

Service Incident Report: CERN-ASGC link issue

London, 07 July 2010
edoardo.martelli@cern.ch

Incident



Description

Atlas users experienced unreasonable low data transfer rate from CERN to ASGC. CERN network team also experienced small size packet lost with ping test. But, ASGC network team identified no quality issue between Taipei and Amsterdam. Small packet size ping test and history diagram also shown no packet lost occurred from Taipei to Geneva. However, there is significant packet lost with larger packet size ping test from Amsterdam to Geneva.

Impact

Thousands of Atlas data processing jobs were queued waiting for efficient network to deliver to ASGC.

Analysis

After segment by segment network tests were implemented, a low power issue was identified on GEANT Frankfurt pop.

Low power issue is a very tough network problem to identify. In this case, users suffered from unreasonable slow data transferring, but ASGC network team experienced no packet loss on small packet size ping. The ping test and its history result have ASGC network team delay 48 hours to check the network issue. A simple and straight forward file transfer test scheme on LHCOPN network for engineer to run test might be helpful to verify such kind of problem in the future.

The 13/5 was public holidays in Geneva and Amsterdam, thus the delay of the Surfnet NOC to open the ticket.

Timeline



All time in UTC

2010/05/12 07:44: jwhuang@twgrid.org found single file transfer rate from CERN to ASGC was less than 50kbps. He warned that there might be a network issue.

2010/05/12 07:55: after check to ASGC smoke ping server, there is no significant packet lost occurred on the smoke ping history. Network issue was ignored by chenws@twgrid.org.

2010/05/12 15:09: neteng@cern.ch is notified of an issue affecting all the transfer from CERN to ASGC

2010/05/12 16:25: after checks on the CERN routers, neteng@cern.ch open a ticket with Surfnet, the circuit provider

2010/05/12 17:14: neteng@cern.ch moves the ASGC traffic to the CERN-ASGC backup link. The transfers improves, but the bandwidth is limited at 1.8Gbps

2010/05/13 10:50: jwhuang@twgrid.org confirmed again. There is no problem on the server pools. There must be a network problem.

2010/05/13 12:04: chenws@twgrid.org found the returning path from CERN to ASGC were routed to Chicago backup link. He also identified the 2gbps peering link is fully occupied. chenws@twgrid.org informed that the issue was caused by congestion. chenws@twgrid.org changed the route from ASGC to CERN back to 10g primary link, but the returning path is remaining on the backup link.

2010/05/13 17:45: chenws@twgrid.org was informed that data communication between ASGC and CERN was down, and the routing path change was done in purpose by CERN to by pass the issue. chenws@twgrid.org was asked to change the route back to backup link. After the route change, jwhuang@twgrid.org reported that data transfer is recovered.

Timeline



2010/05/13 17:45: After some data transferring test with jwhuang@twgrid.org, Edoardo confirmed that stickinking on backup path was the best solution although it was 50KB per second for that moment. He also opened a ticket with Surfnet which provide the link to the ASGC router in Amsterdam.

2010/05/13 18:02: chenws@twgrid.org confirm that there was no quality issue between Taipei and Amsterdam. He also found there was a significant packet lost occured with 8973Byte ping test on router interface to router interface. chenws@twgrid.org wondered there might be an MTU issue between Amsterdam and CERN.

2010/05/13 18:17: Edoardo replied that the MTU setting is correct. His test shown normal ping test with small size packet between router interfaces is having packet lost issue. He believed it was a link issue between Amstedam and CERN and suggested waiting for SURFNet response.

2010/05/14 07:30: For 12 hours no updated information, chenws@twgrid.org sent an email to 'helpdesk@nic.surnet.nl' and noc@netherlight.net to ask them speed up the process.

2010/05/14 08:18: Surfnet open the ticket and starts investigations. The ticket name is Netherlight-20100514-1

2010/05/14 09:45: Netherlight Peter Tavenier [Peter.Tavenier@sara.nl] replied no error found on Netherlight switches. He requested to check Taipei-Amsterdam IPLC link status.

2010/05/14 10:33: Netherlight Peter Tavenier proposed to attached an iperf server between Amsterdam and CERN to run performance test. 10:46 chenws@twgrid.org confirmed would do iperf test.

2010/05/14 11:13: chenws@twgrid.org provided IP information to Peter Tavenier. 11:33 Peter replied the circuit has detached, and a iperf server was set toward to Amsterdam router.

Timeline



2010/05/14 12:05: Netherlight Peter and chenws@twgrid.org both confirm ping test with packet size bigger than 1504 all failed. We were running into a maximum MTU of 1504 somewhere in Amsterdam

2010/05/14 14:05: chenws@twgrid.org identified that there was an MTU mismatch on ForthTen switch behind Amsterdam router. After the MTU mismatch fixed, the iperf test and big size packet ping test runs well between Amsterdam router and the iperf server of Netherlight.

2010/05/14 14:38: Peter Tavenier attached the iperf server toward to Geneva to run iperf test. He found that it did be a packet lost issue between Amsterdam and Geneva. His conclusion is: a lots of packet loss, the bigger the packet size the worse is gets. He decided to call GEANT service

2010/05/14 15:36: According to Netherlight ticket, GEANT found a low power issue in Frankfurt, engineer is on the way, no ETR yet.

2010/05/14 22:54: Geant reports to have fixed the problem

2010/05/14 23:05: ASGC check the links and finds it OK

2010/05/15 00:03: ASGC moves the traffic back to the primary link

Report

<https://twiki.cern.ch/twiki/bin/view/LCG/SIRCernAsgcLinkApril2010>

Tickets:

https://gus.fzk.de/ws/ticket_info.php?ticket=58178

https://gus.fzk.de/ws/ticket_info.php?ticket=58192

https://gus.fzk.de/pages/ticket_lhcopn_details.php?ticket=58197

Improvements



Tickets

Parent-child relationship between GGUS and LHCOPN tickets.

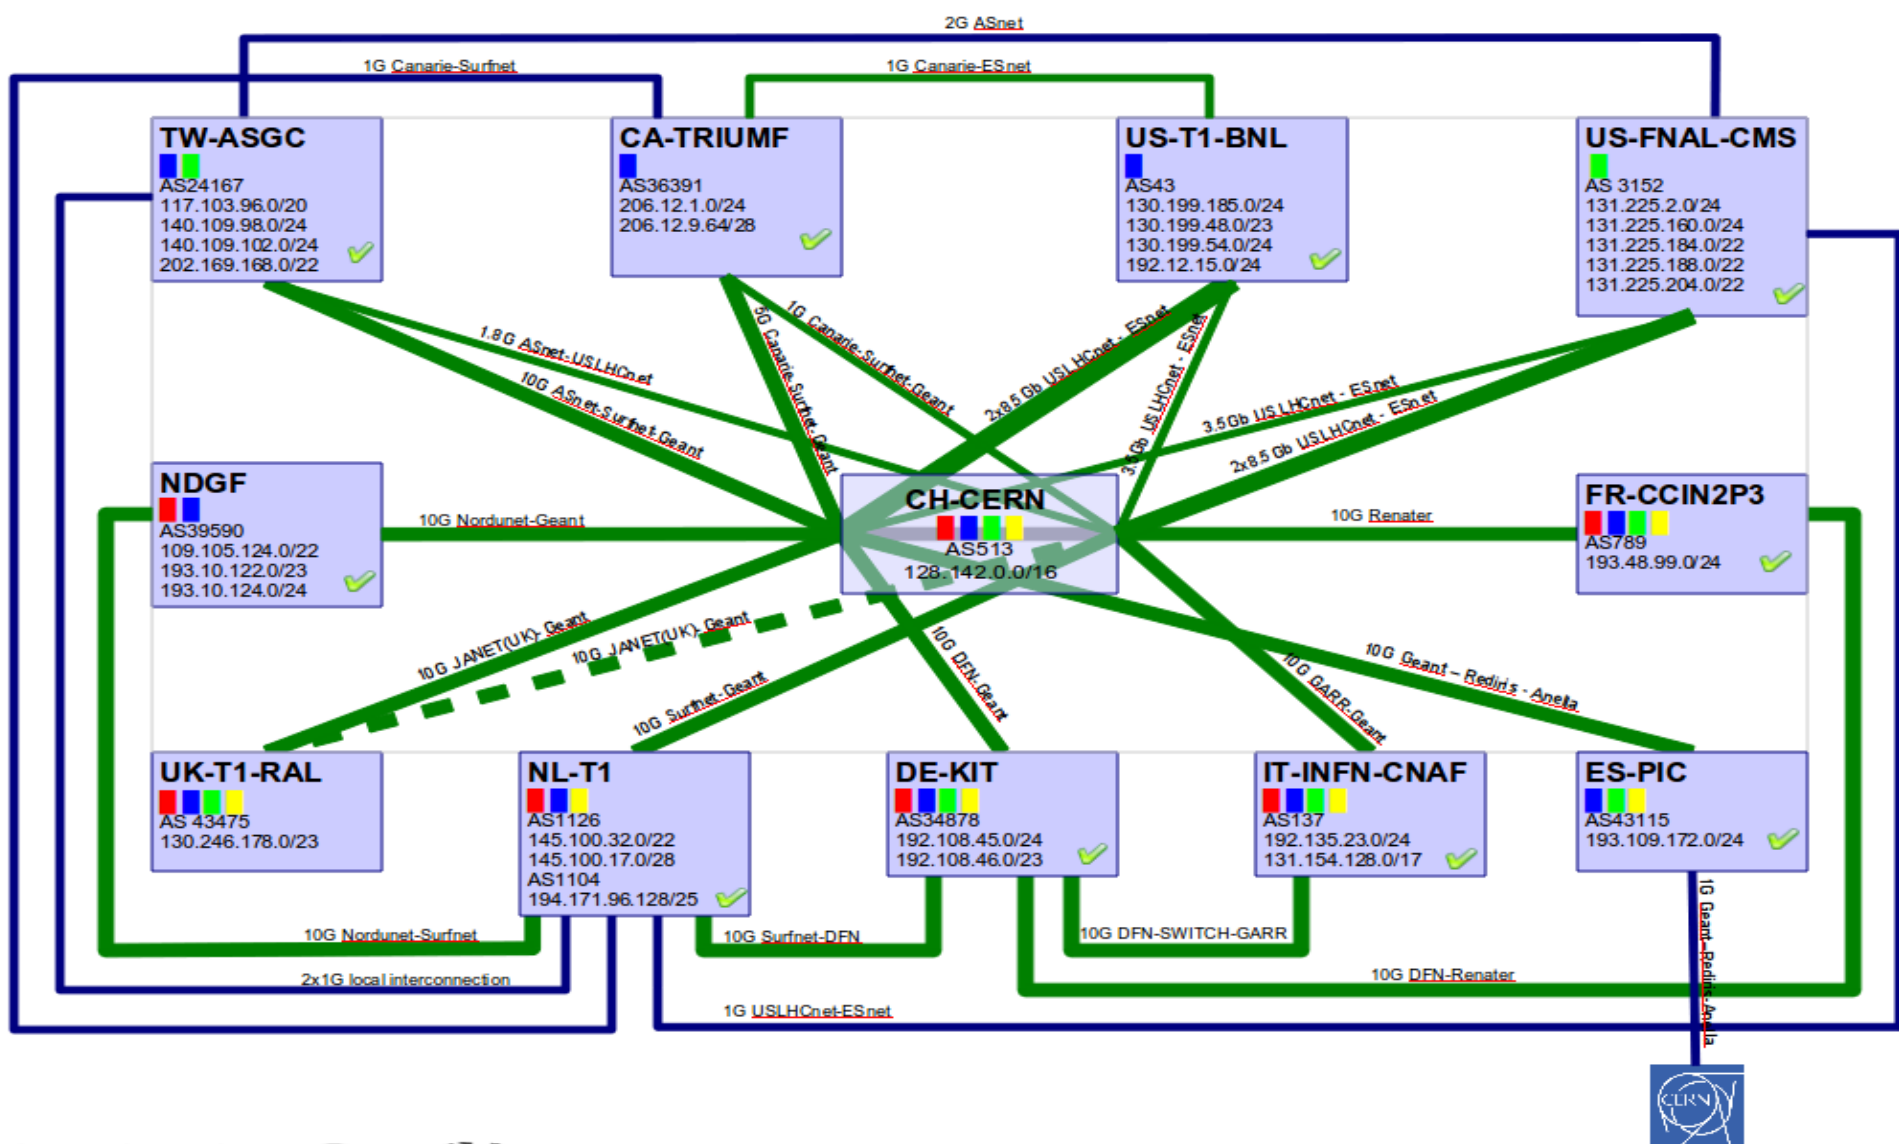
Monitoring:

LHCOPN dashboard:

- Status matrix between any pair of T1/T0 sites. First version based on packet loss and One Way Delay variation only.
- Interface for integration with other dashboards.
- Prototype: October 2010

Questions?

LHCOPN



LHCOPN

	T0-T1 and T1-T1 traffic		= Alice		= Atlas
	T1-T1 traffic only		= CMS		= LHCb
	Not deployed yet		= internet backup available		
	(thick) >= 10Gbps	p2p prefix: 192.16.166.0/24			
	(thin) < 10Gbps	edoardo.martelli@cern.ch 20 100525			