Site report: Wisconsin

Sridhara Dasu
Dan Bradley, Ajit Mohapatra, Will Maier
{dasu,dan,ajit,wcmaier}@hep.wisc.edu

University of Wisconsin - High Energy Physics

USCMS T2 Workshop - LIGO Livingston, 2009.03.03

1 Overview

2 2008-2009 Facilities Status

- Software
- Batch
- Storage
- Network

3 2009 Deployment Plans

- Software
- Batch and storage systems
- Network

4 Facilities Planning

- Cooling and Power

5 How we try to do it

- Emphasize local analysis
- Redundancy

Component	Software
OS	Scientific Linux 4.4
Batch	Condor 7.0.5
Storage	dCache 1.9.0-8
Grid	VDT 1.10.1r

Table: 2009 Software Status

- OSG updates have become much more reliable, though we still can't relocate our installations
- Significant performance improvements (especially scaling) in recent Condor
- dCache running well, but replication is still a major problem

CPU Class	KSI2K	Slots
2 x 2.80 GHz Xeon	210	110
2 x 3.00 GHz Xeon	10	10
4 x 1.80 GHz Opteron	170	240
8 x 2.66 GHz Xeon	180	120
8 x 3.00 GHz Xeon	250	260
Dedicated	820	740
Opportunistic	-	2460

Table: 2009 Batch Status

Since last workshop:

- Added 32 8 x 3.00 GHz Xeon nodes
- Decommissioned 32 2 x 2.4 GHz Xeon nodes

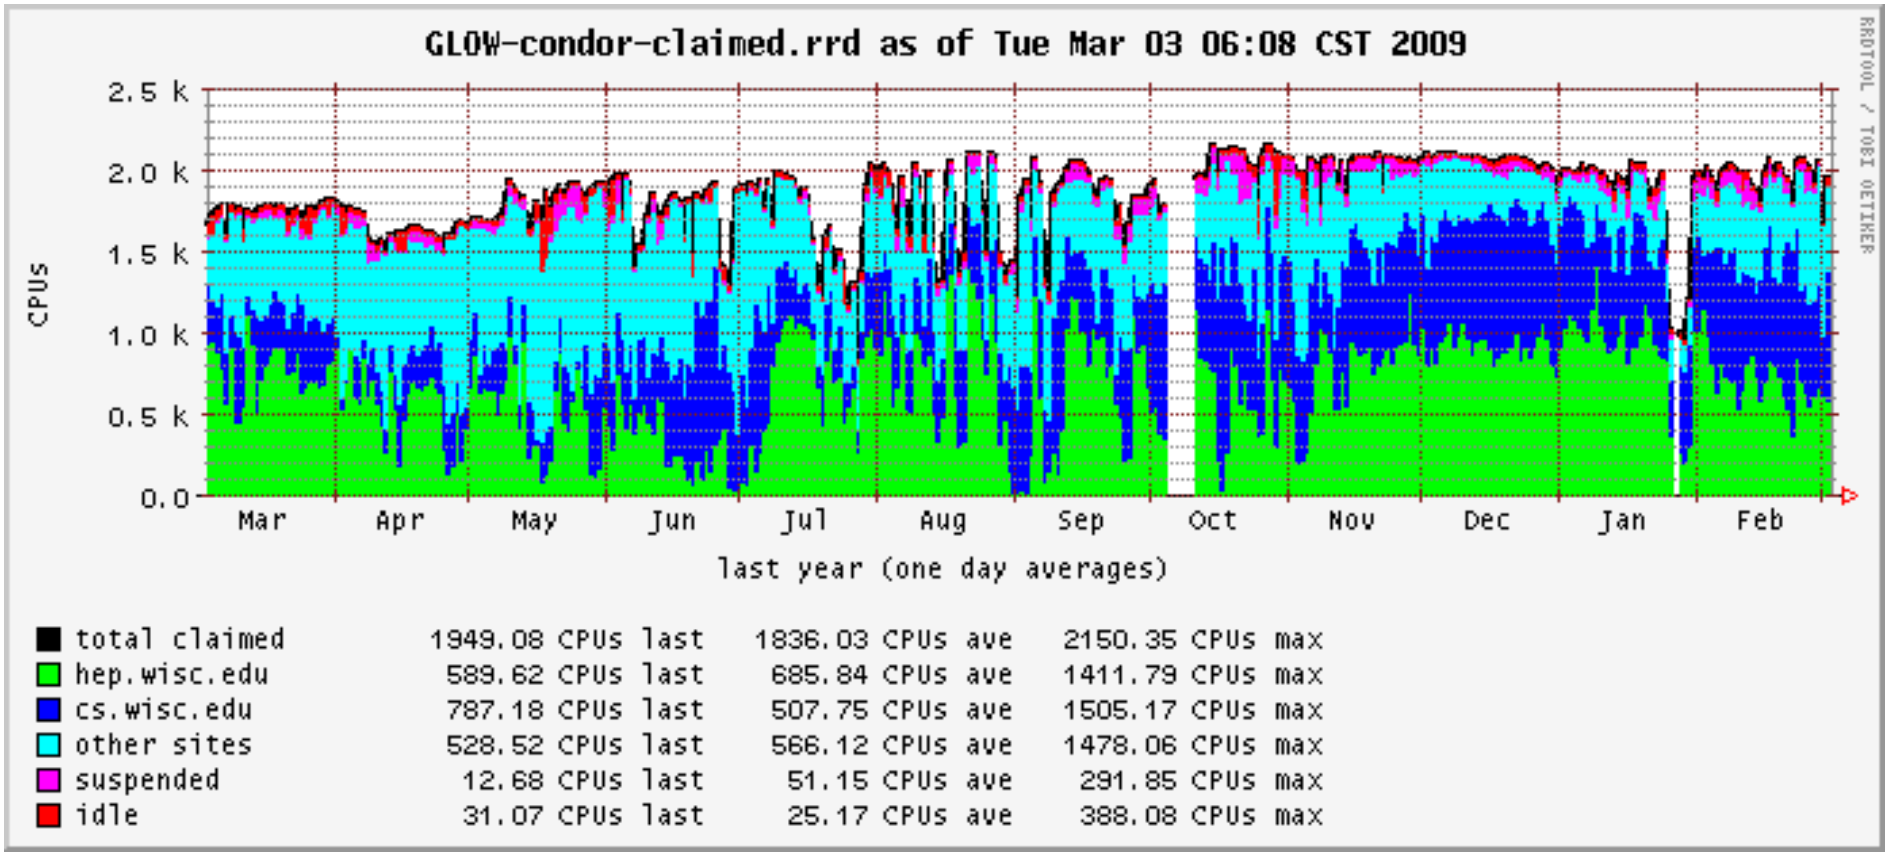

Figure: Dedicated and opportunistic slots, 2008-2009

Dedicated	280 TB
Dual-purpose	200 TB
Raw	480 TB
Replicated	240 TB

Table: 2009 Storage Status

Since last workshop:

- Added ∼300 TB in dedicated and dual-purpose machines
- Switched from dCache ReplicaManager to PFM

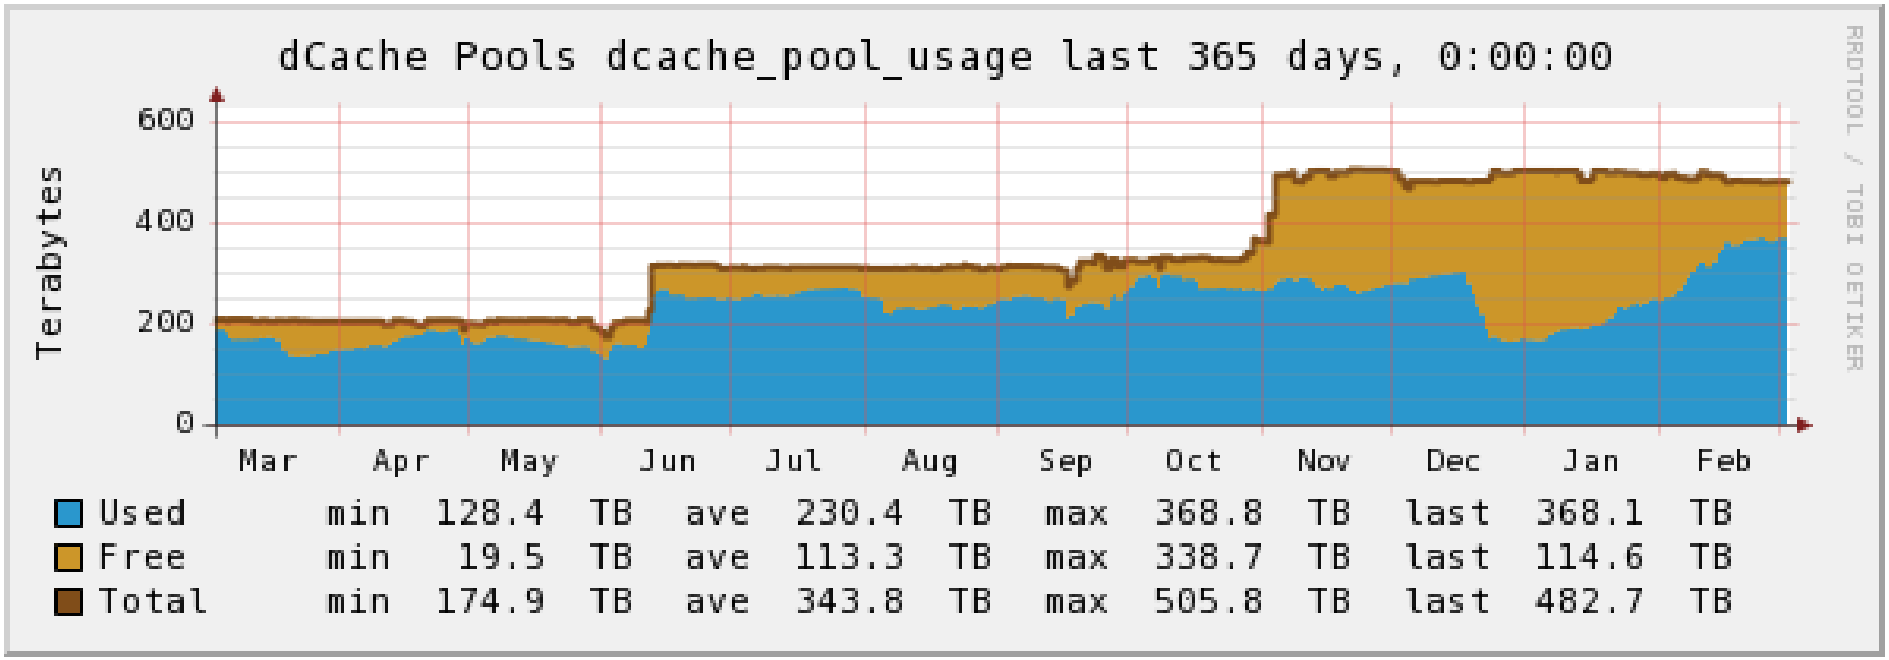

Figure: dCache usage, 2008-2009

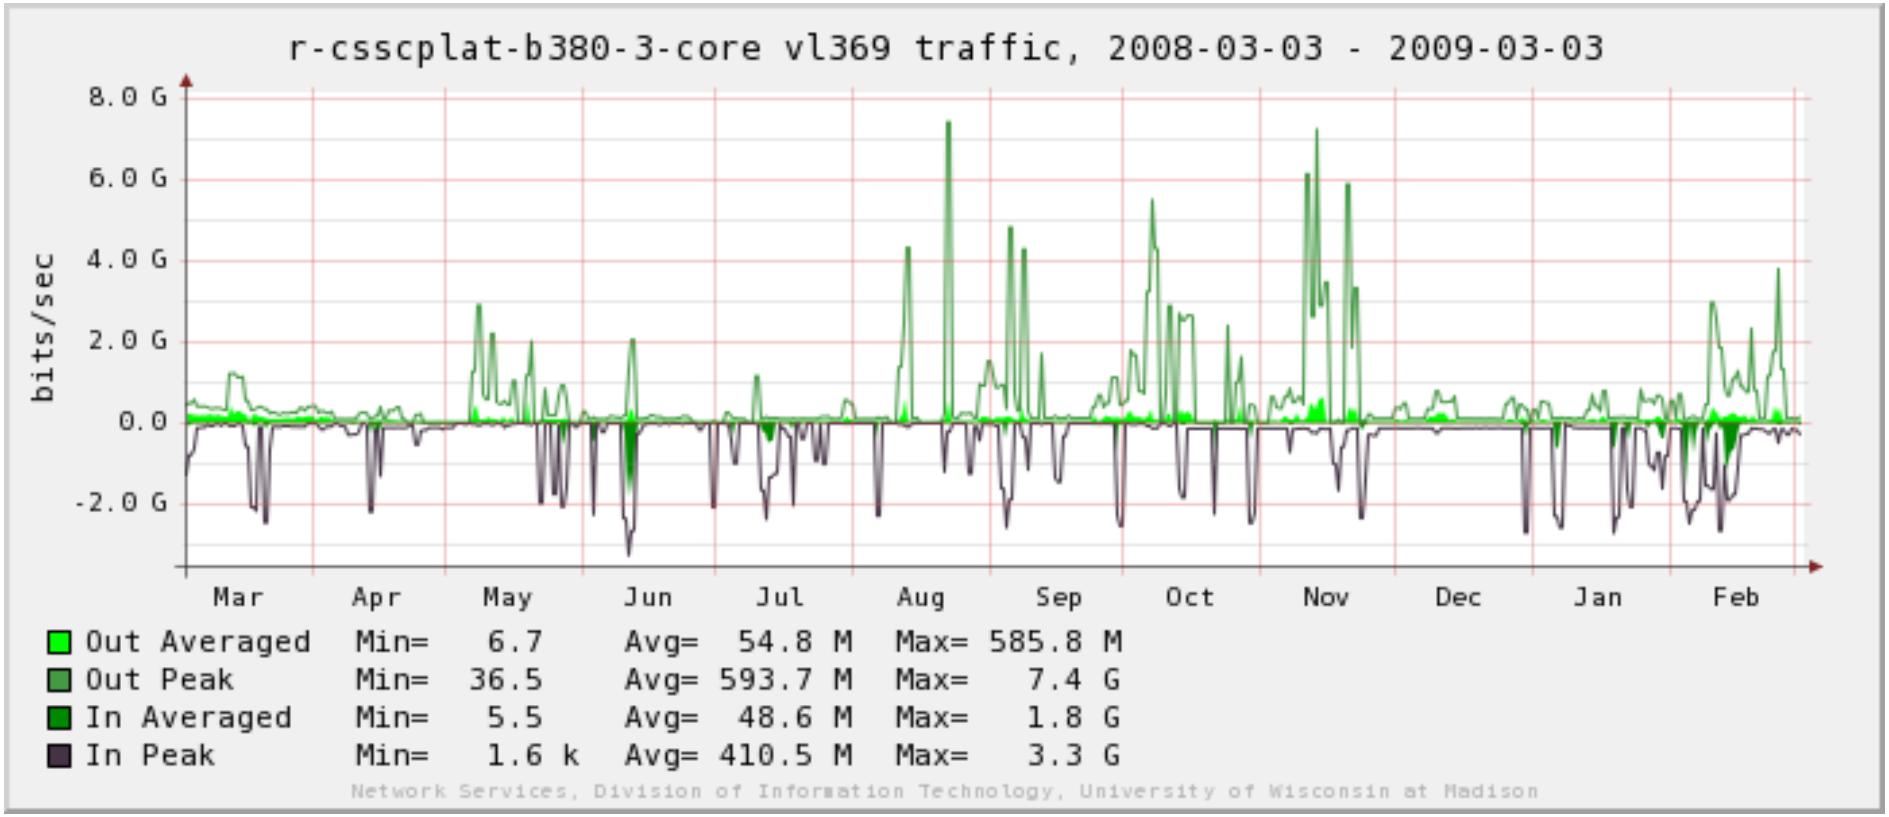

Figure: WAN usage, 2008-2009

- Tier2 traffic accounts for the majority of egress traffic from Wisconsin
- Ingress peaks due to PhEDEx transfers

March-April	Complete Scientific Linux 5.2 testing and begin deployment
May-June	Test BeStMan, dCache'?
July	Upgrade to dCache feature branch, Chimera *OR* Begin switch to dCache'

Table: 2009 Software Deployment Plan

Month	Storage (TB)	Slots
March	100	200
June	300	-
September	100	200
December	300	-
New in 2009	800	400
Total in 2009	1280	1200

Table: 2009 Hardware Deployment Plan

- $100kUSD budget for each round of alternating purchases:
  - Dedicated: 8 x 2.66 GHz Xeon, 16 GB RAM, 24 x 1 TB disk, LSI RAID
  - Dual-purpose: 8 x 3 GHz Xeon, 16 GB RAM, 4 x 1 TB disk

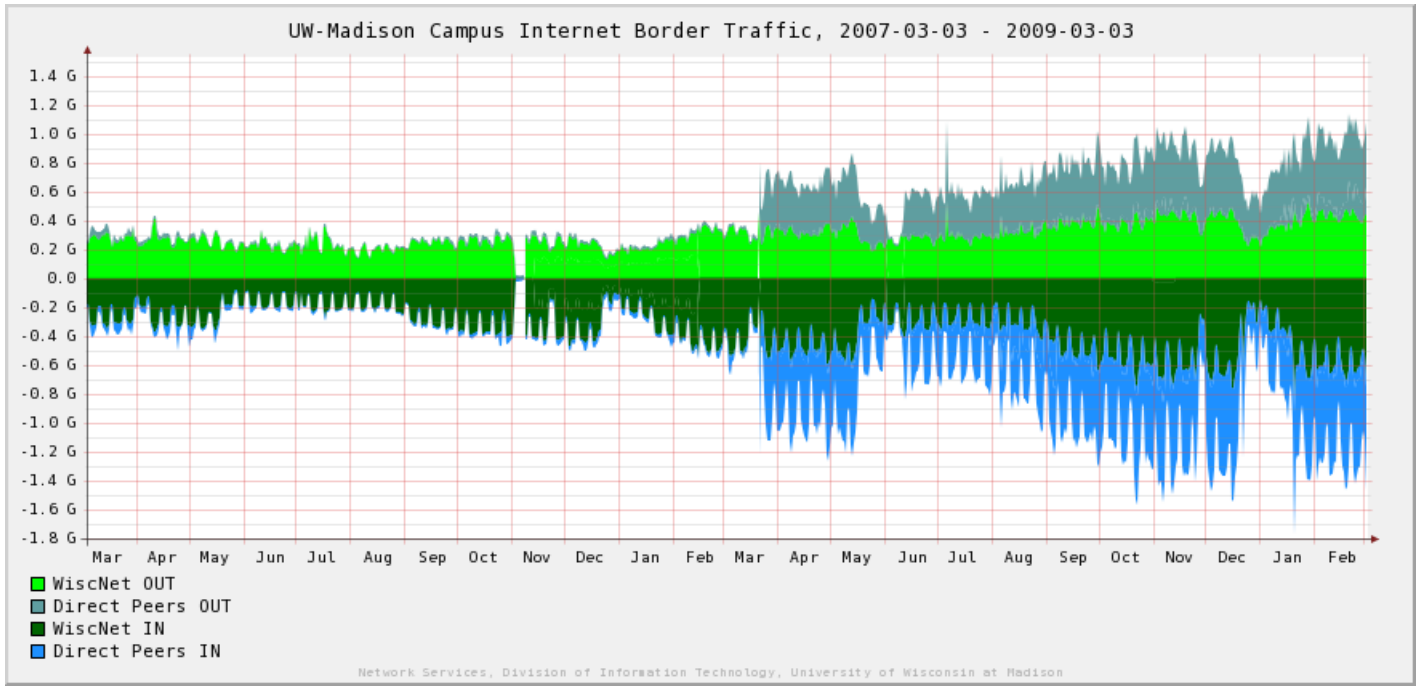

Figure: Wisconsin campus border traffic, 2007-2009

- Dynamically provision extra fiber dedicated to Tier2
- Teach PhEDEx agents to request new fiber before large transfers?

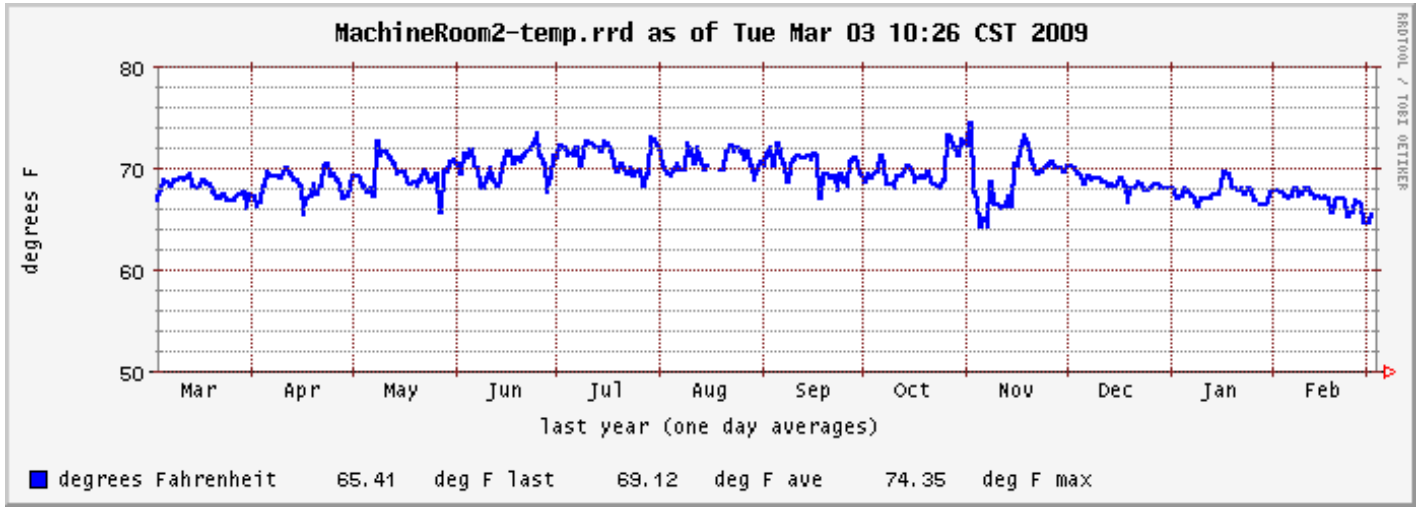

Figure: Machine room temperature (F), 2008-2009

- At cooling capacity in new room; at power capacity in old room
- New room with shared space and power upgrade in old room forthcoming

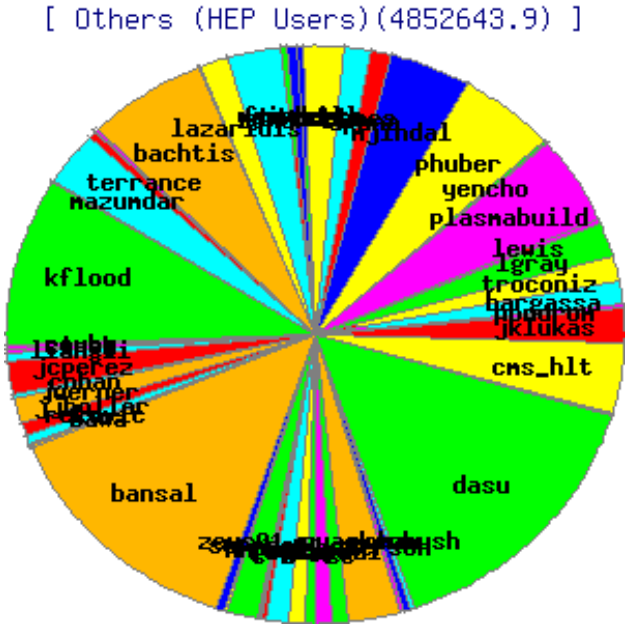

Figure: CPU hours by local username, 2005-2009

- 140 local users, ∼50 running now
- Developed local scripts for quick analysis with central registration (Dan Bradley)
- Treat global production as if it were local (Ajit Mohapatra)

- Aggressive replication of all data (PFM) on small pools and commodity hardware
- Many GridFTP doors
- Dedicated central servers for dCache and GUMS (with hot spares)
- Multiple CEs with easy fail over
- Monitor the monitors: Nagios, central monitoring (JobRobot, SAM, RSV, et al)