



On the future of the TTM service

Andrei Robachevsky
RIPE NCC

RIPE52, 1 May 2006



Introduction

- TTM
 - Initial design done early 1997
 - Became a service in 2000
- 9 years later, time for a critical look
 - What is good?
 - What is not so good?
 - How can we improve



Assumptions when designing the system

- Delay and loss occurs on loaded links
 - Still true, property of the underlying layers
- Network capacity is an expensive resource
 - 1997: Buy as little as possible, wait for delivery
 - 2006: Prices have dropped dramatically, readily available
- Subscribers line is slow
 - 1997: 28k modem was state of the art
 - 2006: ADSL and Cable are cheap and standard

Assumptions when designing the system (2)

- Measurements of applications are not interesting as the link to the user is the bottleneck
 - 1997: Last mile was slow
 - 2006: No longer true
- Internet is sparsely connected, routing has not been optimized
 - 1997: We see traffic going across the Atlantic between two sites a few km apart. Delays indicate non-optimal routing
 - 2006: Well connected Internet



Assumptions when designing the system (3)

- Real Time Traffic Engineering based on performance measurements cannot be done
 - 1997: It is acceptable to measure and have numbers tomorrow
 - 2006: Requirement for measurements with results on the spot
- IPv6
 - 1997: Only experimental
 - 2006: Production
- Stratum 1 NTP server is a useful thing to have
 - Still true

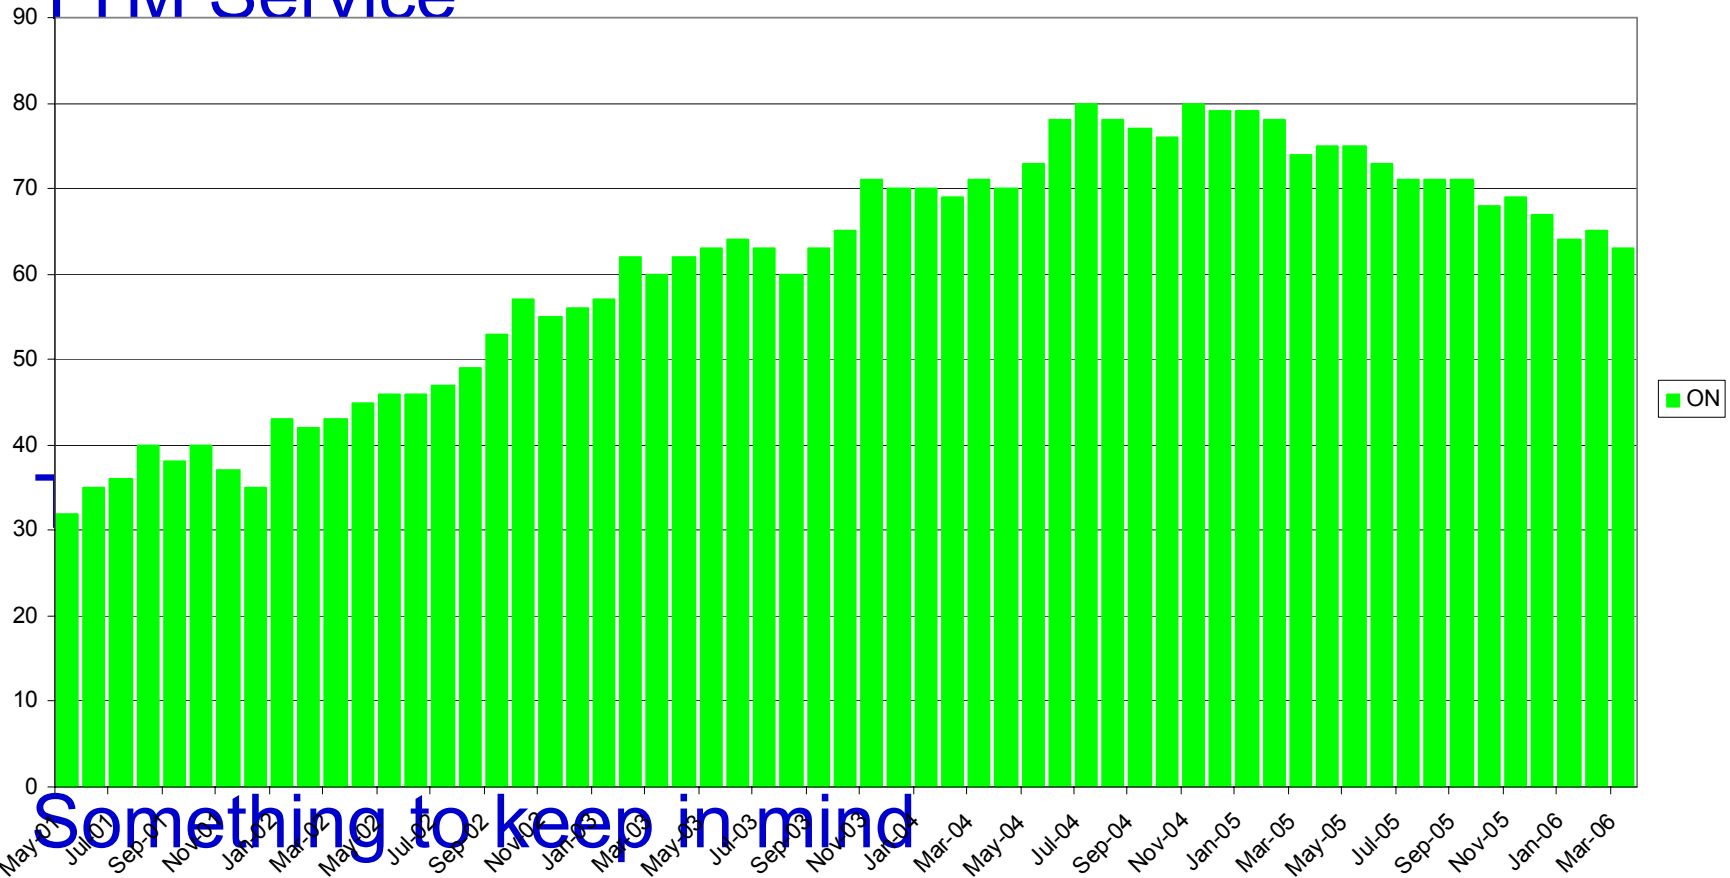
Deployment and maintenance issues

- Installation of the GPS antenna
 - CAT5 help, but still complicates installation
 - No support for other clock sources
- No replacement strategy for boxes
 - The older the hardware, the more problems
 - We cannot force people to upgrade their equipment
- Losing contact with the boxes
 - People leave companies, lose interest
 - No remote hands, hard to debug some problems

Reality Check

- TTM Service

ON



- Something to keep in mind

- Developing a new service takes time, we should build something that is useful for the next 2-5 years

How can we improve?

Simplify and Refocus:

- From network to application
 - More composite measurement, but also more relevant to the end user
- On real time monitoring
 - Small data set that is of immediate use
- From long-term stable observations to ad-hoc measurements
 - Only measure on demand
 - Platform for one time measurements



Some rough ideas

- Measure applications
 - Examples: DNSMON, protocol beacons
- Global one-way measurement platform
 - Measurements on demand
 - Data immediately available for the user
 - Example: OWAMP server network
- General purpose platform
 - Proposal, experiment, analyze, done
 - Example: K-root anycast studies, multicast

Possible architectural model

- Full control of the infrastructure by the RIPE NCC
 - Replacement strategy
 - Better control of locations with TB
 - Similar to k-anycast network
- Multi-tier architecture
 - Simple collectors and more complex aggregators
- TSC or NTP instead of GPS
 - Accept some degradation in accuracy but lower cost and easier deployment

How do existing proposals fit in

- 2005-10/CBM
 - Uses the infrastructure
- 2005-11/Multicast
 - Something to run on demand
- OWAMP
 - Measurement on demand
 - Initiator gets the data immediately after the measurement
 - Can analyze the data and act
- TSC: light weight “TB”.

How to proceed?

- Should we change?
- Is this presentation a good starting point?
- How to develop an elaborated proposal?
- What is the timeframe?

Questions?