

Mail Order

Setting up an enterprise mail system is a monumental task, especially when it involves migrating from data based on proprietary systems to a totally new infrastructure based on open source. But the IT team at GTL managed to do this complex task, smoothly and efficiently.

Messaging and Internet services are vital to the day-to-day operations of most organisations. Just imagine a day without them! Often, communications delivered late can lead to disastrous consequences. Today, technology provides us with a number of options in the form of proprietary and open source solutions.

Let's take a look at how GTL has improved its mail server systems in order to future-proof them in line with organisational needs.

THE ORGANISATION

GTL Limited is a well-known network engineering and IT services company, with branches all over the globe. Being a solution provider, it has invested heavily in the latest technology and has been innovating constantly in order to give the best services to its clients. GTL's IT set-up is entirely managed and operated in-house.

Says general manager, technical support enterprise solutions, Ashok Rumde, "We plan and test every deployment carefully and choose the right approach depending on the complexity involved. Usually a simulation is done before porting or implementing solutions and applications."

THE NETWORK

The GTL network covers 21 locations across the country and spans 11 international locations via Virtual Private Networks (VPN). All major operations are remotely controlled through its main hub located at the company's headquarters at Mahape, Navi Mumbai. In line with today's requirements, GTL has made provisions for its

services to be available on a 24x7 basis. GTL has also provided for redundancy through its disaster recovery centre located at Pune.

Prior to migrating to Linux, GTL maintained an assortment of applications and databases for their messaging and Internet services. This included the Microsoft Exchange 2000 server on the Windows 2000 platform to provide SMTP, POP and IMAP services for their mailing requirements.

SMTP, POP and IMAP, being critical services, require backup servers. Hence, GTL standardised on clustered servers with an external data storage approach for storing mail information such as mailboxes, mailing lists and LDAP databases. This would also allow GTL to scale up to meet its projected expansion targets.

LDAP was seen as a high-performance, open-standards solution, which would meet the architectural requirements of the company's mail system services. The initial operation had 1,800 entries, but this is expected to grow to over 3,500 in the near future.

THE GROUNDWORK

GTL's IT team, in consultation with Red Hat, conducted an IT audit of the existing mailing infrastructure and evolved an approach involving an open source technology based solution. Table 1 summarises network services in use on GTL's network for internal use based on the audit and lists the software packages that were proposed to provide each service after migration.

Customer services and services affecting end-user productivity have been declared critical services. The column labeled 'critical' indicates



which services should have redundant and certified hardware in the data centre and should be configured so that in the event of a service failure in the main server, the service would be up and running again from the backup server almost instantaneously.

Table 1

Service	Current Software Package	Critical
SMTP	Qmail-1.03	Yes
IMAP	Courier-imap-2.1.2	Yes
POP	Qmail-1.03	Yes
LDAP	Openldap servers 2.0.27-2.7.3	Yes
File & Print	Samba 2.2.7a-7.9.0	No
Proxy	Squid 2.4.STABLE6-6.7.3	No
DNS		Yes

UNIQUE FEATURES

By building on its core features of remote administration, anti-spamming support, messaging and security, the latest version of the software—the Qmail system—could be integrated with new open source Web calendaring solutions. In addition, SSL authentication policies create a more secure method of sending and receiving e-mail. Enhanced Web-based administration provides administrators with the ability to view real-time server activity, cutting down on spam and enhancing capacity planning.

Vilas Wagle, senior manager—IT, clarifies, "Through the implementation of the Qmail server, we have been able to communicate effectively and our mail system has become very reliable. Besides, users are very pleased as they experience faster downloading of mail and can access their accounts through the Web".

He adds, "Coming from a Microsoft e-mail environment, end users were accustomed to several features used in Outlook/Exchange. So we felt we needed to make as many of those features available as possible".

MIGRATING E-MAILS FROM WINDOWS TO QMAIL

For this migration, new mail servers in cluster configuration were first created

The hardware/software setup at GTL

GTL hardware:

- Dell Power edge 2580 (2 Nos) (running Q-mail, web-mail server, IMAP, POP, Slave LDAP)
- Dell Power Vault 2208 (Storage)
- Dell 2450 (2 Nos.) (DNS)
- HP - LPR (LDAP Master & Web service & File & Print)
- HP - LHR (Slave LDAP & CHAT)
- Compaq Proliant 5500 (Proxy & Slave LDAP)
- Pentium IV machines

GTL Software:

SMTP—Qmail, Courier IMAP/POP, Auto-responder, Spam Assassin, Qmail-scanner, Razer 2, AV, PGP, IMP Herda, SOAP Web services, Password Migration utilities and tools; Web-based administration—Webmin; Proxy—Squid; File and print sharing—Samba; Directory—OpenLDAP and Mailbox Migration scripts.


in active-passive mode. Red Hat Enterprise Linux AS 2.1 was deployed, OS hardening was carried out and all the necessary mailing solutions consisting of Qmail, spam control and anti-virus were deployed and tested.

The migration of the entire mail solution revolved around two major systems, the LDAP authentication system and the mail server that consisted of the solution for POP, IMAP, SMTP services, AV, queue management and the mail store. Since mail data and user data is continually changing, de-linking between LDAP and mail-store migration processes is almost impossible. Hence, the team had to consider both LDAP and mail-store migration as one integral process.

During the migration stage, it was decided to clean up the mail user data of any errors and bad records. So, the user data from the active directory was extracted in text files and used as the source for migration scripts. Moreover, Perl scripts based on the IMAP protocol were used to synchronise the data from the old mail server to the new mail server as well as data from the active directory to the LDAP server. In the migration process, these Perl scripts also transferred the user's data from mailboxes, including default as well as user-specific IMAP folders.

(The Perl scripts created necessary 'Maildir' structures on the Linux-based Qmail server along with any user-specific folders).

Network set up at GTL



This module was extended to migrate the passwords smoothly to the Linux-based LDAP database by integrating some PHP code with an Apache Web server. A similar Web service based on SOAP protocol was also written on the Windows side. The process involved requesting users to change their passwords through the Web interface and synchronising the password in both Windows and LDAP systems at the same time. Says Vilas Wagle, "Through this process, we also got users to strengthen their passwords."

What were the other steps involved in the process? It is quite obvious that in the case of complex systems, log management is an important task. Mr Wagle says, "We noted that in case certain log files grew beyond limits, the system could crash. To avoid such a problem, we set up all log files to be governed by appropriate log file rotating scripts using standard as well as special scripts."

PLANNING ACTUAL MIGRATION FROM WINS TO DNS

User Access Planning was a very important task in itself. GTL discovered that the user's PC configurations needed a few changes.

So how was this done? The user's PC configurations were based on the WINS name resolution. The GTL team had to change the name resolution configurations on the PCs from WINS to DNS, making sure that all users could specify fully-qualified domain names (FQDN) to access all mail related services whether they were on LAN or WAN (Internet). To achieve this, Vilas Wagle states, "We had to configure a DNS server, which would specifically resolve queries coming from LAN-based users to internal IP addresses. This helped us to migrate users smoothly to a new scenario where they could access their mails over any LAN or WAN seamlessly without having to change their PC

configurations when they moved between the corporate LAN or WAN."

The above configuration was deemed necessary as GTL users worldwide use various methods to access mails. The access methods include Web-mail using intranet/Internet-based URLs both pointing to the same Web-server, POP-based access and the IMAP-based access via Internet-based and LAN-based Fully Qualified Domain Name (FQDN).

AUTHENTICATION

Allowing speedier access to the users was an important consideration. In view of this, making sure that authentication happens quickly was necessary.

- To have quick authentication, GTL used multiple LDAP slave servers and one master LDAP server. The master server was used to update any changes on the LDAP database and multiple slave servers would allow access to users and various other processes for authentication.
- For the slave LDAP servers, GTL used non-dedicated servers. It was ensured that enough bandwidth was available for the replication of the LDAP data from master to slave servers in the smallest possible

time. This also ensured better availability of the LDAP service to the users. A scheme of IP addresses (linked to FQDN) was clearly defined for internal as well as external access and private administrative and management access for all servers.

IMPORTANT STEPS TAKEN DURING MIGRATION

- A backup copy of the existing mail server (especially the mail store) was created and the backup mail store was used as a source to actually transfer the user's mailboxes, keeping the real mail server intact. Keeping a backup of the entire solution and intermediate scripts along with various versions was also routine for the entire duration of the project.
- The entire operation had to be completed with zero effective downtime for the external world. This ensured that no incoming mail was rejected due to the non-availability of the mail server. Hence, the GTL team had to set up a mail relay server with a considerable memory to store all the messages received until the

What's done at the GTL data centre?

Implementation of enterprise CDN using open source Squid

At the data centre, engineers have used Squid in conjunction with some smart routing policies in order to implement a CDN (Content Delivery Network). In this typical scenario, a Squid server is placed at each major GTL location and Web traffic is transparently diverted to these servers. At regular intervals, Web content from specified sites is pulled into these caches in order to avoid the lookup delay that would be incurred by regular users.

Implementation of open source VoIP solution

GTL implemented the SIP-based VoIP server software from the open source Vovida project. A central SIP server—Session Initiation Protocol (SIP)—routes calls from SIP clients across the organisation. The server was modified in order to restrict the number of calls that can be made from a specific GTL location.

File storage appliance

Engineers at the data centre have deployed a file server based on the Linux LVM for volume management on top of hardware RAID. (Logical Volume Manager allows several physical partitions to be represented as a single block device, amongst other things). In this scenario, XFS has been used as the file system as it offers journaling and ACL features. The ever-popular Samba server does the file serving. This server has joined the existing Windows 2000 domain, thus authenticating users transparently without requiring them to have separate passwords. GTL has gained as it has allowed them to have consolidated storage for user files and also provide scalability at the same time through the use of the LVM. GTL has four Squid servers in four different locations. These are very fast because of the caching done on them. They are so robust that they have not been rebooted for the last two years.

What is a data centre?

A data centre is a centralized repository for the storage, management, and dissemination of data and information organised around a particular area or body of knowledge.

This kind of data centre may contain a Network Operations Centre (NOC), which is a restricted access area containing automated systems that constantly monitor server activity, Web-traffic and network performance and report even slight irregularities to engineers so that they can spot potential problems before they surface.

migration process was completed, that is, almost 18+ hours.

- It was also necessary to make sure that the migration scripts ran in the background and were not associated with any terminal or any monitoring PC so that even an accidental shutdown of the PC would not result in the stopping of the process. For this, in order to divert mail traffic from the live mail server to the relay server, the team had to change the priority of the mail exchanger records appropriately.

AT THE DATA CENTRE

Besides being a network engineering, enterprise solutions and customer management solutions company, GTL is also a Category A Internet Service Provider (ISP) that sets high standards for itself and its products. This starts with the performance of its Internet connections and ends with providing secure e-mail accounts.

The best example of this is GTL's own Internet network. In order to provide customers with the highest transmission rates, the network is directly connected with critical Internet intersections all over the globe.

One of the executives, Komal Tekchandani, at the GTL data centre is modest in her approach when she talks about the Qmail applications and the work going on at the centre. "Currently, there are 14 Linux servers all over the country running on Red Hat enterprise Linux with various versions. Most of them are Compaq servers with multiple CPUs with 1 GB RAM."

She further clarifies, "Our Qmail application is unique as it has support for multiple domains/virtual domains, has very good Spam control, is virus

scanning enabled and has a rich Web-user interface. Besides, there is a patch for 'Forward Mails To' option (for roaming users) that comes with Qmail. This does not work by default. For this particular part, we modified the PHP code of the Qmail and made it work. Currently, there are many customers on this messaging setup and it's quite stable."

For Web-hosting it has Apache as a Web-server, JBoss-Tomcat as an application server and MySQL as the database. To enable a trouble-free environment, all machines are behind CISCO PIX firewall, where spam control is enabled and IP-based website hosting is done for enhanced security. Being a data centre there is heavy Net traffic and things need to be streamlined in terms of bandwidth management.

She adds, "We make use of the CISCO Netflow report, which provides valuable information about who is using the network, what applications are being used, when the network is utilised and where traffic is going on the network." There was a typical case of a GTL customer from Kolkata who complained that their bandwidth was choked and they were getting a very slow response on the network. How was this problem rectified?

She says, "After analysing with Netflow, it was found that the data transfer was almost 2 GB per hour and

the maximum traffic was found to be on the ICMP port. When the customer blocked that port, the network traffic was controlled. The customer was happy that his services were working perfectly and that there was no loss of bandwidth."

The Netflow system runs on Linux with MySQL database and is integrated with awk/shell script over all the seven locations. With this in place, things move very fast with additional scripting and customisation. For example, once a new customer is added, within an hour the script analyses the data. It takes the updated database and gives a report for the newly added customer. So the billing lifecycle is reduced from 24 hours to one hour.

So Netflow technology optimised on Linux efficiently provides the metering base for a key set of applications including network traffic accounting, usage-based network billing, network planning, Denial of Services (DOS) monitoring capabilities, as well as network monitoring for both service providers and enterprise customers.

THE ROAD AHEAD

As GTL has an immense talent pool, the company now covers a complete software portfolio in the networking, Internet and security areas. Its mission is to secure the profitability of its customers' IT investments by providing knowledge about new technologies to the market. As the company's CIO, Savio Furtado, states, "For customers to benefit from their investments, choosing the right solution is essential. We at GTL are constantly innovating. Now with open source we are sure we can deliver much more to our esteemed clients."

In a Nutshell

The organisation: GTL is a network engineering and IT services company which has effectively deployed customised solutions for itself and its numerous clients.

The need: To develop a robust Web solution that facilitates the administration of Web servers, mail servers, DNS servers and the like.

The solution: Using Linux as the software platform, GTL has achieved enough flexibility in its work and has been able to deploy powerful and effective solutions.

The benefits: GTL has saved a lot in costs, besides helping its entire team develop more powerful applications. 