Connect. Communicate. Collaborate

GÉANT2 – an overview of the network and activities

I2 Conference
San Diego, 9 October 2007

Cathrin Stöver, DANTE (cathrin@dante.org.uk)

Connect. Communicate. Collaborate

GÉANT2 Topology

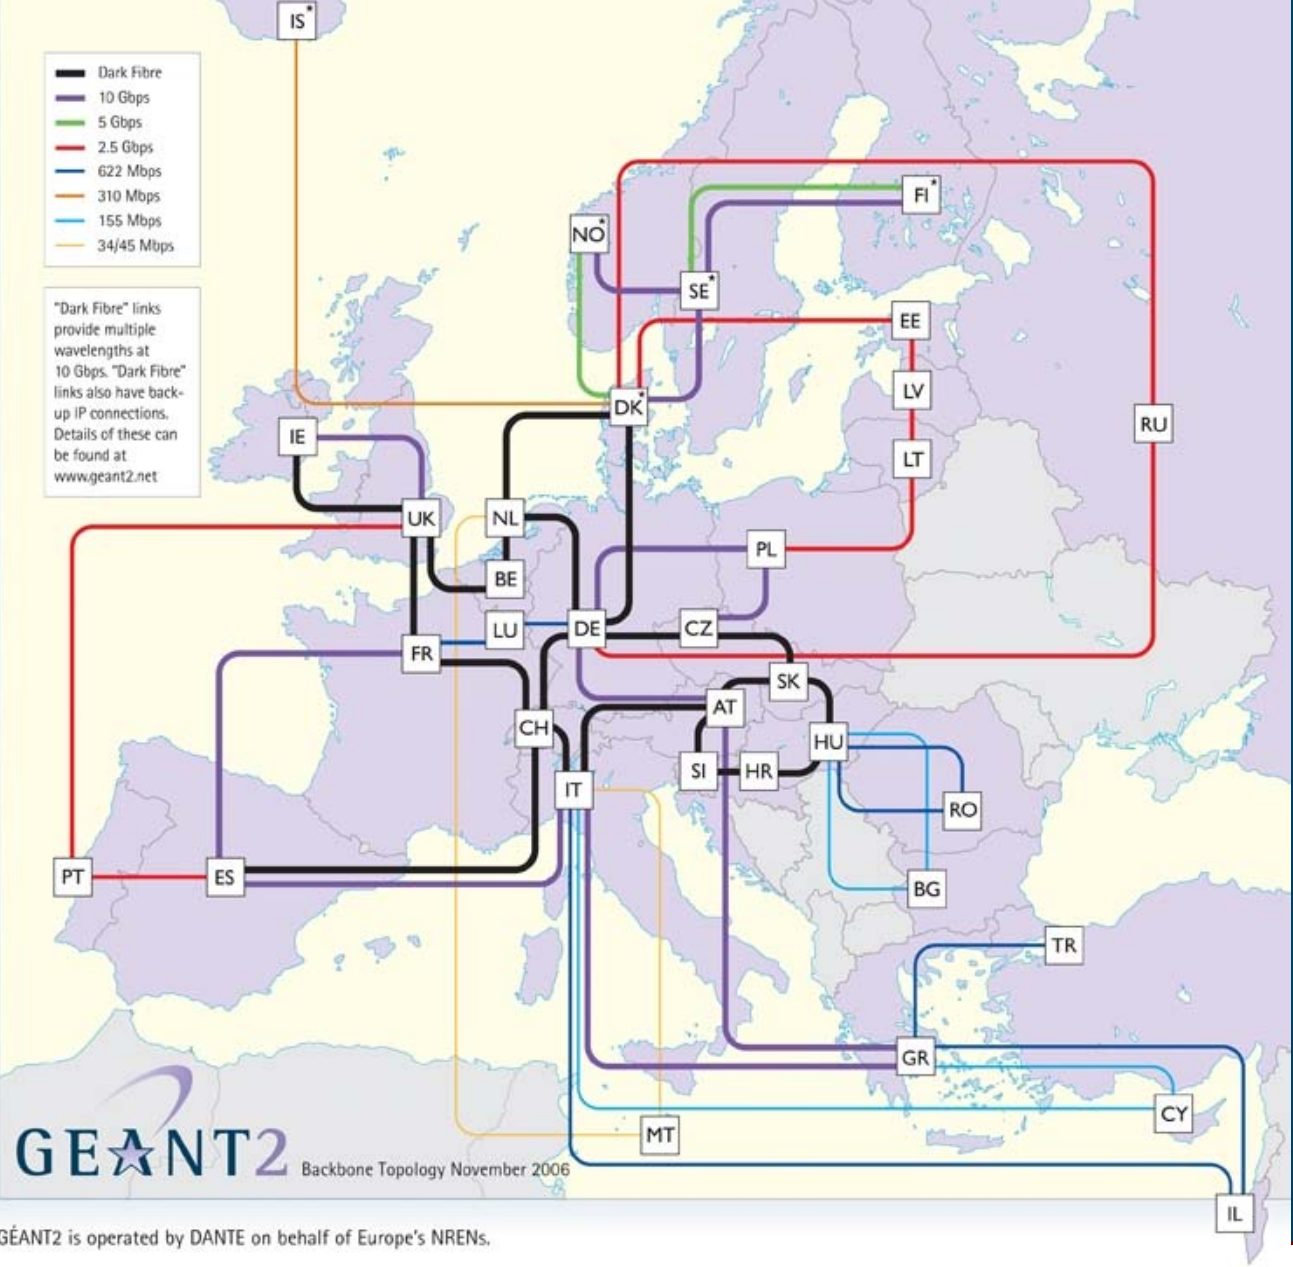

GÉANT2

DANTE

Connect. Communicate. Collaborate

What is GÉANT2?

GÉANT2

Connect. Communicate. Collaborate

- 7th generation of pan-European research network infrastructure
- Funded jointly by European NRENs and European Commission
- Project partners include 30 of Europe's national research and education networks (NRENs), DANTE and TERENA
- Connects 34 European countries and serves over 3500 research and education establishments across Europe
- Service Activities and Joint Research Activities
- Project timescale September 2004 - August 2008
  - Extension to Q2 2009
- Four year project, GEANT3 planned from Q3 2009 to Q2 2013

GÉANT2

DANTE

GÉANT2

Connect. Communicate. Collaborate

GÉANT2

Janet

Cambridge University

RedIRIS

Universidad de Granada

Connect. Communicate. Collaborate

European Research Networks connected via GÉANT2

- GÉANT2 interconnects the European NRENs with each other and globally
- From Russia to Portugal
- From the Nordic Countries to Israel
  - 34 countries across Europe
  - many millions of users

GÉANT2

DANTE

DANTE

Connect. Communicate. Collaborate

- **D**elivery of **A**dvanced **N**etwork **T**echnology to **E**urope
- Established in 1993 and based in Cambridge, UK
- Not for profit Organisation
- Created and Owned by Europe's National Research and Education Networks (NRENs)
- 38 Members of Staff (14 nationalities)

GÉANT2

DANTE

Connect. Communicate. Collaborate

DANTE's Role

- DANTE operates GÉANT2
- DANTE is actively involved in developing new network services to support the European research and education community
- DANTE also leads other projects developing new regional networks across the World:
  - EUMEDCONNECT (Mediterranean)
  - ALICE (Latin America)
  - TEIN2 (Asia-Pacific)

GÉANT2

DANTE

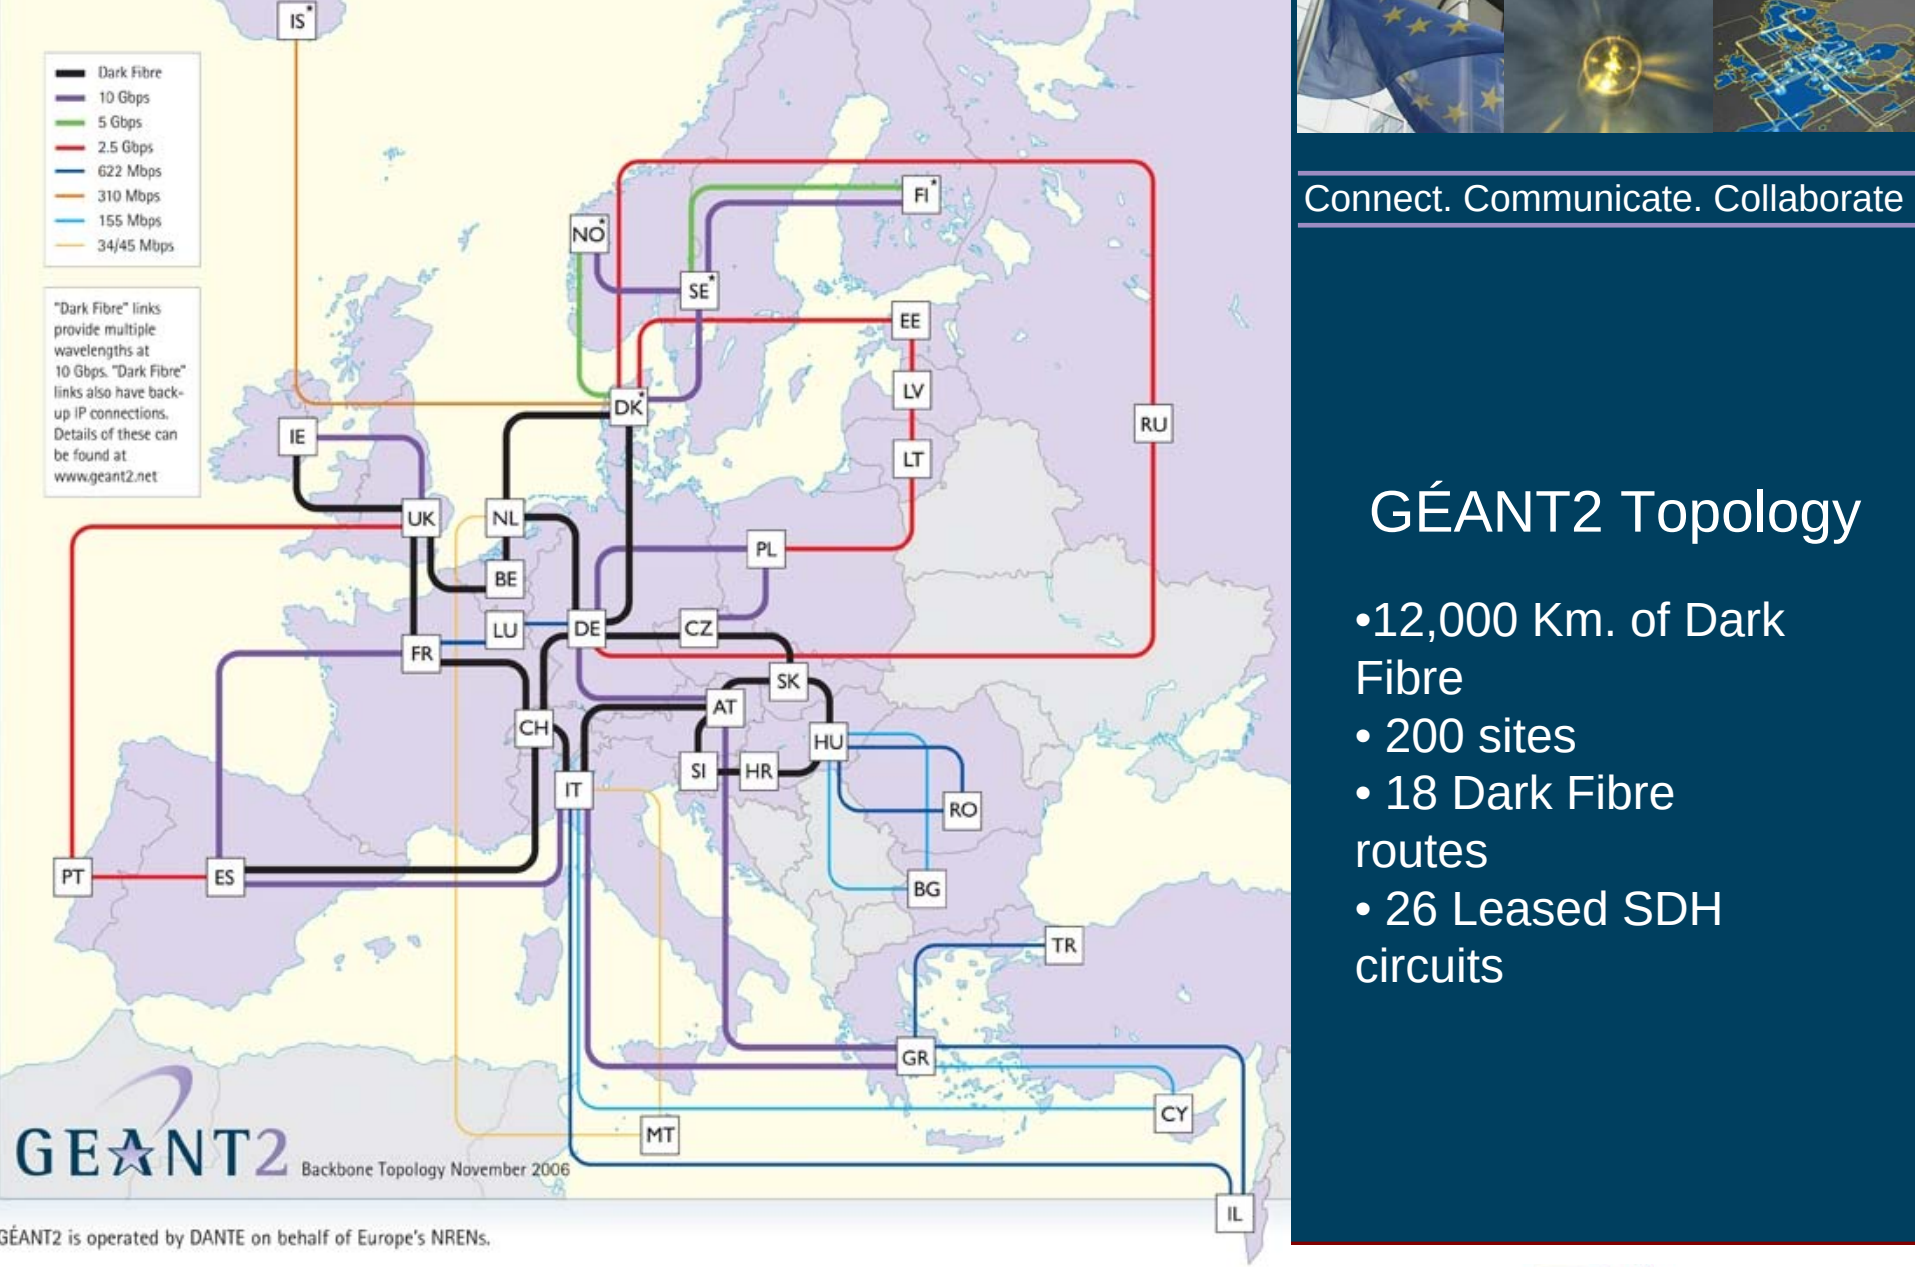

GÉANT2 is operated by DANTE on behalf of Europe's NRENs.

Connect. Communicate. Collaborate

GÉANT2 Topology

- 12,000 Km. of Dark Fibre
- 200 sites
- 18 Dark Fibre routes
- 26 Leased SDH circuits

GEANT2

DANTE

Connect. Communicate. Collaborate

GÉANT2 Objectives

GÉANT2 Objectives

Connect. Communicate. Collaborate

- Provide a gigabit-speeds infrastructure to support European research and education
- Deploy the first international ***hybrid*** network: routed IP traffic combined with switched point-to-point circuits
- Implement end-to-end QoS provision
- Provide a research infrastructure for network technology developments
- Develop a wider range of network services
  - Performance monitoring
  - Security
  - Bandwidth on demand
  - Testbed facility
  - Mobility and roaming

GÉANT2

DANTE

GÉANT2 Objectives (cont'd)

Connect. Communicate. Collaborate

- Provide user support and consultancy
- Benchmark and support development of NRENs
  - Compendium of NRENs
- Coordinate RTD activities
  - Task Forces (TFs)
- Extend geographic reach of the network
- Conduct strategic studies into the future of European research networking (eg. SERENATE, EARNEST)
- Disseminate benefits and achievements of the network

Connect. Communicate. Collaborate

IP and Point to Point services

IP and point to point services in GÉANT2

Connect. Communicate. Collaborate

- All NRENs receive IP access
  - On Fibre group receive standard 10 Gb/s IP access port
  - Off Fibre group access at capacity determined by requirements, technical availability and budget
- All NRENs have possibility of Point to point services as well
  - On Fibre group receive second 10 Gb/s port for GE (typically) point to point services anywhere on fibre cloud
  - Off Fibre group can reserve dedicated path within existing IP access port

IP and point to point services in GÉANT2 (contd.)

Connect. Communicate. Collaborate

- For the On fibre group the 10 Gb IP plus the 10 Gb point to point access is bundled as the minimum access available
  - GEANT2+ service
  - Fixed cost subscription
- Additional (primary and backup) 10 Gb paths across the network are also available for an additional marginal fee
- Internet2 and GEANT2 have agreement in place for international point to point services

GÉANT2

Connect. Communicate. Collaborate

Services and Joint Research Activities

Connect. Communicate. Collaborate

Services and Joint Research Activities

- SA3 End to End QoS
- JRA1 Performance Monitoring
- JRA2 Security
- JRA3 Bandwidth on Demand
- JRA4 Testbed, Cross Border Fiber
- JRA5 Roaming, AAI

From joint research to common services

Connect. Communicate. Collaborate

End to End QoS

- Performance & Allocated Capacity for End users (PACE). The purpose is to help end-users get the network performance they need
- Consists of the following main parts:
    - Performance Enhancement & Response Team (PERT)
    - Multi-domain Premium IP service
    - GÉANT2 Measurement Point network
- Main partners are DANTE, DFN, FCCN, GRNET, PSNC and SWITCH

GÉANT2

DANTE

The PERT

- A virtual team of European NREN network engineers who investigate reports of sub-optimal end-to-end network performance
- Provide a “8x5” service by means of a Duty Case Manager
- Manage open cases with the PSNC-developed PERT Ticket System
- Provide general advice through the publicly available PERT Knowledgebase website

Performance Monitoring

- The Performance Measurement and Management main objective is to build a multi-domain monitoring framework which is the basis to offer a Multi-Domain Monitoring (MDM) Service.
- Consists of the following main parts:
    - Design and develop the framework (perfSONAR).
    - Integrate measurement tools and databases within the perfSONAR framework.
    - Build user visualisation tools using the perfSONAR framework.
- Main partners are DFN, CESNET, CARNet, PSNC, NORDUnet, DANTE
- External collaboration with ESnet, Internet2, RNP, UoD
- The Service offered to the users is a Multi-Domain Monitoring (MDM) Service.
  - Users get access to monitoring data, measurement functionalities and visualisation applications.
  - Authorisation is user type based.
  - Covering Europe targeting to reach end-institution accesses.

Who will make use of PerfSONAR?

Connect. Communicate. Collaborate

- NOCs (NRENs, GÉANT2, ESnet, RNP, etc).
- PERT.
- L2 project users (LHC OPN, DEISA).
- L3 project users (EGEE).
- End-users when appropriate tools will be made available.

- Activity in the process to become a Service to GEANT2 community

Security

- The main purpose is to make the GÉANT2 community as secure as needed

- Consists of the following main parts/strands:
  - Security policies/recommendations
  - Tool development
  - Operational framework
  - Liaison

- Main partners are: CESNET, DANTE, GARR, GRNET, SURFnet, SWITCH
- Collaboration/liaison with TERENA's TF-CSIRT

Bandwidth on Demand

Connect. Communicate. Collaborate

- The main objective is to engineer, automate and streamline the inter-domain setup of dedicated capacity end-to-end paths in a multi-technology environment
- It consists of the following Work Items:
    - Management
    - Requirements Gathering and SoA Surveys
    - Service Specification
    - Implementation
    - Testing & Service Validation
- Main partners are: DANTE, GARR, GRNET, HEAnet, PSNC

Testbed

- The testbed activity consists of three main strands:
  - Design, construct & operate a general purpose standalone testbed
  - Undertake a selected technology testing programme
  - Evaluate technical aspects of using so-called “Cross-Border Fibres” (CBFs)
- Main partners are DFN, CESNET, PSNC, DANTE
- Collaboration/liaison with DANTE ops, JRA1, JRA3

Roaming and Authorisation

Connect. Communicate. Collaborate

- to build a European Roaming Infrastructure based on eduroam
- to pilot and build the federated support for existing Authentication and Authorisation Infrastructures for Research and Education, called eduGAIN
- The combination of the two will allow for access to network and to services with a single login (SSO)
- Advanced technologies will be integrated into these infrastructures where appropriate
- Partners are ARNES, CARNet/Srce, CESNET, Dante, DFN, FCCN, GRNET, HEANET, HUNGARNET, ISTF, NORDUnet (CSC, UNI-C, UNINETT, University of Umea), RedIRIS, RESTENA, SURFnet, SWITCH (different involvement in project parts)

- Collaboration/liaison with
  - many groups: TF-Mobility, TF-EMC2, GN2 activities (JRA1, SA3, JRA3), international groups like eduroam gwg, SALSA FWNA (Internet2), MACE, TF-NGN, DICE, GGF, eConcertation
  - and projects: Akogrimo, EGEE2, Lobster

Connect. Communicate. Collaborate

How does GÉANT2 support demanding users?

GÉANT2 Applications

Connect. Communicate. Collaborate

GÉANT2 has a userbase of 30 million+ researchers spread across 34 countries in Europe….and many more in other world regions….

....they range from primary school children to nuclear physicists….from weather forecasters to cancer specialists….

…..their requirements range from a few Kbps to many Gbps….from email to dedicated circuits….

GÉANT2

DANTE

Connect. Communicate. Collaborate

Information transmitted

Tradicionally :

- Capacity of traffic: Kbits/sec
- Amount of users at the institution: Many
- Applications that are not very sensitive to underlying network (ie. e-mails, FTP, etc.)

Specialised:

- Capacity of traffic : More than 1GBit/sec
- Amount of users at the institution: dozens
- Applications between very specific points and very sensitive to throughput (ie. eVLBI, CERN LHC or DEISA)

Specialised Applications:

- High Energy Particle Physics (HEP)
  - LHC (@ CERN) operational since 2007
  - Needs multiple 10G links between the different nodes for the data processing
  - Real time data processing
- Radio Astronomy (eVLBI)
  - Very Long Baseline Interferometry
  - Flows of 1Gbps in real time (http://www.evlbi.org)
- Supercomputing Centres
  - e.g DEISA project (http://www.deisa.org)
- GRIDS
  - Like EGEE (http://www.eu-egee.org)

A GÉANT2 User Group: eVLBI

Connect. Communicate. Collaborate

- What is eVLBI?
  - A radio-astronomy technique that uses great distances between telescopes to vastly increase their resolution
- Why do they use research networks?
  - To send transfer data from the telescopes to a central correlator whilst the measurement is ongoing. Vastly improves the speed for analysis and allows Window of Opportunity observations

GÉANT2

DANTE

What does GÉANT2 do for eVLBI?

Connect. Communicate. Collaborate

- Monitoring: counters set up and eVLBI weathermap created
- Planning (EVN-NREN forum)
- PERT case –aiming to optimise end-to-end performance
- Point-to-point connections have been implemented
- New connections –ZA, CN, AU, AR–made possible by the global reach of GÉANT2 -with TEIN2, ALICE/CLARA and ORIENT

PERT in Action

- Research Network Support for eVLBI Tracking of the ESA's SMART-1 Spacecraft
  - moon-orbiting spacecraft scheduled to descend to the moon's surface in autumn 2006. To track this descent, telescopes used radio signals to pinpoint the craft.
  - Latin America best positioned to view the event, a telescope in Chile (TIGO) to send data relating to the descent to JIVE (eVLBI, NL) for analysis
  - not a real-time transfer of data and network transfer was intended purely to improve analysis time.
  - Following a request from JIVE, a series of measures were instigated to resolve data-transfer rate issues between TIGO (CL) and JIVE (NL). The effort was coordinated by DANTE and the GÉANT2 PERT and involved input from GÉANT2, RedCLARA, REUNA, SURFnet, TIGO and JIVE engineers.
  - 3 months activity extending beyond the date of descent

PERT In Action - 2

- Initial tests confirm no packet loss in SURFnet, GÉANT2 or REDCLARA networks
- TIGO access to the REUNA backbone was rate-limited to 5Mbps. This was lifted to facilitate the transfers
- Traffic shaping applied and shown to improve transfer rates over REUNA
- A 90Mbps rate-limit on the REUNA-RedCLARA access circuit was raised to the 155Mbps maximum.
- The burst-limit on the access router was found to be set unusually low. This was reset to an appropriate level.
- 28Mbps was achieved between Europe and TIGO using iPerf testing
- Only 2Mbps was achieved on the real transfers of the SMART-1 data

PERT In Action - 3

The PERT trouble ticket was held open beyond the date of the SMART-1 descent to investigate the causes of the poor transfer rate and to facilitate future transfers TIGO-JIVE:

- Tests centred on application optimisation reveal that significant data throughput improvements could be made (to 15Mbps) from TIGO to JIVE if SACK (Selective Acknowledgement) was implemented on the JIVE PC.
- Similarly, improvements could be made if the TIGO kernel was updated.
- Minor local packet loss appears to be a consistent feature of transfers, the most likely cause being congestion, and it is expected that the imminent REUNA upgrade will significantly improve matters.
- More significant packet loss has also been noted sporadically and may not be improved by the upgrade. The cause of this is as yet unknown, but may be due to a router/firewall or frequent traffic rerouting.

SMART-1 and PERT

- Proof of concept and soundness of PERT procedures
- Proof of value of an expert team to address performance issues
- Importance of dedicated case manager in PERT
- Reinforcement of collaboration between radio astronomers and network engineers
- Starting point to achieve a stated objective of the EXPReS project: to connect TIGO to Europe for real-time eVLBI sessions
- Requirements to connect sites over long network paths must be investigated as a complete solution
- Recommendations from PERT to eVLBI community

Connect. Communicate. Collaborate

E2E Coordination Unit

- Triggered by requests coming from end users in Europe
- Active monitoring of an Inter Domain OPN (Optical Private Network) and management of the information flow related to the status of an end to end connection
- For doing so
  - E2ECU will gather network information from the Inter domain monitoring tools provided or directly from the entities involved.
  - E2ECU will periodically relay network information to all the relevant parties

E2E Coordination Unit Set Up

- Entity appointed for the role
- Process Definition
  - Fault Report and Service restoration
  - Hours of Coverage
  - Escalation Procedures
  - Periodic Reports
- Tools to develop
  - Monitoring Tools
  - Trouble Tickets System
  - Data Base

Process

- Fault Reports or Maintenances specifications
  - Via a web interface in the GEANT2 site
  - Via a dedicated phone number
  - Via e-mail address (e2ecu@noc.geant2.net)

- Communication back to the entities involved
  - Via phone or e-mail in case of queries
  - E-mailed Trouble Tickets for relaying updated information

Process II

- Coverage
  - From 8.00 – 22.00 CET Monday to Friday
  - From 9.00 – 18.00 CET Saturday to Sunday

- Escalation Procedures: From the entities involved in a project to the E2ECU and vice versa!

- Monthly Reports to be provided describing the e2e links availability and tickets opened for each specific project

Monitoring Tools

- Each network domain shall use PerfSonar infrastructure to store network status information
- Applications have been developed to access that information and
  - Show the circuit's status => Visualization tool
  - Integrate it with more general applications used by the E2ECU (ie. NAGIOS):
    - SNMP traps ready
    - Dashboards showing alarms that the help desk engineer will acknowledge
    - …..

Extended Features

- Trouble Ticket Systems
  - Able to send e-mails to specific community of users depending on the fault's impact
- Data Base
  - Containing information relating: Links / Projects / Contact information of the network administrators

Time Scales

- Processes under development
- Pilot E2ECU service by Beginning of December
  - Focused in LHC support
  - Will test the processes and tools in place
- Full set up by the end of January 2007

Connect. Communicate. Collaborate

Thank You!

cathrin@dante.org.uk

Information Material available at:

http://www.dante.net/server/show/nav.126

GÉANT2

DANTE