

Tier 3g Infrastructure

Doug Benjamin
Duke University

Infrastructure examples

- Infrastructure is not glamorous. Understanding your needs and capabilities is critical to well running Tier 3
- Examples of Infrastructure include:
 - Networking
 - Physical space and associated hardware (Racks)
 - Electrical Power and Cooling
 - Computer security / data security
 - System administration and maintenance

Physical Space

- Prior to making your computer purchases determine where you will put your hardware.
- Issues to consider are:
 - 1 Rack of computers is heavy > 1000 lbs
 - Rack of computers is noisy and generates a lot of heat
 - Does your University department have a computer room that you can use part of.
 - Do you have space for eventual expansion?
 - Do you have easy access to machines for repairs?

Electrical Power

- What type of electrical power is available? (110 or 220 V) How much current? (number of circuits)
- Each R710 draws 300W (max) 200W (nominal)
- I.e. 10 servers in a rack will draw 3000W
- Consider other equipment as well. E.g. UPS.
- Check with local safety--50-70% of the total circuit capacity can normally be assigned.

Cooling

- Sufficient cooling important to operation of your cluster.
- In next talk Walker Stemple of Dell will show numbers for Power/cooling of various Tier 3 configurations. I am using two examples here for illustration:
- Case 1 - 30K\$ - storage on worker nodes – **4745 W** (@220V) ~ **16000 Btu/hrs** ~ **1.4 tons/AC** (1 ton AC = 12000 Btu/hrs)
- Case 2 – 41K\$ - (storage on worker node+ extra centralize storage- 96 TB total) – **5245 W** ~ **17800 Btu/hrs** ~ **1.5 tons** of AC

Networking Questions

(prior to purchases)

- Determine who is the network responsible for your department? Is she/he responsible for your cluster also?
- Who is the Campus network responsible? Meet them if you can?
- Determine the available bandwidth between your computers and campus backbone?
- Determine the available bandwidth across the campus backbone?
- Determine the available campus bandwidth to Internet 2?

Networking questions (continued)

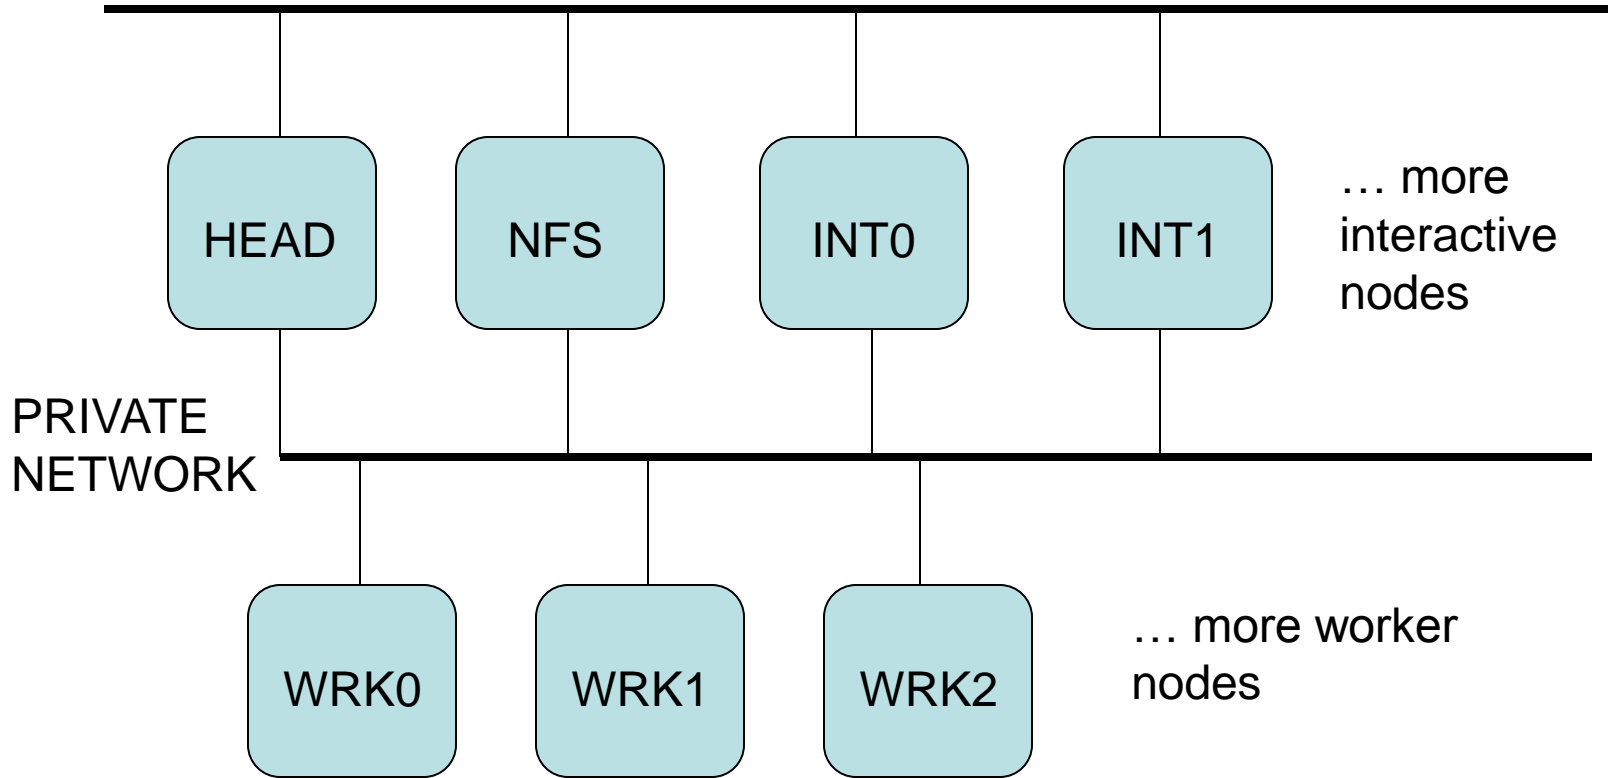
- Is the amount of available bandwidth sufficient for your needs? (100 MB/s ~ 1 TB /day)
- Determine how much networking infrastructure you will have to purchase? Can you use Dell managed switches? Does your campus require Cisco or another vendor?
- Will you have to pay for bandwidth used?

Public/Private compute networks

- How many public IP address can you get?
- What is the campus firewall policy?
 - Some places (like ANL) have several networks - green network – available to general internet via specific ports , yellow network – general campus network. – visitor network – more restricted
- Do you need a private network for your cluster?
- Tier 3g baseline has public and private networks
 - Added complexity with advantages
 - No firewall on private network.

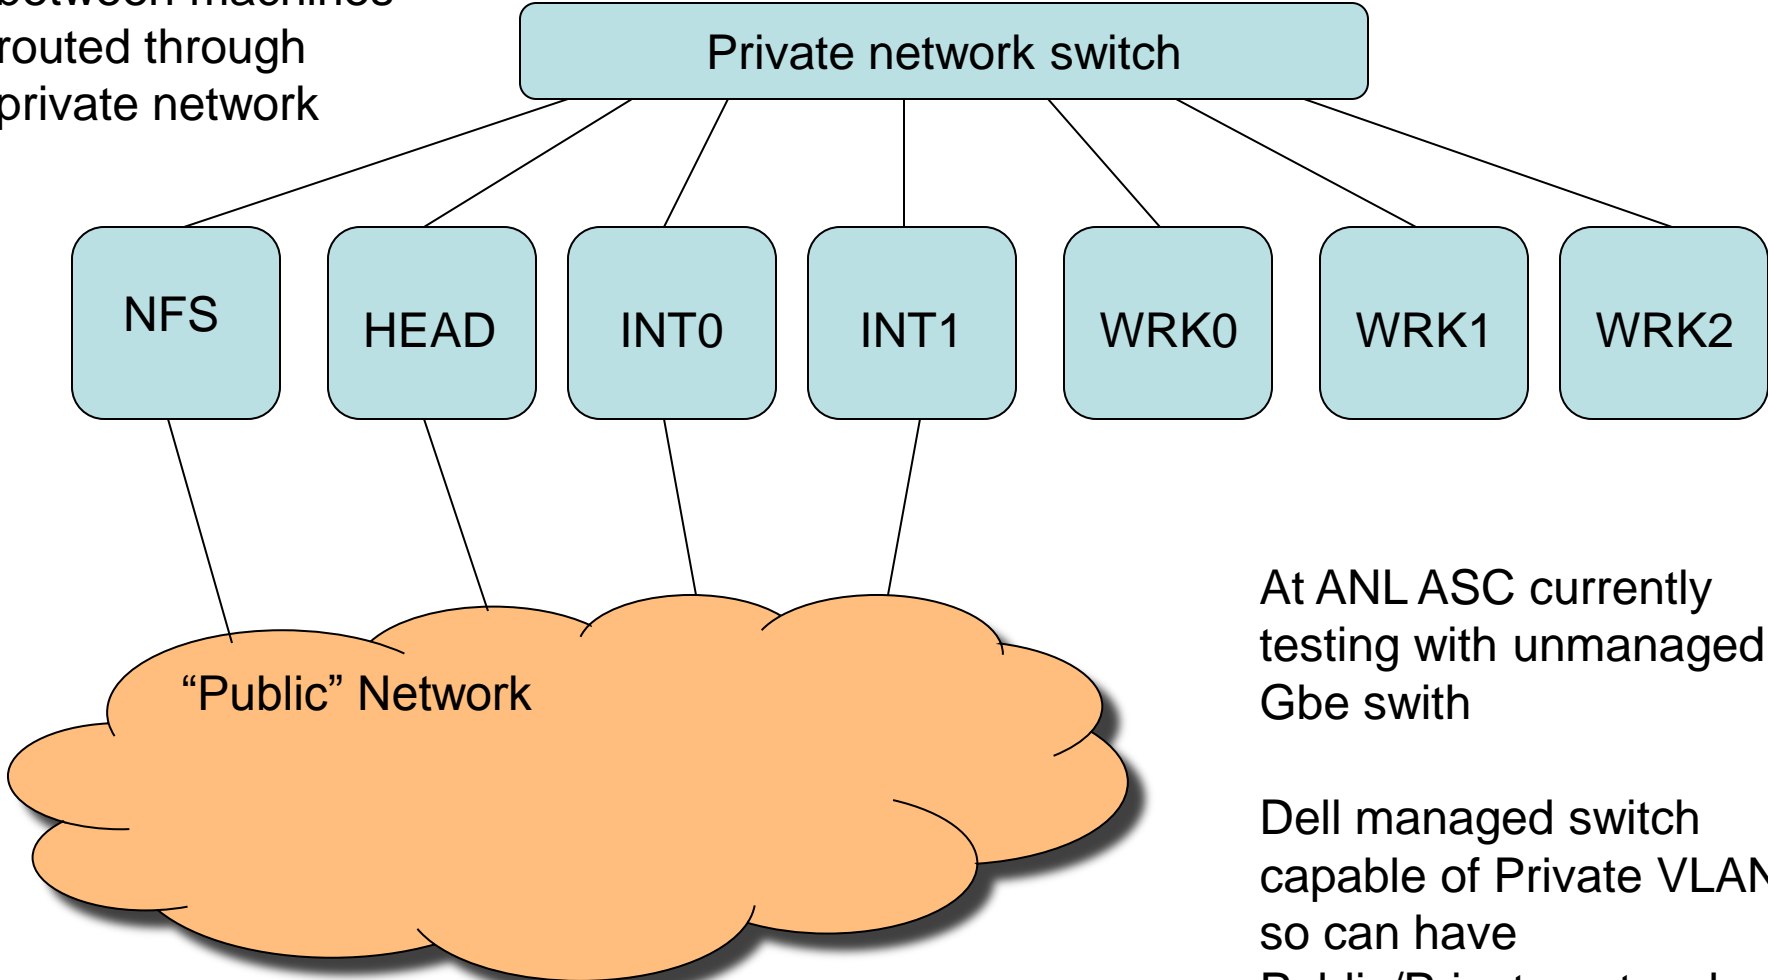
BASELINE T3g NETWORKING

“PUBLIC” NETWORK



BASELINE T3g NETWORKING

All network traffic between machines routed through private network





At ANL ASC currently testing with unmanaged Gbe switch





Dell managed switch capable of Private VLAN so can have Public/Private networks on same switch

T3g NETWORK at ANLASC





“GREEN”
NETWORK



“YELLOW”
NETWORK



PRIVATE
NETWORK



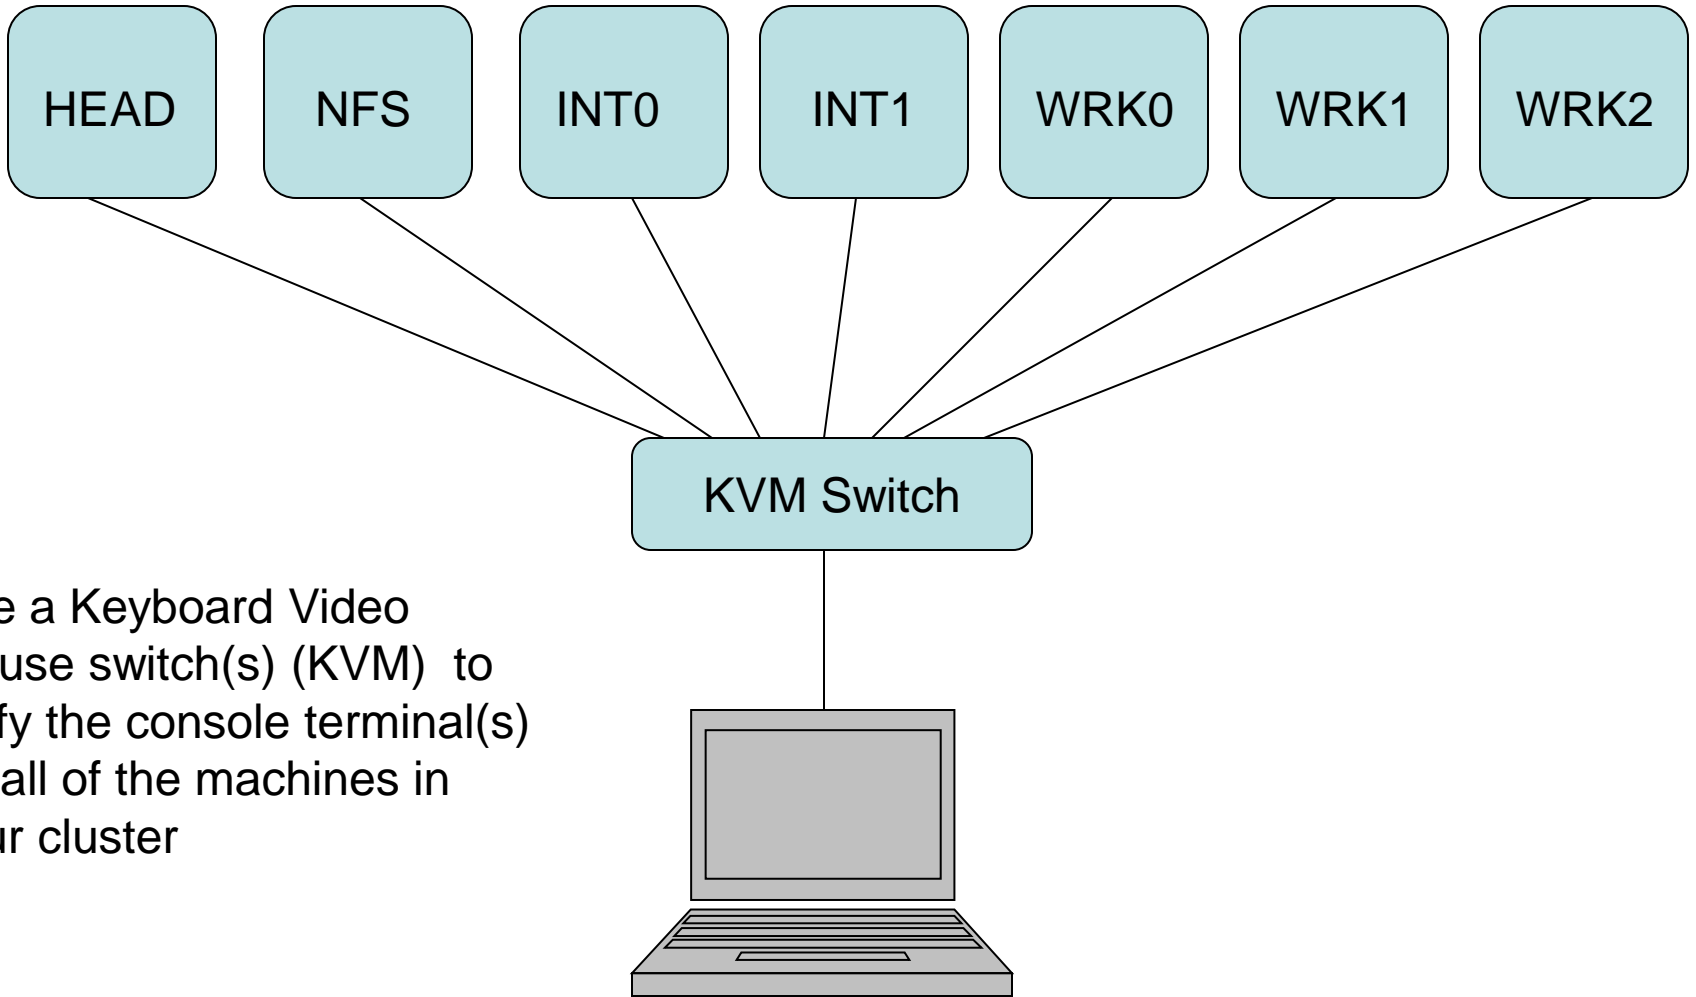
Networking/Cluster design issue

- Think how you will get data to your site?
 - Currently “dq2-get” is safest way to fetch data
 - Recent events so that Failures at a Tier 3 (disks filling up) can have negative effect at Tier 1 – system needs to be made more robust
- Do you want a centralize data space? (good idea)
- How will you access data within your site?
 - NFS SL Linux access insufficient
 - At Duke 1 person running Athena jobs from NFS mounted data put high load on NFS server – client jobs starved for data
 - *Data mover is required for reliable operation*
 - Data on worker nodes (XRRootD) reduces network load - **most efficient data access** – worker nodes will need sufficient disk spindles

System administration issues

- Does your department have a system administrator(s) who can help you?
 - Can they administer the machines (OS/accounts etc)?
 - Will you have to do it all but they provide expert guidance?
- Who is responsible for machine up keep (hardware and software)?
- What is your data preservation plan? What is your backup strategy? (We are missing this piece from the Tier 3 instructions)

KVM to unify console terminal



Use a Keyboard Video Mouse switch(s) (KVM) to unify the console terminal(s) for all of the machines in your cluster

Computer security

- Secure computers are vital to our ability to produce the physics results.
- What are your campus/department computer security policies?
- Who is the department computer security contact? Meet with them.
- What will be your role for your cluster?
- We do not want to be the weak link in the computer security chain. - Computer security should not be ignored.

Conclusions

- Infrastructure is not glamorous but necessary
- Consider the Infrastructure costs and issues before making computer purchases. Save money for it.
- Plan for reasonable cluster expansion.
 - (we might be lucky and get more funds in the future)
- Some forethought now will save you headaches in the future