

RHIC Computing Facility

T. Throwe, BNL
(Deputy Head of RHIC Computing)



DOE RHIC Program Review

July 9-11 , 2003

RHIC Computing Facility (RCF) Mission

- ✍ Formally Initiated In 1997
- ✍ Supply Computing Infrastructure for RHIC Exper.
 - Including code development, repository, & distribution
- ✍ Supply resources & systems for the production processing of RHIC data
 - Raw data recording
 - Reconstruction
 - Data mining
 - Analysis
- ✍ Production simulation *not* part of the RCF mission

Context

- ✍ The RHIC & ATLAS Computing Facilities (RCF & ACF) are currently co-located and co-operated on the raised floor area of Bldg 515, ITD
- ✍ Fully integrated support staff of 25.5 FTE's
 - 20 FTE's from RHIC
 - 4.5 FTE's from ATLAS
 - 1 FTE's DOE MICS/Data Grid
- ✍ ACF is US Tier 1 Regional Center for ATLAS
- ✍ Emphasis here will be RCF, mission as described
- ✍ Strong synergy between RCF & ACF in the form of
 - Shared expertise and operational support
 - Shared infrastructure components

RCF/ACF

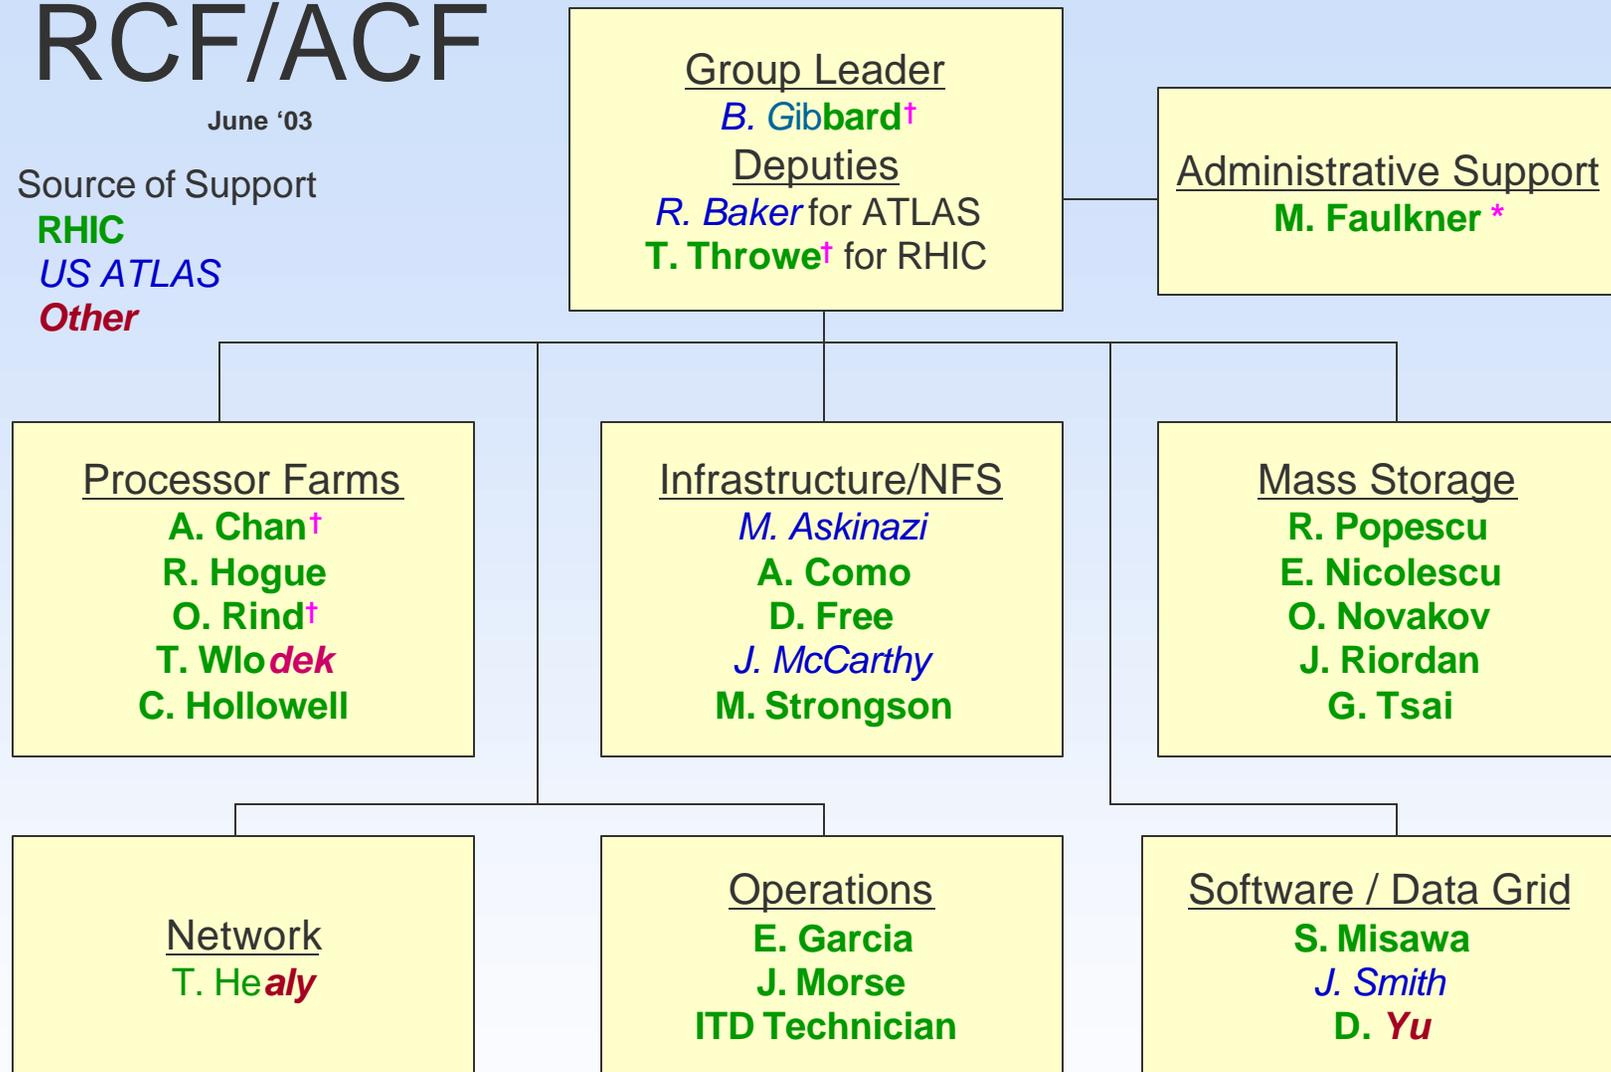
June '03

Source of Support

RHIC

US ATLAS

Other



† Scientific staff

* Administrative staff

All others, IT professional staff

10 July 2003

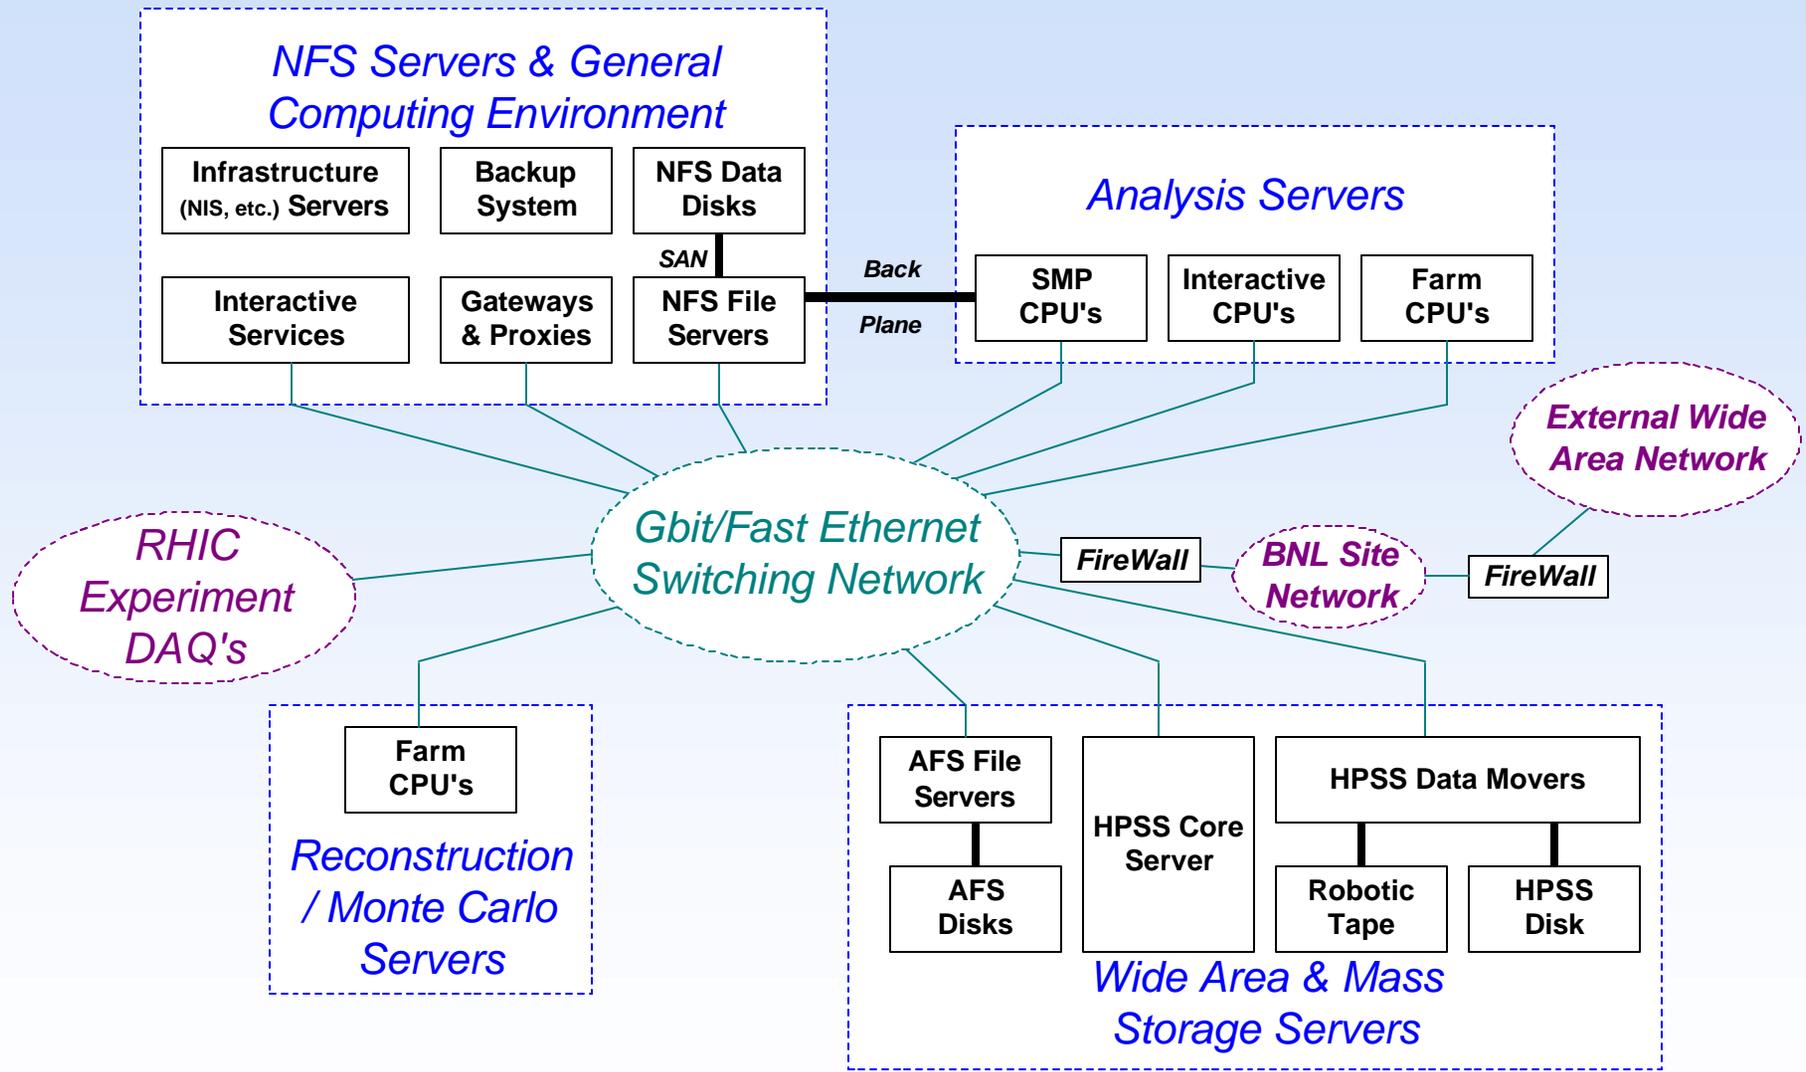
T. Throwe

RHIC Program Review
July 9-11, 2003

Summary For Past Year

- ✍ RCF in a Production Mode
- ✍ Provided Storage and Processing for RHIC Run-3
- ✍ Implemented RCF Upgrades

RCF/ACF Schematic



Mass Storage Subsystem

✍ Hierarchical Storage Management by *HPSS*

✍ 4 StorageTek robotic tape silos

- ~24,000 tape cartridges
- 4.5 PBytes

✍ Tape drives

- 1000 MB/sec



10 July 2003

T. Throwe

RHIC Program Review
July 9-11, 2003



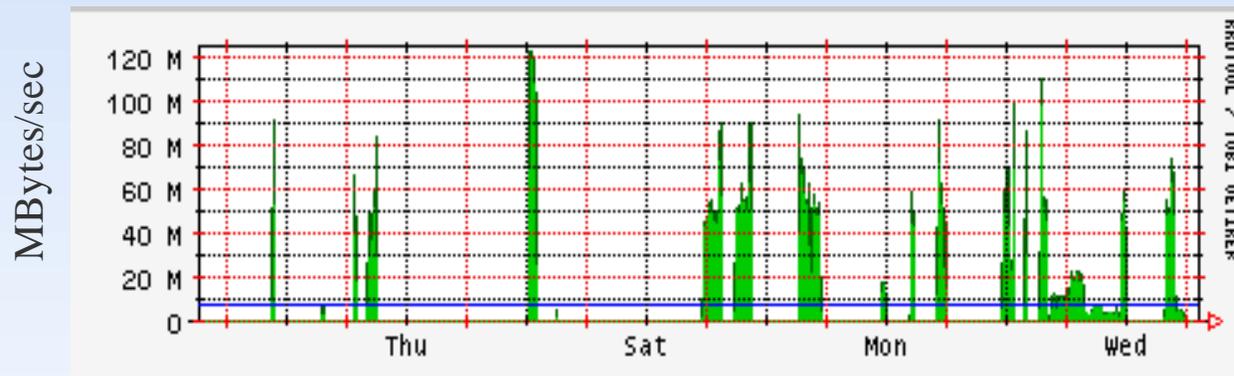
Mass Storage (cont.)

RHIC Run-3

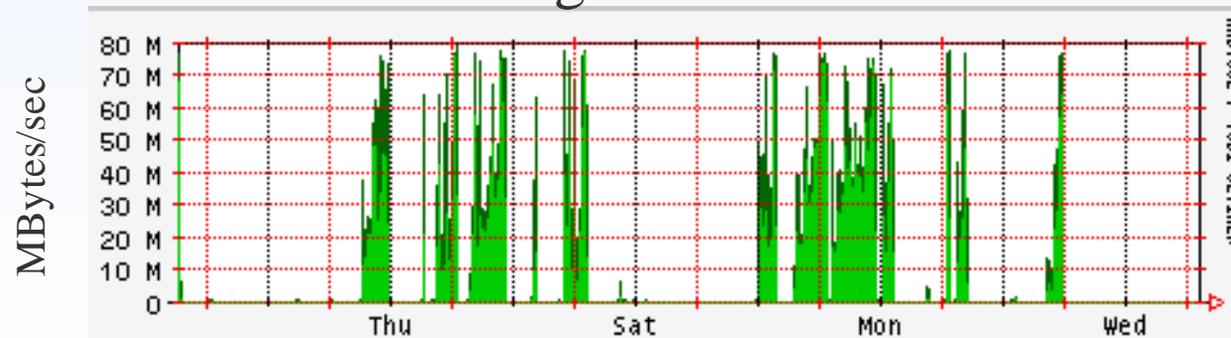
- Raw data recording at rates up to 120 MByte/sec
- RAW data volume was 180 TB
- Total of ~236 TB to date including derived data, but still in the early stages of reconstruction and analysis.
- System ran well

Mass Storage (cont.)

PHENIX Data Sinking – May 7 to May 14



STAR Data Sinking – March 5 to March 12



10 July 2003

T. Throwe

RHIC Program Review
July 9-11, 2003

Processor Farms Subsystem

- ✍ Hardware is soft partitioned by experiment according to allocation
- ✍ CRS Farms have restricted access and locally produced control software
- ✍ CAS Farms have interactive access and batch queues managed by LSF
- ✍ Hardware
 - ~2050 CPU's
 - ... total capacity ~100 kSPECint95 (2+ TFLOPS)



10 July 2003

T. Throwe

RHIC Program Review
July 9-11, 2003



Processor Farms (cont.)

For Run-3

- Approximately 335,000 jobs processed in CRS Farm to date
- Approximately 3,353,000 jobs processed through LSF in the CAS Farm.

RAID Disk Subsystem

FibreChannel Storage Area Network

- Brocade Switches 1x64 + 2x16
- ~30 RAID controllers 2-3 GBytes/sec

118 TBytes of RAID 5 storage



10 July 2003

T. Throwe

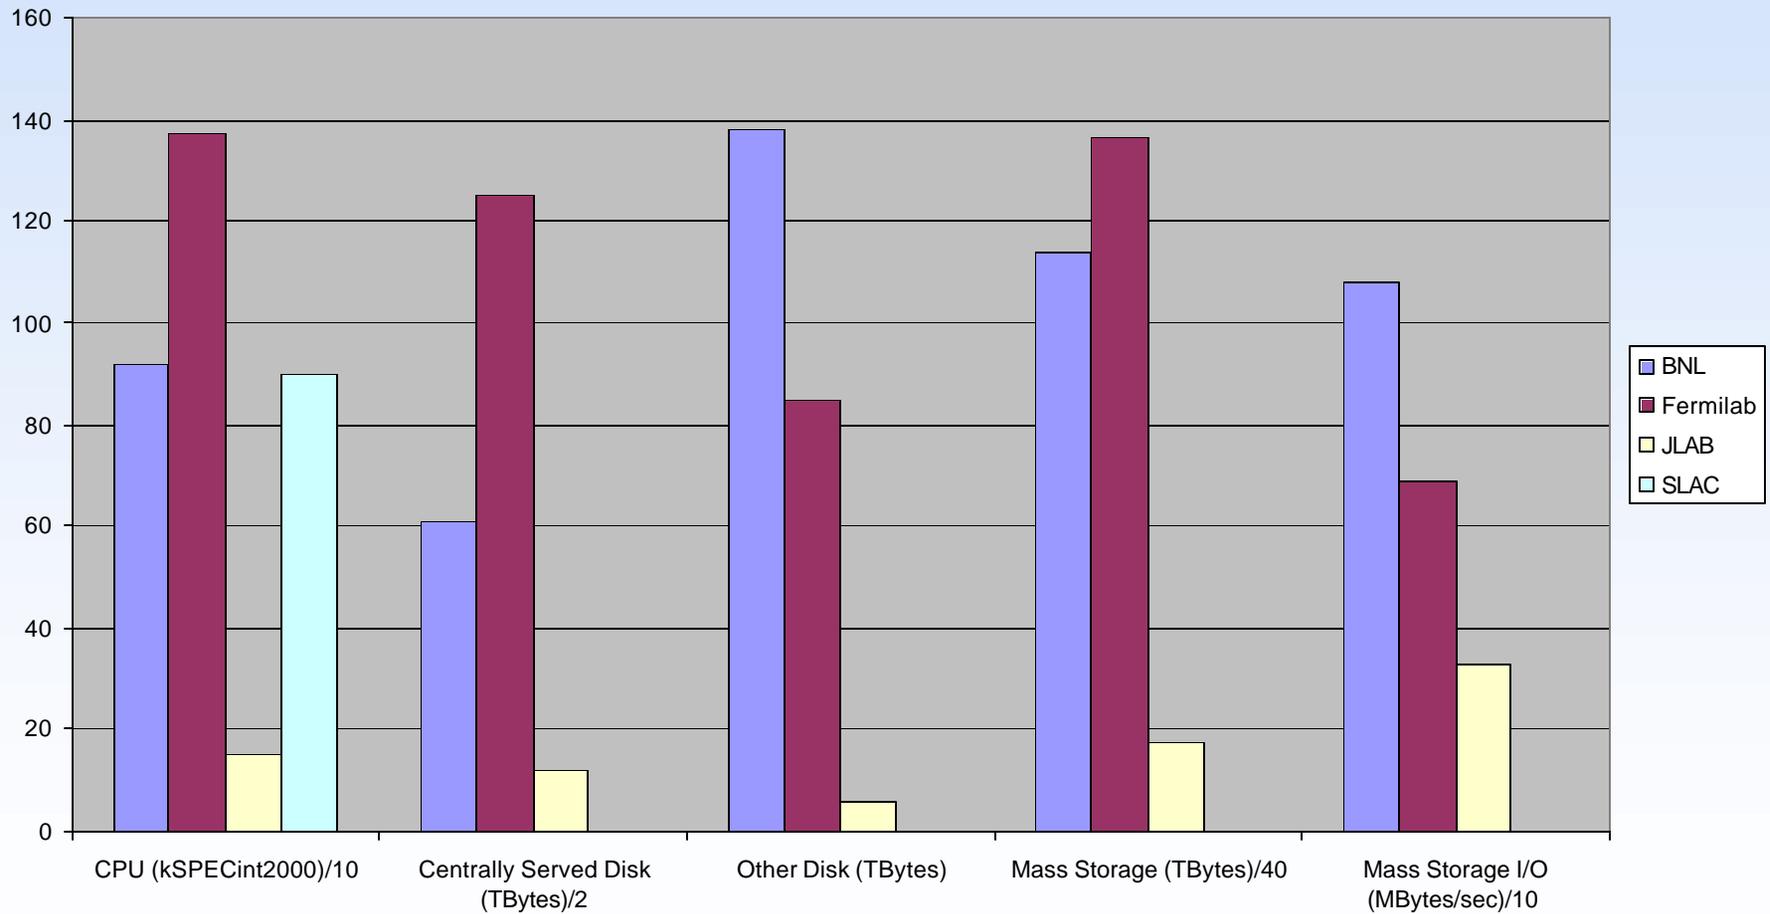
RHIC Program Review
July 9-11, 2003



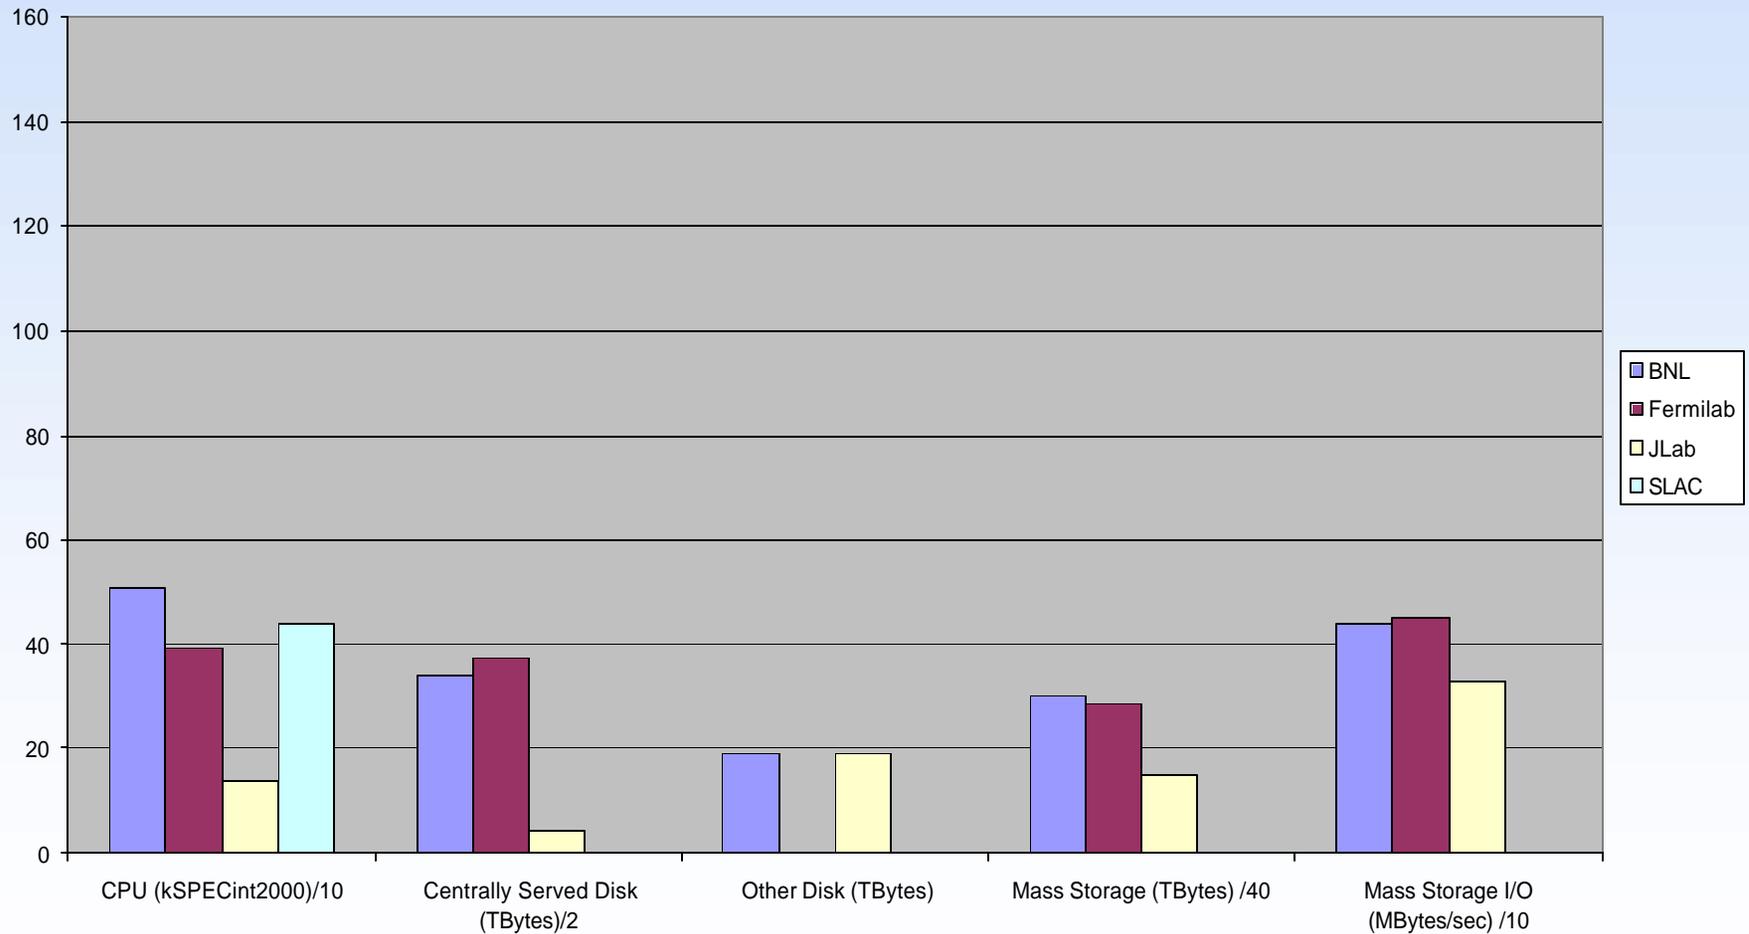
Capital Procurement Plan

- ✍ Majority of FY'03 funds committed at the beginning of the year to prepare for Run-3
- ✍ Example rate plots show we could have sustained substantially higher rates expected during the run
- ✍ Detailed plan for FY'04 not yet available – we will begin next round of “requirements” meetings soon
- ✍ FY'04 will see significant fraction of funds going into replacing old equipment

Computing Capacity Comparisons (Apr '03)



Computing Capacity Comparisons (Jul '02)



Staffing

✍ Comparisons of computing staffs (Fermilab has comparable mission)

Fermilab('02)	SLAC('02)	RHIC('03)
70	16.5	20.0

✍ Report of DOE Facility Operations Review of RHIC, Feb. 5-7, 2002:

“The RHIC computing facility is the one area that does immediately need increased resources. We recommend the addition of ~3 FTE to ensure that the ever-growing analysis and logging capacities are available to the users in a timely fashion.”

✍ RCF/US ATLAS Computing Facilities Effort By Funding Source

Funding Source	Initial RCF Plan	FY 2002 Actual	FY 2003 Actual	FY 2004 Plan
NP Operating *	18	19.5	20	22
ITD/BNL Contrib.	8	0	0	0
RHIC Experiment Contrib.	8	D	D	D
TOTAL RHIC	34	19.5	20	22
HEP/US ATLAS		4.5	4.5	6.5
MICS/Data Grid		0.5	1	0.5
TOTAL	34	24.5	25.5	29
	D	Not well defined, probably of order 2 FTE's		
	*	Includes 1 FTE of hardware support from ITD		

RCF Staffing Issues

- ✍ RCF staff still lean
- ✍ While so far RCF has functioned well,
 - Availability and reliability ... has been good
 - Capacities ... have been more than adequate...once again, lower than anticipated demands during Run-3 resulted in less stress on the facility
- ✍ There have been, and continue to be, significant areas receiving inadequate attention

Where Staffing is Needed

Large Scale Online Storage (RAID Disk)

- Performance/reliability tuning of existing systems
- Integration of new acquisitions
- Technology tracking (NAS, IDE RAID)

CPU Farms

- Increased 3rd party product support
- Evolution of cluster management to increase flexibility
- OS & Kernel evolution
- System image export to desk tops etc.

Wide Area Data Distribution

- Tape copy
- WAN tuning
- Direct secure remote HPSS access
- Data Grid deployment

Facility Cyber Security

- ITD, peer site, and Data Grid compatibility

RHIC Computing Facility Summary

- ✍ RCF has performed well during past runs
- ✍ It is in position to effectively support analysis of data in hand
- ✍ Success to date is, in part, a result of moderate data pressure & judicious application of very dedicated but limited personnel
- ✍ February 2002 DOE Operations Review of RHIC recommended an increase of three FTE's, which should help in a major way, but will still fall short by 1.5 FTE of the 24 FTE's we estimated the facility requires