

PN-ACK-013

106641

U.S. AGENCY FOR INTERNATIONAL DEVELOPMENT

GLOBAL TECHNOLOGY NETWORK

EVALUATION OF THE GTN

HARDWARE PLATFORM AND TRADE LEAD SYSTEM

June 2000

Prepared by

NetDefender

For

DEVTECH SYSTEMS, INC.
1629 K Street, NW, Suite 1000
Washington, DC 20006
Tel: 202/296-8849
Fax: 202/296-4884
www.devtechsys.com

Contract No. PCE-C-00-97-00002-00
USAID

A

GLOBAL TECHNOLOGY NETWORK

EVALUATION OF HARDWARE PLATFORM

PRIMARY NOTES SERVER

Dell Power Edge 2300

- Windows NT 4.0 SP5
- Lotus Notes Server 4.6a

Technical Specifications	
Microprocessor	
Microprocessor types	dual Intel Pentium II micro-processor with an internal operating frequency of 400 MHz and an external operating frequency of 100 MHz.
Internal cache	32-KB L1 cache, 512-KB L2 cache
Math coprocessor	Internal to microprocessor
Expansion Bus	
Bus type	PCI, ISA bus
Expansion slots	two full-length and two half-length dedicated PCI; two shared, full-length PCI or 8- or 16-bit ISA
Memory	
DIMM sockets	four 168-pin sockets
DIMM capacities	64- and 128-MB unbuffered DIMMs or 256-MB registered DIMMs; must be rated for 100-MHz operation
RAM Installed	512 MB (in 4 chips)
Maximum RAM	1 GB
External cache	none
Drives	
Diskette drive	one 3.5-inch, 1.44-MB diskette drive included with standard system
Tape drive	12 /24GB DAT Tape Drive
SCSI hard-disk drives	Two 9.1 GB HDD 10,000RPM
SCSI devices	built-in Ultra2/LVD SCSI host adapter supporting up to six SCSI hard-disk drives in internal bays; built-in Ultra/Narrow SCSI host adapter supporting up to three SCSI devices in externally accessible front bays
CD-ROM drive	one SCSI CD-ROM drive included with standard system
Ports	
Externally accessible:	
Serial	two 9-pin connectors
Parallel (bidirectional)	one 25-pin connector
Video	one 15-pin connector
PS/2-style keyboard	6-pin mini-DIN
PS/2-compatible mouse	6-pin mini-DIN
Server management bus daisy-chain connectors	two modular 8-pin connectors
Internally accessible:	
Ultra2/LVD SCSI controller	68-pin connector

Ultra/Narrow SCSI controller	50-pin connector
Diskette drive	34-pin connector
SDS_SMB connector	6-pin header connector
Video	
Video type	ATI RAGE PRO AGP video controller; VGA connector
Video memory (standard)	2 MB
Power	
AC power supply:	
Wattage	300 W
Voltage	115 V at 60 Hz/230 V at 50 Hz
System battery	CR2032 3-V lithium coin cell
Physical	
Height (with support feet)	43.9 cm (10.3 inches)
Width	26.1 cm (17.3 inches)
Depth	59.2 cm (23.3 inches)
Weight (maximum configuration)	25 kg (55 lb)
Environmental	
Temperature:	
Operating	10 to 35C (50 to 95F)
Storage	-40 to 65C (-40 to 149F)
Relative humidity	8% to 80% (noncondensing)
Maximum vibration:	
Operating	0.25 G at 3 to 200 Hz for 15 min
Storage	-0.5 G at 3 to 200 Hz for 15 min
Maximum shock:	
Operating	6 shock pulses in the positive and negative x, y, and z axes at 50 G for 2 ms
Storage	6 shock pulses in the positive and negative x, y, and z axes at 92 G for 2 ms
Altitude:	
Operating	-16 to 3048 m (-50 to 10,000 ft)
Storage	-16 to 10,600 m (-50 to 35,000 ft)

This server is scaled and sufficient for the immediate need. Moving the system to a web browser-based system, as suggested, should lengthen the service life of this equipment.

SOFTWARE RECOMMENDATIONS

Improved R5 Performance

The Domino Server has had phenomenal market success, due in part to its unmatched performance and scalability. Internal Testing and Benchmarks have shown some major improvements over release 4.6 including:

- Memory usage reduced by 30%
- I/O utilization 10-20% better

- Response time improved by 75%
- More users per server - less traffic
- Reliability, availability, scalability

Your production results will vary based on your production environment.

HARDWARE RECOMMENDATIONS

The only hardware recommendation for this server is to set up a fault-tolerant hard disk system.

RAID

Short for Redundant Array of Independent (or Inexpensive) Disks, RAID is a category of disk drives that employ two or more drives in combination for fault tolerance and performance. RAID disk drives are used frequently on servers but aren't generally necessary for personal computers.

There are a number of different RAID levels. The three most common are 0, 3, and 5.

Level 0: Provides data striping (spreading out blocks of each file across multiple disks) but no redundancy. This improves performance but does not deliver fault tolerance.

Level 1: Provides disk mirroring.

Level 3: Same as Level 0, but also reserves one dedicated disk for error correction data. It provides good performance and some level of fault tolerance.

Level 5: Provides data striping at the byte level and also stripe error correction information. This results in excellent performance and good fault tolerance.

For this system RAID level 1 or level 5 would be the most appropriate. With either, the total amount of hard drive space will decrease with the additional fault tolerance. RAID 1 will take the hard drive space to half the usable space because it will make a mirror on the second drive. RAID 5 will require the purchase of an additional drive and the overall drive space will again be reduced to a subset of the total physical space due to the addition of the fault tolerance.

With RAID 5 in place an entire drive could fail and the system should not be interrupted. Additionally, the type of RAID should be investigated. Since Windows NT 4.0 Server, the current operating system, will provide software RAID and may prove to be adequate for the current need, a hardware solution is the preferred solution. Dell sells a variety of RAID controller cards that will operate very well with the existing core server.

Primary Web Server

Custom-Built Computer

- Dual PII 450Mhz
- 256 MB Ram
- four 2.0 GB Hard Drive
- Windows NT 4.0 SP5
- Internet Information Server 4
- Lotus Notes Client 4.6a

This server's primary tasks are Web Server (usgtn.org) and Lotus Notes Client for replication duties. Again, a fault-tolerant system could benefit this machine, though it is less critical than for the primary Notes Server. The primary function of this server could be duplicated easily with an off-site server using NT replication or any third-party website mirroring application. There are too many to list here and many are very good applications. The choice as to which product to use should be based upon other applications on the machines, security, and reliability.

Standard Workstations

Micron Workstations

- Pentium 133-166
- 32-64 MB Ram
- Windows NT Workstation 4.0 SP5

These machines are severely under-powered. When utilizing a machine for Lotus Notes, various word processing applications, spreadsheets, databases, and standard office applications while logged into Banyan Vines or Windows NT networks utilizing Windows NT Workstation, these machines may perform poorly as compared to a standard corporate workstation. A standard corporate workstation configuration is suggested by Dell.com for a company of 400 or fewer employees.

Dell Dimension 220

- PIII 600-933
- 64 MB Ram
- 9 GB Hard Drive

GLOBAL TECHNOLOGY NETWORK

EVALUATION OF TRADE LEAD SYSTEM

BACKGROUND

The primary trade lead and communication system is built around a Lotus Notes 4.6 platform and consists of several interdependent database modules whose elements are accessible via Notes' client software or through a web browser. The main databases include the following:

- Trade Lead Submission Database
- Buyers Database
- Company Registration Database
- Registered Company Database
- Trade Lead Dissemination/Tracking Database
- Trade System Codes Database
- Email/Calendar Database
- User Group Address Book Database

Each database provides a specific function supporting the overall trade lead process: development, submission, quality control, matching, dissemination, tracking, and program communications. There are two interface environments through which features of this system are made available to users: the Notes Client and the World Wide Web (using standard web browsers). In some cases additional programming code is necessary to extend client features to the web environment; however, not all features are extendable to the web environment.

The Lotus Notes Pump server software is used throughout the system to seamlessly move and share information between databases within the system. Notes Pump operates in the background to ensure that information is moved from one database to another in a timely and proper manner. Individual databases can be viewed as containers, which hold various amounts of function-specific data in document form. These documents can be viewed and categorized in a variety of ways. All of these databases reside on a Lotus Notes Domino 4.6 server. The GTN system is designed to support three user groups—overseas technical representatives, U.S.-based trade analysts, and domestic small and medium-sized firms.

TRADE LEAD SUBMISSION DATABASE

This database contains a template form for trade lead development and submission and is linked to both the buyer database for importing repetitive buyer information and the codes database for assigning appropriate GTN codes for new leads. New trade leads remain in the trade lead submission database until they are approved or deleted. This database provides the option to review draft, pending, and approved leads. It also provides, on the basis of user rights, the option to edit, approve, and/or disapprove leads. If disapproved, the lead remains in this database until the lead is completed and/or deleted. Once approved, lead documents are moved to the trade lead tracking system via Notes Pump for matching and dissemination.

BUYERS DATABASE

The buyers database contains buying company information organized by country.

- There are three ways to enter data into the buyers database. First, by entering a lead into the trade lead submission database, a buyer profile is automatically stored into the buyer database making it available for subsequent leads by that company. Second, a buyer profile can be manually created by entering data into the buyer company database via the Client or through a web browser. Third, databases of companies (buyers) can be mapped and imported into the buyer company database.
- The primary use of this database is to make buyer information available for use within the trade lead submission database and to provide an exhaustive list of buyers, by country, that have used the GTN system.

COMPANY REGISTRATION DATABASE

This database is accessible via the Notes Client and web browser. Companies may register with GTN's system by completing and submitting the form. The form remains in the registered company database until it is reviewed, approved, and/or deleted. Once approved, the company information is moved via Notes Pump to the registered company database. If disapproved, the company remains in this database until the company is contacted and approved or deleted.

REGISTERED COMPANY DATABASE

The database contains company profiles on approximately 16,000 U.S. and foreign firms. The main purpose of this database is to provide a marketplace for firms to be matched and sent trade leads electronically in a targeted manner. This database is fully indexed and provides the following features:

- Key word, code, and form searching
- Creation of trade lead communications data-sets (email and fax)
- Generation of outgoing communications to companies regarding specific trade leads
- Companies are added to this database as they are approved through the company registration database via Notes Pump
- While this database is accessible via the Client software and a web browser, the majority of the features in this database are only accessible via the Client.

TRADE LEAD DISSEMINATION/TRACKING DATABASE

This database is a repository of all approved trade lead documents. Currently there are approximately 3,000 approved trade leads in this database. Trade lead documents are viewable from a web browser and the Client based on user access rights. There are two types of documents in this database—trade leads and company datasheets.

The trade lead is the document that was pumped from the trade lead submission database and is a “parent” document. There is a “child” document associated with this document, which is the trade status document. Specific features of the trade status document are as follows:

- Assigns lead to specific users for follow-up
- Notifies targeted audiences of actions taken
- Captures updated follow-up information by users (date, time, and user stamp) in chronological order
- Generates System Generated Emails (SGEs) to notify a targeted group of lead updates

The company data sheet documents are also associated with particular leads; however, they are also parent documents and for a given lead there can be several data sheet documents, all of which are parent documents. This child document is the company datasheet status document. Specific features of the company datasheet status document are as follows:

- Assigns company to specific users for follow-up
- Notifies targeted audiences of actions taken
- Captures updated follow-up information by users (date, time, and user stamp) in chronological order
- Generates SGEs to notify a targeted group of updates. These updates include SGEs to a Letter of Interest group, a follow-up group, a grants group, and a finance support group.

TRADE SYSTEM CODES DATABASE

This database contains GTN field office codes, GTN technology codes, SIC code, and NAICS codes. The trade lead submission database, the registered company database, and the company registration database use this database. The functions of the codes are as follows:

- Field Office codes – Each overseas office is identified with a unique two letter alpha code that identifies the trade lead source.
- GTN technology, SIC, and NAICS codes – These codes are used for trade lead definition and matching purposes.

EMAIL/CALENDAR

The system supports a wide range of email functionality in the form of System Generated Emails and scalable broadcast emailing. Lotus Notes Mail is the email system used by the project. Notes Mail also provides the project with an on-line (Client and web browser) calendar that provides personal and group calendars as well as the management of resources such as conference rooms. As a component of the email system it allows for email meeting invitations and reminders.

USER GROUP ADDRESS BOOK DATABASE/SECURITY

System security is controlled in the address book database and users are provided with access rights to various levels including read, author, and editor.

FINDINGS

Import/Export

At this time, the process requires programmers' assistance because there is not currently an import/export function built into the system. The problem is complicated further if the data to import is not in a format that is ideal for the application it is to be imported into. For example, if the data is currently held in a relational type database, it does not always smoothly import into a flat file type database. With Lotus Notes you are working with a proprietary database that is not structured even in a normal flat file format. It stores records as if they are individual documents. These documents are similar to input screens in more traditional databases.

The challenge, then, is to take data from a completely alien format and convert it into a format that Notes will understand. This can be accomplished by using several of the advanced features of Lotus Notes. Another problem is that many of these advanced features were in early development stages for Notes 4.6 and were mostly external, whereas in Lotus Notes Domino 5.X they are integrated into the core application. A few examples of these applications are Notes Pump and Notes SQL. Both of these applications were available as external applications, but were difficult to set up and even more difficult to manage. The creating of import templates and using triggers and agents to format and migrate the data through the application would also assist in the automation of the import process.

Movement of Data Through the System

The current system is fairly intensive on human management. Data must often be moved from place to place manually. This was an important step in the past when the system's architecture was changing and evolving, to ensure data integrity. It is now a weakness, since data input into one part of the system requires approval and then manual movement to another. This step is often overlooked, and results in the same problems it was designed to protect against. This process should be simplified for the Trader Staff, so that they verify the data is accurate and then simply hit a button to approve the addition to the main database. The present system does not remove the records from the pending queue. This queue must be cleared by hand. Though this does offer a certain amount of redundancy and backup, it was not created for that purpose. If a backup or rollback is required, then such a process should be created in the existing workflow. It should move data from the areas where backup is required or desired, and automatically archive this data off the main system at set increments to avoid excessive use of drive space. Databases that are currently affected by this issue are the registered company, submission, and company registration database. All of the databases, however, could benefit from the archiving of older data. This archive space could be a tape drive, a CD-R, or any drive not used in production of the Lotus Notes Server.

BACKUP AND RESTORE

As mentioned earlier, a backup procedure needs to be enacted for the existing system. This backup should be done in several levels and steps. Most, if not all, of these steps can be automated. A full backup of the core files should be performed via a tape drive and these tapes kept in a secure off-site repository. There are commercial off-the-shelf (COTS) backup solutions that will provide adequate tape archives of the databases. Examples of these COTS solutions are Arcadia Backup-Exec and Computer Associates' ArcServe. Both of these options are Windows NT-based, menu-driven, automated backup utilities. Additionally, for systems that have a very high uptime requirement, such as this application, a second server is often suggested. This can be a very costly requirement, but a single catastrophic failure could put the system down for hours, if not days, by a simple hardware failure, power failure, or misconfiguration. With these failures in mind, often the best choice for a second server is locating it at a different physical site than the primary server. The secondary server simply replicates at a given frequency all of the pertinent databases. The replication should also only be set to "PULL" until such a time that the second server becomes the primary. This process of switching between the two can be accomplished at the ISP (Internet Service Provider) level rather than at the local site.

LOAD BALANCING/FAILOVER

As the system continues to grow, load balancing may prove a required function. This too can be accomplished in several ways. Currently, the systems are hosted on Window NT 4.0 Server-based systems. Load balancing and clustering can be accomplished using Microsoft Tools or other third-party clustering tools. As with the suggestion for Backup and Restore, using multiple computers at separate locations can be used for load balancing as well, with the traffic management falling on the ISP rather than the local staff. This is often more desirable since clustering is a very advanced application and can be extremely difficult to manage. Failover is another benefit that could be gained by having multiple servers. The GTN Network as it exists today is hosted on a single server, on a single Internet connection from a single ISP, in one building, all of which are on the same power grid. This puts a great deal of faith in areas that cannot be controlled by GTN Staff for a service that now has a worldwide footprint. Currently, if the Ronald Reagan Building has an emergency of some type that causes a power interruption or if the ISP experiences a failure in their systems, a worldwide service would be unavailable. By creating a mirror of the system on another machine at a different location hosted on separate bandwidth (ISP), an emergency failover system could be enacted to provide for the times when a failure occurs at the primary site. The server located at the alternate site could be live, taking submissions as an added resource at all times to assist with the load, or it could be dormant, waiting for its need. Lotus Notes has a feature called replication that could keep the databases in synch with one another.

DOCUMENTATION

Currently, there is well-designed, well-written documentation for the end user and even the Trader Staff. The shortcoming of the documentation seems to fall on the administration and support of the system. This system, as mentioned before, is a fairly complex layout of interconnected databases that mimic the abilities of a relational database with the strength of Notes WorkFlow and Groupware. To provide technical support for end users, a good understanding of the system must be gained, as well as a working knowledge of Lotus Notes in general. But for the administration of the system, with the system as is, it will take advanced Lotus Notes administration experience, some Notes development experience, experience in Lotus Script, intermediate NT administration experience, and TCP/IP networking experience.

REPORTS

The reporting provided by the system currently depends upon some very creative Notes development and some imagination. Fortunately for the system, there was never a paper version of the system prior to its creation so the reporting requirements were not stringent to start with. At this point, the system has been running for more than a year. There have been requests for reports that the system does not provide. Also, standardized forms and customizable reports have been identified as desirable features. These features are available through the Lotus Notes Client using Notes Reporter. The implementation of Reporter could be most smoothly integrated into the existing system by creating new reports using the product and gradually migrating the existing reports to it. The existing reports may be reevaluated for purpose and content.

WEB BROWSER VS. NOTES CLIENT

There has been a great deal of care taken in the development of the forms so that they are usable from both clients. The only problem noted on the input forms as a whole is the length. The length does not play in to the equation as much from a Notes Client approach because of the use of subforms and spinners, which are not of the same value on a webbed approach. Lookup tables and code tables could be made into Java Window pop-ups or some other CGI form to allow a shorter form that flows better. Currently, to fill out a section you have to go to a section that is farther down the page. This is done to allow for better aesthetics, but unfortunately it adds to confusion and disrupts flow. In addition, there are some operations that are exclusively carried out via the Notes Client interface, such as trade lead matching. There are advantages to having these processes available to the web browser clients as well, some of which are: expanded system use, cost of the Notes client, support time for a remote user using the Notes Client for Notes-related issues, more rapid deployment, and simplification of the training to teach one interface. The functionality required to perform these processes is currently available to some extent in Notes 4.X but has been expanded greatly in Notes Domino 5.X. Matching could be developed to work in a popup Java window for instance. The development of these features should be explored and planned to offer the same functionality without adding undue complexity to an already complex system.