



Test-Traffic Project Status and Plans

Henk Uijterwaal, Fotis Georgatos,
Johann Gutauer, Daniel Karrenberg, René Wilhelm
RIPE-NCC New Projects Group

RIPE-36, Budapest, May 2000



Overview

- Manpower, Email
- Status of the measurement network
- Turning TTM into a regular service
- The next series of test-boxes
- Analysis and results
- Plans for the next months



Manpower

- Group of 5
- Johann Gutauer left on April 30
 - M.Sc. Student
 - Finishing his thesis
 - Some results from his thesis later on
- Second Network Engineer starting August 1
 - Installation and support for new test-boxes
 - System programming



Email addresses

- **tt-ops@ripe.net**: Test-box operators at the NCC
 - Operational issues
 - Do **not** use our personal addresses
- **tt-host@ripe.net**: Mailing list for test-box operators
 - Please remember this list when people leave
- **tt-wg@ripe.net**: Mailing list for the working group



Overview


- Manpower, Email
- Status of the measurement network
- Turning TTM into a regular service
- The next series of test-boxes
- Analysis and results
- Plans for the next months



Status of the measurement network

- Status on May 11, 9:00 GMT
- 43 boxes in the field:

- Off: 2
- Setup: 5
 - (Re-)installation
 - Waiting for host: 5
- Watch: 6
 - Data not useful
 - Waiting for NCC: 1
 - Waiting for host: 5
- On/Taking Data: 30





Problems with the older boxes

- Series A boxes (tt01..tt27, tt47) occasionally end up in an undefined state
 - Series B does not show this problem
- Requires a reboot
 - Contact operator
 - Creates downtime
- The problem is a combination of
 - Memory
 - O/S



Upgrade of the series A boxes

- Solution:
 - Increase memory
 - Update O/S
- All problems disappear
- Memory is being sent to TB-hosts
 - Please insert it
 - 10 minutes of work, instructions on the web
- O/S will be updated in the next weeks
 - Host has to swap the disks
 - Instructions will follow



Installation and maintenance procedures

- We now have 2 years of experience with operating 40 boxes at remote sites
- Time to review our procedures and see if they will scale to 100+ boxes
- Some do, some don't
- 2 problematic areas:
 - Keeping the software up to date
 - Check if all processes are still running



Is a box working as it should?

- A number of processes should be running
- 1 job running on a central machine checks this
 - Disadvantages:
 - Does not scale
 - Remote checks cannot solve connectivity problems
 - Hard to handle small differences in setup
 - Solution:
 - Tell boxes which processes should be running
 - Have a remote process check this
 - At the NCC, only check if the box is still reachable
- “There must be a tool to do this”



CFEngine

- Generic tool for system and network maintenance
 - Highly configurable
 - Easy to use
 - Replaces home-grown scripts and crontab entries
- Runs on the test-boxes
 - Box is told what should run on there
 - CFEngine will make sure that the processes are there
- Useful for many other sys-admin applications
 - Homepage & manuals: <http://www.iu.hioslo.no/cfengine/>



CFEngine (2)

- Advantages:
 - Scalable
 - Easy to deal with exceptions
 - Check more often and do more tests
 - A box will continue to take data even the connection to the NCC is lost
 - Reduced work-load for the operator
- Has been introduced successfully for system processes
- Data-taking next



Overview

- Manpower, Email
- Status of the measurement network
- Turning TTM into a regular service
- The next series of test-boxes
- Analysis and results
- Plans for the next months



Turning TTM into a regular service

- RIPE-35: Consensus on the model:
 - Pay for a box (€ 3000)
 - Pay a service fee starting in 2001 (€3000/year)
- Producing a “service contract”
 - Draft exists
 - Still needs feedback from a number of people
 - Turn into RIPE-document after RIPE36



Turning TTM into a regular service (2)

- Current draft:
 - <http://www.ripe.net/test-traffic/RIPE36/note.{ps,pdf,html}>
- TTM “glossy” for your manager
 - Ideas discussed with NCC-COMMS department
 - NCC-COMMS will help with contents and layout
 - Finalize after the service contract is finished



Installation support and site-surveys

- Standard:
 - Host site finds a suitable spot for the box
 - Host site installs the box
- Extra's:
 - Site-survey: Look for a suitable spot for the antenna
 - Installation support
 - Install box and antenna
 - Check connectivity, firewall configuration, etc.
- Any interest?
 - Needs extra resources
 - Not for free!



Overview

- Manpower, Email
- Status of the measurement network
- Turning TTM into a regular service
- The next series of test-boxes
- Analysis and results
- Plans for the next months

The new test-boxes



- 2000 version of the current machines

Antenna: Trimble Palissade



- 1 unit, “Plug and play”
- Standard UTP cable
 - 100’s of meters of cable
- Uses existing drivers
- Working on mounting bracket



To do list for the new boxes

- ➡ Test new PC hardware and finalize specifications
- ✓ Produce a mounting bracket for the antenna
- ➡ Prepare cables

- ➡ Order all hardware
- ✓ Update installation documents
- ✓ Call for sites interested in hosting a box

- Estimate 6...10 weeks, depending on vendors



Overview

- Manpower, Email
 - Status of the measurement network
 - Turning TTM into a regular services
 - The next series of test-boxes
 - Analysis and results
 - Plans for the next months
 - Conclusions
- Raw data
 - Daily plots
 - Routing Vectors
 - Network alarms
 - Trends in the data
 - Network performance scores
 - Next steps




Daily Plots

- Complete re-write
 - Faster
 - Better layout
 - More statistics
- Test-version at:
 - <http://www.ripe.net/test-traffic/Plots/New>
- Move to production version after RIPE36




Daily Plots/Main new features

Hop Count



User selectable time-interval

Delays from tt02 to tt14. Start: 2000-02-16 00:00 End: 2000-02-17 00:00 GMT



Statistics:
Histogram
Delay & Entries:
Entries: 1884
Overflow: 48
Underflow: 1
2.5 Perc.: 43.521
Median: 69.111
97.5 Perc.: 200.043
Mean delay: 79.07
RMS delay: 39.44
Min. hops: 14



Packets sent/valid:
Total: 1994
Valid: 1884 = 94.5 %
SendOK: 0 = 0 %
RecOK: 0 = 0 %
Clocks wr.: 0 = 0 %
Loss: 110 = 5.52 %

Packets arrival/lost:
2.5 Perc.: 0
Median: 0
97.5 Perc.: 1
Uptime: 92.86 %

Over-all statistics:
Number of events
in chain: 67802
Number of routing
vectors: 4
Time-frame: 1 D
Number of bins: 168
Minutes/bin: 8.36667

More Statistics

Absolute Time-scales





Daily Plots/Plots on demand

- Current situation:
 - Generate all plots for all connections
 - Generate plots for 1, 7 and 30 days
- This is not perfect:
 - CPU-time scales as N^2
 - A lot of the 5000 plots that we produce is never looked at by a human
 - People want other time intervals



Daily Plots/Plots on demand

- Only generate the most important plots
 - First guess, will be fine-tuned
- Generate other plots on demand
 - Web-form with a cgi-program and a cache
 - 5...60 seconds depending on the time-interval
- The most popular plots available early in the morning again



Routing Vector Database

- Performance did not scale
 - Completely re-write
 - Out-sourced to the NCC S/W Group
- Importing 1 day of data:
 - Old: 5 hours
 - New: 10 minutes
- Queries faster:
 - Web form faster (seconds instead of minutes)
 - Analysis using routing information faster
- To be switched on after RIPE-36




Network alarms

- Running smoothly
- Updated statistics
- Looked for noisy channels
- Further development will depend on the feedback that we get from you




Number of alarms per day

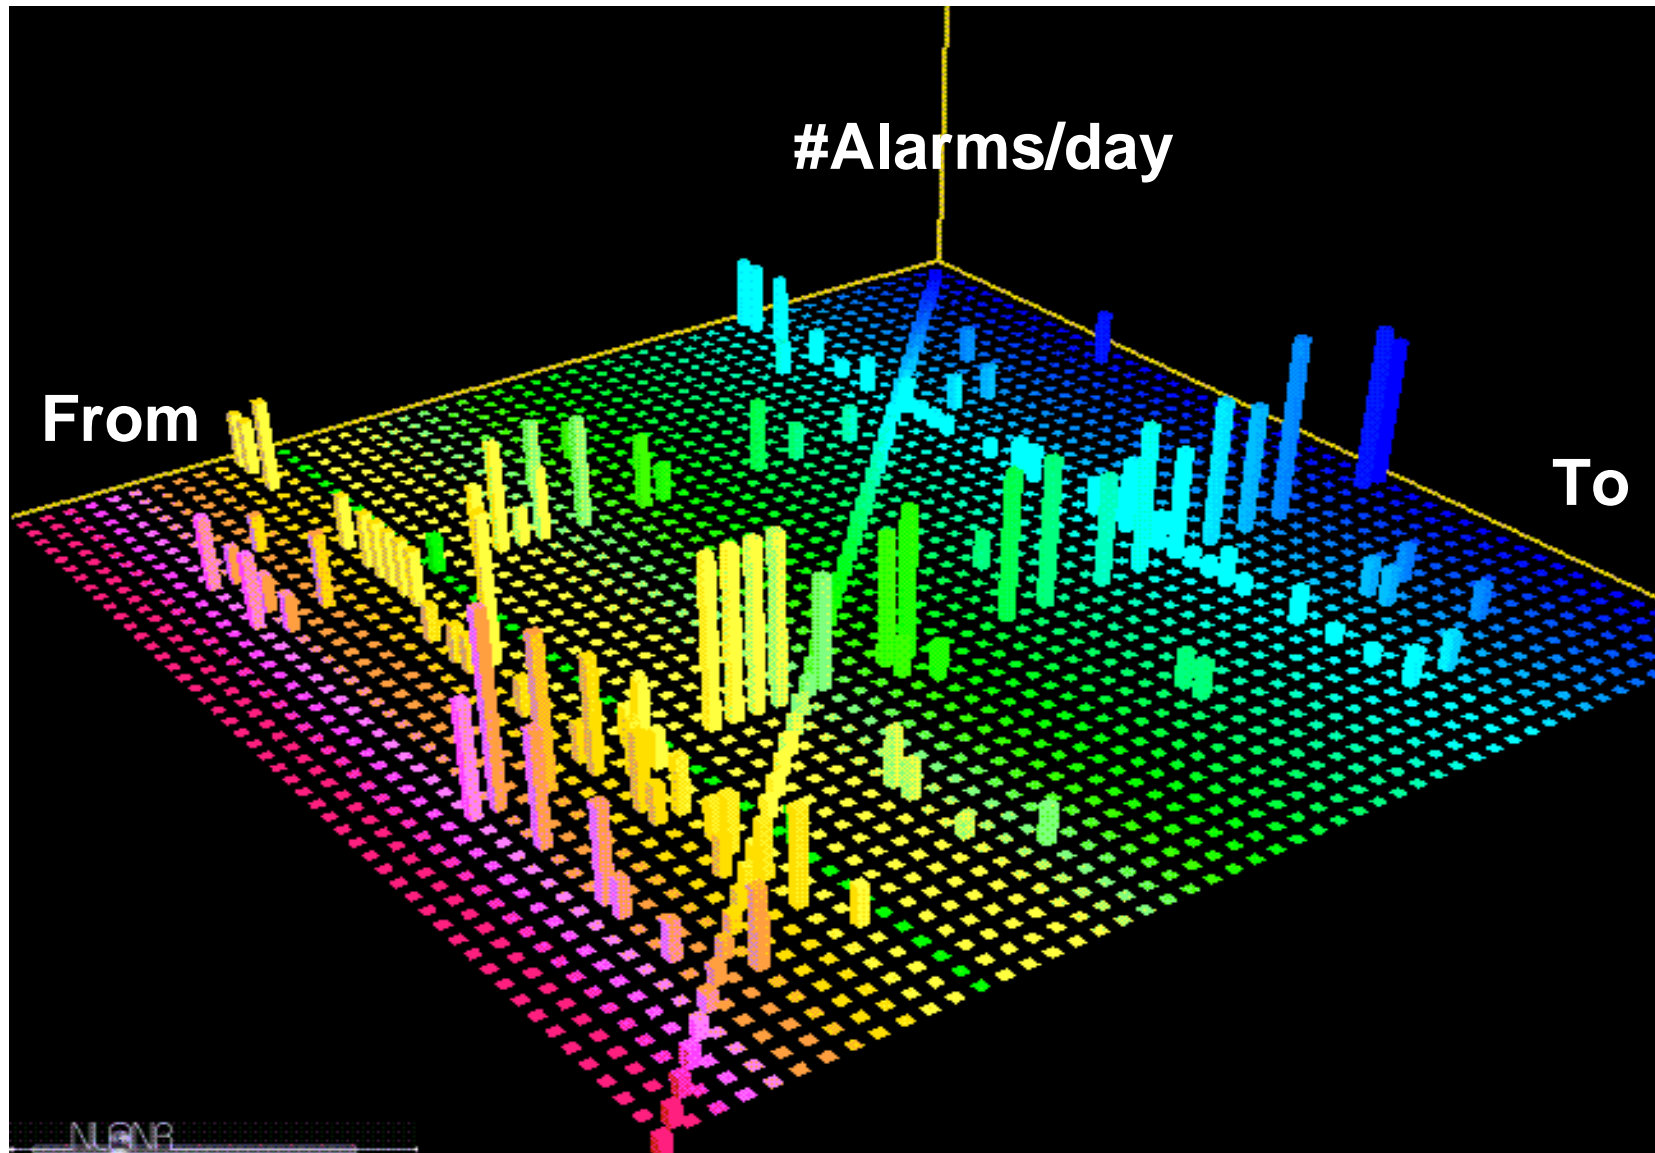




Number of alarms per host



Noisy channels?





Noisy channels?

- Alarms tend to be in a few paths on any given day
- However, there are no paths that consistently produce more alarms than others




Trends in the data

- How do delays develop over longer periods?
- Hard to look at all individual measurements
- Summarize in percentiles:
 - 2.5%: best case
 - 50%: normal case
 - 97.5%: worst case
- Handful of numbers each month
- Some interesting examples on the next slides



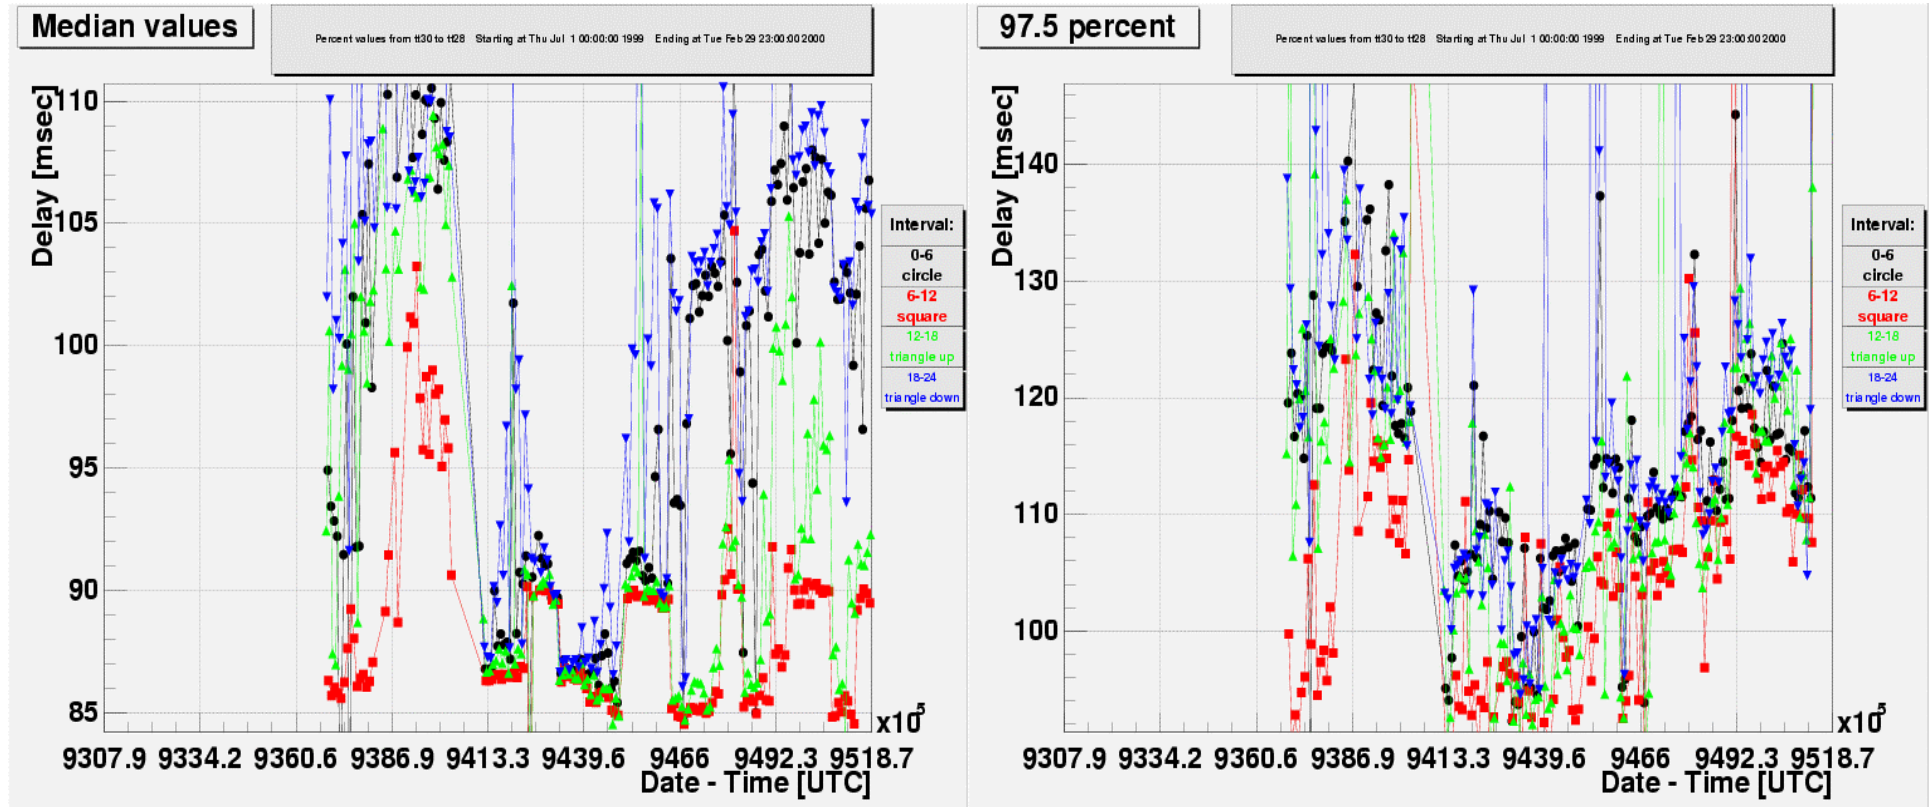
A connection that is getting better



- July 1, 1999-Feb 20, 2000.
- Minimum delay reduced
- Less instances of a saturated link




A connection that is getting worse




- Connection improved at first, then
- Slowly increasing again

Satellite or ground link?



- 350 ms: Satellite link 50, 150 ms: Fiber

Day/Night effect





Trends in the data/Next steps

- Finish research and produce M.Sc.Thesis
- Turn experimental code into production software
- Make results available on the web:
 - Plots and raw data
 - List of plots that deserve human attention
- Correlate with input from NOC's
 - First attempt by Johann
 - How to do this in a more structured way?



Future analysis topics

- Network Performance Scores
 - Postpone discussion until we have first results
 - RIPE37
- Delay variations
 - Measure of short-term jitter on delays
 - Important for applications where packet should arrive at regular intervals
 - IPPM ID seems to have converged
 - On our list of projects for a student



Future analysis topics (2)

- Throughput
 - Treno: abandoned by the author
 - Pathchar-like approach: maximum throughput
 - IPPM does not seem to agree on the method to measure throughput
- Relation delays and traceroutes
 - Modeling of data
 - UMEEPI proposal from Delft University
- Discussion at RIPE36



Overview

- Manpower, Email
- Status of the measurement network
- Turning TTM into a regular service
- The next series of test-boxes
- Analysis and results
- Plans for the next months



Summary of plans for the remainder of 2000

- Continue to move TTM from a experimental project to a regular service
- Getting ready to roll out the next series of test-boxes
- Improving and adding products based on the TTM data

Questions, Discussion

