

Science DMZ

John Hicks - Internet2
SPACI

June 7, 2018

Motivation and Introduction

- Research often requires moving large amounts of data between institutions and facilities using long-haul WANs
- A typical data transfer between two academic institutions in the U.S. requires a minimum of 5 networks each with a separate administrative domain
 - Campus-regional-national backbone-regional-campus
- Backbone and regional networks are usually designed to handle large data transfers
- In most cases, TCP is used as the transport protocol
- TCP is robust but it is not designed for high performance
- TCP's congestion control is designed to be fair by backing off when packet loss is detected

Feature

Enterprise data flow

Research data flow

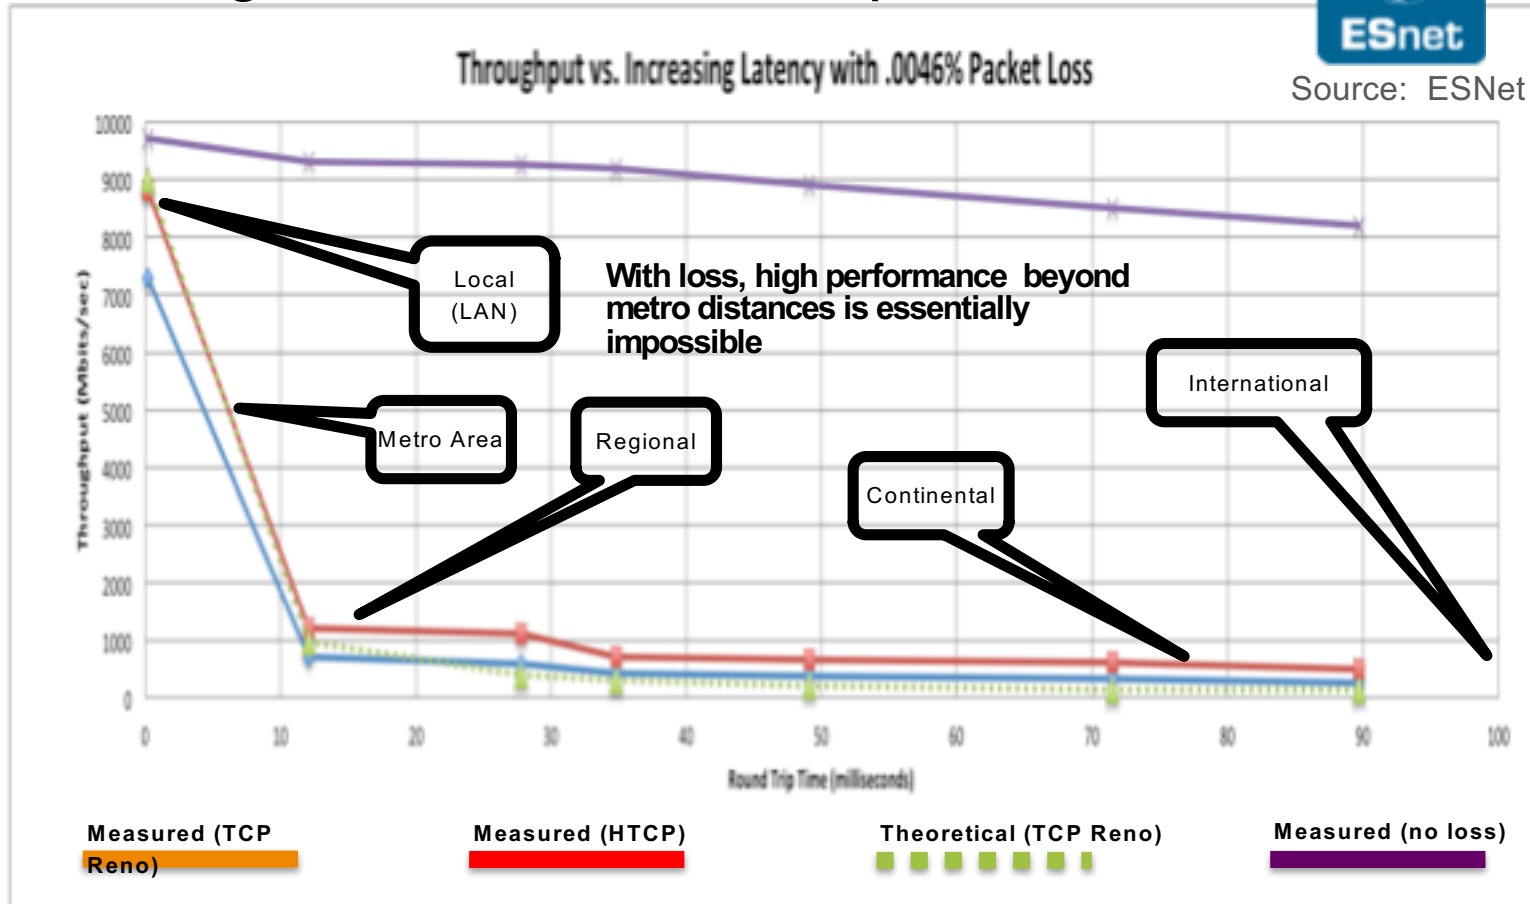
Feature	Enterprise data flow	Research data flow
Duration	Short	Long
Data size	KBs to MBs	TBs and above
Nature of the data	Large variety: web, email, media content, database-related, mobile applications, streaming	Files
Packet loss	Less sensitive	Very sensitive
Latency	Less sensitive	Sensitive
Throughput	Less sensitive	Very sensitive
Concurrent flows	Thousands to millions of flows per second	One to a few flows per second

Source: Jorge Crichigno - Northern New Mexico College

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



Source: ESNet



Science DMZ

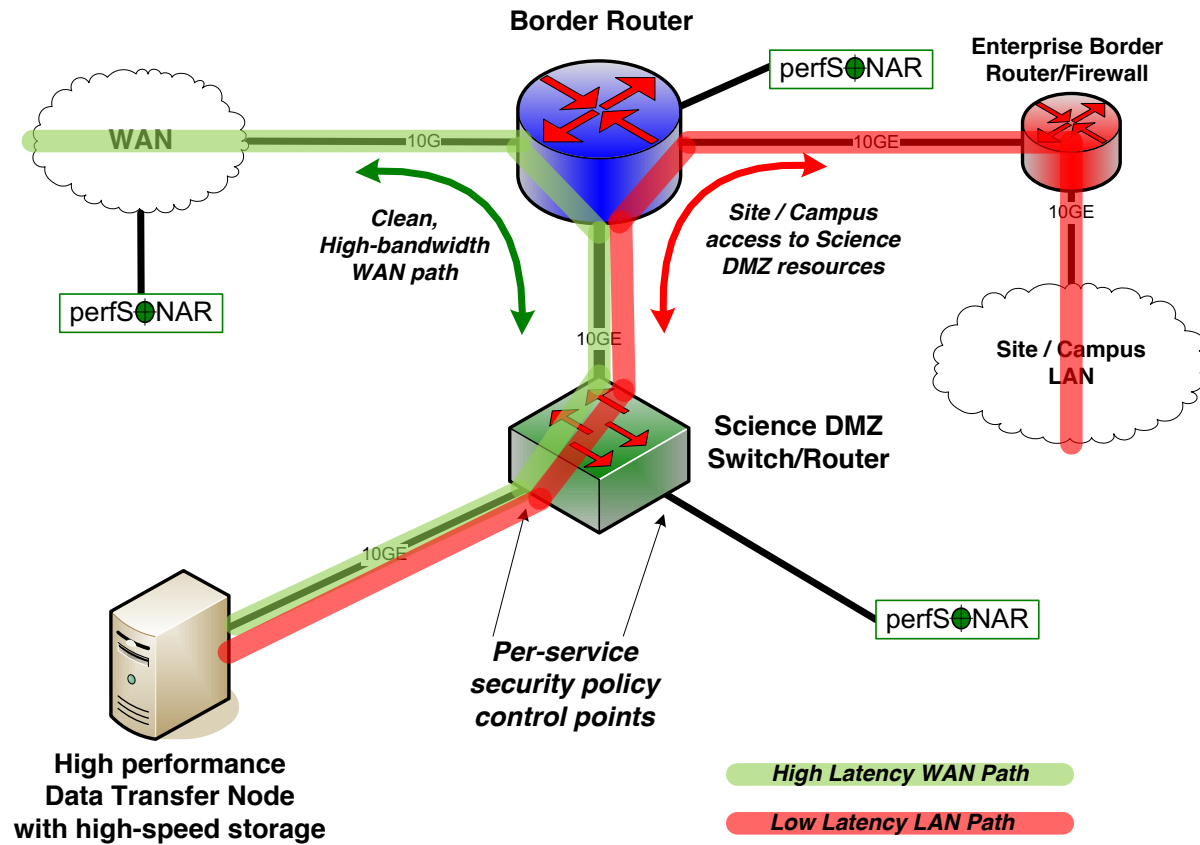
- Following an ecosystem model we are interested in improving end-2-end performance with an emphasis on the end
- To be affective, the end facility needs to have high-performance network devices, a high performance data movement device (DTN), and a monitoring system to debug problems (perfSONAR) *(the latter is out of scope for this talk)*.
- DTN (data transfer node), dedicated transport device(s) tuned for high performance WAN data transfers (e.g. FIONA – PRP)
- By focusing on a concentrated (small?) network segment, you can ensure the network devices and links are clean in addition to the DTN and perfSonar
- The Science DMZ is not a turn key solution, it is a blueprint

Science DMZ (In One Slide)

- Consists of 3 key components, all required:
- “friction free” network path
 - Highly capable network devices (wire-speed, deep queues)
 - Virtual circuit connectivity option
 - Security policy and enforcement specific to science workflows
 - Located at or near site perimeter if possible
- Dedicated, high-performance data movers
 - a.k.a.: Data Transfer Node (DTN)
 - Optimized bulk data transfer tools such as GlobusOnline/GridFTP
- Performance measurement/test node
 - perfSONAR
- Details at: <http://fasterdata.es.net/science-dmz/>

Source: B. Tierney @ ESnet

Local And Wide Area Data Flows



Source: ESNet

Thanks!

jhicks@internet2.edu