



LBNL/NERSC/PDSF
Site Report for HEPiX

Catania, Italy

April 17, 2002

by

Cary Whitney

clwhitney@lbl.gov



LBL/NERSC/PDSF



Who we are

Staffing changes

Software

Hardware

Challenges

What's New



LBL/NERSC/PDSF



National Energy Research Scientific Computing (NERSC)

Largest non-classified compute facility in U.S.

NERSC is a division inside of Lawrence Berkeley National Lab (LBNL).

PDSF is a group within the computational systems group (CSG)

PDSF serves mostly HEP users.

PDSF is a general purpose facility



Staffing



1 FTE Project Lead/Admin

2 FTE Admin

.5 FTE User services support/Nuclear Physics Support
(USG)

Changes here: Steven Chan has left the group lead position to do grid coordination for NERSC and Shane Canon has taken over group lead responsibilities.



Software



Mirroring of both the Star and Atlas environments

This is probably one of the most important pieces of software. It has mostly removed our reliance of AFS to BNL and other sites.



Software - Linux



Linux kernel **v2.2.19**

Linux kernel 2.4 investigation for disk vault usage

Large swap space - 2 GB on each Linux box.

Using Kernel based NFS servers

Openssh v2 for everything when possible

NERSC mandates no clear text passwords to be used where possible; telnet to be totally disabled by Jan 1, 2000

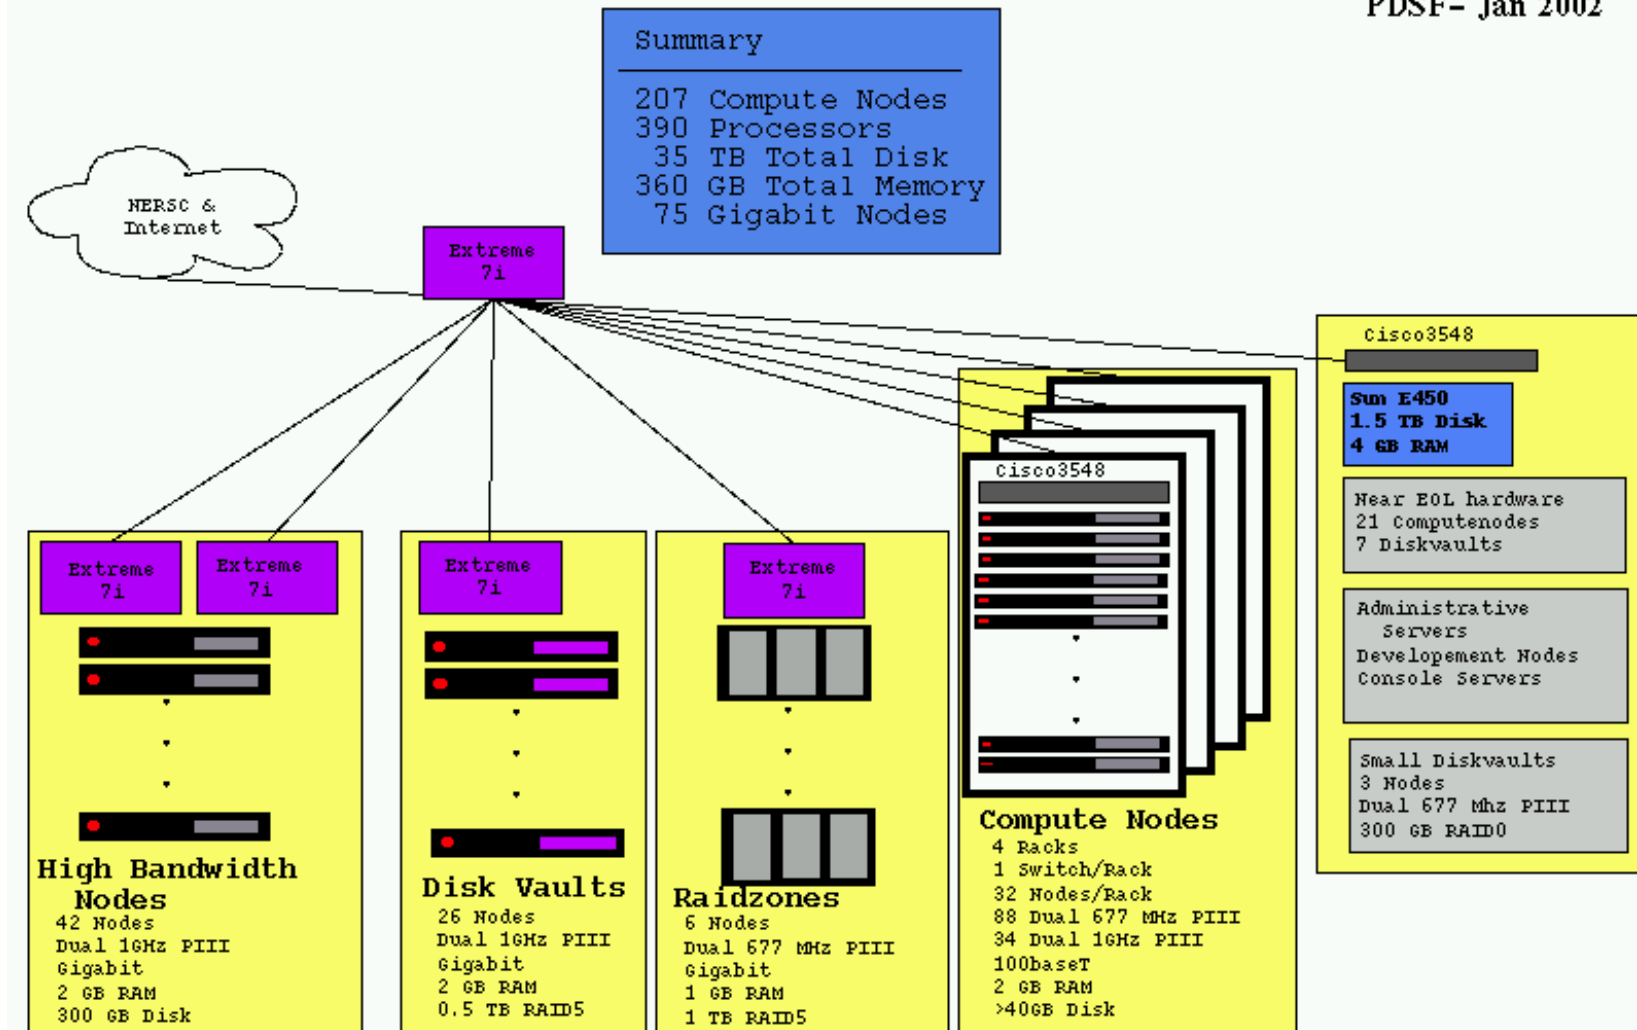
Now we have disabled even clear text ftp access.

AFS via NFS only. Uses Large AFS cache on **pdsfsu05** (5gb).

OpenAFS installed on some interactive nodes to test performance problems with AFS->NFS gateway problems

PDSF Layout

PDSF- Jan 2002



Challenges for PDSF

HENP Computing Challenges

Experiment	Data	Compute
E895 (AGS)	10 TB/yr	600 SPECint95
BaBar (SLAC)	400 TB/yr	5,000 SPECint95
STAR (RHIC)	266 TB/yr	10,100 SPECint95
PHENIX (RHIC)	700 TB/yr	8,500 SPECint95
D0 Run II (FNAL)	280 TB/yr	4,075 SPECint95
CDF Run II (FNAL)	464 TB/yr	3,650 SPECint95
ATLAS (LHC)	1100 TB/yr	20,000 SPECint95



New to NERSC



Grid

PDSF has had grid resources available to over a year now, but NERSC is now working with IBM to get grid working with the SP2

Future Technology Group (FTG)

Working on Linux Check-point restart

Infiniband technologies

Advance Systems Group (ASG)

Global Unified File System (GUFS)

GPFS for Intel



NERSC Cont



Networking

Advancements in bro and bro monitoring

Mass Storage Group

Linux movers