

DTN – Data Transfer Node

Xavier JEANNIN, RENATER
Damir REGVART, CARNET

STF meeting, 22-23 Oct. 2019,
Copenhagen

Public

www.geant.org

GÉANT 4-3 WP6 Task1: Network Technology Evolution

- Data Plane Programming
- **Data Transfer Node Stimulation**
- LoLa
- Optical Time & Frequency Network
- Quantum Key Distribution
- White-Box for research and education
- Router for Academia, Research and Education (RARE)

- By its nature, this task assesses the interest in different topics, and during GN4-3 some of these could be stopped and new ones, if proposed or identified, and if there is evidence of end users, can be started

Data Transfer Node Stimulation in WP6 Task1

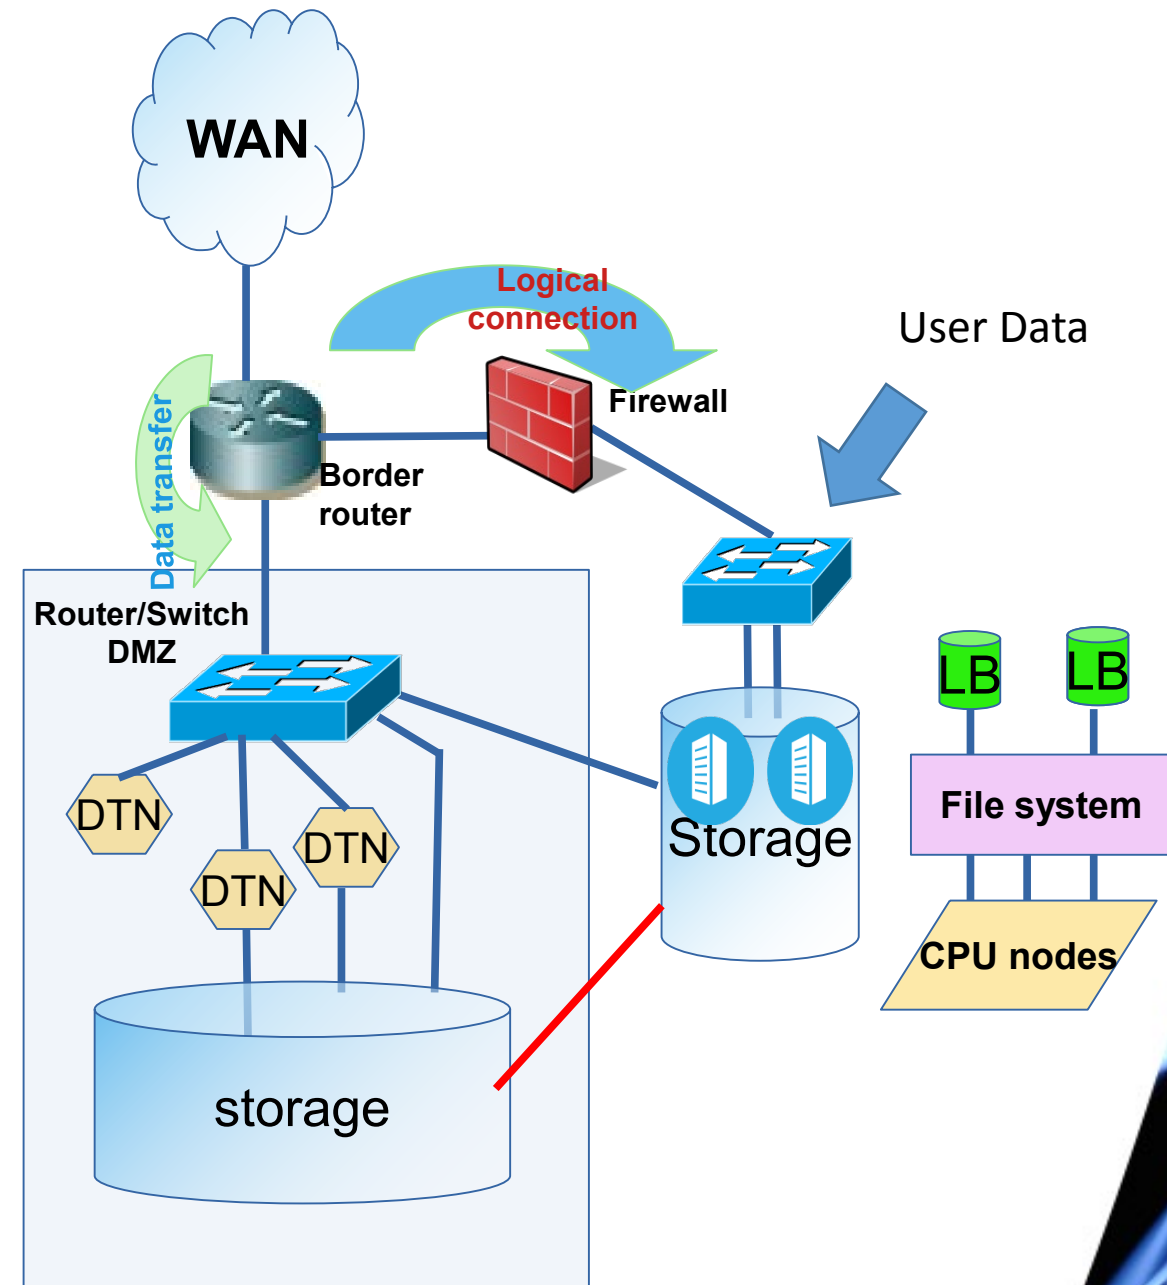
- Milestone 6.4, review of DTN Activities in the Community
- Initially there was not a strong engagement from the NRENs
 - There was a proposal for work put forward in the project preparation phase
 - But there were very few applications from NRENs to do the work, including the NREN that proposed it
- Data Infrastructure survey, Oct. 2019 in order to identify NREN interest in supporting large scale data transfers, especially using a DTN solution

Why large scale data transfers matter for NRENs?

- It is a very common use case for the NRENs, supporting data-intensive science
 - Short distance transfer
 - Long distance transfer
 - Very sensitive to packet loss
 - Even a fraction of 1% packet loss leads to poor performance
 - Congestion / provisioning
 - Optimize network (Optimize network path, MTU, ...)
 - Deploy Science DMZ at the end sites / campuses - <https://fasterdata.es.net/science-dmz/>
- Other providers have the same concern, e.g., Google created an algorithm “Bottleneck Bandwidth and Round-trip time (BBR)”, which is far less sensitive to packet loss
- Established data-intensive transfer users (mainly the CERN community and WLCG) have already addressed this concern (File Transfer Service, and the same principles as Science DMZ).
- Long tail science, and emerging science disciplines, are subject to a poor experience

What is a DTN?

- Mission: Move (BIG) data fast
- DTNs are custom made servers (Linux), optimized for data transfer
- Maximize utilization of network capacity
- DTNs sit at campus edge (Science DMZ model); outside campus firewall
- Day-to-day campus traffic handled differently to “big science” traffic
- Forward large amount of data **from** and **to** the site resources (fast and efficient transport)
- Now supports 100 Gbps flows, tomorrow 1 Tbps



DTN explained

- Speaking about DTN you can find multiple definitions
 - Hardware – physical machine and its specification
 - Software – the transfer tools used
 - Orchestration – managing the flows, possibly configuring bandwidth on demand
- All functions must be « *tuned* »
 - CPU is dedicated for data transfer
 - File transfer (large buffers/chunks of data), tuning TCP parameters
 - Logical (sequential) order in sending data



Example of DTN server (Caltech)

Design options for High Throughput DTN Server



1U SuperMicro Server (Single CPU)
Single 40/100GE NIC
Dual NVME Storage (LIQID 3.2TB each)
~90 Gbps disk I/O using NVME over Fabrics

2U SuperMicro Server (Dual CPU)
Single 40/100GE NIC
Three NVME Storage (LIQID 3.2TB each)
~100 Gbps disk I/O using FDT/NVME over Fabrics

2U SuperMicro (Dual CPU)
Single/Dual 40/100GE NICs
24 NVME front loaded 2.5" drives (U.2)
~200Gbps of disk I/O using FDT/NVME over Fabrics



DTN implementation challenges

- Performance
 - Troubleshooting
 - Tuning
- Monitoring
 - perfSONAR – network measurements
 - Application logs – disk to disk measurements
- Authentication
- Authorisation
- Important that the tools are easy-to-use for end users
 - Needs to be intuitive, not require specialist knowledge
 - Software development

Examples of data transfer projects

- AENEAS project (<https://www.aeneas2020.eu/>)
 - Federated European Science Data Center (ESDC) to support the astronomical community in achieving the scientific goals of the Square Kilometer Array (SKA)
- Process (<https://www.process-project.eu/>)
 - Creating data applications for collaborative research (Exascale learning on medical image data, SKA/LOFAR, Ancillary pricing for airline revenue management, Agricultural Analysis based on Copernicus data)
- Big Data Express (<https://bigdataexpress.fnal.gov/>)
- Pacific Research Pplatform (<https://prp.ucsd.edu/>); Petascale DTN Project
 - Uses FIONA devices; SSD disk with GPU compute built in. Effort to improve data transfer performance between the DOE ASCR HPC facilities at ANL, LBNL, ORNL, NCSA
 - 4.4TB of cosmology simulation data
- Others ...

Existing data transfer tools

- Amazon Web Services - Amazon S3 or Amazon Simple Storage Service (<https://aws.amazon.com/s3>)
 - Object storage service
 - API control/management, batch operations
- Globus Online (<https://www.globus.org>)
 - Non-profit group within the University of Chicago
 - SaaS (Software as a Service) principle. It allows users to move data between two GridFTP servers or between a GridFTP server and a user's machine
 - "DTN Orchestrator"
 - User has "Agent application" that controls and monitors flow of data
- Others...

Survey results 1/4

- 29 NREN responses (from the APM contacts)
- User groups that move data intensively thanks to the NRENs/GÉANT networks:
 - Physics : (HEP) LHCONE, Astro-physics, LOFAR | HPC: PRACE | Astronomy | Biology: human brain, ELIXIR ... | Environment and climate research: CMIP6, Copernicus
 - Some are more established than others
- NRENs identified well-known big projects
- Some NRENs see their role as transport capacity providers



Survey results 2/4

- Success stories:
 - Same well-known projects: LHCONE (CMS, ATLAS), LOFAR, PRACE, Astro-physics, HPC, Astronomy, Biology ...
 - Tuning the campus/LAN (Science DMZ)
 - GÉANT and AARnet - demonstration that NREN networks provide better performance than commercial networks for long haul data transfers
- Issues
 - Network
 - Long distance transfers, firewalls, last mile network, bandwidth capacity
 - Poor network performance and difficulty to troubleshoot
 - Tune campus, LAN and local systems
 - Difficult to implement Science DMZ and security – need to apply security policy differently for day-to-day and “big science” traffic
 - IAAS usage without coordination of NREN
 - Low expectation – so see transport of large volumes of data transfer using hard drives!

Survey results 3/4

- Performance tools
 - perfSONAR, Netflow, Iperf
- Ways in which NRENs can help users move data
 - Traffic engineering (dedicated circuits, VRFs, dynamic provisioning, etc)
 - FileSender service
- Other forms of assistance
 - Support, advice, collaboration
 - Promoting Science DMZ – 3 NRENs
- Large scale data transfers with international organisations
 - Same well-known projects



Survey results 4/4

- Would you be willing to work with us (GN4-3 WP6) on improving data transfer infrastructures for the community? - 11 = Yes
 - The big question is – doing what? What collaborative work would be most beneficial?
- Other comments
 - The problem is often due to local network capacity, the last mile
 - Need more documentation of best practices



Conclusions

- There is some evidence that NRENs already work with their users, but that this can also be improved
- How can GTN4-3 WP6 T1 help NRENs?
 - **What would be the best things to focus on for our work in Year 2 of the project? (Jan – Dec 2020)**
 - Perhaps maintain a wiki of best practices, data transfer projects, DTN configuration examples?
- From WP6 T3 we have tools to detect and localise performance issues
 - Provide training on performance troubleshooting using those tools?
- Is there any software development that would add value?

- Your thoughts?

Thank you

Any questions?



© GÉANT Association on behalf of the GN4 Phase 3 project (GN4-3).
The research leading to these results has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 856726 (GN4-3).