



National Energy Research Scientific Computing Center (NERSC)

NERSC Site Report

Cary Whitney (clwhitney@lbl.gov)

NERSC Center Division, LBNL

05/24/2004



NERSC Outline

- **Security**
- **Networking**
- **Storage**
- **GUPFS**
- **Other systems**



Security

- **Increase use of sniffed passwords**
 - Network
 - Local via tty
- **Email viruses and worms on the rise**
- **Forged email looking more legit**
- **Reading local email especially to root**



Security 2

- NFS exploits
 - no_root_squash
 - NFS v3
 - NFS exported on non privilege port
 - <http://www.deter.com/unix/software/nfsshell.c>
 - AIX defaulted to insecure
 - Linux defaults to secure



Security 3

- **What to do? Our view.**
 - **Good sys admin practices check each other**
 - **Keep systems patched**
 - **Latest exploits are old problems**
 - **Fan out logins instead of chaining**
 - **Tripwire, rpm verify and similar setup**
 - **.rhosts .authorized_keys .known_hosts .netrc checking**
 - **Change ssh keys and grid certs on compromised systems**



Security 4

- Communicate (A lot of the latest exploits span sites. You may be clean, but the hacker may still have accounts at other sites.)
- Encrypted email about exploits
- Local exploits = remote exploits
- Our contact: security@nersc.gov



Security Projects

- **NERSC LDAP server**
- **One-time password**
 - Proactive response (DOE may mandate)
 - User and admin usage
 - Multi-site specification
 - Oak Ridge, PNNL, Argonne, LBL currently
 - BNL and SLAC also contacted
 - FERMI consulted
 - Keep users from having to carry several keys
- **Contact: sychan@lbl.gov**



Bro

- **IDS which is protocol savvy**
- **Forensic tool (collects connection data)**
- **Started as network analyzer**
- **Grid enabled (gsiftp, gatekeeper, cactus aware)**
 - **Tracks authentication layer**
 - **Records gatekeeper commands**
 - **Gsiftp analyze much like ftp (example, site exec)**



Bro Lite

- **DOE Funded**
- **Simplify Bro**
 - **Configuration (GUI)**
 - **Output filters**

Available: July/Aug

- **Beta slots available**
- **Contact: security@nerosc.gov**



Networking

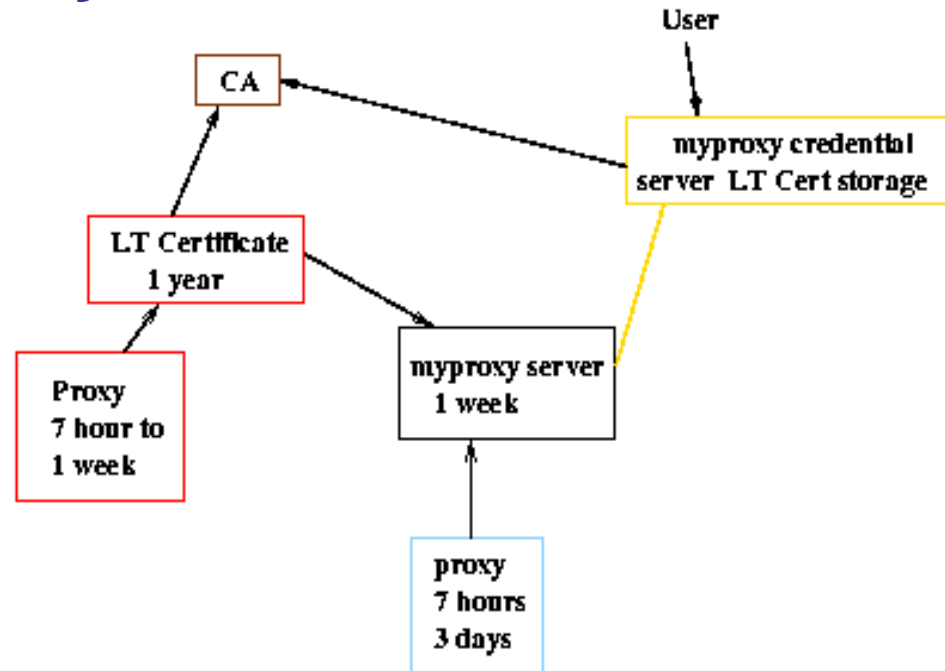
- **Oak Ridge**
 - OC12 to ATM bug cause drop packet thus spoiling single stream performance
- **BNL**
 - Cisco PIX: Unsetting tcp_window_scaling fixed via upgrade version ~6.2




Networking 2

- **Wisdom**
 - 90% host tuning issues
 - 8% first hop issues
 - Rest is in the WAN
- **Every problem is magnified with latency**
- **Multiple streams poor substitute to overcome network problems**
- **Passive taps are good at critical points (~\$500)**
- **Drop counters lie!**
Contact: net@nersc.gov

- Grid myproxy

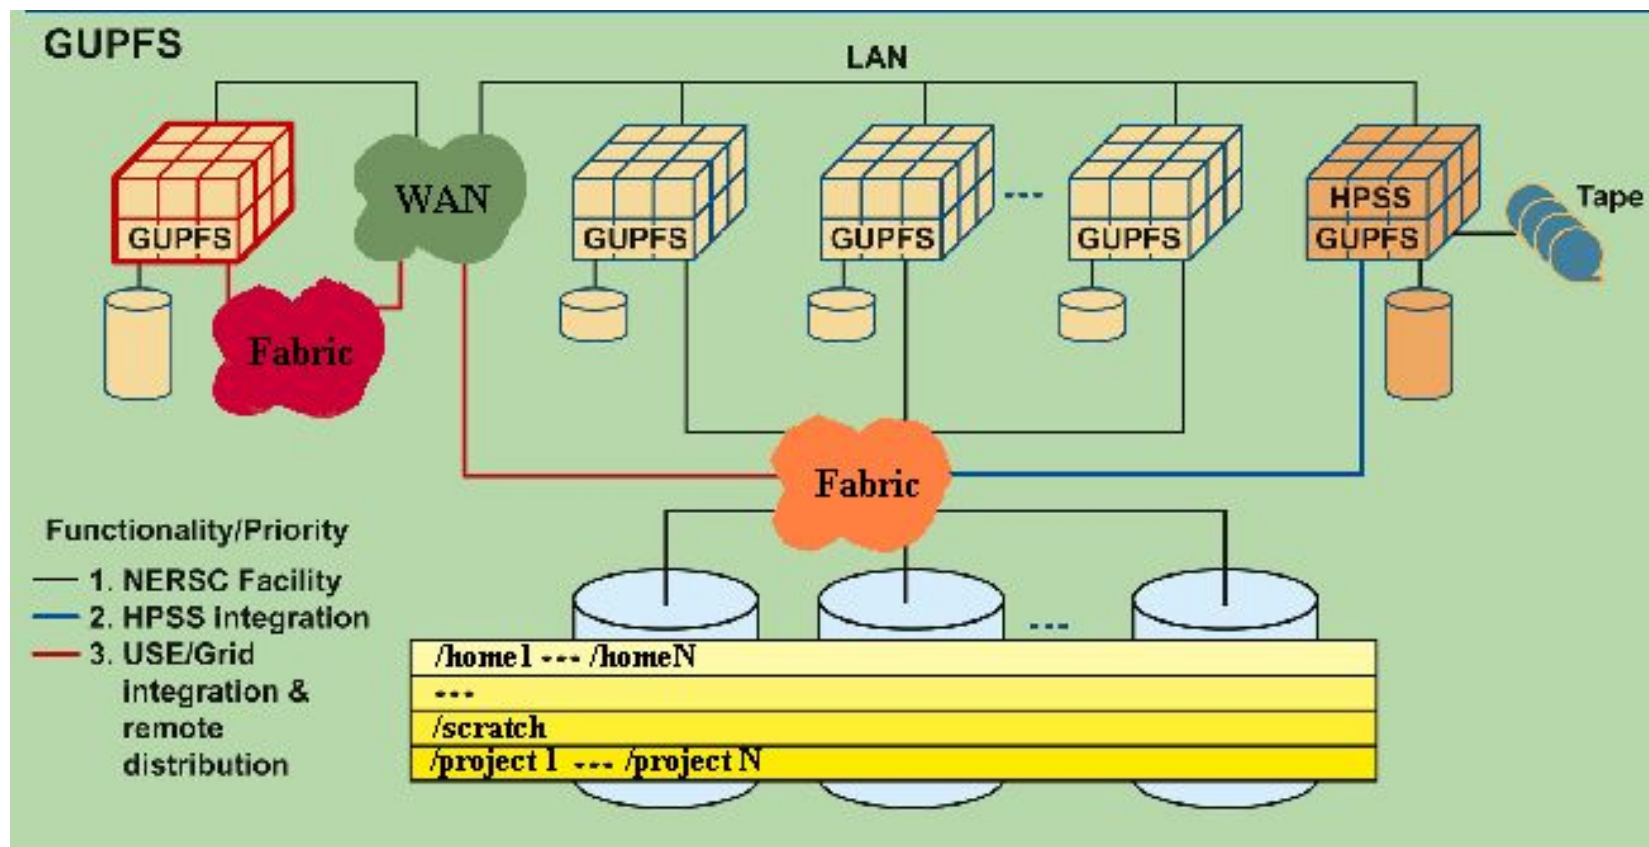


Current non myproxy
With myproxy
Near term myproxy path



- **Problem**
 - Still no way to revoke a proxy cert
 - OCSP – Online Certificate Status Protocol
- **Contact: sychan@lbl.gov**

GUPFS Layout





GUPFS Tested

- **File Systems**

- Sistina GFS 4.2, 5.0, 5.1, and 5.2 Beta
- ADIC StorNext File System 2.0 and 2.2
- Lustre 0.6 (1.0 Beta 1), 0.9.2, 1.0, 1.0.{1,2,3,4}, 1.2.1
- IBM GPFS for Linux, 1.3 and 2.2
- Panasas

- **Fabric**

- FC (1Gb/s and 2Gb/s): Brocade SilkWorm, Qlogic SANbox2, Cisco MDS 9509, SANDial Shadow 14000
- Ethernet (iSCSI): Cisco SN 5428, Intel & Adaptec iSCSI HBA, Adaptec TOE, Cisco MDS 9509
- Infiniband (1x and 4x): InfiniCon and Topspin IB to GE/FC bridges (SRP over IB, iSCSI over IB),
- Inter-connect: Myrinnet 2000 (Rev D)

- **Storage**

Traditional Storage: Dot Hill, Silicon Gear, Chaparral

New Storage: Yotta Yotta GSX 2400, EMC CX 600, 3PAR, DDN S2A 8500



GUPFS Observations

General

- Many vendors are involved and have various products available. But are more focused on commercial market.
- Need standardized, centralized monitoring and management for fabric and storage.

Filesystem

- Progress but still need more work, remains the component with the highest risk
 - Stability
 - Parallel I/O performance and functionality
 - Scalability
- Not enough filesystems can or plan to support many platforms/OS's.
 - ADIC SNFS greatest variety now
 - IBM Storage Tank has many now and more planned
 - Lustre plans to support OS-X
 - IBM GPFS supports AIX and Linux



GUPFS

- **Planned deployment late 2005**
- **Unified filesystem spanning all NERSC systems (NCS, Seaborg, PDSF)**
- **Possible candidates**
 - **GPFS, ADIC, Lustre, Panasas, Storage Tank**
- **Results:**
<http://www.nersc.gov/projects/GUPFS>
- **Contact: gupfs@nersc.gov**



Storage

- **htar**
 - HSI tar utility – creates/retrieves tar archive on hpss.
- **Portal/web access to HPSS**
- **GPFS DMAPI**
 - Linux
 - AIX