

Lotus knows.

Smarter software for a Smarter Planet.

JMP107 New Flavors arrived: “Operating System Soup” with New Recipes - 2010

Daniel Nashed – CEO Nash!COM

Marc Luescher – IBM Technical Support and Geek

Agenda

- Introduction
- Basic Soup : Current Technical Questions and Challenges
- 32 or 64 Bit Soup Flavours
- SAN/NAS or what kind of ingredients do you need ?
- Windows , Linux®, Unix® or what else ?
- Admin or Developer Soup : What do you need to know about different platform options ?
- Migration Soup : what can you expect during a migration ?
- Q&A



- **Daniel Nashed**

- Nash!Com - IBM/Lotus Advanced Business Partner/ISV
- Member of The Penumbra group
 - an international consortium of selected Business Partners pooling their talent and resources
- focused on Cross-Platform C-API, Domino® Infrastructure, Administration, Integration and Troubleshooting
- Platform Focus: Win32/Win64, Linux®, AIX® and Solaris®
- Regular speaker at International Conferences
- DNUG Enthusiast IBM Lotus® Domino Infrastructure

- **nsh@nashcom.de**

- **<http://www.nashcom.de>**

- **Marc Luescher**

- IBM Accelerated Value Professional aka PSM for two large Swiss / American customers
- Technical expert for large Unix, Linux based server architectures and client deployment(s)
- Strong focus on large Domino environments running on Solaris®, AIX®, Linux® and Windows® farms
- Specialist for virtual platforms on UNIX® and Microsoft Windows®
- Regular speaker at Local and International Conferences



- **marc.luescher@ch.ibm.com**

Current technical questions (1/4)

- Memory of Domino Server
 - I was told that Domino can only use 2 G of memory
 - The system guys tell me that our servers are under-utilized but our user complain of performance as every year.
 - What will I gain when using Domino 64 Bit on a 64 Bit OS
- Disk I/O throughput of target system
 - Our storage architects found that very cheap NAS solution and want to use it with Domino. Is this a wise idea or should I continue to use local disks as I did the last 10 years?
 - How much throughput does a Domino system require?
- Backup, Restore Constraints and Service Levels
 - In our SLA we promise mailbox restore in less than 4 hours. We just realized that doing this for a 62 GB mail file did not work out?
 - What will be the performance impact if one of my two Domino clusters will go down. Will this have an impact on our SLA and what can shall I do?
 - Is an active / active configuration the better solution or are options? What are my risks?
 - Is Snapshot technology a valid option for Backups?

Current technical questions (2/4)

- Number of 3rd party tools and kind
 - I need to host 5000 Blackberry users. Can they all fit onto one Domino server running VMWare?
 - My company bought that highly sophisticated archiving solution which needs add-ins. What does this mean?
- Availability of server
 - The OS architects tell me that UNIX has the best system availability with a lot of “9” and that all other OS are less. What does this mean for running Domino?
 - Do I still need a cluster or should I use OS built in features?
- Costs
 - Management told me to reduce cost for the Domino server farms. I was tasked to run a large consolidation effort and to change platform?
 - What are my options?
 - How can I reduce operation costs and what are the biggest cost drivers?
 - How can I charge back real system usage to my end users?
 - What platform is the cheapest to run?

Current technical questions (3/4)

- Number of recommended Domino partitions
 - How many partitions should be hosted on a single operating system instance to make best use of the available resources?
 - What did those architects recommend last time we spoke to them?
- Purpose of Domino server
 - Is the server used for mailbox hosting or application hosting
 - Is the server used for Internet protocol access e.g. HTTP and LDAP
 - Is the server used for 3rd party applications e.g. BES
- User Access mode (on server or replication)
 - Are users working on the server directly to access mail and applications or are they using a local replicas of the mail file which is the recommended option

Current technical questions (4/4)

- Internal disks
 - 10K or 15K that is the question? Or can I still use those old internal disks nobody else wants?
 - Can Domino Transactional Logs be hosted locally?
 - How many OS partitions can be shared on a local disk?
- SAN/NAS
 - SAN support is currently very common.
What should I ask from the SAN guys for my future Domino architecture?
We do not seem to speak the same language?
 - How large should my file systems be and is RAID5 the best option?
 - What do I need to know about DAOS and my SAN?
 - Is NAS an option and is this supported?
- Network
 - How many NICs do I really need and what is the recommended speed?
 - How many connections can my system handle?

32 or 64 Bit Soup

- 32 Bit Domino on 32 Bit OS
- 32 Bit Domino on 64 Bit OS
- 64 Bit native Domino on 64 Bit OS

32 Bit / 64 Bit Basics

- Without “tricks” a 32 Bit OS can only address at most 3-4 GB Memory
 - That's why 32 Bit Process can at most allocate 4 GB Memory
 - Each platform has different limits – see details next slide
- Domino uses 2 memory types
 - Local Memory – separate for each process
 - Shared Memory – shared between all Domino processes
 - Prominent Example: NSF Buffer Pool (usually 512 – 800 MB)
 - All Shared Memory is mapped to all processes
- Address Space Limits
 - 32 Bit = 2^{32} = 4 GB
 - 64 Bit = 2^{64} = 18,446,744,073,709,551,615 = 18.45 Exabytes
 - That's more than we will “ever” need ...

Domino 32bit on a 64Bit Operating System

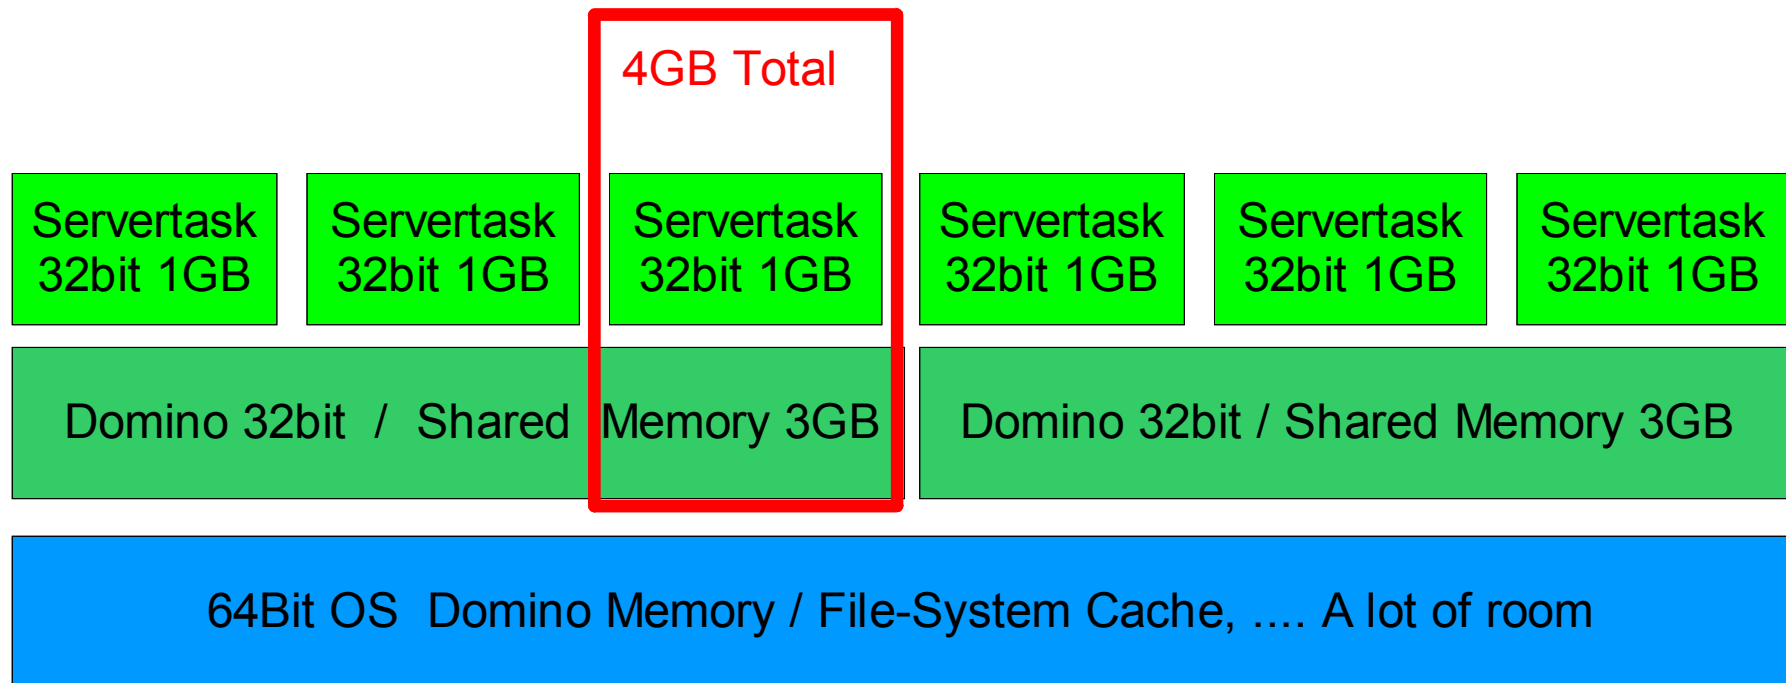
- Domino 8.x - 32Bit on AIX (always 64Bit mode)
 - Up to 9 segments a 256MB = 2,3 GB Shared Memory
 - Up to 2 segments a 256MB = 512 MB Local Memory per Process/Task
- Domino 8.x – 32Bit on Solaris (always 64Bit mode)
 - Up to 3 GB Shared Memory
 - Up to 1 GB Local Memory per Process/Task
- Domino 8.x - 32Bit on Linux (32Bit mode)
 - Up to 3 GB Shared Memory
 - Up to 1 GB local Memory per Process/Task
- Domino 8.x – 32Bit on Windows (32Bit mode)
 - Up to 2 GB Local + Shared Memory absolute limit
 - Windows splits the Memory in 2 GB for System + 2 GB for Applications
 - So Windows 32Bit Memory is quite limited

Domino 32bit on a 64Bit Operating System

- Supported Platforms
 - since Domino 6 on AIX64 + Solaris64
 - since Domino 7.0.1 Linux (SLES 9/10 + REHL 4/5 64Bit)
 - since Domino 7.0.2 on Win2003 R2 64Bit
 - Since Domino 8.5 on Windows 2008 R2 64 Bit
- 64Bit OS allows multiple partitions with dedicated 32Bit address space per process
 - Very good for consolidation of multiple Domino servers
- Changed Limits
 - Win2003 64Bit R2 or SP2
 - 4 GB Total Memory per Partition (Local + Shared Memory)
- The remaining memory is used by the OS 64Bit
 - File caching, buffers, internal resources
- Running 32Bit Domino on Linux64/Win64 gives you already most of the performance and scalability benefits

Domino 32bit on a 64Bit Operating System

- Total Memory per Process is 32Bit = 4 GB
- Router / HTTP uses most local process memory
- NSF Buffer Pool is the biggest Shared Memory block (512 MB)



Native Domino 64bit

- Supported Platforms
 - since Domino 8.0.1 on AIX64
 - since Domino 8.0.1 on Win2003 R2 64Bit
- Domino 64Bit Application allows almost unlimited memory
 - On AIX still 256MB segments but not limited to 16 segments
 - But Domino 8.0.x / 8.5 64Bit does only allow more memory but is not specially optimized for larger memory
 - Tests have not show much performance improvement – only scalability is improved
<http://www.ibm.com/developer works/lotus/library/domino8-64bit/>
- Constraints
 - All add-on applications need to be ported to 64Bit
 - Also all Lotus Script API calls need to be ported
 - Each 64Bit port is a separate platform which needs separate attention and testing!
 - There have been issues with the first 64Bit C-API toolkit version

How much memory is needed ?

- 4 GB Memory is sufficient for most Domino Servers
 - Domino is well optimized for this amount of memory
 - Only very very large servers hosting many users may need more memory
 - NSF Buffer Pool, Router, HTTP need most resources
- The limit is per partition on a 64Bit OS!
 - Multiple partitions can use up to 4 GB Memory each
- So probably Domino 32bit on a 64Bit is what best fits you
 - For Linux there are no special requirements
 - But on Win64 there are some important details to take care about

32 Bit Domino on 64 Bit Windows

- You need Win2003 R2 or SP2 64Bit
 - Domino needs a special Win64 API call to limit the Windows OS Level Cache
 - Domino 8 ships a tool “cacheset” to set the cache size
 - See TN #1270452 for details
 - Still not automatically invoked in Domino 8.0.2 / 8.5
 - You might to rename cacheset.exe to ncacheset.exe (fixed in 8.5.1)
 - Default is 30% of memory
 - Can be tuned via notes.ini MEM_FSCachePercentMem
- Domino 8.5 also supports Windows 2008 32Bit and 64Bit

32 Bit Domino on Windows 64 over 2 GB

- Add-On Applications need to be recompiled and linked with Visual Studio .Net 2003 with link flag /LARGEADDRESSAWARE
 - Else if any process (servtask) exceeds 2 GB limit (local + shared memory) it would crash!
 - Check via dumpbin
 - Example: dumpbin /headers nshdbcat.exe

Result:

FILE HEADER VALUES

Application can handle large (>2GB) addresses

- Default Max Shared Memory is set to 2GB
 - notes.ini ConstrainedSHMSizeMB=3072 gives you a maximum of 3GB
 - You should still use

Different Types of Storage

- Local Disks, Direct Attached Storage
 - SCSI and SAS.
 - Most servers come with hardware RAID support
 - Can be very cost affective
- Centralized Storage
 - Central Appliance/Box (EMC, NetApps, Hitachi, Sun, ...) connected to many machines
 - SAN or NAS storage has different level of integration
- Today the line between NAS and SAN is blur
 - iSCSI also uses network cards instead of expensive FC cards




Storage soup

- Local Disks
- SAN Disks
- NAS Disks

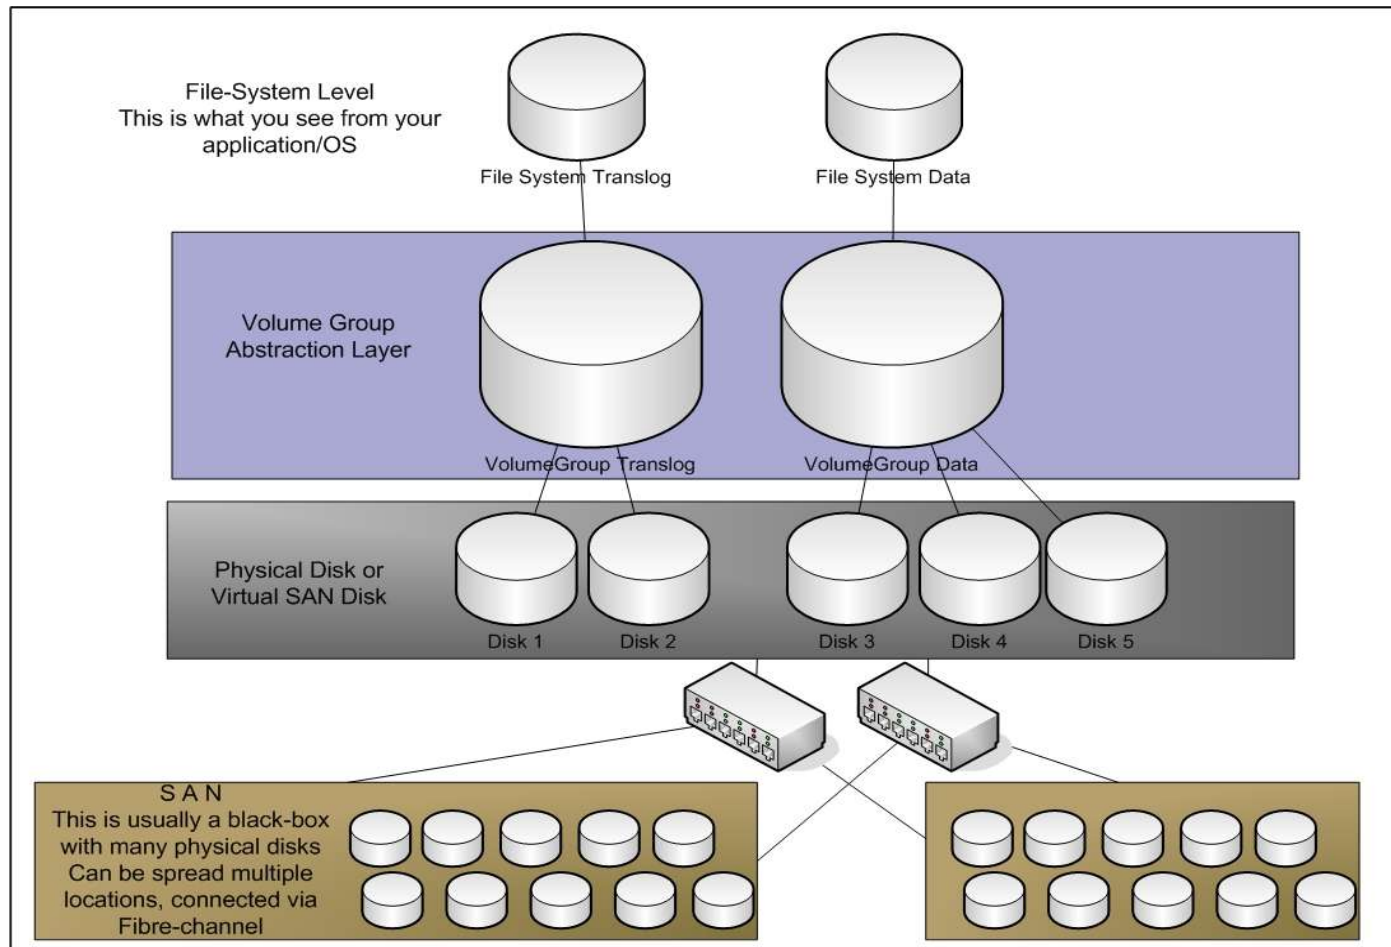
Let's get started ...

SAN / NAS

- SAN – Storage Area Network
 - Dedicated Hardware (expensive FibreChannel cards)
 - Separated into logical units (LUNs) usually presented as virtual hard-drives to the OS
- NAS – Network Attached Storage
 - File-System Level Integration into OS – not on virtual disk level
 - Sun Network File System (NFS) and Common Internet File System (CIFS)



SAN – Storage Area Networks



SAN / NAS

- For NAS and SAN you have to apply best practices and use the latest possible drivers firmware etc for all components to get the right performance
 - Same is true often for local disks (RAID controller firmware, disk firmware, etc)
- Similar rules for performance apply for SAN/NAS Disks
 - You should use separate LUNs for DATA and Translog
 - The more fast physical disk you have in your SAN RAID Set the better your performance
 - Domino needs very fast storage – Enterprise Level SAN for good response times
 - Domino is very demanding in small I/O operations (many I/Os per second)

Local Disks

- Current Server hardware (2U size) ships with RAID controllers with cache and up to 8 bays to add disks with up to 500 GB each
 - If that is not sufficient you can add enclosures and additional RAID controllers
 - No need for an expensive FibreChannel Adaptor
- Local Disks are the same physical disks you would expect on a SAN
 - But often certified/supported SAN disks are more expensive
 - Same rules for local and server disks apply – including the same RAID Sets
- Current Disk Speed is either 10K or 15K
 - The more disks you have in a RAID set the better your performance is
 - Also e.g. RAID 10 is faster than RAID5 but needs more disks
- In many cases local disks can be more cost effective with dedicated machines
 - But usually you cannot fully take benefit on the server hardware with a single Domino partition

Key Requirements to Consider

- All Storage Devices have to be high available
 - You need RAID for every device
 - The cost to recover is much higher than the additional disk and a disk failure is not unlikely!
- Performance of Disks
 - RAID Level (RAID5 -> RAID 1 -> RAID 10)
 - Fast Disks with 15k instead of 10k, More smaller disks
- Performance Disk Connectivity
 - FibreChannel (FC) 1Gbit is too slow (4/8 Gbit)
 - You should ensure you have dual attached cards (if possible in load-balancing mode)
 - iSCSI (Network Cards have 1 Gbit usually) --> New 10 Gbit cards might be an option
- High Availability
 - Do we need disk mirroring with a different location?
 - Or is Domino Clustering used?

General File System Considerations

- Mis-configured I/O Subsystem leads to performance bottlenecks!
 - Use RAID1/RAID10 instead of RAID5 at for Transaction Log
 - Have separate disks for translog, view-rebuild, data
 - Put different file-systems on different controllers/ multi channel controllers for large servers
- Have always at least 20-30% free space on file-systems for data
- SAN (Storage Area Network) is faster than local disks in most cases
 - It is recommended to have different file-systems in different SAN volumes
- Use LVM (Logical Volume Manager)
 - Acts as a layer between the physical disk and file-systems
 - Enables you to resize file-systems and add more disks/SAN space
 - Depending on SAN use multiple virtual disk per Volume Group
- Limit the use of UNIX symbolic links and even more dir-links
 - Use mounted file-systems instead – also on Windows

Mounted File Systems on Windows

- Mount Points are called “Junctions” on Windows
- Generally supported but has some limitations
- Known Limitations (current state as of Domino 8.5.x)
 - Junctions are not recognized by Domino Stats
 - Free space check for accelerated replica creation uses free space of main drive (drive letter)
 - Unmount Problem when deleting the last file in mounted directory has been solved
 - NO need for a special notes.ini parameter to disable deletion of empty folders any more
 - Some 3rd party backup applications do not support junctions

File System Recommendation “NOTESDATA”

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

File System Recommendation “TRANSLOG”

- Size: 5 GB for Circular Translog
 - For archive-style translog size highly depends on your data
 - For linear style translog size should be in the range of 8-20 GB
- Recommended: RAID1
- Allocation Size: 4K
 - Use larger disk block size and matching Stripe size
- Access Pattern:
 - TX logging writes fixed sequential 4k blocks
 - Transaction log files are opened in a synchronous mode
 - OS file system cache is not used
 - NVRAM cache in the disk subsystem helps - make the size of the largest write to the transaction log files.
- On SAN use 4Gbit Fibre Channel rather than 1 GB.
- For local drives have separate controller or at least dedicated channel

File System Recommendation “DAOS”

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

Summary Storage

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

What platform fits best for Domino ?

- There is no best platform
- 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

Three groups of operating systems

- **Microsoft Windows**
 - Most common operating system for Domino
 - Works fine for small to “medium” sized servers in 32bit
 - With current 64bit OS much more scalable
- **Unix**
 - 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 (aka AS/400), zLinux (Linux on zSeries)
 - Mostly used by customers with existing iSeries and zSeries environments
 - Very scalable and high available

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
 - Customers migrating from W32 have problems with mixed case file names
 - Best Practice: keep all directories and file names ASCII lowercase!
- 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

Domino on Microsoft Windows - Pro

- Domino on Windows is most commonly used server platform
 - But Linux is fast growing
- Performance it good for medium sized environments
 - Up to 1000-1200 users
- 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, ...
 - Those tools are developed on Windows, best tested on Windows and most used on Windows...

Domino on Microsoft Windows - Contra

- Few OS level tuning options
 - Scalability is limited for older Windows releases, File-System scalability is limited
 - NTFS has problems with large file-systems > 500GB (e.g. Heavy file fragmentation)
- Fixpack dependences
 - You can run into interoperability issues when applying required OS security patches, etc
- Performance Monitoring is not comparable without 3rd party tools
- Partitioning is supported but not recommended from Best Practices point of view
 - Domino uses the system Account (multiple partitions share same unrestricted user account)
 - You can only one Domino release at the same time
- Licence Costs on user side for OS (CAL = Client Access License)
 - Unless you have the right Office version licensed
- Windows is “less secure” than Unix/Linux

Domino on Linux - Pro

- Open Standards Based Platform
 - Huge dedicated Linux Developer community
 - Many excellent "free" tools (e.g. OS level admen tools, ...)
- Scalability becomes comparable to AIX and Solaris since Domino 7
 - 64Bit OS support since D7.0.1
 - Fully uses x86 Hardware (hyper-threading/multi-threading, no idle memory, etc)
- “Rock Solid Security”
 - Used by may customers in DMZ environments, Built from beginning with Security Design
- Linux® is the fastest-growing server OS
- No Licence Costs on user side for OS (CAL = Client Access License)

Domino on Linux - Contra

- Needs an Enterprise Server Linux Distribution with current Service Packs
 - SuSE Enterprise Server 10
 - RedHat Enterprise Server 5
 - But server OS licence costs are usually lower than for Windows servers
- You need Linux skills to get it up and running
 - But it is getting more and more easier
- Needs OS and Domino Tuning to be scalable
 - Has become much easier in the current releases
- Not all 3rd party products are available for all Domino releases
- Case-Sensitivity of the File-system

Domino on Sun Solaris - Pro

- Solaris 9 or 10 can be used with or without a Solaris Zones setup to make better use of available hardware
 - Solaris is a robust 64 Bit operating system
 - Excellent support for different file system types and sizes

Domino on Sun Solaris - Contra

- You need Solaris skills to get it up and running
 - You should only use that platform if you have those skills in house
- Needs OS and Domino Tuning to be scalable
 - Has become much easier in the current releases
- Not all 3rd party products are available for all Domino releases
- Hardware and Software maintenance is more expensive compared to a Intel® based system
- Case-Sensitivity of the File-system

Domino on AIX - Pro

- AIX 5.3 and AIX 6.1 further improve the use of the latest CPU families of pSeries
 - AIX is a robust 64 Bit operating system
 - Excellent support for different file system types and sizes
 - Support for LPAR's and Micro Partitions
- Domino is available as 64 Bit version to leverage some of the benefits of a 64 Bit operating system

Domino on AIX - Contra

- You need AIX skills to get it up and running
 - You should only use that platform if you have those skills in house
- Needs OS and Domino Tuning to be scalable
 - Has become much easier in the current releases as best practices are available
- Not all 3rd party products are available for all Domino releases and platforms
- Segmented memory model with some smaller limitations
- Currently under limited availability program for Domino 32 Bit

Domino on iSeries - Pro

- Domino on iSeries is rock solid and very scalable
 - 128 Bit OS
 - 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
 - Very good I/O scalability and I/O sub-systems
- Rock solid security
 - e.g. OS system components are signed
- Each LPAR is controlled complete separate

Domino on 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 more expensive than on Intel platform
- 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

Domino on zLinux - Pro

- Brings together mainframe and Linux advantages
- High availability and scalability combined with flexibility and up to date technology
- What Linux brings to zSeries
 - Large portfolio of applications, tools and enablers
 - The ability to enhance an existing zSeries infrastructure
 - Large numbers of trained programmers and administrators
- What zSeries brings to Linux
 - The most reliable hardware available anywhere.
 - Best of breed Virtualization Technology.
 - Designed to support multiple diverse workloads / Complete workload isolation
 - Unmatched scalability
 - Simplified systems management
 - The ability to run many Linux servers on a single hardware platform
 - Green IT - Less machines, Less power consumption, ...

Domino on zLinux - Contra

- -Only makes sense for customers running zSeries as a strategic platform
 - Needs special OS level and hardware skills
- Not many add-on tools available
 - Porting is a 1:1 port from Intel Linux but you need a zSeries machine for porting and testing
- You should only run native zLinux on the machines running Domino to benefit from discounted hardware IFLs (CPUs only allowed to run Linux)

UNIX/Linux File System Differences

- Unix & Linux are designed from scratch to support multi-user, multi-tasking environments!
 - A lot of Linux/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
 - Dedicated system accounts which run in user context separating partitions completely
- Some Windows specific functionality like OLE, DDE, ... is not supported on Unix
- File-systems and path names “look” different (details next page)

Best Practices - 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_TcplpAddress=0,192.168.1.42:1352
- Don't use too many Domino partitions on a single OS, Best Practices is 3-4 for UNIX per LPAR

Hardware Example HP DL380 G5

- Affordable mid size hardware – 2U size with local Disks
 - Up to 2 Intel Xeon 5X00 sequence Dual-Core or Quad-Core processors with Intel VT technology
 - Up to 64GB of 667MHz DDR2 Fully-buffered DIMMs; with interleaving, mirrored memory and online spare capability
 - 4 available PCI-Express expansion slots standard with optional PCi-X/PCI-X Express riser
- RAID Controllers
 - HP Smart Array P400 512MB with BBWC (Performance models)
 - HP Smart Array P400 with 256MB read cache (Base models)
 - HP Smart Array E200 with 64MB read cache (Entry model)
 - Expandable storage for support of up to eight (8) SAS or SATA small form factor drives



Customer Linux Example Two Locations - 6 Boxes

- Spread multiple Domino Partitions (DPARs)
On multiple physical machines

- Each machine has one counter part on a different physical box

- 3 Partitions each


- Balance two busy and one lower profile DPAR

- Leverage Domino Clustering

- No SAN mirroring!



- 2 Quad-Core CPUs
- 16 GB RAM
- SAN disk for data
- Local Disks for TL



Different Types of Virtualization

- Desktop Virtualization
 - Run multiple Virtual Machines on your Desktop PC
- Application Virtualization
 - Run an application in a “cloud”
- Presentation Virtualization
 - Run an application on a server and have the presentation layer on the client side
 - E.g. Terminal Services, Citrix
- Profile Virtualization
 - User Profile / User Data Roaming etc...
- Server Virtualization
 - Run multiple Virtual Machines with a guest operating system on a single hardware

Short History of Virtualization Techniques

- Virtualization is around for a long time
- Mainframe virtualization (VM/390), zSeries LPAR technology
- IBM pSeries LPAR technology
 - Hardware Virtualization
- Deskview from Quarterdesk (same company invented Quemm)
 - First desktop virtualization products
- Virtual PC 2004
 - Acquired by Microsoft turned into Virtual PC 2007 and Virtual Server 2005 R2 SP1
 - New Product Windows 2008 R2 Server 64bit with Hyper-V
- VMware
 - Has been around for a while with Server and Desktop products
 - Entry level versions are free
 - ESX Server is most commonly used enterprise virtualization product

Drivers for Virtualization

- Cost reduction and Consolidation
- Current Intel/AMD hardware became much more scalable
 - Single Domino partition cannot leverage mid size hardware
 - Consolidation of smaller machines into one bigger hardware
 - Hardware Virtualization Support: Intel VT and AMD-V
- Less custom 3rd party operating system drivers needed
 - Reduces the risk of driver dependencies and improves stability – also on desktop ;-)
- More flexibility in management and providing new servers
 - Server Cloning from image, Moving servers, etc.
- Resource balancing and better use of hardware resources
- Easier disaster recovery (DR) with enterprise virtualization products
 - e.g. VMware Motion
 - Take care for all those kind of options very strict SAN/NAS requirements apply

Best Practices for Virtualization

- Allocate fixed Memory, Fix Disk Size
 - Do not share volumes
- Either set fixed 1 CPU or 2 to n CPUs
 - Guest Operating systems have different kernels for single and multiple CPUs
 - With low CPU demand configuring one fixed CPU should give best performance
- Use VMFS Volumes instead of directly allocated physical volumes
- Enable Hyperthreading on Host machine
- If not needed turn Vmotion off (causes overhead)
 - Important servers should be clustered instead
- Small machine example
 - 32 Bit
 - 1 CPU
 - Max 3 GB RAM (more is not effective on a 32Bit machine)

IBM/Lotus Support for VMware

- Domino is supported on VMware
 - Best effort will be made to help in case of problems
 - In case of performance problems you get only very limited help
 - TN #1106182 VMware product support information for IBM Lotus Domino-based server products
 - No direct support statement for other virtualization solutions
 - But WebSphere has a support statement for other solutions like XEN, Microsoft etc
- Generally it is assumed that there are no functional differences
 - But in special cases you have to re-create a problem on a not virtualized environment
- There have been some performance issues with high CPU peaks
 - Have been fixed with current versions of VMware, Guest Additions Operating System versions

Licensing

- IBM Subcapacity Licensing

- Before Subcapacity Licensing was introduced you had to license based on the possible CPUs in a machine
- Now you have basically to licence for the configured CPUs per machine
- Available for all Lotus products now with the original product numbers
- Details: <http://www.ibm.com/software/lotus/passportadvantage/subcaplicensing.html>
 - One main PDF with multiple associated PDFs
- ILMT Tool to calculate licences (IBM Licence Metric Tool)
 - <http://www.ibm.com/software/tivoli/products/license-metric-tool/>

- Microsoft Licensing

- Microsoft offers also flexible, fair licensing for virtual environments
- Details: <http://www.microsoft.com/licensing/about-licensing/virtualization.aspx>

Which kind of servers should be virtualized?

- VMware works best with Windows
 - For Linux partitioning as shown before fits best for Linux
- All smaller server types
 - Small Mail Servers (100 – 200 users), Applications servers with lower I/O
- Sametime, Blackberry, Quickr, SMTP gateways, ...
- Development Servers, Test Servers, ...
- Admin Server (with high availability option e.g. Vmotion)
- The less resources a Domino Server uses the more likely it makes sense to virtualize the server

Clustering vs. Virtualization High Availability

- Standard High Availability Solutions only work in a failover mode and protect you against hardware failures or disasters
 - Needs a mirrored SAN in two location and uses the double disk space anyway
 - Idle resources, Costs extra Licences
- Domino Clustering
 - Active Clustering, Client is aware of the cluster and can failover if needed
 - Failover is on database level – e.g. You can take a database off-line
 - You are protected against much more issues
 - Server crash
 - Database corruptions usually do not replicate
 - You can take down one server for maintenance
 - You can balance resources on both cluster mates
 - Clustering is not limited to two servers
 - Clustering is included in the Domino Server license

Memory Overhead ESX 3.0 for each VM

Virtual CPUs	Memory	Overhead for a 32-bit virtual machine	Overhead for a 64-bit virtual machine
1	1 GB	84 MB (8%)	180 MB (18%)
1	8 GB	139 MB (1%)	236 MB (2%)
1	16 GB	203 MB (1%)	300 MB (1,5%)
2	1 GB	101 MB (10%)	300 MB (30%)
2	8 GB	221 MB (3%)	413 MB (5%)
2	16 GB	349 MB (2%)	541 MB (3%)
4	1 GB	141 MB (14%)	523 MB (52%)
4	8 GB	222 MB (3%)	605 MB (7%)
4	16 GB	350 MB (2%)	734 MB (4%)

- Source: http://www.vmware.com/pdf/vi_performance_tuning.pdf

Virtualization Tips

- In case you install productive servers on Linux check “Time Keeping”
 - www.vmware.com/pdf/vmware_timekeeping.pdf
 - Very detailed information about timer implementation
 - Disable NTP Services etc. Only use Guest Addition Time Sync
- Use current ESX Server releases and also guest OS
 - E.g ESX 3 has better support for Linux multi-threading (NPTL threads)
- Ensure you run the current Guest Addition Software
- Apply best practices on all levels
 - Your environment is so stable and performs as your weakest component
- You need an VMware expert tuning your VMware environment
 - Good starting point for VMware http://www.vmware.com/pdf/vi_performance_tuning.pdf
- Domino tuning for virtualized and physical machines is basically the same

Section 2

- Admin Soup
- Interfaces to OS people an admin needs eg backup, monitoring, root access, scripting, trouble shooting
- Developer Soup
- Interfaces to OS people a developer needs and what tools etc he needs to be present

System Interfaces Backup/Restore

- Different vendors have different tools
- Choose a tool which allows you to backup the Domino transactional logs using an Add-In task
- Backup and Restore performance usually is depends on :
 - LAN interfaces (yes Gigabit and dedicated is recommended)
 - Number of parallel streams allocated to tools (don't go crazy)
 - File-system size and IO during backup
 - Technology and architecture of backup system used

System Interfaces Defragmentation Tools

- Depending on the file system in use you might need to consider tools to defragment the file system
- Windows based systems are the most affected. Defragmentation of file-systems used by Domino will improve :
 - Backup and Restore times
 - Domino housekeeping task times (compact, updall)
 - Agent Manager Run Times
 - End User Performance

System Interfaces Monitoring

- Make sure you can use the existing monitoring tools like Patrol, Tivoli or others :
 - You need platform statistics to be able to assess CPU and memory load
 - You need a good understanding of disk and network statistics
 - You need real time and historical data
 - You need to define thresholds and warning levels
 - All systems will act different

Application Interfaces 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...)

Domino 3rd party Tools Other Add-In Tasks

- Typical running 3rd party tools on Domino servers are :
 - Backup Interface
 - Archiving Interfaces
 - Anti-Virus Solutions
 - Anti/Spam Solutions
 - User Management Solutions (IDM)
 - All kind additional gateways

AppDev Challenges when Porting to 64bit

- New compiler version and vendor for Win64
 - Switch from Microsoft to Intel compiler
- Pretty straight forward porting for plain C-API applications
 - But specially complex add-on tools need many libs (e.g. Virus scan engines etc)
- New data-types in Notes and all Domino handles need to be changed
 - OS switches to 64bit Handles, Domino stays with 32bit HANDLE
 - New DHANDLE = Domino Handle needs to be used
 - On AIX DHANDLE is now 32bit (instead of 16bit for some handles)
- The cleaner the application is developed, the easier the porting
- Data-Type changes
 - Size of parameters changes, needs detailed checking and casting etc ...

AppDev New Compiler needed on Win64

- New build environment
 - Win2003 R2 64bit + new compilers and Domino 64bit for testing
 - Microsoft 64bit compiler was not “ready”
 - IBM switched to the Intel compiler
 - Microsoft compiler still needs to be installed to get the Windows headers
 - You can install a trail version and keep the headers
- Windows Compiler details listed in C-API toolkit
 - Microsoft (R) C/C++ Optimizing Compiler Version 14.00.40310.41 for AMD64
 - Intel® C++ Compiler for Windows version 9.1
 - Note: Intel® C++ Compiler for Windows version 9.1 is the only supported compiler.
 - It works on top of SDK of Microsoft (R) C/C++ Optimizing Compiler Version 14.00.40310.41 for AMD64. C API 64bits users need to install Microsoft (R) C/C++ Optimizing Compiler Version 14.00.40310.41 for AMD64 as well.
- Compile environment for AIX and zLinux stays the same

AppDev Changes when building the Application

- Notes 8.5 C-API contains 32bit and 64bit sample makefiles
 - First downloadable version was “broken”
- “cmp” directory in C-API toolkit contains original IBM make options for all platforms including 64bit
 - You should check the sample makefiles and the txt file in “cmp” for details
- On Win64 Domino structures are not packed any more
 - Different byte alignment for structures – used the natural 64bit alignment
 - So don't use "/Zp" option on Win64 !!!!
- On AIX no dataseg is used for changing the number of data segments
 - Don't use "-bD:0x20000000" option to set the data segment size.

AppDev

Example Data Type Changes

- Check C-API User Guide for porting tips
“Lessons learned during Domino 64bit development”
- Domino Data-Types (e.g. LONG) stay the same
- OS Data-Types like int, long (on Unix) change
- OS calls like strlen returns size_t which has changed
 - Usually WORD is used as a result in Notes which has not changed
 - Results in many warnings
- C-API from Lotus Script
 - Declares stay the same
 - There is an issue in first 8.x versions – Fixed in 8.0.2 FP2
 - SPR# ADEE7JRLJQ - On Windows 64, API's pull their arguments in 8 byte alignment, while LotusScript was stacking these arguments in 4 byte alignment causing issues with these calls. Fix is to make sure everything is in 8 byte alignment on Windows 64

Migration Soup

- How should the data be moved?
- What do I need to consider afterwards

Migrating to a different Platform

- ODS is the same on all platforms
- a.) Shutdown and copy the files (e.g. Ftp) to the target machine
 - Recommended: run fixup -J -f + updall -R or better compact -D afterwards to ensure consistency
- b.) Or replicate databases / use Adminp to move users to a different server
- In both cases take care about case-sensitivity when moving to Unix/Linux
- Directory structure should not change – it's all relative to the data-dir
- Take care about native code e.g. C-API from Lotus Script or OS calls
 - Needs porting for each platform
- Some special functionality needs porting or might not work
 - Direct access to file-system might change, OLE functionality only on Windows, etc
- Tuning on Domino side is almost the same but OS-level tuning changes
 - See details for each platform at the end of the presentation

Migrating to a virtual platform

- Be careful about requirements for CPU, RAM and IO needs
 - CPU and I/O are most critical resources
 - Storage and I/O requirements are very often misunderstood
- Do not share the same boot or temp disk amongst too many virtual hosts.
 - This will come with a performance penalty
- Do not undersize the network interfaces and be sure to not overcommit those resources and share them for end user access, replication and backup.
 - Backing up a virtual server using virtual interfaces will take much longer and will need system resources
- SAN and NAS connectivity can quickly become a system bottleneck
 - Those interfaces are shared amongst all instances
- Mix busy, medium and low servers on the same host, even when that means that a DEV and TEST server might be running next to a productive instance.

Q&A

- We hope you enjoyed the session ...
- Please fill out your evaluations!
- Questions?
- Presentation Updates
 - <http://www.nashcom.de/lotusphere>
- Contact Daniel
 - nsh@nashcom.de
 - <http://www.nashcom.de>
 - <http://blog.nashcom.de>
- Contact Marc
 - marc.luescher@ch.ibm.com

Legal Disclaimer

© IBM Corporation 2009. All Rights Reserved.

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

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

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.

IBM, the IBM logo, Lotus, Lotus Notes, Notes, Domino, Quickr, Sametime, WebSphere, UC2, PartnerWorld and Lotusphere are trademarks of International Business Machines Corporation in the United States, other countries, or both. Unyte is a trademark of WebDialogs, Inc., in the United States, other countries, or both.

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

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

Intel, Intel Centrino, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

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

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