# LOW-LEVEL NKN ARCHITECTURE

Luke Sheneman

NKN Technical Talk

09/24/2015

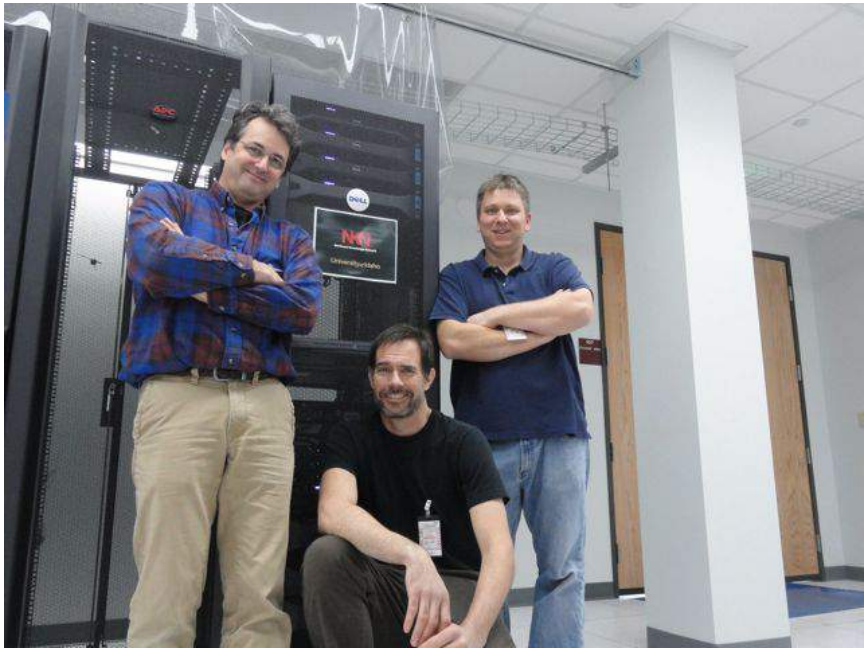

# Origins of the NKN Architecture

Northwest Knowledge Network (NKN)
Systems Architecture, Requirements, and Recommendations

Luke Sheneman, Greg Gollberg

03/24/2011

ABSTRACT

More and more research data are being generated every day and the dollar investment in data generation and acquisition is staggering. While much of this data collection has been funded with federal dollars, public access to these data products is largely nonexistent. Technologies across disciplines accelerate research data acquisition (e.g. LIDAR and other remotely sensed geospatial datasets, higher resolution model outputs, etc.) and at a statewide level research data storage needs are growing accordingly. In Idaho, much of these data are widely distributed and they exist in uncoordinated, isolated and uncurated datasets. Grant agencies such as the NSF increasingly expect data management plans, data sharing policies, and normalized metadata from grant applicants. Data must not only be acquired and stored; data must also be discoverable, interoperable, sharable, and in a form that lends itself to true research collaboration. We describe an architecture for storing statewide research data in a distributed, secure, discoverable and interoperable form that enables a transformative level of data replication, collaboration and analysis. This proposed systems architecture involves high capacity, highly available data storage hardware along with the related server systems that enable the core functionality of the Northwest Knowledge Network (NKN).

## 1 Introduction

The Northwest Knowledge Network (NKN) will be a large-scale regional clearinghouse for research data that serves students, faculty, researchers, government agencies, stakeholders and decision makers. These data will span many distinct collections. Initially, these data will likely come from EPSCoR-sponsored research and therefore be

# History of NKN Hardware Acquisition

- NKN had ***no*** hardware infrastructure prior to June 2011
- NKN purchased a development server with 24TB storage in June 2011 through EPSCoR Track 1 Funds
- NKN purchased our existing hardware in Fall of 2011 and was fully deployed at INL and UI by early 2012 (birth of the modern NKN)
- NKN upgraded network components through 2013 NSF CC*NIE Award and WC-WAVE EPSCoR funding
- NKN added storage in early 2014 to INL and UI

# Overall Design Goals

- Distributed Datacenters
  - Backup, Recovery, Load Balancing, Failover
    - University of Idaho
    - Idaho National Laboratory
- Scalable Enterprise Storage
- Flexible Virtualized Server Environment
- Entirely Redundant Components
- Security:  Tightly Controlled and Managed

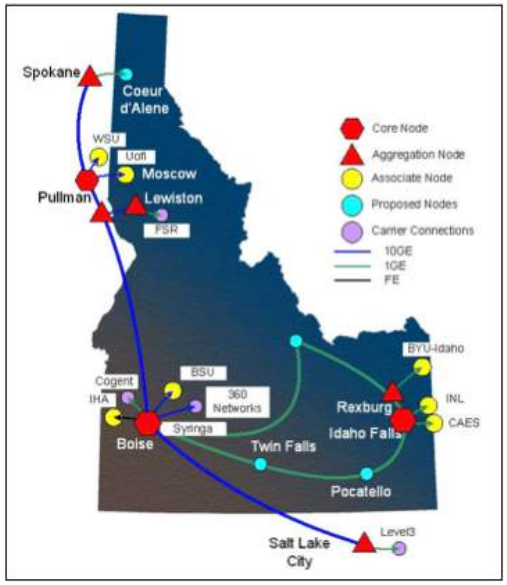

# High Level NKN Network Architecture

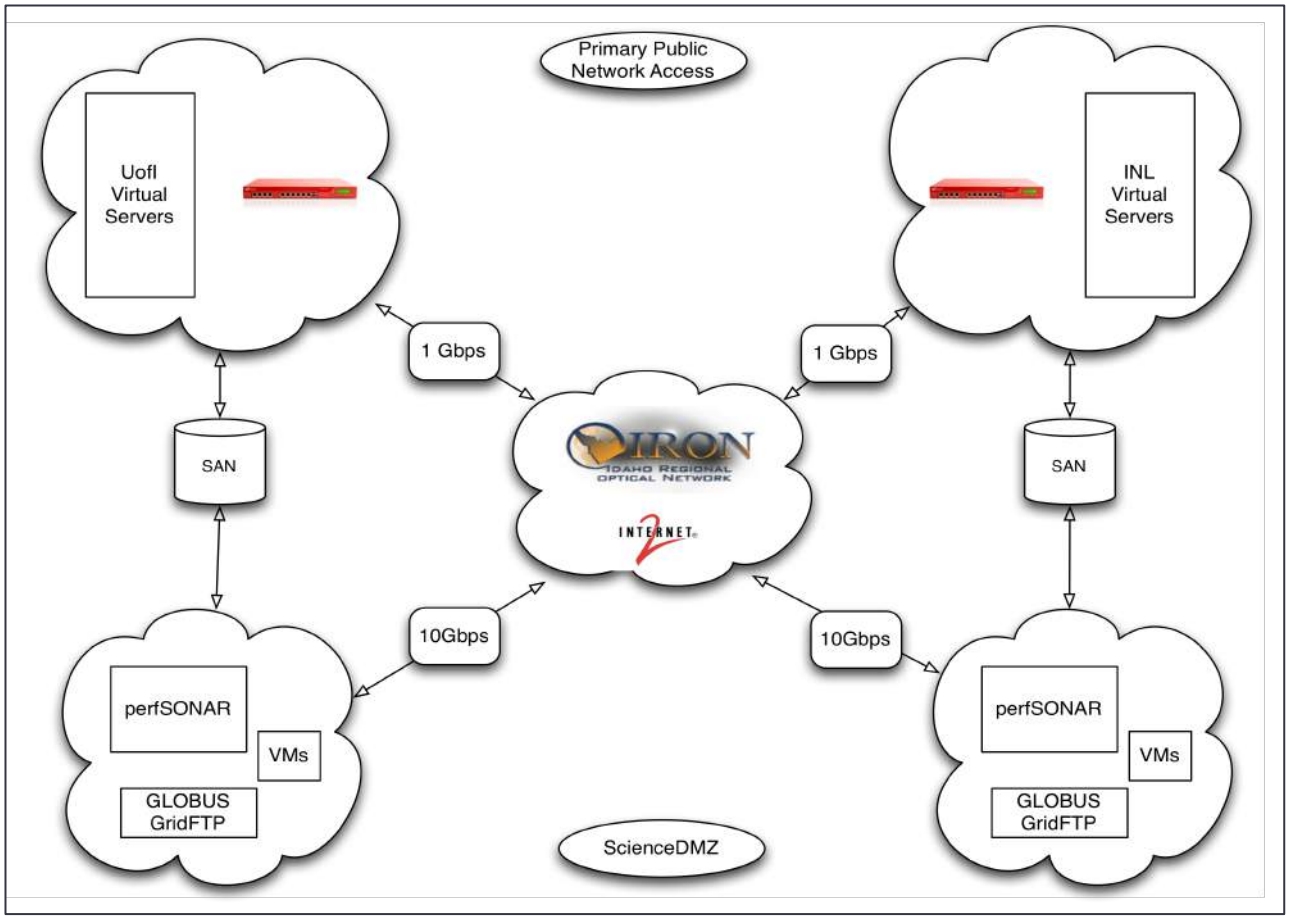


Highlights:

- Dedicated Firewalls (Watchguard)
- Science DMZ
  - perfSONAR
  - Dedicated File Transfer
  - No Firewall
- Secure Shared SAN
- NSF CC-NIE

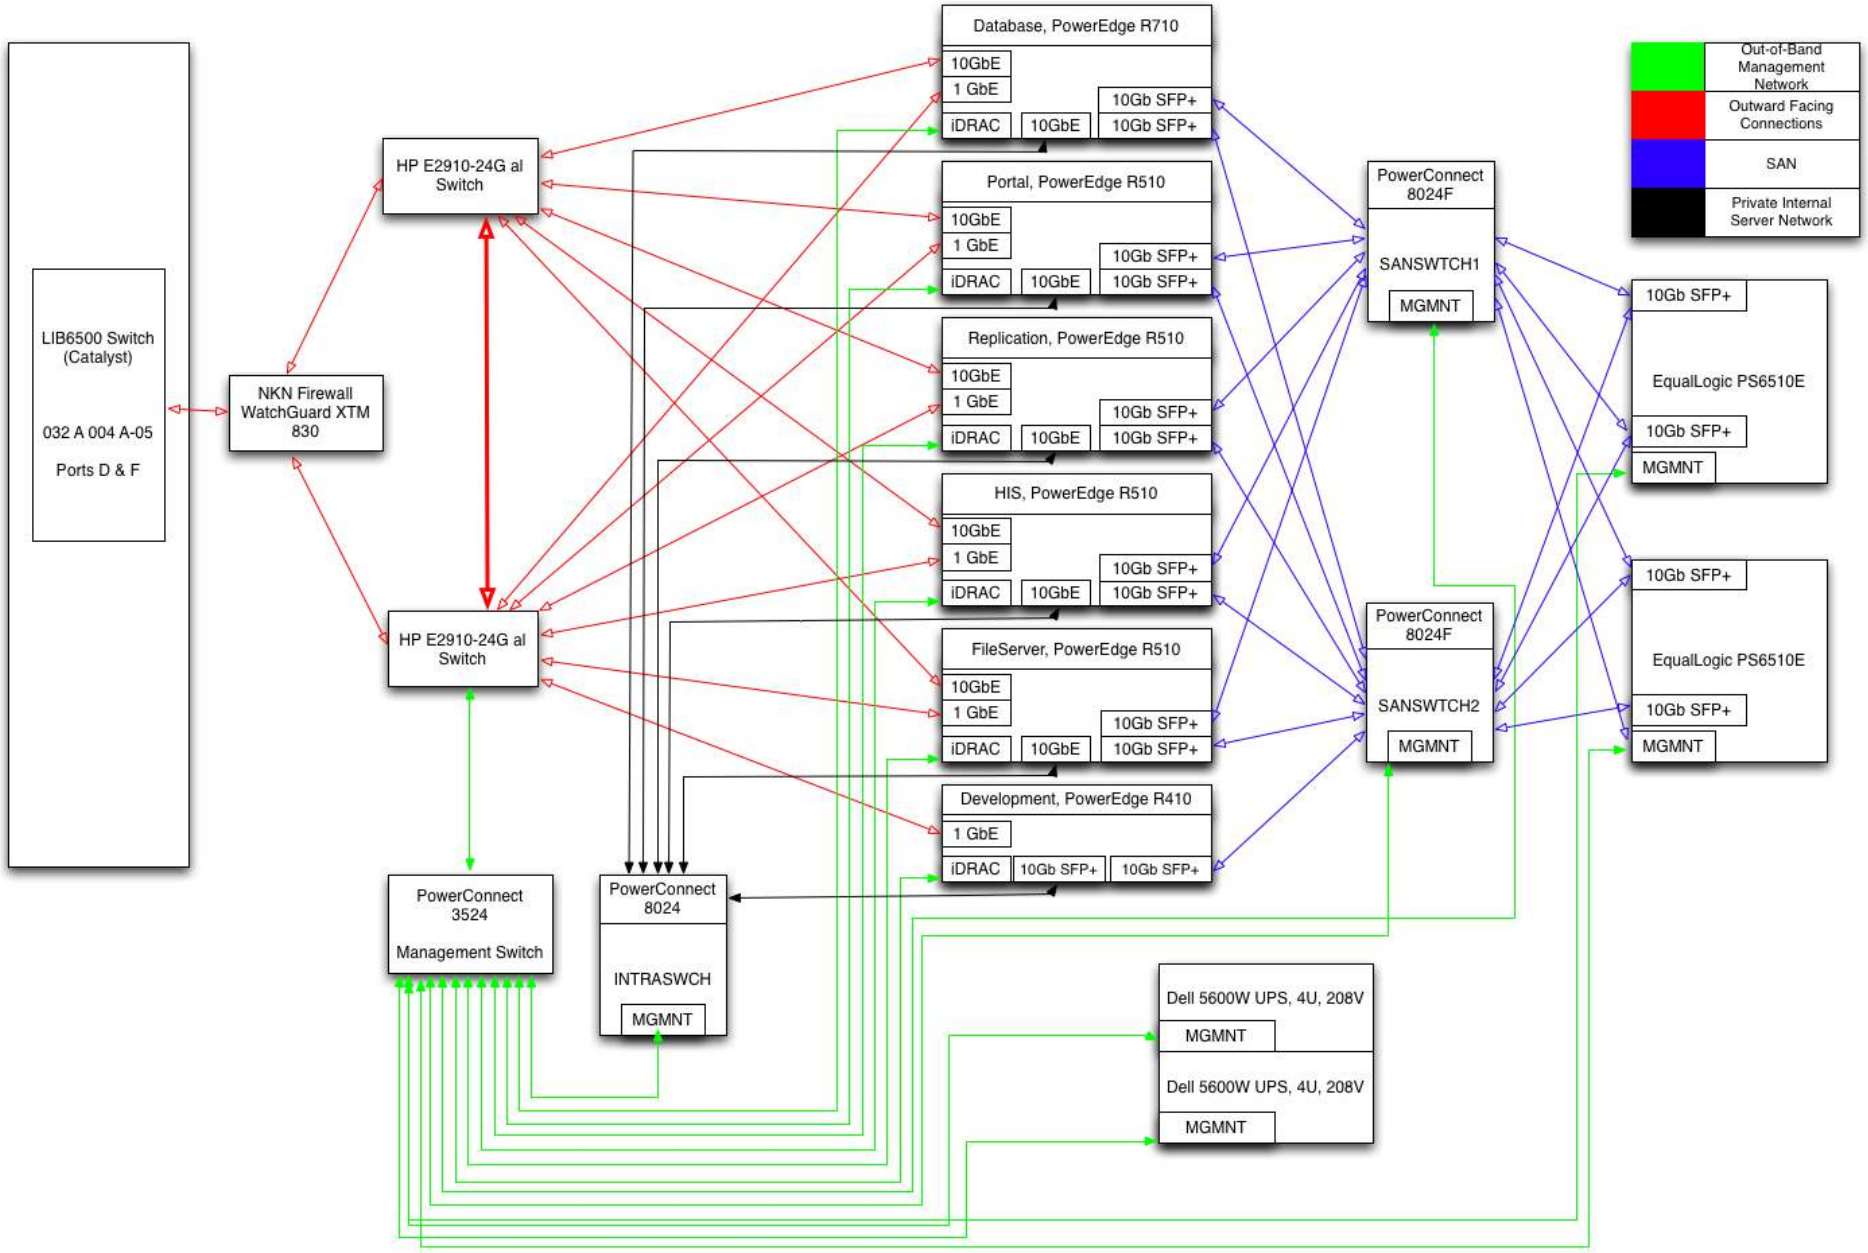
LIB6500 Switch (Catalyst)
032 A 004 A-05
Ports D & F
NKN Firewall WatchGuard XTM 830
HP E2910-24G al Switch
HP E2910-24G al Switch
PowerConnect 3524
Management Switch
PowerConnect 8024
INTRASWCH
MGMNT
Database, PowerEdge R710
10GbE
1 GbE
iDRAC
10GbE
10Gb SFP+
10Gb SFP+
Portal, PowerEdge R510
10GbE
1 GbE
iDRAC
10GbE
10Gb SFP+
10Gb SFP+
Replication, PowerEdge R510
10GbE
1 GbE
iDRAC
10GbE
10Gb SFP+
10Gb SFP+
HIS, PowerEdge R510
10GbE
1 GbE
iDRAC
10GbE
10Gb SFP+
10Gb SFP+
FileServer, PowerEdge R510
10GbE
1 GbE
iDRAC
10GbE
10Gb SFP+
10Gb SFP+
Development, PowerEdge R410
1 GbE
iDRAC
10Gb SFP+
10Gb SFP+
PowerConnect 8024F
SANSWTCH1
MGMNT
PowerConnect 8024F
SANSWTCH2
MGMNT
10Gb SFP+
EqualLogic PS6510E
10Gb SFP+
MGMNT
10Gb SFP+
EqualLogic PS6510E
10Gb SFP+
MGMNT
Dell 5600W UPS, 4U, 208V
MGMNT
Dell 5600W UPS, 4U, 208V
MGMNT
Out-of-Band Management Network
Outward Facing Connections
SAN
Private Internal Server Network

# NKN Datacenters and IRON

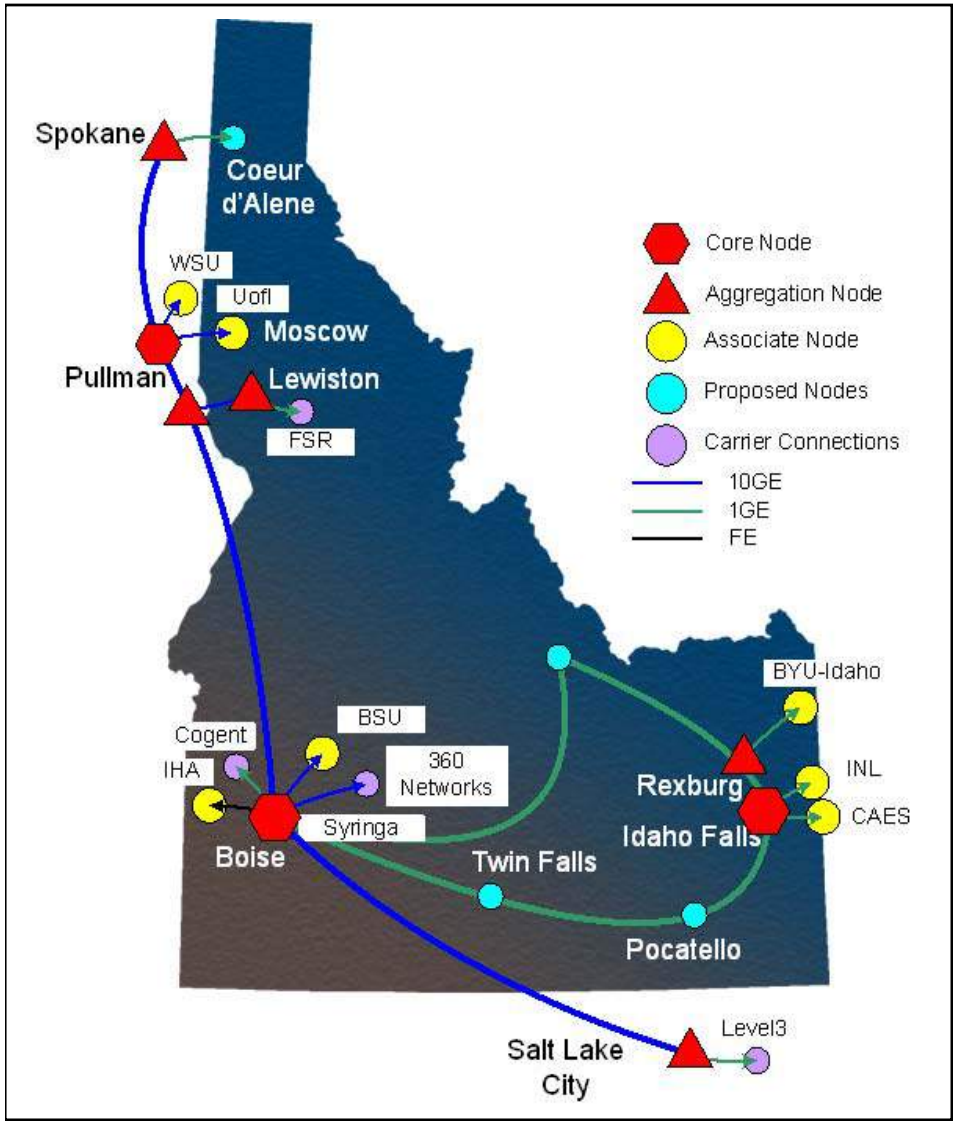


- Geographically diverse, redundant datacenters
- Leverage statewide IRON infrastructure with 10Gbs intra-state backbone
- Excellent connectivity to surrounding regional optical networks
  - Pacific NW GigaPOP
  - Oregon GigaPOP
  - Intermountain GigaPOP
  - Northern Tier Network Consortium
  - Front Range GigaPOP

# High Availability Near-Line iSCSI EqualLogic Block Storage

- 3x Dell EqualLogic PS6510E iSCSI Storage Arrays per datacenter
- 48 spindles per array, 2TB and 3TB disks = nearly 750TB
- RAID 6, redundant controllers and power
- Data automatically load balanced across arrays in group
- Snapshot backups – automatic @ block level

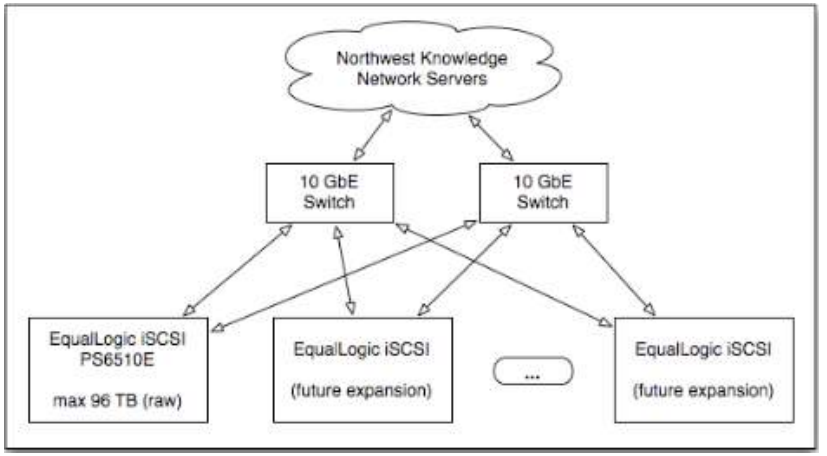


# EqualLogic PS6510E

- **Dell EqualLogic PS6510E**
  - High drive density (48 3TB disks in a 4RU Chassis)
  - 10Gbps NICs – Copper or SFP+
  - iSCSI – Initiators from ESXi Hypervisor or VM level
  - Simple Configuration and Management
  - Linear performance scalability – each chassis has multiple controllers
  - Each Chassis Configured RAID6 with 7200 RPM SATA
  - Scalable to several petabytes within one management group
  - Expansion requires no further networking (SAN Switch) investment
  - Same HW vendor (Dell) – support and integration tools
  - Good VMWare Support (host integration tools, EqualLogic multipath extension modules)
  - Snapshots, thin-provisioning, strong ACL and auth support, monitoring tools

# NKN Data Replication

- Multi-Level Replication
  - Database Level (multi-master)
    - MS SQL Server
    - PostgreSQL / MySQL
  - Filesystem Level
    - GlusterFS (distributed file system/cloud storage)
    - rsync
  - Data Transfer Nodes (DTNs) – ScienceDMZ
    - Globus/GridFTP


# Performance

- Physical Dell PowerEdge Servers
  - In Each NKN Datacenter
    - 4x R510 with 8-12 cores per server, 128GB RAM
    - 1x R710 with 12 cores, 192GB RAM
    - Running VMWare ESXi v5/v6 Hypervisors
    - Multiple guest VMs per Host
    - Multiple 10Gbps copper/twinax NICs per host

- 10Gbps fiber and copper networking
- Dedicated 10Gbps server-to-server LAN (“rocket.net”)

- Firewall: 10Gbps line speed, >20Gbps aggregate

# NKN can move data!

- Joint USGS/NKN data transfer test in February 2012 using Globus Online / GridFTP
- Moved 250GB of USGS GAP data between NKN datacenters (UI -> INL) at 1Gbps (the maximum at the time)
- Totally ***saturated*** one IRON fiber segment between Boise and Idaho Falls
  - Outcome of this:  We have upgraded links now to 10Gbps

# Scalability

- Drop-in storage architecture
  - > 2 PB of storage without additional network equipment.
  - Performance scales linearly
- Drop-in Server design
  - Virtualization with cookie-cutter physical server resources
- All major networking components are 10Gbs ready

# Manageability

- Multi-path remote management
  - Dedicated NKN VPN
  - NKN management VLAN
  - Out-of-band management VLAN
  - Virtual console access via vSphere
  - iDRAC interfaces on PowerEdge Servers
    - Remote Console
    - Power Control
    - Environmentals and Diagnostics
- Centralized syslogs
- vSphere, EqualLogic management clients
- Full RHEL6/7 Entitlements and Update Management

# Rich Internal Linux Infrastructure

- Approx. 55 VMs between UI and INL
- Run primarily RHEL 6.5 – Academic licenses
- Some Windows Server 2008 and Ubuntu
- OpenLDAP on RHEL VM
  - NKN LDAP auth enables: Unix accounts, firewall, web applications (Drupal), databases, SAMBA SMB/ CIFS, ArcGIS, OwnCloud
- NFS & Samba, NTP, SSH, FTP, Jabber (openfire), rsync, Subversion, SMTP (postfix), IMAP (dovecot), OPeNDAP (THREDDS)
- Dozens of websites (Apache httpd, Tomcat, Drupal)
- Databases:  MySQL, PostgreSQL