

## READER-EXPERT I/O

### You've got mail!

**Our e-mail storage needs are growing like wildfire. We keep adding users, expanding capacity, and now are about to add MS Exchange servers to break up the load. Backups take way too long to run, and we are nervous about the time it would take to restore. We cannot be alone in this problem! Can you make it go away?**

Indeed, you are not alone. E-mail is a tough beast to contain. Heck, my personal e-mail stash is nearly 10GB!

While you can try to contain the problem by setting storage quotas (Try telling the CEO that he/she has to delete old e-mails if he/she wants to continue to receive



**Jacob Farmer** is the CTO of Cambridge Computer Services. He can be reached at [jacobf@cambridgecomputer.com](mailto:jacobf@cambridgecomputer.com).

new ones. For that matter, try telling sales or customer service!), you could also try the “hit-it-with-a-hammer” approach. Invest in a storage area network (SAN) infrastructure or a big network-attached storage (NAS) server. Each gives you the tools to add storage on-demand and to tackle backup.

The NAS approach gives you storage pooling so that you can accommodate multiple Exchange servers. If you use snapshot technologies, you could be running LAN-free backups while your system is in production.

But be careful: NAS for Microsoft Exchange is still a relatively new phenomenon, and the jury is still out as to how elegant and scalable it actually is. So, at minimum, you'll want to make certain that

the NAS box has been specifically certified for MS Exchange. I know of only one product line that is.

As for SANs, they'll give you centralized storage pooling and limitless capacity. Backups can run over the SAN to eliminate the bottlenecks associated with LAN backup. Most backup software products allow you to back up Exchange while it is in production. If not, you could use a conventional open-file manager. I am a big SAN storage fan, so I would lean in this direction. The only problem is cost, but SAN prices are moving aggressively downward.

But the problem with throwing either SAN or NAS at your e-mail system is that both cater to the false notion that it is okay to allow content to pile up in e-mail systems. Remember that e-mail systems are first and foremost messaging systems, which means they are designed to facilitate message transfer, not long-term storage.

Some people store messages forever. Some never purge deleted messages. Presumably they do so because they think they will need to refer to these messages one day. For these people, old messages should be classified as “content,” and you should be looking at content management tools to solve the problem.

The solution is to install an e-mail archive, or at the very least, an attachment manage-

ment system. These products interface with e-mail servers and move content from e-mail servers to dedicated archives. Some systems move just the attachments, leaving behind pointers. Retrieval is transparent to the users. Other systems move the entire message based on a set of rules. Both types of systems weed out duplicate messages and duplicate attachments.

The benefits of e-mail archiving are many. From an infrastructure perspective, it can greatly reduce the load on e-mail servers (in terms of bandwidth, CPU, and storage). And it can also greatly improve your backup-and-restore capability. Not only are backup loads lighter (most of the content has been removed to the archive), but the actual backup process is easier because content is static and therefore can be backed up using conventional methodologies.

E-mail archiving also eliminates the need for “brick-level” backup products, which handle individual mailboxes. Those types of products tend to be very slow and cumbersome and are managed by the IT department. With an e-mail archive, users help themselves to old messages.

E-mail archiving also has business benefits that can improve ROI. The most obvious beneficiary is the user who regularly refers to old messages. The search-and-retrieval tools of a dedicated archive are far superior to those found in conventional e-mail client software. Search criteria is more flexible (you can search for content in both messages and attachments), and traffic in and out of the organization can be mined and monitored by group managers, select co-workers, compliance officers, etc.

And, lastly, if a company is asked to pro-

## READER-EXPERT I/O continued

vide copies of past e-mails for legal reasons, e-mail archives can save a fortune.

In short, e-mail systems are not content management systems, and if they are forced to play this role, management nightmares will likely abound. The better

solution is to implement either a powerful storage management infrastructure (e.g., SAN or NAS) or install an e-mail archive. Of course, the best solution is to do both. But if your budget is tight, the archive is probably the better value. □

---

If you have a question you would like to ask one of our experts, please e-mail Heidi Biggar at [heidib@pennwell.com](mailto:heidib@pennwell.com).