.Compact.Unsuccessful.TimeLimit 0
  - DBMT.Compact.BackLog 0
  - DBMT.Compact.Fixup 0
  - DBMT.Updall.Began 04.06.2016 21:27:44
  - DBMT.Updall.Finished 04.06.2016 21:28:03
  - DBMT.Updall.Processed 36
  - 04.06.2013 21:28:05 Database compactor process shutdown

Multiple DBMT Processes

- DBMT takes as input for database(s) to act on:
  - A single database (.nsf)
  - An indirect file (.ind)
  - A directory name
  - No specification - Runs against all databases

- Restricted in order to prevent overlap
  - Only one DBMT process running against all databases is allowed to be loaded at any given time.
    - Any additional DBMT process launched to run against all databases will not be allowed and will simply exit
  - Administrators are allowed to run multiple DBMT processes by specifying a database, indirect file or directory name, but consideration should be given before doing so to avoid overlap
  - DBMT by default, does not compact system databases, unless the system database (.nsf) is specified explicitly on the command line or explicitly in an indirect file (.ind)

Domino 9 Updall

- Updall performs the following tasks by default.
  These are also tasks that the database maintenance tools performs:
  - purges deletion stubs
  - expires soft deleted entries
  - updates unread lists

- New -nodbmt option
  - Skips those operations for updall if you have DBMT configured
  - Faster if you run updall for other reasons

Domino 9 Compact – New Options

- **-# nn**
  - Execute up to nn compactions in parallel using different threads.
  - Up to a maximum of 20 Threads with a default of 1.
- **-W nn**
  - Only compact databases which haven't been compacted successfully in last nn days. Ordered by last compact time.
- **-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.
- **-***
  - Include *.ntf and *.box when processing a directory

Upgrading ODS for DB with old DB Class

- If you used explicit file versions for a database like test.ns7 you cannot use compact to bring the database to a newer ODS level
  - Even when you change the extension of the database, the internal database header has a fixed version which will prevent the ODS change
- "-upgrade Upgrade databases created with older DB classes to the most recent class."
  - For example: **compact -C -upgrade help/** will change the ODS for all help databases
  - Tip: **-ODS** Perform a copy-style compact only if the current ODS is less than desired default ODS.
- Extract from C-API header file
  - **#define DBCLASS_NOTEFILE 0xff01**
  - #define DBCLASS_V6NOTEFILE 0xff10
  - #define DBCLASS_V8NOTEFILE 0xff11
  - #define DBCLASS_V85NOTEFILE 0xff12

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 for 32bit and 1024 MB for 64bit Domino

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 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 – Platform Disk

- Used to check Disk Performance

- Check Server Stats
  - Platform.LogicalDisk.1.AvgQueueLen
  - Platform.LogicalDisk.1.PctUtil

- Interpretation
  - Good < 1 < AvgQueueLen > 3 > 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