Planning, Installing, and Configuring the Lotus Domino Server on the RS/6000 SP

IBM Poughkeepsie - RS/6000 SP
Parallel Subsystem Integration Team
Version 1.0 - January 28, 1997

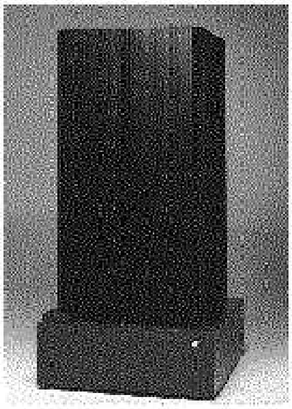

http://www.ibm.com
http://www.lotus.com
RS/6000 SP
IBM

Table of Contents

Special Notices

Next Chapter

RS/6000 | Solutions | Hardware | Software | Support | ReSource | sitemap
IBM | Order | Search | Contact IBM | Help | © | ®

Table of Contents Previous Chapter

Overview - Lotus Domino Server and the RS/6000 SP

What is Lotus Domino Server Powered by Notes?

The Lotus Domino Server Powered by Notes (hereafter referred to as the Lotus Domino Server, or Domino) is an industry leading integrated messaging and groupware product. It is a client/server platform for developing and deploying groupware applications, a platform for highly specialized applications, and an open environment on which independent software developers can integrate their own products.

Groupware is a convergence of three technologies:

- Communication - Email, messaging, and routing;
- Collaboration - includes topics such as shared databases, orderly replication, discussion groups, access control, and Internet access. Collaboration is a framework for exploiting the many ways users need to communicate and share information;
- Coordination - Workflow-like control of tasks and documents, such as orderly review and sign-off processes.

Lotus Domino Server combines three powerful group support technologies - messaging, distributed object store, and a rich application development environment - to form a component architecture that serves as the basis for a large variety of groupware applications. Advanced Domino functionality *enables* you to combine distributed databases with powerful communications, replication, and security layers--all tailor-made for geographically distributed corporations. The ability to integrate Domino applications with critical legacy systems provides for a powerful, flexible, and secure enterprise application development system. From front-end to transaction processing applications, Domino Rich-Text presentation interface brings new looks and function to your old character-based applications. This integration, combined with Dominos' support of Internet standards enables extremely powerful, interactive Web applications that leverage and exploit enterprise information.

Domino is an excellent choice for wide-scale information sharing - other big software developers are frantically working on Notes alternatives, but its two closest rivals, Microsoft's Exchange and Netscape's Collabra Share, still can't match the power of Domino. Domino is the original ***Intra***net, and you can extend this connectivity seamlessly, with the Domino http server, into the ***Inter***net.Domino's recognition in the industry as both a groupware and internet application has been extremely positive - so much that Domino 4.5 finished second on PCWeek's top 10 product list.(1)

For more information on Domino please visit the Web site at *http://www.lotus.com.*

What is the RS/6000 SP?

The IBM RS/6000 Scalable POWERparallel System (RS/6000 SP) is a distributed parallel computing system comprised of up to 512 processing nodes. Each node is a complete RS/6000 workstation, with CPU, disk, and memory, running a separate copy of the AIX operating system. The nodes are all connected via the SP Switch, a non-blocking packet switch which allows high-bandwidth, low-latency Internet Protocol (IP) communication between all 512 nodes in the system. All nodes are connected to a single point of control - the RS/6000 SP Control Workstation - for ease of installation and management. To state it simply: The RS/6000 SP system is a cluster of AIX(2) *servers packaged with a high-speed interconnect between them - in a highly reliable frame - with superior system management capabilities.*

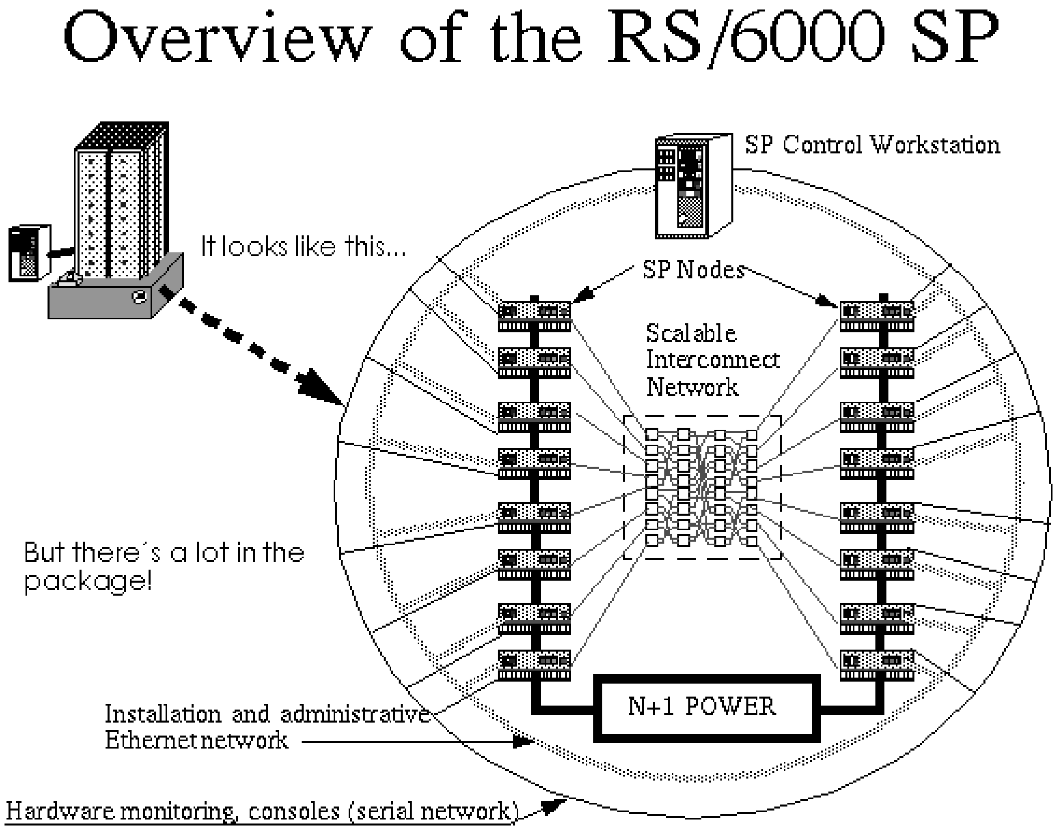

Advantages of the RS/6000 SP system over usual LAN server implementations are

- ***Scalability*** - incremental additions of SP nodes increase capacity and workload while virtually eliminating any performance bottlenecks. Performance scalability is enabled by the SP switch, since the total bandwidth of the switch increases as nodes are added.

- ***Manageability*** - non-disruptive installation, configuration, and upgrades. Management of the collective of SP nodes is from a single point of control.

- ***Reliability and availability*** - key redundancies in SP frame power and SP switch provide a highly reliable hardware infrastructure. IBM’s industry-leading HACMP for AIX *(High Availability Cluster Multi-Processing)* software allows the configuration of the SP system so that failures of node hardware can be recovered.

- ***Existing deployment* -** there are thousands of SP installations worldwide, where the SP is a preferred choice in deploying an enterprise-level server implementation.

- ***World-class service and support* -** IBM is there to guide and support you through its world-wide service and support organization. And you get the benefit of IBM's experience in providing leading-edge technologies such as symmetric Multi-Processing (SMP) nodes in the SP, PowerPC processors, and upgrades to the SP Switch.

For more on the IBM RS/6000 SP system, see the IBM RS/6000 SP system web page at *http://www.austin.ibm.com/software/sp_products/sp2.html.*

The integration of Domino and the RS/6000 SP

With the introduction of the Lotus Domino Server Powered by Notes, customers can now realistically configure systems (or networks of systems) to support thousands of users from a single location. As the capacity of your Domino servers increase, the degree to which your business depends on the availability of the Domino server complex will increase dramatically; bringing with it additional requirements and expectations. The RS/6000 SP system provides a scalable, complete solution for your Enterprise Domino Server.

A typical Domino configuration consists of mail, database, and replication servers. These servers are grouped into centrally administered domains which make up an enterprise-wide Domino environment, as pictured below.

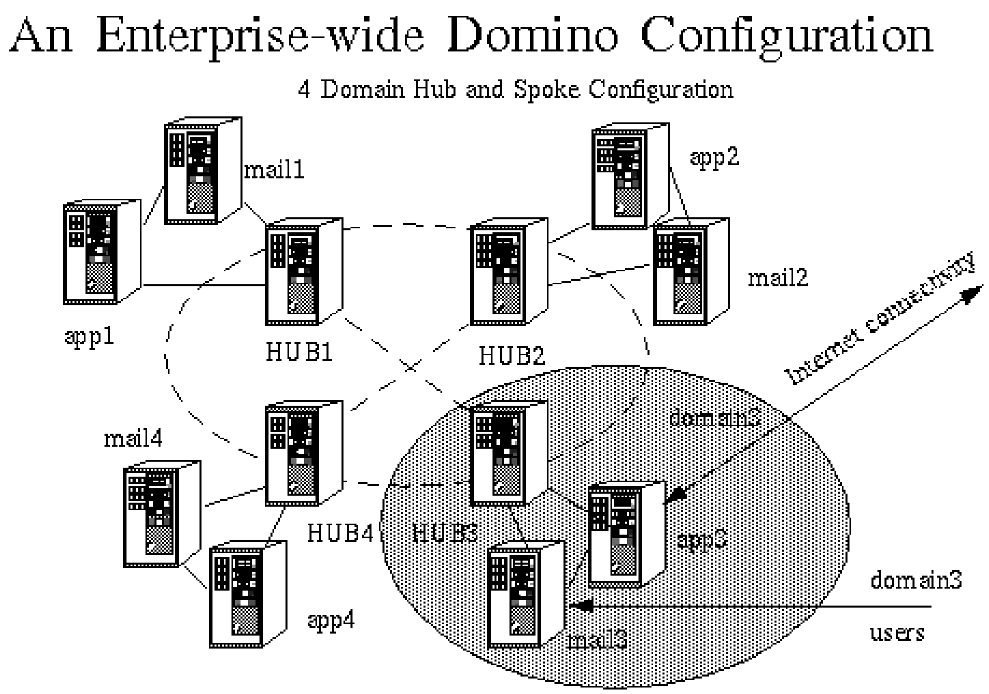

The picture illustrates a hub-and-spoke topology, which is common in large sites and is the most adaptable for system growth and change. Within each domain, a dedicated hub or multiple hub servers are used to replicate information and route mail across domains. Within each domain, there are typically three types of servers: one type dedicated to data replication and/or mail routing; another as user mail servers; and the third for serving data and applications. Organization of the servers within a domain is shown in the following picture.

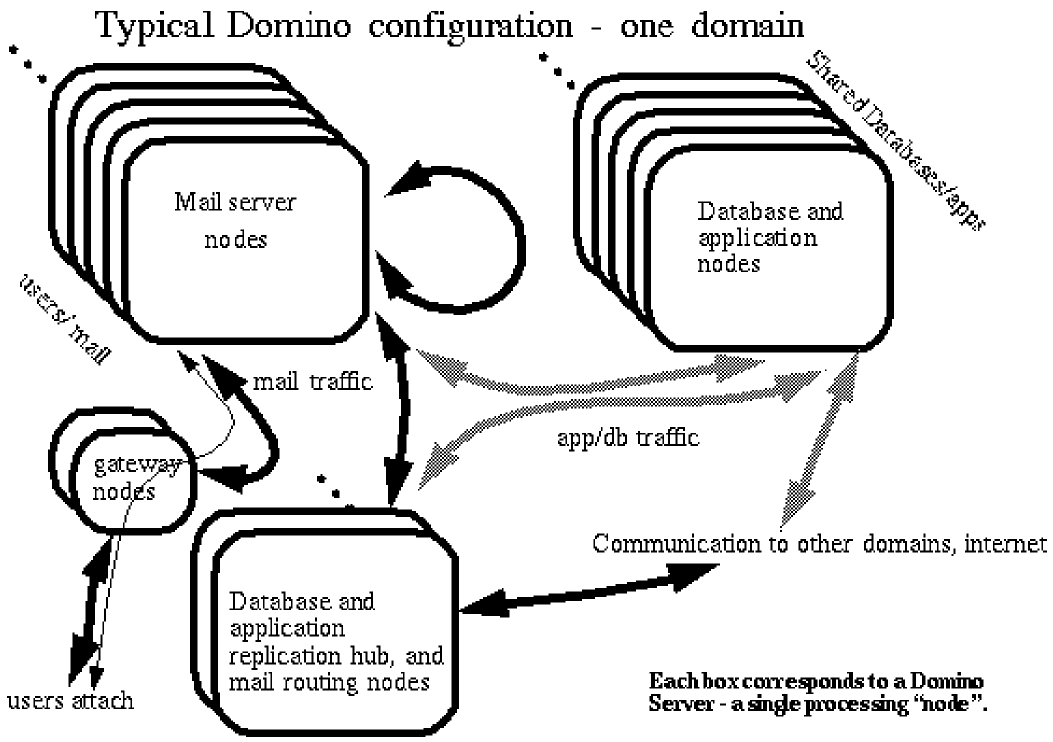

Database and application server nodes provide the shared database/collaborative applications for the Domino environment, **which includes internet and intranet connectivity**. Domino users are registered across the mail server nodes, which serve as a home base for users, physically holding their mail and routing mail to its destination. The replication server, or "hub", replicates and receives database/application updates between servers inside and outside this domain. In large Domino configurations, there may be dedicated mail routing hubs as well as database replication hubs, relieving mail servers from this task.

Many questions arise when planning a large configuration of Domino servers:

- As I add more clients and servers, can my networks handle the increase in network traffic?
- As I add more servers, how many administrators must I add to manage them? Can I manage the collective from one place?
- What if a server goes down? Can I detect it? Can I recover from it? How long does it take and what is the impact to my business?

- How do I integrate this configuration with the rest of my enterprise, and the internet?

Some readers may question the use of dedicated servers for the mail, database, and replication functions - after all, any one server could be configured to do all of these functions. However, in a large enterprise configuration, the ability to incrementally add function - applications, users, or availability of certain users or applications ***as needs change*** must be addressed in an orderly and organized way. For a large organization, the configuration pictured is best suited to this model as it compartmentalizes functions such that they can each be incrementally configured and administered.

This document will detail the planning, installation, and configuration of the IBM RS/6000 SP system as a Domino Enterprise server.

Footnotes

(1) PC WEEK, December 16,1996
(2) AIX: IBM's UNIX operating system

Table of Contents Next Chapter

Table of Contents Previous Chapter

Planning a Domino Server Configuration on the RS/6000 SP

Model for the Domino/SP configuration

The following picture illustrates the SP Domino configuration. This section will describe how to plan for implementing it depending on your needs.

Domino on SP - Logical Model Showing External Connectivity

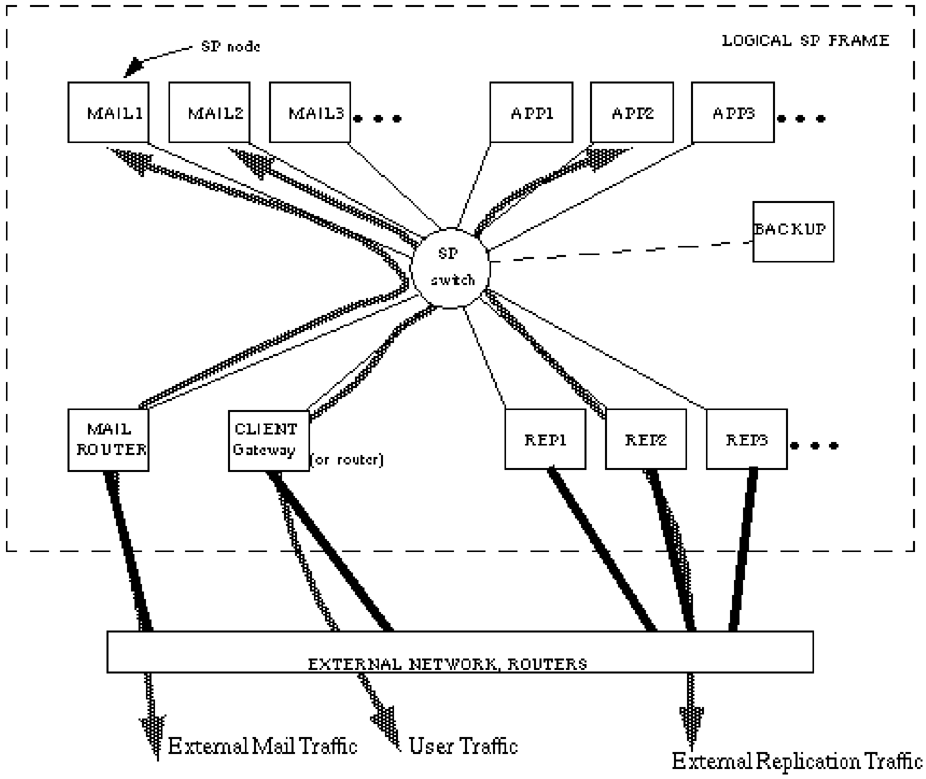

In the picture, SP nodes, rectangles in the picture, are one of the following types indicating their Domino server function;

- A set of mail servers where users are registered;
- A designated mail routing node or nodes;
- A set of application servers where shared databases and applications reside - note that these can include internet servers;
- One or more replication server hubs, for transferring copies of application databases to remote servers; and
- Additional nodes if required to provide network connectivity for the entire configuration. These nodes do not perform Domino function but simplify the system and reduce the amounts of network hardware required on other nodes in the system.

- Backup Servers if required;

This model will serve as the basis for planning, installing and configuring the SP in this document.

Planning your basic Domino configuration

These are the questions you want to ask, when planning a Domino configuration on the SP.

- What types of SP nodes will I be using?
- How many users, and how much mail disk space per user?
- How many applications/ databases, their sizes, and how many users will access?
- How will I be replicating to other sites?
- What are my availability requirements, for user access, applications, mail, etc?
- What will be my SP administration point or points of control?
- How will I backup and restore the configuration?

Planning SP Domino Servers - SP and AIX specifics

Planning SP Administration

The control workstation is generally regarded as the central point of the SP configuration, for SP configuration data (SDR) and installation images, etc. If you are planning a multi-rack SP it is recommended that you use a node in each rack as an install server (for AIX and applications). Also, you may be planning using other enterprise monitoring tools such as NetView and NotesView, as well as backup tools like ADSM.

Depending on the size of your installation and what other administration tools you choose to use, you may want to pick a node or nodes that you centrally administer either the whole SP or racks of the SP.

It is recommended that you have at least one dedicated administration node per SP system.

Planning SP Node Types

SP nodes are available in 3 types: High, Wide or Thin nodes. The High node is a 2-to-8 way PowerPC (604) SMP processor (112 MHz), and can up to 2048MB of memory, 3 internal SCSI2 or SCSI2 Fast/Wide disks, and 16 microchannel slots. Thin nodes are equivalent to 39x class POWER2 (66 Mhz) or POWER2SC (120 MHz) RS/60000 processors, and can have up to 1024MB of memory, 2 internal SCSI2 Fast/Wide disks, and 4 microchannel slots. Wide nodes are equivalent to 59x POWER2 (66,75 MHz) or POWER2SC (135 MHz) RS/6000 processors and can accommodate up to 2048MB of memory, 4 internal SCSI2 or SCSI2 Fast/Wide disks, and 8 microchannel slots. Wide nodes are twice as wide as Thin nodes, and High nodes are twice as high as Wide nodes. 16 Thin nodes fit in an SP frame, or 8 Wide nodes, or 4 High nodes. Node types can be mixed. Each node has an SP switch adapter. Thin nodes have an integrated ethernet adapter for the SP ethernet and the Wide and High nodes have a ethernet microchannel card, and Wide and High nodes have a SCSI2 F/W microchannel adapter for SCSI2 Fast/Wide internal disks. In a typical SP configuration Thin nodes will have 3 available microchannel slots for network or external disk configuration, Wide nodes will have 5, and High nodes will have 13.

Planning SP AIX User specifics

An AIX user is required for Domino installation and administration. A group must be created which has root/system authority for these users. The group created should be called notes, as should the installation user. The choice of user name is up to you, but should also be in the group you create. It is recommended that the AIX user name be unique on each node, and also that the user home directory and Domino data Directory (where notes.ini resides) is unique, and local to each server node. This will ensure that the server is not dependent on the availability of a remote filesystem, and so that availability can be easily added.

Configuring the users should be done centrally on the control workstation or administration node. This allows the user definitions to be propagated to all the nodes. This process will be described further in the next chapter.

If you need to create a central repository (nfs directory or amd user home directory) for keeping common server files and notes data files, create a separate user or directory. It is recommended that you keep all Domino server AIX users and data directories uniquely named and local (not nfs mounted).

The recommended user configuration is illustrated below.

Recommended AIX User Configuration on the SP nodes

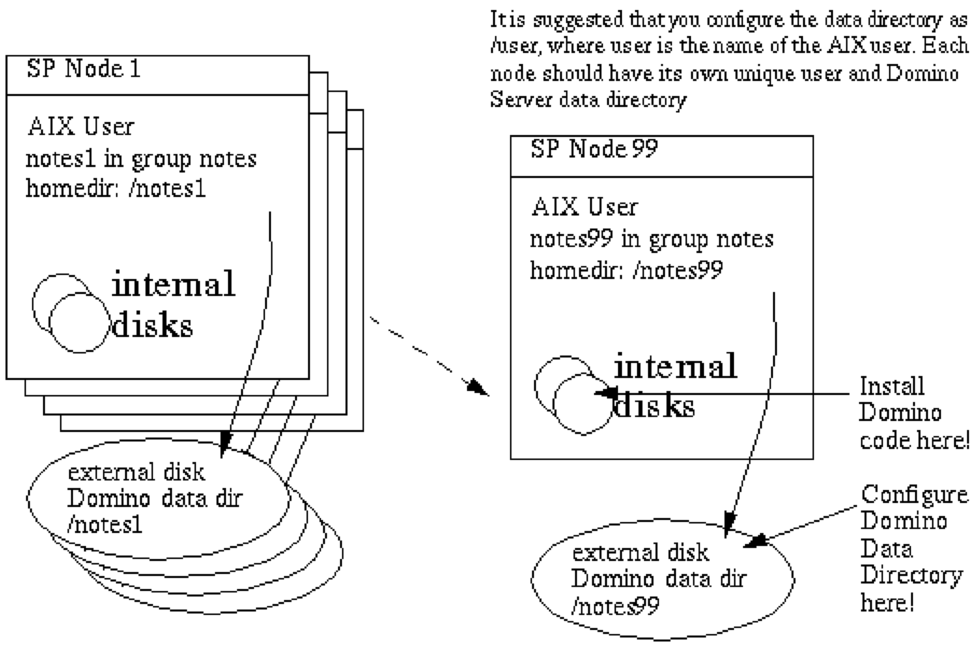

Planning Domino Filesystems and SP External Disk Requirements

The Domino Server application uses two filesystem paths:

- A path where the Domino executable code is installed. This is user selectable, but is ALWAYS symbolically linked by the Domino install process to /opt/lotus/bin. A usual choice for this filesystem is /usr/lpp/lotus, which fits in with AIX lpp naming conventions.

- A path where the Domino data directory, notes.ini, other control files, and user mail databases exist. This is configured when the Domino GUI is run for the first time on a new server. This is referred to as the Domino data directory, or just the data directory.

It is strongly recommended that the Domino Server executable code be installed on SP internal disks, and the data directories be installed on external disks, that are uniquely named for each server, and are configured with a mount point of ***/user***, where user is the AIX user for the Domino Server on that particular node.

External disks on your Domino Server nodes will give you the capability for more capacity, better performance, flexibility, and higher availability. Your Domino system will benefit from the external disks in the following ways:

- RAID or mirrored disks that can be concurrently serviced should a disk fail;

- an I/O subsystem that has higher performance (since there are more high performance disks that can be attached externally);

- An availability configuration where a backup node can "takeover" a server node's external disks should that node fail; and

- The capability to non-disruptively grow your I/O configuration

Finally, any server in an HACMP availability configuration MUST employ external disks for notes data directories. Any server or filesystem where data loss or corruption is absolutely to be avoided should use mirrored or RAID adapters and external disks. The use of HACMP to provide takeover capability of your disk subsystems will affect how you plan to cable your external disks.

Planning Disk Space and Configuration

Internal disk - Domino Executable Code

The file system where Domino is installed needs to be at least 130MB to hold the Domino server code.

External disk - Domino Server data directory

For all servers the Domino Server data directory, where the key databases reside, will need to be large -- at least 2GB. (The IBM phone directory database is over 1GB.) It is suggested to break out the data directory as a separate file system.

- **Mail servers:** allocate a base amount of space for the data directory (notes.ini, names.nsf, and name and address books). Add to that the amount of mail space allocated per user (typically 10-25 MB per user) times the number of users. It is suggested separating the base data directory and user mail databases, for ease of administration.
- **Application servers:** allocate a base amount of space for the data directory and add to that the size of the application databses.
- **Mail Routing Servers:** the external disk space will be used for mail in flight - this is typically in the tens of Gigabytes. Although under normal circumstances this space will not be fully utilized, remember to allow sufficient space to hold backed up mail in the event a network connection is lost to a routing destination (or a routing destination failure occurs).
- **Replication servers:** allow sufficient space to hold all of the database files that will be replicated outside the domain.

Because of the way Domino works, the JFS log logical volume(s) will be very heavily used -- for performance reasons consider putting it on a disk by itself.

Disk mirroring or RAID will increase your disk requirements, be sure to plan accordingly.

Cabling multi-node SSA disk rings for availability

IBM SSA (Serial Storage Architecture) disk subsystems are an example of a robust I/O subsystem for Domino. The SSA subsystem's high capacity (up to 432 GB per adapter), high throughput (up to 35MB/sec sustained), and flexibility of configuration make it a desirable choice to use with Domino. The ability to hot-plug both cables and disks without affecting system access to data through alternate routes makes this choice even more desirable for the enterprise. The SSA loop architecture makes it possible to physically configure a large pool of physical devices to be accessible by one of up to eight (8) nodes - this is important for configuring highly-available groups of servers. Logical reconfiguration of devices to specific volume groups and file systems can be done dynamically through the AIX logical volume manager (LVM), without having to disrupt the running applications. For more on SSA and other IBM disks subsystems please see the SSA storage Web page at *http://www.storage.ibm.com/storage/hardsoft/ssaovu.htm.*

In multi-server SSA disk configurations, you want to plan your configurations for performance and availability. It is recommended that you plan to configure each server

- with mirrored SSA adapters and disks

- with each servers' disks directly accessible to it, for performance

The following picture illustrates this point.

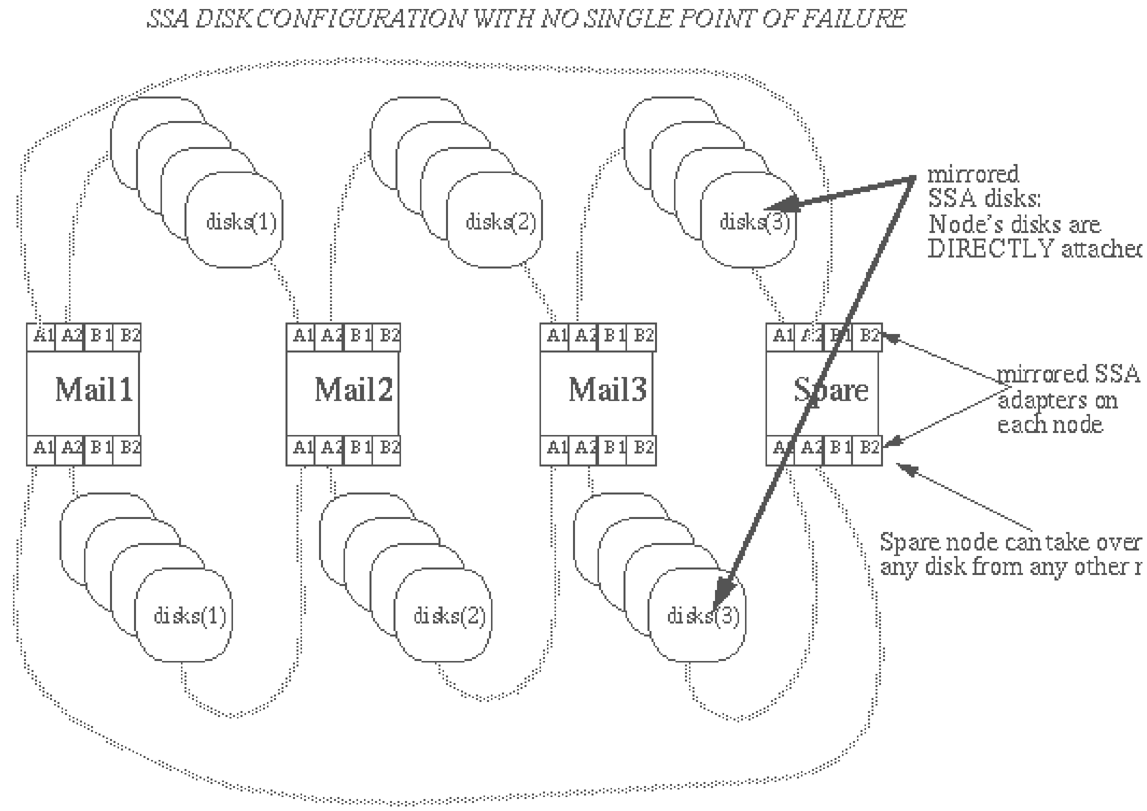

Why you should use RAID or mirror external disks

RAID or mirrored external disks will provide the following advantages in your configuration:

- Data protection and availability - disks can be concurrently repaired while the Domino server(s) stay up;
- Performance advantage for read-only access to data and filesystem logs; and
- Backup flexibility - You can take a mirrored copy of data off-line for backup, as opposed to taking a server off-line.

If you choose not to mirror all Domino external disks, it is recommended that you mirror at least the Domino Server data (notes.ini, name and address books, and filesystem logs). This will generally require only one extra disk per server and will ensure that even if a disk is lost the server will remain available.

Optimizing Domino Server Data Directory on External Disks

One key advantage of Domino Servers in the AIX environment is the ability to spread your application data across multiple disk devices which can be effectively managed as a single I/O subsystem. And this can change dynamically as needed.

You will want to consider the following issues when planning space and topology of the Domino Server data directory on external disks;

- mirrored disks - how much and what data
- volume group definitions - use one volume group with multiple logical volumes and filesystem (jfs) logs or multiple volume groups with single logical volumes and filesystem logs
- filesystem (jfs) log placement for performance. Because of the way the Domino Server works, the filesystem log logical volume(s) will be very heavily used -- for performance reasons consider putting it on a disk by itself.
- Simplifying the administration of the filesystems involved.

The following 3 pictures illustrate the best options for designing your data directory.

Option 1 - One Volume Group, filesystem log, multiple filesystems for users

Assume that on a Domino Server you have an external disk filesystem called /notesX and your mail files are therefore in /notesX/mail, the /notesX/mail filesystem in AIX can be in a separate file system (and also a different disk). Taking this one step further, you could do some manual load balancing by having /notesX/mail2, 3, 4, and 5 (an so on) across different AIX logical volumes, as illustrated in the picture on the following page.

By keeping sets of users that are in different departments or functional groups in their own AIX logical volume and filesystem, you can compartmentalize their space and requirements, while also balancing load across different disks. Keep in mind that the filesystem (jfs) log could be mirrored and on separate disks, (for higher performance and availability), although that this might not be cost-effective for a 4MB log to be on a separate disk. In this example, by mirroring the filesystem log for the volume group it is eliminated as a single point of failure.

The other logical volumes contain "sets" of mail users. You can see that a single disk failure will not affect the server (it's key files and filesystem log are mirrored), and at worst only a set of users might be affected by the failure. This is much better than having a disk failure take the whole volume group off-line.

Data Directory Option 1: One Volume group and filesystem log

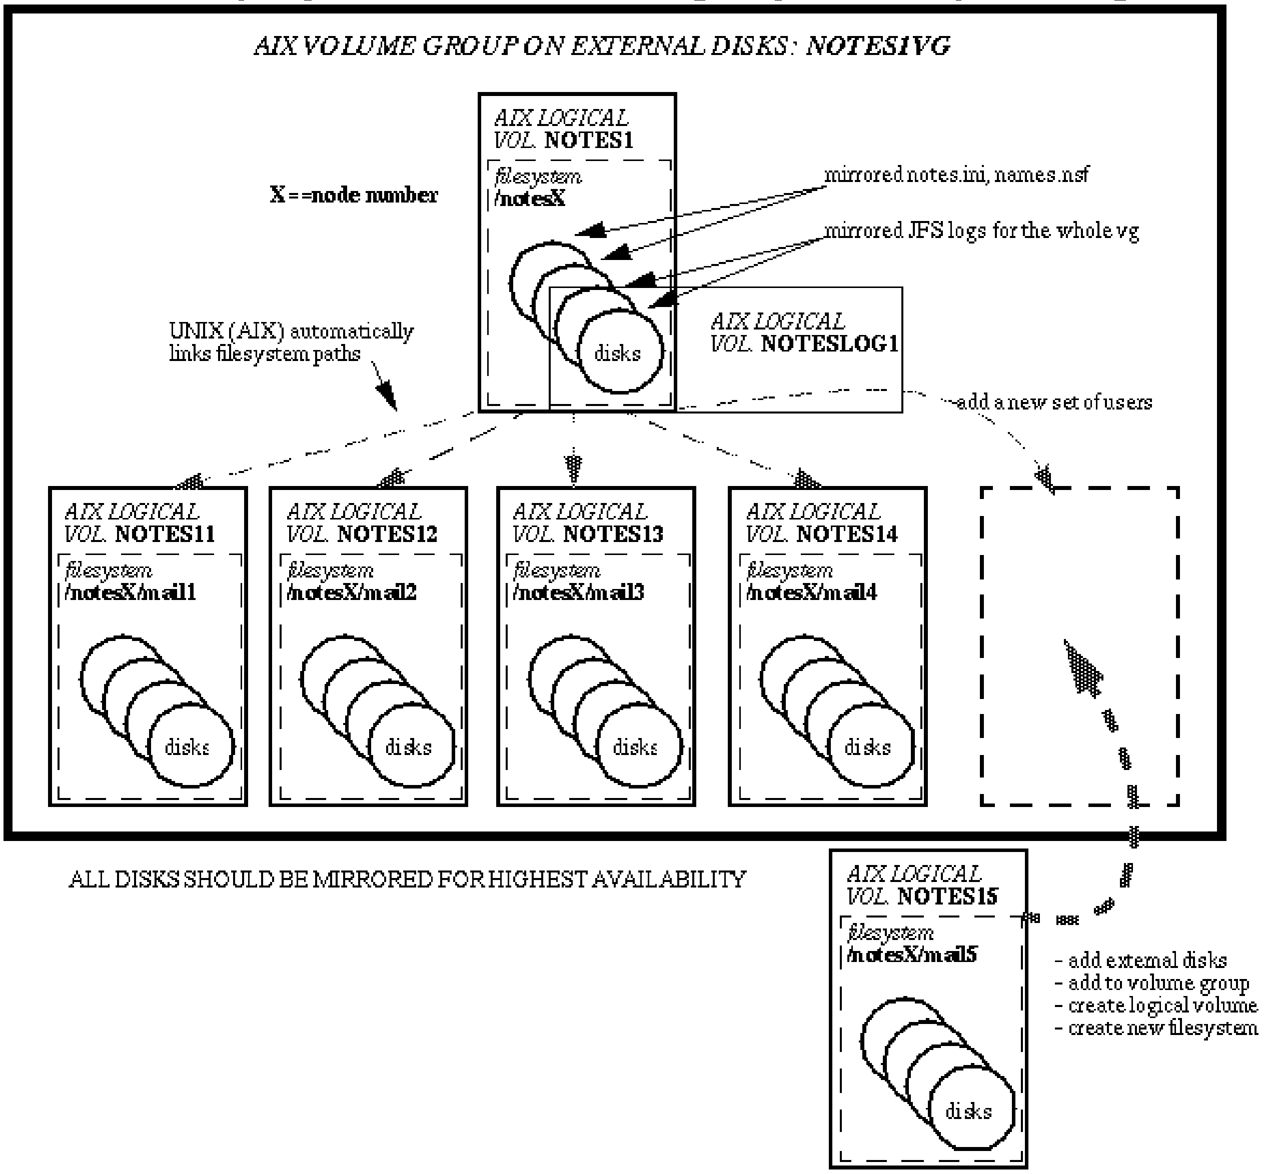

Option 2- Multiple Volume Groups, filesystem logs, filesystems

This option places one disk in each volume group (two if mirrored). Each volume group has its own filesystem (jfs) log, which may be mirrored and on separate disks. Keep in mind that putting one 4MB filesystem log on a dedicated disk might not be cost-effective; so that likely the filesystem log will be on the same disks as the filesystems.

In this implementation a failure of one disk in the user mail directory area will only affect the users corresponding to that disk or log - as opposed to losing the whole set of users in all mail directories (if you only had one volume group and filesystem and log. This is because each volume group has its own filesystem (jfs) log. This comes at the administrative expense of having to configure multiple volume groups.

Data Directory Option 2: Multiple Volume Groups

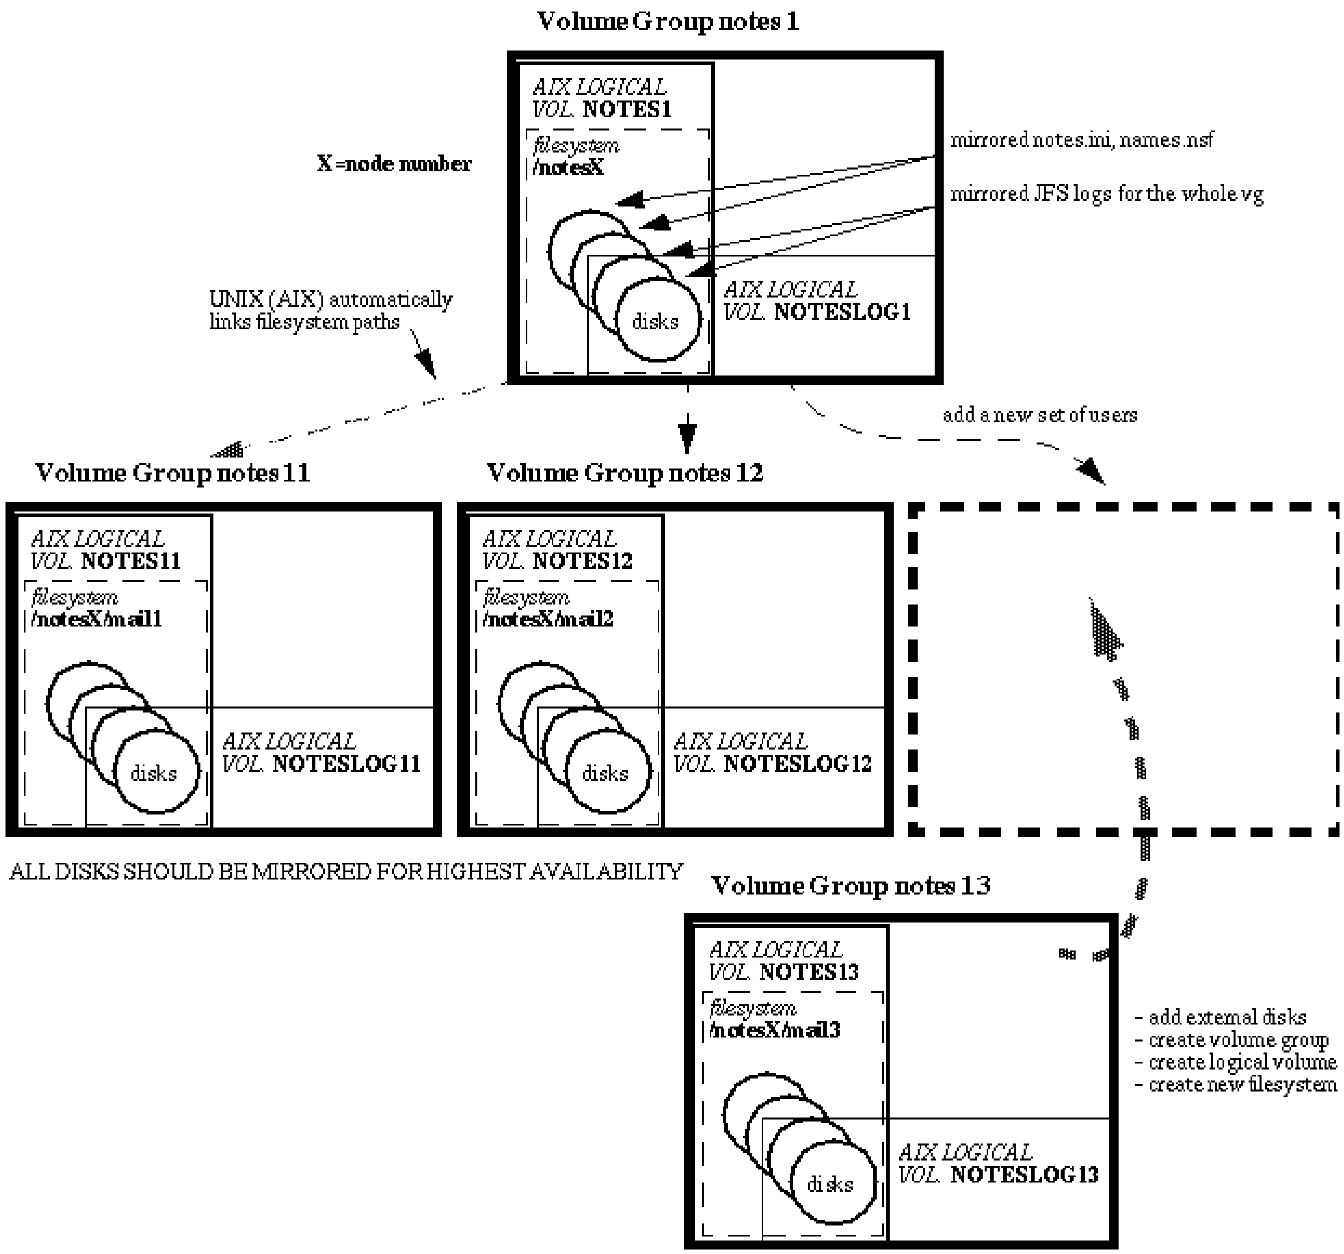

Option 3- One Volume Group, multiple filesystem logs, filesystems

This option is the same as option 1 but now there is a separate filesystem (jfs) log logical volume created for each filesystem. These filesystem logs can be mirrored and on separate disks (for availability and highest performance). Keep in mind that putting each 4MB filesystem log on a dedicated disk might not be cost-effective; so shown here is multiple 4MB logs for each filesystem on one disk (two if mirrored).

This implementation is very similar to a multiple volume group implementation in that a failure of one disk in the user mail directory area will only affect the users corresponding to that disk or log - as opposed to losing the whole set of users in all mail directories (if you only had one volume group and filesystem and log.

Data Directory Option 3: One Volume group and multiple filesystem logs

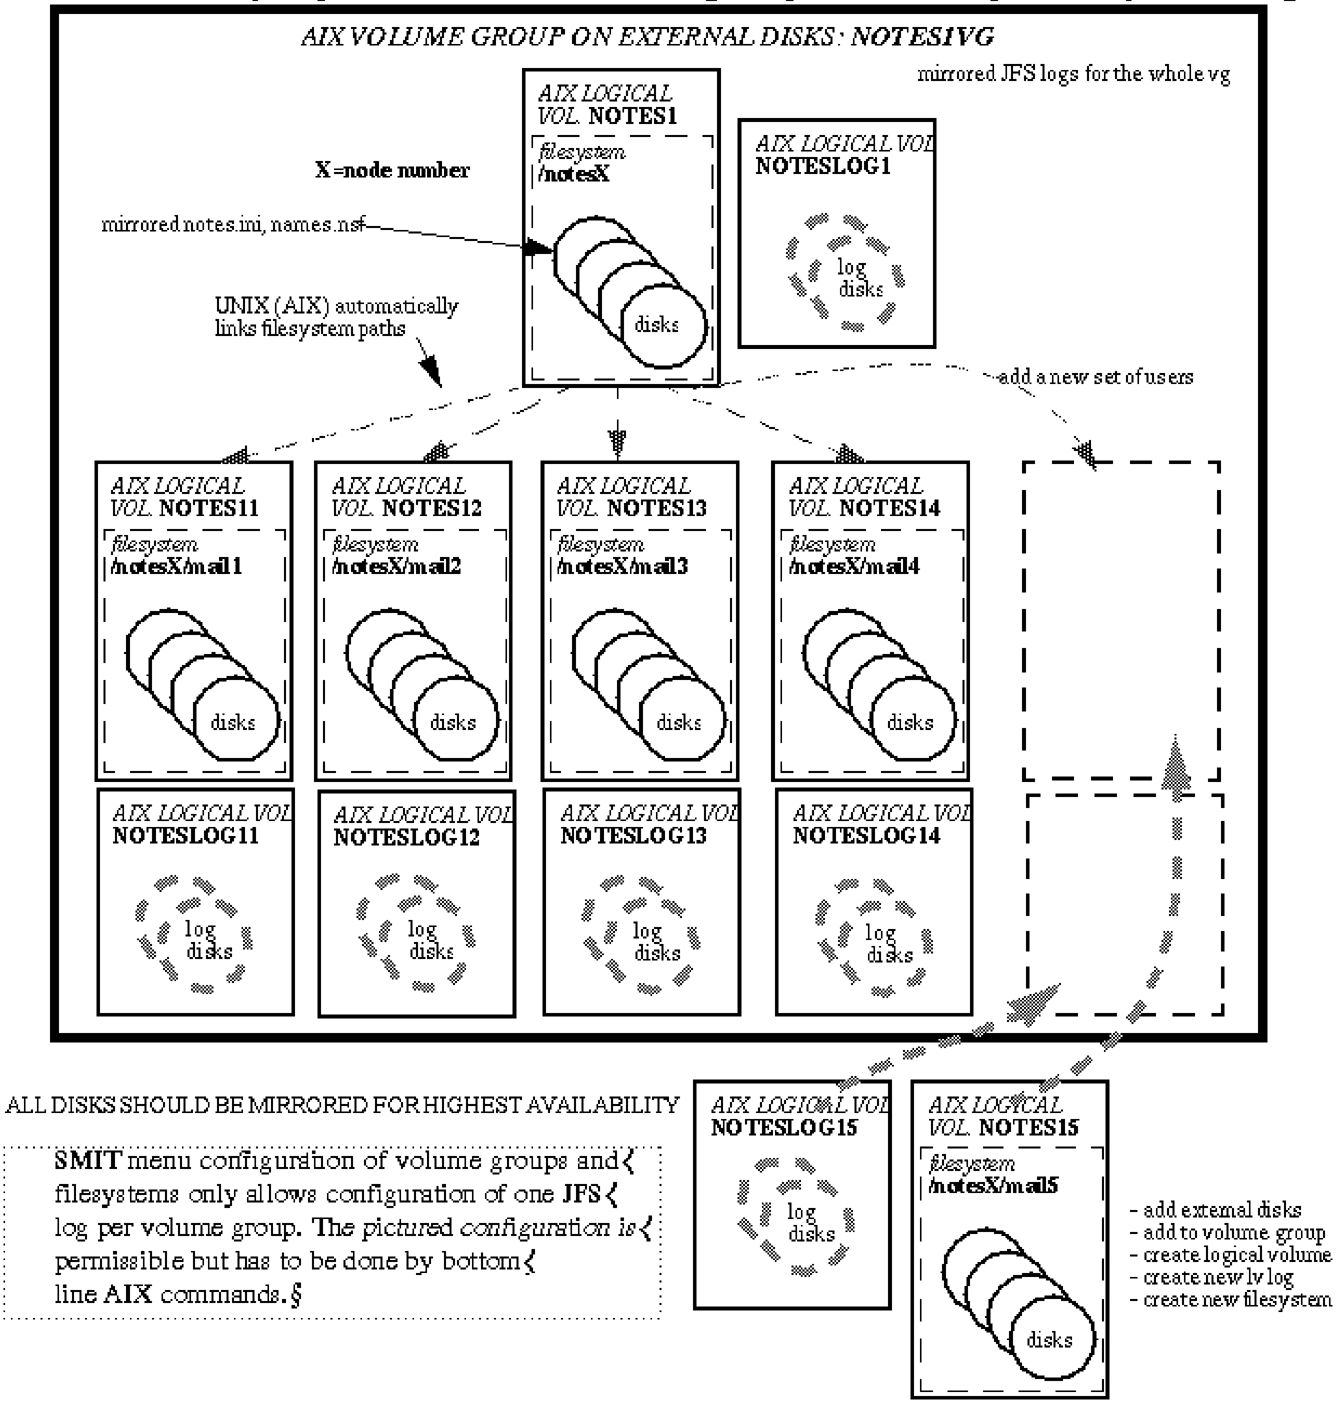

Each of the options shown, if configured judiciously, should yield about the same performance. At the time of the writing of this document, no data exists to identify the highest performance option, but each should be within a few percent of each other.

The most important objectives to remember in configuring your data directories are

- making them easy to administrate. It should be easy to add users/space so that the access is uniformly spread across the disks.
- making them available, by mirroring external disks or at least the jfs logs; and
- maximizing performance, by creating jfs logs on separate disks if possible..

The three options pictured offer the best alternatives for maximizing your data directories' performance and simplifying its administration. If you did not implement one of these three options for your data directories, you would probably end up with a mail directory setup as in the following picture, which is **NOT RECOMMENDED**;

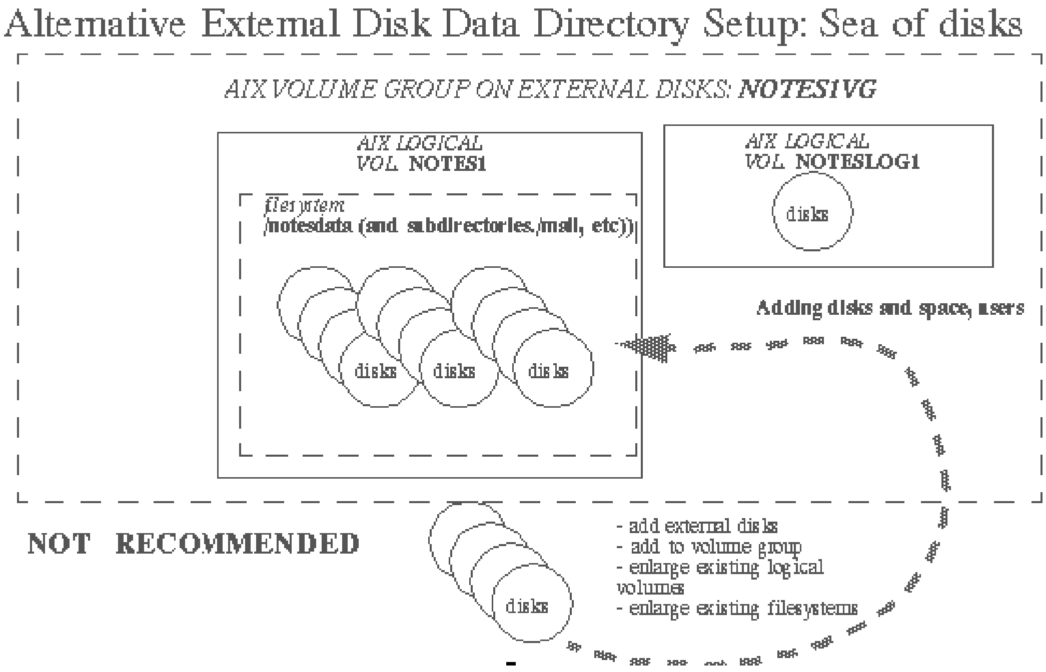

This alternative configuration appears easier at a glance but keep in mind that logical volumes and filesystems will have to be made bigger in order to use the extra disks. Also, there is no guarantee of any spread of users or files across the disks, and the access to the JFS filesystem log may become a bottleneck. This may also impact backup and restore times for sets of users and create a longer outage if a disk failure occurs.

Regardless of how you implement the data directories, for performance reasons it is recommended that you keep the filesystem log for the data volume group(s) on separate disks, and mirror at least the filesystem log(s) logical volume and Domino Server logical volumes that correspond to the "base" of the data directory (notes.ini, name/address book). This will ensure that a disk failure in the volume group(s) will keep the server available.

Planning external network connectivity for users

There are many ways to configure external networks and connect users with them to the SP. Consider the following two options for attaching an external network to the SP:

- Have an adapter card in each node to attach to the external network
- Have a dedicated node or "nodes" with adapters to be gateways into the SP switch network, or use a router into the SP switch network.

The first option becomes cost-prohibitive as the number of nodes in your configuration increase.

The second option is more scalable and is able to leverage the bandwidth of the SP switch to attach clients to your system. This is more attractive for the following reasons:

- Availability - one can configure gateway node pairs with pairs of intelligent routers to route client traffic to and from a number of mail and database nodes to provide some network traffic and load balancing (split the network traffic by sub-netting) and availability / failover (via HACMP) between the two gateway nodes. You could place these two gateway nodes in separate frames of the SP to increase the availability, as well as placing the two routers on different power boundaries (separate circuit breakers).
- Cost Effectiveness - it reduces the number of external network connections into the SP - fewer parts are easier to manage. Also network hardware requirements for the SP and site (ATM switch ports for example) will be reduced.
- Simple management and easy growth - when you want to add a Mail (or Database) Server node - the network infrastructure does not need to change (until your gateway capacity is exceeded) - you can simply add a node to the SP and start putting users there. (SSA simplifies this as well by making it easier to physically share disks among a group of nodes and re-assign disks to different nodes as the need arises)
- Performance - Large network pipes (ATM) may not be practical on each node, but large pipes on gateways into the high-speed switch should provide better throughput, also greater stability since the more traffic that gets taken off the external network, the more stable the external network should be.

The SP switch network is a highly reliable network. Inside the AIX IP device driver for each SP node’s switch connection, there are 4 possible hardware paths established (at initialization time) between each pair of SP nodes. Should one path fail, another will be used instead, without impact on higher level applications. Redundant cooling fans and power supplies are implemented, which make the SP switch an enclosed network that is exceptionally reliable.

Dedicated external network connectivity for other server nodes

Replication and Mail Routing nodes in the configuration should have their own external network adapters so that replication and mail traffic does not interfere with user (client) network connections. Additionally, application servers that are internet servers may require their own external network connection. Typically these connections are high bandwidth FDDI or ATM connections from the Domino server(s) to the external network.

Planning SP Domino Servers - Domino specifics

Planning Mail Servers

Refer to the following table for planning mail servers. At this time, the Thin Nodes are the most cost-effective, fitting the most users in the least amount of space (keep in mind that a High node takes up 4 times as much space as a thin node).

Node Capacity and Workload for Mail Servers

(*) Data for the POWER2SC THIN nodes is being collected, and not available at this time.
(**) Multi Server (MS) HIGH node configuration is undergoing certification - not supported at this time.
(***) These numbers are presently being used by IBM in their rollout, and assume no more than 30-40% concurrent users. Benchmark measurements are approximately 85% of the total registered users per server.

	THIN (39H * 512 MB)	**MS - HIGH (R40 8way 1024 MB) ** **
Recommended users registered per server	1000	3600 ***

Planning Domino replication hub and mail routing hub servers

It is recommended that SP Wide nodes or High nodes (at least 1024MB of memory) be used as replication servers, as workloads have shown better replication performance with larger main memory.

Typical configurations of replication servers employ multiple servers, each handling different classes of replication performance. For example, use of three replication servers with priorities high, medium, and low can be implemented for traffic priorities and sizes that will guarantee replication performance. For Mail Routing Servers, be sure to allow sufficient directory space such that the server will not "fill up with mail" should a network to a destination fail (or a destination server itself).

Planning Application Servers

Domino Application Server requirements will depend on the type of database, server, and application, so at this time guidelines are not readily available. The best advice is to monitor your configuration (AIX and Domino performance statistics), and plan accordingly.

Planning Domino Server Names and connectivity

You will need to choose names for each of your servers so that other Domino Servers and users will be able to communicate with them. To understand this completely, it is important to understand how the Domino Server resolves a server name into a TCP/IP address.

A Domino Server resolves a server name in the following way:

- The Name and Address book is checked with the name of the server to be resolved. If a connection record for that server name exists, then that is used to resolve the name into a TCP/IP address.
- If there is no connection record for the server in the previous step, then the Domino Server uses standard name resolution on the server it runs on (DNS, /etc/hosts, or NIS) to resolve that name into a TCP/IP address. Note that if there are no connection records for a server and the name is also not resolvable by TCP/IP name resolution then the server will not be reachable (unable to find path to server).

Plan your server names and coordinate these with your choice of IP address names - follow this recommended procedure.

- For servers without external network connectivity, choose SP switch IP names as your Domino Server Names. Use names that are significant - like SITEMAIL01 as the first mail server, SITEAPP01 as the first app server, etc...(in this example SITEMAIL01 and SITEAPP01 would resolve to switch IP addresses). Note that these servers should be configured such that their default routes (TCP/IP) are either gateway nodes or a router attached to the switch - this allows these servers to be communicate with networks out side the SP. (NOTE - when you do this you must also add a static route such that communication from SP nodes back to the control workstation is routed properly - also add a static route for the SP ethernet back to the control workstation). Typically mail and application servers are configured in this way.
- For servers directly connected to external networks, typically mail routing, replication, or internet hubs, choose their Domino Server names as the IP names corresponding to their external network interface(s). This allows these servers to be connected to by all servers from other sites, by TCP/IP resolution. To make DNS name resolution between Domno servers in different DNS/Domino domains more efficient, use the search parameter in your /etc/resolv.conf file to specify all nameservers to search for in resolving Domino Server IP names.
- The final step is to add connection records on all servers in the first step, to point to the servers in the second step, through the SP switch. This allows mail and app servers to replicate through the switch to the replication hubs that are connected to other sites.

Make sure you fill out the Domino Server names worksheet in the back of this chapter, properly recording names, connection records, and default routes that need to be configured.

Planning for backup

If your backup server will be in the SP, you will need to plan for it, and tape or disk peripherals for backup. If you will be using ADSM, consult the ADSM Administration Guide.
Typically an SP Wide node is used as an ADSM backup server, which provides connectivity for up to 6 attached tape devices. The ADSM agent for the Domino Server, which provides incremental backup capability for Domino Server data, is recommended. For more information on ADSM, see the web page at *http://www.storage.ibm.com/storage/software*.

Planning Domino Servers for High Availability

Use of HACMP for AIX (High Availability Cluster Multi-Processing) or Lotus Domino Advanced services to provide high availability to components of your environment can be incrementally added as needs designate.

Where to use HACMP or Domino Advanced Services

Both HACMP and Domino Advanced Services can provide availability for Mail, application, and replication servers. As such, there is confusion over where to implement what.

HACMP (a separate product for AIX) can be configured (through use of directories on external disks) in up to 8 node clusters, where a designated node will detect and restart a Domino server on another node, either a hot spare or active node. This process is called failover. HACMP can be used with any of the mail routing, mail, replication, and application servers in your configuration.

In the Domino Server Powered by Notes release 4.5, Domino Advanced Services (a separate product for Domino) provides event driven replication of any selected database, between all nodes in a configured (up to 6 node) cluster. Should a server of one of the database replicas fail, Domino Advanced Services will failover the user to a server of another replica. This is very similar to an HACMP failover, with the added benefit of user load balancing of across the cluster.

Since Domino advanced services provides nearly identical function to HACMP, is there any need for HACMP in your implementation? The answer is up to you, but the following recommendations are those that are the easiest to implement and most cost-effective.

Domino Server Function	Preferred Availability Solution	Rationale
Mail	either HACMP or Domino Advanced Services	Domino Advanced Services will support Mail user failover if a mail server is unavailable. However, shared mail is not yet supported and HACMP may be easier to implement for Mail availability.
Mail Router	HACMP *	Domino Advanced Services does not fail over mail routing functions. HACMP can failover / restart and recover this functionality.
Application	Domino Advanced Services	Domino Advanced Services will load balance and failover users to a backup for properly replicated databases.
Out-of-Domain Replication	HACMP **	Domino Advanced Services does not failover out-of-domain replication. Therefore if it is required to keep a dedicated replication server available, use HACMP.
Internet	HACMP***	Domino http/IP addresses not failed over by Domino advanced Services as with HACMP
Internotes	Domino Advanced Services	Internotes Servers are part of the advanced services functions/servers that can be failed over

(*) Note: HACMP may not be required. Your requirements may not demand that the Mail Router be available all of the time, since no data will be lost (if mail spool disks on this server are mirrored). Mail-in-flight will be delayed until the Mail Router is repaired.

(**) Note: HACMP may not be required, since you might not demand that a replication server be available all of the time, since replication can be completed when the server is repaired, or there might be multiple replication servers each covering different priority replication which will cover the replication on a different time scale.

(***) Note: If you are using load balancing (LoadLeveler ISS) across multiple Domino servers with the same data, HACMP might not be required.

Planning SP Domino Servers - Domino Advanced Services

Refer to the Domino Advanced Services release notes for configuring Domino Advanced Services. Use the SP switch as the private LAN recommended by this documentation for the cluster replicator.

Planning SP Domino Servers - HACMP

Additions to the your Domino configuration required for HACMP are;

- External disks. Database(s) and data directories to be made available must be migrated from internal to external disks if not already on external disks.

- Backup node(s). If you implement a hot standby configuration with HACMP, at least one backup node for at most 7 server nodes is required. This node must be cabled into networks and disks for **each** node it is to back up. Depending on how many servers to one backup you require will determine your external disk choices and whether or not a thin or wide node is required.

- HACMP software. This can be incrementally added on only the nodes that need it.

- Redundant network hardware. This is required only if an external network is being made available (refer to the previous section on configuring a reliable external network connection for users);

- HACMP start and stop scripts. These are scripts required so that HACMP can start and stop Domino servers. A sample set of scripts is provided in the appendix of this document.

HACMP failover configurations

Hot Standby

For hot standby configurations for Domino servers, an N+1 HACMP configuration using HACMP rotating or cascading resource groups is recommended. In this implementation, N+1 (max N=7) SP nodes are connected to external disks, where Domino Server data resides. Should any server fail, the backup or spare node will assume identity of the failed server by acquiring it's IP addresses and external disks, and restart the Domino Server.

This configuration will yield the best performance of the system after a failover, since a dedicated spare node is the failover node. Spare HACMP servers can be configured to normally perform other tasks - such as a system administration, NetView, or ADSM backup/restore. Should failover occur, these primary functions could be shut down and the server would perform its backup server function until the failed server is repaired.

Mutual Takeover

For mutual takeover configurations for Domino Servers, a mutual failover configuration using HACMP is recommended. In this configuration, up to 8 active SP nodes are active, with any of the nodes being able to takeover a failed node's IP address and disks and restart the Domino Server.

This configuration may cause some degradation in server performance after a failover because the server failed over to will now be running two instances of the Domino Server.

For the client gateway nodes, an HACMP mutual takeover configuration is recommended using cascading resource groups - where both gateway nodes are active and either can assume IP addresses and continue should the other fail.

For more information consult the HACMP version 4 documentation, and complete the planning worksheets in that documentation.

Creating your Domino/SP planning worksheets

SP/AIX Installation Worksheet

Use this worksheet to configure AIX users for the installation and administration of Domino and their home directories. It is recommended that the data directory name chosen be /user, where user is the name of the AIX user you create.

SP Node	Domino AIX user	User Home Directory (also data directory)
1	notes1	/notes1
2	notes2	/notes2
3	notes3	/notes3
4	notes4	/notes4

Domino Server Filesystems Worksheet

Use this worksheet to specify your external disk configuration, and where the Domino Server data directories on each node will be.

SP node	Domino Data Volume Group Name	Domino Data Logical Volume Name	mirror (y or n)	Domino Filesystem Mount Point	Filesystem Log Name	Filesystem size in PPs	hdisk list	
1	notesvg1	noteslv11	y	/notes1	noteslog1	8	250	hdisk2
1	notesvg1	noteslv12	n	/notes1/mail1	noteslog1	8	1000	hdisk3
1	notesvg1	noteslv13	n	/notes1/mail2	noteslog1	8	1000	hdisk4
1	notesvg1	noteslv14	n	/notes1/mail3	noteslog1	8	1000	hdisk5
2	notesvg2	noteslv21	y	/notes2	noteslog2	8	250	hdisk2
3	notesvg3	noteslv31	y	/notes3	noteslog3	8	250	hdisk2
4	notesvg4	noteslv41	y	/notes4	noteslog4	8	250	hdisk2

Domino Server Names Worksheet

Use this worksheet to specify your Domino Server names for each SP node. be sure to record which IP names and addresses will correspond to each server name. Don't forget to include any default routes for user connectivity to servers that aren't connected to external networks (also add connection records on these servers to use the switch to connect to servers on the external networks).

SP node	Domino Server Name	switch IP name	switch IP address	external network IP name	external network IP address	Connection Records (server/IP name)	Default Route IP name	Default Route IP address
1	SPapp1	SPapp1	1.2.1.1	N/A	N/A	SPrep1/SPsw3, SPMR1/SPsw4, SPinet1/SPsw5	gate1	9.9.9.1
2	SPmail1	SPmail1	1.2.1.2	N/A	N/A	SPrep1/SPsw3, SPMR1/SPsw4, SPinet1/SPsw5	gate1	9.9.9.1
3	SPrep1	SPsw3	1.2.1.3	SPrep1	8.8.8.1	N/A	gate2	9.9.9.2
4	SPMR1	SPsw4	1.2.1.4	SPMR1	8.8.9.1	N/A	gate2	9.9.9.2
5	SPinet1	SPsw5	1.2.1.5	SPMR1	7.7.7.1	N/A	gate2	9.9.9.2

Table of Contents Next Chapter

Table of Contents Previous Chapter

Installing AIX and Domino on the SP Nodes

.domino Running on RS/6000 SP

Installation of SP hardware - requirements

You should ensure that each SP node has at least 256MB of physical memory and 2.2GB of internal disk.

Installation of SP nodes, AIX software/config, and Domino

It is strongly suggested that the reader refer to the web-available Domino Server Release notes, for the level of the code you install, at *http://www.lotus.com*. Refer to the Notes utilities in Appendix B of this document for utilities that allow performing many of these steps in parallel from your control workstation or control node.

1. Install SP AIX on the SP nodes via the SP AIX installation process. Use the AIX 4.1.4 image shipped with your SP, and apply the additional PTF updates required in the Domino Server Release Notes.
2. Install PSSP 2.1 release at latest release level on the SP nodes.
3. Install external disks and adapters, and any additional PTF updates required in the Domino Server Release Notes.
4. Create filesystem /opt on each node (mkdir /opt) where Domino Server code is to be installed.This will contain a link for the Domino executable code, placed there by the Domino install process.
5. Create the filesystem where the Domino executable code is to reside. This filesystem must have at least 130 MB of free space. It is suggested that this be on each node's internal disk(s).
6. Create volume groups and filesystems (from your worksheets) on external disks if present, for the notes data directory. The size of this (base) filesystem must be large - on the order of 2GB, to hold the notes name and address book. Follow the planning instructions from the previous chapters to create the volume groups and logical volumes to provide highest availability and performance.
7. Create your Domino Server AIX group and user for each node (from your worksheets). Typically this is user notesX (node number=X) with group notes. Note that for availability reasons it is recommended that the home directory for the user be local and the unique data directory name that you have chosen for each node (for example, make the data directory on node7 /notes7, and the home directory for the user notes7 you create on that node). Change ownership of the home directories (from the previous step) to the appropriate AIX user (chown user.group /dir).
8. Add the following Domino Server variables to each nodes AIX Domino user's .profile:
   export Notes_SHARED_DPOOLSIZE=250000000
   export NSORDER=local,bind,nis
9. Increase paging space on each node such that there is as 2 times as much paging space as there is physical memory.
10. Customize the SP nodes to add necessary addresses and default routes (from your Domino Server Name worksheets) for the server names and addresses you have chosen. Do not forget when you do this to add a static route so that communication from the SP nodes back to the control workstation through the SP ethernet network is maintained. Configure TCPIP on SP nodes so that the Domino Server names you have chosen are resolvable IP address names, by way of the name

resolver (DNS, /etc/hosts, NIS). Verify with the ping command your IP connectivity between servers. Update the search parameter in your /etc/resolv.conf file if you are using DNS, specifying all nameservers in the DNS/Domino domains that this configuration will communicate with.
11. Install the Domino Server executable code on all nodes, via the Domino installation process. If you have the Domino code on the control workstation, install to the nodes from there. Specify the Domino AIX user you created on your worksheets and install to multiple nodes to speed the process. Specify the directory that you created in step 5 as the directory for the installation.

You are now ready to configure the Domino Server software.

Table of Contents Next Chapter

RS/6000 | Solutions | Hardware | Software | Support | ReSource | sitemap
IBM | Order | Search | Contact IBM | Help | © | ®

Table of Contents Previous Chapter

Configuring Domino on the IBM RS/6000 SP

Configuring Domino Servers

It is strongly suggested that the reader refer to the web-available Domino Server Release notes, for the level of the code you install, at *http://www.lotus.com*. See Appendix B for utilities that will help you perform some of these steps in parallel on multiple nodes.

1. Install the data directory (from the worksheets you created) via the Domino server GUI (type notes& as the AIX user you created) **on one node**. Make sure that you specify the AIX users's home directory (/user) as the data directory for the Domino Server installation.
2. Start the Domino server on this system by (as the Domino AIX user), typing the command "server". Verify that the server starts without problems.
3. Don't forget to add connection records, for each server that is not connected to external networks, the names of servers (from your Server Names Worksheets) attached to the external networks, specifying the switch IP addresses of these externally connected servers. ***Note:*** It is useful to use the Trace Connection function to verify connectivity between servers (this is in the File-tools-User Preferences-Ports menu on the Notes GUI).
4. Follow the procedure(s) in the Domino Server Release notes, for the level of the code you install, to configure that server, and other servers.

The servers are now ready for any further configuration such as Domino Advanced Services, or HACMP.

Table of Contents Next Chapter

RS/6000 | Solutions | Hardware | Software | Support | ReSource | sitemap
IBM | Order |

Table of Contents Previous Chapter

Adding / Configuring Availability (HACMP)

This chapter assumes you have determined which server(s) you wish to make available via HACMP, and completed the HACMP planning worksheets in the HACMP Planning Guide. For more on HACMP see the web page at *http://www.austin.ibm.com/software/Apps/hacmp.html.*

1. **Add and/or configure hardware if required.** On the nodes to be installed with HACMP, external disks / disk adapters may need to be installed if not already. Make sure that the takeover node(s) are also cabled to the disks they need to take over. Follow the procedure in the HACMP Administration guide for setting up the shared volume groups (from your Notes Data Filesystems worksheets), and turning off auto-varyon, etc.
   Note: it is strongly recommended that RAID or mirrored disks and adapters are used in the external disk configuration.
   Add any redundant network adapters if your connection to an external network is to be made available. Additional adapters are not required for any SP switch IP address failover.
2. **Install HACMP** on the SP nodes where HACMP is to be configured. Follow the directions in the HACMP Installation guide.
3. **Configure HACMP Cluster, Nodes, Adapters, and topology**. This can be done from one node in the cluster. Remember that the SP uses alias IP addresses as boot and service addresses in HACMP. You will need three Notes SP Switch IP addresses for each Notes Server that will be recoverable via HACMP - (1) the Switch "base" address that is defined in the SP Control Workstation SDR, (2) the Switch boot IP address for HACMP, and (3) the Switch service IP address for HACMP ***(this is the Notes server address).***
   Note: If you are adding HACMP to an existing SP Notes Server configuration, to change the existing "base" (or SDR) switch IP address that the Notes server is using into an service address that you can configure to HACMP for IP address takeover, follow the following procedure;

   - You must come up with a boot address for HACMP, and a "dummy" IP address that you will replace the switch "base" (or SDR) address with. The old switch "base" address will become an HACMP service address. Make sure that all these IP addresses are resolvable on all nodes, vis DNS or /etc/hosts (whatever nameservice your installation uses).

   - On the control workstation, edit the switch interface for the HACMP nodes that will have their exiting switch "base" addresses become HACMP service addresses so HACMP can control them. Do this via the SP SMIT menus, effectively changing the switch "base address to the "dummy" IP address you have chosen, and customize these nodes to have them reconfigure these addresses.

   - Configure the boot and service addresses to HACMP as you would normally.

   - Note that on spare Servers you configure to HACMP you only need a "dummy" address for the switch "base" address, and a HACMP service address. Customize the node to pick up the "dummy" IP address, and configure HACMP on that node for only the service address.

4. **Create HACMP Start and Stop Scripts.** You will need to create start and stop scripts for starting and stopping the Notes server processes on the HACMP nodes. Sample scripts are provided in an Appendix A of this document, and Appendix B provides utilities to set this up on SP nodes in parallel.
5. **Configure HACMP Resource Groups.** Follow the procedures in the HACMP manuals for

configuring rotating resource or cascading resource groups from your HACMP worksheets for backing up the Notes servers. Specify the IP addresses, volume groups, and Lotus servers to be failed over. Follow the process in the HACMP Installation and Administration Guides.

6. **Configure HACMP Network Up/Down Events.** Switch adapters can become single points of failure for a node. To eliminate these as points of failure for a server, use the HACMP network up/down events to initiate failover via "/usr/sbin/cluster/utilities/clstop -grsy", if the switch network does not reactivate in a certain amount of time after an outage.
7. **Update /.rhosts** for all nodes in the cluster to include all SP ethernet, switch (SDR and alias) address names with root privilege.
8. **Synchronize and Verify the configuration.** Follow the HACMP menu procedures to distribute the cluster information for topology and resources to all other nodes in the cluster, and verify the configuration. See the HACMP Administration Guide for more information.
9. **Update /etc/inittab and /usr/sbin/cluster/etc/harc.net** for all nodes in the cluster. Change tcpip and nfs processes in /etc/inittab to run level 2 (from run level a) on any nodes configured for IP address failover. Uncomment the line in /usr/sbin/cluster/etc/harc.net which starts nfs (near the bottom of the file).
10. **Validate the configuration.** Start the HACMP cluster, and monitor the information on each cluster node's /tmp/hacmp.out file Failover the nodes in the cluster, monitoring the output in /tmp/hacmp.out. Correct any mistakes.

Sample HACMP configurations

The following four pictures illustrate three sample SP / Domino Server configurations tailored for HACMP recovery of these servers or network connections.

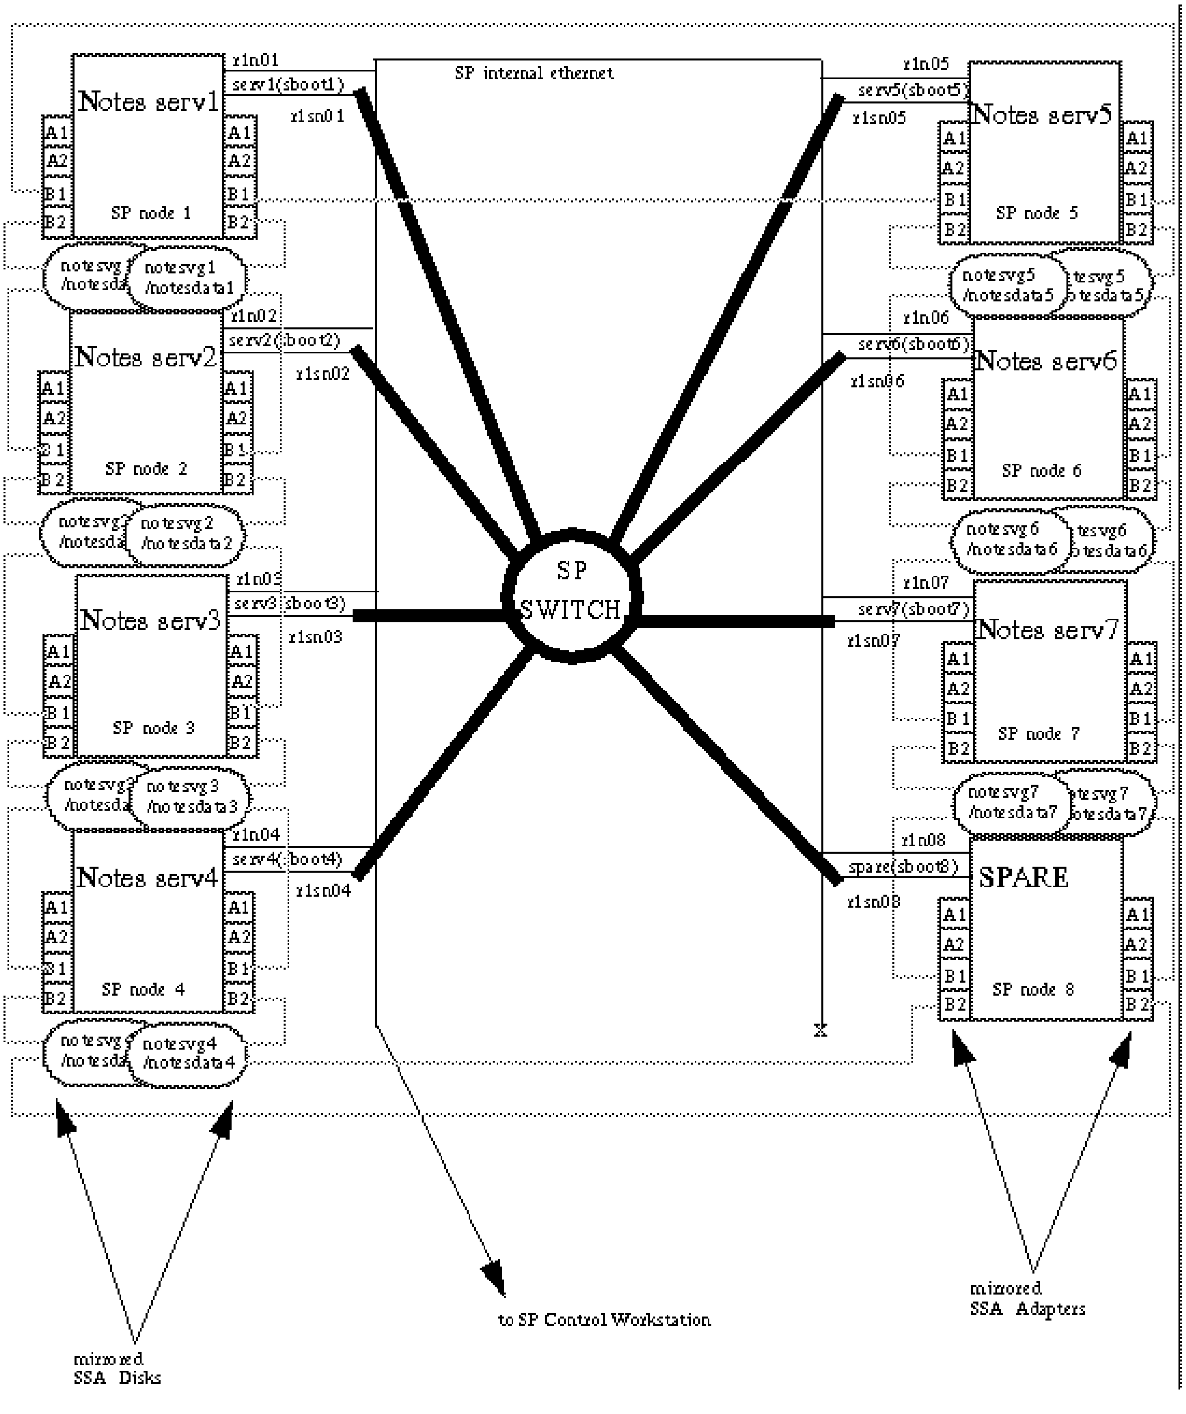

The previous picture illustrates an HACMP cluster configuration where one spare SP node can be a backup server for up to seven (7) Domino Mail or Application Servers. Should any of the servers fail, the spare node will assume the failed node's SP switch IP address and external volume groups and filesystems, and restart the Domino Server.

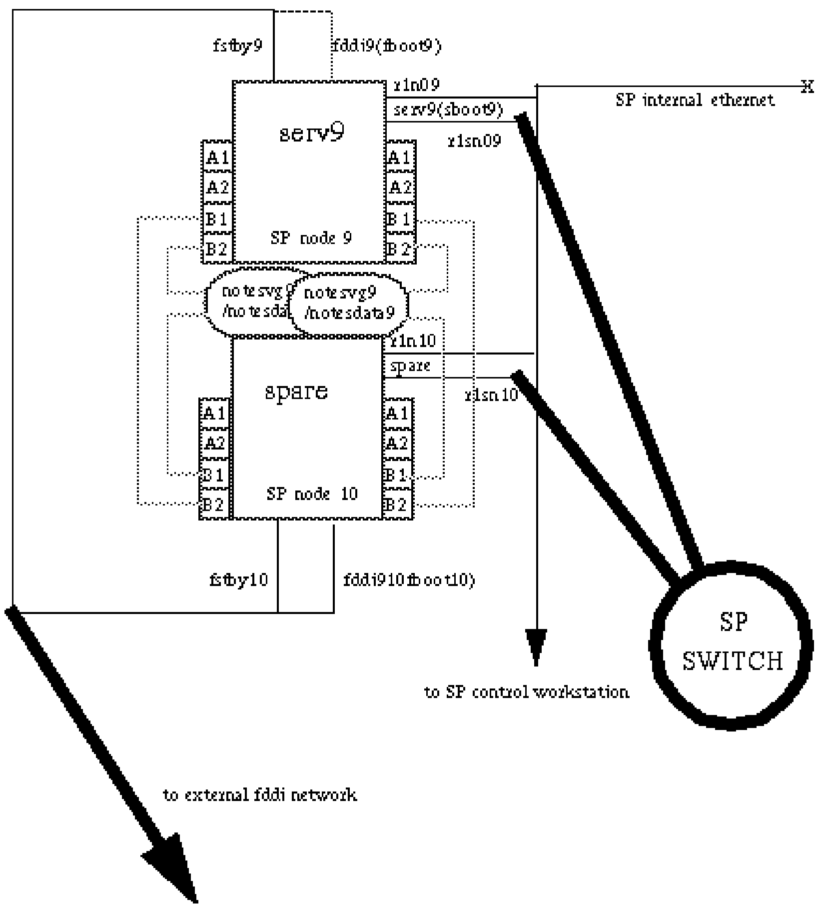

This picture illustrates an HACMP cluster configuration where one spare SP node can be a backup server for a Replication or Mail Routing Server - one that has external network connections. Should any of the servers fail, the spare node will assume the failed node's SP switch and external network FDDI IP address as well as external volume groups and filesystems, and restart the Domino Server.

Note that this configuration shows only one backup per active server. You can configure a backup SP node for up to seven (7) active servers - this will depend on external network requirements and choice of SP nodes.

.

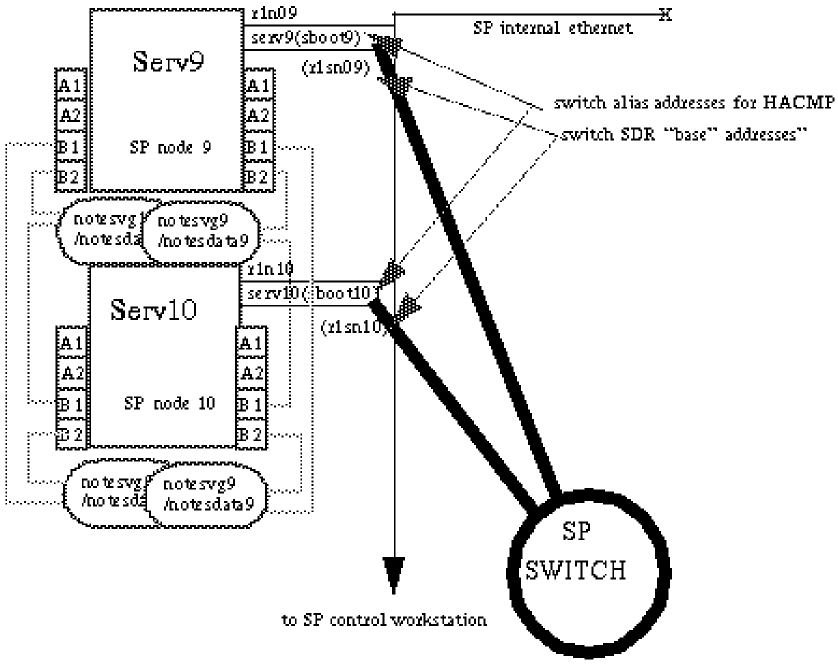

This picture illustrates an HACMP cluster configuration where an active SP node can be a backup server for another server - one that can have external network connections (not shown). Should either of the servers fail, the designated backup node will assume the failed node's SP switch (in addition to its already active address), external Domino volume groups and filesystems, and restart the Domino Server.

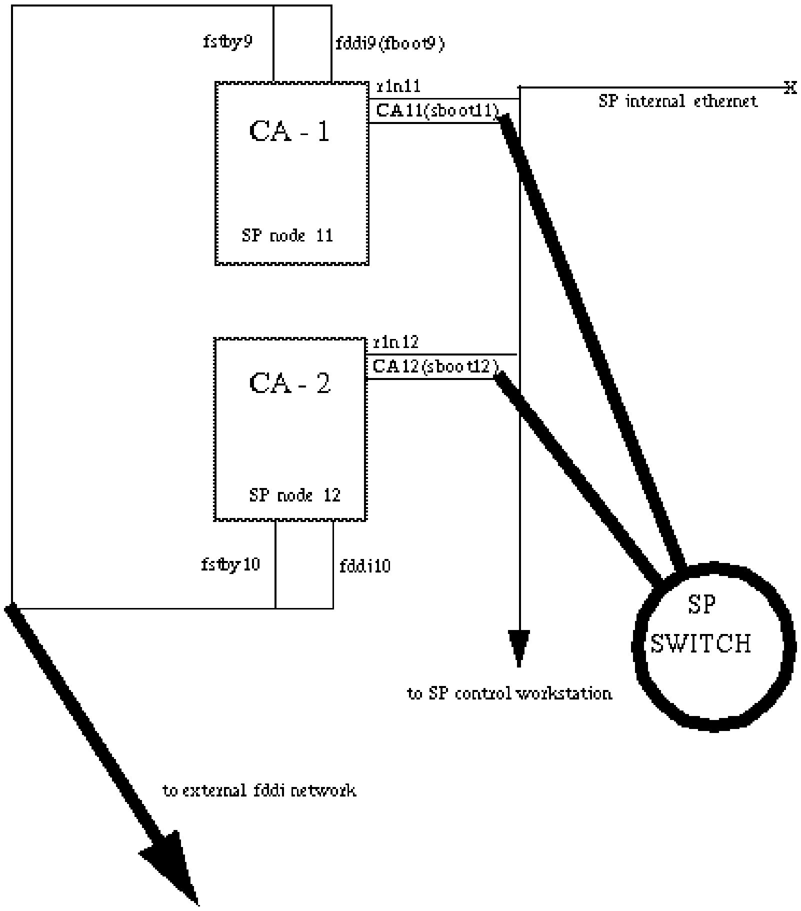

This final illustration shows how external client gateways can be made available via HACMP. Both client access (CA) nodes in the picture are active and are gateways from external networks into the SP Switch. Should either server fail, HACMP will detect the failure and the other server will assume the IP address of the failed server.

Table of Contents Next Chapter

 Search | Contact IBM | Help | © | ®

© Copyright IBM Corp. 1994, 1995, 1996. All rights reserved.
Last modified: Tue Oct 08 09:03:03 US/Eastern 1996

APPENDIX A: Scripts to Start and Stop the Domino Server

These scripts are provided as is and can be used as HACMP start and stop scripts, or to run servers in the background so a dedicated console window is not required.

Start Script (For Version 4.5 Servers):

```
# Lotus Notes Server Startup Script
#
NotesUser=spnotes
DataDir=/spnotes
#
OUTPUT_LOG="/tmp/notes_log."$NotesUser
INPUT_FILE=$DataDir/.notes_input
chown $NotesUser $INPUT_FILE $OUTPUT_FILE
export Notes_SHARED_DPOOLSIZE=60000000
export NSORDER=local,bind,nis
chown -R $NotesUser.notes $DataDir/*
cd $DataDir
su - $NotesUser " -c

PATH=$DataDir:/opt/lotus/bin:$PATH

if [ ! -x /opt/lotus/bin/server ]
then
echo "Notes : Cannot access server command - exiting"
exit 1
fi

if [ ! -d $DataDir ]
then
echo "Notes : Cannot access notes data directory - exiting"
exit 1
fi
```

```
if [ -w $DataDir/.secure_hidden_passwd ]
then
cat $DataDir/.secure_hidden_passwd > $INPUT_FILE
echo " show server" >> $INPUT_FILE
else
echo " show server" > $INPUT_FILE
fi

cp -p $OUTPUT_LOG $OUTPUT_LOG.‘date +"%m%d%y_%H%M"‘
cd $DataDir

echo "Starting Lotus Notes V 4.11a server ... Waiting 25 secs."
/opt/lotus/bin/server < $INPUT_FILE > $OUTPUT_LOG 2>&1 &

#do not delete the quote down below

"

###########################################################################
```

Stop Script

```
#!/bin/ksh
##########################################################################################

# this is the corresponding script to STOP a Notes server which had been started via the
# rc.notes script. The script attempts to QUIT the server gracefully, and waits up to 2 minutes
# If Notes has not shut down by then, the plug is yanked. Depending on your environment, you
# may want to change this time-out period.
##########################################################################################

NotesUser=spnotes
DataDir=/spnotes
INPUT_FILE=$DataDir/.notes_input
#
echo "Stopping the Lotus Domino R4.x server ... Waiting for shutdown to complete."
echo " quit " >> $INPUT_FILE
count=0
NOTES_RUNNING=‘ps -fu $NotesUser | grep /opt/lotus | grep -v ibmpow/notes | grep -v grep ‘
while [[ -n $NOTES_RUNNING ]] ; do
sleep 10
count=‘expr $count + 1‘
echo ".... been waiting for "$count"0 seconds"
if [ $count -eq 13 ] ; then
echo "Two minutes and Notes is still running"
echo "now for the ungraceful method ...."
for i in ‘ps -fu $NotesUser | grep lotus | grep -v grep | awk ’{ print $2 }’‘; do
kill -9 $i
done
```

```
mems=‘ipcs | grep $NotesUser | awk ’{ print $1 $2 }’ | awk -F"m" ’{ print $2 }’ | awk ’{ print $1 }’‘
sems=‘ipcs | grep $NotesUser | awk ’{ print $1 $2 }’ | awk -F"s" ’{ print $2 }’ | awk ’{ print $1 }’‘
for j in $mems;do if [ -n "$j" ] ; then ipcrm -m $j;fi;done
for j in $sems;do if [ -n "$j" ] ; then ipcrm -s $j;fi;done
echo "Notes server TERMINATED"
exit
fi
NOTES_RUNNING=‘ps -fu $NotesUser | grep /opt/lotus | grep -v ibmpow/notes | grep -v grep ‘
#
echo "Stopping the Lotus Domino R4.x server ... Waiting for shutdown to comp lete."
done
echo "Notes server ended cleanly"
```

APPENDIX B: Domino on SP Utilities

These utilities will be available on an as-is basis from http://www.rs6000.ibm.com/resource/technology/notesSPcfg. Note that all commands are intended for use on the SP from the SP control point (control workstation or other control node).

Download utilities will soon be available

SPuser - create a set of AIX users for Domino on respective nodes
Syntax: SPuser (nodes) user dir group grouplist

where:
(nodes) is the list of SP nodes where the user is to be configured;
user is the name of the AIX notes user;
dir is the home directory of the AIX Domino user (notes external disk data dir);
group is the name of the primary AIX Domino group, usually notes; and
grouplist is the list of all other groups the user should be in;

Multi node example: **SPuser 1-3,7 notes /nd notes notes staff** - this would create on nodes 1 through 3 the users notes1,notes2,notes3, with home directories of /nd1,/nd2,/nd3, respectively. Note how SPuser adds in a node number in both the user name and home directory.

Single node example: **SPuser 7 notes7 /nd7 notes staff** would create just node 7’s notes7 user and /nd7 home directory, in group notes and staff. Note that when you specify only one node number, SPuser will not add that number to the user name and home directory.

Note: SPuser copies the file ./default.notes.profile to each users home directory as .profile. /etc/passwd, /etc/group, and /etc/security/passwd are replicated to all of the nodes specified.

pcrio - parallel external filesystem i/o configuration
Syntax: pcrio <filename> or pcrio configure

where:
<filename> is a configuration file showing external disk volume groups, logical volumes, and filesystems for a set of nodes, or configure specifies that a previously entered configuration is to be

implemented on the specified nodes and disks.

The easiest way to understand how this command works is to observe the files in the utilities directories which are sample configuration files for notes filesystems. Below is an abbreviated example;

```
fs.notes---------------------------------------------------------
vg 3-5,16
fs /nd 2,3 250 log mirror isolate automount
fs /nd/m1 4-7 50 log isolate automount
fs /nd/m2 4-7 50 isolate automount
-----------------------------------------------------------------
```

File keyword identifier descriptions:

log: Create a separate lv for this filesystem as it's log, do not use an existing jfslog.
(note that this is the default for the first filesystem created in a volume group).

mirror: This will mirror the filesystem logical volume and log, if specified. Note that there must be sufficient disks specified or the command will fail.

isolate: This will position the logical volume and the jfslog on separate hdisks, avoiding contention between them.

automount: Specifies that the volume group to be created will be configured such that it will automatically be varied on-line when a node boots.

The file shown, fs.notes, specifies a volume group to be configured on nodes 3,4,5, and 16. Logical volumes and filesystems and filesystem logs will be created on the specified disks and be part of the volume group. Typing in "*pcrio fs.notes*" would generate the following output BUT WILL NOT CONFIGURE ANYTHING ON THE NODES.

Volume group ndvg3 on node 3 with default partition size of 8MB (autovaryon)

has logical volume ndlv3 which has 250 mirrored partitions on disk hdisk2
and logical volume ndlvlog3 which has 1 mirrored partitions on disk hdisk3
and logical volume ndm1lv3 which has 50 unmirrored partitions on disks hdisk5-7
and logical volume ndm1lvlog3 which has 1 unmirrored partition on disks hdisk4
and logical volume ndm2lv3 which has 50 unmirrored partitions on disks hdisk5-7

and filesystem with mount point /nd3 with log ndlvlog3 (mirrored)
and filesystem with mount point /nd3/m13 with log ndm1lvlog3 (not mirrored)
and filesystem with mount point /nd3/m23 with log ndm1lvlog3 (not mirrored).

...... and so on for the other nodes........

To properly use this command, issue it and verify the output you see is correct.

Then type in "*pcrio configure*", and verify that your filesystems are created properly.

Note: The default Physical Partition size is 8MB in pcrio. To change this, set the ksh variable DEFAULT_PSIZE=xx (to set it to xxMB).

cexecdir and rmexecdir - create lotus notes executable install directory
Syntax: cexecdir (nodes) mount_point partitions disk, or
rexecdir (nodes) mount_point

where:
(nodes) is the list of SP nodes where the user is to be configured;
mount_point is the filesystem mount point;
partitions is the number of partitions in the filesystem;
disk is the disk in the root volume group (rootvg) where the logical volume and filesystem will be.

Single/Multi node example: **cexecdir 1-16 /usr/lpp/lotus 30** - this would create on nodes 1 through 16 a logical volume and filesystem /usr/lpp/lotus in rootvg. This size of this filesystem would be the physical partition size of rootvg (4MB) * 30 or 120MB.

Note that cexecdir also mounts the filesystem and creates an /opt directory, in anticipation of a notes code installation. (/opt/lotus will be linked symbolically to /usr/lpp/lotus).

cpaging and rmpaging - create or remove paging space
Syntax: cpaging (nodes) size vg disk, or
rmpaging name

where:
(nodes) is the list of SP nodes where the user is to be configured;
size is the size of the paging space in partitions;
name is the name of the paging space;
vg is the volume group the paging space is configured in
disk is the disk (hdisk) that the paging space is created on.

Single/Multi node example: **cpaging 1-16 256 rootvg hdisk1**- this would create on nodes 1 through 16 a logical volume and paging space in rootvg. This size of this paging space would be the physical partition size of rootvg (4MB) * 256or 1024MB. Typical recommended paging space sizes are twice system main memory size.

cfg_server_bg - configure servers so they run in background
Syntax: cfg_server_bg (nodes) notes_user

where:
(nodes) is the list of SP nodes where the user is to be configured;
notes_user is the notes user that can start the notes server
data_directory will be configured as /notes_user.

Multi node example: **cfg_server_bg 7-9 notes /ndata** - this would create on nodes 7 through 9 start and stop scripts /usr/bin/start_notes_bg.notesX (X=node_number) and /usr/bin/stop_notes_bg.X. These scripts rely on files /tmp/notes.input.notesX for server input and notes.log.notesX for server output. Note that node numbers are appended to the user names and data directories in this example.

Single node example: **cfg_server_bg 7 notes7 /ndata7** - this would create on node 7 start and stop scripts /usr/bin/start_notes_bg.notes7 and /usr/bin/stop_notes_bg.notes7. These scripts rely on files /tmp/notes.input.notes7 for server input and notes.log.notes7 for server output. Note that node numbers are NOT appended to the user names and data directories in this example.

If the data directory mount point you specified is NOT the same as /user, where user is the user name you created, you must create a symbolic link to point /user to the mount point of the data directory on your nodes. You can use the *addlink* command to do this. Adding the links will permit the use of the *cfg_server_bg*, *ns*, *nc*, *showme*, and *hourly* commands.
Note: use of cfg_server_bg is required to enable use of the server console monitor, *ns* and the *start_servers/stop_servers* programs.

addlink - create a symbolic link on given node
Syntax: addlink node user /datadir

where:
node is the node on which the link is to be created (can only specify one node);
user is the AIX user for the Domino Server on that node;
/datadir specifies the Domino Server data directory mount point.

addlink creates a symbolic link /user that points to /datadir, so that other commands such as *cfg_server_bg*, *ns*, *nc*, *showme*, and *hourly* can still be run.

start_servers and stop_servers - start and stop servers on given nodes
Syntax: start_servers (nodes) or stop_servers (nodes)

where:
(nodes) is the list of SP nodes where the user is to be configured;

This command starts or stops the servers via the background scripts created by cfg_server_bg.

nc and ns - notes client GUI and notes server console monitors
Syntax: nc (nodes) notes_user display:0 or ns (nodes)

where:
(nodes) is the list of SP nodes where the user is to be configured;
notes_user is the AIX user that is configured as the notes administrator;
display:0 is the X display for the nc GUI output if DISPLAY variable not set.

The nc utility starts the client (notes) on the specified server nodes, popping up the Notes GUI. The ns utility pops up an aixterm window that refects the output from the server output file, /tmp/notes.log.X and input from the term window will go to /tmp/notes.input.X (X is the notes user that started the server process).

Multiple node example: **nc 1-16 notes** would pop up 16 windows, for each of the nodes 1-16, of the notes client GUI. Executing **ns 1-16** would pop up the server consoles for the 16 nodes. Note that with nc, the node number is added to the user name for switching to that user on a node to execute the notes program.

Single Node example: **nc 7 notes7** would pop up the window, for node7's notes client GUI. Executing **ns 7** would pop up the server console for node 7. Note that with nc, the node number is NOT added to the user name for switching to that user on a node to execute the notes program.

The DISPLAY variable set will carry to nc, otherwise you will have to specify it to nc explicitly as (NAME:0, where NAME is an IP name or address).

Note: Make sure you have your PATH statement (in .profile in the home directory of the notes AIX user) correctly identify the path where notes.ini is- this will avoid errors with nc (if this is not configured correctly you might keep getting the installation GUI).

showme - parallel execution and display of multiple notes server commands
Syntax: showme <command> <parse_field>

where:
<command> is the Domino Server command, single quote delimited (if it includes spaces) to be executed in parallel on nodes specified in /hosts.notes. The hosts in
/hosts.notes contain the SP ethernet names of the nodes where the command is to be run;
<parse_field> is the field to select the output by, single quote delimited (if it includes spaces)

For example, specifying "showme 'sh stat' 'server'" will return (for a /hosts.notes that contains the proper node(k72n07) ;

HOST: k72n07

Server.BootID = 1361509
Server.Name = CN=testbc/O=bctest
Server.Path.Data = /nd7/notesr4
Server.Ports = TCPIP
Server.Time.Start = 01/10/97 22:46:55
Server.Title =
Server.Version.Notes = Release 4.11a

hourly - parallel execution and display of multiple nodes AIX and notes statistics
Syntax: hourly

Note: Statistics shown will be gathered with the "showme" command and other AIX performance monitoring commands, and use the /hosts.notes files configured for that command. Output is placed in

the file
/tmp/notes_hourly_data. This command can be configured with the AIX cron command to provide a timely monitor of all Domino servers that are presently running.

Statistics generated are;

- CPU utilization via the AIX vmstat command;
- IO performance via the AIX iostat command;
- Shared Memory Utilization;
- Domino Elapsed server run time and Peak Usage
- Domino Server mail.box status (pending mail)
- Domino Server Number of Current Users
- Domino Server Memory Allocation and pools

For realtime graphical monitoring of AIX and Domino Server outputs, the Performance Monitor (PTX) Agent for the Domino Servers will soon be available (beta at the time of this writing).
