



R&E Networks Debugging End-to-End International Problems

Case Study on FNAL to DESY Throughput Issue

Authors
Chin Guok

Feb 12, 2006
Version 1.0



Contents

| | |
|--|-----------|
| Executive Summary | 2 |
| Introduction | 2 |
| 1 The Problem | 3 |
| 2 Troubleshooting | 3 |
| 2.1 <i>Initial Troubleshooting and Discoveries</i> | 3 |
| 2.2 <i>Coordinated Debug Session</i> | 6 |
| 2.2.1 Test 1: 16:44 GMT, TCP, 10 parallel streams, 10s..... | 6 |
| 2.2.2 Test 2: 16:50 GMT, TCP, 10 parallel streams, 180s..... | 7 |
| 2.2.3 Test 3: 17:10 GMT, UDP, 2 streams, 100Mbps & 300Mbps, 10s & 180s | 9 |
| 2.2.4 Test 4: 17:16 GMT, UDP, 1 stream, 700Mbit/s, 120s | 9 |
| 2.3 <i>Conclusion</i> | 10 |
| 3 Workaround/Solution | 10 |
| 4 References | 12 |

Executive Summary

Network connectivity to an end user is typically viewed as an end-to-end service. In actuality, connectivity between two hosts may span multiple continents and require the seamless transport of packets over several separately managed and operated networks. This is usually transparent to the user. When a network problem arises, a user will typically report the issue to his/her network provider. This is regardless of the fact that the actual problem may be in a network that is managed by a completely different entity. It is not unusual for a provider network to simply close the case if it determines that the problem lies beyond its infrastructure. In ESnet, a trouble ticket (generated when a problem is reported) is tracked and followed-up on until the issue is resolved, regardless of where the problem is.

Introduction

Fixing a network problem may be as simple as plugging in a connector that has come loose, or correcting a duplex mismatch on an interface. Locating the network problem however, is typically much more involved even if the fix is simple. This is compounded by the fact that a problem observed (and reported) by an end user may be in a network that is managed and operated by a different organization. It is often the case that distinct networks (or Autonomous Systems (ASes)) interact with one another as peers and not customers. As such, they are not obligated to assist in debugging a network problem reported by a peer AS, but do so only voluntarily. Even when assistance is offered by a peering AS, effective troubleshooting is often hampered by scheduling (different time-zones), varying levels of expertise, and differing operational and management procedures.

Below is a case study of a network problem that spanned five countries (USA, UK, France, Switzerland, Germany), seven time-zones, and five ASes (FNAL, ESnet, GEANT, DFN, DESY).

1 The Problem

On Nov 25, 2005, a user from DESY sent an e-mail (see Figure 1.1) to FNAL, GEANT, and DFN reporting poor throughput between FNAL and DFN. This was subsequently forwarded to ESnet trouble@es.net on Nov 28, 2005.

```
From: Michael Ernst [mailto:Michael.Ernst@desy.de]
Sent: Friday, November 25, 2005 9:22 AM
To: vyto@fnal.gov; Chris Welti; Robert Stoy
Cc: 'Dan Yocum'; 'Michael Ernst'; bobyshev@fnal.gov; 'Phil DeMar'; 'Donna
Lamore'; wilhelm@dfn.de; ifisk@fnal.gov
Subject: FNAL => DESY data transfer problems

Dear Vyto, Chris and Robert,
  thank you very much for agreeing to investigate the Fermilab => DESY
data transfer problem.

As most of you know this problem persists for quite a while, and, despite
the valuable effort of a good number of talented people, has never been
resolved. This has become a burning issue since DESY is in the process of
building a Tier2 center for ATLAS and CMS and is heavily involved in the
currently running Service Challenges. As a result of the poor network
performance between Fermilab and DESY, DESY is currently unable to fully
participate in this worldwide effort in a way we are expected to.

Therefore DESY is kindly asking you to spend any affordable effort that
is required to find an explanation as to where and why this bottleneck
exists and what it would take to resolve the issue.

According to my understanding the following people representing their
respective organizations are willing to work on the problem

Vyto      <vyto@fnal.gov>      Fermilab
Chris     <welti@switch.ch>   GEANT
Robert    <stoy@noc.dfn.de>    DFN

Thank you very much again for your help,

Michael
```

Figure 1.1: Trouble report from Michael Ernst

2 Troubleshooting

The following bullets are an outline (in chronological order) of the actions that were taken in troubleshooting the issue. For full details please refer to ESnet trouble tickets 13594, 14191, 14233, and 14265.

2.1 Initial Troubleshooting and Discoveries

Nov 28 2005

An ESnet Performance Center [1] account was created for Chris Welti (SWITCH) to run throughput tests. His initial tests between a host on the SWITCH network (ezmp3.switch.ch) to the Performance Center test node in Chicago (chi-pt1.es.net) yielded 837 Mbps. This was done using Iperf [2] with 10 parallel streams and a TCP window size of 2 MB.

Dec 2 2005

Chris Welti (SWITCH) performed more throughput tests and achieved the following results:

SWITCH (cemp1.switch.ch) to ESNET (chi-pt1.es.net): 983 Mbs

SWITCH (cemp1.switch.ch) to DFN (ws-stu2.g-win.dfn.de): 933 Mbs

DFN (ws-stu2.g-win.dfn.de) to SWITCH (cemp1.switch.ch): 930 Mbs

DFN (ws-stu2.g-win.dfn.de) to ESNET (chi-pt1.es.net): 84 Mbs

The results were confusing because the path from DFN to SWITCH, and SWITCH to ESnet, was completely covered by the DFN to ESnet path (see Table 2.1.1). This was also the case for the return route from ESnet to DFN (see Table 2.1.2).

| Traceroute from DFN to SWITCH | Traceroute from SWITCH to ESnet | Traceroute from DFN to ESnet |
|---|--|--|
| 1 cr-stuttgart1-ge5-1 (188.1.76.69) 0.109 ms | | 1 cr-stuttgart1-ge5-1 (188.1.76.69) 0.085 ms |
| 2 cr-frankfurt1-po3-0 (188.1.18.69) 3.531 ms | | 2 cr-frankfurt1-po3-0 (188.1.18.69) 3.356 ms |
| 3 * * * | 1 swiCE2-G5-2 (130.59.35.129) 0.254 ms | 3 * * * |
| 4 * * * | 2 switch.rtl.gen.ch.geant2.net (62.40.124.21) 0.250 ms | 4 * * * |
| 5 swiCE2-10GE-1-1.switch.ch (62.40.124.22) 11.533 ms | 3 ch.fr1.fr.geant.net (62.40.96.30) 8.601 ms | 5 ch.fr1.fr.geant.net (62.40.96.30) 19.753 ms |
| 6 cemp1-eth1.switch.ch (130.59.35.130) 11.397 ms | 4 fr.uk1.uk.geant.net (62.40.96.90) 15.623 ms | 6 fr.uk1.uk.geant.net (62.40.96.90) 26.781 ms |
| | 5 uk.nyl.ny.geant.net (62.40.96.169) 85.466 ms | 7 uk.nyl.ny.geant.net (62.40.96.169) 95.448 ms |
| | 6 esnet-gw.nyl.ny.geant.net (62.40.105.26) 84.319 ms | 8 esnet-gw.nyl.ny.geant.net (62.40.105.26) 95.814 ms |
| | 7 chicr1-oc192-aoacr1.es.net (134.55.209.57) 104.320 ms | 9 chicr1-oc192- aoacr1.es.net (134.55.209.57) 115.450 ms |
| | 8 perf-if1-ge-chi.es.net (198.124.238.26) 104.203 ms | 10 perf-if1-ge-chi.es.net (198.124.238.26) 115.389 ms |

Table 2.1.1: Traceroute from DFN to ESnet

| Traceroute from SWITCH to DFN | Traceroute from ESnet to SWITCH | Traceroute from ESnet to DFN |
|-------------------------------|--|--|
| | 1 chi-ge-perf-if1 (198.124.238.25) 0.245 ms | 1 chi-ge-perf-if1 (198.124.238.25) 0.258 ms |
| | 2 aoacr1-oc192-chicr1 (134.55.209.58) 20.227 ms | 2 aoacr1-oc192-chicr1 (134.55.209.58) 20.236 ms |

| Traceroute from SWITCH to DFN | Traceroute from ESnet to SWITCH | Traceroute from ESnet to DFN |
|---|--|---|
| | 3 esnet.nyl.ny.geant.net (62.40.105.25) 20.324 ms | 3 esnet.nyl.ny.geant.net (62.40.105.25) 20.318 ms |
| | 4 ny.uk1.uk.geant.net (62.40.96.170) 89.002 ms | 4 ny.uk1.uk.geant.net (62.40.96.170) 88.990 ms |
| 1 swiCE2-G5-2 (130.59.35.129) 0.260 ms | 5 uk.fr1.fr.geant.net (62.40.96.89) 96.142 ms | 5 uk.fr1.fr.geant.net (62.40.96.89) 96.028 ms |
| 2 switch.rtl.gen.ch.geant2.net (62.40.124.21) 0.244 ms | 6 fr.ch1.ch.geant.net (62.40.96.29) 104.262 ms | 6 fr.ch1.ch.geant.net (62.40.96.29) 104.270 ms |
| 3 so-7-2- 0.rtl.fra.de.geant2.net (62.40.112.22) 8.341 ms | 7 swiCE2-10GE-1-1.switch.ch (62.40.124.22) 104.320 ms | 7 so-7-2- 0.rtl.fra.de.geant2.net (62.40.112.22) 112.355 ms |
| 4 dfn- gw.rtl.fra.de.geant2.net (62.40.124.34) 8.274 ms | 8 cempl-eth1.switch.ch (130.59.35.130) 104.254 ms | 8 dfn- gw.rtl.fra.de.geant2.net (62.40.124.34) 112.294 ms |
| 5 cr-stuttgart1-po3-0.g- win.dfn.de (188.1.18.70) 11.599 ms | | 9 * |
| 6 ws-stu2.g-win.dfn.de (195.37.209.114) 11.387 ms | | 10 ws-stu2.g-win.dfn.de (195.37.209.114) 115.422 ms |

Table 2.1.2: Traceroute from ESnet to DFN

To further investigate this anomaly, ESnet offered to configure filters specifically to count packets between ws-stu2.g-win.dfn.de and chi-pt1.es.net. These filters were placed at the ESnet-GEANT DMZ, and also on the router interface (on chi-cr1.es.net) connecting to chi-pt1.es.net. By comparing the counters, it could quickly determine if packets were being dropped within ESnet.

Vyto Grigaliunas (FNAL) suggested that the end-to-end path be checked for traffic classification/shaping/policing.

ESnet confirmed that there was no traffic classification, shaping, or policing applied to this traffic.

Chris Welti (SWITCH) conversed with Robert Stoy (DFN) and found out that the two links from Frankfurt to Berlin, and from Berlin to Hamburg were MPLS label switched, although they were both 10 Gbs links, with link utilization of only about 10%. Additionally, these links were totally transparent and there was no traffic classification, shaping or policing applied.

Dec 5 2005

Toby Rodwell (DANTE) offered to assist and configures identical packet counters in DANTE.

Michael Ernst (DESY) acquired an account on a host in CERN (lxshare220d.cern.ch) from the LCG Deployment Team and performed throughput testing. His results were similar to Chris Welti's (SWITCH):

CERN (lxshare220d.cern.ch) to DESY (dcache31.desy.de): 776 Mbps
 DFN (ws-stu2.g-win.dfn.de) to ESnet (chi-pt1.es.net): 83.9 Mbps

Chris Welti (SWITCH) decided that the end points for the throughput tests should be fntst-1.fnal.gov and dcache33.desy.de because they better characterized real end systems. Additional filters were configured in ESnet, DANTE (Toby Rodwell), DFN (Robert Stoy), and DESY (Michael Ernst) to count the traffic flow between the test host pair. A traceroute from dcache33.desy.de to fntst-1.fnal.gov (see Table 2.1.3) showed links in common with the traceroute from DFN to ESnet (Table 2.1.1).

| | | | | |
|----|---|--------------|------------|------------|
| 1 | rt-80-4 (131.169.80.4) | 0.258 ms | 0.249 ms | 0.232 ms |
| 2 | rt-11-16 (131.169.11.16) | 4.431 ms | 4.340 ms | 3.854 ms |
| 3 | cr-hamburg1-ge5-3.x-win.dfn.de (188.1.47.41) | 2.192 ms | 2.685 ms | 2.726 ms |
| 4 | cr-berlin1-po0-0.x-win.dfn.de (188.1.18.109) | 57.770 ms | 39.140 ms | 46.980 ms |
| | MPLS Label=476 CoS=0 TTL=127 S=0 | | | |
| 5 | cr-frankfurt1-po13-0.x-win.dfn.de (188.1.18.54) | 14.450 ms | 16.250 ms | 16.090 ms |
| 6 | dfn.rtl.fra.de.geant2.net (62.40.124.33) | 18.983 ms | 17.615 ms | 15.016 ms |
| 7 | so-6-2-0.rtl.gen.ch.geant2.net (62.40.112.21) | 22.597 ms | 24.258 ms | 21.775 ms |
| 8 | ch.fr1.fr.geant.net (62.40.96.30) | 31.274 ms | 31.101 ms | 32.724 ms |
| 9 | fr.uk1.uk.geant.net (62.40.96.90) | 37.110 ms | 36.944 ms | 36.995 ms |
| 10 | uk.nyl.ny.geant.net (62.40.96.169) | 106.461 ms | 108.397 ms | 110.199 ms |
| 11 | esnet-gw.nyl.ny.geant.net (62.40.105.26) | 106.962 ms | 107.137 ms | 116.487 ms |
| 12 | chislsdn1-chicr1.es.net (134.55.207.34) | 127.075 ms | 126.421 ms | 127.321 ms |
| 13 | chicr1-chislsdn1.es.net (134.55.207.33) | 126.662 ms | 127.260 ms | 126.489 ms |
| 14 | fnal-pos-chi.es.net (134.55.209.37) | 177.842 ms * | 229.786 ms | |
| 15 | ge2-1.r-s-bdr.fnal.gov (198.49.208.4) | 127.637 ms | 129.497 ms | 127.525 ms |
| 16 | vlan313.r-s-hub-fcc.fnal.gov (131.225.15.62) | 127.860 ms | 128.162 ms | 126.979 ms |
| 17 | vlan302.r-s-starlight-cd.fnal.gov (131.225.15.29) | 128.158 ms | 129.925 ms | 127.473 ms |
| 18 | fntst-1.fnal.gov (131.225.2.49) | 127.738 ms | 128.627 ms | 127.097 ms |

Table 2.1.3: Traceroute from dcache33.desy.de to fntst-1.fnal.gov

2.2 Coordinated Debug Session

With the appropriate filters (to capture packets between dcache33.desy.de and fntst-1.fnal.gov) configured in ESnet, DANTE, DFN, and DESY, a conference call was setup on the ESnet audiobridge for Dec 8, 2005 16:00GMT to coordinate the troubleshooting efforts.

2.2.1 Test 1: 16:44 GMT, TCP, 10 parallel streams, 10s

Date of test: Dec 8, 2005 16:44GMT

Duration: 10 seconds

Test profile: 10 parallel TCP streams

Achieved Bandwidth: 30 Mbps

Total packets sent from FNAL to DESY: 183,783

Total packets sent from DESY to FNAL: 103,682

Results: 2 packets lost on the way from FNAL to DESY (see Table 2.2.1.1)

No loss found on the way from DESY to FNAL (see Table 2.2.1.2)

| Network | Description of Counter | Packets |
|---------|------------------------|---------|
|---------|------------------------|---------|

| | | |
|-------|--|-----------------------------|
| ESnet | FNAL (US) -> ESnet (US) | 183,783 |
| | ESnet (US) -> GEANT (US) | 183,783 |
| GEANT | ESnet (US) -> GEANT (US) | 183,783 |
| | GEANT (US) -> GEANT (UK) | 183,783 |
| | GEANT (UK) -> GEANT (France) | 183,781 ^{*2.2.1.1} |
| | GEANT (France) -> GEANT (Switzerland) | 183,781 |
| | GEANT (Switzerland) -> GEANT (Germany) | 183,781 |
| | GEANT (Germany) -> DFN (Germany) | 183,781 |
| DFN | GEANT (Germany) -> DFN (Germany) | 183,781 |
| | DFN (Germany) -> DESY (Germany) | 183,781 |

Table 2.2.1.1: Packet Counts for fntst-1-fnal.gov -> dcache33.desy.de Test1

| Network | Description of Counter | Packets |
|---------|--|-----------------------|
| DFN | DESY (Germany) -> DFN (Germany) | 103,682 |
| | DFN (Germany) -> GEANT (Germany) | 103,682 |
| GEANT | DFN (Germany) -> GEANT (Germany) | 103,682 |
| | GEANT (Germany) -> GEANT (Switzerland) | 103,682 |
| | GEANT (Switzerland) -> GEANT (France) | 103,682 |
| | GEANT (France) -> GEANT (UK) | 103,682 |
| | GEANT (UK) -> GEANT (US) | 103,682 |
| | GEANT (US) -> ESnet (US) | 103,682 |
| ESnet | GEANT (US) -> ESnet (US) | 103,682 |
| | ESnet (US) -> FNAL (US) | 0 ^{*2.2.1.2} |

Table 2.2.1.2: Packet Counts for dcache33.desy.de -> fntst-1-fnal.gov Test1

^{*2.2.1.1} 2 packets were lost due to SDH framing errors (loss rate of 0.0000108).

^{*2.2.1.2} Due to misconfiguration, packets at the ESnet -> FNAL border were not counted.

2.2.2 Test 2: 16:50 GMT, TCP, 10 parallel streams, 180s

Date of test: Dec 8, 2005 16:50GMT

Duration: 180 seconds

Test profile: 10 parallel TCP streams

Achieved Bandwidth: 150 Mbps

Total packets sent from FNAL to DESY: 2,498,589

Total packets sent from DESY to FNAL: 1,277,772

Results: 2 packets lost on the way from FNAL to DESY (see Table 2.2.2.1)

No loss found on the way from DESY to FNAL (see Table 2.2.2.2)

| Network | Description of Counter | Packets |
|---------|--|-------------------------------|
| ESnet | FNAL (US) -> ESnet (US) | 2,498,589 |
| | ESnet (US) -> GEANT (US) | 2,498,589 |
| GEANT | ESnet (US) -> GEANT (US) | 2,498,589 |
| | GEANT (US) -> GEANT (UK) | 2,498,589 |
| | GEANT (UK) -> GEANT (France) | 2,498,543* ^{2.2.2.1} |
| | GEANT (France) -> GEANT (Switzerland) | 2,498,543 |
| | GEANT (Switzerland) -> GEANT (Germany) | 2,498,543 |
| | GEANT (Germany) -> DFN (Germany) | 2,498,543 |
| DFN | GEANT (Germany) -> DFN (Germany) | 2,498,543 |
| | DFN (Germany) -> DESY (Germany) | 2,498,543 |

Table 2.2.2.1: Packet Counts for fntst-1-fnal.gov -> dcache33.desy.de Test2

| Network | Description of Counter | Packets |
|---------|--|-----------|
| DFN | DESY (Germany) -> DFN (Germany) | 1,277,772 |
| | DFN (Germany) -> GEANT (Germany) | 1,277,772 |
| GEANT | DFN (Germany) -> GEANT (Germany) | 1,277,772 |
| | GEANT (Germany) -> GEANT (Switzerland) | 1,277,772 |
| | GEANT (Switzerland) -> GEANT (France) | 1,277,772 |
| | GEANT (France) -> GEANT (UK) | 1,277,772 |
| | GEANT (UK) -> GEANT (US) | 1,277,772 |
| | GEANT (US) -> ESnet (US) | 1,277,772 |
| ESnet | GEANT (US) -> ESnet (US) | 1,277,772 |
| | ESnet (US) -> FNAL (US) | 1,277,772 |

Table 2.2.2.2: Packet Counts for dcache33.desy.de -> fntst-1-fnal.gov Test2

*^{2.2.2.1} 46 packets were lost due to SDH framing errors (loss rate of 0.0000184)

2.2.3 Test 3: 17:10 GMT, UDP, 2 streams, 100 Mbps & 300 Mbps, 10s & 180s

Date of test: Dec 8, 2005 17:10 GMT

Duration: 190 seconds

Test profile: 100 Mbps UDP stream for 10 seconds, followed by 300 Mbps UDP stream for 180 seconds

Effective outgoing rate: 92.6 Mbps, 265 Mbps

Effective received rate: 92.3 Mbps, 264 Mbps

Total packets sent from FNAL to DESY: 2,787,179

Results: 10,352 packets lost*^{2.2.3.1} (as reported by receiving host) on the way from FNAL to DESY (see Table 2.2.3.1)

63 packets re-ordered (as reported by receiving host)

| Network | Description of Counter | Packets |
|---------|--|-------------------------------|
| ESnet | FNAL (US) -> ESnet (US) | 2,787,179 |
| | ESnet (US) -> GEANT (US) | 2,787,179 |
| GEANT | ESnet (US) -> GEANT (US) | 2,787,179 |
| | GEANT (US) -> GEANT (UK) | 2,787,179 |
| | GEANT (UK) -> GEANT (France) | 2,787,074* ^{2.2.3.2} |
| | GEANT (France) -> GEANT (Switzerland) | 2,787,074 |
| | GEANT (Switzerland) -> GEANT (Germany) | 2,787,074 |
| | GEANT (Germany) -> DFN (Germany) | 2,787,074 |
| DFN | GEANT (Germany) -> DFN (Germany) | 2,787,074 |
| | DFN (Germany) -> DESY (Germany) | 2,787,074 |

Table 2.2.3.1: Packet Counts for fntst-1-fnal.gov -> dcache33.desy.de Test3

*^{2.2.3.1} 10,247 packets lost between DFN and dcache33.desy.de

*^{2.2.3.2} 105 packets were lost due to SDH framing errors (loss rate of 0.0000377)

2.2.4 Test 4: 17:16 GMT, UDP, 1 stream, 700 Mbs, 120s

Date of test: Dec 8, 2005 17:16 GMT

Duration: 120 seconds

Test profile: 700Mbps UDP stream

Effective outgoing rate: 497 Mbps

Effective received rate: 486 Mbps

Total packets sent from FNAL to DESY: 5,076,311

Results: 98,035 packets lost*^{2.2.4.1} (as reported by receiving host) on the way from FNAL to DESY (see Table 2.2.4.1)

41,304 packets re-ordered (as reported by receiving host)

| Network | Description of Counter | Packets |
|---------|--|-------------------------------|
| ESnet | FNAL (US) -> ESnet (US) | 5,074,727* ^{2.2.4.2} |
| | ESnet (US) -> GEANT (US) | 5,074,727 |
| GEANT | ESnet (US) -> GEANT (US) | 5,074,727 |
| | GEANT (US) -> GEANT (UK) | 5,074,727 |
| | GEANT (UK) -> GEANT (France) | 5,076,088* ^{2.2.4.3} |
| | GEANT (France) -> GEANT (Switzerland) | 5,076,088 |
| | GEANT (Switzerland) -> GEANT (Germany) | 5,076,088 |
| | GEANT (Germany) -> DFN (Germany) | 5,076,088 |
| DFN | GEANT (Germany) -> DFN (Germany) | 5,076,088 |
| | DFN (Germany) -> DESY (Germany) | 5,076,088 |

Table 2.2.4.1 Packet Counts for fntst-1-fnal.gov -> dcache33.desy.de Test3

*^{2.2.4.1} 96,228 packets were lost between DFN and dcache33.desy.de

*^{2.2.4.2} 1,584 packets were lost between fntst-1.fnal.gov and ESnet

*^{2.2.4.3} 223 packets are lost due to SDH framing errors (loss rate of 0.0000439)

2.3 Conclusion

It was determined after the coordinated troubleshooting session that the GEANT link (serviced by COLT) between UK and France was to be taken down to resolve the SDH framing errors issues.

3 Workaround/Solution

On Dec 9 2005, COLT shutdown the GEANT link between UK and France to swap out a suspect card in one of the switch nodes. This caused a reroute through Austria and Italy (see Table 3.1).

| | | | | |
|----|---|------------|------------|------------|
| 1 | vlan200.r-s-starlight-cd.fnal.gov (131.225.2.200) | 0.430 ms | 0.372 ms | 0.220 ms |
| 2 | te4-3.r-s-starlight-fnal.fnal.gov (198.151.133.234) | 0.602 ms | 0.272 ms | 0.242 ms |
| 3 | 198.151.133.185 (198.151.133.185) | 1.496 ms | 2.411 ms | 1.487 ms |
| 4 | chicr1-chislsdn1.es.net (134.55.207.33) | 1.559 ms | 1.419 ms | 1.434 ms |
| 5 | aoacr1-ocl92-chicr1.es.net (134.55.209.58) | 21.572 ms | 21.485 ms | 21.475 ms |
| 6 | esnet.nyl.ny.geant.net (62.40.105.25) | 21.649 ms | 21.921 ms | 21.928 ms |
| 7 | ny.at1.at.geant.net (62.40.96.110) | 127.871 ms | 127.970 ms | 127.978 ms |
| 8 | at.de2.de.geant.net (62.40.96.58) | 140.363 ms | 140.386 ms | 140.730 ms |
| 9 | de.it1.it.geant.net (62.40.96.62) | 149.709 ms | 149.605 ms | 149.532 ms |
| 10 | so-6-3-0.rtl.gen.ch.geant2.net (62.40.112.33) | 156.871 ms | 156.739 ms | 156.783 ms |
| 11 | so-7-2-0.rtl.fra.de.geant2.net (62.40.112.22) | 164.918 ms | 164.820 ms | 164.846 ms |
| 12 | dfn-gw.rtl.fra.de.geant2.net (62.40.124.34) | 164.803 ms | 165.030 ms | 164.807 ms |
| 13 | cr-berlinl-pol-0.x-win.dfn.de (188.1.18.53) | 177.862 ms | 177.757 ms | 177.942 ms |
| | MPLS Label=444 CoS=6 TTL=1 S=0 | | | |
| 14 | cr-hamburg1-pol0-0.x-win.dfn.de (188.1.18.110) | 178.001 ms | 178.135 ms | |

Table 3.1: Traceroute from fntst-1.fnal.gov to dcache33.desy.de

Michael Ernst (DESY) reported a jump of throughput from 30Mbps to 160Mbps when the route changed.

In Jan 2006, GEANT switched the route to utilize a transatlantic line between the US and the Netherlands. This caused traffic from FNAL to DESY to bypass the links in UK and France (see Table 3.2), and as a result, the throughput in increased to 600-700Mbps (as report by Chris Welti (SWITCH) on Feb 06, 2006).

| | | | | |
|----|--|------------|------------|------------|
| 1 | chicr1-chislsdn1.es.net (134.55.207.33) | 0.517 ms | 1.110 ms | 1.065 ms |
| 2 | aoacr1-ocl92-chicr1.es.net (134.55.209.58) | 20.510 ms | 20.504 ms | 20.494 ms |
| 3 | esnet.nyl.ny.geant.net (62.40.105.25) | 85.800 ms | 20.638 ms | 20.644 ms |
| 4 | so-7-0-0.rtl.ams.nl.geant2.net (62.40.112.133) | 153.452 ms | 103.777 ms | 103.784 ms |
| 5 | so-6-2-0.rtl.fra.de.geant2.net (62.40.112.57) | 111.043 ms | 111.022 ms | 110.999 ms |
| 6 | dfn-gw.rtl.fra.de.geant2.net (62.40.124.34) | 110.980 ms | 110.937 ms | 110.914 ms |
| 7 | cr-berlinl-pol-0.x-win.dfn.de (188.1.18.53) | 128.030 ms | 128.073 ms | 128.016 ms |
| | MPLS Label=556 CoS=0 TTL=1 S=1 | | | |
| 8 | cr-hamburg1-pol0-0.x-win.dfn.de (188.1.18.110) | 128.304 ms | 128.421 ms | 128.225 ms |
| 9 | 188.1.47.42 (188.1.47.42) | 128.287 ms | 128.352 ms | 128.322 ms |
| 10 | * * * | | | |
| 11 | * * * | | | |
| 12 | * * * | | | |

Table 3.2: Traceroute from chi-sl-sdn1.es.net to dcache33.desy.de (via US <-> Netherlands link)

4 References

- [1] ESnet Performance Center: <https://performance.es.net>
- [2] Iperf: <http://dast.nlanr.net/Projects/Iperf>