


# The Blue Waters



Network Ecosystem

# Overview

- Who are NCSA Neteng?
- What does the networking to BW look like?
  - Within NPCF
  - Outside NPCF
- Performance Issues
  - Where to start looking?
  - Typical Suspects
- How to use the perfsonar nodes



# NCSA Neteng

5 person group responsible for all aspects of networking at NCSA

- Blue Waters

- HPC Systems

- Office networks

- Enterprise and Datacenter / Innovation projects

- Supporting services

2 staff members assigned to Blue Waters for system and user support


- Provide assistance with resolving network related issues

  - poor transfer performance

  - protocol communication issues

  - network solutions for novel computing approaches

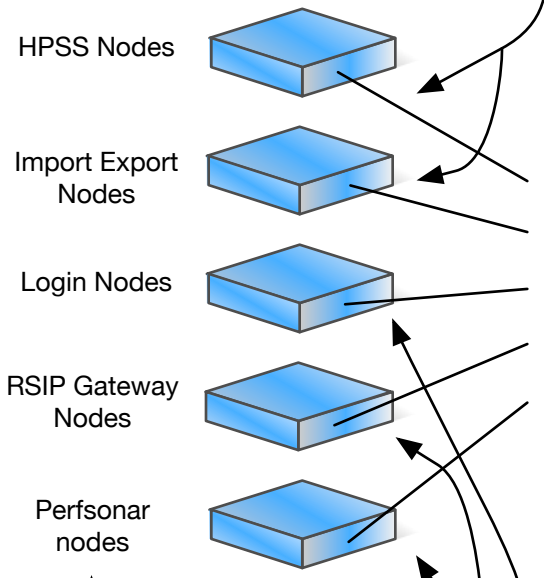
# The 10k ft view



When you are interacting with Blue Waters your traffic is taking one of many paths to reach it.

# Within NPCF


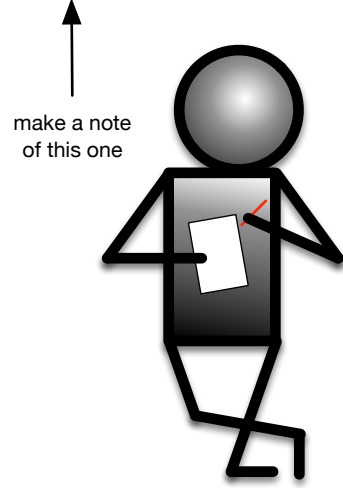
Blue Waters external facing nodes all on the 141.142.176.0/24 network



These are all 40G attached

These are each a crazy 24x10G bundle

These are all 10G attached



BW-CORE




CORE-WEST




CORE-EAST


These are all 5x10G bundles



EXIT-WEST



EXIT-EAST



Here is what the physical network looks like within the NPCF datacenter

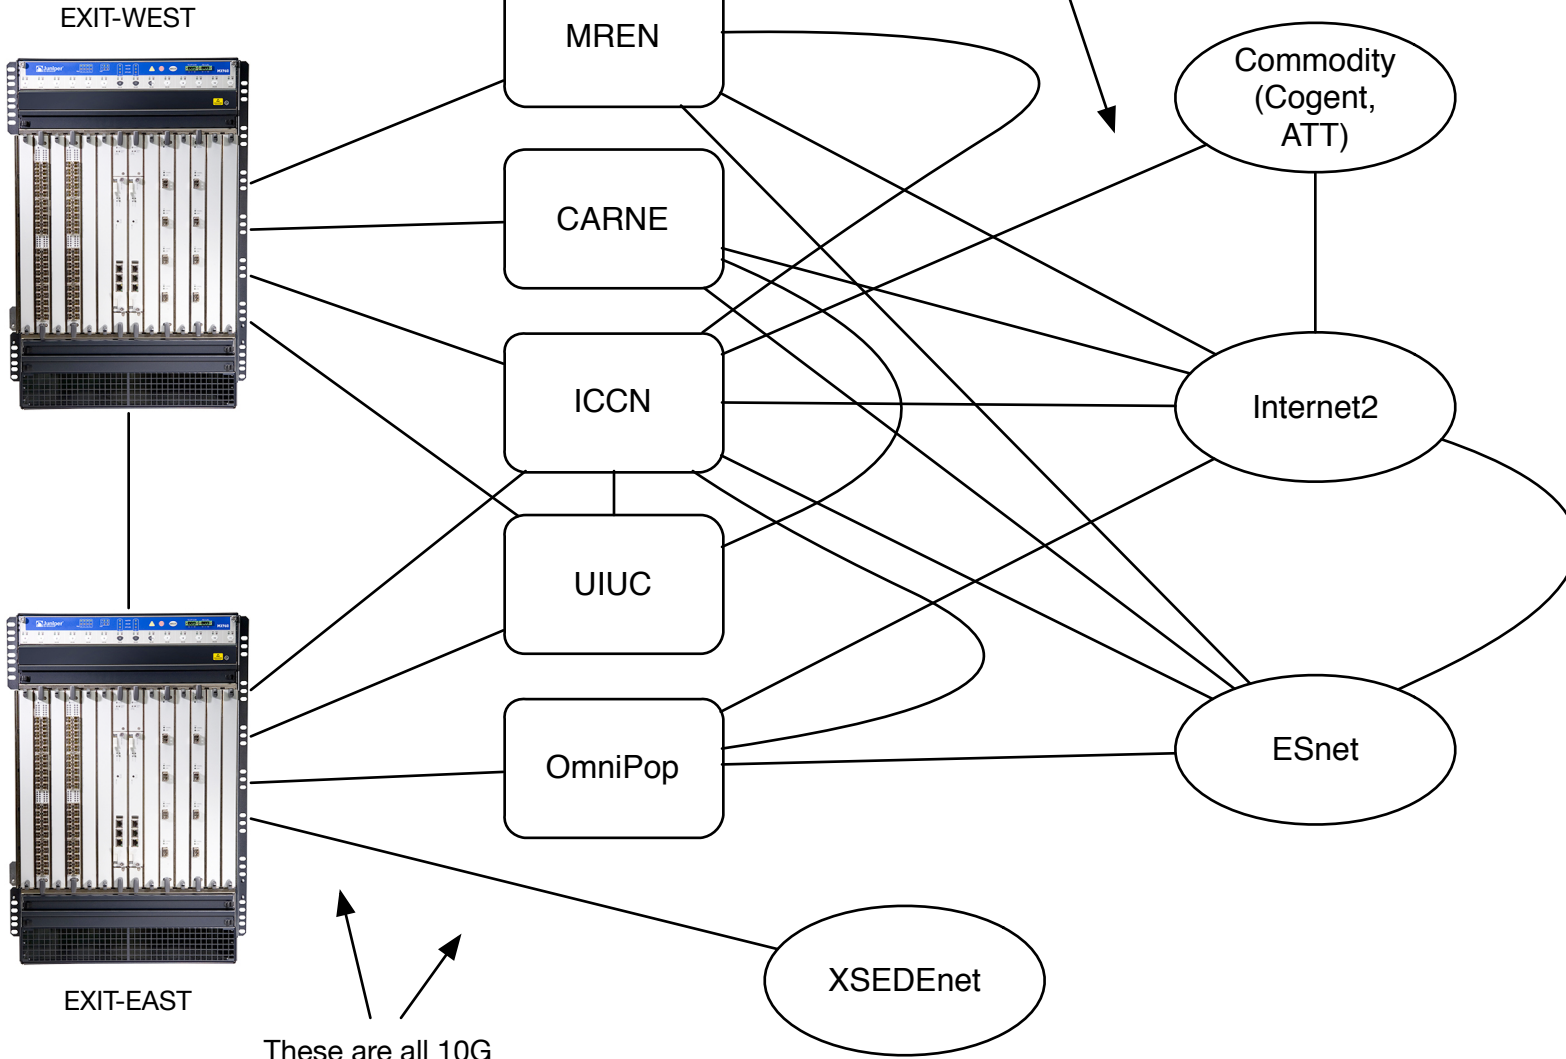
*Someday all these 10G links will be replaced with a few 100G links*

# Outside NPCCF

Regional Networks

Some of these are 10G  
Some are 100G

National Networks



What networks do you connect to?

What networks do those networks connect to?



# Performance issues



Highly connected networks are a blessing and a curse  
nice for resiliency

problematic from a routing policy standpoint

what works for one destination may not work for another

every end site makes their own policy on what to announce  
and where to announce it.

We frequently end up looking at the network path when troubleshooting a transfer issue. If there is something we can do at our end to adjust the path, we will do it. Otherwise we work with the other providers or the end site to understand their policy and make any adjustments.

# Performance Issues Cont'd

There are many different elements that can affect network throughput. ESnet has a great slide deck covering all the main points:

<http://fasterdata.es.net/assets/fasterdata/JT-201010.pdf>

## Quick Summary:

Application - SCP, sFTP, FTP (not so good), GridFTP/GO (best)


if you must run SCP: <http://www.psc.edu/index.php/hpn-ssh>

First hop - are you on a fast network, does the equipment have adequate buffering? Can you run 9k MTU?

Tuning - Most OS's today auto-tune. See ESnet deck for details

Firewalls - Is there one in the path?

Is your disk system the bottleneck?





# PerfSonar nodes

PerfSonar nodes are test points hosted at many organizations and network providers. They can gather data on throughput, latency and packet loss.

<http://stats.es.net/ServicesDirectory/> ← For finding PS nodes

They also provide end users with tools to analyze their system's network stack and get an idea of what transfer speeds to expect. (NPAD and NDT)

BW has two perfsonar nodes available for use

<http://psb-bw.ncsa.illinois.edu> :

For bandwidth testing, TCP stack tuning

<http://psl-bw.ncsa.illinois.edu> :

For latency and packet loss testing

NCSA perfsonar node on the border router: <http://ps.ncsa.illinois.edu>

# Questions?

