



## When e-mail stops, teamwork counts

*April 5, 1998*

By 7:15 Sunday night, September 13, IT Services E-mail Team members were sure the Exchange server repairs were complete. The server was back online and processing e-mail normally. The only thing left to do was notify the customers. But it had been a very long week.

### Hitches in the routine

The trouble started around 4:30 p.m. Tuesday, September 8, when an Exchange server administrator on the E-mail Team noticed unusually large mail queues on Exchange server DLEE02. A mail queue is a group of messages waiting to be delivered. Generally, a large mail queue is not serious, but the admin decided to take some routine steps to make sure the queues would clear out. When one of those routine steps refused to complete, the admin took another fairly routine step: he rebooted the server. By 6:00 p.m., the reboot was finished, and the server was running normally, so the admin headed for home.

Unfortunately, he didn't stay there long. At 7:45 p.m., he got paged. DLEE02 had stopped routing mail. Customers began receiving Non-Deliverable Reports (NDRs) on e-mail they sent, and the Exchange Information Store showed unusually high CPU usage. The admin again started the basic steps to reset the services and at 8:30 p.m. had to reboot the server again. By now, he was concerned, so when the system came back up at 9:00 p.m., he stayed in his office and monitored it carefully until 10:15.

When the beeper under his pillow went off at 12:10 a.m., Wednesday, September 9, the admin began to suspect the problem was not the minor and routine glitch it had at first seemed. Back at the office, he found history repeating itself: DLEE02 had stopped routing mail, customers were receiving NDRs, and CPU usage was pegged.

All his experience said that the service restarts and/or reboot he had performed earlier ought to have fixed these symptoms. Since the symptoms were back, it was time to look for other solutions. The E-mail Team's primary focus is to run and maintain the e-mail service. Top priority was to get the server back to processing e-mail. Twelve fifteen in the morning is no time to be thinking through a complex problem on your own. With other Exchange Team members out of town at a Microsoft conference, the admin decided to contact Microsoft and consult with a technical support representative.

After some discussion, the support rep and admin diagnosed the problem: corruption in the database that contained customer message stores. They had two options: restore the server database from a back-up or rebuild the existing database to repair the corruption. The tape from the scheduled back-up at 9:00 the night before was available, but the two technicians were concerned about the quality of the back-up. Since the server was showing symptoms of a problem before the back-up ran, was it possible that the tape contained corrupted data?

Rather than spend time restoring a possibly corrupt database, they elected to rebuild the existing one. At 1:20 a.m. the admin started the rebuild utility, believing it would take five hours to complete it and run another utility to correct any corruption left over from the rebuild. Thus, the server should be back up by 6:30 a.m., before the majority of DLEE02 customers arrived at work. Unfortunately, this was not the case.

### Best-laid plans

On Wednesday morning, the rebuild utility finally finished at 11:40 a.m. A DOS error on the server forced a reboot, and the admin started the second utility at 12:15 p.m. The second utility returned an error, but by this time the server had been down for far too long. Customers were unable to exchange e-mail, so the admin decided to try to start the services again in the hope of getting mail flowing. Even this turned out to be impossible because the left over database corruption prevented the services from starting.

By 1:15, the frustrated admin contacted Microsoft again, this time getting a representative who recommended installing a hot fix (a software fix) before trying to run the rebuild utility again. Downloading the hot fix from Microsoft's Web site took some time, but by 2:15 p.m., it was installed. In the meantime, the

Microsoft representative realized that restoring from a tape back-up was still a possibility. If the tape back-up finished, the database might contain corruption but would still be usable.

Abandoning plans to try the rebuild, the admin began preparing for the restore. At 4:00 p.m., he started a file-level back-up to capture changes since the last full back-up, and at 5:30 p.m., he started to restore from the most recent back-up tape. The restore finished at 9:00 p.m., and other E-mail team members applied the earlier file-level back-up to restore the further changes.

By this time, Microsoft had a representative on-site. Around 11:00 p.m., he decided to run a utility that would show what kinds of errors were corrupting the database. That utility finished running at 3:30 Thursday morning. It showed a database that needed defragmenting, but one that was workable for the time being.

Defragmenting is the process of rearranging chunks of data in the database so that all the data is compacted into the smallest amount of storage space. When applications store information in a database, they use the next available block of space that is big enough for the piece of information in question. Over time, this can result in small amounts of free space being "trapped" between larger chunks of information. Defragmenting arranges the large chunks of information so that no free space is trapped between them. This leaves the previously trapped space available for more data storage.

At 5:00 a.m., with the fragmented but workable database restored, team members began preparing to bring DLEE02 back online. It was a struggle and required a reboot, but by 8:30 a.m., the server was back up and running. The original admin, back after a few hours of sleep, started a back-up on the server around 11:00 a.m. The back-up froze after an hour and twenty minutes, but the server continued to run normally, so the team elected to leave it running to keep e-mail in service. Meanwhile, they performed an audit on the server's operating system to make sure nothing in the operating system could be causing problems. They found nothing.

At 5:00 p.m., the admin stopped the frozen back-up procedure in the hope that the daily back-up, scheduled for 9:00 p.m. would run normally. At 9:00, that daily back-up started and froze after an hour and twenty minutes.

By 7:30 Friday morning, one of the Exchange team members at the Boston Microsoft Conference called with a temporary fix to allow the back-up to run. With the fix applied, the admin started a complete back-up at 9:00 a.m. Since the server can run during a back-up, the customers were able to continue using e-mail.

### **Planning the work, working the plan**

With the server at least functioning, the team had a little breathing room. Friday afternoon at 2:00, team members met to develop a plan to repair the corrupted database on DLEE02. Working from information provided by Microsoft, they put together a schedule for a weekend of marathon repair. They would start with a complete back-up, then install several software patches, and finally run utilities to defragment, repair, and verify the database. At each major phase, they would run another back-up to capture the state of the database.

If team members were on site to start each phase of the repairs as soon as the previous one finished, they estimated they would have everything completed sometime Sunday morning. With the goal of having the server back online and fully repaired by 6:00 a.m., Monday, they had a wide margin for error. They sent a message to customers announcing the down time and synchronized watches.

### **Let the games begin**

At 6:30 p.m., team members began the first stage of the repair marathon: a complete back-up. This was their fail-safe. If something went wrong and they couldn't complete the repairs, they could restore this back-up for Monday morning. The database would still need repair, but at least the service would be back online.

The back-up completed at 8:10 p.m., and the team began installing software patches. They had some trouble with one installation and elected to do it later and proceed to the defragment utility. They started it at 9:30 p.m. By 11:30 p.m. team members arriving back early from the Boston conference were at the Lewisville Worldwide Command Center to monitor progress. With the defrag utility running, there was nothing to do but wait.

Saturday morning at 8:11, the team arrived back at Lewisville. The defrag utility was nearly finished. By 9:39 a.m., they were ready to start the next utility, the one that would repair any corruption left after the defrag, but the utility wouldn't start. A call to Microsoft yielded information about a registry setting that needed to be changed to let the utility run. The Microsoft rep also recommended two more procedures not included in the original list.

By 1:15 p.m., the two extra procedures were complete, and the team was back on track. They now started the utility to repair left over corruption. They would need to run this utility multiple times, until it reported that it performed no repairs during its run. By 8:14 p.m., the first run was finished in 6 hours, 45 minutes, reporting nine fixes. The team started the utility again.

At 3:30, Sunday morning, a team member arrived at Lewisville to check on progress. She found that the repair utility was complete and reported no fixes. The team had expected to run the utility at least three times, so this was good news. The next step was another complete back-up, which completed at 6:00 a.m. Now it was time for the verification utility, which ran with no errors. DLEE02 was now ready for testing.

Problems cropping up on other Dallas area Exchange servers delayed testing on DLEE02, but the team sorted them out and turned attention back to DLEE02. At 10:30 a.m., they fired up the server. Everything worked--except Outlook Web Access (OWA), the simple Web-based e-mail reader that gives basic access to some Outlook/Exchange functions.

The team went to work on OWA. By noon, they had determined that part of the problem was with the way some file permissions were set on the server. But even after they corrected the file permissions, OWA still didn't function correctly. Because the Dallas area Outlook/Exchange servers work in concert, OWA was available from the other two servers. DLEE02 customers could use it from one of those locations to view their e-mail. At 7:00 p.m., the team elected to re-boot DLEE02, notify the customers of a small, temporary change in the way they gained access to OWA, and call it a weekend. The server was back online by 7:14 p.m., and the customer message went out, after some refinements, at 10:47 p.m.

Further OWA repairs and detailed analysis of the entire situation could wait until later in the week. For now, customers were back in business, and the team beat the estimated repair time by ten and a half hours. It was a small victory at the end of a long, frustrating week.

During the following week, the team, working with Microsoft analysts, began investigating ways to prevent the problems from reoccurring. While no one ever wants problems like this to crop up, the dedication and skill of a team working in concert is what makes it possible to solve them when they arise. That kind of teamwork is fundamental to keeping e-mail flowing at TI.

### **The latest**

Since the September outage, representatives from IT Services, the Outlook/Exchange customer base, and Microsoft have formed a task force to diagnose and correct software and support problems that caused that outage and other intermittent Outlook/Exchange service outages in recent weeks. Plans already in motion include moving the Exchange service to new hardware with freshly installed software. Outlook/Exchange customers will receive more detailed information about the move process in the coming days.

Other long-range plans include creating an expanded test lab for the Exchange environment, implementing regular down times for scheduled maintenance on the Exchange servers, and practicing data recovery methods. Additionally, refinements in support procedures should help to reduce the margin for human error.

"I want to assure our customers that we understand the business critical need for e-mail at TI. We operate all the e-mail systems with that in mind and 7x24 service as our goal," said Fred Berlack, product manager for the IT Services E-mail team.

The team remains committed to running and maintaining TI's e-mail service with maximum uptime.