Lotus. software

Lotusphere 2003

the essential human element

IBM

Lotus software

Best Practices
Lotus Domino on UNIX and Linux

Daniel Nashed (Nash!Com)
Marc Luescher (IBM Switzerland)

BP 108 / Lotusphere 2003

IBM

About the speakers

- **Daniel Nashed**
- since 1999 Nash!Com own small company
- first German member of Penumbra group

- focused on Cross-Platform C-API Dev., Domino Infrastructure,
- Administration, Integration & Troubleshooting
- on W32, Linux, AIX & Solaris

- Technical writer for German Groupware Magazine

- **Marc Luescher**
- working for IBM Support Switzerland as Technical Lotus Service Manager with a strong UNIX focus

- background of 12 years Lotus and Domino
- co-author of some security and UNIX papers & Redbooks
- part of Domino on UNIX subject matter experts for AIX, Solaris and HP-UX

- good friend of Daniel :-)

IBM

Agenda

- Introduction
- General concepts Domino for Unix
  Special platform considerations

- Designing/Sizing your system
- Tuning your Domino installation
- Troubleshooting your Domino installation

- Questions and Answers

IBM

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
- there is no "best" -- platform choose what best fits into your infrastructure

- Win32, iSeries (AS/400), zSeries (S/390)
- **Linux, AIX, Solaris, HP-UX**

IBM

Which platform fits into my environment?

- Domino scales to the limit of the underlying platform
  - Linux is a good choice for small/medium workload
  - AIX, Solaris & HP-UX are the right choice for larger systems
- A lot of people are using Linux in DMZ for security reasons
  - (IPTables stateful packet filter in kernel, SSH, ...)
- You could use a mixed environment with Linux for smaller servers and Unix for larger machines
- With actual Kernel 2.4 Linux is an excellent alternate platform for Windows based Intel machines
- D6 with 2.6 kernel will really rock!
- AIX 5L on Regatta Servers is great technology for consolidated server infrastructures - we have already customers using it.

Linux on Microsoft Xbox

- http://xbox-linux.sourceforge.net/

Xbox Linux Login Screen

IBM.

Domino Console on Xbox ;-)

IBM.

Full Xbox Environment

IBM.

Which platform fits into your environment?

- You get good ROI using Partitioning & Clustering specially on Unix (-> see BP104 Clusters are cheaper to run ...)
- Hints: Cost of not having a service available, Clustering included in Domino 6 standard server license

- It can make sense to keep Add-on products like LEI, Sametime, QuickPlace on separate machines/partitions.
  - even if centralizing infrastructure to larger machines run add-on products on separate machines/partitions
  - a number of add-ons are Domino release dependent
  - some add-on software is native developed on Win32, than ported to Unix/Linux and has better support on Win32
  - it can make sense to keep those machines on Win32

Domino for Linux

- Operating system is free - you could even download the full installation for example from SuSE or RedHat
  - professional versions support more hardware (RAID controllers, ...)
  - OS level support and professional versions are not free!
- It's one of the most secure and best maintained OS because a lot of very dedicated people who develop Linux
- It is catching up very rapidly and IBM also supports Linux on iSeries and zSeries - but not yet for Domino
- It's one of the fastest growing platforms
- Scalability is still not at large Enterprise level
- But Kernel 2.4 gives you already great performance

Linux Kernel 2.4 Enhancements

- Max number of processes per user has been increased
- File handle limit/Open files increased
- Rewritten kernel I/O subsystems
  - Filesystem
  - Networking
  - Asynchronous I/O
- Better SMP support & process scheduling
  - up to 8 processors
- Supports up to 4GB system memory on Intel
- Network Enhancements
  - Stateful Firewall - Netfilter (iptables)
  - Improved Network Address Translation
- Better Software & Hardware RAID support

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

- filesystems and path names look different (details next page)

Filesystem differences Unix

- There are no drive letters in Unix like C:\
  - Everything is mounted into the root tree /
  - 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!!!
- 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)

General Filesystem/Disk Considerations

- Use RAID1 instead of RAID5 at least for Transaction Log
- Have separate disks for translog, view-rebuild, data
  - Put different filesystems on different controllers/ multi channel controllers for large servers
- Have always at least 20-30% free space on filesystems for data
- If you use SAN (Storage Area Network) it is still recommended to have Transaction Log on fast local disks for performance reasons (depending on your SAN type & configuration).
- Don't use UNIX symbolic links within the data directory
  - Use mounted filesystems or db dirlinks instead
- Choose Journaled filesystems (like JFS, ext3, QFS, Veritas, ...)

Unix Filesystem Structure

- /local/notesdata is standard for single partition
- proposed filesystems per partition - each on different disk !

  /local/notes1/
  - notesdata
  - translog
  - viewrebuild/

- Enable Transaction Log with "Favor runtime"
- Use view_rebuild_dir=/local/notes1/viewrebuild/ (2-4 GB)

- Additional
  - /var for system log files
  - swap file -> usually 2 times main memory
  - Tip: Multiple swap files with same size on different physical disks

Journaled Filesystems

- It's like Transaction Log for Domino but on OS level for filesystem
  - Changes are grouped into atomic transactions, they either happen completely or not at all
  - Operations are first written to a journal file before any changes are made to the filesystem
  - If system fails during commit the transaction can be replayed from the journal file
- Journaled filesystems are faster in most cases
  - specially when not unmounted in clean state (crash)
- Each OS support different filesystems
  - Linux: ext3, ReiserFS, JFS, XFS
  - AIX: JFS, JFS2
  - Solaris: QFS, Veritas
  - HP-UX: HP-UX JFS 3.x
- Some Filesystems are not Large File enabled by default!

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 ...
- 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
- Take care about resource splitting (PercentAvailSysResources=n is your friend; more details later in Marc's part ...)

Filesystem Structure for Domino Binaries

- /opt/lotus/bin (/opt/lotus only required for pre-6.0)
  - 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/lotus/notes/latest/platform contains binaries
  - linux
  - ibmpow
  - sunspa
  - hpux
  - latest is a link to the actual version (e.g. latest -> 60000)
- /opt/lotus/notes/latest/platform/res/C
  - contains resource files
  - Resource files are taken from Windows compile
  - Domino has own run-time environment for interpreting resources

Best practices Unix start scripts

- Redirect console output stream to a file for problem analysis
  - some debug messages are only written to console
  - crash info will only written to console
- Redirect console input stream from a file to allow local console
  - Example: /opt/lotus/bin/server < console.in >> $server.log 2>&1 &
- Don't use cconsole. Attach to console.in and server.log instead (see script in LKB #161094)
  - Also check out Java Domino Console in D6.0
- Write full startup environment to server.log for troubleshooting
  - e.g. set >> server.log; ulimit -a >> server.log; no -a >> server.log
- Export Domino specific environment settings in start script
- Test if script is executed with the right user!

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

- Sync Timezone/DST settings on Unix level and Domino!
  - Warning: Domino for Unix does not get the DST/TimeZone information from OS. It interprets OS time using Domino settings
  - use DSTLAW (Example for EMEA: DSTLAW=3,-1,1,10,-1,1)

- Also locale settings are not read from OS (check notes.ini)
  - DateOrder=DMY
  - ClockType=24_HOUR
  - DateSeparator=.
  - TimeSeparator=:

Check patches with checkos

- cross platform Unix script to check Unix patch-level
  - ships with Domino 6
  - actual version at http://www.lotus.com/ldd/checkos
  - already helps for R5
- It will tell you in details which patches you need to apply
- In general you should have actual OS releases and current patch levels (e.g. AIX 4.3.3 M10, AIX 5.1, Solaris 8, SuSE 8.1, RedHat 8)
  - AIX 5.2 is planned post 6.0.1, Solaris 9 is planned for 6.0.1
- Always check the patch lists in readme.nsf of each actual release

How to install a Domino Server on Unix?

- we will not cover setup/installation ...
  ... but there is a great new Redbook about Domino 6 for Linux covering all details about general Unix/Linux installation
- it's for Linux but it works in the same way on all Unix platforms
- SG24-6835-00 - Lotus Domino 6 for Linux
  - http://www.redbooks.com
  - Installing RedHat, SuSE, and Domino 6 for Linux
  - Improving the performance of your Domino server
  - Administering Domino and Linux

IBM

Domino Infrastructure

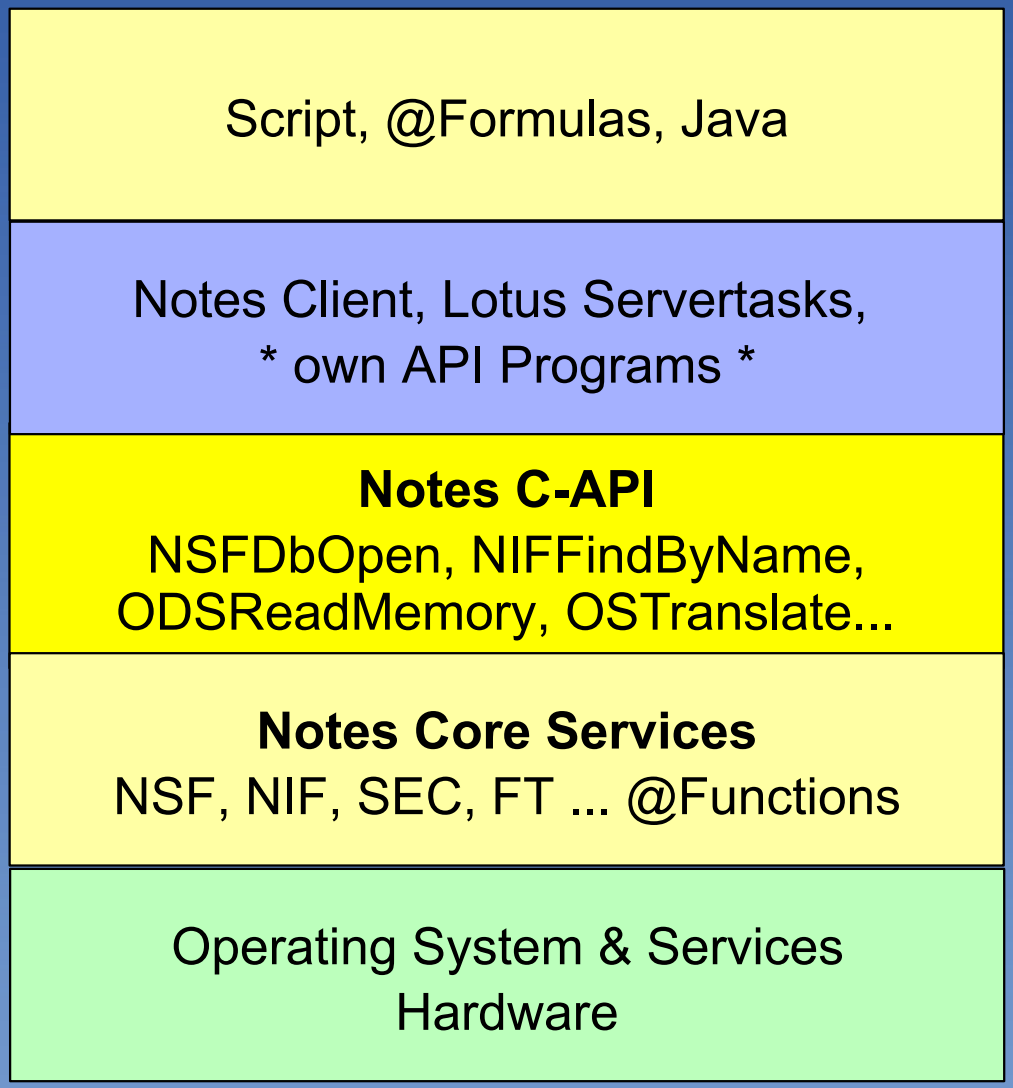

- You may also call C-API from Script
- Lotus core Servertasks are also built using published & unpublished C-API
- Platform independent Notes C-API calls based on lower level core functions & routines
- Notes Core Services which encapsulate Notes and OS-specific functionality
- Platform specific calls and services like NLS, memory, disk, network
- Example: UNIX sockets

IBM

Internal Platform differences
No RISC No Fun ;-)

- ODS (On Disk Structure) is the same for all platforms but in memory presentation of structures might differ
  - the API layer has special calls for reading structures from ODS into memory (e.g. ODSReadMemory) because RISC has other alignment than Intel
  - Constants in C-API global.h take care of platform specificdata type sizes
  - Handles have different sizes on different platforms
- RISC uses Big-Endian (low/high), Intel Little-Endian (high/low)
  - Network-Byte-Order is Big-Endian => no conversion is needed on RISC
- Each platform has some special compile options
  - e.g. AIX uses -qnoro (string literals as read/write for example for BSAFE)
  - Read platform specific C-API documentation very carefully
- Platforms use different native character sets
  - Internal storage is LMBCS (Lotus Multibyte Character Set)
  - Notes has build-in functions for conversion (e.g. OSTranslate)

Example calling C-API from Script on Unix

- Different platform use different extensions for DLLs/Libs
  - *.dll on Windows
  - *.so on Linux & Solaris
  - *.a on AIX
  - *.sl on HP-UX
- Take care: Different platforms use 32bit others 16bit Handles
- For AIX export LDR_CNTRL=IGNOREUNLOAD (LKB #183090)
  - this ensures that Libs are not unloaded in wrong order
  - not needed in D6 any more
- See details in the example "Remote Console in Script"
  - Domino supports a Remote-Console call in Lotus Script ;-)
  - Custom script class in R5 to get the same functionality
  - also uses LBMCS for international characters :-)

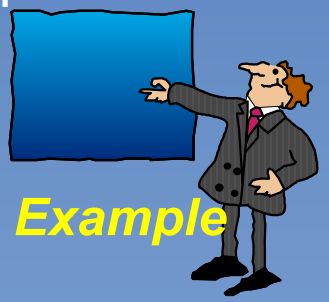

Special AIX Considerations

- Memory
  - AIX programs can only use 11 segments with up to 256MB each
    - Since R5.0.9 Domino uses 8 instead of 10 shared memory segments
    - Some tasks like HTTP had the need for more than 256MB private memory; 2nd segment is now allocated if needed (max. 512MB)
  - If you use third party add-on software check if it already supports the new memory model
  - if not use dataseg utility to patch binaries (see LKB #189972)
- Install & enable IOCP device
  - bos.iocp.rte I/O Completion Ports API
  - make sure it is not just "defined" it must be "available"!
- Install debugger (bos.adt.debug) for NSD to run properly
- Make sure JFS is "Large File Enabled"

Best Practice Unix/Linux Security

- Disable all services not needed
  - e.g. disable inetd or restrict services
- Restrict FTP access and Login to the box
- Get file ownership and permissions right
- Use OpenSSH (Secure Shell) instead of Telnet
  - available for all platforms (see http://www.openssh.org)
  - a great free Telnet & SSH client is Putty (see http://www.chiark.greenend.org.uk/~sgtatham/putty)
- Additional
  - Put Machines into separate Network/DMZ is recommended
  - Install Intrusion Detection Tools (e.g. from Tivoli)

General Tuning Tips UNIX & Linux

- Run Unix in 64Bit Mode, even though Domino application is 32Bit
- Use Multiple Mail.Boxes (2-6 in large systems)
  - Mail_Number_Of_Mailboxes
  - 1 mail.box per 350-500 registered users without exceeding 6 per server.
- Run compact -B once a week to reduce file size
- Use Default_Index_Lifetime_Days=n (default: 45 days) to ensure not needed indexes are discarded
- Run updall at night to have view index up to date in the morning
- Use view_rebuilt_dir=/.../viewrebuild/ on a separate disk to avoid disk I/O contention with log.nsf, names.nsf and mail.boxes

Configuration Overview UNIX

- Standard Domino Configuration
  - notes.ini
  - Server & Configuration Document
  - Special Domino files like httpd.cnf
- Domino on UNIX uses additional areas where to setup configuration values. We will mention a lot of settings in different areas for each platform
- Special settings for UNIX
  - environment (profile)
  - special files (host file, /etc/security/limits, ...)
  - Kernel Parameters

Part Two
Tuning and Troubleshooting your Domino Installation

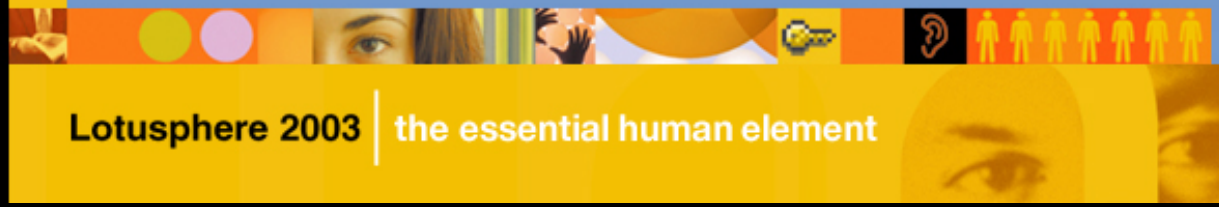

IBM.

What we would like to cover in the 2nd part

- Tuning your Domino Installation
  - My adapted "golden rules" for tuning a system
  - Monitoring your system with basic UNIX tools
  - Operating System Tuning
    - AIX
    - HP-UX
    - Linux
    - Solaris
  - Domino Tuning
- Troubleshooting your Domino Installation
  - Crashes, Hangs, Performance Issues
  - NSD's and other data required

Tuning your Domino Installation UNIX&Linux

- Tuning Best Practices
  - A few tips, mainly valid in the kernel tuning areas but which can easily be adapter for all tuning parameters
    - 1) Monitor your system over an extended period of time
    - 2) Exercise caution when doing changes
    - 3) Read and understand documentation
    - 4) Keep records of all changes at any time
    - 5) Only one change at a time
    - 6) Be aware of potential side effects
    - 7) Evaluate results over time
    - 8) Apply next setting

IBM.

Tuning your Domino Installation UNIX&Linux

Monitoring Performance

- Before doing any changes on your system it is a good advise to monitor your system for a certain period of time
- Short Term Monitoring (see next slides for more details)
  - iostat -x io statistics and activity
  - vmstat virtual memory activity and some CPU statistics
  - mpstat detailed look on CPU statistics
  - netstat -i summarizes network activity
- Long Term Monitoring (check with your OS vendor for already provided long temr monitoring scripts)
  - use the platform provided performance history scripts e.g. /etc/init.d/perf for Solaris to keep a 30 day history of system behavior.

Tuning your Domino Installation UNIX&Linux

- Monitoring Performance
  - Short Term Monitoring (Snapshots)
    - iostat -x reports disk statistics
    - Watch the %b column, values > 20 % need to be checked against high svt_t column.

extended device statistics

device	r/s	w/s	kr/s	kw/s	wait	actv	svc_t	%w	%b
sd0	0.3	10.9	4.0	141.3	0.0	0.1	9.1	0	6
sd1	0.0	0.0	0.7	0.2	0.0	0.0	12.2	0	0
sd2	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0	0
sd15	0.3	10.9	3.9	141.3	0.0	0.1	9.2	0	6
sd16	0.0	0.0	0.8	0.2	0.0	0.0	11.2	0	0
ssd12	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0
ssd13	42.1	4.4	3437.7	67.9	0.0	0.5	11.7	0	19
ssd14	47.7	11.1	3621.4	194.4	0.0	0.7	11.5	0	23
ssd15	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0
ssd16	19.0	37.1	1162.3	3624.2	0.1	2.2	40.5	1	19
ssd17	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0
ssd18	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0
ssd19	23.2	23.6	1251.2	2663.2	0.0	0.7	14.2	0	15
ssd20	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0
ssd21	19.5	23.7	1354.3	2662.5	0.0	0.6	14.0	0	16
ssd22	12.5	34.0	709.3	3217.9	0.0	0.7	15.9	0	16
ssd23	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0
ssd24	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0
ssd25	43.6	5.7	3380.1	88.0	0.0	0.6	11.4	0	20
ssd26	47.1	9.9	3602.8	163.0	0.0	0.6	10.3	0	23
ssd27	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0
nfs1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0

Tuning your Domino Installation UNIX&Linux

- Monitoring Performance
  - Short Term Monitoring (Snapshots)
    - vmstat summarizes virtual memory activity and some CPU statistics.
    - Monitor the sr column to keep track of the scan rate.
    - Watch us, sy and id columns to check CPU usage, keep track of the r column to see how many threads are contenting for CPU time, if it remains higher then about 4x the number of CPU's then reduce server concurrency.
    - Watch out for a high number of Context Switches (cs)

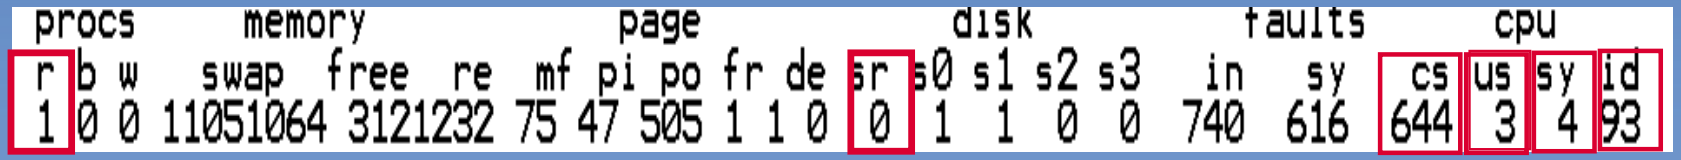

procs			memory			page						disk				faults		cpu			
r	b	w	swap	free	re	mf	pi	po	fr	de	sr	s0	s1	s2	s3	in	sy	cs	us	sy	id
1	0	0	11051064	3121232	75	47	505	1	1	0	0	1	1	0	0	740	616	644	3	4	93

IBM.

Tuning your Domino Installation UNIX&Linux

- Monitoring Performance
  - Short Term Monitoring (Snapshots)
    - mpstat 60 gives a more detailed look on CPU statistics
    - watch out for high numbers in usr, sys or wt

CPU	minf	mjf	xcal	intr	ithr	csw	icsw	migr	smtx	srw	syscl	usr	sys	wt	idl
0	14	3	670	405	299	598	11	38	307	0	610	6	8	4	82
1	14	4	485	195	188	603	13	41	278	0	642	7	8	5	81
2	14	3	631	689	683	553	11	38	315	0	619	7	9	4	81
3	14	3	711	187	181	596	12	41	282	0	639	7	8	4	81
4	13	2	16	181	176	602	9	39	278	0	589	5	7	3	85
5	11	3	716	125	120	628	10	40	279	0	596	5	7	4	85
6	11	3	14	105	99	635	10	40	280	0	592	5	7	4	85
7	10	2	30	135	129	631	10	40	278	0	586	5	7	4	85
8	10	3	91	117	111	657	11	39	280	0	597	5	7	4	84
9	10	3	672	106	100	649	11	40	282	0	588	5	7	4	84
10	10	3	628	106	99	653	10	40	284	0	585	4	7	4	84
11	10	3	25	106	99	639	10	40	280	0	589	5	7	4	84
12	12	4	649	107	100	658	13	43	291	0	647	6	8	5	82
13	13	4	24	107	100	657	13	43	289	0	649	6	8	5	82
14	14	3	90	107	100	647	12	43	288	0	655	6	7	4	82
15	14	3	81	107	100	650	12	44	288	0	655	6	7	4	82

Tuning your Domino Installation UNIX&Linux

- Monitoring Performance
  - Short Term Monitoring (Snapshots)
    - netstat -i summarizes network activity
    - watch out for ierrors, oerrors and collisions

Name	Mtu	Net/Dest	Address	Ipkts	Ierrs	Opkts	Oerrs	Collis	Queue
lo0	8232	loopback	localhost	17530502	0	17530502	0	0	0
ge0	1500	se57do01ge0	se57do01ge0	2841178086	0	1188181633	0	0	0
ge1	1500	se57do01ge1	se57do01ge1	3488506	0	127	0	0	0
qfe0	1500	se57do01en0	se57do01en0	1128512	0	89209960	0	0	0
qfe1	1500	se57do01en1	se57do01en1	38584398	37	79222613	0	0	0
qfe2	1500	se57do01en2	se57do01en2	90254130	0	100756013	0	0	0
qfe3	1500	se57do01en3	se57do01en3	46128	0	838	0	0	0
qfe4	1500	se57do01en4	se57do01en4	9384	0	545	0	0	0
qfe5	1500	se57do01en5	se57do01en5	30936689	12	838	0	0	0
qfe6	1500	se57do01en6	se57do01en6	87213275	0	2845	0	0	0
qfe7	1500	se57do01en7	se57do01en7	45780	0	422	0	0	0

Tuning your Domino Installation UNIX&Linux

- Monitoring Performance
  - Long Term Monitoring
    - use the platform provided performance history scripts et /etc/init.d/perf for Solaris to keep a 30 day history of system behavior.
    - or use the SE toolkit www.setoolkit.com (freeware) as a platform for more intelligent system monitoring
    - written by Rich Pettit
    - and Adrian Cockcroft
    - two well known Solaris gurus

Tuning your Domino Installation UNIX&Linux

- Domino Platform Tuning for AIX 1/3
  - AIX Kernel Settings
    - MaxUProc
      - This AIX setting determines the maximum number of processes that a single user account can run. By default this is to small for a Domino server. The recommendation is to set this to at least 128.
      - Check the current value with :
      - lsattr -E1 sys0 -a maxuproc
  - AIX Domino Server Account Settings
    - AIXTHREAD_SCOPE=S , MALLOCMULTIHEAP=1 and
    - Since 4.3.2 AIX has its own threadpool which shows big performance increases. In the Domino user's .profile enter :
    - export AIXTHREAD_SCOPE=S (default in Domino 6)
    - export MALLOCMULTIHEAP=1

Tuning your Domino Installation UNIX&Linux

- Domino Platform Tuning for AIX 2/3
  - AIX Domino Server Account Settings (to get started with...)
    - /etc/security/limits
      - data = (min 3 x default value)
      - stack = (min 3x default value)
      - rss = (min 3 x default value)
      - nofiles = (min 3 x default value) better unlimited
  - Page Space adjustments
    - /etc/inittab/vmtune (to get started with....best is to have no paging)
      - Page space should be by default about 2x the amount of RAM of your server until you exceed 512 MB RAM. if the server shows more than 20% paging space used with lsps -a command then reduce vmtune as shown below :
      - /usr/samples/kernel/vmtune -P 50 (down to 20)

Tuning your Domino Installation UNIX&Linux

- Domino Platform Tuning for AIX 3/3
  - AIX Network Specific
    - Network Tuning
      - Network interface adapter tuning in order to ensure that large amounts of data can be transferred and received by the system. The minimum of tx_queues should be 512 or higher if supported by hardware.
      - To view current settings :
        - lsdev -Cc adapter -s available (list adapters)
        - lsattr -e1 <adapter> ¦grep size (list queue sizes)
    - Default Socket Buffers (to get started with...)
      - Set the send and receive window size to 16384 Bytes. you can check them with the no -a command.
        - tcp_sendspace=16384
        - tcp_recvspace=16384
        - sb_max=65536

Tuning your Domino Installation UNIX&Linux

- Domino Platform Tuning for Solaris 1/3
  - Solaris Memory Specific
    - File descriptors limits
      - To increase the number of file descriptors for the Domino server in /etc/system
      - set rlim_fd_max=65536 (mandatory)
    - Memory tuning will improve the system response time when the file system is heavily used. The cachefree settings should be set to about 1/16 of physical memory divided by 8192 to get 8 KB pages. It can be set up to a value of 1/6 of physical memory.
      - set priority_paging=1 and cachefree=8192 (for a 1GB)
      - Pre Solaris 8 setting (only for Solaris 6 and 7)
    - High filesystem page in rate or better performance = ON !
    - If high file system page in rates are happening set the following value in /etc/system. It adjusts the percentage of memory that the kernel will map into address space for the file system cache.
      - set segmap_percent=20 (default 12 up to 60)

Tuning your Domino Installation UNIX&Linux

- Domino Platform Tuning for Solaris 2/3
  - Excessive filesystem flush daemon overhead
    - On very busy Domino servers with a lot of I/O you may want to limit the CPU cycles that the file system flush daemon can use.
      - set autoup=600
      - set tune_t_fsflushr=1
      - Those settings wake up the fsflush daemon once a second and allow older pages to remain in memory before being flushes to disk
  - Throtted page daemon I/O
  - When using multiple swap disks or 10'000 RPS disk you should increase the I/O rate of the page daemon.
    - set maxpgio=16384

Tuning your Domino Installation UNIX&Linux

- Domino Platform Tuning for Solaris 3/3
  - No risk no fun (experimental settings no golden rule)
  - In case of unpredictable intermittent slowdowns in network response time from a consistently loaded server try investigating into setting sq_max_size parameter by adding into /etc/system
    - set sq_max_size=512

Tuning your Domino Installation UNIX&Linux

- Domino Tuning for HP-UX 1/3
  - Kernel threads
    - Splitting processes into threads allows to run simultaneously on multiple processors and is therefore improving performance for Domino a lot. Set values as below :
    - max_thread_proc = (#Domino Users +500)
    - nkthread = (#Domino Users +1000)
  - Streams
    - An I/O pipeline which passes serial data between the HP-UX operating system and a kernel driver associated with a raw like a terminal. Set the value as below :
    - NSTREVENT = (#Domino Users) Min 50

Tuning your Domino Installation UNIX&Linux

- Domino Tuning for HP-UX 2/3
  - Asynchronous I/O
    - POSIX asynchronous I/O allows a process or thread to start multiple simultaneous reads and/or writes to multiple files. With POSIX async I/O that can overlap some elements of computation and processing. Set the values as below :
      - aio_max_ops = (2 x #Domino Users)+500
      - aio_phsysmem_pct = 50 (max % of sys memory for POSIX)
      - fs_async = 1 (asynchronous instead of 0 synchronous)
  - Files (OS file limits to be raised for Domino)
    - maxfiles = (4x #Domino Users)
    - maxfiles_lim = (4 x #Domino Users) +1
    - nfile = (5x #Domino Users) + 500
    - nflocks = (#Domino Users + 50)
    - ninode = (5x #Domino Users + 500)

Tuning your Domino Installation UNIX&Linux

- Domino Tuning for HP-UX 3/3
  - Semaphores
    - Interprocess Communication semaphores are used to keep processes synchronized and to prevent collisions when accessing shared data structures. Set the value below for improved Domino performance :
      - semmap = (256x Max#DominoPartitions)+256)
      - semmns = (1024x Max#DominoPartitions)
      - semmnu = (512x Max#DominoPartitions)
      - semume = (512x Max#DominoPartitions)
  - Shared Memory
    - Shared Memory is reserved space for storing data structures and data being shared among processes. They are based in the kernel. Set the value below :
      - shmseg = (256x Max#DominoPartitions)+128
      - shmmax = 0x40000000 (max size of shared mem segment)

Tuning your Domino Installation UNIX&Linux

- Domino Tuning for Linux / Kernel Parameter

- Two ways to set kernel parameters
  - Example SuSE: echo "8192" >/proc/sys/fs/file-max
  - Example RedHat: file /etc/sysctl.conf: fs.file-max=49152

- Number of open files
  - fs.file-max=49152
  - fs.inode-max = 3*fs.file-max
    - check actual values: echo /proc/sys/fs/file-no; inode-nr & inode-state
- Increase maximum number of shared memory segments
  - kernel.shmmni=8192
- Increase maximum number of system-wide semaphores
  - kernel.sem=250 18432 32 1024
- Increase maximum number of threads
  - /proc/sys/kernel/threads-max

Tuning your Domino Installation UNIX&Linux

- Domino Tuning for Linux / Limits
- 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
- Number of open files for notes
  - user/etc/security/limits.conf
    notes soft nofile 49152
    notes hard nofile 49152
- Number of processes/threads notes
  - user/etc/security/limits.conf
    notes soft nproc 12500
    notes hard nproc 12500

Tuning your Domino Installation UNIX&Linux

- Domino Tuning for Linux
- IPC/Shared Memory (Domino specific)
  - export Notes_SHARED_DPOOLSIZE=20000000 to enlarge the shared memory segments allocated
- Set the "noatime" attribute flag to disable the file access modified timestamp on all filesystems for Domino, /var, /tmp in /etc/fstab
- Some additional parameters to check
  - /proc/sys/vm/buffermem
  - /proc/sys/vm/bdflush

Tuning your Domino Installation UNIX&Linux

- Domino Tuning for Linux / Hard Disks
  - Default settings for disks are very conservative
  - Tune settings with hdparm
    - hdparm -A1 -a8 -c3 -d1 -Xnn -W1 devicename
    - (-A) sets drive read lookahead flag
    - (-a) sets FS read ahead. 8 sectors (4KB) to 12 sectors (6KB)
    - (-c) sets EIDE 32bit I/O support
    - (-d) enables DMA
    - (-X) sets the DMA mode (see man page for details)
    - (-W) IDE write caching mode (be aware of possible data lost when crash)
    - Example: hdparm -A1 -a8 -c3 -d1 -Xnn -W1 devicename
  - And test Performance before and after with
    - (-t) perform device read timings
    - (-T) perform cache read timings

Tuning your Domino Installation UNIX&Linux

- Domino Tuning - Most Important Overall Server Settings
  - PercentAvailSysResources (notes.ini) Best Setting
    - Per partition memory size, expressed as a percent of total system memory. Should normally be 100% if only one partition or 50 % if less. You might want to experiment with a rule of leaving about 1 GB RAM for the OS and only use the rest for the Domino partitions.
    - e.g. PercentAvailSysResources=50
  - ConstrainedSHMSizeMB (notes.ini) New Setting for D6
    - New D6 parameter to restrict the memory usage of Domino in a D6 architecture, see Eddy in the Lab for more details.
  - NSF_BUFFER_POOL_SIZE(_MB) (notes.ini) only if required by Support
    - Buffer pool allocation per partition Domino server. Might be required if high number of memory > 4 GB RAM or old Domino release < 5.04. Can be 2000 MB at most per partition.
    - E.g. NSF_BUFFER_POOL_SIZE=2000

Tuning your Domino Installation UNIX&Linux

- Domino Tuning - Most Important Overall Server Settings
  - Server_Max_Concurrent_Trans (notes.ini)
    - This settings allows you to limit the number of concurrent transactions a server can handle. Only required on very large systems, default should be fine. e.g 1000
  - Server_Pool_Tasks (notes.ini)
    - In case of a high number of NRPC connections increase the number of Server_Max_Concurrent_Trans to a high value (1000) or -1. At the same time set the Server_Pool_Tasks parameter to a value of 50 to 100 to specify the number of server threads. Best practice to start with 10% of Server_Pool_Tasks value
  - NSF_DbCache_MaxEntries
    - In case of high statistical value in DbCache.OvercrowdingRejections then set this parameter to either the maximum number of concurrent users or the maximum number of databases open (whatever is higher)
    - e.g. NSF_DbCache_MaxEntries=10000

Troubleshooting your Domino Installation

- 1) Contact support
- 2) Open an incident (opt fwd me the incident number by mail)
- 3a) For Crashes
  - Run Full NSD
  - Collect at least last screen of console log
  - Memory check dump if generated
  - Core dump if generated
  - REQ files if HTTP server crash, might need to turn on using "tell http debug thread on" or in httpd.cnf
- 3b) For Hangs
  - Add DEBUG_SHOW_TIMEOUT=1 to notes.ini
  - Add DEBUG_CAPTURE_TIMEOUT=1 to notes.ini
  - Add DEBUG_THREADID=1 to notes.ini
  - Run Full NSD
  - Collect at least last screen of console log
  - Memory check and core dump if generated

Troubleshooting your Domino Installation

- 3c) For Memory Leaks
  - Note: Out of Memory does not mean that you are dealing with a memory leak. Also do not assume the subsystem reporting low memory conditions to be the cause of memory leaks.

    - SET DEBUG_MEMDUMP_TIMEDATE = 1 in notes.ini
    - Show memory dump at regular intervals will create memory.dmp in Notes data directory
    - Multiple NSD's in short intervals (3-5 times in a row)
    - Option to use NSD -noinfo -nomem to get only key data
    - Collect console log for that period
    - You might also use a UNIX shell script using ps -elf and server -m to do the same at more regular intervals

Troubleshooting your Domino Installation

- 3d) For Performance Problems
  - Multiple NSD's in short intervals (3-5 times in a row)
  - Option to use NSD -noinfo -nomem to get only key data
  - Statrep.nsf of Domino Server
  - Vital system monitoring data iostat, vmstat, mpstat
  - on AIX systems :
    - use svmon -P ALL or svmon -G
    - use tprof -skex sleep <intervall> needs bos.perfagent.tools otherwise use ps avwwg
    - try binding partitioned servers to processors with system with large L2 cache sizes and > 4 CPU's
    - use schedtune
  - on Solaris systems :
    - Forte Collector and Analyzer toolkit Version 6 or 7
    - Free with Solaris 8 and 9, 60 day free license for Solaris 7.
    - This toolset is used by Sun and IBM/Lotus to troubleshoot complex performance issues.

Troubleshooting your Domino Installation

- ## Server Console Log with Semaphores and NSD Sample Output

When the server hangs output similar to following will be generated.
THREAD [04162:00063-05398] WAITING FOR RWSEM 0x91BD (@2016C310) (R=0,W=4,WRITER=04162:10538,1STREADER=04162:10538) FOR 30000 ms
THREAD [04162:00063-05398] WAITING FOR RWSEM 0x91BD (@2016C310)

Taking an NSD the stack for the waiting thread will look similar to the following

```
########################################
## thread  63/63  :: server pid=4162, k-id=  15625  , pthr-id=5398
## stack          :: k-state=wait, stk max-size=98304, cur-size=5592
########################################
WaitOnNativeSemaphore(0x203c5210, 0x0, 0x0, 0x0) at 0xd09dee90
OSLockSemInt(??, ??) at 0xd09dca4c
OSLockSemInt(??, ??) at 0xd09dca4c
OSLockWriteSem(??) at 0xd09e09c0
ANSPushContext__9ANSessionFv(??) at 0xd20bc2cc
ANSetSessionContext(??) at 0xd20b37f8
LsxMsgProc(??, ??, ??) at 0xd20adcdc
SendLSXMessage__7DLLNodeFUsUlT2(??, ??, ??, ??) at 0xd12102b0
SendLSBEMessage__17LSIClassRegModuleFUsUlT2(??, ??, ??, ??) at 0xd0a5ed54
SetRunContext__12CLSIDocumentFUsPv(??, ??, ??) at 0xd0e7221c
Run__21CRawActionLotusScriptFP13CDefActionCtxUsPUl(??, ??, ??, ??) at 0xd0e6386c
Run__10CRawActionFP13CDefActionCtxUsPUl(??, ??, ??, ??) at 0xd0e5da64
Execute__10CRawActionFP13CDefActionCtx(??, ??) at 0xd0e5dee4
RunAlone__10CAssistantFP13CDefActionCtx(??, ??) at 0xd0e4f88c
Run__10CAssistantFP21tagASSISTRUNCTXSTRUCT(??, ??) at 0xd0e4c5bc
AgentRun(??, ??, ??, ??) at 0xd0e49364
ServerRunServerAgent(??, ??) at 0x1003f56c
DbServer(0x3b985c, 0x14e0c) at 0x100146a4
WorkThreadTask(??, ??) at 0x10050bc4
Scheduler(??) at 0x10041164
ThreadWrapper(??) at 0xd09dac48
pthread._pthread_body(??) at 0xd00f0358
```

Additional Resources

- AIX
  - AIX Performance and Tuning Guide (IBM publication SC23-2365-04)
  - Performance Toolkit Version 1.2 and 2 for AIX (SC23-2625-04)
- HP-UX
  - Tunable Kernel Parameters (HP-UX 11i) Hewlett Packard
  - Memory windows in HP-UX 11i Hewlett Packard
  - Streams/UX for the HP9000 Reference Manual Hewlett Packard
- Linux
  - Lotus Domino 6 for Linux (Redbook SG24-6835-00)
- Solaris
  - Sun Performance and Tuning, Java and the Internet, Adrian Cockcroft and Richard Pettit
  - Solaris Internals, Jim Mauro and Richard McDougall
  - Lotus Domino R5 for Sun Solaris 8, IBM Technical Support Organisation (Redbook SG24-5969-01)

Additional Resources

- Domino
  - Maximizing Application and Server Performance Whitepaper
  - Inside Notes : The architecture of Notes and the Domino Server
- Web
  - http://www.redbooks.ibm.com
  - http://www.lotus.com/ldd (Lotus Developer Domain aka Notes.Net)
- Domino Platform Homepages
  - http://www.ibm.com/eserver/pseries/solutions/eb/eb_email.html/
  - http://www.lotus.com/dominosolaris
  - http://www.sun.com/lotus
  - http://www.lotus.com/dominosolaris
  - http://www.lotus.com/dominohpux

IBM.

Special thanks

- We want to give special thanks & credit to a number of People who helped collecting and reviewing this information.
- Carolynn McCarthy, Ivan Dell'Era, Eddy Bell, Kenneth Brunsen (Iris), Catherine Stone, Cameron Hildebran, James Grigsby (IBM), Eric Sosman, Daniel Edwin (Sun), George Chiesa (dotNSF), Daniel Cisowski
- Razeyah Stephen's Iris performance team and a number of other people from the differents part of the IBM/Lotus Support organisation and development...

Questions & Answers

- Ask questions now or find us after the session here or at the booth/lab
- "Meet the Developers" (Iris Labs) to meet with Eddy
- daniel.nashed@nashcom.de
  - http://www.nashcom.de/ls2003
  - at dotNSF, inc. booth #116
- marc.luescher@ch.ibm.com
  - at Sun Microsystems booth
- Please fill out your evaluations ...