PHIL PENNOCK

on IMAP service for customers

Phil is an expatriate Briton living in The Netherlands, desperately trying to live up to the job title of Senior Systems Administrator at an ISP. He uses IMAP for handling his email at home and at work.

pdp@nl.demon.net

THE ISSUE OF PROVIDING IMAP SERVICE to customers has been raised more and more frequently. This article attempts to explain the issues involved, why there's no current product at my ISP, and the need to be very careful in promising anything to customers (or potential customers).

I think it important for Sales/Product-Management to understand the issues so that they can be better informed than their competitors.

In Internet mail systems, there's one distinction that most people can happily ignore, even though it affects every message they send or receive. It's the concept of "final delivery." Email messages can be passed through many systems between the sender and recipient (though some systems will decide that 30 is too many and evidence of a mail-loop). But only one of those systems, the last, involves delivery into the user's real mailbox. That's the final delivery.

Up until final delivery, the message is in transit. If the message can't be passed on, then a delivery failure notification ("bounce") message is generated and passed back to the sender. Once it reaches the final mailbox, there will be no bounce. There's no guarantee that the mail will actually be read, but it won't be bounced because it's been delivered to "where it was supposed to go."

IMAP is designed as a system to manage mail once it has reached its final delivery point. SMTP is *before* final delivery. POP3 is a slightly weird hybrid, but not really well suited to use after final delivery. For many customers using POP3, the POP3 retrieval to their Outlook program *is* the final delivery.

Some ISPs do not provide final delivery. They pass mail on to customers. Historically, ISPs just provided SMTP; nowadays you can also find POP3. If a message is not collected within 35 (or whatever) days, the mail systems delete it. If it was read by POP3, it's automatically deleted; if it was not read by POP3, a bounce message is sent back.

Once mail reaches final delivery, it can stay in the system for as long as the user wants. There's no 35-day limit. If someone has an important message they treasure, they might want to keep it for the next 70 years or more—this is to be expected, not something exceptional.

Because ISPs just handle messages "in transit," there are rough limits on how much mail needs to be stored. ISPs engineer the systems for a different set of operating conditions than those used for providing final delivery.

For instance, systems are extremely unlikely to lose email. That doesn't mean that they *can't* lose email. An asteroid can strike the planet, destroying Western

Europe; enough disks can fail in the NetApp, all at once, to lose messages. These are both possible (and the latter is, hopefully, *more* likely). If a customer collects mail regularly, between a few hours' and a couple of days' worth of their email might be lost; if the customer doesn't collect regularly, up to 35 days' worth might be lost. The design makes this unlikely, but there's no possible design that makes losing email impossible. In the worst case, business insurance covers this.

The monetary worth of 35 days of email as opposed to 35 or 70 years of email is an interesting thing to consider . . .

Because of the issues of volume, many ISPs can't offer an IMAP product that scales to all customers. Individual small products can be offered, though, on the scale of what a company of a couple of hundred employees might implement internally—that's doable.

IMAP needs to be considered not as "another way for customers to get their email" but, instead, as "how customers manage email they've received." They each get multiple mailboxes and can move mail between them, delete mail, flag mail, add attributes and keywords, search on message content, and much more besides.

The mail-store of the IMAP system will hold immense amounts of information important to an organization. Only the organization itself knows *how* important. But any strategy needs to handle issues such as backups, replication, archives, and policies on personal mailboxes of departing staff.

But how much information can be held? How much should be held? "All email" can get to be a very large amount of data over time. Much of the information will lose relevance; some will not. From an operations point of view, it's certainly useful to me that I can look over the past three or fours years' worth of list mail, quickly searching for messages matching a few keywords on a subject that I vaguely remember coming up before; with paper memos this would not be a productive use of my time, and I'd be better off figuring out the solution from scratch (or having decent documentation).

Any "IMAP solution" offered needs to consider not just "mail for a few months." It needs to be "email useful to the customer, for the lifetime of its usefulness." It needs to provide for backups, in various forms. It needs true disaster recovery. It needs a lot that is expensive to provide—probably the reason that customers come to ISPs, after their sysadmin/consultant has told them how much it costs to do things properly.

Any IMAP offering that's cheap to implement carries a potential legal and financial nightmare—I really don't think that most ISPs offering IMAP have properly evaluated this, just as many who rushed to offer free accounts didn't evaluate their long-term business plans, either. The ISP market is still young enough that it's filled with cowboys who have much financial backing but simply don't understand what they're doing or what the long-term consequences of their products are, focusing on "get more customers now" instead of "get customers whom we'll keep and who won't be suing us into bankruptcy in three or eight years."

Email has much to offer (including searchability, mentioned above), but it also brings challenges that are typically not addressed. IMAP is an important part of the picture, but not the whole picture. I believe organizations need to understand how email messages and other forms of documentation are used, and establish policies for email retention and migration of information into more formal documentation.

For instance, a policy might resemble the following:

Email messages more than seven years old are archived onto a read-only long-life medium and deleted from the live system. After only six years, someone reviews mail to mailing lists X, Y, and Z for information that looks as though it might still be relevant and collects those messages for review by specialists.

The specialists collate those that still hold relevant important information and ensure that all the information is held in Procedures, Policies, and Guidelines or other formal documentation.

Clearly, this ties deeply into internal information management. Not losing information involves some bureaucracy (the NOC will now boo and hiss). Choosing how to handle this is not easy and may well be different for each customer. Their business processes for information management need to be designed to migrate information of any value from ephemeral communications such as email into more static forms of retaining information such as traditional documentation.

An ISP could offer some standardized services along these lines, with tools and calendaring designed to make it easy for companies to collect the information they need; with automatic burning to DVD (this year—who knows which medium in six years?) of a customer's mail every time a certain volume is accumulated or amount of time passes; with those DVDs being mailed to the customer by recorded delivery or courier. There's a lot that can be done.

But any time that you look at getting this involved in a company's internal processes, what you're actually selling is IT outsourcing services, not Internet access.

I believe that offering "proper" IMAP access is something which intrudes deeply into the market of some much larger companies; it's not something to be taken lightly, and it's certainly outside their usual areas of expertise.

Certainly, if one wished to start moving into ASP (ye olde Applications Service Provider, which we heard so much about a couple of years ago) or IT outsourcing, then that's a different matter, and IMAP is just one of the technologies which would be used to provide services. But this is not Internet access. It's not shifting information around or allowing customers to shift information to others. It's controlling how customers shift information around within their own organization, which is a fundamentally different animal.

"Just providing IMAP access to a mail drop" is a short-sighted viewpoint; unfortunately, many customers won't understand the issues here and will want "just that." Perhaps this article can be massaged into a document issued by ISPs to customers to help them think more deeply about the issues and to realize that the providers want to work with them to offer good solutions, not just take their money for whatever they can without regard to the consequences.

Some ISPs can offer a small IMAP product in the near future, but it's essential that this isn't hyped and that we are proactive in ensuring that customers have information on the limitations described here; and it's equally important that they consider this a stop-gap solution while they develop something that integrates better with their business processes.