



# How Can I Speed Up My Data Transfers to/from NAS?

User Training Presentation, Mar. 7, 2012  
NASA Advanced Supercomputing Division

## Why is my data transfer so slow?

- Firewall
- Network Congestion
- Bad Cabling
- Application
- System Configuration
- System Issues



# Transfer Applications: BBFTP

- BBFTP:

- Pros:

- Fastest option when configured properly
    - Multi-stream file transfers
    - Only authentication is encrypted, less CPU overhead

- Cons:

- Might require software install
    - Bit of a learning curve
    - Need to make sure ports used are open on the firewall
      - Generally, these ports are 5020-5022 for control, and 5000-5011 or 50000-51000 for data; however by default, random ports will be chosen for data
    - Very specific options to work optimally

- Example:

```
bbftp -u username -e 'setnstream 8; setrecvwinsize 4096; setsendwinsize 4096; get filename' -E 'bbftpd -s -m 8' remotehost
```

# Transfer Applications: BBCP

- BBCP

- Pros:

- Similar to BBFTP
    - Multi-stream file transfers
    - Does not require server to be running all the time

- Cons:

- Peer-to-peer software may be blocked by firewalls

- Example:

```
bbcp -s 8 -w 4M filename remotehost:/home/username
```

# Transfer Applications: GridFTP

- GridFTP

- Pros:

- Works with Globus online
    - Multi-stream file transfers
    - Allows custom-defined TCP window sizes

- Cons:

- Can be complex to set up, if using grid certificates (gsiftp)
    - Difficult to install in user environment
    - Firewalls may block the data ports

- Example (non-GSI certificate mode):

```
globus-url-copy -p 8 -tcp-bs 4M file:/home/username/filename  
sshftp://user@remotehost/path/filename
```

# Transfer Applications: SCP

- Secure Copy Protocol (SCP)

- Pros:

- Easy to use
    - No custom setup

- Cons:

- Single-stream file transfer only
    - Potentially large CPU overhead
    - Older versions of SSH (pre-5.x and vendor-specific) have built-in static buffer sizes that greatly limit rates across the WAN

# Other Transfer Applications

- SFTP
  - Pros:
    - Easy to use (FTP-like)
  - Cons:
    - Single-stream file transfer
    - Encrypts data (potentially high CPU overhead)
- Rsync, RCP, FTP, wget, etc.
  - Should be avoided due to security concerns
  - If used, make sure the application being used encrypts your authentication credentials and does identification verification of the remote system

# Tuning



## Why is tuning important?

By default, systems are configured for LAN transfers. WAN transfers have higher latency and require larger congestion windows. Bandwidth Delay Product (BDP) is used to calculate the optimal TCP window size for a host:

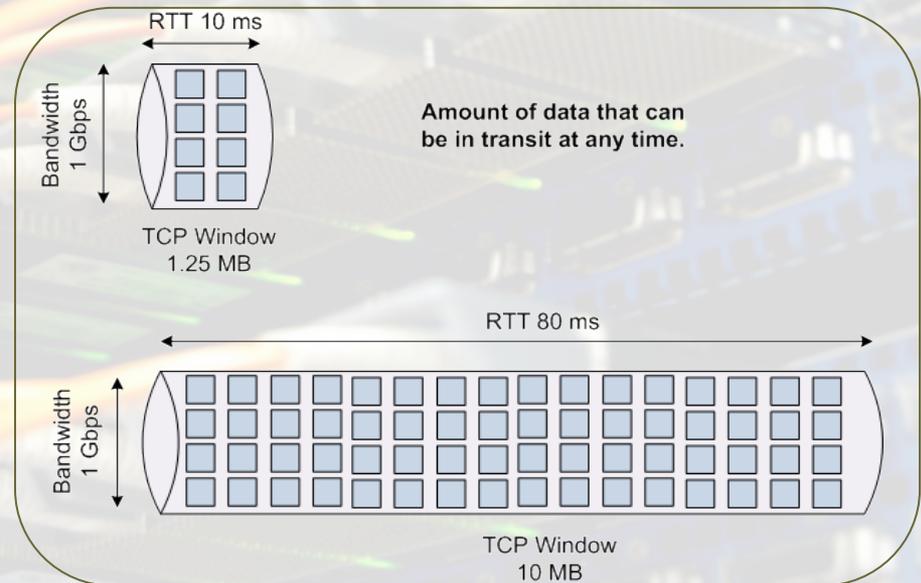
$$\text{Bandwidth} \times \text{Round Trip Time} = \text{Window Size}$$

An example: For a 10 Gbps link across the country, a TCP window of around 100MB is needed.

- For a 10 Gbps and a 80.0 ms delay;
- $1 \times 10^{10} \text{ b/s} \times 8 \times 10^{-2} \text{ s} = 8 \times 10^8 \text{ b}$  or 800Mb or 100MB Window Size

For 1 Gbps link, window of about 10MB is needed.

4MB is usually a good size for a majority of systems on a shared network.



# Tuning (continued)



## Example Linux Configuration:

1. Edit `/etc/sysctl.conf` and add:

```
net.core.wmem_max = 4194304
net.core.rmem_max = 4194304
net.ipv4.tcp_rmem = 4096 524288 4194304
net.ipv4.tcp_wmem = 4096 524288 4194304
```

2. Load with `sysctl -p` or reboot system

- Newer Macs come fairly well tuned by default.
- Windows requires special software to download to change the registry settings.

## For more information:

[http://www.nas.nasa.gov/hecc/kb/TCP-Performance-Tuning-for-WAN-Transfers\\_137.html](http://www.nas.nasa.gov/hecc/kb/TCP-Performance-Tuning-for-WAN-Transfers_137.html)

[http://www.nas.nasa.gov/hecc/kb/Optional-Advanced-Tuning-for-Linux\\_138.html](http://www.nas.nasa.gov/hecc/kb/Optional-Advanced-Tuning-for-Linux_138.html)

<http://fasterdata.es.net>



# Tricks

- Compress data!
  - If many small files, use `tar` first
  - Some applications support compress on the fly:
    - `scp -C`
    - `bbftp -c`
    - `bbcp -c`
- Use 10G bridge nodes instead of PFEs
  - PFEs are login nodes and not designed for transfers
  - PFEs often fully utilized and will give poor transfer rates
- If using `scp`, use low overhead cipher
  - e.g. `scp -carcfour file user@remotehost:`

# Tools

- Network Diagnostic Server:

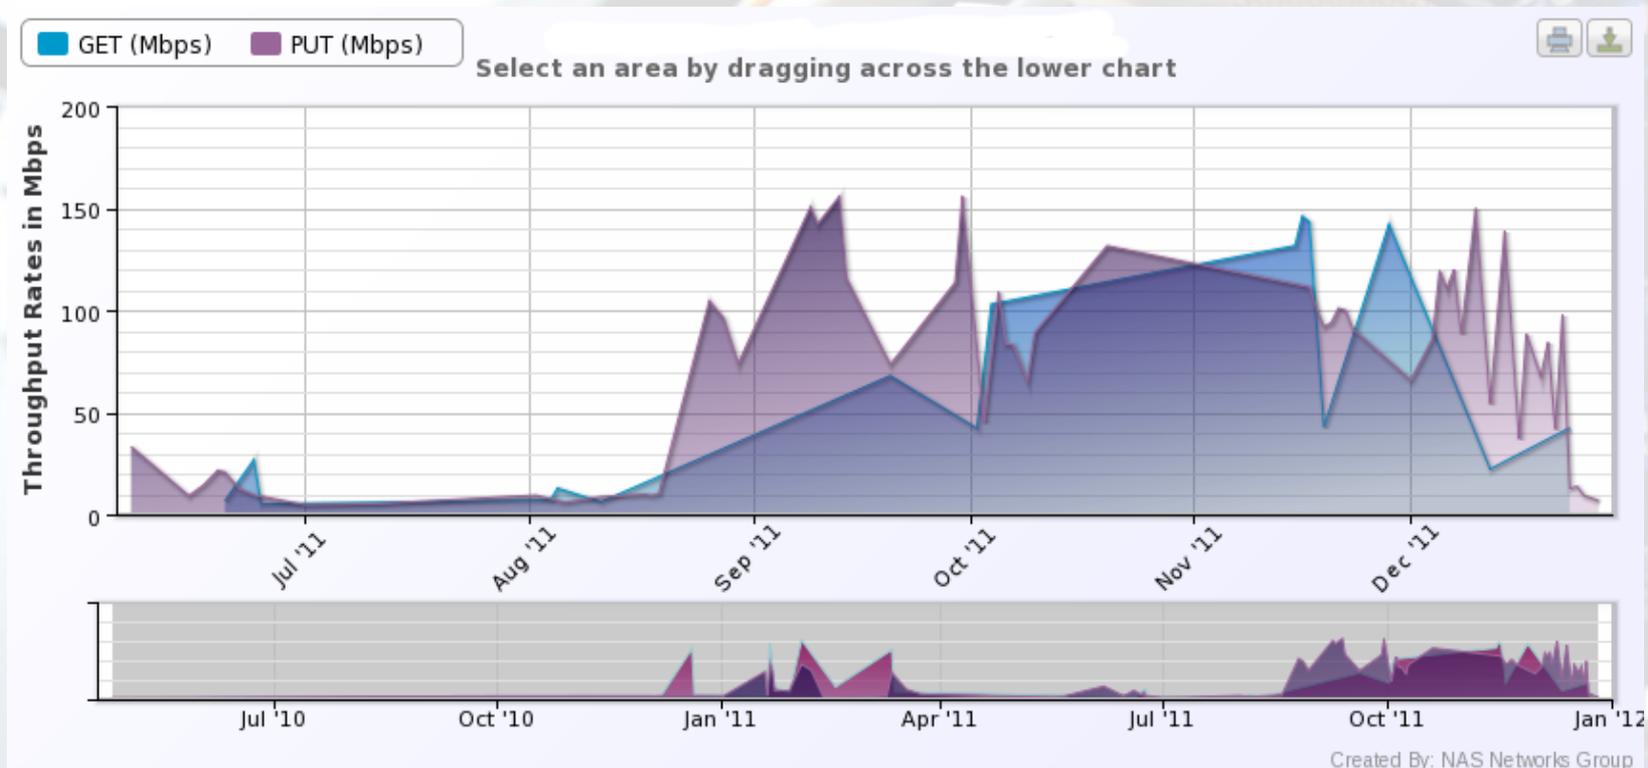
<http://npad.nas.nasa.gov>

(Java web software or downloadable CLI tool)

- Network Utilization Display:

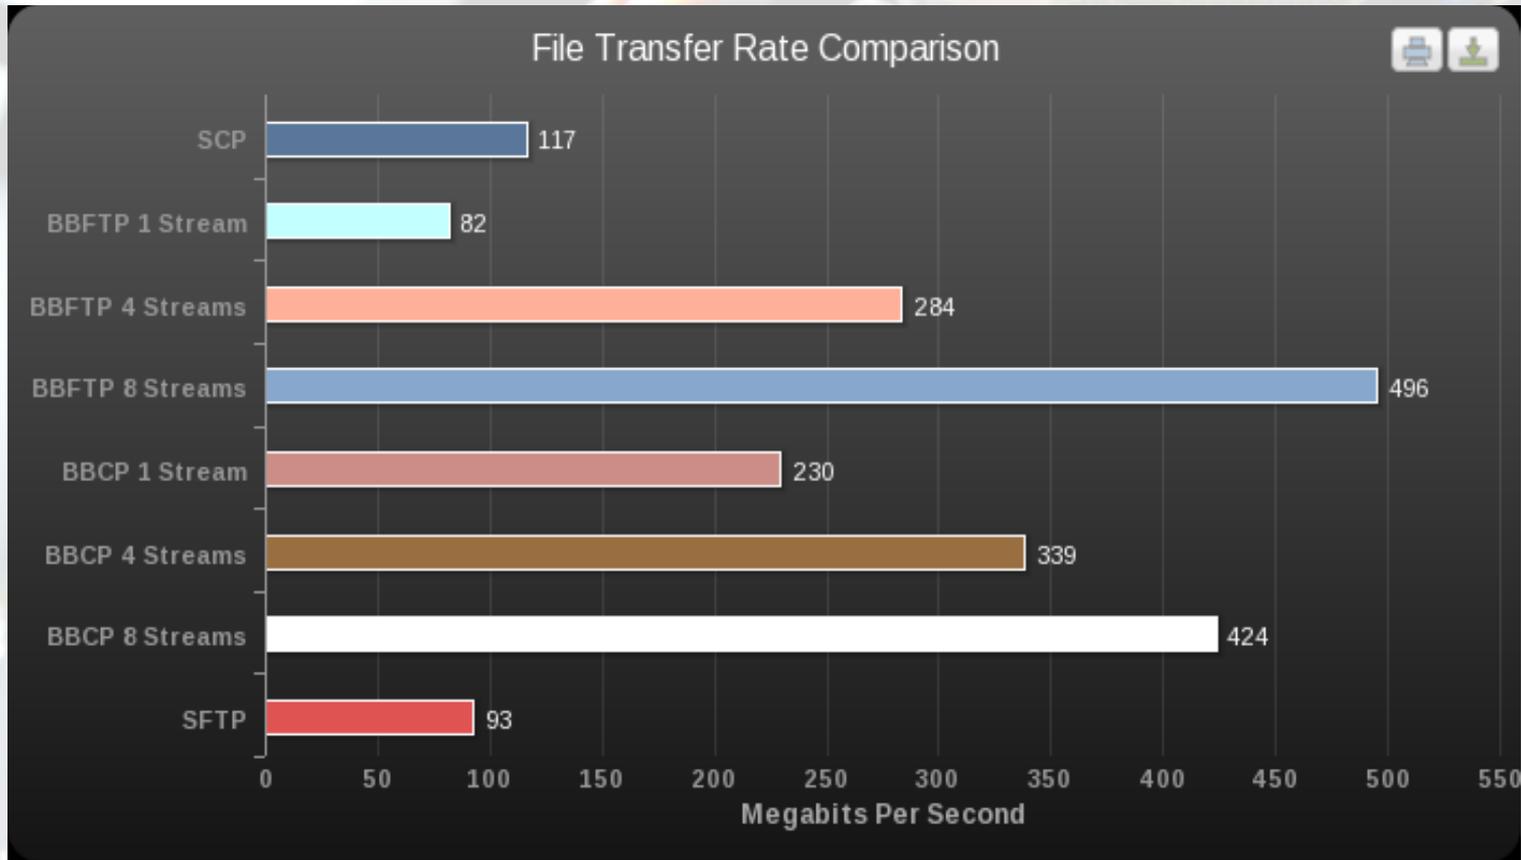
[http://www.nas.nasa.gov/hecc/support/pfe\\_utilization.html](http://www.nas.nasa.gov/hecc/support/pfe_utilization.html)

# Graph Showing System Tuning Improvement



The bottom graph shows a timeline with the larger graph zooming in on the last six months. System had about 10X performance increase from tuning. The timeline graph shows that it lost its tuning configuration (due to an upgrade). NAS engineers identified this and contacted the user to have it re-tuned.

# Graph of File Transfer Application Speeds



\* Host used in testing was on a 1 Gbps WAN link with 20 ms latency. All tests were done with the same systems and tuning.



# Support

NAS proactively monitors traffic flows:

- Looks for flows meeting specific criteria
- Works with users to help improve their transfer rates

For support contact the NAS Control Room at:

Phone: 800-331-8737

Email: [support@nas.nasa.gov](mailto:support@nas.nasa.gov)