



ESnet

ENERGY SCIENCES NETWORK

Data Mobility

ESnet Science Engagement
Lawrence Berkeley National Laboratory

Data Mobility Break Out
Hands-On Workshop Networking Topics-EPOC,
NYSERNET, UoSC

April 5, 2022

Ken Miller ken@es.net



U.S. DEPARTMENT OF
ENERGY
Office of Science




Performance At Different Data Scales

Data set size					
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
		1 Minute	5 Minutes	20 Minutes	1 Hour
		Time to transfer			

This table available at:

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

A Better Approach: Science DMZ Design



DTN Design, History & Purpose

- Original concept came from initial Science DMZ Design (~2012)
- Basic idea:
 - Host(s) dedicated to the task of data movement (and only data movement)
 - Limited application set (data movement tools), and users (rarely shell access)
 - Specific security policy enforced on the switch/router ACLs
 - Ports for data movement tools, most in a 'closed wait' state
 - Nothing to impact the data channel
 - Typically 2 footed:
 - Limited reach into local network (e.g. 'control channel': shared filesystem, instruments)
 - WAN piece that the data tools use (e.g. 'data channel')
- Position this, and the pS node, in the DMZ enclave near the border

Why a DTN?

- DTN = Data Transfer Node
 - Purpose built server to transfer data
 - Tuned to the performance as necessary
 - A tuned 10G is better than 25,40,100G untuned servers
 - DTNs can have local storage, connected storage, or both
 - Multiple DTNs can be setup for various projects
 - Also, Multiple DTNs can be clustered together
-
- **Match the DTN to the capabilities of the site and wide area network infrastructure**

DTN Design Considerations

- Single Resource for moving data
- Largest possible NIC to match needs and cover data transfer speed Typically 10G or 25/40 for campus.
- At a minimum, connect directly to border router with R&E connectivity or in a Science DMZ environment.
- Fast CPU of 3.3 Ghz or greater to support higher speeds
- Multiple CPU cores of 8+ to scale with parallel data transfers
- Sufficient local storage and options to connect external for growth if necessary

Reference Architecture or Use cases:

<https://fasterdata.es.net/science-dmz/DTN/hardware-selection/>

DTN Reference Architecture - wash-dtn1.es.net

- CPU
 - 2 x Intel Cascade lake Xeon Gold 6246
 - 12 cores each
 - 3.3GHz 165W TDP processor
- Memory
 - 12 x 16G DDR4 2933 ECC RDIMM (192G total)
- Disk
 - 10 x Intel P4610 1.6TB U.2/2.5” PCIe NVMe 3.0 x4 Drives
 - 2 x Enterprise 960G 2.5” SATA SSD (OS, onboard Intel SATA Raid 1)
- Network
 - Mellanox ConnectX-5 EN MCX516A-CCAT 40/50/100GbE dual-port QSFP28 NIC
- Application
 - Globus
 - https://app.globus.org/file-manager?origin_id=2a6a759c-5cfe-4402-ac5e-a06d9d7f7c37&origin_path=%2F

Data Mobility Benchmark

- Try to benchmark your DTNs and Data Architectures monthly or after any changes.
- Download ESnet data Climate Data Sets from Wash-DTN1.es.net or another ESnet server to test your write speeds
 - https://app.globus.org/file-manager?origin_id=2a6a759c-5cfe-4402-ac5e-a06d9d7f7c37&origin_path=%2F
 - Climate-Small, ~245GB, 1496 files, 305 folders
 - Climate-Medium, ~245GB, 117 files, 1 folder
 - Climate-Large, ~245GB, 11 files, 1 folder
 - Climate-Huge, ~245GB, 2 files, 1 folder
- For larger systems, try the DME datasets:
 - https://app.globus.org/file-manager?origin_id=5837354e-7087-4d0d-b7bc-e3655f883899&origin_path=%2F
 - ds08, ~1TB, 30076 files, 1 folder
 - ds10, ~1TB, 100 files, 1 folder
 - ds16, ~1TB, 4 files, 1 folder
- Once downloaded, you can re-upload to test your read speeds.

Data Transfer Scorecard with Rates by Audience

Host Transfer Rates	$\frac{1}{6}$ PetaScale (Minimum)	$\frac{1}{3}$ PetaScale	$\frac{1}{2}$ PetaScale		PetaScale: 1 PB/wk	PetaScale: 1 PB/day
	10G Capable DTN				10xG, 25G, 40G, 100G DTNs	
Data Transfer Rate/Volume (Researcher)	1 TB/hr	2 TB/hr	3 TB/hr		5.95 TB/hr	41.67 TB/hr
Network Transfer Rate (Network Admin)	2.22 Gb/s	4.44 Gb/s	6.67 Gb/s		13.23 Gb/s	92.59 Gb/s
Storage Transfer Rate (Sys/Storage Admin)	277.78 MB/s	555.54 MB/s	833.33 MB/s		1.65 GB/s	11.57 GB/s

A benchmark table is provided to gauge data architecture performance, which can vary depending on number of files, folders, size of files, distance between sites, CI performance (network, server, disk/filesystem), as well as data transfer tool.





ESnet
ENERGY SCIENCES NETWORK

Questions?

EPOC contact
epoc@iu.edu

Ken Miller
ken@es.net



U.S. DEPARTMENT OF
ENERGY
Office of Science

