



It Is Internet (Pty) Ltd T/A isoho.st
Registration: 2008/021004/07
Director: E. S. Peer
55 Gordon Court
Pierre van Ryneveld
Centurion, 0157

Incident report: 2015-05-22

Description

After a successful switch-over of isoho.st's primary Internet break-out from Hetzner to Mweb on Thursday evening (21 May), network traffic to customer IP addresses for virtual machines not yet migrated to the new platform at Mweb was directed over a tunnel configured between the two clusters. Due to a Linux kernel bug [1] affecting the tunnel end-point at Hetzner, the path MTU [2] associated with routes via the tunnel was decreasing from the normal 1476 bytes to about 500 bytes at approximately 1.5 hour intervals for 20 minutes at a time. Some customer virtual machines were not configured to do path MTU discovery, which caused packets larger than approximately 500 bytes to be dropped on the tunnel, resulting in the failure of TCP connections to the affected VMs.

Analysis

The kernel bug causing the collapsing path MTU was most likely triggered immediately after the cut-over was completed around 23:00 on the 21st, however, the ping tests we were using to verify connectivity did not detect the fault, due to the default size of ping packets being smaller than the reduced path MTU, and the fact that TCP continued working on isoho.st test VMs that were employing PMTUD to correctly "follow" the effective MTU of the tunnel. Since the same tunnel was working for many weeks prior, the reason for the bug only being triggered after the cut-over is unknown.

We became aware of a network problem at 9:40 on Friday 22 May when we noticed TCP connections to one of the isoho.st internal VMs failing (we later discovered this VM was not configured to perform PMTUD). The first customer complaint was received at 10:45. Troubleshooting was hindered by the complexity of the redundant tunnel infrastructure, BGP and policy routing that was required to enable the migration from Hetzner to Mweb. The intermittent nature of the fault also meant that we only had the opportunity to track down the cause ~22% of the time. The cause of the fault was identified at 22:49. We implemented a work-around preventing the bug triggering within the next 30 minutes.

Impact

Customer VMs not performing PMTUD experienced a total loss of TCP connectivity for about 22% of the time, for the majority of the day. PMTUD must be enabled and requires ICMP *fragmentation needed* messages to be accepted by firewall rules. It is not known exactly how many client VMs were affected, but some clients were using a common OS image without functioning PMTUD.

Actions

With the migration to the isoho.st platform in Johannesburg complete as of Sunday 24 May, the tunnel to Hetzner is no longer required for production network traffic. The new cluster does not, and will never rely on the configuration that triggers this particular kernel bug, eliminating the possibility of a similar incident re-occurring.

References

- [1] <http://www.spinics.net/lists/netdev/msg178790.html>, <http://www.spinics.net/lists/netdev/msg180287.html>
- [2] http://en.wikipedia.org/wiki/Path_MTU_Discovery

Please direct any questions or comments to support@isoho.st