The end-to-end network performance troubleshooting cookbook

Vincenzo Capone, Richard Hughes-Jones
GÉANT
Email: vincenzo.capone@geant.org, richard.hughes-jones@geant.org

Introduction

The growth in size and geographical distribution of scientific collaborations, while enabling researcher to achieve always higher and bolder results, also poses new technological challenges, one of these being the additional efforts to analyse and troubleshoot network flows that travel for thousands of miles, traversing a number of different network domains. While the day-to-day multi-domain monitoring, fault detection and handling procedures are firmly established and agreed on by the network operators in the R&E community, a cleverer end-to-end traffic analysis and troubleshooting is still something users are in need of, since the network providers not always have specific tools in place aimed to deal with this category of problems.

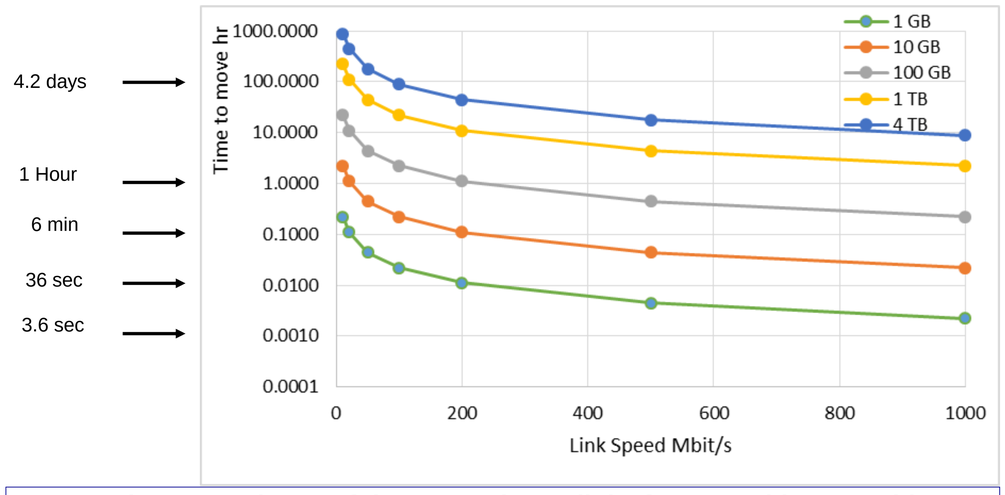

1 Gbyte to 4 Tbytes of data moved over links from 10 Mbit to 10 Gbit

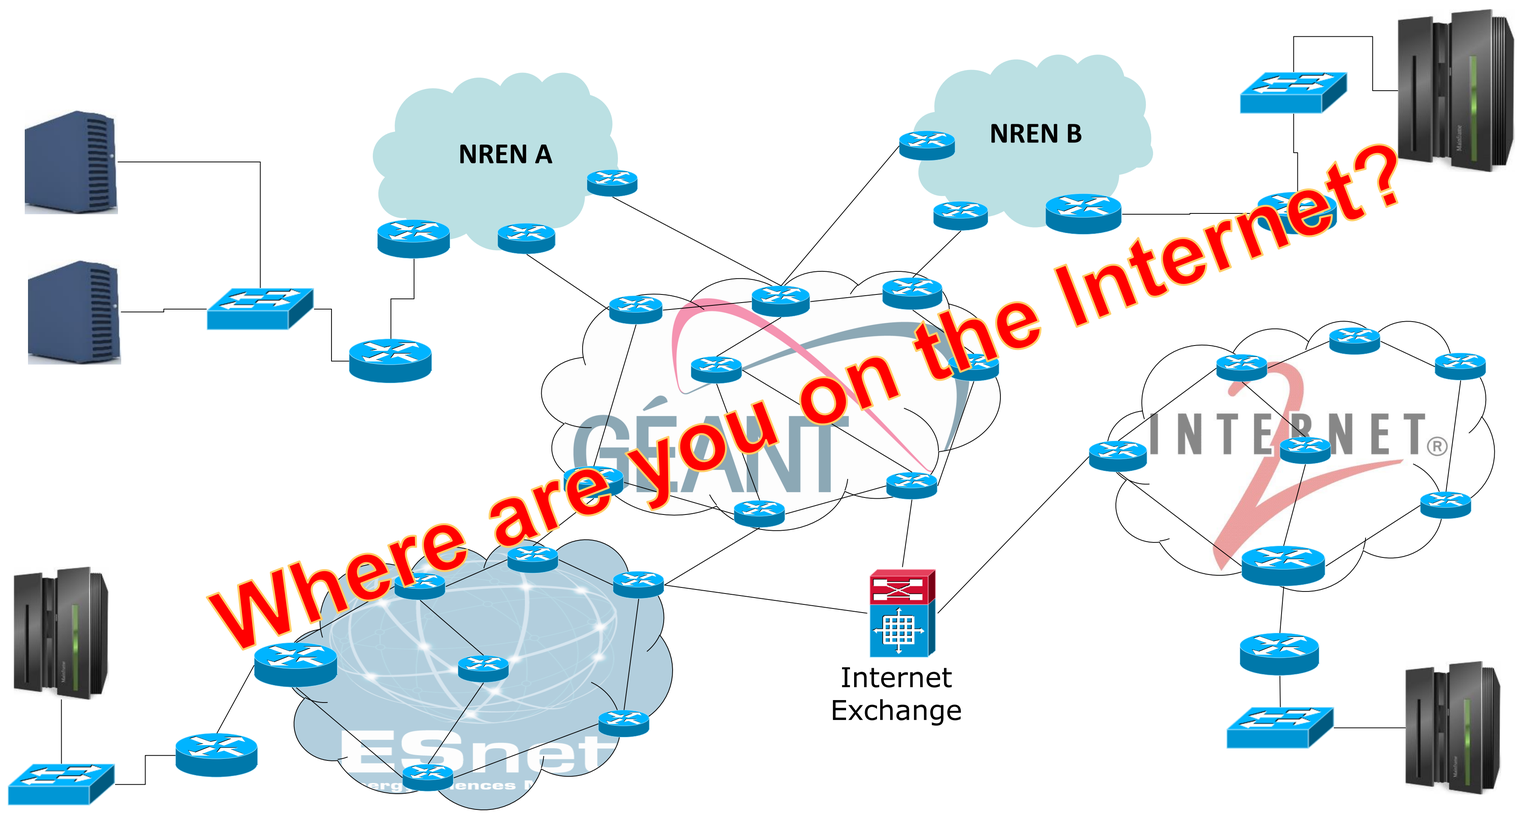

Pre-requisite

Before starting to dig into the test, we need some measurement tool in place. There are various choices, that mainly depend on the resources available (money, skills, manpower). We can identify two groups:
The "**bare minimum**", which means that anything below this will not be really useful to properly assess the performances

- network monitoring measuring traffic usage and error counts (e.g. MoultiRouterTrafficGrapher, *CACTI)*
- end-host I/O performance (measure the purely storage performance - e.g. disk to memory)

The "**nice to have**", being some additions that can give more insight:

- NetFlow-based network monitoring (not an alternative to CACTI, but a complement – tracks the data flows in/out the network)
- Application-embedded benchmarking/monitoring

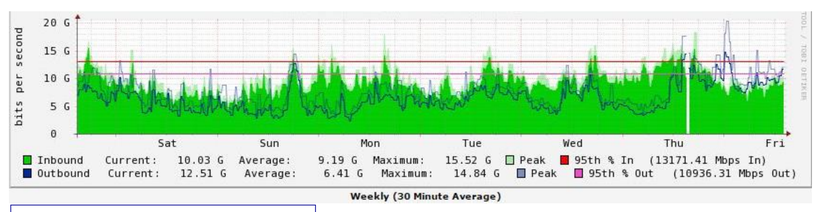

CACTI graph example

Testing scope

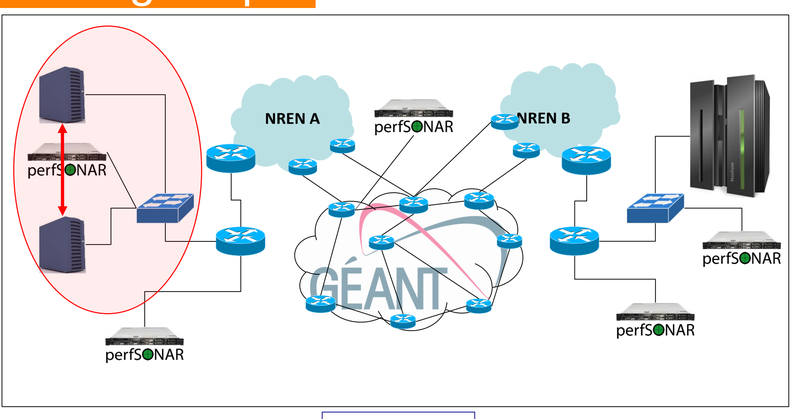

Intra-LAN

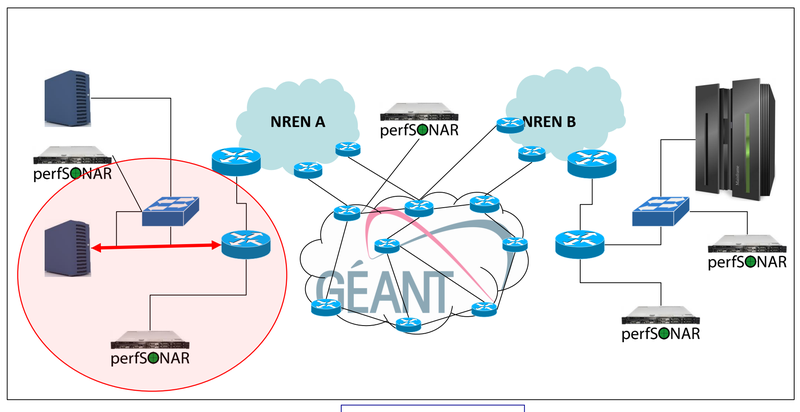

LAN to edge

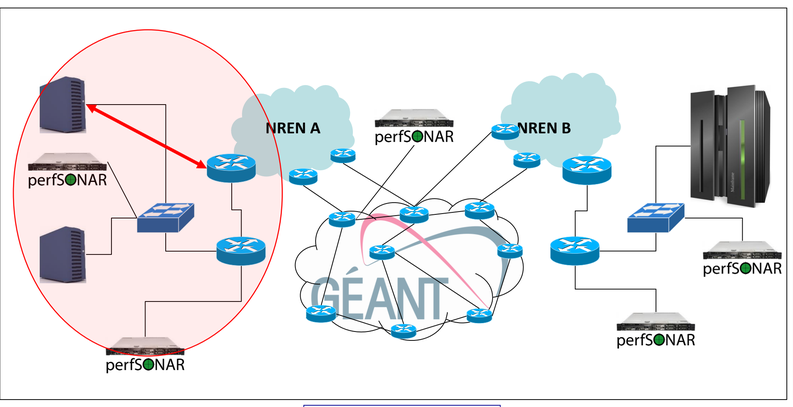

LAN to NREN

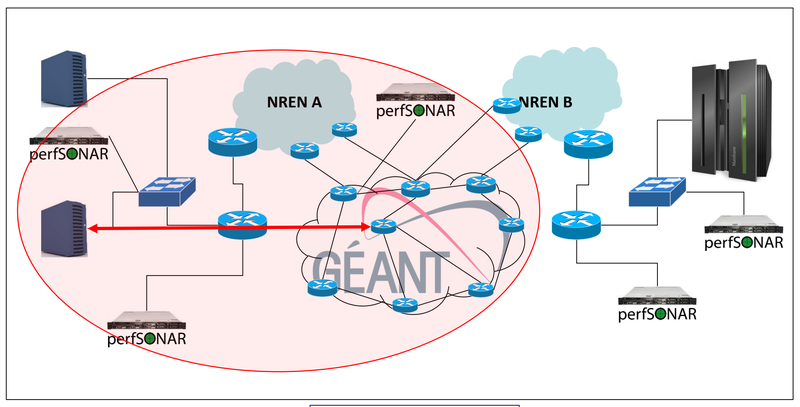

LAN to GEANT

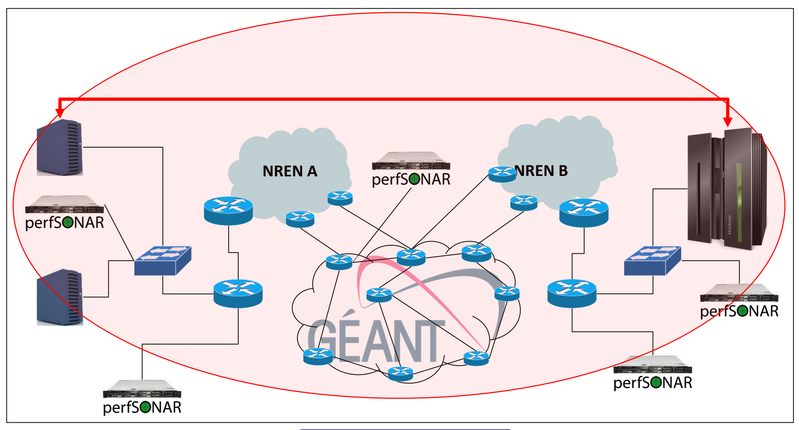

end-to-end

Testing methods

- perfSONAR (the "Swiss knife" of the network manager)
  - Link *quality*
    - *Packet loss, latency, jitter*
  - Routing symmetry and stability
- Memory to memory transfers
  - CPU/NIC-bound
    - *Doesn't involve the storage sub-system*
  - Maximum possible throughput (from **that** specific server)
  - IPERF, UDPmon, etc.
- Disk to disk transfer (end to end, not local)
  - Storage area configuration and capabilities
  - Disk controllers/buffers/RAIDs
- Data application transfers
  - Science-specific
  - Analysis of the data distribution model and tools
  - Authentication and Authorization time-issues

Low level → High level

Test matrix

What you test/measure	Intra-LAN	LAN to edge	LAN to upstream NREN	LAN to GÉANT	end2end
perfSONAR	• Local switching fabric • packet loss • max achievable throughput	• Routing/switching equipment • packet loss • max achievable throughput	• Routing/firewall equipment • Access link • packet loss • max achievable throughput	• NREN/GÉANT connection • packet loss • max achievable throughput	• Full network path • Routing symmetry and stability • packet loss • max achievable throughput
Memory to memory	• NIC drivers • Local switching fabric • max achievable throughput	• TCP/IP settings • switched topology • max achievable throughput	• Routing/firewall equipment • Access link • max achievable throughput	• NREN/GÉANT connection • max achievable throughput	• Full network path • max achievable throughput
Disk to disk	• Storage infrastr. (r/w) • Max storage transfer rate	n/a	n/a	n/a	• Storage infrastr. (r/w) • Max storage transfer rate
Application	• Application settings • Max storage transfer rate	n/a	n/a	n/a	• Application settings • Max storage transfer rate

Conclusions

To some extent, end-to-end network performance measurement has not yet reached the maturity of a fully automatable process, and still requires a lot of manual work. The imfornation presented in this poster are not expected to be fully exhaustive of the topic, instead try to offer that kind of "at-a-glance" view that can help the site administrator to be aware of the issues and challenges, and to efficiently liaise with the network people to deal with long-distance data transfer problems.
The GÉANT eduPERT group (http://services.geant.net/edupert/Pages/Home.aspx) is a good starting point if you want to have more information and hints.