



“That problem we were not having yesterday, is it still there?”


## Introduction to the RIPE NCC Test Traffic Measurements Service

Henk Uijterwaal  
RIPE NCC  
8 February 2003

# Test Traffic Measurements Service


- Our Mission:

“Doing **objective, accurate and reliable** one-way performance measurements on the Internet using IETF standardized metrics”



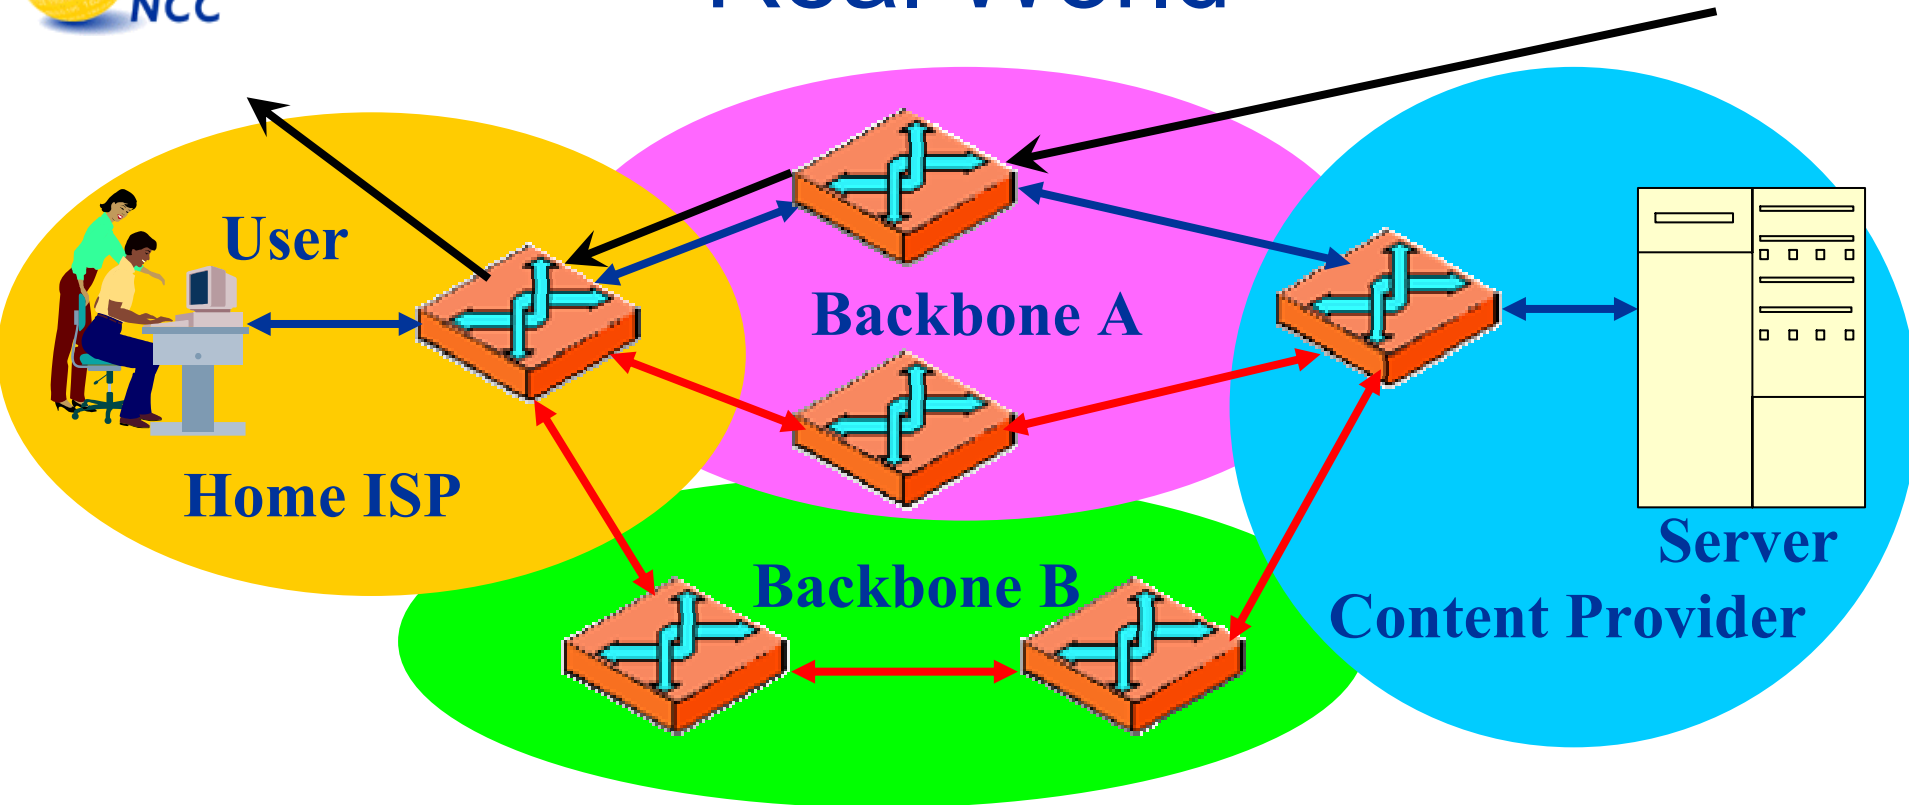
- But hey, **why** should we measure something?

 User: “The Internet is slow today!”



- People are interested in end-to-end performance, i.e. “from my PC to that www or ftp site”
- Maybe we know properties of the individual components, e.g. link capacity & hardware specifications, but *can we simply add them up?*

# Real World



- Implications of load on hardware, cross-traffic, bandwidth utilization etc. (“**dynamic interference**”)
- Network equipment and infrastructure are owned by **different** providers: What are the properties of individual components? What **actually** is installed?

# Service Provider's Perspective

- An ISP wants to deliver a good service to its customers
- **Questions:**
  - Is my equipment working properly?
  - Is the equipment of our upstream working properly?
  - Do we meet our SLA's?
  - Do they meet their SLA's?
  - Where do we need more capacity?
- Answering these questions is **hard**





# Why Measure on the Internet?

- Because...
  - The user is interested in **perceived** end-to-end performance
  - Performance that the user sees is **not** the sum of the individual components (“dynamic interference”)
  - A single provider does **not** have access to all components in a path

*Solution: Measure end-to-end performance parameters*

# Parameters To Be Measured (IPv4/IPv6)

- *Routing information:*
  - Router/Interface level
  - AS-level
- *Delay or Latency*
- *Packet Loss*
- *Bandwidth*
- *Derived quantities:*
  - IPDV or Jitter
  - Packet reordering
  - Protocol specific performance
- *Suggestions from TTM users ;)*



# One-way or Round-trip?

- Why are **one-way** measurements better?
  - Routing is often asymmetrical
  - Problem can be in the incoming or outgoing path
  - Trivial to combine two one-way measurements
  - However, it's **impossible** to decompose a round-trip measurement into two one-ways

*Haven't we forgotten something...*





# One Needs a Meter To Measure...

- That's why you need the RIPE NCC Test Traffic Measurements Service
  - We run and maintain the “meter” for you (TTM test-box)
  - We produce objective and accurate data for you to look at



# One Needs a Meter To Measure

- **Our Goals:**

- Offer a “Black Box Service” that monitors the connectivity of your network to other parts of the Internet
- One-way end-to-end measurements with 0.01 ms accuracy
- Dedicated measurement devices and infrastructure
- Active measurements emulating “real traffic”
- Follow well defined IETF standards (RFCs 2330, 2678-2681)
- We focus on external networks, but techniques **can** be used on internal networks as well
- Both IPv4 and IPv6 supported

Now let's assume that we've installed a test-box...

# Features of the TTM service

- Proactive network alarms
  - notification by e-mail of unexpected changes
- Regular traceroute measurements
  - to learn how test-traffic is (most likely) routed
- Performance measurements
  - Delays and losses:
    - Online: WWW interface on the boxes.
      - <5 minutes delay
    - Offline: Full statistical analysis
      - The following morning
  - IP Delay Variations (IPDV a.k.a. Jitter)
  - Bandwidth

# Alarm Messages

Date: Fri, 19 Apr 2002 05:55:14 GMT  
From: Test traffic project <ttraffic@ripe.net>  
Subject: Testbox ALARM SET on tt01.ripe.net

TB 69 at 1019195711 ALARM SET

TB 69 at 1019195711 Long: 2226 2.0/ 2.5/ 3.0  
Short: 59 10.0/10.5/11.0

[...]

To see how the delays developed in the last days, open this URL: [http://tt01.ripe.net:10259/cgi-bin/multiple.cgi?&tt69=213.181.58.101&delay=delay&loss=loss&RRD\\_STA\\_RT=now-2days&RRD\\_END=now](http://tt01.ripe.net:10259/cgi-bin/multiple.cgi?&tt69=213.181.58.101&delay=delay&loss=loss&RRD_STA_RT=now-2days&RRD_END=now)

Inbound Delays - Netscape 6

http://tt01.ripe.net:10259/cgi-bin/multiple.cgi?tt69=213.181.58.101&delay=delay&loss=los

**Inbound Delays**

Test Box Home | Documentation | Status | Help | Configuration | Measurements | TTM Web Site

Inbound Delay Help | Inbound Delays

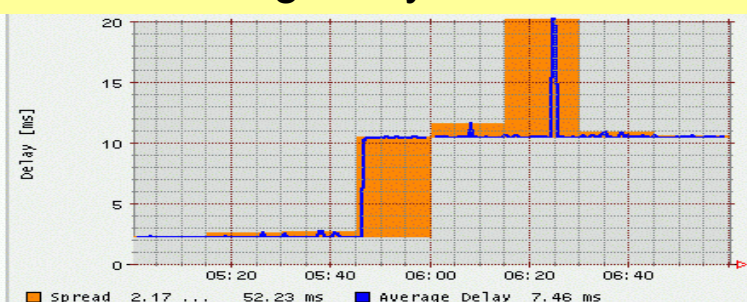
**Select Test-box(es):**

All testboxes

tt02    tt03    tt04    tt07    tt08  
 tt12    tt13    tt16    tt17    tt20  
 tt21    tt22    tt23    tt25    tt26  
 tt27    tt28    tt31    tt32    tt34  
 tt35    tt36    tt39    tt40    tt42  
 tt44    tt46    tt47    tt49    tt52  
 tt53    tt54    tt55    tt56    tt57  
 tt58    tt59    tt64    tt69    tt70  
 tt71    tt72    tt73    tt75    tt76

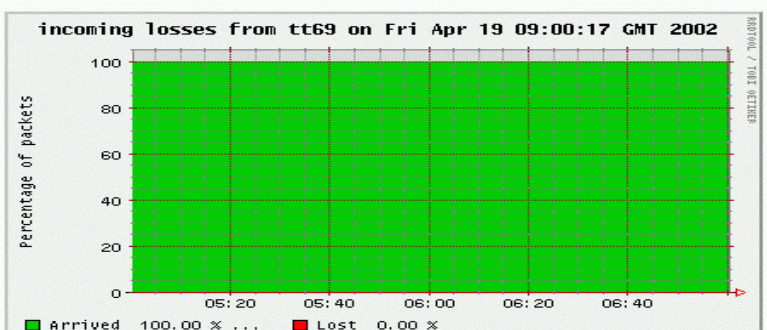
**Delays:**   
**Losses:**   
**Start:** now-4hours   **Length:** start+2h  
**Height:** medium   **Width:** narrow  
**Y-range delay:** 0 to 20   **Loss plots:**  Loss Only    Loss and Arrived

**Incoming delays from tt69**



Spread 2.17 ... 52.23 ms   Average Delay 7.46 ms

**Incoming Losses from tt69 on Fri Apr 19 09:00:17 GMT 2002**



Arrived 100.00 % ...   Lost 0.00 %

Document: Done (1.248 secs)

# Offline analysis: summaries

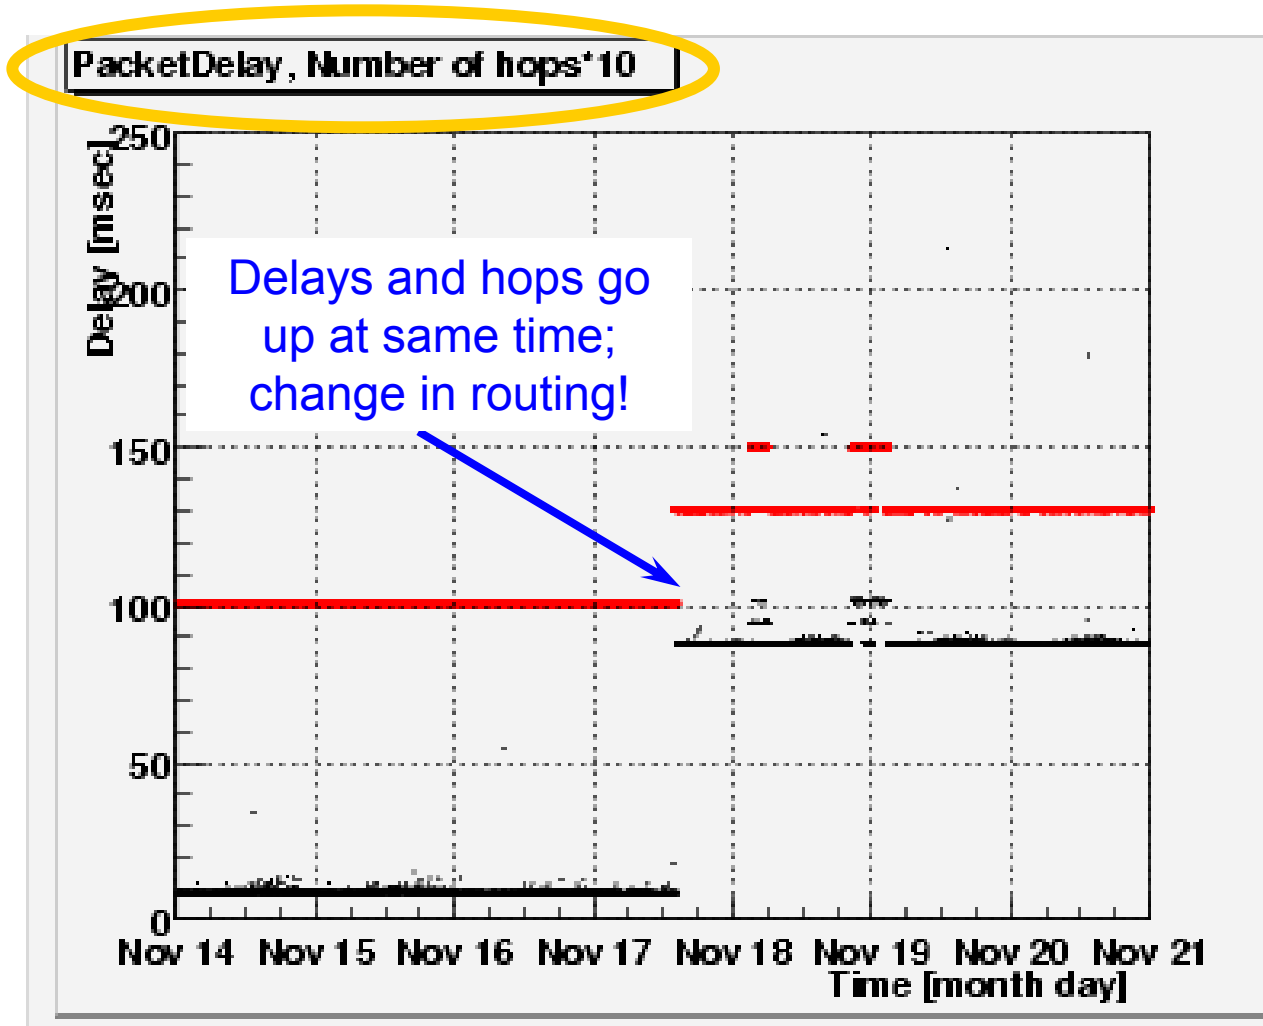
- One page for each box
  - highlight relatively big changes

## RIPE NCC at AMS-IX, Amsterdam, NL

Box name, organisation location	Incoming delay and loss				Outgoing delay and loss			
	Minimum	Median	Maximum	Loss	Minimum	Median	Maximum	Loss
tt59: NIKHEF Amsterdam, NL	<a href="#">0.2210</a> (0.2140)	<a href="#">0.2910</a> (0.2820)	<a href="#">0.8800</a> (0.5320)	<a href="#">1.2</a> (1.2)	<a href="#">0.6350</a> (0.2660)	<a href="#">0.7220</a> (0.4150)	<a href="#">0.8550</a> (0.8330)	<a href="#">0.9</a> (1.1)
tt35: HEAnet Dublin, IE	<a href="#">10.69</a> (10.69)	<a href="#">10.73</a> (10.74)	<a href="#">11.05</a> (10.98)	<a href="#">1.2</a> (1.0)	<a href="#">87.71</a> (8.361)	<a href="#">87.83</a> (8.499)	<a href="#">88.40</a> (9.718)	<a href="#">1.6</a> (1.6)
tt42: NTUA Athens, GR	<a href="#">37.21</a> (37.26)	<a href="#">37.30</a> (37.35)	<a href="#">37.55</a> (37.58)	<a href="#">1.1</a> (1.2)	<a href="#">53.23</a> (37.19)	<a href="#">53.39</a> (37.34)	<a href="#">53.95</a> (37.71)	<a href="#">0.3</a> (0.3)
tt81: @Home Benelux B.V. Amsterdam, NL	<a href="#">0.7180</a> (0.7150)	<a href="#">0.7880</a> (0.7860)	<a href="#">5.857</a> (1.187)	<a href="#">2.0</a> (1.2)	<a href="#">0.7180</a> (0.7150)	<a href="#">0.7880</a> (0.7860)	<a href="#">5.857</a> (1.187)	<a href="#">2.0</a> (1.2)
tt23: Internet2/UCAID Ann Arbor, MI, US	<a href="#">62.00</a> (62.00)	<a href="#">62.05</a> (62.07)	<a href="#">134.6</a> (62.29)	<a href="#">1.5</a> (1.2)	<a href="#">62.80</a> (62.80)	<a href="#">62.88</a> (62.88)	<a href="#">63.08</a> (63.08)	<a href="#">0.1</a> (0.1)
tt75: GlobalSoft.com Ltd. Nicosia, CY	<a href="#">59.17</a> (126.6)	<a href="#">129.6</a> (136.9)	<a href="#">219.7</a> (263.7)	<a href="#">1.3</a> (1.3)	<a href="#">56.28</a> (66.04)	<a href="#">62.86</a> (67.28)	<a href="#">128.7</a> (175.3)	<a href="#">0.1</a> (0.6)

**10x higher delay  
to Dublin**

# Zoom in: Delay vs Time Graph



# Routing at AS level

(from collected traceroutes)

Dates	ASpath Id	Occurrences	Hop	AS-num(s)	AS-name
From Mon Nov 11 21:33:00 2002 To Sun Nov 17 14:07:19 2002	69764	1341	1	<a href="#">3333</a>	RIPE NCC
			2	<a href="#">1200</a>	Amsterdam Internet Exchange (AMS-IX) Peering AS
			3	<a href="#">702</a>	UUNET - Commercial IP service provider in Europe
			4	<a href="#">1849</a>	UUNET UK (formerly PIPEX, Public IP EXchange)
			5	<a href="#">1213</a>	HEAnet national network

Dates	ASpath Id	Occurrences	Hop	AS-num(s)	AS-name
From Sun Nov 17 14:21:07 2002 To Mon Nov 18 04:59:09 2002	158065	143	1	<a href="#">3333</a>	RIPE NCC
			2	<a href="#">1200</a>	Amsterdam Internet Exchange (AMS-IX) Peering AS
			3	<a href="#">9057</a>	Level 3 Communications
			4	<i>unknown</i>	<i>no response OR no matching AS found</i>
			5	<a href="#">1213</a>	HEAnet national network




# Routing at IP Level

- Look at processed results of traceroutes
  - more detail, but also many more entries, routes changing due to load balancing
- Host names of intermediate hops suggest routing via New York City
  - 2x 40ms crossing the Atlantic, not good
  - Policies, BGP dynamics of upstream ISP
  - RIPE NCC re-tuned their BGP configuration, other upstream, traffic stays within Europe again

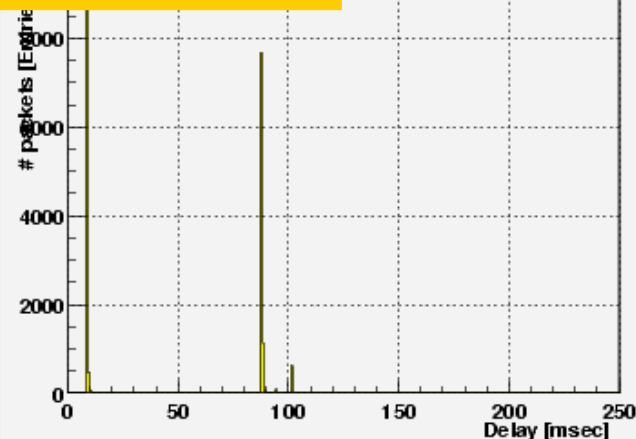
# Full TTM delay plot

(day, week, month, or user specified range)

Packet Delay, #hops\*10



Packet Delay



**STATISTICS:**

**Delay & Hops:**

Entries: 19662  
 Overflow: 14  
 Underflow: 0  
 2.5 Perc: 8.4ms  
 Median: 9.3ms  
 97.5 Perc: 101.2ms  
 Mean: 47.6ms  
 RMS: 40.2ms  
 Min. hops: 10  
 Max. hops: 15

**Packets sent/valid:**

Total: 19992  
 Valid: 19662 = 98.3 %  
 Send bad: 17 = 0.085 %  
 Recv bad: 0 = 0 %  
 2 Clocks bad: 0 = 0 %  
 Lost: 313 = 1.6 %

**Packets lost:**

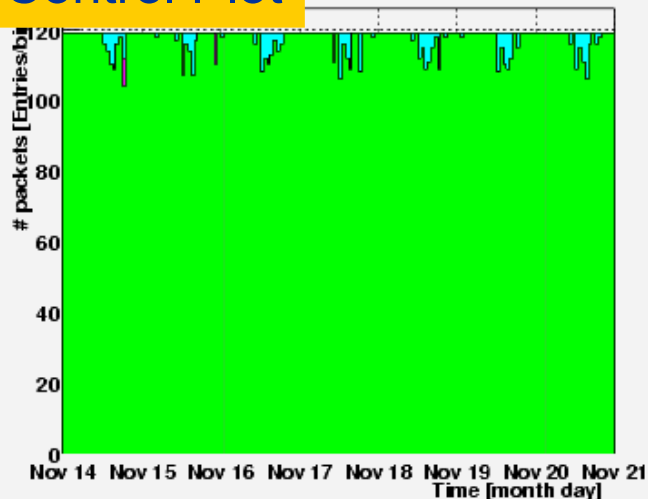
2.5 Perc: 0.0%  
 Median: 0.0%  
 97.5 Perc: 9.2%  
 Uptime: 100 %

**Over-all statistic:**


Time period: 7 days  
 Number of routing vectors: 8  
 flaps: 17  
 Number of bins: 168  
 Minutes/bin: 60

**Statistics**

Control Plot



Packets Arrived





# Test Traffic Measurements

## Plots on Demand

Date Format: yyyyymmdd

Time Format: hh:mm

SRC:  Start Date

DST:  End Date

Min Delay (ms):  Plot Format:  gif  PostScript

Max Delay (ms):  Plot Type:  delay  trends  jitter

tt01.ripe.net -> tt34.ripe.net collected traceroutes from 20020415 10:00 to 20020415 18:00

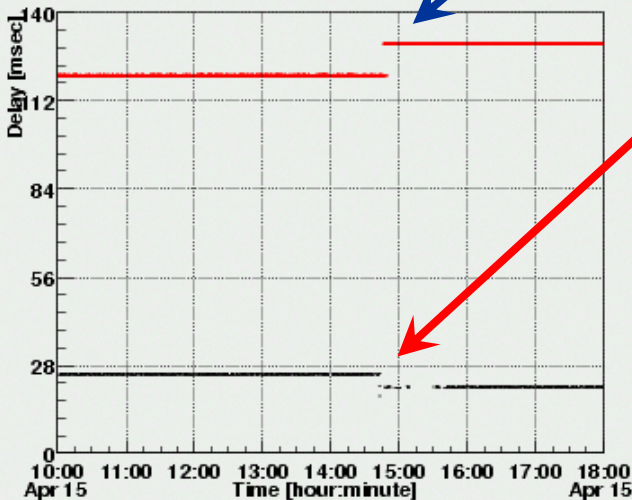
# Change in Routing

Around 14:45


#Hops up from 12 to 13

Delays down 25 to 21ms

### Packet Delay, #hops\*10



### Packet Delay



#### STATISTICS:

##### Delay & Hops:

Entries: 872  
 Overflow: 0  
 Underflow: 0  
 2.5 Perc: 21.0ms  
 Median: 25.1ms  
 97.5 Perc: 25.2ms  
 Mean: 23.7ms  
 RMS: 1.9ms  
 Min. hops: 12  
 Max. hops: 13

##### Packets sent/valid:

Total: 960  
 Valid: 872 = 90.8 %  
 Send bad: 85 = 8.9 %  
 Recv bad: 0 = 0 %  
 2 Clocks bad: 0 = 0 %  
 Lost: 3 = 0.31 %

# Routing

- From Mon Apr 15 07:12:41 2002
- To Mon Apr 15 14:36:20 2002

- X.Y.0.238 g0013.nikrtr.ripe.net
- X.Y.15.44 amster-dam11.z.z
- X.Y.103.105 *Unknown hostname*
- X.Y.96.114 nl.uk1.uk.z.z
- X.Y.96.125 uk.se1.se.z.z
- X.Y.103.118 nordunet-gw.se1.z.z
- X.Y.252.130 s-gw.z.z
- X.Y.68.42 fi-gw.z.z
- X.Y.252.50 Unkn
- X.Y.255.190 funet6-p21-csc0.z.z
- X.Y.187.46 csc1-e00-funet6.z.z
- X.Y.4.105 tt34


- From Mon Apr 15 14:47:12 2002
- To Tue Apr 16 14:00:07 2002

- X.Y.0.238 g0013.nikrtr.ripe.net
- X.Y.15.97 Asd-nr10.NL.z.z
- X.Y.96.81 r1-Gi1-2-1.12.Ledn-KQ1.z.z
- X.Y.230.14 r1-Se0-1-0.0.ffm-KQ1.z.z
- X.Y.230.110 r1-Se0-3-0.0.hmbg-KQ1.z.z
- X.Y.230.150 r2-Se1-1-0-0.Sthm-KQ1.z.z
- X.Y.119.213 sw-gw.z.z
- X.Y.252.130 s-gw.z.z
- X.Y.68.42 fi-gw.z.z
- X.Y.252.50 Unknown address
- X.Y.255.190 funet6-p21-csc0.z.z
- X.Y.187.46 csc1-e00-funet6.z.z
- X.Y.4.105 tt34

Old: NL - UK - SE  
 New: NL - DE - SE  
 shorter physical path wins  
 from more hops

Anonymous and edited to fit on the screen

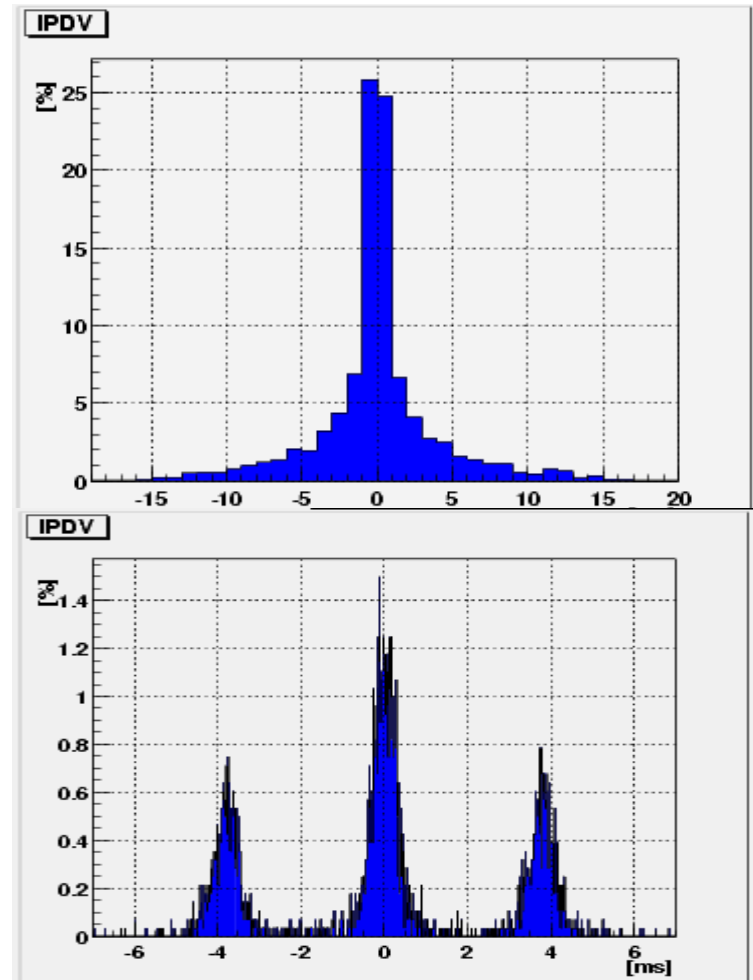
# Trends in the data



- Median delay over 6 months
  - Night • Morning • Afternoon • Evening
- Saturated link in the path from A to B
  - combine with traceroute data to identify the link
- Intended for capacity planning

# IP-Delay Variations or Jitter

- For some applications, absolute delay does not really matter
- But packets **should** arrive with constant intervals
  - Voice over IP
  - Video on demand
- IPDV metric
  - **Difference** in delays of consecutive packets



# Bandwidth

- The next measurement to be added
- We will measure 2 Parameters:
  - C: Total Capacity
    - how much data can be send over a link per unit time
  - A: Available Bandwidth
    - how much data can be send without effecting *existing, current* traffic
- Final stages of  $\beta$ -testing

# How do *you* get TTM results?

- Install our measurement device
- RIPE NCC Test-Box
  - 1u rackmount PC
  - GPS antenna
- Install antenna on roof or wall
  - Up to 250m of UTP cable between box and antenna
- Sign up for the service
- Plug and play





# Service model

- **You buy the hardware once**
  - € 2500
- **You get a service contract for the machine**
  - Maintenance of the machine
  - Measurements between your machine and all others
  - Analysis of the data by the NCC staff
  - Software upgrades, access to new features
  - User support
  - € 3000/year 1<sup>st</sup> machine, volume discounts

*All you have to do, is to look at the results*

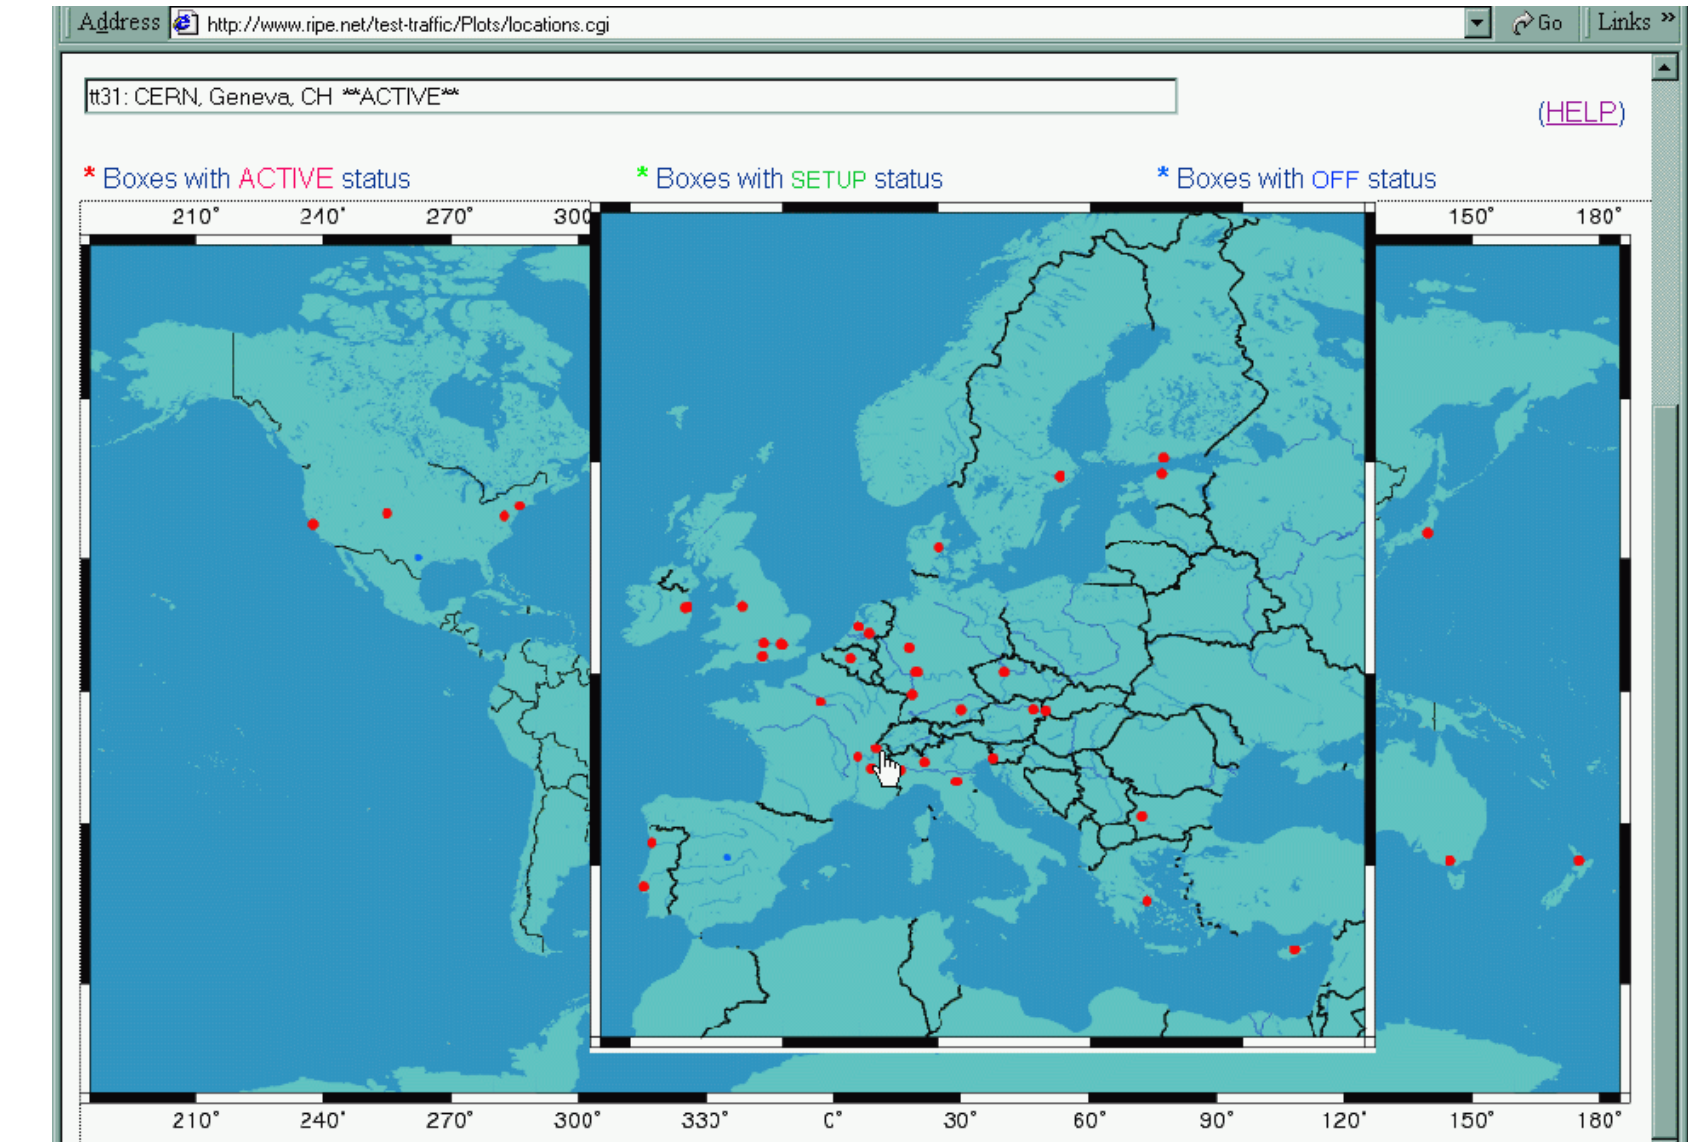


# Why get this service from the RIPE NCC?

- Neutral
- Independent
- Correct results without a commercial bias
- Access to PoP's of competitors
- Load sharing



# Where are all these boxes located?



# Some other features

- IPv6-support?
  - Yes, on all boxes shipped after 1/1/2003
  - Upgrading existing boxes
- Adjust packet rates?
  - Yes, WWW-interface
- Access to raw data?
  - Yes, build your own tools
- Packet reordering?
  - Algorithm being developed in the IETF
- QoS/DiffServ support?
  - Under development

# Conclusion

- Measurements are necessary to get a good overview of end-to-end performance
- Measuring parameters like delay, loss, bandwidth requires a measurement device
- The RIPE NCC test-box is a plug-and-play device to get these numbers, all you have to do is look at the data



# URL's, Contact Addresses

- <http://www.ripe.net/test-traffic>
  - Papers
  - Presentations
  - “For future test-box hosts”
- [ttm@ripe.net](mailto:ttm@ripe.net) : TTM Crew @ NCC
- [tt-wg@ripe.net](mailto:tt-wg@ripe.net): RIPE WG on this topic (Majordomo)

# Questions, Discussion

