DominoPower Magazine
The Power Magazine for Lotus Notes & Domino Professionals

INTERNET PLANNING

Why Internet mail gets out of control and how to avoid going postal

By Ron Herardian

This article is the first in a series of three articles on Internet email. In this first installment, I'll explain the central issues surrounding capacity planning for Internet email and give you practical advice, pointers and potential pitfalls. I've also included an example capacity model for network traffic projections.

Capacity planning has often been viewed as more of an art than a science. Examples of failure abound, ranging from severe inadequacies leading to catastrophic system failures (like the California highway system and the San Francisco Bay Bridge) to grossly over-engineered systems that waste millions dollars. Internet email is no exception. In fact, recent years have seen a radical increase in Internet mail and most companies have had difficulty keeping up.

Most companies start out thinking of Internet email as a simple gateway service for a few users to communicate with a limited number of business partners and clients. This way of thinking stems from early email gateway services like MCI Mail, Compuserve Mail, SprintMail and AT&T Mail that provided business-to-business email communications through proprietary infrastructures at a relatively high cost compared to Internet email. These services can now be thought of as legacy services.

Within a few months of installing or upgrading Internet email facilities, many companies discover that they haven't invested enough in their email infrastructure and, as their systems become less reliable and begin to fail, they feel the pain of making do with inadequate resources. The resources are of course, hardware, software, network bandwidth, and the human resource represented by the available time of network administrators.

Top 10 planning issues for Internet mail

In the last few years, since approximately 1995, most companies have seen a dramatic and continuing increase in Internet email traffic. Planners have often been caught off guard by increases of 500% or 1000% inside of 24 months. There are many reasons for this but here are the top 10:

1. Changing role of SMTP

The role of SMTP (Simple Mail Transport Protocol) tends to change from a simple gateway service to an enterprise email hub linking business units, clients, and business partners. To understand a company's growth pattern, it's necessary to understand its business, particularly in terms of potential mergers and acquisitions, or other potential forms of rapid growth. You need to be able to map business developments to technology requirements (at least as much as possible without venturing into a fantasyland). For most companies, this means that business and technology planners need to work more closely together. As the economies become increasingly Internet-centric, this will become the norm rather than the exception.

2. Mission critical e-business

Internet email is a key component in electronic commerce. This often means that not only will the volume of Internet email go up but the importance of Internet email increases. When a business stands to loose revenue because of interruptions of Internet email, it's time to rethink the SMTP infrastructure.

3. Rapid adoption of standards

The messaging industry and business at large has selected Internet standards for messaging making the SMTP, MIME, POP3, IMAP4, LDAP, S/MIME, X.509 and SSL standards increasingly important. This means that

companies are implementing systems supporting these technologies at a record pace.

4. Business-to-business Internet
Traditional, proprietary business-to-business email communications (MCI Mail, Compuserve Mail, SprintMail, AT&T Mail) are being dropped at a rapid pace in favor of Internet-based communications. It is worth noting that the overall volume of business-to-business email tends to increase and this trend is likely to continue. In part, this is because proprietary email systems are now connected to the Internet at large. In the late 80s, you chose your proprietary provider and mostly communicated with others on that service. Now, everyone can reach everyone else and the service itself has become far less visible in the overall email infrastructure.

5. Thirst for bandwidth
Often the network infrastructure for SMTP, in terms of server backbones, internal WANs, and connections to the Internet turn out to be unable to handle the load. Companies do not always realize the interdependency of these three aspects of networking and their impact on Internet email reliability.

6. Competition for Internet access
One common problem is that SMTP email services are often in competition with the Web for access from internal networks to the Internet. This means that corporate email to and from the Internet can be interrupted by an excess of casual browser users, outside visitors, or even PointCast subscribers. To avoid this, network planners must anticipate this issue. [Here at ZATZ, we had noticed a different problem. Whenever there's a cute animation or video going around via email forwarding, our pipe often gets clogged when someone decides to send a 24MB video to 200 of her closest friends. We've been able to discourage this practice through training and we see far fewer "hairballs" in the pipe. -- DG]

7. Unsolicited bulk email
Most companies are virtually helpless against unsolicited commercial email (colloquially referred to as "SPAM"). Often corporate backbones are flooded by unauthorized SMTP relaying and users are sometimes inundated with SPAM email messages. Spammers use a variety of techniques to probe networks and to discover user email addresses. Companies that fail to aggressively combat spammers could become victims at any point.

8. Larger and larger messes
When looking at email statistics, network planners often ignore the inevitable increase in message sizes. Also, statistics for internal email are often used to make projections of network traffic for Internet email when, in fact, Internet email messages often have many more and sometimes larger attachments than internal email messages.

9. Over-simplification
Capacity plans based on current statistics even allowing 100% increases in 12 to 24 months often turn out to be woefully inadequate. Email doesn't always behave like other systems when it comes to capacity planning because a variety of factors can cause sudden increases and because the peak-to-average ratio for network traffic can vary tremendously from one group of users to another.

10. Lack of planning
The absence of capacity planning exacerbates all of these issues and is almost always a central factor in catastrophic failures where systems are down, or highly unreliable, for extended periods of time. To make matters worse, when a system begins to crumble, it's not always clear what the problem is. Tactical fixes are often layered on over time instead of properly planning and undertaking a complete redesign of a company's SMTP infrastructure.

[Don't forget another aspect of email usage: hard disk storage. Because hard drives are so cheap, we tend to encourage our internal users to save all their email for potential future reference. However, over time, email can take up hundreds of megabytes of disk space, requiring upgrades of disk drives. The drives themselves cost little, but the time to perform the upgrade, the downtime of the upgrade system during the upgrade, and the idle time of the user can often have a considerable cost. -- DG]

Number one warning sign

Poor performance is the number one warning sign for Internet email. It's important to remember to test or to

attempt to duplicate reported problems such as delivery failures during peak hours.

When a gateway or MTA falls behind in processing messages during peak hours messages queue up faster than they are sent or delivered. When overloaded, the performance of most MTAs and gateways will degrade leading to a rapid reduction in efficiency. The gateway or MTA may catch up slowly during non-peak hours but when users begin complaining that recipients on the net aren't getting messages in a timely manner it's time to start looking for the bottlenecks.

Despite the many issues and risk factors for SMTP, Internet email bottlenecks are often found in the internal email system. In other words, users may complain about slow Internet email but, of course, they don't understand the email system. The actual problem could be in the internal email routing. This should be ruled out before any changes are made or hardware is purchased.

A "peak" at capacity planning

Network planners and technical staff often don't know how to predict email traffic on the network or how to justify needed resources to management. Capacity planning for Internet email involves two fundamental things: server sizing and network engineering. Here are a couple of pointers:

Typically there are two major utilization peaks for a messaging system associated with the business hours of a given group of users within a given time zone. For a company where most employees work from 9 to 5 the peaks tend to occur at approximately 10:30AM and 1:45PM and last as long as 1 hour. Keep in mind that your actual mileage may vary. Note that the apparent length of peaks could be longer for systems that have insufficient capacity. It's also important to predict overlapping peaks across time zones if Internet email services are provided centrally for multiple regions.

One maxim of capacity planning is that a server should be sized, or a network engineered, to accommodate the peaks rather then the averages. In other words, if messaging statistics tell us that there are 50,000 email messages per day to and from the Internet and, for example, we anticipated that peak traffic would be 200% of average then we have to engineer a system that could handle a peak of approximately 3.5 MPS (messages per second) given an 8 hour day. What that means for servers and networks depends on the email technologies involved, the number of recipients per message, and of course the sizes of the messages themselves.

Table A is an example model for network traffic projection.

Measure	Number	Units
Number of users (theoretical)	2,500	users
Total messages per day sent and received per user	20	Messages (est.)
Percentage of all messages to and from Internet	5.00%	of messages
Internet messages per day sent and received	2,500	messages
Number of recipient domains per outbound message	1.5	recipients
Additional outbound messages for multiple recipients	625	messages
Internet messages per day sent and received	3,125	messages
Average size of messages to and from Internet	67	KB
Total message bytes transferred daily	209	MB
TCP/IP, DNS, SMTP, protocol and transmission error overhead	10.00%	overhead
Total bytes transferred	230	MB
Length of business day (excluding overlapping time zone hours)	16	hours
Network traffic to and from the Internet	2.35	MB

Number of time zones	5	time zones
Number of overlapping time zone peaks	2	
Average network traffic	32	Kbps
Peak network traffic is n% of average	300.00%	of average
High average network traffic with non-overlapping time zone peaks	96	Kbps
Peak network traffic for overlapping time zone peaks	134	Kbps

What's next?

This article only scratches the surface for Internet email and capacity planning but it does set the stage for solid analysis and planning by pointing out key considerations and practical tips. In a future issue, we'll have a look at coexistence and migration of SMTP email services for Lotus cc:Mail and Domino.

Special thanks to D. Kahvedjian for inspiring this series of articles.

Glossary

Here's helpful glossary of terms:

- CSMA/CD: Carrier Sense Multiple Access/Collision Detection
- IMAP4: Internet Mail Access Protocol version 4
- LDAP: Lightweight Directory Access Protocol
- MIME: Multi-purpose Internet Mail Extensions
- MPS: Messages Per Second, a relative measure of email traffic used for server sizing
- POP3: Post Office Protocol version 3
- S/MIME: Secure MIME (see MIME)
- SMTP: Simple Mail; Transfer Protocol
- SSL: Secure Sockets Layer, a standard for the encryption of IP datagrams
- X.509: A standard for digital certificates and encryption used in S/MIME

Ron Herardian is CEO & Chief Systems Architect at Global System Services (GSS). You can reach Ron via email at rherardi@gssnet.com, or via his web page at http://www.gssnet.com.