Overview of ATLAS Computing Facility

Michael Ernst

DOE Annual HEP Program Review

16 - 18 April 2007

office of high energy physics

BROOKHAVEN NATIONAL LABORATORY

ATLAS Computing at BNL

- Principal objectives
  - Fulfill role as principal U.S. center (Tier-1) in the tiered ATLAS (and LHC) computing model
    - Supply capacity to ATLAS as agreed in the MOU (~23%)
    - Guarantee the capability and capacity needed by US ATLAS physics program
  - Establish a critical mass in computing and software to support ATLAS physics analysis at BNL and elsewhere in the U.S.
  - Contribute to ATLAS software most critical to enabling physics analysis at BNL and the U.S.
  - Provide to U.S. ATLAS physicists expertise and facilities and strengthen U.S. ATLAS physics
  - Leverage projects outside ATLAS, in particular Open Science Grid (OSG), where they can strengthen the BNL and U.S. ATLAS programs

U.S. in the ATLAS Computing Model

- BNL Tier-1 responsibilities (~10 Tier-1's total)
  - Archival shares of raw and reconstructed data, and associated calibration & reprocessing
  - Store and serve 100% of ATLAS reconstruction (ESD), analysis (AOD) and physics tag data (TAG)
  - Physics group level managed production/analysis
  - Resources dedicated to U.S. physicists: additional per-physicist capacity at 50% of the level managed centrally by ATLAS
    - Allow ~ X 2 acceleration of analysis of one 20% data stream
    - Actual allocation of these resource will be done by U.S. ATLAS Resource Allocation Committee
  - U.S. Tier-2's have a complimentary role
    - Bulk of simulation and end user analysis support
    - Store and serve 100% of AOD, TAG and subset of ESD
  - Tier-1 and Tier-2's both support institutional and individual users
    - Primarily end user analysis

Revision of Required Capacities

- Many adjustments in requirement estimates during past year
  - Change in LHC startup schedule
  - Identification and removal of Heavy Ion contribution from HEP requirement
  - Inclusion of efficiency factors for resources utilization
    - Chaotic CPU – 60%
    - Programmatic CPU – 85%
    - Disk – 70%
  - Adjustment in author count based calculation of contributions to Tier 1 and Tier-2's to reflect that CERN contributes Tier 0 and CAF but not otherwise
    - US 20% (~actual US author fraction) ⟶ 23% (US fraction of non-CERN authors)

ATLAS (and LHC Comp. Grid) Commissioning Schedule

Experiments

- Continued testing of computing models, basic services
- Testing the full data flow
  DAQ→Tier-0→Tier-1←→Tier-2
- Building up end-user analysis support

ATLAS Full Dress Rehearsal (FDR)
Exercising the computing systems, ramping up job rates, data management performance, …

1st April is the target to have required services in place to prepare for Dress Rehearsals!

ATLAS has already started preparations on 26th February

2007

2008

Services

Introduce residual services
File Transfer Services for T1-T2 traffic
Distributed Database Synchronization
Storage Resource Manager v2.2
Virtual Org. Management Services (VOMS) roles in site scheduling

Commissioning the service for the 2007 run– increase performance, reliability, capacity to target levels, monitoring tools, 24 x 7 operation, ….

01Jul07 - service commissioned - aiming at full 2007 capacity

first LHC collisions

LHCC Review Conclusions - Paul Dauncey

- A lot of progress since the last Comprehensive Review
  - Jamie Shiers: ***"Despite the problems encountered – and those yet to be faced and resolved – I believe that it is correct to say we have a usable service (not a perfect one)"***
- Several critical components still be to deployed
  - Without disrupting services
  - With a somewhat uncertain schedule
- Service problems seen are amorphous and not easy to categorise
  - Many one-offs, so progress will be slow in fixing them

- **The bottom line is that we do have a service – need to build on this and steadily improve it…**

What is the Full Dress Rehearsal (FDR) ?

A complete exercise of the full chain from Trigger to Distributed Analysis, to be performed in Summer 2007, a few months before data taking starts

- Generate O(10^7) events: few days of data taking, ~1pb^{-1} at L=10^{31}
- Mix and filter events to get correct physics mixture as expected at HLT output
- Pass events through G4 simulation ("as installed misaligned distorted" detector geometry
- Run LVL1 simulation
- Produce byte streams => emulate raw data format
- Send raw data to Point 1, pass through HLT nodes and SFO, write out events into streams, closing files at boundary of luminosity blocks
- Send events from Point 1 to Tier-0; manipulate/merge files according to final model
- Perform calibration & alignment at Tier-0 (and possibly also outside)
- Run reconstruction at Tier-0 (and maybe Tier-1s) => produce ESD, AOD, TAGs, DPDs
- Distribute ESD, AOD, TAGs, DPDs to Tier-1s and Tier-2s; replicate databases
- Perform distributed analysis, use TAGs, produce additional group-specific DPDs, etc.
- Run Data Quality Assessment at all levels of data production

Steps BNL Tier-1 Center is involved

FDR - Timescale & Duration

- Timescale?
  - Between July and November 2007
- One problem is that the FDR competes for resources with the ongoing ATLAS detector commissioning which by then will be reaching its final stages
  - Primarily in the TDAQ, Tier-0 and Data Quality Monitoring
  - Will require careful scheduling
- Series of 1 week “runs” separated by 2-3 weeks of analysis and preparation for the next one?
  - Would allow for 3-4 runs prior to low energy running
- These runs will provide good testbeds to exercise the ATLAS global shift operations infrastructure
  - In fact these might be the schedule driver to have this in place
  - Procedures in place to ensure all shift slots are covered, on-call rotas, etc.

ATLAS Data Flow

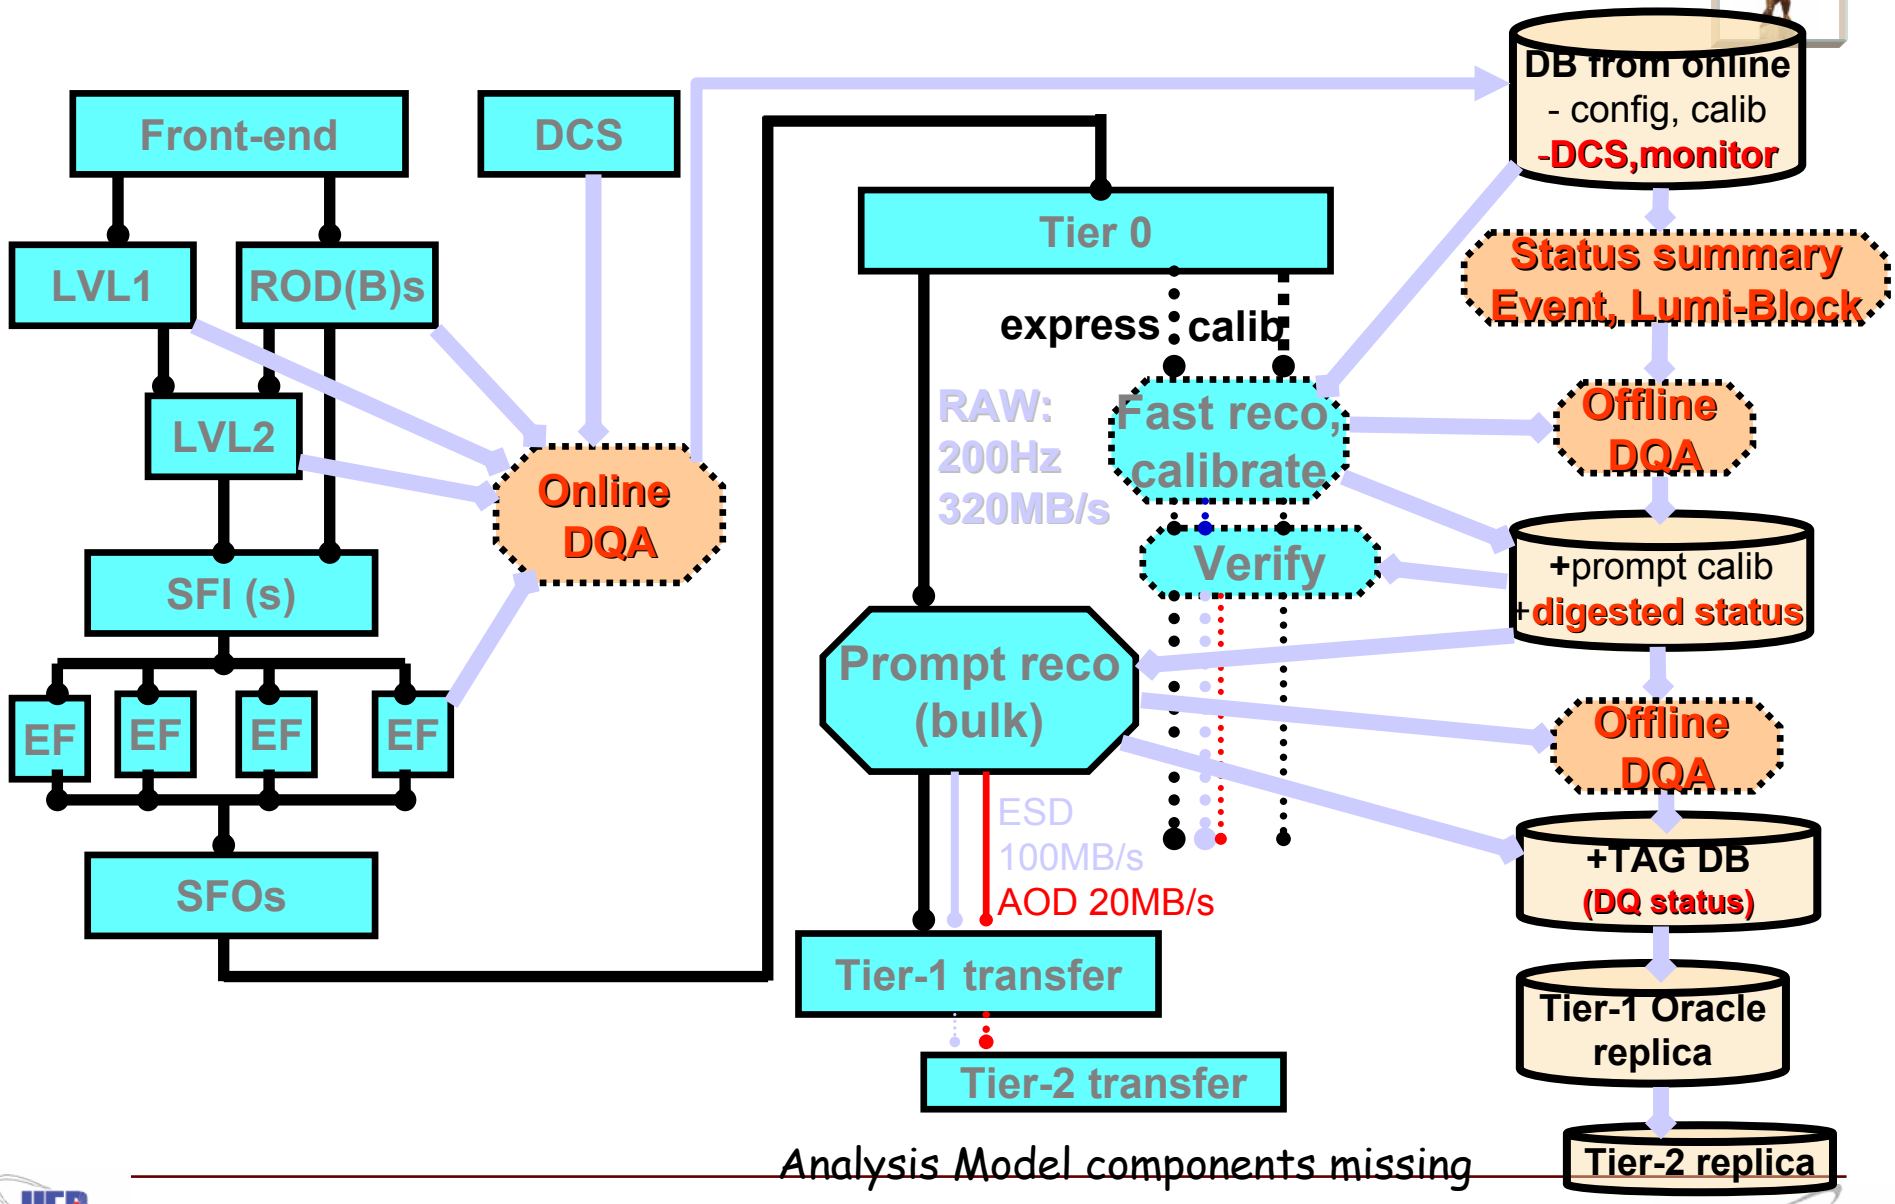

Analysis Model components missing

U.S. ATLAS Tier-1 Data Flow (2008)

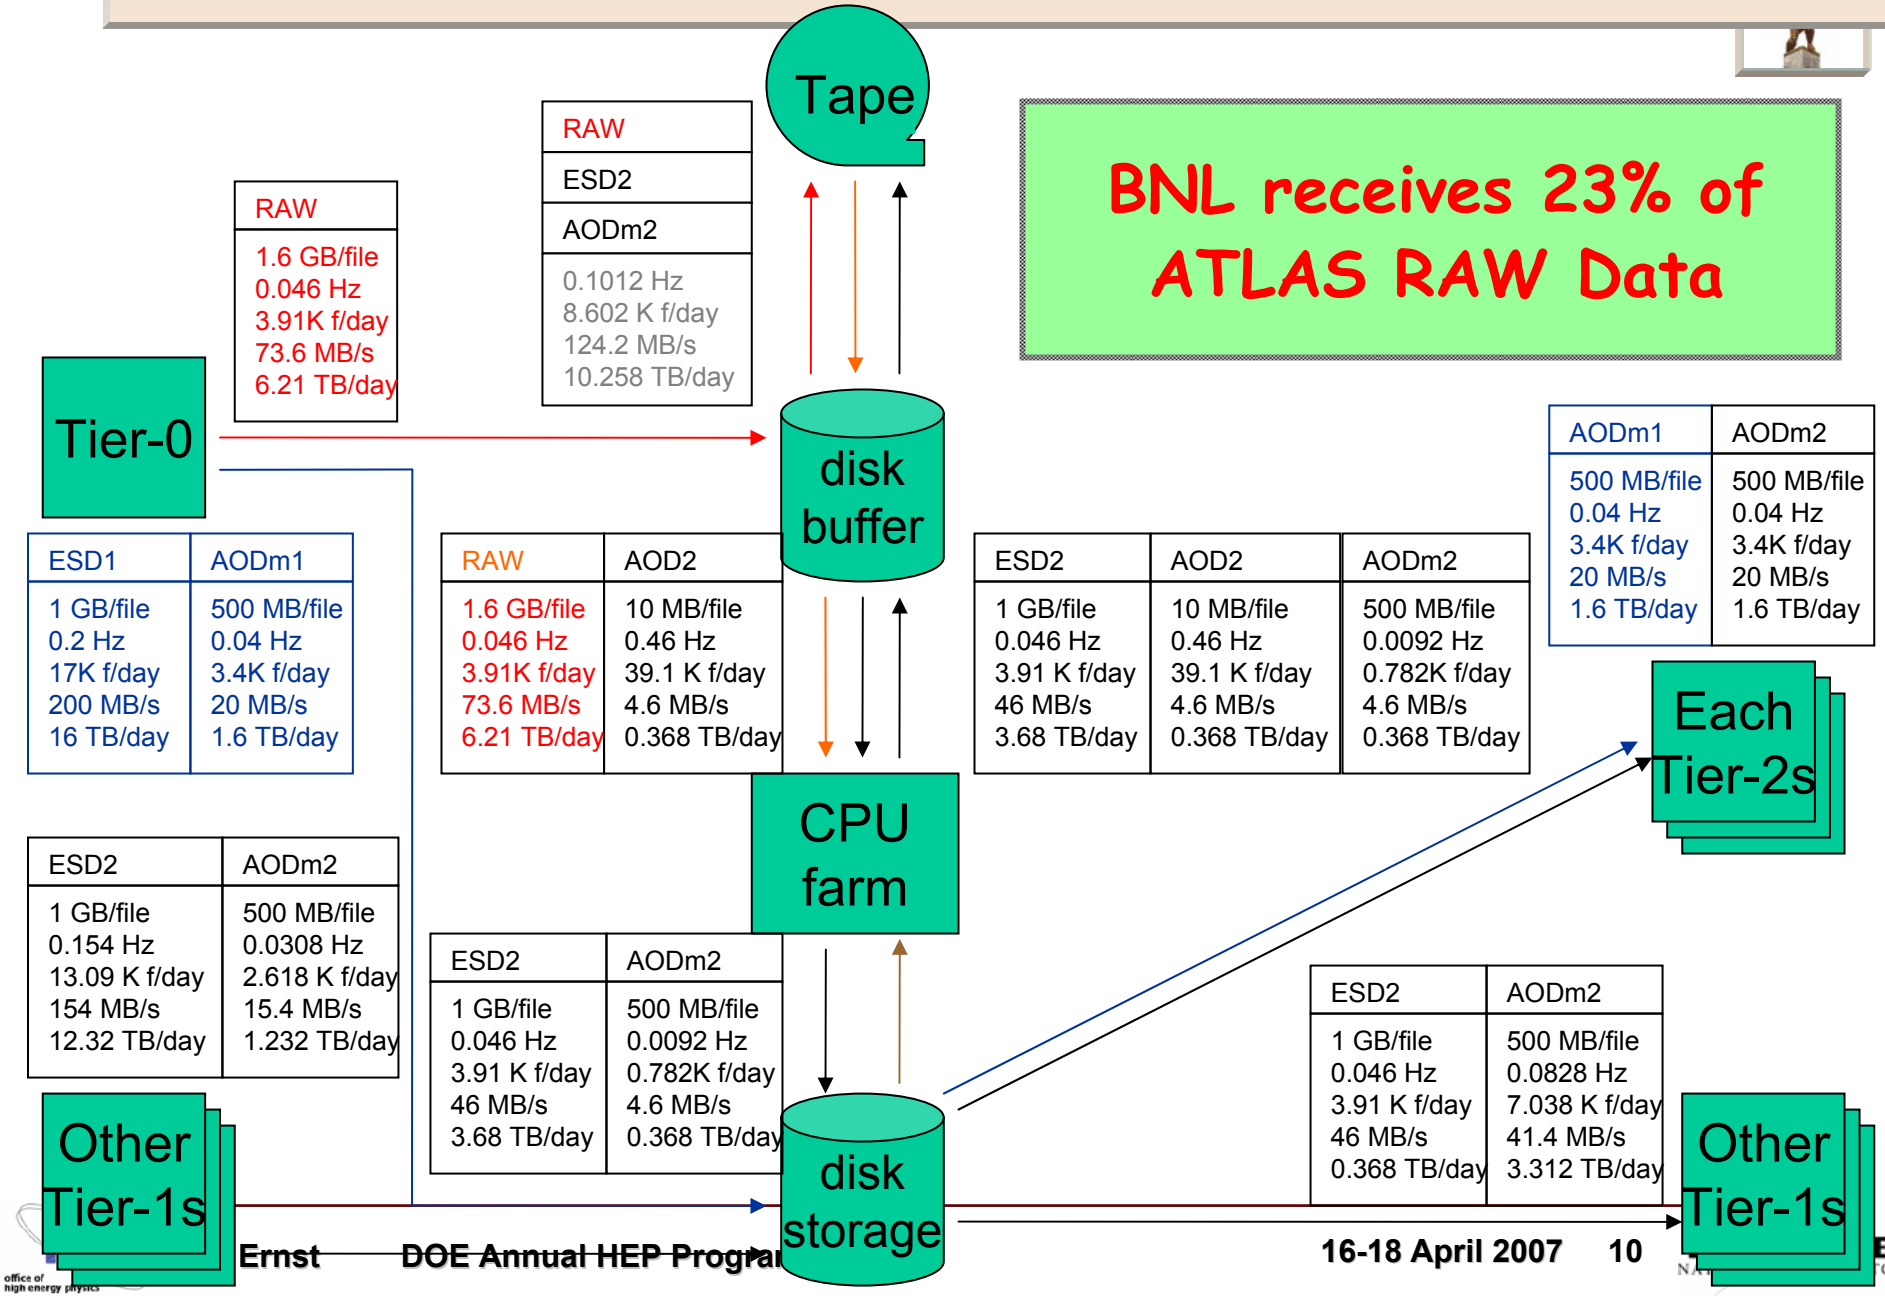

BNL receives 23% of ATLAS RAW Data

Computing System Commissioning (CSC) 2006 Production

Total Production

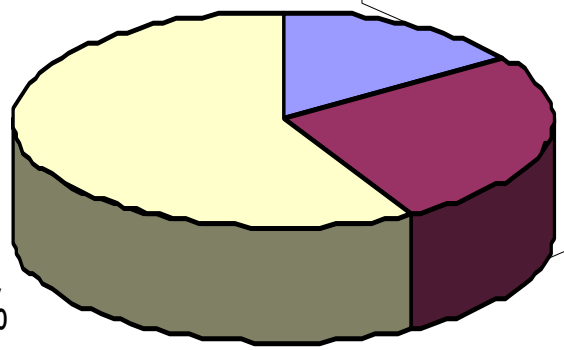

US Production

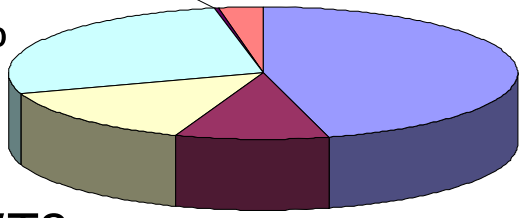

Tier 1 Utilization

- BNL Tier 1 is largest ATLAS Tier 1 and is delivering capacities consistent with this role

WLCG Accounting: ATLAS Tier-1's + CERN Apr - Oct 2006

	CPU use		disk occupancy		tape occupancy	
	KSI2K-days	% of total	TB at end of period	% of total	TB at end of period	% of total
CERN Tier-0 + CAF	95,858	28%	182	48%	469	35%
ASGC	13,413	4%	20	5%	13	1%
BNL	88,184	26%	48	13%	357	27%
CC-IN2P3	24,264	7%	15	4%	153	12%
CNAF	20,108	6%	18	5%	95	7%
FNAL	4,619	1%	-	0%	-	0%
FZK-GridKA	23,195	7%	26	7%	115	9%
NDGF	18,761	6%	28	7%	-	0%
NL LHC/Tier-1	14,574	4%	10	3%	18	1%
PIC	6,207	2%	8	2%	54	4%
RAL	27,672	8%	14	4%	54	4%
TRIUMF	1,876	1%	7	2%	-	0%
TOTAL	**338,731**	**100%**	**376**	**100%**	**1,328**	**100%**

Tier 1 Facility Capacity and Cost Estimate

- Revised Tier-1 Facility Capacity Profile

YEAR	2007	2008	2009	2010	2011
CPU (kSI2k)	2,834	7,140	11,598	18,838	26,875
Disk (TB)	1,556	4,610	8,921	17,262	24,427
Tape (TB)	993	3,284	6,276	11,996	18,781
WAN	2 x λ	2 x λ	3 x λ	4 x λ	4 x λ

wLCG Plan to pledge US model					
CPU (kSI2k)	2,560	4,844	7,337	12,765	18,193
Disk (TB)	1,100	3,136	5,822	11,637	16,509
Tape (TB)	603	1,715	3,277	6,286	9,820

- Revised Tier 1 Facility Cost Profile

Projected Tier 1 Cost Profile (@ Year k$)

($ Items below include overheads)	2007	2008	2009	2010	2011
On Program Staff Level (FTE's beyond '07)	*20*	*20*	*20*	*20*	*20*
Labor (Fully loaded salaries)	2,892	3,855	4,048	4,250	4,463
MST (travel, maint, licen, etc)	927	1,220	1,394	1,833	1,552
Facility Space & Power	248	356	469	562	598
Capital Equipment	2,826	3,762	3,902	4,989	3,124
Total	**6,893**	**9,193**	**9,813**	**11,634**	**9,736**

Tier 1 Facility Evolution Plan for FY '07

- Addition of 5 People
  - Increase from 15 in '06 to 20 during '07, which is anticipated steady state staffing level
    - Fabric/Grid/Network infrastructure: 2
    - Storage Management (dCache): 1
    - ATLAS Distributed Data Management Operation (DDM/DQ2): 1
    - Operations/user support: 1
  - Working with Stony Brook University to attract candidates
- Equipment upgrade
  - CPU: Linux Farm: 1.6 MSI2k → ~2.9 MSI2k
  - Disk: Distributed: 0.52 PB → ~1.5 PB
  - Modest Mass Storage Upgrade
    - Tape Drives, Licenses (RTU) for existing slots in Robot, dCache Write Pools
- Facility Expansion Plans
  - Shared physical infrastructure for RHIC & ATLAS Facility
  - Will run out of power and space
    - 2-phase expansion planned to add 4500 sq.ft. in 2008 and 2009
    - 2-phase UPS extension planned to add 2 MW in 2007 and 2010 .

Tier 1 Resource Procurement Schedule

- Procurement for 480 new CPU cores (120 nodes) in progress
  - Delivery announced for last week in April
  - Adds 1,1 MSI2k, 2.7 MSI2k total usabel for ATLAS in May
    - A bit more than we have pledged to wLCG for 2007 (to be delivered by 1 July)
  - Total job slots now at 1288, will be total of 1768 in May
  - 4.5 TB (gross) / 3 TB (net) disk space / node, 740 TB (net) total usable for ATLAS
    - Includes 70% efficiency factor as defined by ATLAS & wLCG
    - Still have to add 500/350 TB in second half of 2007
- Work on tripling the disk capacity
  - Increasing imbalance between CPU power and disk capacity per worker node is forcing us to look into alternate storage architectures
  - Model using distributed dCache managed disk space installed on compute nodes seems no longer viable nor cost-effective
  - Have launched a storage evaluation project
    - Benchmark and analyze a spectrum of solutions, conclusions expected in June/July
    - Well connected with similar activities through other labs (e.g. Fermilab) and HEPiX
    - Foundation for storage procurement to be conducted in July
  - Capacity will still be well beyond capacities requested by current or projected ATLAS production operations with residual used by local US ATLAS user community

Expected Computing Capacity Evolution

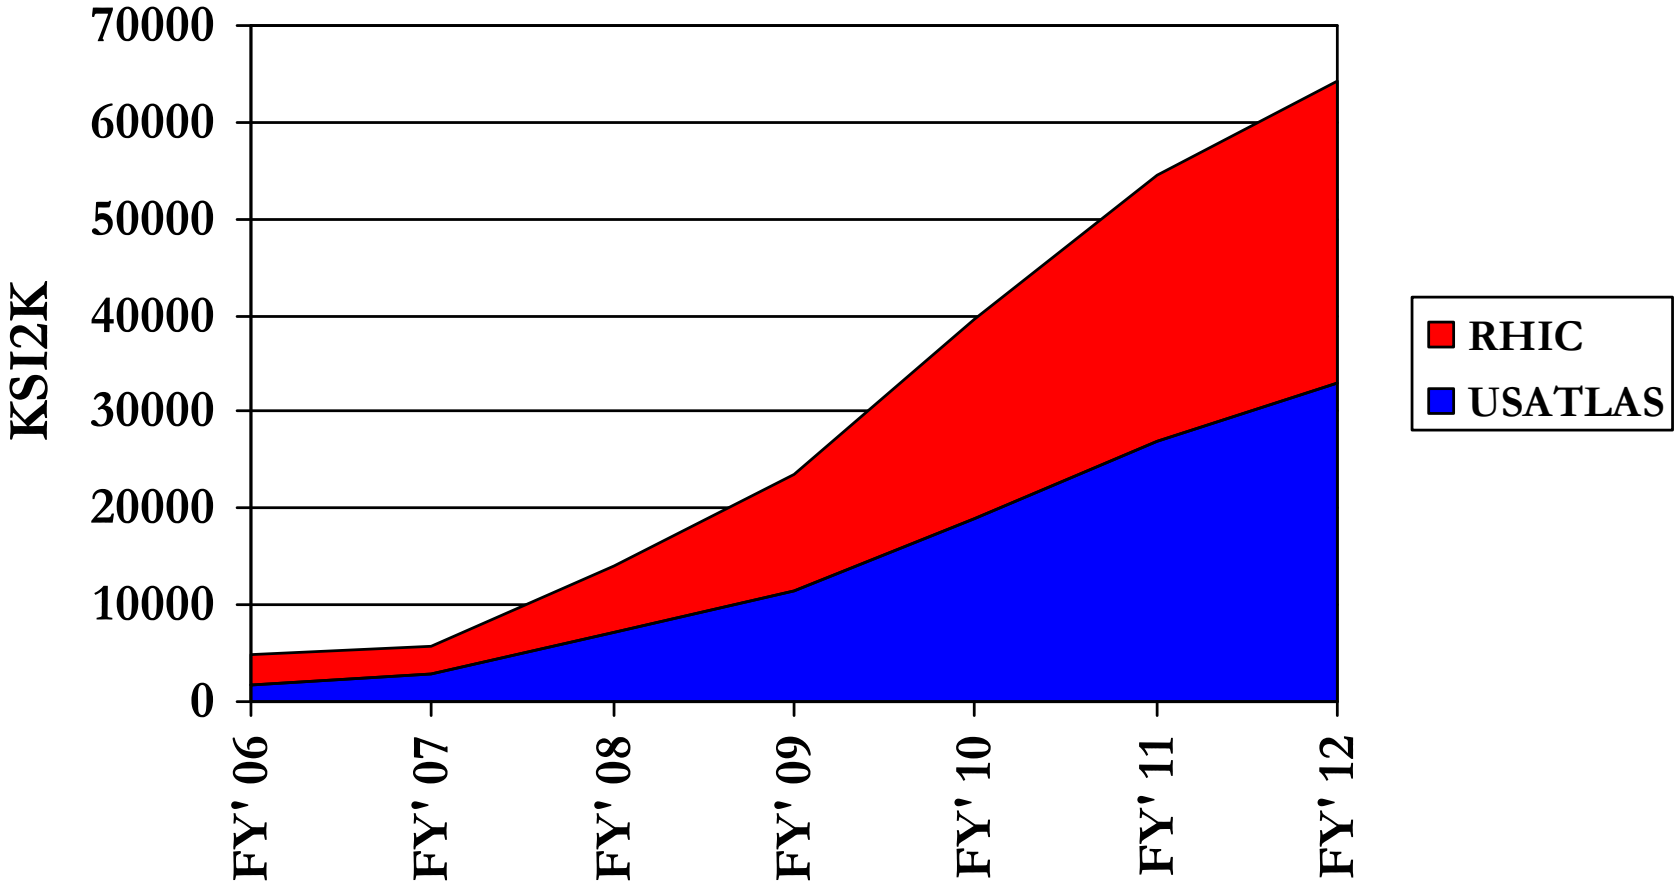

Expected Storage Capacity Evolution

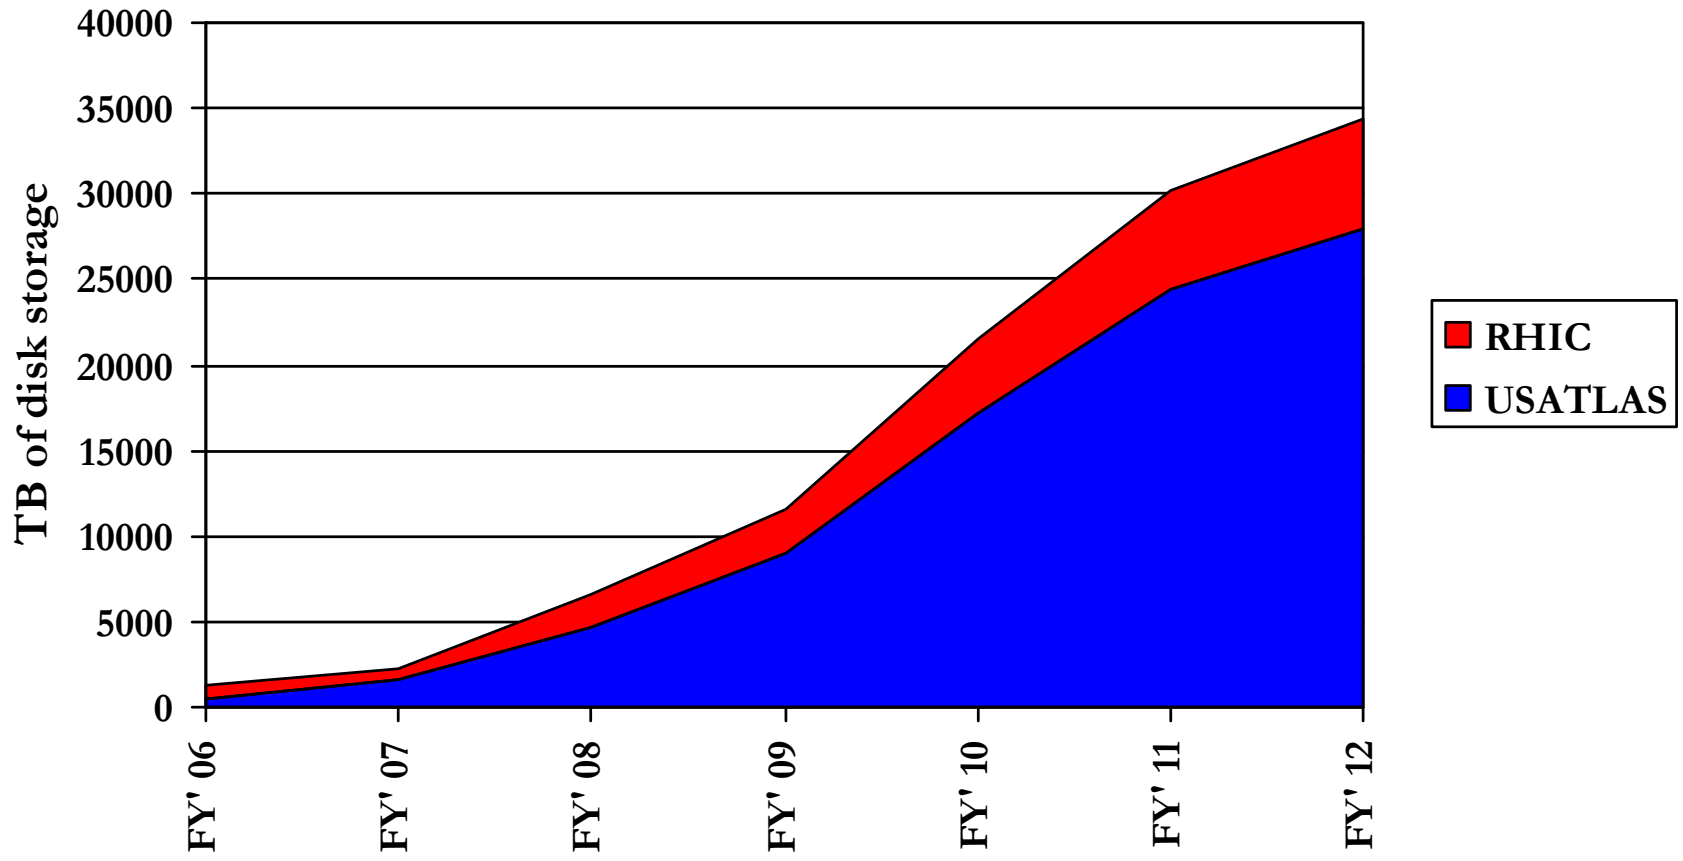

SI2K/CPU Evolution (quad-core in FY' 08)

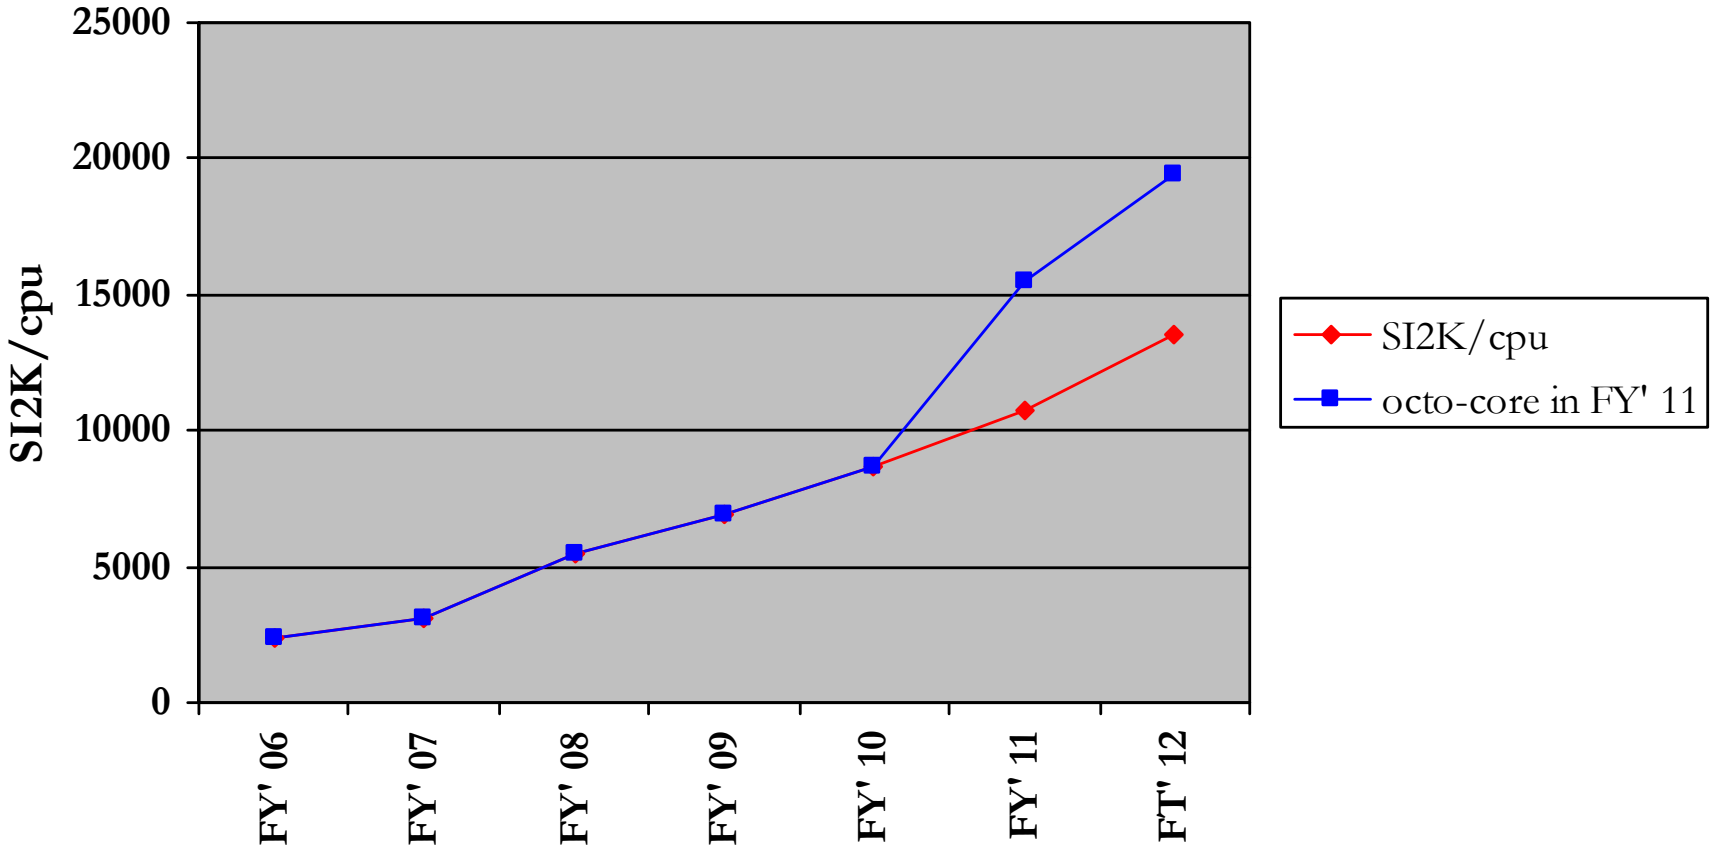

SI2K/Watt Evolution (quad-core in FY' 08)

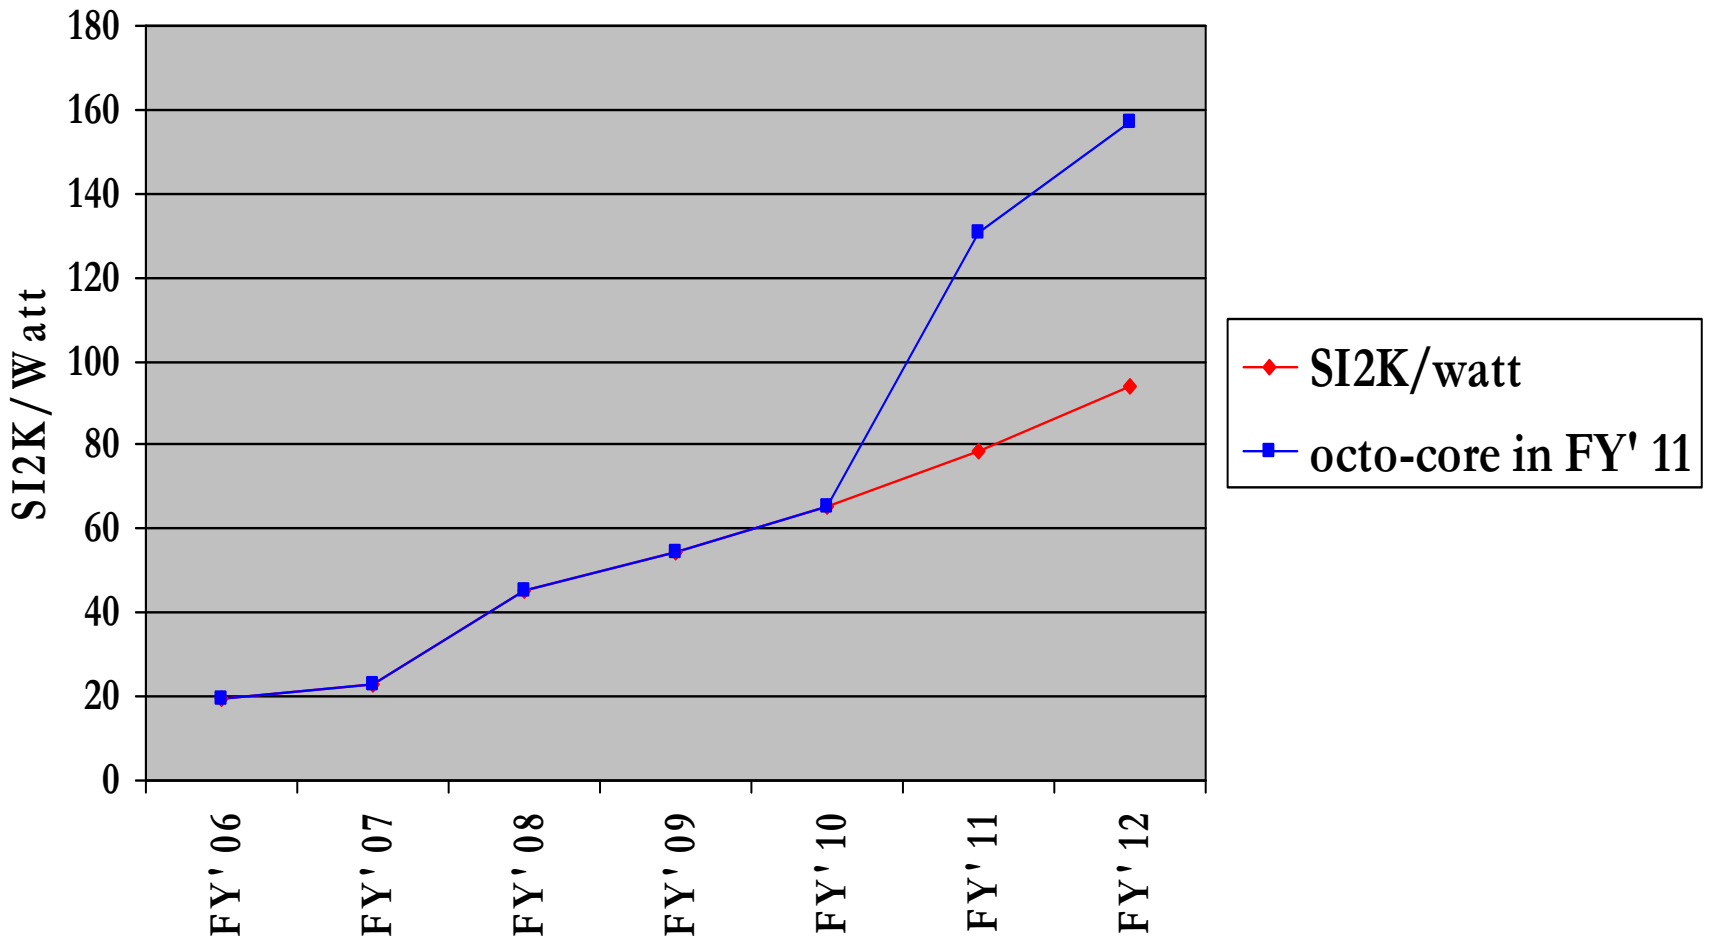

Storage Evolution Scenarios

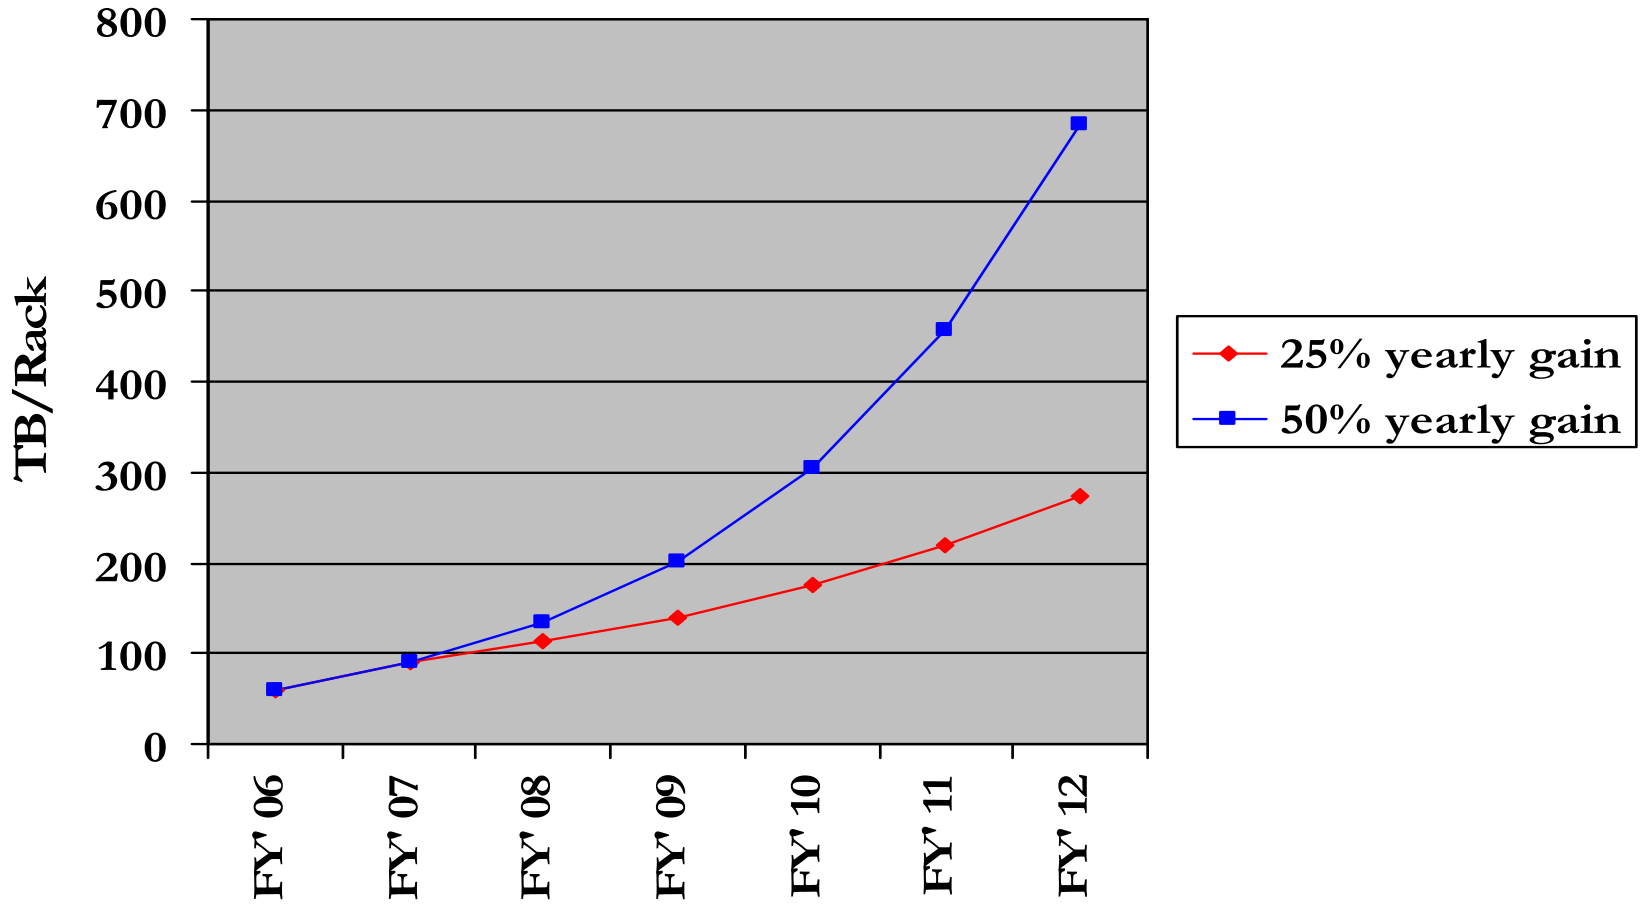

Rack Count

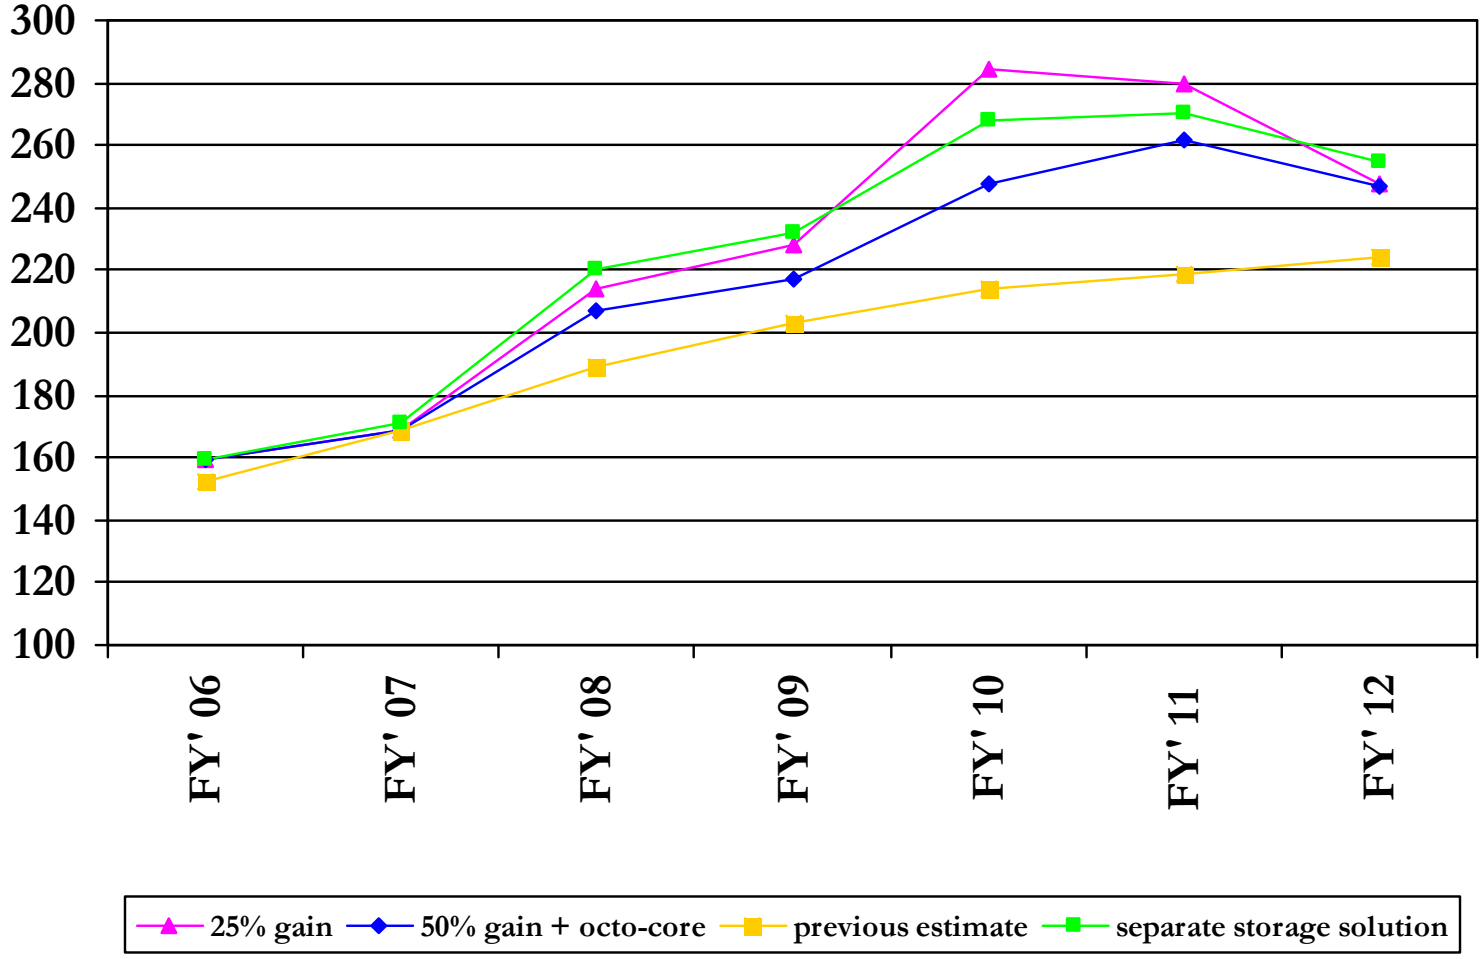

Note: Numbers include equipment replacement after 4 years

Projected Space Requirements

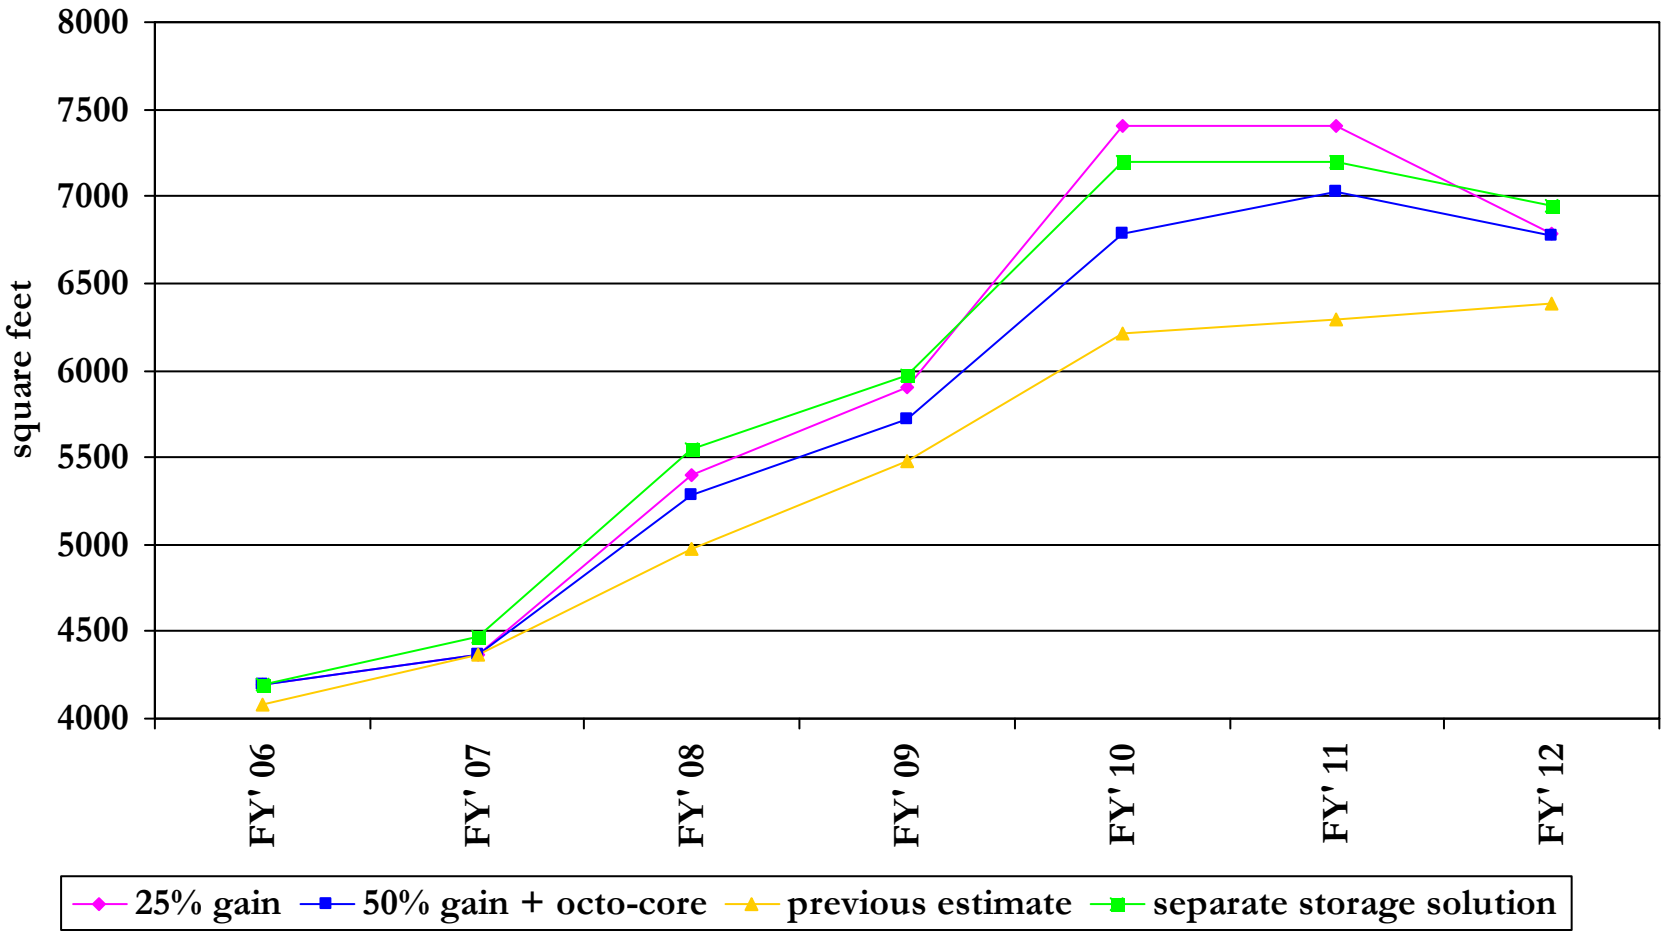

Projected Power Requirements

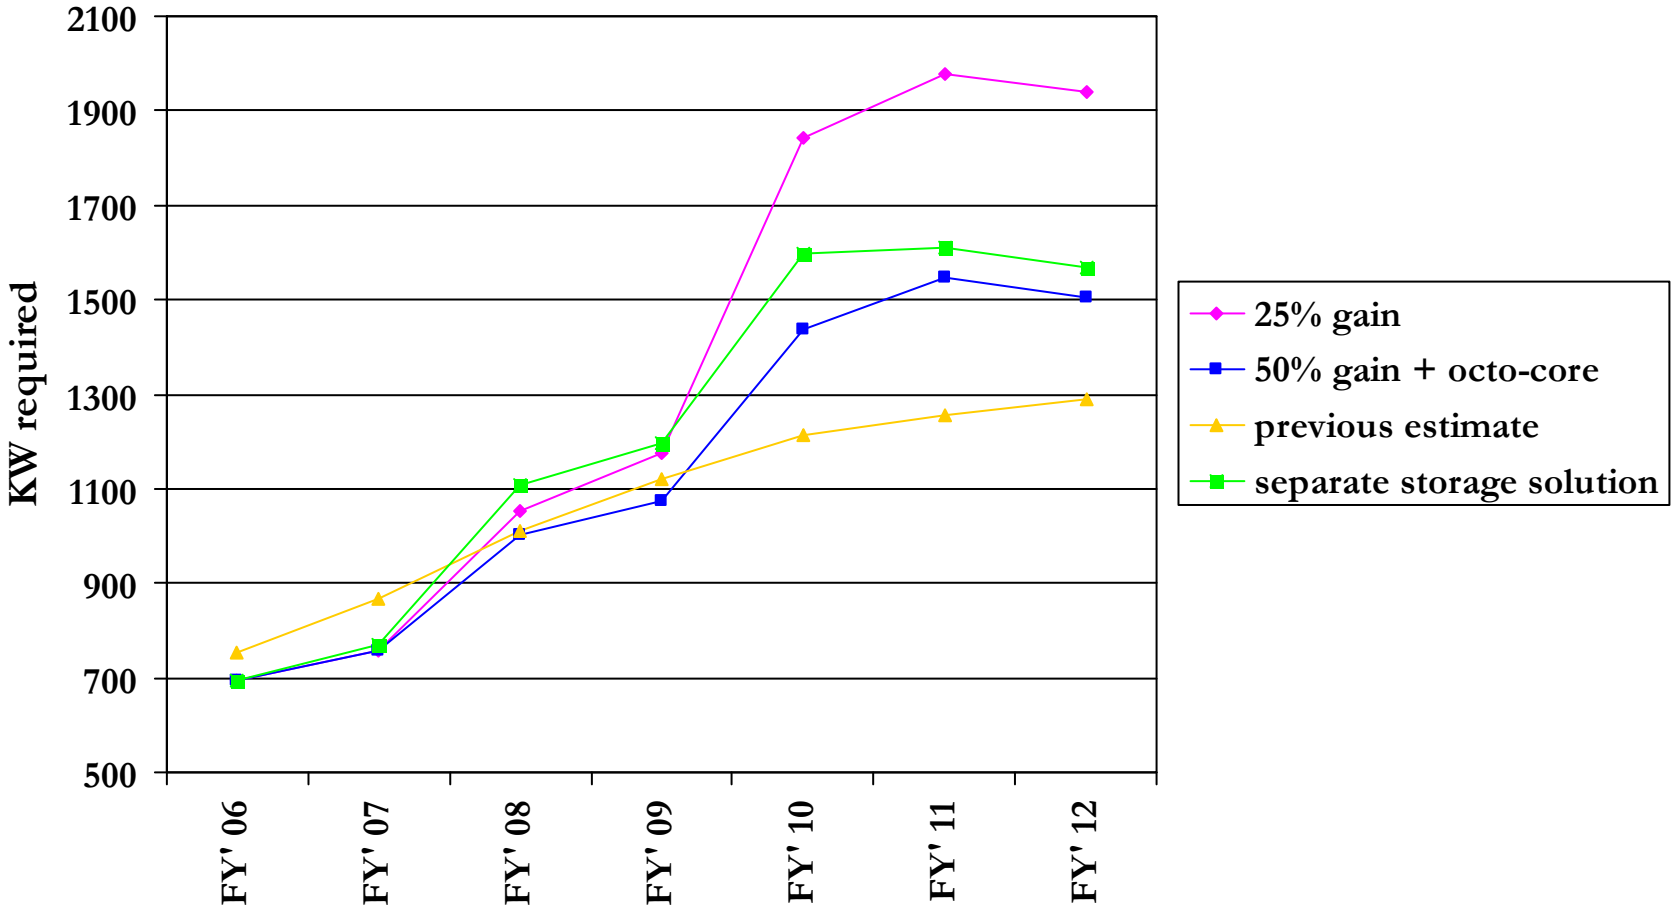

Infrastructure Planning at Tier-1

- Currently available space filled

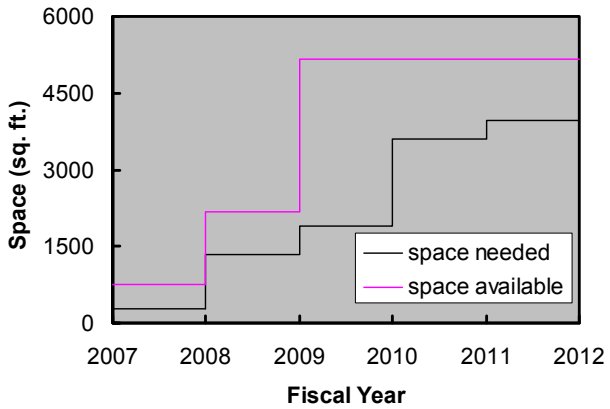

- Soon running out of Power

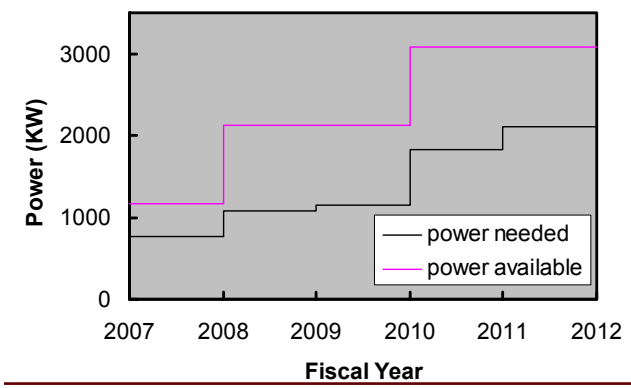

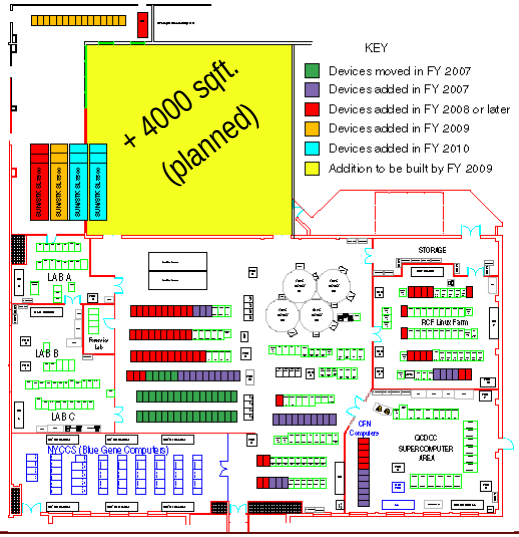

More development in 2007

Given the LHC / ATLAS Milestones

- This year is all about moving from development/deployment into stable and continuous operations
  - Deployment of full-scale hardware by 4/1/08
  - Operations
    - Monitoring and alarming Framework in place to support 24x7
    - Significant number of detailed checks and automation
    - In the process of defining levels of support and rules to follow, depending on the issue/alarm (building on RHIC operations)
  - Still more work to do on performance
    - Did we make good choices in site architecture?
    - Does it all scale? >75% use of resources through Grid Interfaces
    - Data Transfer has been/still is an issue
    - Storage System performance, including demonstrating higher data-serving rates to applications for ATLAS data analysis
      - Looking also into xrootd (have setup a tesbed, working with U.S. ATLAS Software developers and SLAC)

Transition to Operations

We need to run our sites with reasonable amount of manpower!

Stability is important, maybe more than performance

- Define milestones for uptime, success rates as measured by Site Availability Monitoring tests and DDM data replication exercises
- The Tier-1 center at BNL is tightly coupled to the Tier-2's in the US
- Tier-2's soon to act more like the Tier-2's of the Computing Model
  - Carrying the load of production MC
  - Hosting datasets for analysis
  - Hosting the work of various analysis groups
  - Supporting "local communities"
  - The effort to produce physics results and FDR will be important tests of our readiness

Issues (1/2)

- On the Critical Path
  - While (PanDA managed) Production is coming along very nicely Data Transfer (and Storage?) and Data Replication is on the critical path
  - DDM/DQ2 (ATLAS Distributed Data Management) is vulnerable and apparently not up to the performance level required
  - Data Management and Transfers least transparent component / functionality
    - Dashboard monitoring informative to some extent but not really helpful in case of problems
      - No transfers, or slowly moving – why?
      - Trouble at source or destination?
      - Nature / reason of trouble?
  - Very complex situation – Diagnosis difficult and requires expert-level knowledge in multiple areas
    - Currently limited to few experts – not scalable and does not allow site admins to assess how their site is doing
    - Often requires access to distributed log files

Issues (2/2)

- (Still) on the Critical Path
  - Storage Systems and Tier-1 / Tier-2 Sites
    - No technology baseline in U.S. ATLAS
    - Impact on operational readiness and interoperability unclear
    - Significant number of technical problems at all levels (FTS, DDM/DQ2, SRM/dCache)

Securing the Facilities' Readiness

- Towards ATLAS Milestones
  - Put an Integration Program in place which aims at building an Integrated Virtual Computing Facility that we need to support LHC Analysis in the US
  - With exercises designed to verify sites' readiness, stability and performance
  - Coordinated by U.S. ATLAS Facility Manager (M.Ernst / BNL)
    - In the process of visiting all U.S. ATLAS Tier-2 Centers
    - Organizing quarterly F-to-F meetings w/ Tier-2's, now incl. Tier-3's
- Exploit commonality and establish (technology) baseline whenever possible
  - Synergy allows to bundle resources (development and operations)
- Site Certification
  - Site admins are asked to install well defined software packages and to make needed capacities available to the Collaboration
  - We will continuously run use-case oriented exercises and will document and archive the results
    - Heartbeat – Data Transfers on a basic level
    - Dataset replication based on high-level functionality (DDM/DQ2)
    - Processing (Analysis job profile)
      - Grid Job submission (PanDA) – distribution based on data affinity
      - Local data access (from SE)
    - Monitor and archive results from exercises

Coordination & Communication

- Find U.S. Production Operations and Deployment Manager and put him in charge of
  - Coordination of Production activities at U.S ATLAS sites
  - Coordination of deployment of ATLAS Software
  - Member of the Facilities Group
- Weekly Integration Meetings
  - Chaired by U.S. ATLAS Facility Manager
  - Activity driven agenda
    - Integration News & Planning – a summary of issues and upcoming activities w/ computing milestones and deliverables
    - Reports from Production & DDM operations (technical)
    - Site issues – summaries
      - Move details to site help sessions (people call in and help with particular issues)
    - Standing items
      - Development status of tools & aids

Tier-1 Cyber Security

- Tier-1 is within BNL IT Division managed cyber security domain
  - Site Firewall – partitions BNL including Tier 1 from world
  - Tier 1 Firewall – similarly partitions Tier 1 / RHIC enclave from rest of BNL
- ATLAS Tier 1 / RHIC facility is among most aggressive organizations at BNL in cyber security
  - Tom Throwe, facility deputy director, chaired Lab's Cyber Security Advisory Comm. for last year & involved in site wide configuration auditing system
  - Tier 1 / RHIC first at BNL to establish firewall protected internal enclave
    - Affording protection against cyber security problems elsewhere within BNL
  - Moved this past year from simple ssh to 2 factor authentication
    - => ssh keys and/or one-time passwords (ssh keys phase out expected)
    - This action addresses only facility intrusions of last 2 years
      - which were a result of externally sniffed passwords

Site Cyber Access and Incident Handling

- ATLAS Login Account Granting
  - Requires BNL Guest Number and cyber security training
  - Currently granting temporary accounts with application for Guest Number and phone verification from known ATLAS sponsor
- Grid Access Does Not Require Local Account
  - Does require Grid certificate and proper Virtual Organization membership
- User Traceability Required
  - Grid users are mapped to unique local accounts
  - US ATLAS PanDA will use “glexec” (part of Condor startd) to allow job re-authentication as payload user
- Cyber Security Incidences Response
  - WLCG/OSG security protocols exist to protect against damage and contain problems
  - Bob Cowles (SLAC) is U.S. ATLAS cyber security officer
  - BNL cyber security as well as BNL Tier-1 are involved in defining and acting when these protocols are invoked
  - BNL Tier-1, with local full ESD copy, has reduced sensitivity to disruption by an incident elsewhere in the WLCG

BNL Optical Private Network (OPN) Security

- Transfer Bandwidth Performance Requirements necessitate By-Passing of Firewalls for Select Connections
  - Currently only for CERN => BNL Tier-1
- Alternate Security Controls
  - Router security controls
    - Currently using Access Control Lists (ACL)
    - If list grows, port range limitation can be added
  - Host-based security controls
    - Only a limited number of dual homed machines (dCache doors) are accessible via OPN
    - Only required transfer associated ports/services are enabled on these machines
    - These servers are very carefully patched and well monitored
    - BNL cyber security team approves configurations, receives logs and has root access

BNL Tier 1 Wide Area Network Storage Interfaces

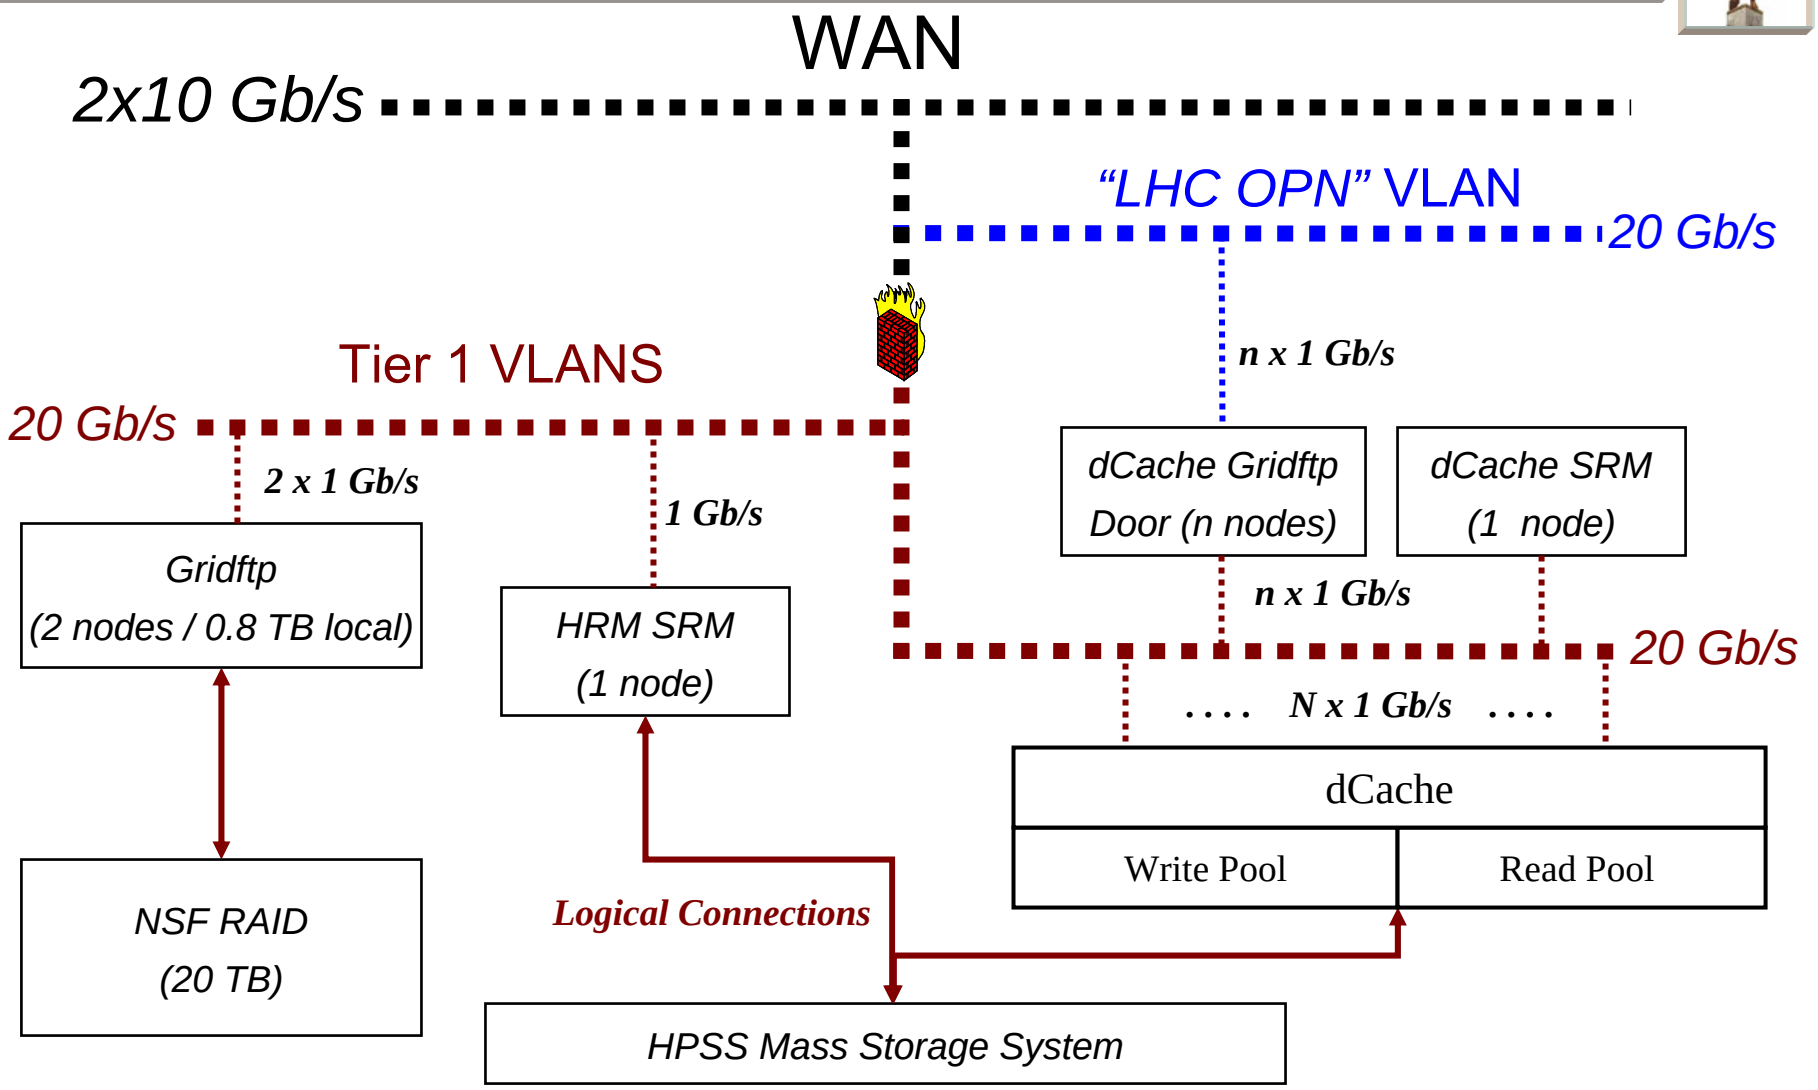

User Interaction & Feedback

- Web site (currently under revision) and mailing lists
- Trouble Ticket System (CTS {*Locally developed*} migrating to RT)
- Weekly: US ATLAS Grid teleconference
  - Production, configuration and infrastructure issues
  - To be re-assigned to Computing Integration/Operations Meeting
- ~ Monthly: BNL ATLAS Computing Meetings
  - Brings together management, facilities, and (power) users to address concerns related to production, (software) infrastructure and non-production utilization
- US ATLAS *Computing* Resource Allocation Committee (RAC)
  - Primary source of input on priorities

Summary

- The BNL Tier-1 serves as the hub and principal center of the US community, with scale-up for data taking underway
- US ATLAS Tier-1 facility at BNL is on track to meet the performance and capacity requirements of the ATLAS computing model augmented to supply appropriate additional support to US physicists
- The facilities, both Tier-1 and Tier-2's, have performed well in both ATLAS computer system commissioning and WLCG service challenges
  - An Integration Program is needed to ensure readiness in view of the steep ramp-up
- Staffing pressure continues at the Tier-1 due to:
  - … rapid growth in capacity and usage, both programmatic and chaotic
  - … support role for Tier 2's (and Production development & operations)
  - … need for in-house expertise in growing number of critical areas
    - Database admin and usage optimization this year added to
      … data storage, management and transfer, user and access management, etc.
  - Staff growth to address this is critically important so recruiting will continue to be an major activity this year
- Space, Power & Cooling at the Tier-1 center on the critical path