



LotuspHERE[®] 2007

JMP103 Operating System Soup - Which Flavor Tastes Best with IBM Lotus Domino?

Daniel Nashed, CTO - Nash!Com Germany

Marc Luescher – IBM Technical Support - Switzerland

Lotusphere® 2007



IBM®

About the Speakers

- Daniel Nashed – Nash!Com - Germany
 - ◆ Nash!Com - IBM/Lotus Advance 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 and Troubleshooting
 - ◆ Platforms: W32, Linux, AIX® and Solaris®
- Marc Luescher – IBM Technical Support - Switzerland
 - ◆ Working for IBM Support Switzerland as Advisory Lotus Service Manager with a strong UNIX and Linux focus
 - ◆ Background of 15 years Lotus and Domino
 - ◆ Co-author of some security, UNIX papers, RedBooks and technotes
 - ◆ Part of Domino subject matter experts for AIX, Solaris and Linux

Agenda

- Operating Systems and Domino
 - ◆ Operating system specific strength
 - ◆ Operating system specific weaknesses
 - ◆ Operating system requirements for Domino
- Best Practices for Administrators
- Best Practices for Developers
- Migrations of Domino between different operating systems
- Domino Tuning, Performance and Scalability
- Questions and Answers

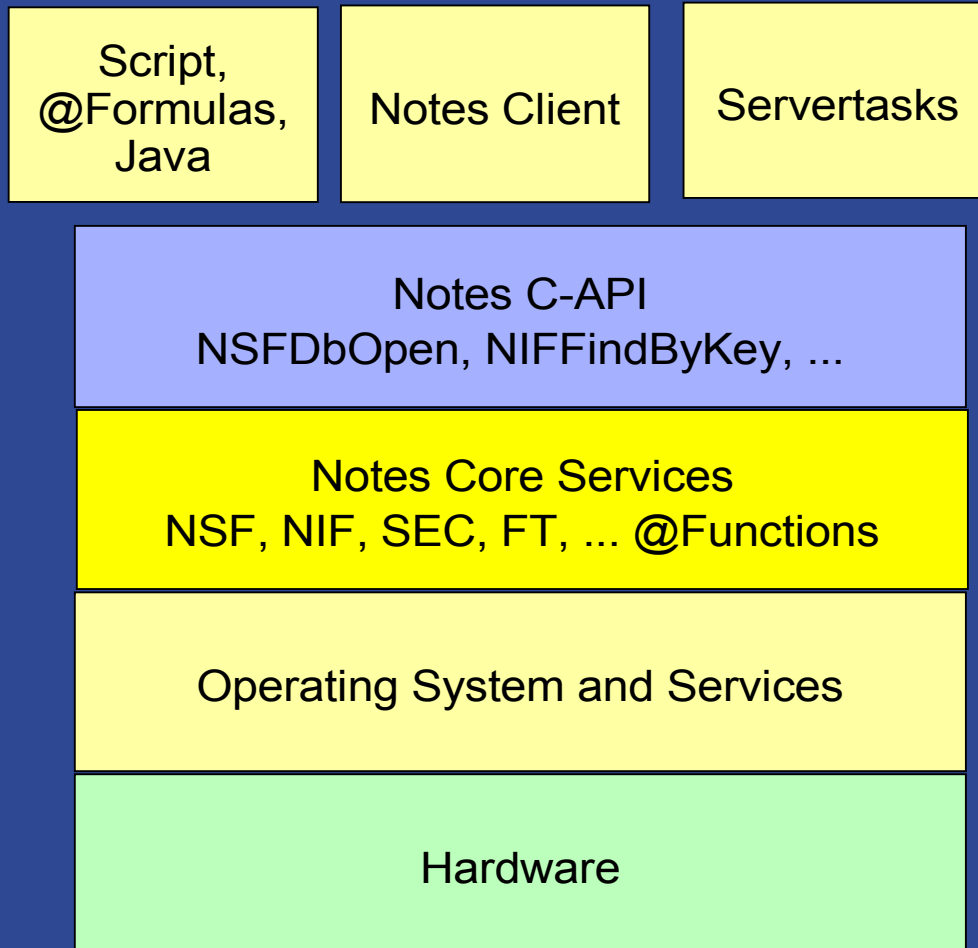
Part 1

- Operating Systems and Domino
 - ◆ Operating system specific strength
 - ◆ Operating system specific weaknesses
 - ◆ Operating system requirements for Domino
- Best Practices for Administrators
- Best Practices for Developers
- Migrations of Domino between different operating systems
- Domino Tuning, Performance and Scalability
- Questions and Answers

Domino is cross platform

- Underlying infrastructure and APIs are designed for multi-platform support
 - ◆ 98% of features work on all platforms
 - ◆ even NSF (ODS) is the same!
 - ◆ published APIs are build to support ISVs to design cross platform solutions
- Supported operating systems give you a wide range of choices
 - ◆ Win32, Linux, Linux, AIX, Solaris, iSeries (AS/400), zSeries (S/390), zLinux (Linux on zSeries)
 - ◆ Each platform has it's strength
 - ◆ There is no "best" platform
 - ➔ choose what best fits into your infrastructure

Domino Infrastructure Overview



Platform independent Notes C-API calls based on lower level core functions & routines

Notes Core services which encapsulate Notes and OS-specific functionality

Platform specific calls and services like NLS, memory management, TCP/IP, file services

Memory, Disk, CPU, other I/O devices

Which platform fits into my environment?

- The real question is: What fits best into my environment?
- This usually depends on your corporate IT platform strategy
 - ◆ Available platform know how, Existing infrastructure,
 - ◆ System Architecture, Server sizes, SAN Topology, ...
- But you need also to take into account your Domino requirements
 - ◆ Some requirements might be better served by some Operating Systems
 - ◆ From the field experience we recommend:
 - ➔ Sametime, Quickplace, LEI, Domino.Doc on Win32
- We will provide details in this session about platform differences

Three Groups of operating systems

- Windows
 - ◆ Most common operating system for Domino
 - ◆ Works fine for small to medium sized servers
 - ◆ Microsoft tries to position Windows as a “Enterprise Operating System”
- Unix/Linux
 - ◆ AIX, Solaris
 - Commonly used in large enterprise environments and for consolidation
 - ◆ Linux
 - Very fast grown platform
 - Since Domino 7.x with 64bit OS and IOCP support large enterprise ready
- IBM Midrange / Mainframe
 - ◆ iSeries (AS/400), zSeries (S/390), zLinux (Linux on zSeries)
 - Mostly used by customers with existing iSeries and zSeries environments
 - Very scalable and high available
 - Those platforms are build and supported by separate teams

Terminology

- S/390
 - ◆ Older IBM Mainframe. 31Bit Architecture
- zSeries
 - ◆ New IBM Mainframe 64 Bit Architecture
- zOS
 - ◆ Traditional Mainframe OS running on zSeries
- xLinux/zLinux
 - ◆ xLinux running on Intel Machines (xSeries)
 - ◆ zLinux running native on zSeries
- LPAR
 - ◆ Logical Partition. First introduced in S/390 Mainframe
 - ◆ Now also used in pSeries (since "Regatta") mainly running AIX, but supports Linux
- DPAR
 - ◆ Domino Partition running on dedicated hardware or inside a LPAR

Domino on xLinux vs. zLinux

- Domino on xLinux and zLinux is almost the same
 - ◆ From Domino and Linux administrative perspective
 - ◆ No special skills needed for zLinux unless you are looking into the deep details of zLinux hardware platform interfacing (5-10% of back-end differences)
 - ◆ You can hardly see no difference. “Feels” like a normal Linux machine
- Same SuSE Enterprise server (SLES8/9) is used
- There might be specific tuning due to machine size
 - ◆ OS-level e.g. kernel parameters, security limits, Host / VM Tuning, ...
 - ◆ Domino level e.g. SharedMemory sizes, BufferPools, PercentAvailSysResources
 - ◆ Number of Domino Partitions (DPARs) per LPAR
- zSeries performance highly depends on the right choice on disk storage
 - ◆ Choice of hardware
 - ◆ Right configuration of Domino files-systems among Domino Partitions

Domino on 64-Bit OS and Hardware

- Currently Domino works in 32-Bit mode on all platforms beside iSeries
 - ◆ iSeries has native 64-bit support and OS already has 128-bit pointers
- Domino runs in a dedicated 32-Bit address space on a 64-Bit OS
 - ◆ AIX, Solaris starting with R5, Linux since 7.0.1, Win64 since 7.0.2
 - ◆ Depending on the platform this allows to use up to 3 GB of shared memory + 1 GB local process memory per Domino Partition (DPAR)
 - ◆ This allows multiple instances of Domino (aka Partitions) to benefit from large memory configurations
- Plans to support Domino 64bit
 - ◆ Each of those implementations will be treated as a separate platform!
 - ◆ Current plan is to have a 64-Bit Beta around Domino 8 GA for ISVs (W64 & UNIX)
 - ➔ Should ship in a later dot release of Domino 8
 - ➔ Currently no plants to support AMD64
 - ◆ Standard Lotus Script/Java/Formula code will just work
 - ➔ C-API applications need recompile!

Domino Memory Limits

- AIX Domino 6/7
 - ◆ 2 Segments a 256MB = 512MB Local Process Memory per servertask
 - Totally sufficient for local memory of all servertaks (HTTP most critical part)
 - ◆ 9 Segments a 256MB = 2.25GB Shared Memory for all processes in common
 - This will change with new configuration parameters for “very large memory model”
- Solaris Domino 6/7
 - ◆ More than 1GB of Local Process Memory
 - ◆ 3GB Shared Memory for all processes in common
- Linux Domino 6
 - ◆ 1GB Local process memory
 - ◆ 2GB Shared Memory for all processes in common
- Linux Domino 7 even with 32Bit Kernel
 - ◆ 1GB local process memory
 - ◆ 3GB Shared Memory for all processes in common (by changing the mapped base)

Domino 7.x Memory Limits Windows

- Domino 7.x on Win32
 - ◆ 2 GB in total (Local Process Memory + Shared Memory)
 - The other part of the 32Bit address space is reserved for system/kernel
 - ◆ /3GB switch would increase this to a 3GB limit (and leave 1GB for system)
 - But this leads to 10-15% performance hit and does not make much sense
- Domino 7.0.2 on Win64
 - ◆ Does not need the /3GB switch and allows Domino to use up to 3GB RAM for Shared and Local Process memory without performance impact of /3GB switch

Partitioning

- What is partitioning?
 - ◆ Run multiple Domino servers (Partition = DPAR) on the same hardware with separate IP addresses, file-systems, etc ...
- Benefit:
 - ◆ Current hardware cannot be fully used by a single DPAR
 - ◆ Higher scalability by running medium sized DPARs on the the hardware
 - ◆ Each partition can use up to 4GB RAM (depending on OS)
- You get good resource usage and higher availability using Partitioning & Clustering
 - ◆ Hints: Cost of not having a service available, Clustering included in Domino 6 standard server license

Domino on Windows - Pro

- Domino on Win32 is most commonly used server platform
- Performance is good for medium sized environments
- Not much Tuning needed for smaller environment
 - ◆ But also not much tuning possible
- Domino Admins are familiar with "Windows"
 - ◆ GUI, File-System, Ctrl-Alt-Delete, Reboot ;-) ...
- Best choice for most add-on products like Sametime, Quickplace, LEI, Domino.Doc, ...

Domino on Windows - Contra

- Few OS level tuning options
- Scalability is limited
- File-System scalability is limited
 - ◆ NTFS has problems with large file-systems > 500GB (e.g. Heavy file fragmentation)
 - ◆ Total Maximum file-system size 2.1 TB – but that is not really recommended
- Fixpack dependences
 - ◆ You can run into interoperability issues when applying required OS security patches, etc
- Performance Monitoring is not comparable to Linux/Unix without using 3rd party tools
- Partitioning is supported but not recommended from Best Practices point of view
 - ◆ Domino uses the system Account
 - ◆ You can only one Domino release at the same time
- Win32 is less secure than Unix/Linux :-)

Domino on AIX - Pro

- Domino on AIX most commonly used Large Enterprise Platform
- True 64Bit OS with dedicated 32bit address space per DPAR
- Native LPAR support + Load Manager since AIX 5.3
 - ◆ With micro partitioning since AIX 5.3 on Power5 CPUs.
- IBM JFS/JFS2 File-system is fast and scalable
- Logical Volume Manager support
 - ◆ Re-Size File-systems on the fly with new hard disks or SAN storage, fault tolerance...
- Good SAN integration
- Good and easy to use Admin Tools (smit)
- Current Power5 CPUs support SMT (simultaneous multi-threading)
 - ◆ Power5 is 20% faster than Power4, SMT adds another 20% of performance gain!
- OS level support for AIX is directly from vendor
 - ◆ Problems can be escalated to IBM Lab (Austin) with additional AIX PMR

Domino on AIX – Contra

- Needs Tuning to perform good
- Segmented memory model
 - ◆ The only Unix OS using a segmented memory model
 - ◆ Causes complicated memory tuning and limits shared mem to 2.25 GB!
- Hardware and Maintenance is quite expensive
- Add-on tools need to be ported to AIX
 - ◆ But most add-ons are available on AIX

Supported AIX Versions

- Domino 6.5.x
 - ◆ AIX 5.1 latest, AIX 5.2 latest and AIX 5.3 latest
- Domino 7.0.x
 - ◆ AIX 5.1 latest, AIX 5.2 latest and AIX 5.3 ML03 or above, ML05 recommended
- Required APAR's
 - ◆ IY68374 (NIS issues in ML03)
 - ◆ IY58691 (Ipv4 and Ipv6 issue in ML00 and ML01)
 - ◆ IY45484 (easier administration if ML00 installed)
- Required Filesets
 - ◆ bos.adt.debug (dbx required for NSD)
 - ◆ bos.perf.tools (required to run topas)
 - ◆ bos.adt.samples (required to run vmtune - since AIX 5.2 vmo is used)

Domino on Solaris - Pro

- Domino on Solaris commonly used Large Enterprise Platform
- 32Bit or 64Bit OS with dedicated operating system
- Excellent 3rd party file-systems available for Solaris 7 and 8 like Veritas Volume-Manager and File-systems
- Solaris UFS and QFS are high performance file systems since Solaris 9
- “Logical volume” support in Solaris 9 and 10 for local disks or SAN's
 - ◆ Re-Size File-systems on the fly with new hard disks or SAN storage
- Native “Zones” support and micro partitioning since Solaris 10
 - ◆ Comparable to AIX LPAR's, but a slightly different concept
- Excellent support in combination with security products and in DMZ configurations
- Long lasting relationship with IBM

Domino on Solaris – Contra

- Needs Tuning to perform good in high end servers with a lot of CPU's and RAM
- Hardware and Maintenance is more expensive compared to Intel based products
- Solaris OS updates (patches) seem to be undertaken very slow
- Older Solaris Releases use 3rd party file systems and volume managers which make tuning much more difficult
- Add-on tools need to be ported to Solaris
 - ◆ Most common add-ons are available on Solaris

Supported Solaris Versions

- Domino 6.5.x
 - ◆ Solaris 8 FCS10/2005 or later and Solaris 9 Level 10 or higher
- Domino 7.0.x
 - ◆ Solaris 9 Level 10 or higher and Solaris 10
 - ◆ Solaris 10 zones support

Domino on Linux - Pro

- Open Standards Based Platform
 - ◆ Huge dedicated Linux Developer community
 - ◆ Many excellent "free" tools (e.g. OS level admin tools, ...)
- Scalability becomes comparable to AIX and Solaris since Domino 7
- 64Bit OS support since D7.0.1
- Fully uses x86 Hardware
 - ◆ Cheaper hardware with Enterprise level OS
 - ◆ Fully takes benefit of hyper-threading on Linux
- "Rock Solid Security"
 - ◆ Used by many customers in DMZ environments, Built from beginning with Security Design
- Good SAN Integration for current Linux Enterprise Servers
- Linux® is the fastest-growing server OS
 - ◆ IDC: Linux server market share based on unit sales rise from 24% to 33% in 2007.
 - ◆ Forrester Research survey: 52% of business users say they are replacing Windows servers with Linux.

Domino on Linux Contra

- Needs an Enterprise Server Linux Distribution with current Service Packs
 - ◆ SuSE Enterprise Server (SLES 8,9,10)
 - ◆ RedHat Enterprise Server 4.0
- You need Linux skills to get it up and running
 - ◆ Needs OS and Domino Tuning to be scalable
- Scalability was limited in D6.x because of missing IOCP Support
- Not all 3rd party products are available for all Domino releases
 - ◆ e.g. Trendmicro for Domino 7 on Linux is still not available

Supported Linux Versions

- Domino 6.5.4/6.5.5
 - ◆ Red Hat Enterprise Linux 2.1 - Update 2
 - ◆ Red Hat Enterprise Linux 3.0 - Update 1
 - ◆ UnitedLinux 1.0/Powered by UnitedLinux 1.0 - Service Pack 2
 - ◆ D6 is not supported on SLES9
- Domino 7.0/7.0.1
 - ◆ SLES 8 - Service Pack 3 (or higher).
 - ◆ pstack-1.1.7.IBM-1 ftp.suse.com/pub/projects/pstack/sles8-i386
 - ◆ SLES 9 - Service Pack 2 (or higher).
 - ➔ gdb (GNU Debugger) is used instead of pstack
 - ◆ Also supports SLES 9 in 64bit Mode!
- Domino 7.0.1
 - ◆ RedHat Enterprise Server 4 (RHEL 4)
- Domino 7.0.2
 - ◆ Adds SLES 10 support! :-)

Checking Minimum Version Level

- /opt/ibm/lotus/.install.dat contains information about OS levels needed for installed Domino release
 - ◆ Copied from install file linux/domino/sets/data/nui.cfg
 - redhat_linux_ent3_os_kernel_minver = "2.6.9-5"
 - redhat_linux_ent3_os_minimum_ver = 4
 - suse_linux_sles8_os_kernel_minver = "2.4.21-138"
 - suse_linux_sles8_os_minimum_ver = "8.1"
 - suse_linux_sles9_os_kernel_minver = "2.6.5-7.139"
 - ...
 - ul_linux_os_kernel_minver = "2.4.21-138"
 - ul_linux_os_minimum_patch = 3
 - ul_linux_os_minimum_ver = "1.0"

iSeries - Pro

- Domino on iSeries is rock solid and very scalable
 - ◆ It is highly optimized
- Build in OS level performance monitoring
 - ◆ Uses the iSeries performance monitoring with build-in hooks into Domino
- True LPAR support for each Domino Partition
 - ◆ Can run multiple instances with different Domino releases
- Very good I/O scalability and I/O sub-systems
- Domino on iSeries is currently the only native 64bit implementation
- Rock solid security
 - ◆ e.g. OS system components are signed

iSeries - Contra

- For most people iSeries is an alien with "green screen" interface
 - ◆ The new i5 has a nice Windows GUI with almost all configuration options
- Very few people know about Domino and iSeries and can support you in case of problems
- Hardware and Maintenance is quite expensive
- Support and Development thru a separate team
 - ◆ Delayed shipment of releases
- Not many add-on tools available
 - ◆ C-API development is quite painful if you don't know the platform in detail

Part 2

- Operating Systems and Domino
 - ◆ Operating system specific strength
 - ◆ Operating system specific weaknesses
 - ◆ Operating system requirements for Domino
- Best Practices for Administrators
- Best Practices for Developers
- Migration of Domino between different operating systems
- Domino Tuning, Performance and Scalability
- Questions and Answers

Key Differences: Windows vs. Unix

- Most Unix systems are based on RISC architecture
- Windows & Linux (most) are based on Intel architecture
- Unix & Linux are designed from scratch to support multi-user, multi-tasking environments!
 - ◆ A lot of Unix services are already implemented on kernel level
 - ◆ Security is essential part of the OS core services
- Domino partitions run with different users and don't use a system account
- Some Windows specific functionality like OLE, DDE, ... is not supported on Unix
- File-systems and path names "look" different (details next page)

UNIX/Linux File System Differences

- There are no drive letters in Unix like C:\
 - ◆ Everything is mounted into the root tree /
 - ➔ Very flexible additional file-systems can be added inside the data directory without using directory links etc.
 - ◆ Take care about applications using full path
- "/" and "\" work on all platforms
 - ◆ Good practice: use / and relative path anywhere if possible
 - ◆ or build path using notes.ini directory entry
 - ◆ when migrating from W32 also check Config/Server document & notes.ini!!!
- Files and directories have owners and file permissions
 - ◆ Make sure your Domino user owns all files in the data directory and the directory itself (also true for translog and other directories)
 - ◆ `chown -R notes:notes /local/notesdata`

UNIX/Linux File System Differences

- Path names on OS level are case sensitive
 - ◆ Best Practice: keep all directories and filenames ASCII lowercase!
 - ◆ Customers migrating from W32 have problems with mixed case filenames
 - ◆ There is an open SPR to make Domino case insensitive on UNIX!
 - But it is not yet committed
 - ◆ In the mean time Nash!Com has a (workaround-)solution based on an Extension-Manager routine (free on Linux; commercial on other platforms)

Filesystem Structure for Domino Binaries

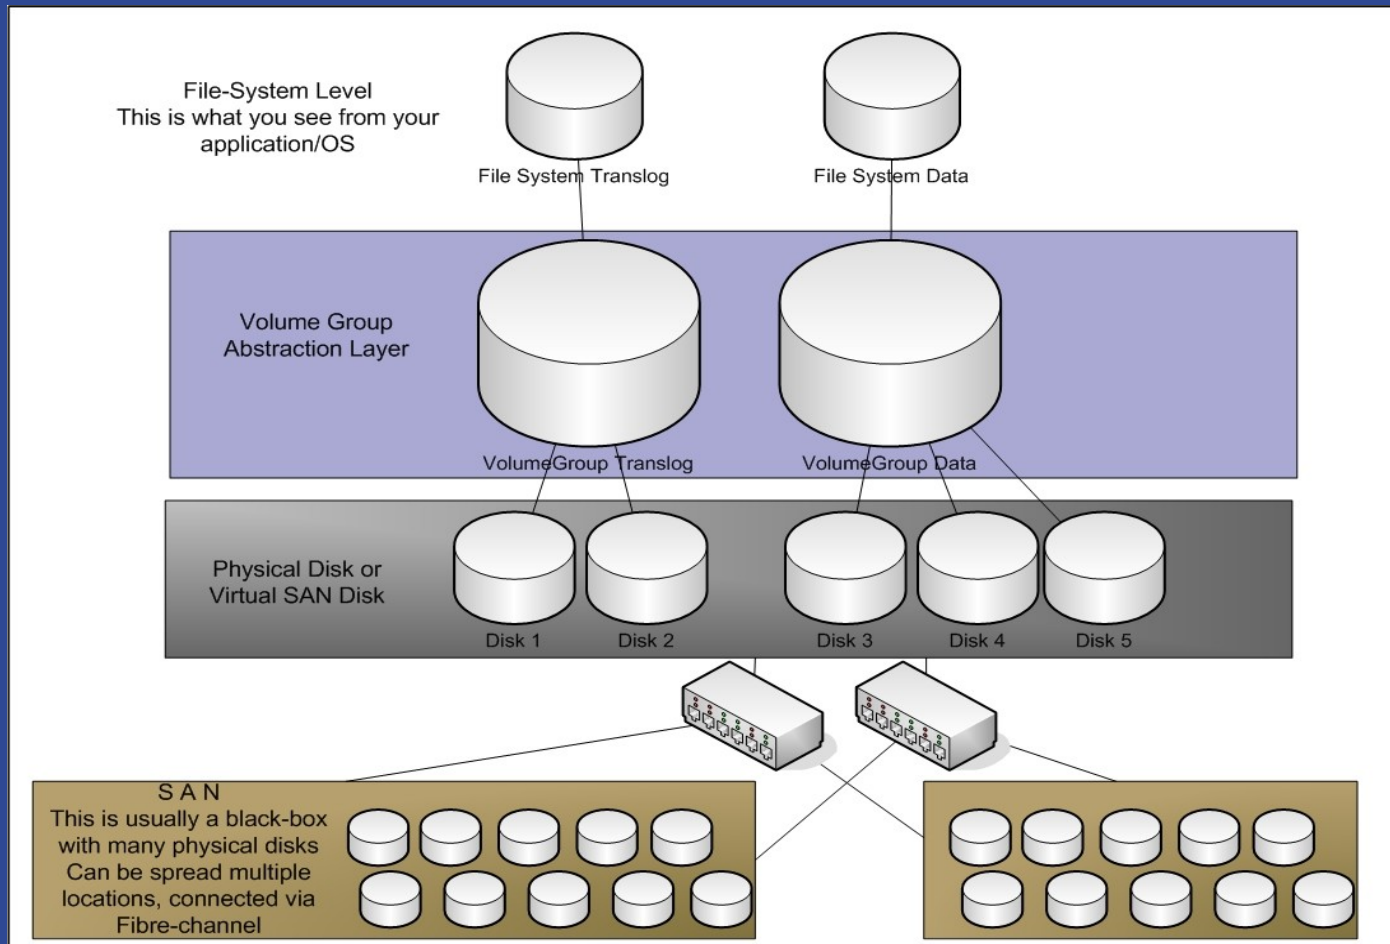
- /opt/ibm/lotus/bin
 - ◆ Contains links to Servertasks and main server binary
 - ◆ Servertasks need a link to tools/startup (ln -s tools/startup taskname)
 - ◆ Start /opt/lotus/bin/server
- /opt/ibm/lotus/notes/latest/"platform" contains binaries
 - ◆ "Platform" could be linux, ibmpow, sunspa
 - ◆ Latest is a link to the current version (e.g. Latest -> 70000)
- /opt/ibm/lotus/notes/latest/"platform"/res/C
 - ◆ Contains resource files
 - ◆ Resource files are the same than what we know from Windows compile
 - ◆ Domino on Unix/Linux has own run-time environment for interpreting resources

Unix File-System Structure

- /local/notesdata is standard for single partition
- Proposed file-systems per partition - each on different disk !
 - ◆ /local/notes1/
 - notesdata
 - translog
 - viewrebuild/
- Enable Transaction Log with standard settings
 - ◆ Use view_rebuild_dir=/local/notes1/viewrebuild/ (3-4 GB)
- /var for system log files
- Swap file (up to two devices)
 - ◆ Have at least the size of the physical memory

Simple SAN Example

- Basic Example to demonstrate how it works



General File-System/Disk Considerations

- Mis-configured I/O Subsystem leads to performance bottlenecks!
- Use RAID1 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

File-System Recommendations - "Notesdata"

- Size: depends on what you want to do on that box for testing.
 - ◆ With LVM and SAN you can increase the size later on
- Recommended: RAID5
- 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 disabled
 - ◆ Unrequired read-aheads can increase the I/O traffic, as well as filling the SAN's I/O cache buffer with unwanted blocks
- 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
 - ◆ Unless you have archive-style translog
- 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.
- Use 2 GB Fibre Channel rather than 1 GB.
 - ◆ Have a dedicated channel and avoid using data switches

Best Practice Partitioning

- Partitioning allows you to optimize the usage of your hardware
 - ◆ Some internal Domino resources do only scale beyond a certain limit
 - ◆ (View/FT-Index, Amgr, Semaphores, Shared Memory, ...)
 - ◆ Don't have too many users per partition – good number is 1500 users
- Use different Unix accounts per partition & get file permissions right
 - ◆ Name Unix user like CN of the Domino Server
- Have separate IP addresses per partition plus one IP for the box
 - ◆ Bind all OS Services to the primary IP of the box
 - ◆ Bind all Notes Services to the service IP of the Domino partition
 - ◆ e.g. TCPIP_TcpIpAddress=0,192.168.1.42:1352
- Take care about resource splitting (memory tuning)
 - ◆ PercentAvailSysResources=n or ConstrainedSHM= ...

Remote Setup Partitioned Servers on Unix/Linux

- Domino on Win32 stores partition information in Registry
- Domino on Linux/Unix stores partition information in /opt/ibm/lotus/.install.dat

ls -la is needed to list files starting with a dot

```
data_directories {  
  "/local/notesdata" {  
    data_UNIX_gid = 101  
    data_UNIX_group = "notes"  
    data_UNIX_uid = 100  
    data_UNIX_user = "notes"  
    ...  
  }  
  ddir_data {  
    "/local/notesdata" {  
      size = 604309335  
      ...  
    }  
  }  
}
```

Install Incremental Installers/Hotfixes on Unix/Linux

- Hotfix is only aware of the standard location for binaries and data
 - ◆ If you switch chose a different binary directory you need an environment setting to specify the location
 - ◆ Example: `export NUI_NOTESDIR=/IBM/domino/lotus`
 - ◆ This allows the installer to find `.install.dat`
- `.install.dat` contains all information about the installed Domino version

Locale Settings on Unix

- Locale defines Date, Time, Currency and other settings ...
 - ◆ You need to set it up properly to get the right results within your Domino applications
- Set Unix Language
 - ◆ Example: `export LANG=de_DE` or `en_US`
- Sync Timezone/DST settings on Unix level and Domino!
 - ◆ Domino for Unix does not sync DST/TimeZone information from OS in all releases
 - ◆ Use DSTLAW (example for EMEA: `DSTLAW=3,-1,1,10,-1,1`)
 - ◆ Check "Show locale" and "Show timezone" via console for proper setup
- Also some locale settings are not read from OS (check notes.ini)
 - ◆ `DateOrder=DMY`
 - ◆ `ClockType=24_HOUR`
 - ◆ `DateSeparator=.`
 - ◆ `TimeSeparator=:`

Unix/Linux Start/Stop Scripts

- Redirect console output stream to a file for logging
 - ◆ Some debug messages are only written to console
 - ◆ Crash info will only written to console
 - ◆ Only use Console Log if needed
 - Uses file-handle per task for writing the log file (overhead)
 - ◆ Java Controller uses Java-Code around Domino main process to control the server (more complexity)
- Redirect console input stream from a file to allow local console
 - ◆ Example: `/opt/lotus/bin/server < console.in >> server.log 2>&1 &`
 - ◆ Attach to input and output files for a kind of "life console"
 - ◆ More a troubleshooting mode. Better use remote console
- Free cross platform start script
 - ◆ <http://www.nashcom.de/nshweb/pages/startscript.htm>
 - ◆ More than a start script.
 - ◆ Allows to start/stop,monitor & troubleshoot your server (NSD...)

Linux Distributions

- SuSE, RedHat and UnitedLinux 1.0 Enterprise are the only supported distributions
 - ◆ Different Releases of Domino support different Versions of each distribution
 - ◆ Take care that you use the right distribution else your Domino server might not run and is totally unsupported
 - ◆ Different versions use different libraries, stack sizes, Java runtime, ...
 - ◆ Using LD_ASSUME_KERNEL=xyz is not really a solution
- Only Enterprise Servers are supported
- Lotus Development build environment for D6.x is still the older egcs++ 2.91 compiler on RedHat 7.2
 - ◆ That's why the compat libs are needed
- Domino 7 uses gcc++ 3.2.2 (see details next slide)

NSD & pstack

- NSD uses pstack for debugging
 - ◆ D6 ships pstack (/opt/lotus/notes/latest/pstack)
 - ◆ D6 on RH 3 AS/ES uses /usr/bin/pstack
 - ➔ need to update it to pstack 1.2-3 or newer for complete call-stacks
- D7 does not ship pstack any more needed for SLES8
 - ◆ SLES8 does not include pstack
 - ◆ download at <ftp://ftp.suse.com/pub/projects/pstack/sles8-i386/>
- SLES9 / RH ES 4.0
 - ◆ Kernel 2.6 always uses the GNU debugger (gdb)
 - ◆ Needs to be installed for NSD

Part 3

- Operating Systems and Domino
 - ◆ Operating system specific strength
 - ◆ Operating system specific weaknesses
 - ◆ Operating system requirements for Domino
- Best Practices for Administrators
- Best Practices for Developers
- Migration of Domino between different operating systems
- Domino Tuning, Performance and Scalability
- Questions and Answers

Writing portable C-API Applications

- Use the C-API instead of platform specific functions
 - ◆ For example no atoi on AS400
 - ◆ Use toolkit data types (like DWORD, ...)
- Do not use W32 specific coding
- Use conditional defines if needed for platform dependent code
- Use ODS-Memory functions for ODS-Structures
- Use OSTranslate and LMBCS for international cross platform applications

API from Lotus Script

- OS level calls are OS level
- C-API calls depend on the OS too
 - ◆ Win32: nnotes.dll
 - ◆ AIX: libnotes_r.a
 - ◆ Linux/Solaris: libnotes.so
- Not just different names but also different size for handles depending on platform!
 - ◆ Not simple to port!
- If you want an example in Lotus Script for cross platform coding send a mail to Daniel@nashcom.de

C-API Platform Specific Compiler

- You need a platform specific C/C++ Compiler
- Windows
 - ◆ Start with Visual C++ 6.0 Std. Edition
- Linux
 - ◆ GNU C Compiler for D7
- Solaris
 - ◆ Sun Forte C 6.0 Update 2
- AIX
 - ◆ AIX Visual Age C++ 5.0
- iSeries
 - ◆ Native ILE C or C++ compiler!

Domino for Linux C-API and Build Platform

- The only supported compilers and build platforms!
- Domino 6.5.x / 6.x / 5.x
 - ◆ Linux egcs++ compiler. Applications must be built on Red Hat 7.2 with the Linux egcs++ compiler (egcs 2.91.66 19990314 which is egcs 1.1.2 release). C API applications built on other Linux environments are not supported.
 - ◆ Make sure you use egcs++ in your make files and to install the right compiler
 - ◆ RH 7.2 install CD images can be downloaded from RH homepage
- Domino 7.0 for Linux
 - ◆ United Linux 1.0 SP3 (-> SLES8 SP3)
 - ◆ GNU Compiler Collection (gcc) version 3.2.2.
 - ◆ Make sure to use g++ in your makefile also for standard C programs
 - ◆ New link-option: -rpath-link \$(NOTESDIR) \$(LIBS)

Part 4

- Operating Systems and Domino
 - ◆ Operating system specific strength
 - ◆ Operating system specific weaknesses
 - ◆ Operating system requirements for Domino
- Best Practices for Administrators
- Best Practices for Developers
- Migrations of Domino between different operating systems
- Domino Tuning, Performance and Scalability
- Questions and Answers

Moving data-bases to a different OS

- You have to decide to transfer the data either
 - ◆ Via Replication
 - Could cause issues with Unread-Marks and private Design (e.g. Folders)
 - ◆ Via file-transfer (FTP, Backup/Restore, ...)
 - Could cause database corruptions depending on file transfer
 - Only copy NSF files and build FT index from scratch
 - Rebuild transaction log
 - Recommended to use: `fixup -f -J` and `compact -D`
 - This will fixup database corruptions and also discards the view index
 - Ensures the databases are consistent
 - Downside: You have to build all indexes by opening the databases manually
 - Workaround: Nash!Com Tool "nshrun" that builds indexes, captures names of build up views, runs multi-threaded, ...
 - ◆ Ensure all databases file-names are US ASCII lowercase !

Other Recommendations

- Keep individual file-systems below 500GB
- Don't use directory links or symbolic links
- You should think about if you want to take over the server name and IP address or build a new server
 - ◆ Tip: Take benefit of clustering between different releases and operating systems during migration
- Take care about the platform depending differences we have mentioned in the development area

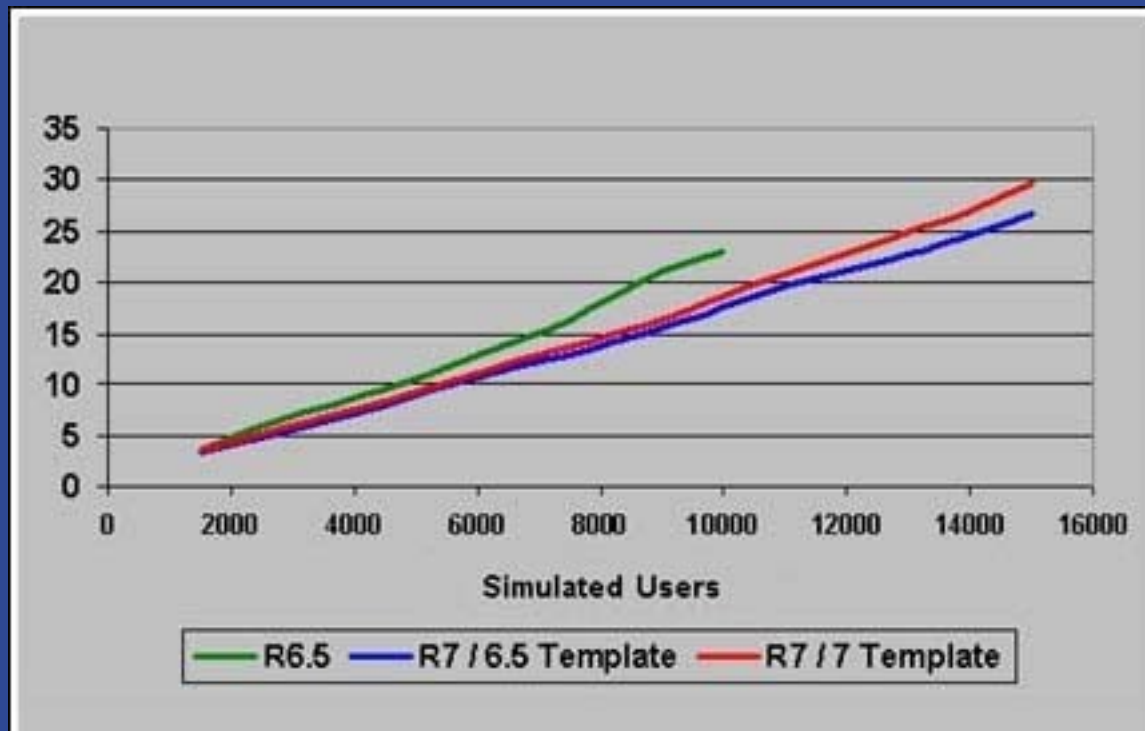
Part 5

- Operating Systems and Domino
 - ◆ Operating system specific strength
 - ◆ Operating system specific weaknesses
 - ◆ Operating system requirements for Domino
- Best Practices for Administrators
- Best Practices for Developers
- Migrations of Domino between different operating systems
- Domino Tuning, Performance and Scalability
- Questions and Answers

Performance Improvements Domino 6/7 for AIX

- Domino 7 reduces CPU on same hardware by 20%

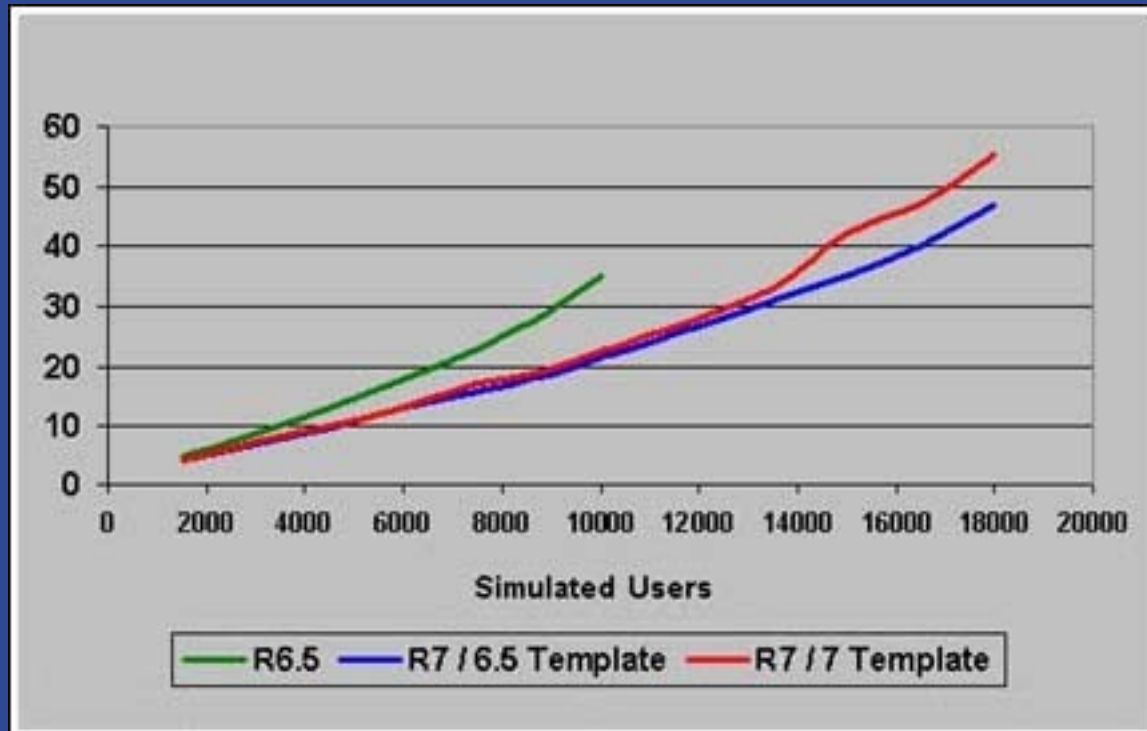
Resource	Domino 6.5	Domino 7	Change (percent)
CPU percent busy	23	18.51	-20
Total disk read KB/sec	1,150,318	1,123,650	-2
Total disk write KB/sec	3,316,083	3,461,609	4
Shared memory used (MB)	1160	1190	3
Process memory used (MB)	23	64	178
Network bytes/sec	2,059,541	2,240,567	9



Performance Improvements Domino 6/7 for Solaris

- Solaris 9 reduces CPU utilisation by almost 35% using same hardware

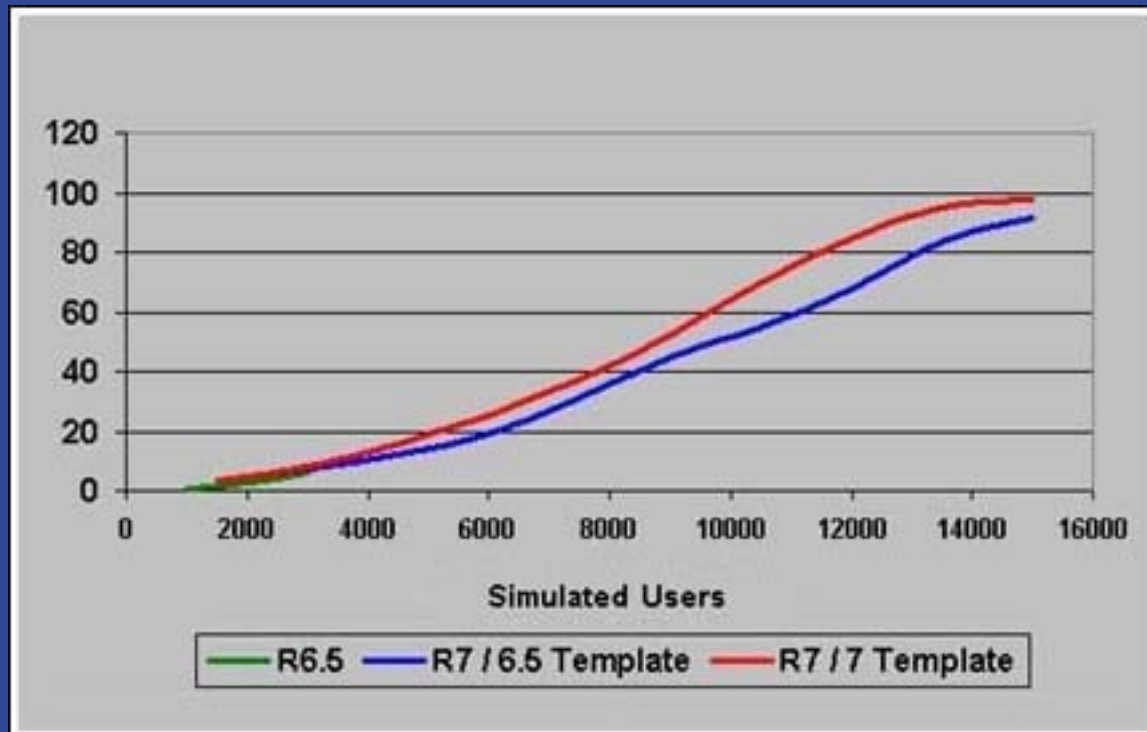
Resource	Domino 6.5	Domino 7	Change (percent)
CPU percent busy	34.9	22.8	-35
Total disk read KB/sec	29,011	30,583	5
Total disk write KB/sec	13,631	13,990	3
Shared memory used (MB)	2706	2173	-20
Process memory used (MB)	18	57	217
Network bytes/sec	1,934,801	1,908,093	-1



Performance Improvements Domino 6/7 for Linux

- Domino on LINUX can now better use the available CPU and memory for better scalability.

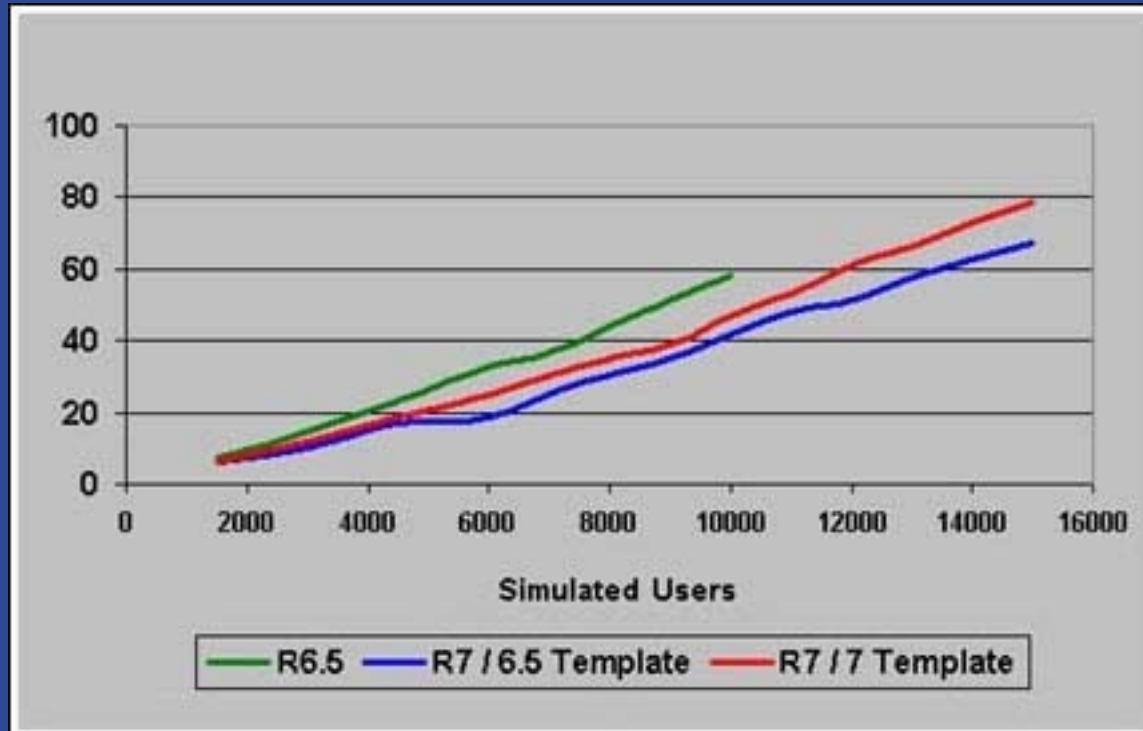
Resource	Domino 6.5	Domino 7	Change (percent)
CPU percent busy	6	8.3	38
Total disk read KB/sec	4082	3770	-08
Total disk write KB/sec	3549	3741	05
Shared memory used (MB)	615	998	62
Process memory used (MB)	947	1,060	12
Network bytes/sec	632,128	725,063	15



Performance Improvements D6/7 for Windows 2003

- Domino 7 reduces CPU requirements by 19% and provides better use of memory.

Resource	Domino 6.5	Domino 7	Change (percent)
CPU percent busy	58	46.9	-19
Total disk read KB/sec	18,651	18,138	-3
Total disk write KB/sec	11,127	11,505	3
Shared memory used (MB)	1294	1184	-8
Process memory used (MB)	23	46	100
Network bytes/sec	2,068,184	2,051,615	1



Performance Improvements by updating to more recent releases of operating systems

- Windows 2000 -> Windows 2003
 - ◆ Updating allows to potentially leverage bigger filesystems
 - ◆ More use of memory for Domino (3 GB) depending on Domino release, etc ...
 - ◆ Use of built in Remote Desktop Console for monitoring
- AIX 5.2 and AIX 5.3
 - ◆ Allow use of 64 Bit mode for Domino which is required with AIX 5.4
 - ◆ Better filesystem and caching support
 - ◆ Improved micro partitioning and LVM support
 - ◆ Far better troubleshooting tools are built in now
- Solaris 9 and 10
 - ◆ Allows the use of built in file systems like UFS Plus and QFS
 - ◆ Allows the use of micro partitioning with Solaris 10 (zones)
 - ◆ Excellent troubleshooting tools available
 - ◆ Default tuning for Domino has become far easier then before
- Linux (SLES 9 and 10 , RedHeat 4)
 - ◆ Far better system drivers (specially SAN) and resource management

Domino Monitoring – show stat platform

- “Show stat platform” is a Lotus Domino console command to give platform related informations.
 - ◆ “Show stat platform” gives statistical details about the current system. Just run those statistics at the beginning and the end of a test period.
 - ◆ On AIX requires perfstat package for better AIX 5L integration.
 - ◆ By default disk statistics are not automatically collected each time the server is rebooted and this must be enabled for each platform
 - ➔ e.g. For Windows run “diskperf -Y”
 - ◆ Use “show platform reset” to reset platform statistics
 - ◆ Platform Statistics are only collected once per minute
 - ◆ For low-level performance troubleshooting you should use platform specific tools

OS-Level Performance Monitoring & Troubleshooting

- Operating System Statistics
 - top, nmon
 - ◆ Information about processes -- nmon AIX and Linux
 - vmstat
 - ◆ Virtual memory activity and CPU statistics
 - iostat -x
 - ◆ I/O statistics and activity
 - netstat -i
 - ◆ Summarizes network activity
 - ◆ You need to install the systat package for iostat and sar
 - Domino Platform Statistics
 - ◆ Great way to get an overview on system performance
 - ◆ Comparable stats cross all platforms (OS level details are described in events4.nsf)
 - ◆ But stats are collected once per minute only

All Platforms Tuning - iostat

- Iostat reports CPU and I/O statistical information.
- You can run it in the background to collect for a specific period by using "iostat 30 120 > iostat.out &" for 30 seconds and 120 times
- Key values :
 - ◆ % tm_act Indicates the time disk was active
 - ◆ Kbps Indicates amount of data that was written per sec
 - ◆ Tps Indicates number of transfers per second.
 - ◆ Kb_read The total number of kilobytes read
 - ◆ Kb_wrtn The total number of kilobytes written

```
System configuration: lcpu=8 drives=6 ent=0.50 paths=12 vdisks=2

tty:      tin      tout      avg-cpu: % user % sys % idle % iowait physc % entc
          0.0      13.6      9.4   5.5  85.2      0.0   0.1  20.0

Disks:    % tm_act    Kbps      tps      Kb_read    Kb_wrtn
hdisk5    0.0        0.0       0.0       0          0
hdisk2    0.0        2.4       0.4       4          8
hdisk1    0.0        0.0       0.0       0          0
hdisk0    0.0        0.0       0.0       0          0
hdisk3    0.0        0.0       0.0       0          0
hdisk4    0.0        0.0       0.0       0          0
```



All Platforms Tuning - vmstat

- For high us/sy columns (user/system CPU) check r (run-queue)
- In case run-queue >> number of CPUs your system is CPU bound
- For high numbers of wa (wait IO in %) check iostat -x to check if machine is I/O bound
- High number of cs (context switches) e.g. more than 5000-8000 context switches per partition you might have a priority inversion

■ System configuration: lcpu=8 mem=8192MB ent=0.50

kthr		memory				page				faults				cpu				
r	b	avm	fre	re	pi	po	fr	sr	cy	in	sy	cs	us	sy	id	wa	pc	ec
1	1	1313504	406653	0	0	0	743	6610	0	149	5106	3665	24	13	62	1	0.23	45.1
2	0	1313504	406081	0	0	0	612	5735	0	150	5581	3663	55	12	32	1	0.39	77.4
1	1	1313504	406576	0	0	0	437	4182	0	204	5446	3560	36	11	52	0	0.29	57.1
1	0	1313504	406157	0	0	0	362	3992	0	137	4840	3482	22	11	66	1	0.21	41.7
1	0	1313503	406454	0	0	0	437	4616	0	102	4940	3375	28	13	59	1	0.25	50.1
1	1	1313496	406362	0	0	0	364	3857	0	194	6602	3788	65	21	13	1	0.49	98.7

All Platforms Tuning - netstat

- Netstat gives an overview of the current adapter configuration with assigned ports and IP addresses.
- It can be used as a first basis for network related troubleshooting before iptrace or perfpmr are being used.

Name	Mtu	Network	Address	Ipkts	Ierrs	Opkts	Oerrs	Coll	
en0	1500	link#2	0.11.25.bf.4f.3f	1500427183	0	526417978	4	0	
en0	1500	172.16.21	gmapdom01	1500427183	0	526417978	4	0	
en0	1500	172.16.21	gmapdom05	1500427183	0	526417978	4	0	
en1	1500	link#3	0.11.25.bf.4f.e5	2376943	0	2736901	3	0	
en1	1500	10.100.84	gmapdom01-d	2376943	0	2736901	3	0	
en1	1500	10.100.84	gmapdom05-d	2376943	0	2736901	3	0	
en2	1500	link#4	3a.2c.40.1.d0.b	4360155	0	6977764	0	0	
en2	1500	172.16.4	gmapdom01-m	4360155	0	6977764	0	0	
lo0	16896	link#1		3489508	0	3506240	0	0	
lo0	16896	127	loopback	3489508	0	3506240	0	0	
lo0	16896	:::1		3489508	0	3506240	0	0	

All Platforms Tuning - mpstat

- “Mpstat” gives an overview on the per CPU usage of the underlying system hardware.
- It can be used as a basic troubleshooting instrument for I/O bound issues

```
System configuration: lcpu=8 ent=0.5
```

cpu	min	maj	mpc	int	cs	ics	rq	mig	lpa	sysc	us	sy	wa	id	pc	%ec	lcs
0	31	37	0	2244	3100	637	0	54	100	3800	76	20	1	3	0.18	36.1	667
1	0	1	0	46	40	13	0	69	100	137	49	2	1	49	0.07	13.3	646
2	0	0	0	9	0	0	0	0	-	0	0	38	0	62	0.00	0.1	19
3	0	0	0	9	0	0	0	0	-	0	0	46	0	54	0.00	0.1	19
4	0	0	0	9	0	0	0	0	-	0	0	37	0	63	0.00	0.1	20
5	0	0	0	10	0	0	0	0	-	0	0	45	0	55	0.00	0.1	20
6	0	0	0	9	0	0	0	0	-	0	0	37	0	63	0.00	0.1	19
7	0	0	0	9	0	0	0	0	-	0	0	43	0	57	0.00	0.1	19
U	-	-	-	-	-	-	-	-	-	-	-	-	0	50	0.25	50.3	-
ALL	31	38	0	2345	3140	650	0	123	100	3937	34	8	1	58	0.25	49.8	1429

AIX Monitoring - topas

- Topas reports selected statistics about the activity on the local system.
- It requires bos.perf.tools and perfagent.tools file sets.
- The mains sections are PAGING, MEMORY and PAGING SPACE.

```

Topas Monitor for host:      gmapdom01      EVENTS/QUEUES      FILE/TTY
Tue Jan  9 16:33:01 2007      Interval:  2      Cswitch      4463      Readch      4565.5K
                          Syscall      7774      Writech     1808.1K
Kernel      18.4      |#####|      Reads        788      Rawin       0
User        62.0      |#####|      Writes       198      Ttyout      946
Wait        3.9      |##|      Forks         0      Igets       0
Idle        15.7     |#####|      Execs         0      Namei       111
Physc =     0.47      %Entc=     93.2     Runqueue     4.0      Dirblk      0
                          Waitqueue    1.0

Network      KBPS      I-Pack      O-Pack      KB-In      KB-Out
en0          421.7     374.1      159.9      340.3      81.4
en1           0.0       0.0       0.0       0.0       0.0
en2           0.0       0.0       0.0       0.0       0.0
lo0           0.0       0.0       0.0       0.0       0.0

Disk      Busy%      KBPS      TPS      KB-Read      KB-Writ
hdisk5    43.8      1305.1     121.5     880.7      424.4
hdisk4    33.3      918.6     127.0     914.6       4.0
hdisk2    17.9     1659.8     132.0     884.7     775.1
hdisk0    10.4      478.2      42.8     279.0     199.3
hdisk3    10.4      243.1      27.4     173.3      69.7
hdisk1     0.4       4.0       1.0       4.0       0.0

WLM-Class (Active)      CPU%      Mem%      Disk-I/O%
dompin01                65        37        15
System                  2         9         0

Name      PID      CPU%      PgSp      Class
server    258056   7.2115.7  dompin01
replica   319694   1.0      24.1     dompin01
lrud      16392    0.2      0.1     System

PAGING
Faults      973
Steals      1142
PgspIn      0
PgspOut     0
PageIn      789
PageOut     369
Sios        1159

MEMORY
Real,MB     8192
% Comp      62.3
% Noncomp   19.4
% Client    19.5

PAGING SPACE
Size,MB     2048
% Used      5.3
% Free      94.6

NFS (calls/sec)
ServerV2    0
ClientV2    0
ServerV3    0
ClientV3    0

Press:
"h" for help
"q" to quit
    
```

AIX Monitoring - perfpmr

- Perfpmr is a set of tools and scripts to collect and analyze in depth the complete OS and hardware setup.
- It allows in depth collection of elements in regards to CPU, Memory, Paging Space, File Systems, I/O, Network Connection etc.
- It requires root access to be run and is used by IBM Support to troubleshoot complex Domino and AIX performance problems.
- A set of perfpmr output can become very large.

AIX Monitoring - truss

- The truss command executes a specified command, or attaches to listed process IDs, and produces a trace of the system calls, received signals, and machine faults a process incurs.
- Each line of the trace output reports either the *Fault* or *Signal* name, or the function name with parameters and return values. The subroutines defined in system libraries are not necessarily the exact system calls made to the kernel.
- Truss does not report these subroutines, but rather, the underlying system calls they make. When possible, system call parameters are displayed symbolically using definitions from relevant system header files.
- By default, undefined system calls are displayed with their name, all eight possible arguments and the return value in hexadecimal format.

AIX Monitoring - smit(ty)

- Smit(ty) is the overall AIX menu guided configuration tool used for configuration changes in a non graphical mode.
- Tip : with "F6 – Command" you always get the command line for the executed command (you can learn from it or modify it and use it in own scripts)

```
Move cursor to desired item and press Enter.
```

```
Software Installation and Maintenance
Software License Management
Devices
System Storage Management (Physical & Logical Storage)
Security & Users
Communications Applications and Services
Print Spooling
Advanced Accounting
Problem Determination
Performance & Resource Scheduling
System Environments
Processes & Subsystems
Applications
Installation Assistant
Cluster Systems Management
Using SMIT (information only)
```

AIX Platform Tuning - sar

- Sar when used with -b option reports buffer activity for transfers, accesses and cache hit ratio per second.
- You can also run it in the background to collect for a specific period by using "sar -b 30 120 > sar.out &" for 30 seconds and 120 times
-
- Key values :
 - ◆ breads Indicates the time disk was active
 - ◆ Kbps Indicates amount of data that was written per sec
 - ◆ Tps Indicates number of transfers per second.
 - ◆ Kb_read The total number of kilobytes read
 - ◆ Kb_wrtn The total number of kilobytes written

```
AIX gmapdom01 3 5 00C1A2DB4C00 01/09/07
System configuration: lcpu=8 ent=0.50
16:29:42 bread/s lread/s %rcache bwrit/s lwrit/s %wcache pread/s pwrit/s
16:29:52 0 0 0 0 0 0 0 0
16:30:02 0 0 0 0 0 0 0 0
16:30:12 0 0 0 0 0 0 0 0
```



AIX Tuning - vmo

- The vmo command is used to fine tune the AIX memory and CPU subsystems in multiple areas.
- It allows to fine tune AIX LPAR for specific application and user load needs
- Key Values :
 - ◆ Minfree Number of Frames on the free list at which VMM starts to steal pages to replenish the free list. Default is 960, values up to 204800
 - ◆ Maxfree Number of Frames on the free list at which page stealing is to stop. Default is 2 x Minfree
 - ◆ Scrub Enables or disables freeing of paging space disk blocks from pages in memory for Deferred Page Space Allocation Policy Settings. Default is 0. Recommended is to set it to 1
 - ◆ ScrubClean Enables or disables the freeing of paging space disk blocks for Deferred Page Space Allocation Policy pages in memory that are modified. Default is 0, recommended is to set it to 1
 - ◆ Lru_File_Repag The LRU daemon begins to run when the number of client pages is within the number of minfree pages of maxclient parameters threshold. The LRU daemon attempts to steal client pages that have not been referenced recently. Default 0 Recommended 1.

AIX Tuning – vmo minperm maxperm

- The operating system takes advantage of the varying requirements for real memory by leaving in memory pages of files that have been read or written. If the file pages are requested again before their page frames are reassigned, this technique saves an I/O operation. These file pages may be from local or remote file systems.
- The ratio of page frames used for files versus those used for computational segments is loosely controlled by the minperm and maxperm values:
 - ◆ If percentage of RAM occupied by file pages rises above maxperm, page-replacement steals only file pages.
 - ◆ If percentage of RAM occupied by file pages falls below minperm, page-replacement steals both file and computational pages.
 - ◆ If percentage of RAM occupied by file pages is between minperm and maxperm, page-replacement steals only file pages unless the number of file repages is higher than the number of computational repages.
 - ◆ Recommendation is to start with minperm%=20 and maxperm%=80 and then adapt slowly as required based on paging and application load.

AIX Tuning - ioo

- ioo command
 - ◆ configures Input/Output tuning parameters. This command sets or displays current or next boot values for all Input/Output tuning parameters. This command can also make permanent changes or defer changes until the next reboot. Whether the command sets or displays a parameter is determined by the accompanying flag.
- Key values:
 - ◆ pv_min_pbuf (default 256 or 512, recommendation 1024)
 - ➔ Specifies the minimum number of pbufs per PV that the LVM uses. This is a global value that applies to all VGs on the system.
 - ◆ j2_maxPageReadAhead (default 128, recommendation 512)
 - ➔ Specifies the maximum number of pages to be read ahead when processing a sequentially accessed file on Enhanced JFS.

AIX Tuning – ioo (2/2)

- Key values (2/2):
 - ◆ j2_dynamicBufferPreallocation (default 16, recommendation 256)
 - Specifies the number of 16k slabs to preallocate when the filesystem is running low of bufstructs. This might be interesting when the speed of i/o load is exceeding the speed of preallocation (on heavy used systems)
 - ◆ j2_nBufferPerPageDevice (default 512, recommendation 2048)
 - Specifies the minimum number of file system bufstructs for Enhanced JFS. This is a further performance boats if above parameter was already set
 - ◆ numfsbufs (default 196, recommendation 930)
 - Specifies the number of file system bufstructs. If VMM must wait for a free bufstruct, it puts the process on the VMM wait list before the start i/o routine is used. This value should be increased on striped volume groups or disk arrays.

AIX Tuning - AIXTHREAD_SCOPE

- AIXTHREAD_SCOPE=S
 - ◆ This AIX environment variable controls the operating system contention scope.
 - ◆ P signifies process based contention scope (M:N)
 - ◆ S signifies system based contention scope (1:1)
- Performance tests have demonstrated performance improvements in response times as well as better CPU utilisation when using the "S" value.
- You can set this variable by including it in the /etc/environment file, but it should also be set through the Domino server start script.
 - ◆ In later releases this is already set by the domino internal start scripts

AIX Tuning - MaxuProc

- MaxuProc=1024 (default 40, recommended 512 or 1024)
- This setting specified the maximum of processes per user ID on AIX.
- To check the current settings use "lsattr -E -l sys0 -a maxuproc" and to change the value either use smitty or use "chdev -l sys0 -a maxuproc=newvalue"
- The low value is usually set to safeguard from creating too many processes of an application. Domino heavily uses that feature by intention.

AIX Tuning - MALLOCMULTIHEAP

- MALLOCMULTIHEAP=32 (default 1 recommended 32)
 - ◆ This environment setting specifies the number of heaps within the process private segment.
 - ◆ Smaller number of heaps can help reduce size of the process. Multithreaded applications like Domino which use the malloc subsystem heavily obtain far better performance by settings this parameter to the maximum value of 32.
- Caution: This setting can cause memory fragmentation when used with HTTP and/or JAVA and should not be used under those conditions.

Solaris Tuning - SEGMAP_PERCENT

- SEGMAP_PERCENT=40 (default = 12, value of 24-60 recommended)
- In order to improve file system caching a portion of the memory can be associated on larger systems for additional file system caching.
- Setting this variable in the etc/system file can greatly improve the performance of Domino's I/O based operations. Improvement in disk performance off up to 30% have been seen.
- Contact the author for a method on how to calculate the best values for this setting, which require domain restart to become active.

Solaris Tuning - Obsolete Parameters

- Solaris Autoup
 - Solaris Tune_t_fshflushr
 - Solaris Maxpgio
 - Solaris Rlim_Fd_Max
- Those SUN recommended tuning parameters are all obsolete if you are using Solaris 9 or 10. The current operating system does not need any further tuning in this area.

LINUX Tuning - Notes_Shared_Dpoolsize

- Notes_Shared_DPoolsizes=xxx
- This environment variable controls the operating system shared memory block size.
- For Solaris this setting is currently set to 8 MB and for Linux it is recommended to set it to 20 MB (20480).
- You can run into memory fragmentation or poor memory utilization when settings this value to high.
- Tip : You can search for "IPC" in a NSD and look in the shared memory section for the segsz and the number of segments.

LINUX Tuning - tune2fs

- tune2fs is invoked as part of the start script and provides the best default tuning parameters
 - ◆ Runs with root privileges (sticky bit) to set tuning parameters
 - ◆ /proc/sys/fs/file-max has been set to "131072".
 - ◆ /proc/sys/kernel/shmmax has been set to "268435456".
 - ◆ /proc/sys/kernel/sem has been set to "250 256000 32 1024".
 - ◆ /proc/sys/net/ipv4/tcp_fin_timeout has been set to "15".
 - ◆ /proc/sys/net/ipv4/tcp_max_syn_backlog has been set to "16384".
 - ◆ /proc/sys/net/ipv4/tcp_tw_reuse has been set to "1".
 - ◆ /proc/sys/net/ipv4/ip_local_port_range has been set to "1024 65535".
 - ◆ /proc/11568/mapped_base has been set to "16777216".
 - ➔ Set the mapped base to allow 3 GB shared memory

LINUX Tuning – Kernel Parameter

- Depending on your start script you need the following line in all /etc/pam.d/* control files like /etc/pam.d/login,su,... to ensure limits can be set
 - ◆ session required /lib/security/pam_limits.so
 - ◆ You can set the limits for the notes user or for * (all users)
 - ◆ SLES10 already has include file concept for all required “session” settings
- Number of open files for notes
 - ◆ /etc/security/limits.conf
 - ➔ notes soft nfile 49152
 - ➔ notes hard nfile 49152
- Number of processes/threads notes
 - ◆ user/etc/security/limits.conf
 - ➔ notes soft nproc 12500
 - ➔ notes hard nproc 12500

Windows Performance Tuning

■ Disable Paging of Code

- ◆ HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session Manager\Memory Management
- ◆ Value Name: DisablePagingExecutive Data Type : REG_DWORD Data : 0 or 1 Default : 0
 - ➔ Enable to stop paging for “pagable” parts of kernel drivers

■ I/O Tuning

- ◆ HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session Manager\Memory Management
 - ➔ Value Name: “IOPageLockLimit”: set to a decimal value of “RAM size in kBytes divided by 8”
- ◆ Increase performance by disabling the last access update file's date.
 - ➔ HKLM\SYSTEM\CurrentControlSet\Control\FileSystem\NtfsDisableLastAccessUpdate
 - ➔ Add key NtfsDisableLastAccessUpdate as REG_DWORD and set it to "1".

Windows Performance Tuning - General

- Deactivate all services you do not need
- Remove all protocols and bindings you don't need
 - ◆ Avoid network port filtering!
 - ◆ No NETBIOS!
- No Screen-Savers - Just lock the console
 - ◆ Screen-Savers can keep the CPU busy
- Depending on your network you might want to check the MTU Size

Domino Tuning - SERVER_POOL_TASKS

- SERVER_POOL_TASKS=20 (default = 20, values up to 160)
 - ◆ This variable defines the maximum number of initial thread pool tasks (IOCP threads) for the Domino server.
 - ◆ By default a value of 20 is used.
- Increment that value in steps of 20 when CPU and I/O load are still low.

Domino Tuning -SERVER_MAX_CONCURRENT_TRANS

- **SERVER_MAX_CONCURRENT_TRANS=100** (default 20 or 40)
 - ◆ This variable defines the maximum number of concurrent i/O threads which can be run at the same time by the Domino server.
 - ◆ It is used as a throttle to reduce CPU usage and Context Switches.
 - ◆ By default a value of 100 is used.
 - ◆ The recommended setting is `SERVER_POOL_TASKS` x # of Domino ports active on server, but at least 2x.
 - ◆ Increment that value in accordance to `SERVER_POOL_TASKS`. Value of up to 1000 have been seen in production for this setting.

Domino Tuning -PercentAvailSysResources

- PercentAvailSysResources=n (2-100)
 - ◆ Assigns a portion of the memory to the Domino server. The maximum amount of memory that a Lotus Domino partition can see is 4 gigabytes.
- It should only be used on systems with more than 4 GB of RAM.
 - ◆ But lacks granularity when adjusting the memory when running on systems with large amounts of RAM (e.g. 16 GB or higher).

Domino Tuning -NSF_BUFFER_POOL_SIZE

- **NSF_BUFFER_POOL_SIZE_MB=xxxx**
 - ◆ Defines the size of the NSF buffer pool, which is the dedicated section of memory dedicated to buffering i/o transfers between Lotus Domino and the operating system.
 - ◆ The size is dependent on the amount of memory available in the system.
 - Good starting point for larger installation around 700-900 MB
- **Check with show stat and look for**
 - ◆ Database.Database.BufferPool.PercentReadsInBuffer
 - ◆ which should have a value in the above 90's. Should this not be the case the NSF_BUFFER_POOL variable needs to be increased.
 - 90 % = OK
 - 95 % = Good
 - 98 % = Excellent

Domino Tuning -NSF_DBCACHE_MAX_ENTRIES

- NSF_DBCACHE_MAX_ENTRIES=n
 - ◆ This variable defines the number of databases for which a server can hold information at one time in its database cache.
 - ◆ This cache contains key information about recently opened databases. If the number is too small older databases will be removed from the cache.
 - ◆ The database cache size defaults to three times the size of the NSF_BUFFER_POOL_SIZE.
 - ◆ Recommendation: Around the number of concurrent users or number of databases on the server -whichever is higher
- Check with show stat and look for Database.DbCache.OvercrowdingRejections
 - ◆ Should have a value close to 0. If not NSF_DBCACHE_MAX_ENTRIES value needs to be slowly increased.

Domino Tuning -NSF_DBUCache_Max_Entries

- NSF_DBUCache_Max_Entries=n
 - ◆ This variable allows the Lotus Domino server to cache the specified number of databases.
 - ◆ The size is dependent on the maximum number of active or concurrent databases per Lotus Domino partition.
 - ◆ Default is

Domino Tuning - ConstrainedSHM

- ConstrainedSHM=1
- Setting this variable will restrict the size of shared memory to a default value per platform.
- The current limits are :
 - ◆ Windows
 - ◆ AIX
 - ◆ Solaris
 - ◆ Linux
- This setting should not be used on AIX 5L systems as it will by default exceed the amount of memory segments available by setting it to 2.25 GB which is over the support 8 segments boundary.

Domino Tuning -ConstrainedSHMSizeMB

- ConstrainedSHMSizeMB=2048
 - ◆ This variable constraints the size of shared memory to a specific amount indicated in megabytes.
- Some third party application which also use shared memory to transfer data might need to use this setting.
 - ◆ Use multiple of 256 MB for the calculated value.
- Tip : You can search for "S_DPOOL" in a NSD and look in the shared memory section for the size and the usage rate of segments.

Domino Tuning -iNotes_WA_ProfileCacheSize

- iNotes_WA_ProfileCacheSize=n
 - ◆ This variable allows the Lotus Domino server to cache the specified number of databases served by the HTTP task.
 - ◆ The size is dependent on the maximum number of active or concurrent databases per Lotus Domino partition used by HTTP.
 - ◆ The effect of this variable are only relevant when using Lotus Domino Web Access.

Part 6

- Operating Systems and Domino
 - ◆ Operating system specific strength
 - ◆ Operating system specific weaknesses
 - ◆ Operating system requirements for Domino
 - Domino Tuning, Performance and Scalability
 - Best Practices for Administrators
 - Best Practices for Developers
 - Migrations of Domino between different operating systems
- Questions and Answers

Resources: Links

- Domino on Linux Homepage
 - ◆ www.ibm.com/lotus/dominolinux
- Domino 7 for Solaris 10 Redbook
 - ◆ <http://www.redbooks.ibm.com/abstracts/sg247162.html>
- Domino7 for Linux Redbook
 - ◆ <http://www.redbooks.ibm.com/abstracts/sg246835.html>
 - ◆ Not completely up to date but good introduction for installation

Resouces: Interesting Free Tools

- A great free Telnet & SSH client: Putty
 - ◆ <http://www.chiark.greenend.org.uk/~sgtatham/putty>
- WinSCP - Windows based copy program using SSH
 - ◆ <http://winscp.sourceforge.net>
- Webmin - Windows based extendable Linux Admin Client
 - ◆ <http://www.webmin.com>
- Nmon
 - ◆ http://www.ibm.com/developerworks/eserver/articles/analyze_aix/

Questions and Answers

- Questions?
 - ◆ Now or send an email after Lotusphere
 - ◆ nsh@nashcom.de
 - ◆ marc.luescher@ch.ibm.com
 - ◆ Updated slides: <http://www.nashcom.de/nshweb/pages/lotusphere.htm>
- Please fill out your evaluations!

© 2007 All Rights Reserved.

- The workshops, sessions and materials have been prepared by IBM or the session speakers and reflect their own views. They are provided for informational purposes only, and are neither intended to, nor shall have the effect of being, legal or other guidance or advice to any participant. While efforts were made to verify the completeness and accuracy of the information contained in this presentation, it is provided AS IS without warranty of any kind, express or implied. Neither IBM nor the speaker shall be responsible for any damages arising out of the use of, or otherwise related to, this presentation or any other materials. Nothing contained in this presentation is intended to, nor shall have the effect of, creating any warranties or representations from the speaker or 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.