



Reseller Incident Report # 1991

Please note all time and date information in this document refers to Eastern Daylight Time (EDT) unless otherwise specified. EDT is GMT-5 hours.

Date of Issue:	August 12 - 15, 2008
Severity Level:	Emergency
Associated Trouble Ticket:	1991
Service Affected:	Cluster A: Webmail, POP, IMAP unavailable
Customer Impact:	50% of customers on Cluster A had no access to their mailboxes
Outage start time:	August 12: 14:44
Issue confirmation time:	August 12: 15:21
Resolution time:	August 15: 16:00
Outage duration:	<p>Overall incident duration: August 12 15:21 – August 15 16:00, 73 hours, 16 minutes.</p> <p>Outage Duration Breakdown: Initial 34 hour period affected 50% of the user base, resulting in no Webmail, POP, or IMAP. Next 20 hours and 37 minute period affected 33% of the user base, given limited Webmail. Final 18 hours, 39 minute period affected 12.5% of the user base affected, still limited Webmail.</p>
Incident details:	<p>Incident summary: A shelf controller hardware failure caused 14 disks to require a rebuild. This resulted in the degradation of multiple storage volumes. This failure affected 50% of customer mailboxes on OpenSRS Email Service - Cluster A. The restoration process was consecutive for the affected devices and therefore took a number of days.</p> <p>Over the course of the 4 day degradation, 30 System Status email notifications and simultaneous System Status website updates were provided to resellers in regular intervals.</p> <p>Incident Details:</p> <p>August 12</p> <p>14:44 The storage controller logs a shelf failure.</p>



	<p>15:01 The Network Operations Center (NOC) monitoring begins to receive alerts.</p> <p>15:05 Technical Support escalates customer examples of “Service Unavailable” error messages obtained for mailboxes on Cluster A.</p> <p>15:21 NOC advises that some customers on Cluster A can not access their mailboxes via POP, IMAP or Webmail. Inbound mail is being queued locally. A System Status notification is sent to resellers.</p> <p>15:27 Technical Operations (Ops) notices a high number of connections on the load balancers and continues to investigate.</p> <p>15:45 A failover to a backup load balancer is done, but does not correct the issue. Ops restarts IMAP instances.</p> <p>16:06 Ops identifies a NetApp (storage device) shelf failure as the probable cause of the issue. The rebuild process automatically started (at 14:44) onto available spare disks. The vendor is engaged.</p> <p>18:00 Senior staff and executives hold briefing meetings to assess the scope and impact of the outage. These meetings continue throughout the incident. To improve access for the unaffected volumes, Ops disables access to all the mailboxes on the failed volumes.</p> <p>18:54 A new shelf controller is delivered by the vendor to our data centre and is physically installed by Ops staff.</p> <p>19:23 Ops confirms that 50% of customers on Cluster A are affected by hardware failure. These customers have no access to POP, IMAP, or Webmail. The remainder of the customers have full mailbox access. After testing, Ops fails back to the new shelf controller. This allows the device to separate and address the affected 3 volumes.</p> <p>19:30 The 50% of Cluster A customers with mailbox access are able to receive previously queued mail. This process took some time to complete.</p> <p>22:30 Ops and NOC continue to closely monitor the restoration process. Concurrently, teams work to strategize alternative solutions to provide customers with safe mailbox access as soon as possible.</p> <p>August 13</p> <p>11:15 Technical Support create a temporary email address for customers to send tickets. As well, Support reports high</p>
--	--



	<p>call volumes, which the team handles in a timely manner.</p> <p>12:15 Ops implements a change to enable access for Forward-only and Filter-only accounts.</p> <p>14:24 Inbound mail is fully confirmed and tested as being delivered for Forward-only and filter-only accounts. The Technical Support ticketing capability is fully restored.</p> <p>18:03 Our Development and Ops teams test a potential interim solution to provide limited Webmail access.</p> <p>19:33 The customer base is notified that the estimated restore time is Saturday, August 16, 2008 20:00 UTC. Ops and Development (DEV) begins implementing limited, Webmail-only access to the affected users on Cluster A. These customers receive an Email Outage communication in their Webmail box and are able to send and receive new messages. They are unable to see old mail stored on our servers.</p> <p>22:11 The Webmail script to give the remaining 50% of customers on Cluster A limited access progresses.</p> <p>22:45 Operations confirms script finished running and all caches flushed. The affected 50% of customer mailboxes have limited access to Webmail.</p> <p>23:13 The first of the affected volumes complete rebuilding. Ops begin testing.</p> <p>August 14</p> <p>00:34 Ops completes testing and confirms that one of the affected volumes returns online. One-third of these customers can gain access to their mailboxes.</p> <p>20:57 The second affected volume is restored. Ops begins testing.</p> <p>21:21 An additional one-third of previously affected customers are brought back online in sequence. Ops starts a mail merge script prepared by DEV to combine the interim limited Webmail access with the customer's restored mailbox. These users have full access via POP, IMAP, and Webmail as well as send/receive mail.</p> <p>August 15</p> <p>02:56 NOC confirms that the mail merge script for the second affected volume is complete. An addition to those customers who have access, another one-third of customers now have full access to their mailboxes.</p>
--	---



	<p>05:13 The rebuilding of the third and final affected volume proceeds ahead of schedule. Based on the timing of volume restores and success of the mail merge scripts, Ops determines the full restore timeline estimate would be soon adjusted.</p> <p>08:56 Ops adjusts the restoration completion time to Friday, August 15, 4:00PM ET (20:00 UTC).</p> <p>09:15 The final volume is restored. Ops begins testing, then commences the mail merge script to restore these customers.</p> <p>14:50 Mail merge script completes. Services are fully restored. Ops conducts extensive testing on the restored volumes to ensure all customers and their services are no longer affected.</p> <p>16:00 Operations confirms service is online. All customer mailbox access via IMAP, POP or Webmail is available. No mail or data was lost. Inbound mail is flowing as per normal.</p>
<p>Workaround/resolution:</p>	<p>The shelf controller was replaced and 14 disks were rebuilt. Temporary mail stores were made available to provide customers mail services. Once the restoration was complete, all mail content (mail and folders) was merged from temporary volumes to the customer’s original mailbox.</p> <p>Before the end of August, we are making storage architecture changes to Cluster A to ensure that we eliminate the chance that a similar event with storage will occur in the future.</p>