



Reseller Incident Report # 2993

Please note all time and date information in this document refers to Eastern Daylight Time (EDT) unless otherwise specified. EDT is GMT-4 hours.

Date of Issue:	October 6 - 11, 2008
Severity Level:	Emergency
Associated Trouble Ticket:	2993
Service Affected:	Cluster A: IMAP, POP and Webmail, inbound mail
Customer Impact:	IMAP, POP and Webmail intermittently offline for approximately 40% of mailboxes (20% of all users). Inbound mail was queued. Outbound mail was functioning.
Outage start time:	October 6, 2008: 13:07
Resolution time:	October 11, 2008: 12:24
Outage duration:	4 days, 23 hours, 17 minutes
Incident summary:	<p>An upgrade to the Linux kernel on our IMAP/POP servers had an obscure interaction with NetApp filer that caused mail server file locks to not be released.</p> <p>Upon reaching the timeout to lock the mailbox index, Dovecot opened the index file for writing regardless of the lock. Collision of writes to the index caused index file corruption for some mailboxes. Any subsequent user access to an account with a corrupted index triggered Dovecot to attempt to reindex the entire mailbox.</p> <p>As more and more indexes got corrupted, we reached a critical mass of corrupted indexes. The reindexing on a large scale caused a high number of reads on the NetApp that resulted in the disks becoming too busy and overwhelmed the system. Normally, the self-healing nature of Dovecot would fix the corrupted indexes by triggering a reindex of the mailbox. However, this expected behavior did not occur effectively due to the Dovecot bug. The result was a growing number of corrupted mailboxes, which were repetitively trying to reindex unsuccessfully. This increased the load on the storage infrastructure by a factor of ten.</p> <p>Over the course of the incident, some customers on Cluster A experienced intermittent to offline mailbox access via IMAP, POP and Webmail. At the beginning of the issue, inbound mail was periodically queued and released for all</p>



	<p>customers. Once the issue was identified on three mailstores, we queued inbound mail for these users for the duration of the incident. Slowness experienced by users interacting with the system caused an increased number of sessions to be occupied for a greater duration. This resulted in some customers receiving intermittent timeouts or “Service is unavailable” error messages.</p> <p>In order to troubleshoot and reduce the load on the systems, we began queuing inbound mail on a different filer. Tier 3 NetApp engineers collaborated with our team to conduct extensive tests on the affected filers. Together we determined that hardware was not the root cause and that the filer was acting appropriately.</p> <p>In order to address the symptoms experienced by users, and while we continued to investigate the root cause, additional NetApp shelves were added to the storage cluster. The team triggered a script to move a fraction of mailboxes to the new storage shelves in order to alleviate the stress on the existing shelves. Once triggered, analysis of initial mailbox migration determined that the length of time to complete the move would be longer than originally expected. We shut down three mailstores on the NetApp filer to alleviate I/O from user access in order to maximize throughput for the move script.</p> <p>These actions did alleviate some of the stress on the system, but because the load was immense (an order of magnitude above normal), the symptoms experienced by users were very much the same.</p> <p>At this point parallel investigation done from the software perspective pointed to the problem arising at approximately the time of a kernel upgrade. Having eliminated hardware as the cause of the problem, our engineers engaged NetApp’s kernel developer to assist in the root cause analysis. The software issue with the Linux kernel was identified and a fix was put in place. This successfully addressed the original locking issue. The team commenced reindexing the affected mailboxes. Three mailstores were taken offline for this process and brought</p>
--	--



	<p>online once reindexing completed. At this stage it was apparent that we were dealing with a secondary issue as symptoms of the original problem were still detected even though the locking problems were resolved. The file locking problem exposed a Dovecot bug, which resulted in indexes never being properly written to disk.</p> <p>We then engaged the primary Dovecot developer to isolate a potential issue with Dovecot. A Dovecot patch was created, tested and implemented. Then, we successfully reindexed the affected mailboxes.</p> <p>All mailstores were brought online and users were able to authenticate via IMAP/POP/Webmail. We proceeded to deliver the previously queued inbound mail.</p>
--	---