SANReN
South African National Research Network

Performance Enhancement Response Team and Big Data Transfer Service Proof of concept

Roderick Mooi – Principal Engineer, SANReN

Kasandra Pillay – Senior Engineer, SANReN

pert@sanren.ac.za

science & technology
Department:
Science and Technology
REPUBLIC OF SOUTH AFRICA

CSIR
our future through science

SANReN
South African National Research Network

PERT

Performance Enhancement Response Team

Team supporting resolution of end-to-end performance problems for networked applications (GÉANT)

"Even with high-quality infrastructure, end users may not always get the performance they might expect from the network" – TERENA compendium 2008

science & technology
Department:
Science and Technology
REPUBLIC OF SOUTH AFRICA

CSIR
our future through science

What is PERT?

Source: GÉANT eduPERT

As networks and networked applications grow and develop the demands they make on performance and reliability grow ever greater. Ultra-high speed data applications, real-time cloud based systems, high definition live video transmission, and remote computing applications all increase the demands on the network and the expectations of the users.

With all these requirements, there is an increasing need for skilled experts to be able to respond rapidly to issues and to use their experience, expertise and knowledge to help solve users' problems quickly and efficiently.

Performance Enhancement Response Teams (PERTs) provide an investigation and consulting service to academic and research users on their network performance issues. If they have an issue (for example low throughput, slow response times, or degradation), typically the user will contact a PERT directly or through standard IT support channels. To diagnose problems a PERT considers the complete end-to-end path, so one of the main functions of a PERT is to question end-users, system administrators, and network operators who have systems on that path.

SANReN
South African National Research Network

Overview

Architectures and tools for optimising big data transfers especially for science and research – how fast can you go?

- Science DMZ network architecture
- Data transfer nodes and tools
- perfSONAR monitoring toolkit
- What is SANReN doing?

SANReN
South African National Research Network

Science DMZ

"A Network Design Pattern for Data-Intensive Science"

- trademark of the Energy Sciences Network (ESnet – USA)
- built at or near lab or campus network perimeter
- optimised for high-performance scientific applications
- not general purpose / everyday business computing
- addresses common network performance problems

following slides used with permission…

(http://fasterdata.es.net/science-dmz/science-dmz-community-presentation/)

Motivation

- Networks are an essential part of data-intensive science
  - Connect data sources to data analysis
  - Connect collaborators to each other
  - Enable machine-consumable interfaces to data and analysis resources (e.g. portals), automation, scale
- Performance is critical
  - Exponential data growth
  - Constant human factors
  - Data movement and data analysis must keep up
- Effective use of wide area (long-haul) networks by scientists has historically been difficult

"but we've always shipped disks!"

ESnet

Data Mobility in a Given Time Interval

Data set size	1 Minute	5 Minutes	20 Minutes	1 Hour
10PB	1,333.33 Tbps	266.67 Tbps	66.67 Tbps	22.22 Tbps
1PB	133.33 Tbps	26.67 Tbps	6.67 Tbps	2.22 Tbps
100TB	13.33 Tbps	2.67 Tbps	666.67 Gbps	222.22 Gbps
10TB > 100Gbps	1.33 Tbps	266.67 Gbps	66.67 Gbps	22.22 Gbps
1TB	133.33 Gbps	26.67 Gbps	6.67 Gbps	2.22 Gbps
100GB 100Gbps	13.33 Gbps	2.67 Gbps	666.67 Mbps	222.22 Mbps
10GB < 10Gbps	1.33 Gbps	266.67 Mbps	66.67 Mbps	22.22 Mbps
1GB	133.33 Mbps	26.67 Mbps	6.67 Mbps	2.22 Mbps
100MB < 100Mbps	13.33 Mbps	2.67 Mbps	0.67 Mbps	0.22 Mbps
Time to transfer				

This table available at:

http://fasterdata.es.net/fasterdata-home/requirements-and-expectations/

ESnet

1 TB of data vs network speed

10 Mbps	300 hrs (12.5 days)
100 Mbps	30 hrs
1 Gbps	3 hrs
10 Gbps	20 mins

= DHL: time to ship… ☹

- The disk subsystem can also be a bottleneck
- Parallel streams
- Don't try saturate the network - be nice
- Rule of thumb: 1/4 to 1/3 of a shared path that has nominal background load
- E.g. **1 Gbps host: target 150-200 Mbps** (20-25 MB/s)

A small amount of packet loss makes a huge difference in TCP performance

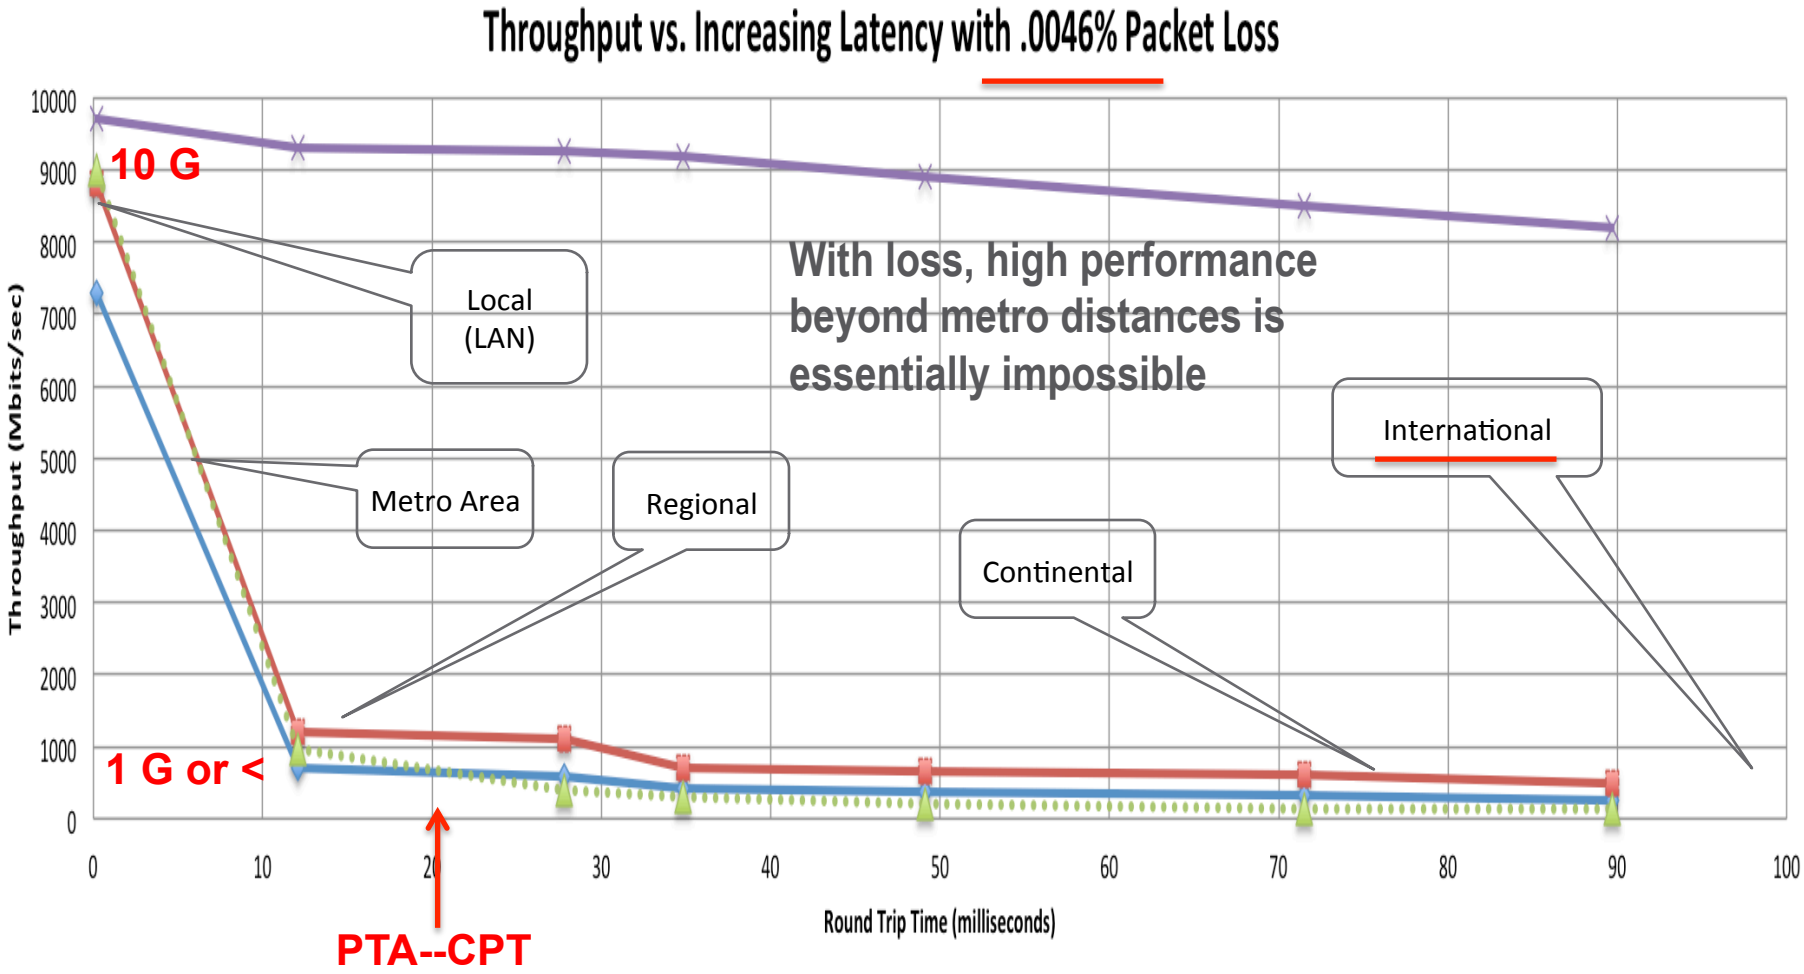

The Science DMZ Design Pattern

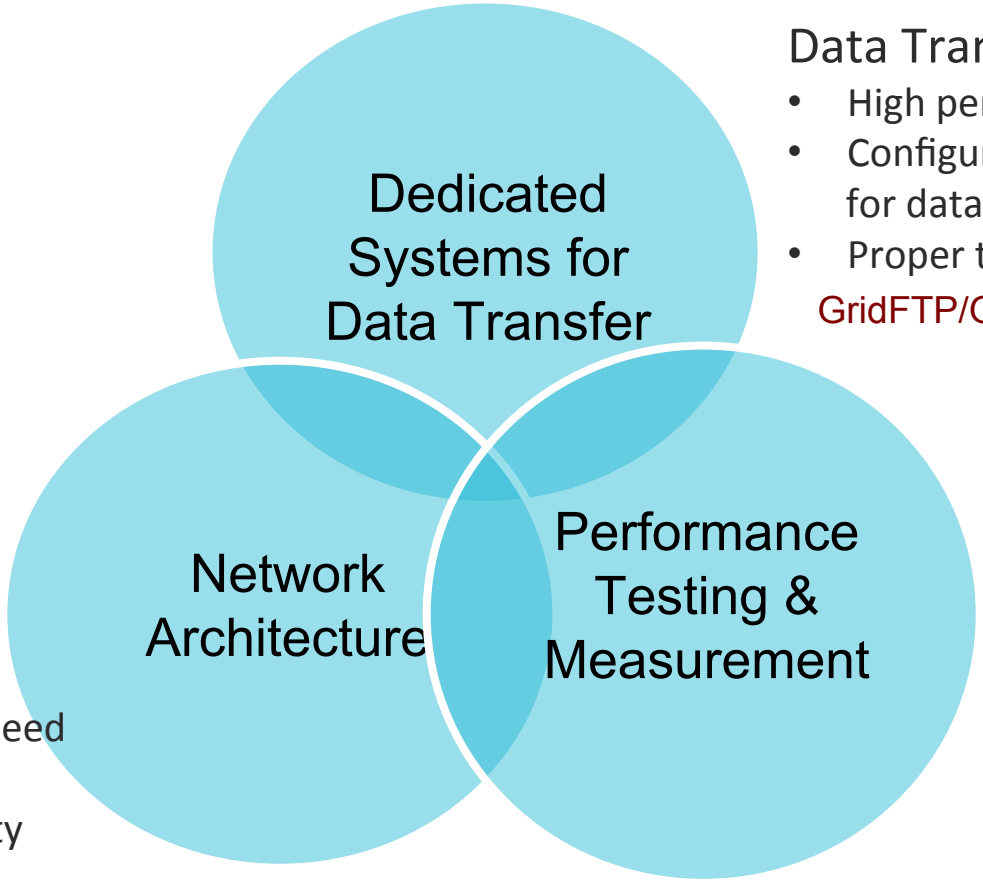

Data Transfer Node

- High performance
- Configured specifically for data transfer
- Proper tools

GridFTP/Globus, etc.

Science DMZ

- Dedicated network location for high-speed data resources
- Appropriate security
- Easy to deploy - no need to redesign the whole network

perfSONAR

- Enables fault isolation
- Verify correct operation
- Widely deployed in ESnet and other networks, as well as sites and facilities

ESnet

Science DMZ Design Pattern (Abstract)

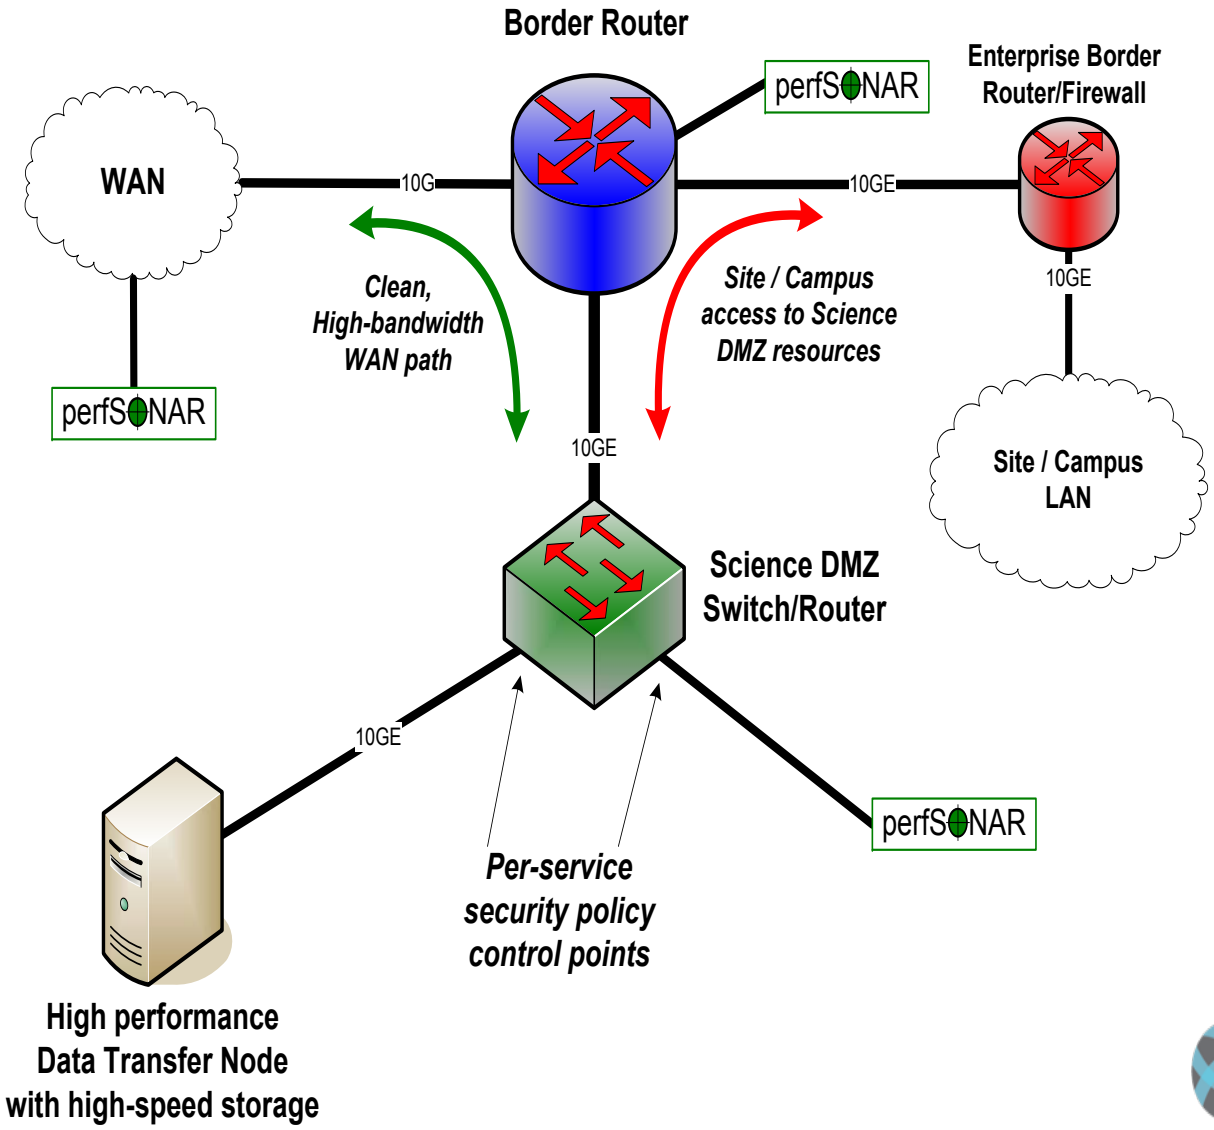

ESnet

SANReN
South African National Research Network

Why Science DMZ?

- **Performance**

Improve performance of data-intensive research

High-speed access to cloud resources

- **Usability**

Software to enable high-speed data transfer, e.g. Globus

- **Cost**

Delay expensive firewall upgrades

High-speed switches, rather than expensive routers

- **Security**

Layered security, applied on both network and host

SANReN
South African National Research Network

POCs / How we engaged

- Contacted all “known” big data transfer users by email to gauge interest (36 odd direct emails, all IT directors)
- Collected information on Big Data Transfer Projects – e.g. via DIRISA, SARIR
- We need help – who do you know?

SANReN
South African National Research Network

Big Data Transfer Projects

- Bioinformatics – UP
- South African Weather Services – Cape Town and Pretoria
- SAEON – Cape Town, Pretoria, Small sites
- H3ABioNet/CBIO – UCT
- Bioinformatics – SUN
- Global Change and Sustainability Institute – WITS

- SKA (Special Case)
- CERN (Special Case)

SANReN South African National Research Network

Geographic Summary

Cape Town	Pretoria	Johannesburg	Stellenbosch	International	Smaller sites
UCT (Various sites)	UP	WITS	SUN	Belgium	SANPARKS, Kimberley
CHPC	SAWS	UJ		US: Uni. Pennsylvania or NRAO in Soccoro, New Mexico.	SANPARKS, Phalaborwa
SAEON - SANBI and Foreshore	SAEON National Office colbyn				Port Elizabeth - Nelson Mandela Bay (Old CSIR Campus)
*Open Exchange point					UKZN Wildlife, Pietermaritzburg

SANReN
South African National Research Network

DTN Placement Plans

3 locations have been chosen for first phase:

- Wits University (Johannesburg)
- CSIR (Pretoria), and
- Teraco Data Centre (Rondebosch, Cape Town)

SANReN
South African National Research Network

Immediate next steps

- Engage with the sites where we would like to place DTN's
- Purchase of Hardware
- Testing performance tools: Globus/Aspera
- Training for SA NREN and community: workshops
- Implementation of Science DMZ architectures
- Run tests

SANReN
South African National Research Network

Thanks!

Questions?

pert@sanren.ac.za

See you at tomorrow's workshop?!!

science & technology
Department:
Science and Technology
REPUBLIC OF SOUTH AFRICA

CSIR
our future through science