

DALNET Horizon Capacity Planning Report

DALNET planned to upgrade its Horizon servers in Spring 2000 to enable the DPL and WSU server sites to better back each other up, to accommodate the new shared files that were to be delivered with our Horizon enhancements, and to add capacity for future growth. To this end, the DALNET Board set aside an equipment upgrade fund for FY2000.

This report from the DALNET staff is the information we have to date to help us plan for needed growth of our Horizon servers. The report takes into consideration these factors:

1. Current Horizon servers and their utilization (Appendix 1)
2. Predicted growth of DALNET's Horizon system, including (Appendix 2)
 - a. Annual growth of member's databases;
 - b. Addition of new members;
 - c. Annual growth for new releases of operating system, Sybase, and Horizon software;
 - d. Addition of shared files with the delivery of our Horizon enhancements.
3. Backup and recovery needs
(Appendix 3 describes DALNET's current backup plan)
4. System performance needs
(Appendix 4 is a description of Sun's high availability/clustering capability for improved backup, recovery and system performance. Iowa State University is using this capability successfully in their Horizon implementation.)

A number of options for increasing the capacity of DALNET's Horizon servers have been identified. The DALNET staff worked with Tim Hyde, *epixtech*, Anna Golubeva and Jefferi Holland, SUN Systems, and Jim Sands and Daniel McLeod, SBC DataComm, to determine the viability and estimated costs of the options. We then reviewed the advantages, disadvantages and other issues of each option with Sun and *epixtech* in order to get their technical evaluation for the DALNET Finance Committee and Board.

DALNET Staff who participated in this investigation include:

Scott Muir, Project Leader
Jeff Trzeciak, WSU/DALNET
Kerry Sanders, DPL/DALNET
Louise Bugg, WSU/DALNET
Tim Cromer, DPL/DALNET
Tim Turner, DPL/DALNET
George Marck, WSU/DALNET
Martin Austin, WSU/DALNET

Report written by Louise Bugg

September 6, 2000 (rev. October 9, 2000)

Options for DALNET's Horizon Servers

This report describes five DALNET Horizon Server options:

- Option 1:** Two production servers at two sites without clustering
- Option 2:** Two production servers at two sites with clustering
- Option 3:** Two production servers at a single site with clustering
- Option 4:** A single production server
- Option 5:** Centralized and decentralized production servers with clustering
- Option 6:** Single production server for Horizon only
- Option 7:** Outsourcing DALNET's server services

Each option includes production/test/development Horizon server needs, WebPAC/iPAC production and test/development server needs, DALNET Website production and test/development server needs, and DALNET's firewall at both the DPL and WSU sites.

The descriptions of each option include these technical factors, as appropriate:

- Downtime to upgrade/user impact
- Staffing requirements
- System backup to provide best availability
- System performance
- Risk in case of disaster
- Horizon design issues for shared databases
- *epixtech* and Sun experience
- Software requirements
- Server site preparation
- Telecommunications requirements

Chart 1 lists the current servers and their upgrade or replacement in each option.

Chart 2 has spreadsheets with the estimated costs for each option—2A is one time costs; 2B is first year annual costs; 2C is second year annual costs.

Chart 3 is a spreadsheet with staffing costs estimated for each option.

Finally, there is a recommendation as to the two best options from a technical viewpoint.

OPTION ONE—Two Servers at Two Sites without high availability/clustering

Upgrade the current Sun E5500 servers at WSU and DPL to be able to handle all DALNET databases as well as have room for growth.

Upgrade the Sun E450 test server to handle all DALNET databases with room for growth. In this option, it would serve as a test server and as an emergency backup for DALNET's Horizon databases for WebPAC/iPAC only.

Other DALNET server needs in this option are:

- The WebPAC/iPAC test server needs disk
- Need a Sun E230 as a Website production server
- Need a Sun Sparc as a Website test/development server
- Need a Sun E250 for a Horizon development server for operating system upgrades and alpha testing with epixtech
- Need a Sun E250 for backup/test Firewall server at WSU site
- Need a Pix firewall backup at DPL site

Advantages:

1. There would continue to be two server sites, in case of major disaster at either site.
2. The staff are already trained in this model.

Disadvantages:

1. There are differing backup tape systems at each site.
2. Without Sun's high availability/clustering option, one server cannot immediately take over for the other in case of a failure. Backup and recovery would be by copying the files from one server to the other using tape backups, which could take several days to do so that it would only be used in cases of lengthy downtime.
3. This option requires more staff support than a single site, more communication among staff, and more duplication of staff to maintain multiple server sites, multiple firewalls, and telecommunications.
4. Software costs are higher for multiple Horizon servers.
5. Horizon design issues are more complex for shared files in a multi-site server environment.
6. Upgrading the servers would require Horizon downtime for DALNET libraries.
7. There would be continuing maintenance of multiple servers.

Other Considerations:

1. The DPL site will need a generator to have continuing power as at the WSU site.

OPTION TWO—Two Servers at Two Sites with high availability/clustering

Upgrade the current Sun E5500 servers at WSU and DPL to be able to handle all DALNET databases as well as have room for growth. Obtain Sun's high availability/clustering capability to be able to switch between these two sites.

Connect the two sites with fiber that is "leased" from SBC DataComm.

Upgrade the Sun E450 test server to handle all DALNET databases with room for growth. In this option, it would serve as both a test and development server because the two production servers would be backing each other up continuously.

Other DALNET server needs in this option are:

- The WebPAC/iPAC test server needs disk
- Need a Sun 230 as a Website production server
- Need a Sun Sparc as a Website test/development server
- Need a Sun E250 for backup/test firewall server at WSU site
- Need a Pix firewall backup at DPL site

Advantages:

1. Can take advantage of Sun's high availability/clustering capability to improve system availability, release migration and performance.
2. There would be servers at two sites, in case of major disaster at either site.

Disadvantages:

1. There are differing backup tape systems at each site.
2. This option requires a fiber link between the two server sites, which will add to the cost and complexity.
3. This option requires more staff support than a single site, more communication among staff, and more duplication of staff to maintain multiple server sites, multiple firewalls, telecommunications and coordination of the cluster.
4. Software costs are higher for multiple Horizon servers.
5. Horizon design issues are more complex for shared files in a multi-site server environment.
6. Upgrading the servers would require Horizon downtime for DALNET libraries.
7. There would be continuing maintenance of multiple servers.
8. Sun's high availability/clustering capability would require significant training of two DALNET staff, and possibly three to be able to handle two sites.
9. Sun support staff in Michigan do not have experience with a remote cluster configuration.

Other Considerations:

1. Sun would need to inspect each server site and determine needed site preparation. The WSU Computer Center has already had such an inspection for its Sun Center for Excellence and the resulting site preparation is expected to comply with the cluster requirements.
2. The DPL site would need to do the site preparation required by Sun.
3. The WSU Computer Center has a generator for backup electricity in case of emergency.
4. DPL would need to get a generator.

OPTION THREE—Two Servers at One Site with high availability/clustering

House both Sun E5500 servers at a single site and upgrade them to handle all DALNET databases as well as have room for growth. Obtain Sun's high availability/clustering capability to be able to switch between these two servers using shared disk.

Upgrade the Sun E450 test server to handle all DALNET databases with room for growth. In this option, it would serve as both a test and development server because the two production servers would be backing each other up continuously.

Other DALNET server needs in this option are:

- The WebPAC/iPAC test server needs disk
- Need a Sun 230 as a Website production server
- Need a Sun Sparc as a Website test/development server
- Need a Sun E250 for backup/test Firewall server if at WSU site
- Need a Pix firewall backup if at DPL site

Advantages:

1. Can take advantage of Sun's high availability/clustering capability to improve system availability, release migration and performance.
2. This option should provide economies of scale with staff to support a single firewall, less complex telecommunications, and a single server site.
3. The two servers could use the same backup system.

Disadvantages:

1. Software costs are higher for multiple Horizon servers.
2. A cold backup site would be needed in case the single site facility had a disaster.
3. Sun's high availability/clustering capability would require significant training of two DALNET staff.
4. There would be continuing maintenance of multiple servers.
5. Horizon design issues are still more complex for shared files on multiple servers, even at a single site.
6. Upgrading and moving the servers would require Horizon downtime for DALNET libraries.
7. Would need to determine the head end frame relay circuit needs and install at the site selected.

Other Considerations:

1. WSU's Computer Center has facilities to handle both servers with a generator for emergency power.
2. Sun needs to inspect cluster server sites and determine needed site preparation. The WSU Computer Center has already had such an inspection for its Sun Center for Excellence and the resulting site preparation is expected to comply with the cluster requirements.
3. The DPL site would need to do the site preparation required by Sun.
4. The DPL site would need a generator.

OPTION FOUR—One Server

Replace both Sun E5500 servers as well as two Sun E450 servers with a single Sun server with capacity to handle all DALNET databases as well as have room for growth. This server could handle all Horizon, WebPAC/iPAC and DALNET Website production, test, and development needs.

Other servers would be needed for the firewall at the WSU site. A Pix firewall backup would be needed for the DPL site.

A trade in of the two E5500s and two E450s would be included. The remaining E450 could be used for the WSU firewall backup.

Advantages:

1. Software costs are lower for a single Horizon server.
2. Would need fewer staff resources to operate and maintain a single server at a single site, with a single firewall, and less complex telecommunications.
3. Sun accepts trade-ins, which helps reduce the cost.
4. Would have a single backup strategy.
5. Less Horizon downtime for DALNET libraries to migrate to the new server than to upgrade existing servers because DALNET could keep the existing servers operational until all libraries had migrated to the new server.
6. No Horizon design issues about shared files across multiple servers and multiple server sites.
7. Could use this single server for all Horizon functions, including production, test and development Horizon software as well as production, test and development WebPAC/iPAC software.

Disadvantages:

1. Single point of failure with server, so would need backup server strategy for those areas that are not redundant within the server.
2. A backup site would be needed in case the single site facility had a disaster.
3. The new hardware and software would require substantial training for DALNET computing staff to learn.
4. This would be forging new ground with *epixtech* for such a large installation on a single, large server.
5. The Sun E10K server is more sensitive to temperature and humidity conditions than the E5500 server, so that the site environmentals have less tolerance.
6. Would need to determine the frame relay circuit needs and install at the site selected.

Other Considerations:

1. Sun needs to inspect E10K server sites and determine needed site preparation.
2. The WSU Computer Center has already had such an inspection for its Sun Center for Excellence and the resulting site preparation will likely comply with the E10K requirements. Additional electrical outlets may be needed at WSU.
3. The DPL site would need to do the site preparation required by Sun.
4. The DPL site would need a generator.

OPTION FIVE—Centralized and Distributed Servers with high availability/clustering

Authorize one or more DALNET libraries to have their own Horizon server at their local site. Upgrade the centralized Sun E5500 servers to be sized for all DALNET shared files and all DALNET Horizon databases including those stored on Sun servers at member site(s). The centralized servers would use the Sun high availability/clustering capability to share disk and switch between servers.

Upgrade the Sun E450 test server to handle all DALNET databases with room for growth. In this option, it would serve as both a test and development server because the two production servers would be backing each other up continuously.

Other DALNET server needs in this option are:

- The WebPAC/iPAC test server needs disk
- Need a Sun 230 as a Website production server
- Need a Sun Sparc as a Website test/development server
- Need a Sun E250 for backup/test Firewall server if at WSU site
- Need a Sun E450 (or smaller) for a decentralized site
- Need a Pix firewall backup if at DPL site

Advantages:

1. Horizon is designed to run as a single database on a single server, so operation could be more efficient on a distributed server.
2. If a distributed server is down, only that database is impacted.
3. Software upgrades should be easier on smaller servers.
4. This allows new options for members to run their own servers.

Disadvantages:

1. Staff support for maintaining multiple servers, multiple server sites, multiple firewalls, more complex telecommunications will require more duplication of effort, more communication and may have both remote and physical access requirements.
2. Horizon design issues are even more complex for shared files on multiple servers in multiple sites.
3. It would be more difficult to keep the shared databases and the software versions in synch across DALNET libraries.
4. Software costs are higher for multiple Horizon servers.
5. Each server site would need site prep for battery or generator electrical backup.
6. There would likely be differing tape backup systems across the server sites.
7. Upgrading and moving the servers would require Horizon downtime for DALNET libraries.
8. Would need to determine the frame relay circuit needs and install at the site selected.

Other Considerations:

Same as for the single site with a cluster (Option Three).

OPTION SIX—Single Server for Horizon Only

Replace both Sun E5500 servers and one Sun E450 server with a single Sun server with capacity to handle DALNET Horizon production, test and development needs and have room for growth.

Other servers would still be needed for the DALNET WebPAC/iPAC production, test, and development needs as well as for the DALNET Website production and test and for the firewall at the WSU site. A Pix firewall backup would be needed for the DPL site.

The two Sun E5500s can be traded in and the Sun E450 can be recycled to eliminate the need to purchase a computer for the WSU firewall backup.

Advantages:

1. Sybase and other software costs may be lower for a single Horizon server.
2. Would need fewer staff resources to operate and maintain a single server at a single site, with a single firewall, and less complex telecommunications.
3. Sun accepts trade-ins, which will help reduce the costs.
4. Would have a single backup strategy.
5. Less Horizon downtime for DALNET libraries to migrate to the new server than to upgrade existing servers because DALNET could keep the existing servers operational until all libraries had migrated to the new server.
6. No Horizon design issues about shared files Horizon files across multiple servers and multiple server sites.
7. Could use this single server for all Horizon functions.
8. Could upgrade this single server later to add capacity for WebPAC/iPAC and DALNET Website.

Disadvantages:

1. Single point of failure with Horizon server, so would need backup server strategy for those areas that are not redundant within the server.
2. A backup site would be needed in case the single site facility had a disaster.
3. The new hardware and software would require substantial training for DALNET computing staff to learn.
4. This would be forging new ground with *epixtech* for such a large Horizon installation on a single, large server.
5. The Sun E10K server is more sensitive to temperature and humidity conditions than the E5500 server, so that the site environmentals have less tolerance.
6. Would need to determine the frame relay circuit needs and install at the site selected.
7. Would not be able to handle WebPAC/iPAC and DALNET WebSite server capacity on the single server without an upgrade, so would still have multiple computers to maintain.

Other Considerations:

1. Sun needs to inspect E10K server sites and determine needed site preparation.
2. The WSU Computer Center has already had such an inspection for its Sun Center for Excellence and the resulting site preparation will likely comply with the E10K requirements. Additional electrical outlets may be needed at WSU.
3. The DPL site would need to do the site preparation required by Sun.
4. The DPL site would need a generator.

OPTION SEVEN—Outsource

Completely outsource the Horizon site to a non-DALNET company (an applications service provider <ASP>) with sizing, performance, backup and availability issues contractually determined.

Advantages:

1. DALNET staff would no longer have to handle the operational issues, backup and capacity planning, site management, telecommunications, etc. except in the development, management and evaluation of the service provider's contract.

Disadvantages:

1. Highly likely there is no vendor with experience running a Horizon site as large and complex as DALNET's, even *epixtech* or another *epixtech* customer.
2. Couldn't take advantage of WSU's educational discounts with Sun Systems and other software vendors.
3. Would have to lay off DALNET staff.
4. ASP vendor market is still in early stages of development and quite "volatile."
5. Might need to dispose of current DALNET servers.
6. Would have less DALNET "control" of the system on a day-to-day basis and would have to negotiate for special support needs that had not been foreseen to include in the contract.
7. DALNET would need to do an RFP to get costs, which would be a major project.

Recommendations

The two best options from a technical perspective, according to our Sun and *epixtech* consultants, are:

1. Single server; and
2. Two servers at a single site with high availability/clustering.

Both options have, in effect, a “hot” backup in their design, with the ability to share disk between “servers” in case of server failure. In the single server option, the computer is divided into “domains” that perform like separate servers.

Both options are at a single server site, which:

- Requires less staff support than for multiple server sites;
- Has less complex telecommunications;
- Has a single firewall; and
- Has a single tape backup strategy.

Both options have the risk of a single server site, so that a cold backup site would need to be evaluated for disaster recovery requirements.

The single server option would enable a phased upgrade for DALNET libraries, but would require that DALNET keep the old servers in place until libraries and functions were migrated and fully operational.

Appendix 1

DALNET Horizon Server Utilization at WSU

Server	%CPU used	Disk Space used	Real Memory Available	Free Memory
E5500 (Prod)	15.43%	78%	4096M	62M
E450 (WebPAC)	2.83%	53%	2048M	1188M
E450 (Test)	2.20%	73%	2048M	31M

Note: no data available for E5500 at DPL

Data as of 9-5-2000; based on one week's activity during a low use period this summer.

L.Bugg/G.Marck
Sept. 6, 2000

Appendix 2

DALNET Horizon Server Database Sizes for Planning Purposes#

Databases*	Union Catalog+	Shared Patron	OS, etc.	Total
186	110	50	11	357
205	121	55	12	393
226	134	62	13	435

All figures in GB gigabytes

Row 1 is current or estimated figure

Row 2 represents 10% growth for year one

Row 3 represents 10% growth for year two

New members the size of Tyndale (30,000 volumes) require about 1.5 GB of space

*Includes Marygrove and William Tyndale

+Includes the Union Catalog and authority resource files

#Database sizes include mirroring

S. Muir

September 1, 2000

Database Backup and Disaster Recovery Scenarios

Overview:

Throughout the day transactional backups are run against the databases. Nightly backups are run on the database server to the "backup" array within the server. At WSU, ADSM backs up that array to tape and manages the tapes including off-site storage. At DPL, DLT backup is used and copies are stored off-site.

Within the server at WSU there are three arrays - one backup and two data. Each array has fourteen drives that are both mirrored and striped. The server at DPL currently has two arrays with backup and data mirrored and then a test partition that is not mirrored.

Recovering from problems:

Problem: A single database has a problem (i.e., doing an upgrade and an error occurs, have to restore)

Solution: Recover from the backup array directly, no tapes are necessary. This is the simplest of the scenarios. This occurred at WSU with a successful recovery of the data.

Problem: A disk within an array goes bad.

Solution: Veritas (the disk management software on the server) takes over. Disks are mirrored within the server. If a disk goes bad, another takes over. We receive notification that there has been a disk failure and we report it to SUN to come and repair.

Problem: The backup array (fourteen disks) goes bad.

Solution: It will have no immediate affect on the operation of the system. We will call SUN and they will replace it. We would stop transactional and regular backups and reformat the drives once installed.

Problem: A single database array (fourteen disks) goes bad.

Solution: Each array has a front and a back. The front of array one mirrors the back of array two. The back of array 1 mirrors the front of array 2. To restore we'd replace the affected array and resynchronize.

Problem: (WSU only) Both database arrays (fourteen disks) go bad. This will result in considerable downtime.

Solution: We'd restore from the backup array.

Problem: All arrays go bad. This will result in considerable downtime.

Solution: WSU would restore from the ADSM tape backup to the backup array and then restore the databases from the backup array. DPL would restore from DLT.

Problem: WSU needs to restore to DPL

Solution: Restore from ADSM to the backup array on the WSU Horizon server or to the Horizon Test server, if the E5500 is down, and then either tar the backup to DLT tapes using the existing single DLT tape drive, or ftp via frame relay - which hasn't been tested.

Problem: DPL needs to restore to WSU

Solution: DLT tapes would be loaded into the WSU server tape drive and the data would be restored.

Clustering

Clustering is a SUN Microsystems, Inc. method of delivering high-availability. A cluster is a group of servers, each running Solaris, that are interconnected to work as a single, highly available system. Clustering minimizes any server or database downtime because one server can immediately take over (failover) for the other through a system of automatic fault detection and recovery.

SUN claims that clusters can deliver near 100% uptime. Employing this strategy recognizes the mission critical nature of DALNET's systems to its member libraries and their users. Iowa State University is currently using clustering as their choice for high-availability of their Horizon system and is satisfied with their results.

Clustering Services:

- Would enable the two production servers to back each other up continuously so that access to Horizon can continue even if one server fails.
- Facilitates disk sharing so each server can access the other's disks.
- Connects servers together so there's no one point of failure as with a single server.
- Facilitates maintenance since one server can be taken off-line for repair or testing, with minimal impact on services.
- Improves sever management since both servers can be managed as a single entity.

Clustering is accomplished by interconnecting two identical servers via a low-end workstation (console) and network cabling, thereby allowing them to work together. DALNET's existing E5500s could be upgraded to be identical and serve as the base for clustering, to keep costs of implementing clustering low.

DALNET's Sun representatives recommend that it is best economically and logistically for the cluster to be at one site. There are sites with clustered servers in separate buildings, which can be up to 10 kilometers (i.e. 6 miles) apart, but there are no "remote clusters" in Michigan.

Remote clustering requires:

- ATM backbone with fiber optic preferable
- Disk upgrade from A5000s to A5200s
- Longwave GBICs (to connect to the ATM fiber)
- 4 Quad Ethernet Cards (2 per server)
- console system

Sun provides an intensive 1-week training program for clustering in Southeast Michigan.

S. Muir/L. Bugg
Rev. Oct. 9, 2000