RCF Operations, Plans and Issues

Michael Ernst

DOE/Nuclear Physics Review of RHIC Science and Technology

22 - 24 July 2009

RCF Mission and Scale

- Mission
  - Online Recording of Raw Data
  - Production reconstruction of Raw Data
  - Primary Facility for Data Selection and Analysis
  - Long time Archiving and Serving of all Data
- Scale
  - Authorized staff of 20 FTE's
  - Historically ~$2.0M/year equipment replacement funding – 2006 limited to $1.3M, 2007 had $1.7M, last year to $1.7M, this year back to $2.0M
    - Funds primarily used to improve storage & network infrastructure, and to address obsolescence
    - Remainder allows for modest expansions of PHENIX and STAR computing and storage capacities

RCF Staff

- Current authorized staff level: 20 FTE's
- Excellent synergy in the context of a co-located ATLAS Tier-1 Center in terms of operations
  - Very high level of commonality
  - A dramatic divergence in technical directions could change this, but this seems very unlikely
- It does not allow for aggressive involvement in new technologies
  - Effort spent primarily on Integration and Operation

	Current FTE's	Target FTE's
Linux Farms	3.0	3.5
Mass Storage	4.2	4.2
Central Disk	2.5	2.5
User Support	2.9	2.9
Fabric Infrastructure	2.6	2.6
Wide Area Services	1.8	1.8
Administration	1.5	2.5
Total	18.5	20.0

Strategies and best practices for RHIC Computing

- Drivers that mandate a strategy of continuous refresh of computing facilities
  - Maintaining old or non-aligned (with the current/future program) tools and software infrastructure is costly in effort
    - Each unique solution costs multiple FTEs (at Experiments and RCF)
  - Robotic storage and tape technology must move forward – costs of robot slots vs tape media cost must figure into the economic model
    - Density of tapes doubles every ~ 3 years
    - Strategy is to migrate data (to new media gen) and keep "online" in robotic storage
      - **Timing is essential to save on media costs (new robot may be needed earlier)**
      - **Requires additional tape drives to copy data**
      - **Transition to LTO Gen 4 (LTO5, 1.6 TB) planned for 2011 (no free media exchange)**
    - 3rd SL8500 library needed in 2010, 4th in 2012

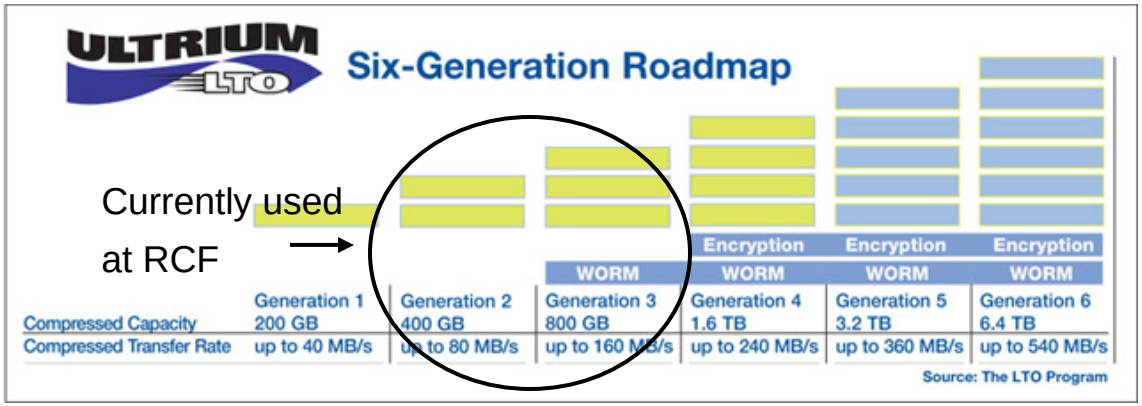

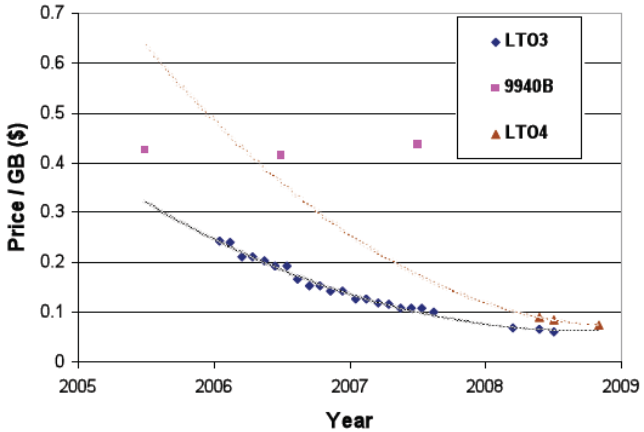

Replenishment of Computing Facilities

- About 25% of the CPUs will have to be replaced each year
  - Expect heavy demand on analysis computing due to PHENIX & STAR Detector and DAQ upgrades (higher trigger rates)
- About 25% of the disks will have to be replaced each year
- Servers and Network equipment will have to be replenished on either a 3 or 4 year cycle

RCF Capacities

End of FY09	PHENIX	STAR
CPU (MSI2K)	~5.7	~4.4
Disk (TB)	1,430	1,123
Tape (PB)	4.1	2.6

- Will meet the required capacities to process Run 9 data
  - ~2x needed in 2010 but majority of equipment funds needed for central infrastructure expansion and replenishment of obsolete equipment

PHENIX Raw Data Volume collected & archived

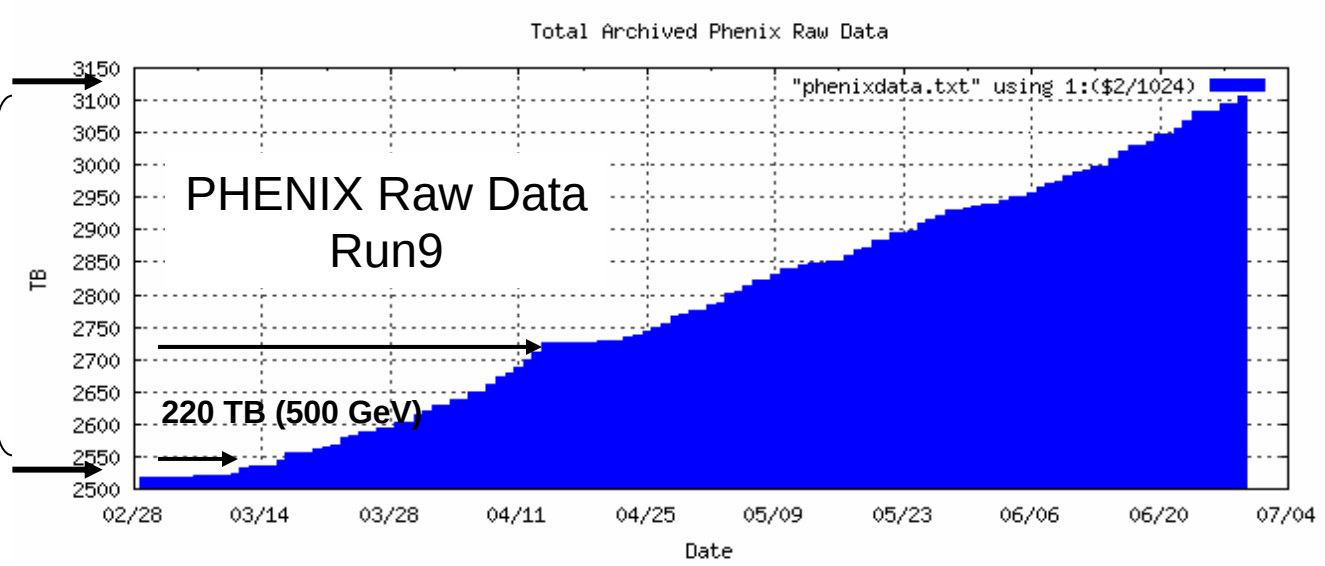

~600 TB

2500 TB

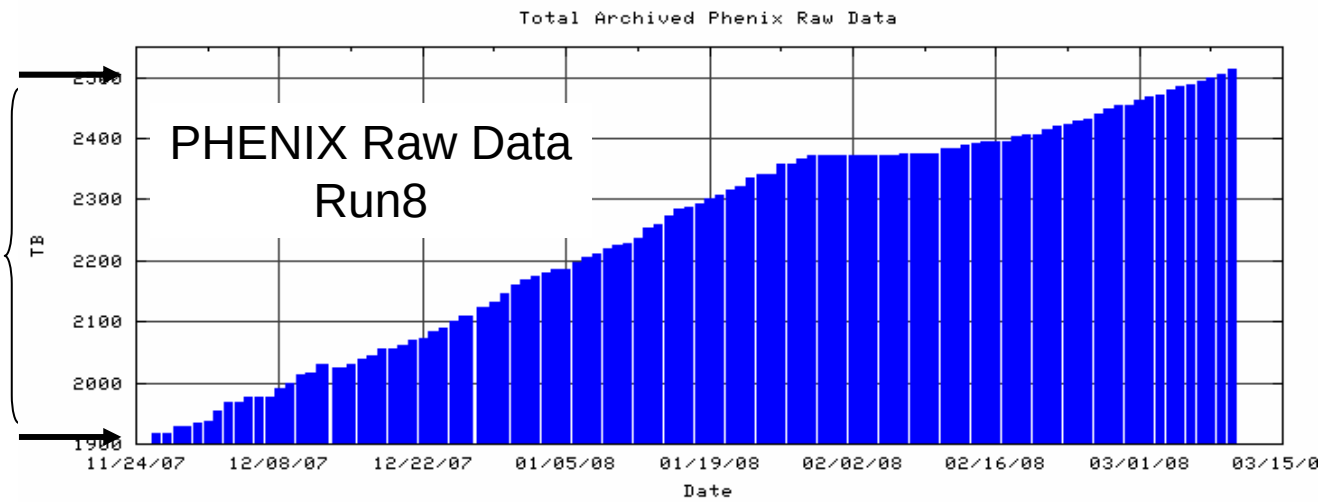

~600 TB

1900 TB

STAR Raw Data Volume collected & archived

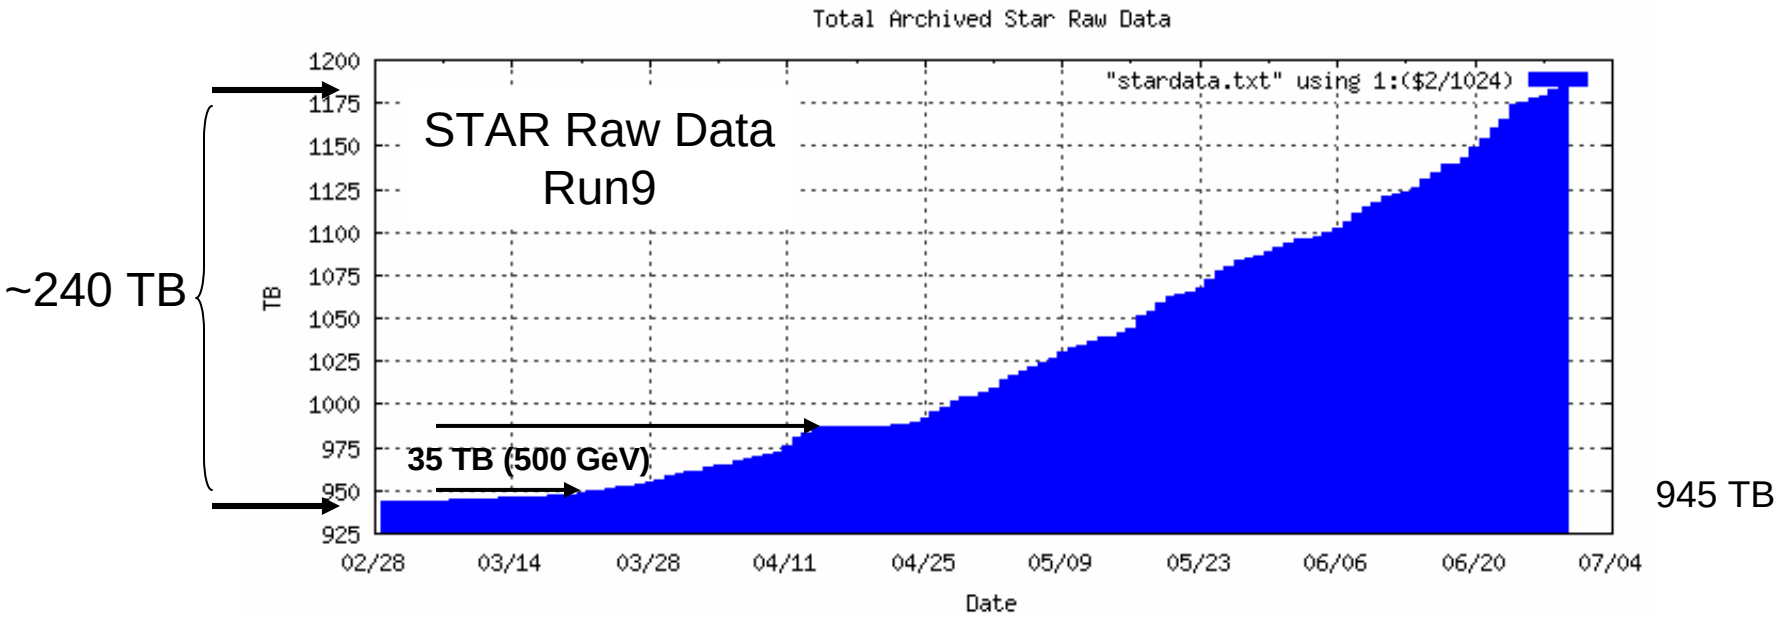

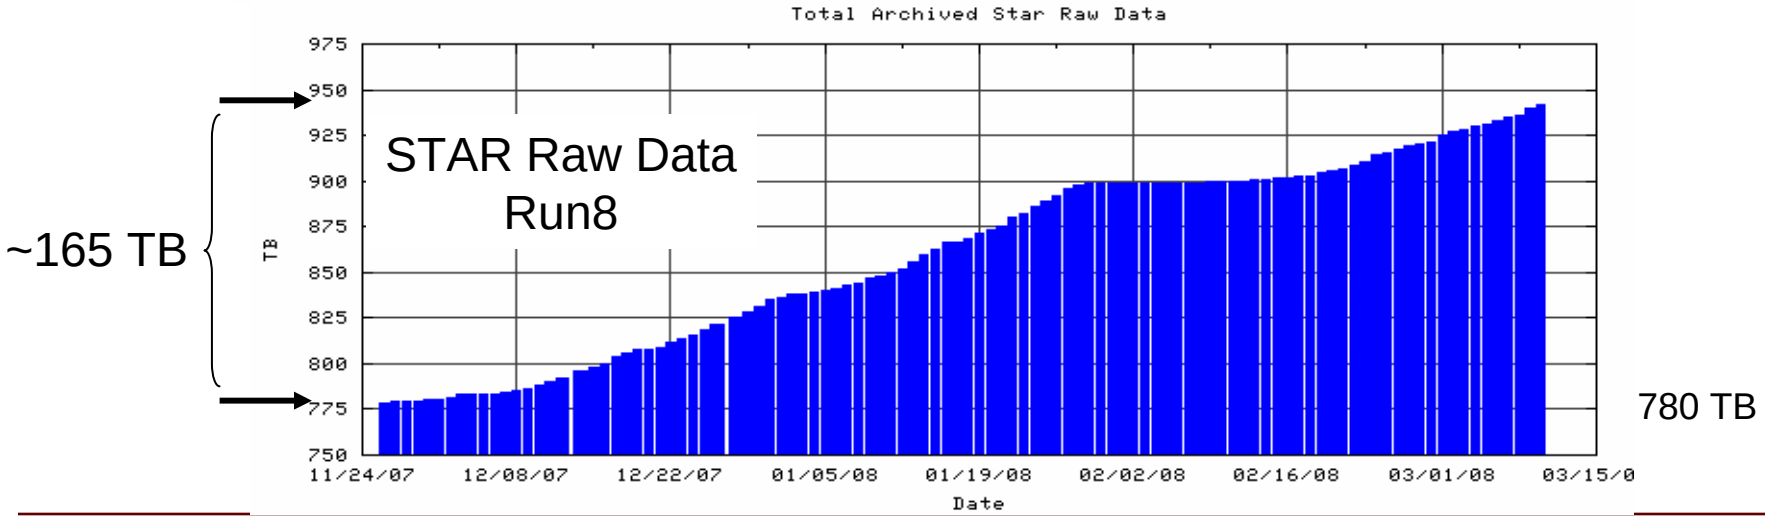

RAW Data collected in RHIC Runs

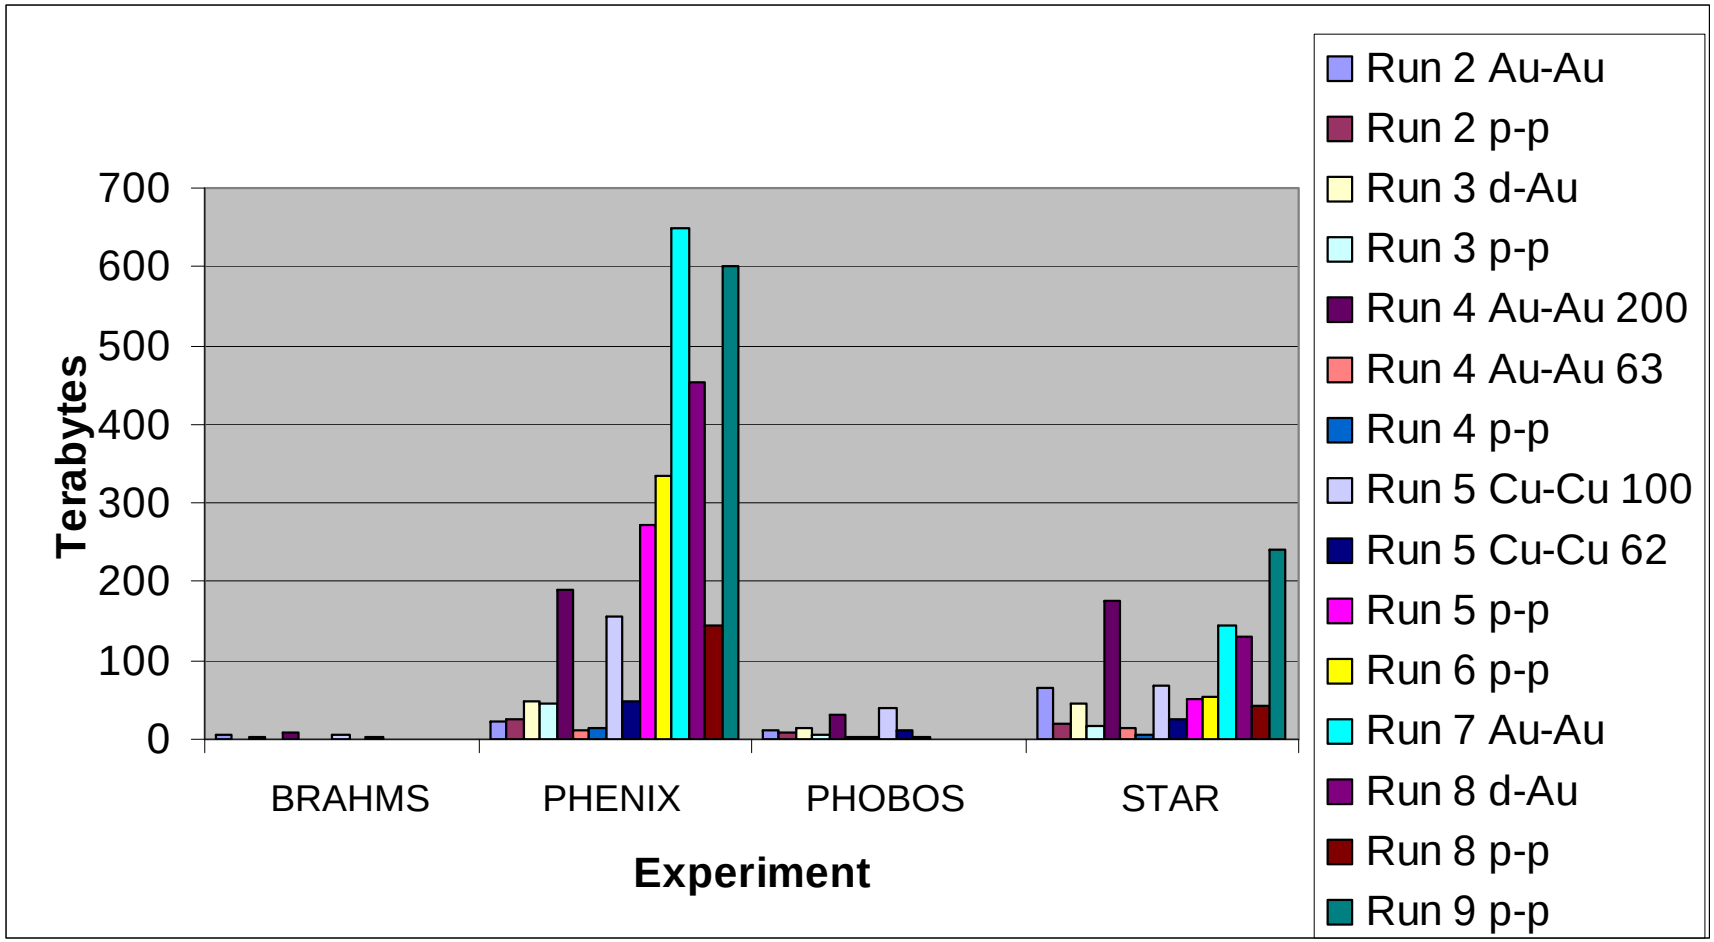

Computing Requirements Estimate

- A Comprehensive Long Range Estimate done by PHENIX, RCF and STAR in Fall/Winter 2005
  - Conclusions published as part of "Mid-Term Strategic Plan: 2006-2012 For the Relativistic Heavy Ion Collider"
- Input is Raw Data Volume for Each Species & Experiment by Year
- Model for Requirements Projection
  - Assume Facility resources need to scale with Raw Data volume
  - With adjustable parameters reflecting expected relative …
    - Richness of data set (density of interesting events)
    - Maturity of processing software
    - Number of reconstruction passes

    … for each experiment, species, and year
  - In view of revised running scenarios, detector upgrade schedules and RCF funding profile we (PHENIX, STAR (have written a comprehensive S&C plan) and RCF) have developed a revised deployment schedule for computing resources to fit the budget guidance for equipment and person power

Projecting Central and Experiment resource needs

- Using tentative RHIC running scenarios
  - Example taken from STAR's revised S&C plan

Run	Species	Purpose	Dataset
2009	p+p 200 GeV	$\Delta G(x)$ with dijets	900 M events, 50 pb^{-1} sampled
2010	Au+Au 200 GeV	Precision Au	600 M events, 2 nb^{-1} sampled
	p+p 500 GeV	First W measurements	250 M events, 10 pb^{-1} sampled
2011	Au+Au 5-40 GeV	Energy Scan	50 M events
	U+U 200 GeV	Highly elliptical zone at high density	650 M events
2012	p+p 500 GeV	Precision W	550 M events, 150 pb^{-1} sampled
	Au+Au 200 GeV	Heavy Flavor, RHIC II	750 M events, 5 nb^{-1} sampled
2013	p+p 500 GeV	Precision W	550 M events, 150 pb^{-1} sampled
	Au+Au low E	Energy Scan	50 M events
2014	p+p 200 GeV	Au reference with HFT	2100 M events
	Au+Au 200 GeV	Heavy Flavor, RHIC II	1200 M events, 10 nb^{-1} sampled

  - We understand that this is the tentative run plan
    - Need to be flexible and nimble regarding changing requirements

Experiment Parameters	2009	2010	2011	2012	2013	2014	2015
STAR sust. Au-Au Data Rate (MB/sec)	0	265	325	273	14	180	180
STAR sust. p-p Data Rate (MB/sec)	61	30	20	67	20	171	171
PHENIX sust. Au-Au Data Rate (MB/s)	0	450	600	600	600	600	600
PHENIX sust. p-p Data Rate (MB/s)	350	380	600	500	500	500	500
RCF Aggregate Rate Au-Au (MB/s)	0	715	925	873	614	780	780
RCF Aggregate Rate p-p (MB/s)	411	410	622	567	567	671	671
Experiment Efficiency (Au-Au)	85%	85%	85%	85%	85%	85%	85%
Experiment Efficiency (p-p)	85%	85%	85%	85%	85%	85%	85%
STAR Au-Au Raw Data Volume (TB)	0.0	519.7	638.0	535.4	27.5	353.0	353.0
STAR p-p Raw Data Volume (TB)	240	58.8	40.0	131.4	40.0	335.3	335.3
STAR Raw Data Volume (TB)	240	578.5	678.0	666.8	67.5	688.3	688.3
PHENIX Au-Au Raw Data Volume (TB)	0.0	882.5	1176.6	1176.6	1176.6	1176.6	1176.6
PHENIX p-p Raw Data Volume (TB)	600	745.2	1176.6	980.5	980.5	980.5	980.5
PHENIX Raw Data Volume (TB)	600	1627.7	2353.2	2157.1	2157.1	2157.1	2157.1
Aggregate Raw Data Volume (TB)	840	2206.2	3031.2	2823.9	2315.9	2845.4	2845.4
STAR Derived Data Volume (TB)	360	810.0	950.0	830.0	103.0	1070.0	1070.0
PHENIX Derived Data Volume (TB)	50.0	400.0	800.0	500.0	500.0	500.0	500.0
Aggregate Derived Data Volume (TB)	260.0	1210.0	1750.0	1330.0	850.0	1570.0	1570.0
STAR Data Volume (TB)	600	1388.5	1628.1	1496.8	170.5	1758.3	1758.3
PHENIX Data Volume (TB)	650	2027.7	3153.2	2657.1	2657.1	2657.1	2657.1
RCF Annual Aggregate Data Volume (TB)	1250.0	3416.2	4781.3	4153.9	2827.6	4415.4	4415.4
RCF Total Aggregate Data Volume (TB)	7250.0	10666.2	15447.5	19601.4	22429.0	26844.4	31259.8
ATLAS Data Rate (Read+Write, MB/s)	800+120	1391+152	1992+180	2344+198	3543+257		
ATLAS Aggregate Data Volume (TB)	3540.0	6429.0	10730.0	14627.0	20945.0		
RACF Total Data Rate To/From Tape (MB/s)	1331.0	2258.0	3094.0	3415.0	4414.0		
RACF Total Data Volume on Tape (TB)	10790.1	17095.2	26177.5	34228.4	43374.4		

Equipment Spending Profile (k$)

Year	2009	2010	2011	2012	2013	2014	2015
Equipment Funds	2000	2500	3000	3000	3000	3000	3000
Robot	0	100	80	200	50	50	50
Tape Drives	0	200	250	140	60	300	150
HPSS Cache	135	500	280	149	100	540	280
Network	250	300	200	200	0	500	200
Central Disk	200	0	400	180	240	270	244
CPU replacement	0	234	175	88	263	40	243
Central Server	100	160	100	75	50	100	50
Total Available to Experiments	1315	1006	1515	1968	2237	1250	1833

- HPSS Cache (disk and servers) needs to be replaced in FY09, obsolete technology
  - Cost is $530k, but only $135k available
  - Submitted a request for supplemental funding to DOE in view of tight FY09 & FY10 budget

HPSS Architecture and Principal Deployment

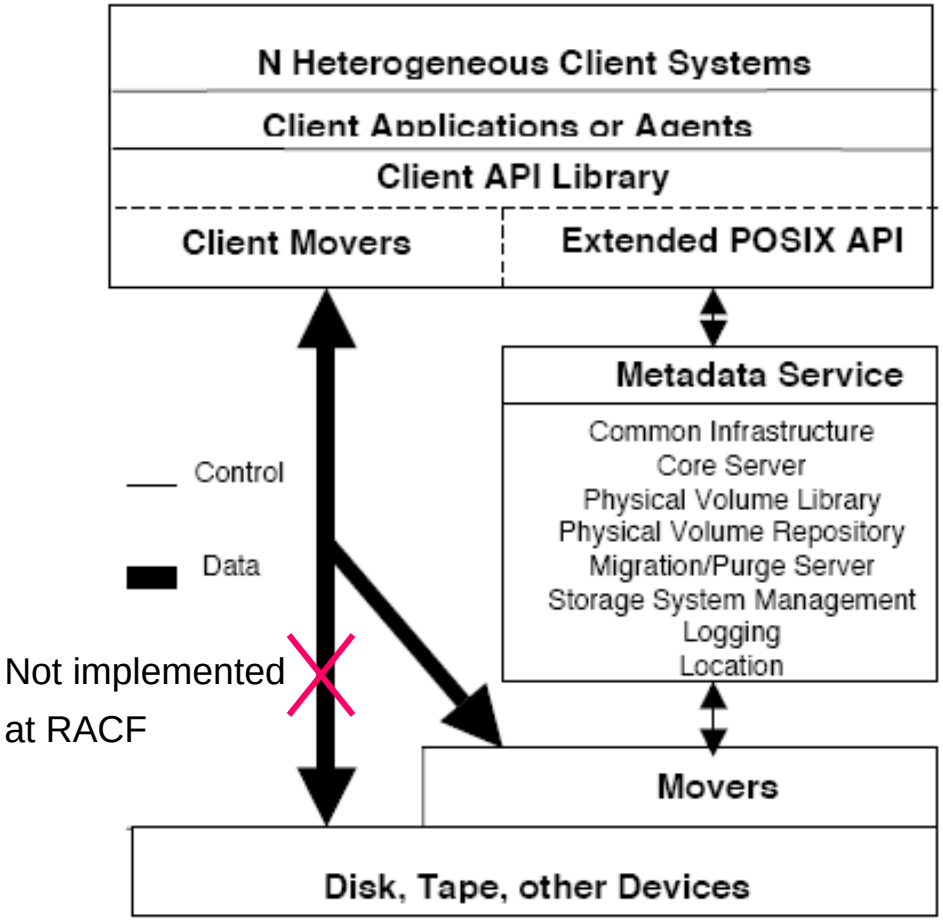

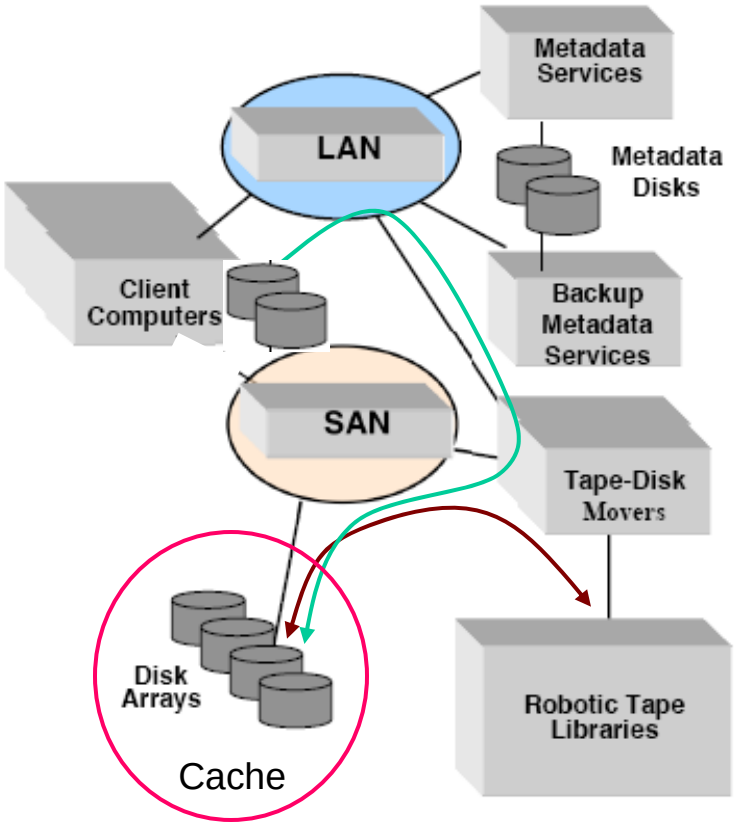

All data to/from tape moved through cache

HPSS Cache Throughput

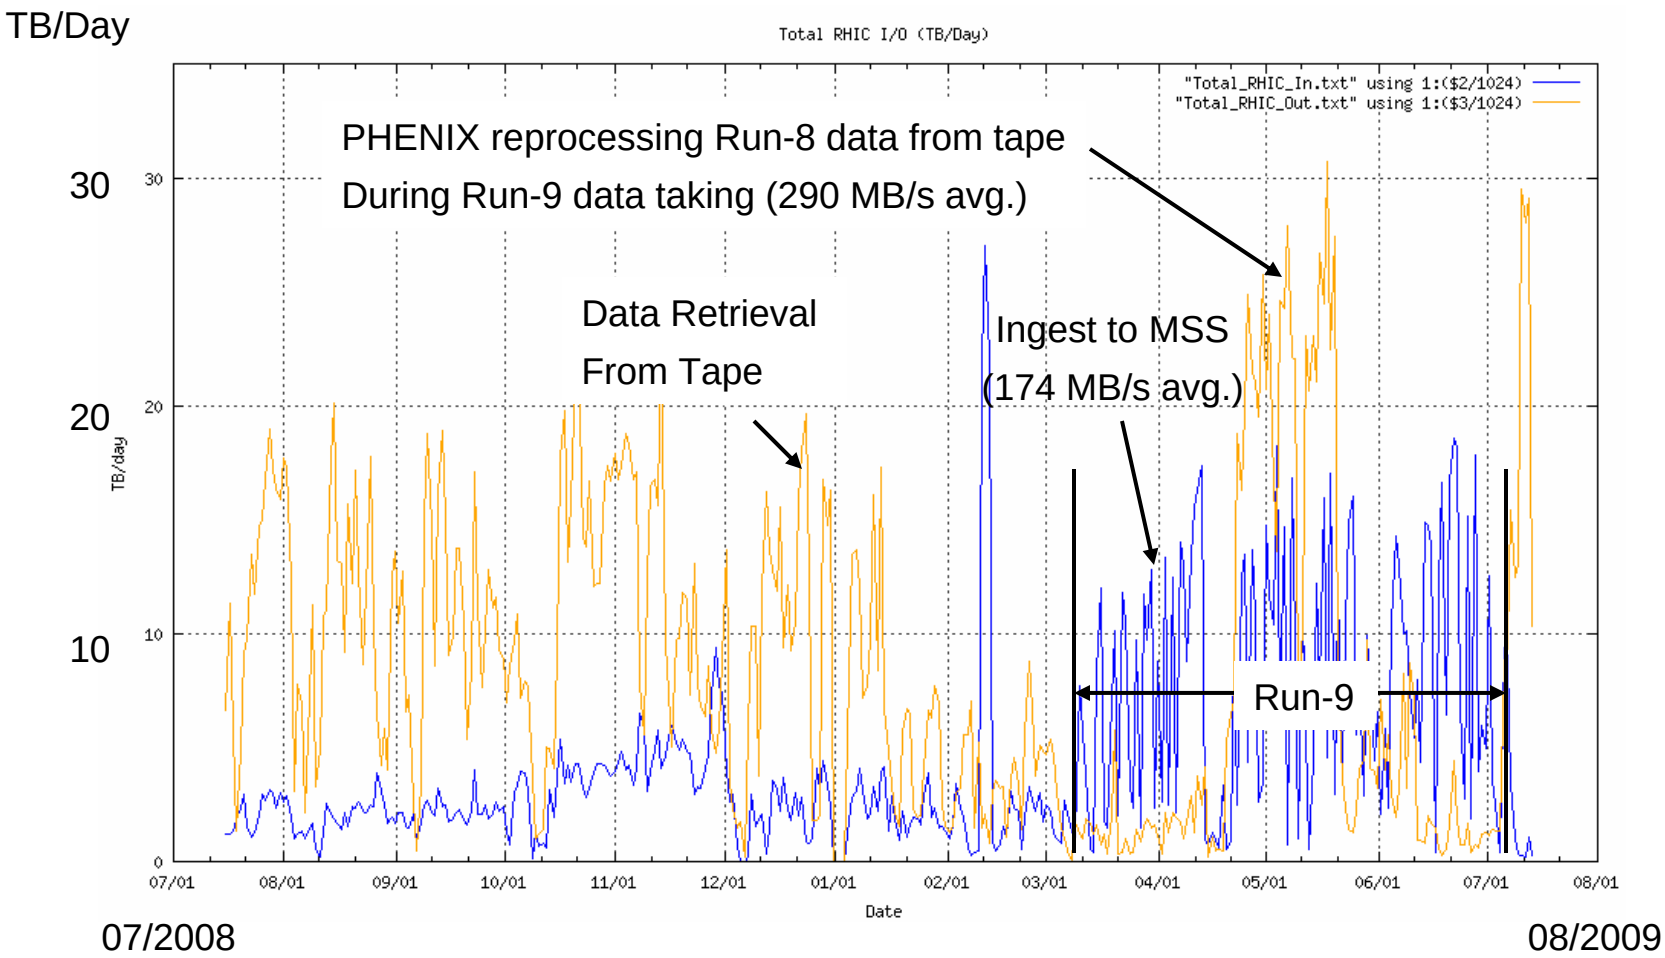

- For the first time overlapping Reads & Writes from/to Tape at high rate

Mass Storage (HPSS) Disk Cache Issues

- Disk Cache acts as an intermediate dataflow adaptation layer between tape drives and distributed storage of PHENIX and STAR compute farms
- Would need 3x more tape drives w/o disk cache
- All 5 RAID systems have reached End Of Life
  - In addition, increasing performance requirements exceed existing capabilities
- 13 out of 20 servers have reached End Of Life
- With LTO4 streaming rate >100 MB/s inter-mover network needs upgrade to 10 GE
- Needed FY09 capital funds for important infrastructure upgrades (network switches, servers and central storage)
  - If we were to purchase the disk cache from remaining funds would leave us with inadequate funds for experiment computing expansion (FY09 and FY10), leading to production competing with analysis

New RCF Network – “Weather Map”

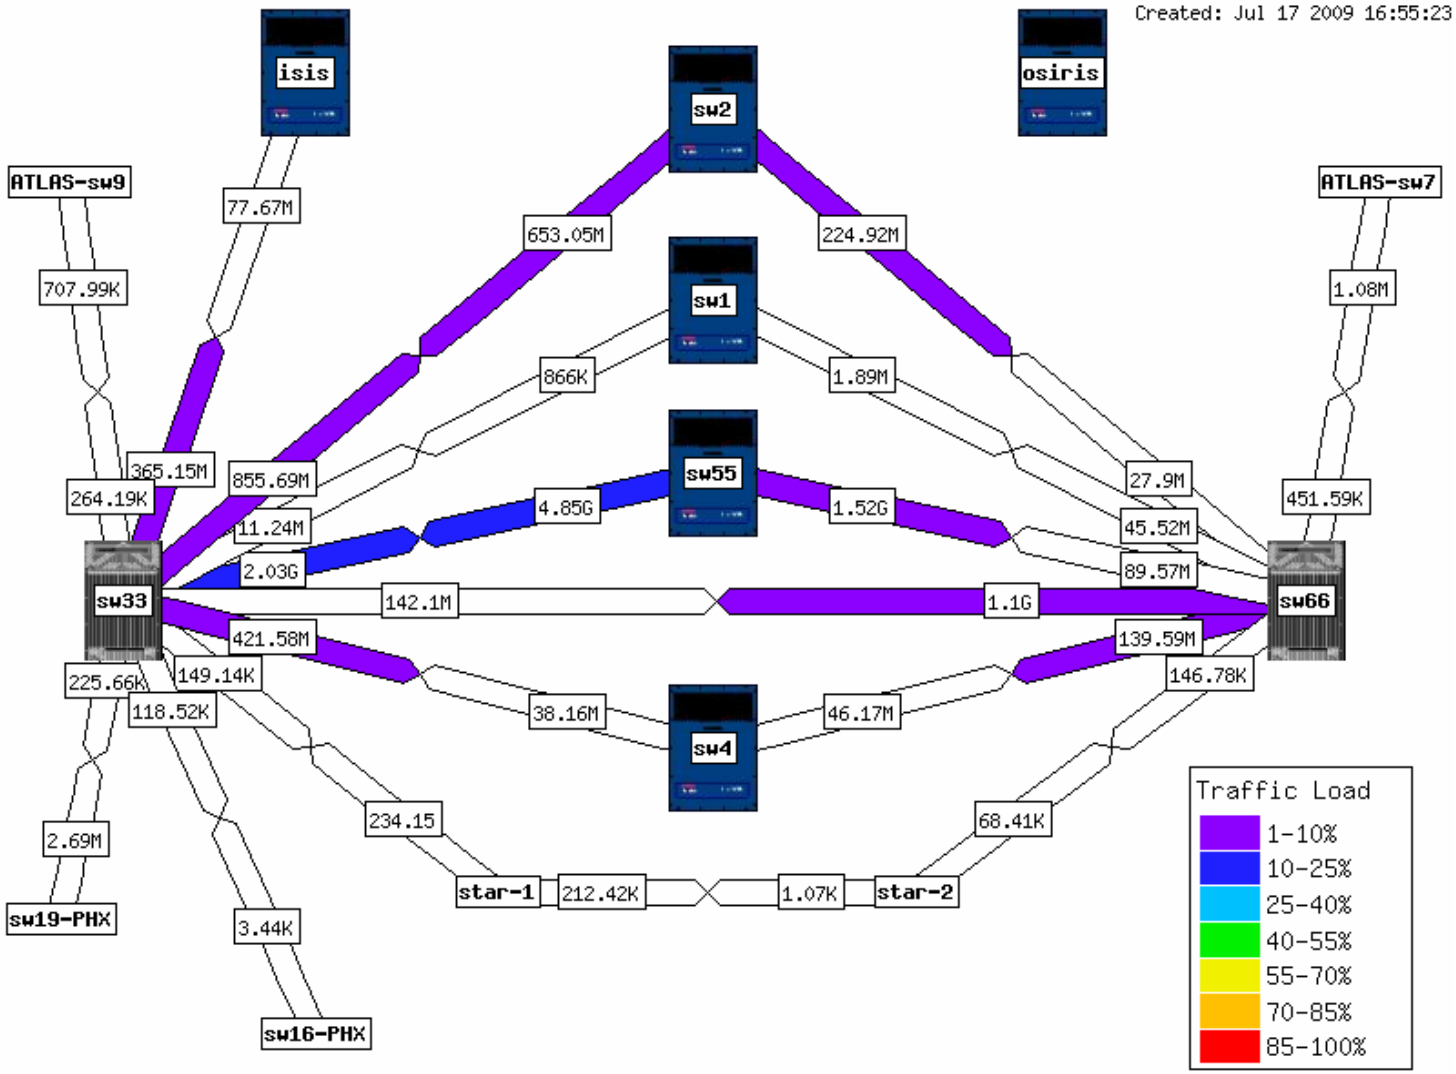

Linux Farm Highlights in 2009

- RHIC computing capacity as to number of computing cores increased by 30%.
- Computing capacity in SI2K increased by 60%.
- Local storage capacity in Linux Farm increased by 58%.
- Commercial LSF batch system finally completely replaced by Condor in 09/08 (saving ~$100k in operating)
- Virtualized PHENIX interactive nodes in 2009 to increase CPU efficiency.

Condor Occupancy (PHENIX and STAR)

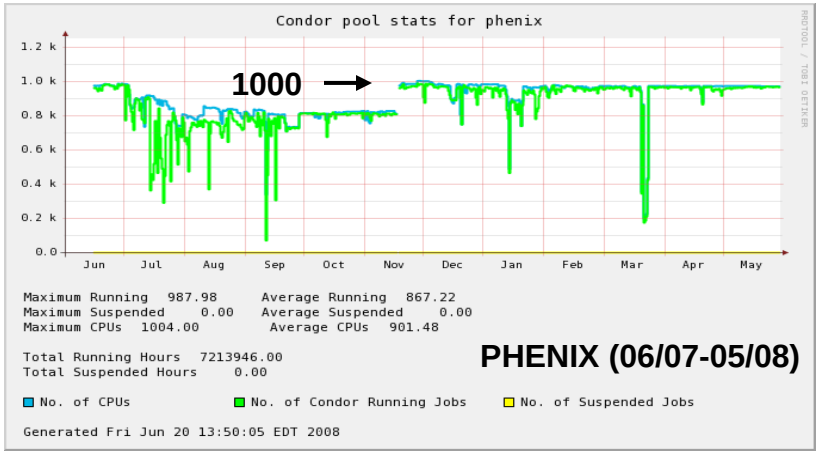

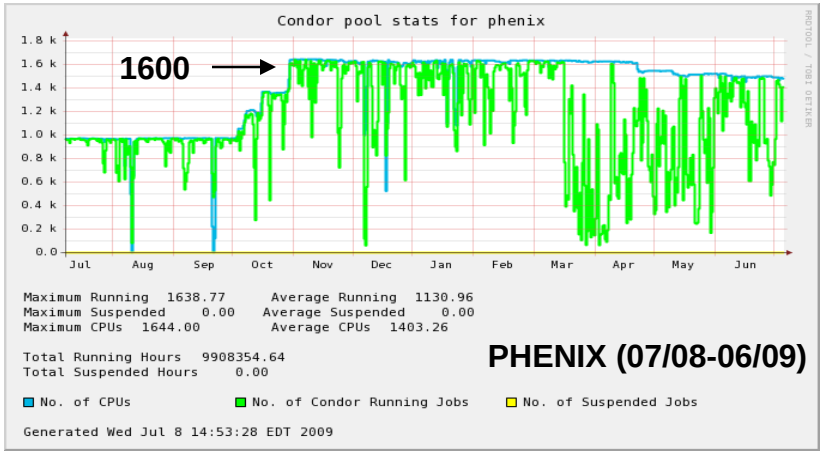

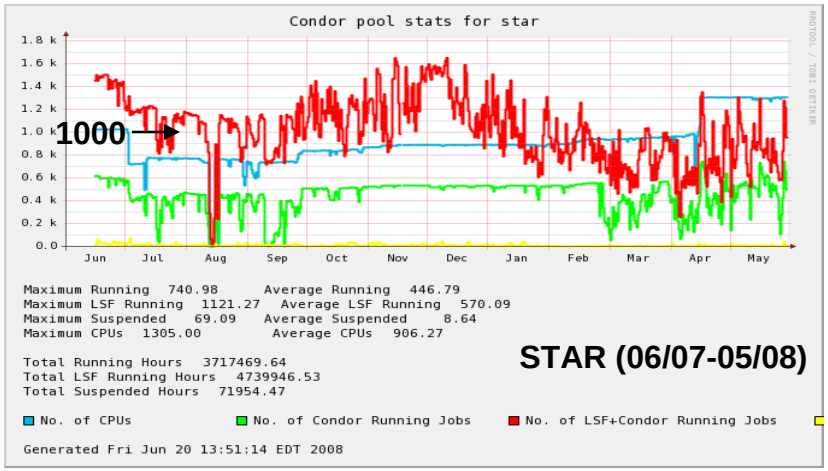

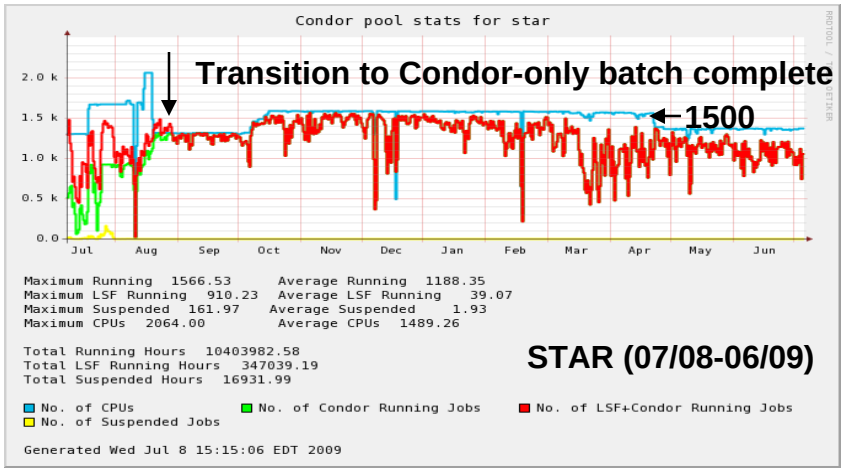

Condor Usage (last 12 months)

Most Visited | Red Hat, Inc. | Red Hat Network | Support | Shop | Products | Training

Statistics from 07-09-2008 to 07-09-2009

no. jobs completed

source \ destination	phenix	phobos	star	brahms	atlas	rcf	total
phenix	4512424	1095239	81850	892842		247747	**6830102**
phobos	100565	825745	1102	27812		4442	**959666**
star	892356	696444	9892692	461492		325825	**12268809**
brahms	1718	354	31423	70345		1356	**105196**
atlas					973188	21	**973209**

no. jobs evicted before completion

source \ destination	phenix	phobos	star	brahms	atlas	rcf	total
phenix	338084	37834	5151	26840		22916	**430825**
phobos	4749	21032	1	1276		651	**27709**
star	23860	16594	292524	8131		52724	**393833**
brahms	23	25	6	1488		41	**1583**
atlas					39600		**39600**

total effective runtime hours consumed by completed jobs

source \ destination	phenix	phobos	star	brahms	atlas	rcf	total
phenix	8253684.81	521375.16	24807.8	253625.74		90192.38	**9143685.89**
phobos	137631.49	2310210.24	150.5	9757.39		3292.23	**2461041.85**
star	152563.63	163842.6	7851764.71	62533.21		46243.17	**8276947.32**
brahms	590.3	311.88	90811.49	90679.27		716.75	**183109.69**
atlas					2042672.27	1.54	**2042673.81**

total ineffective runtime hours consumed (including jobs removed)

source \ destination	phenix	phobos	star	brahms	atlas	rcf	total
phenix	617199.19	142484.28	16163.74	126631.28		69359.36	**971837.85**
phobos	54874.33	196376.5		10622.72		11041.13	**272914.68**
star	43086.32	31240.75	790039.99	15964.74		32753.52	**913085.32**
brahms	21.34	48.08	80.01	20699.99		362.77	**21212.19**
atlas					133125.07		**133125.07**

Done

- Condor usage by RHIC experiments increased by 108% (in terms of number of jobs) and decreased by ~ 4% (in terms of CPU time) since 2008.
  - Increase in number of jobs due to STAR migration to Condor.
- About 33% of PHENIX and 19% of STAR jobs used general queue.
- General queue jobs amounted to 24% of all RHIC Condor jobs during this period (37% in the previous 12-month period).
- General queue efficiency ~89% (ie, rate of ineffective run time is ~11%).

STAR relies on External Resources

- Computing plan relies on external resource to supply:
  - Additional analysis passes
  - Simulation needs
- Simulation
  - Grid resources suitable for basic simulation
  - Exploitation of Cloud Computing could cover for complex simulations
- Key point for dedicated facility
  - User analysis support
  - Permanent storage for archive
  - Regional support
  - Distributed knowledge / no single POF

STAR Cloud Computing activities

J. Lauret

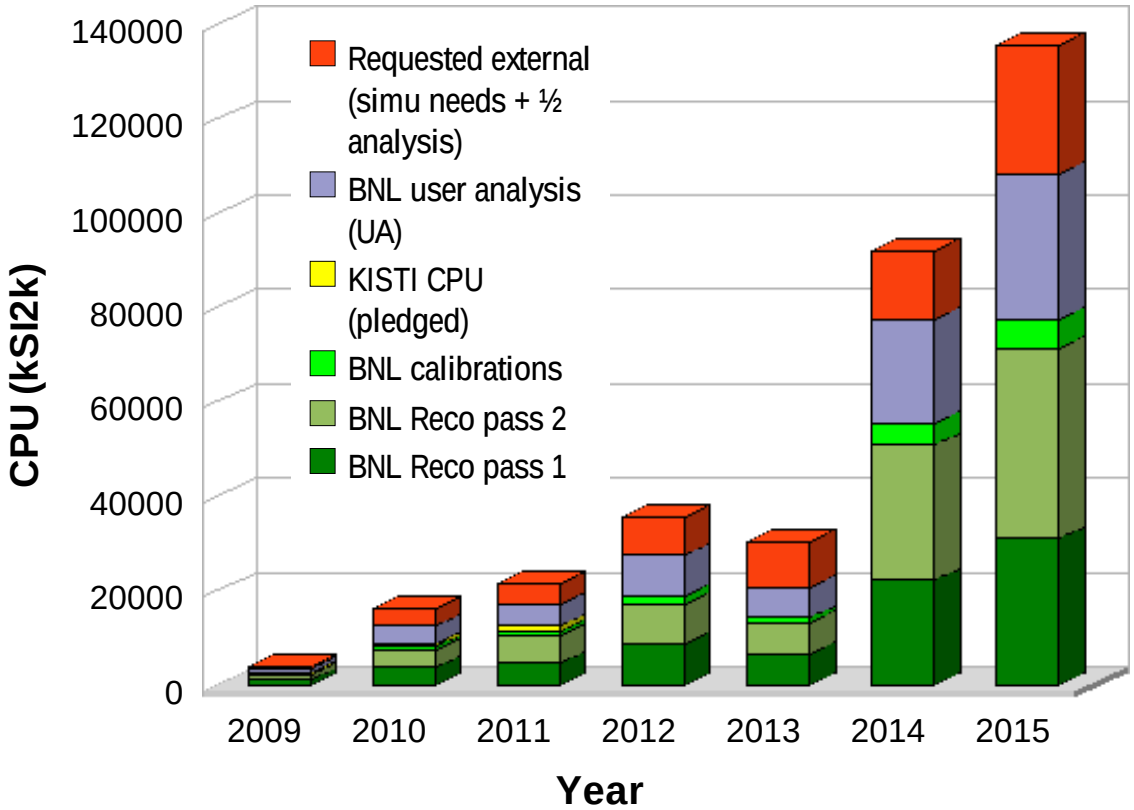

May 29th 2009 - Nimbus cloud project saves brainiacs' bacon
May 2nd 2009 - Number Crunching Made Easy
April 8th 2009 - Feature - Clouds make way for STAR to shine
April 2nd 2009 - Nimbus and cloud computing meet STAR production demands (HPCwire)
August 13th 2008 - Feature - STAR of the show
April 30th 2008 - The new Nimbus: first steps in the clouds

STAR: Storage outlook - Analysis

- Large storage needs
  - Cost prohibitive for Central Disk (CD) solution
  - Inexpensive Distributed Disk (DD) managed by Xrootd
- STAR's target model: light green
- To be expected
  - 2009 no change
  - 2010 no increase on CD, Monte Carlo MUST move to DD
  - 2011 DD used as production buffer
  - 2012 end of storage evolution (no further expansion of CD expected)

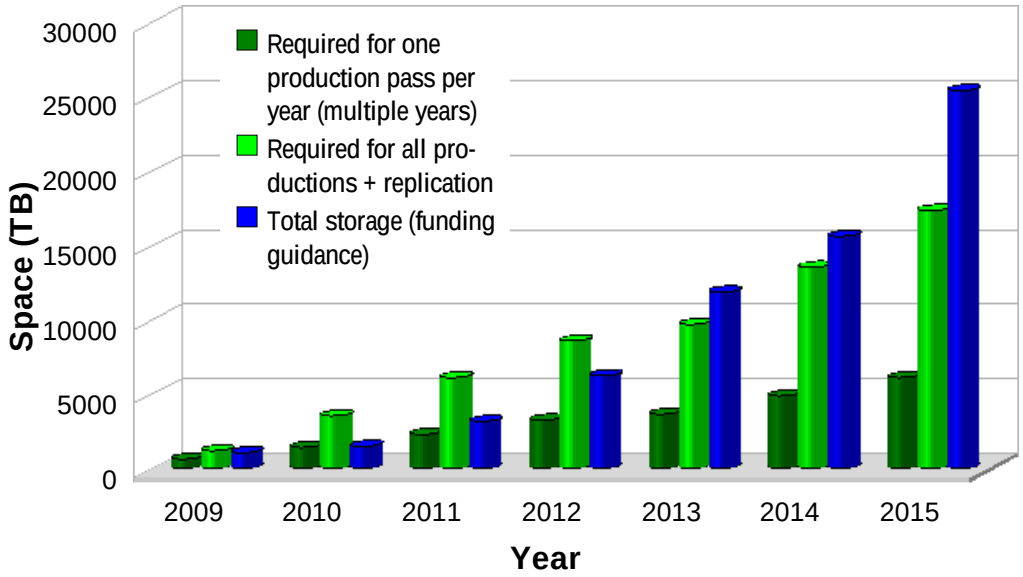

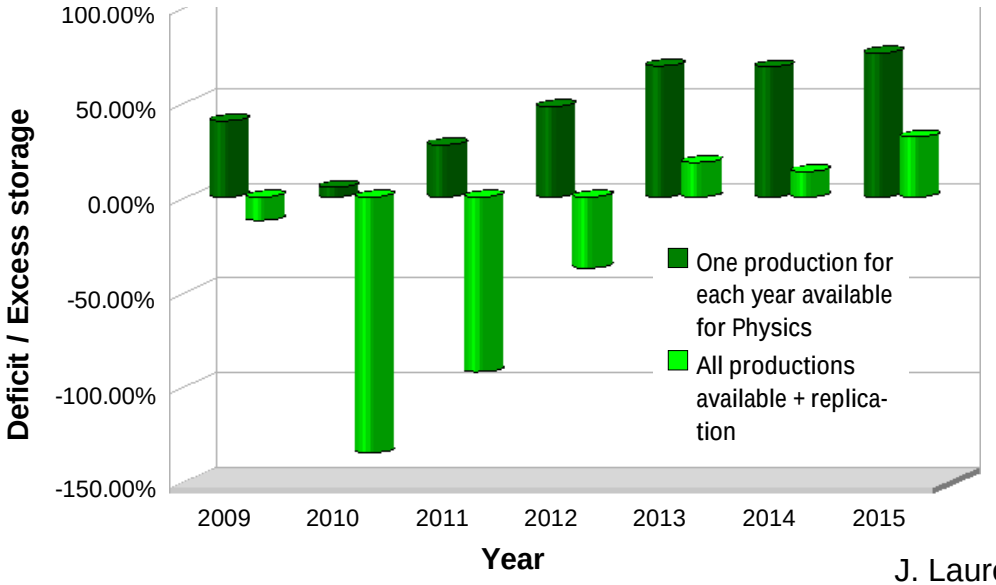

J. Lauret

Progress on Physical Infrastructure for RACF

1. Renovation of existing area of ~2000 sq. ft. adjacent to the computing facility, ready for occupancy in Oct 2008 $1,200k
   - ~~Funding lined up, construction has started~~ Complete & Occupied
2. Data Center Expansion (6,400 sq. ft.), ready for occupancy in July 2009 ~Complete $4,750k
   - ~~Funding lined up, design complete, construction to start ~Oct. 2008~~
3. Purchase and Installation of UPS (1 MW), needed in 2009 $1,250k
   - ~~RACF getting 300KW of UPS Power from NYBlue to cover 2008/9 needs~~
   - ~~Received $950k of supplemental funds from DOE/HEP~~
   - ~~Rest will be paid out of ATLAS Program Funds~~ ~Complete
4. Power upgrades to allow full occupancy in 2011 – 2013 $4,100k
   - Unstable Power on LI requires UPS power for reliable operation of central services and combined compute & storage servers
   - Flywheel and Diesel Generator have long lead times (~18 months)

RACF Space and Deployment Planning

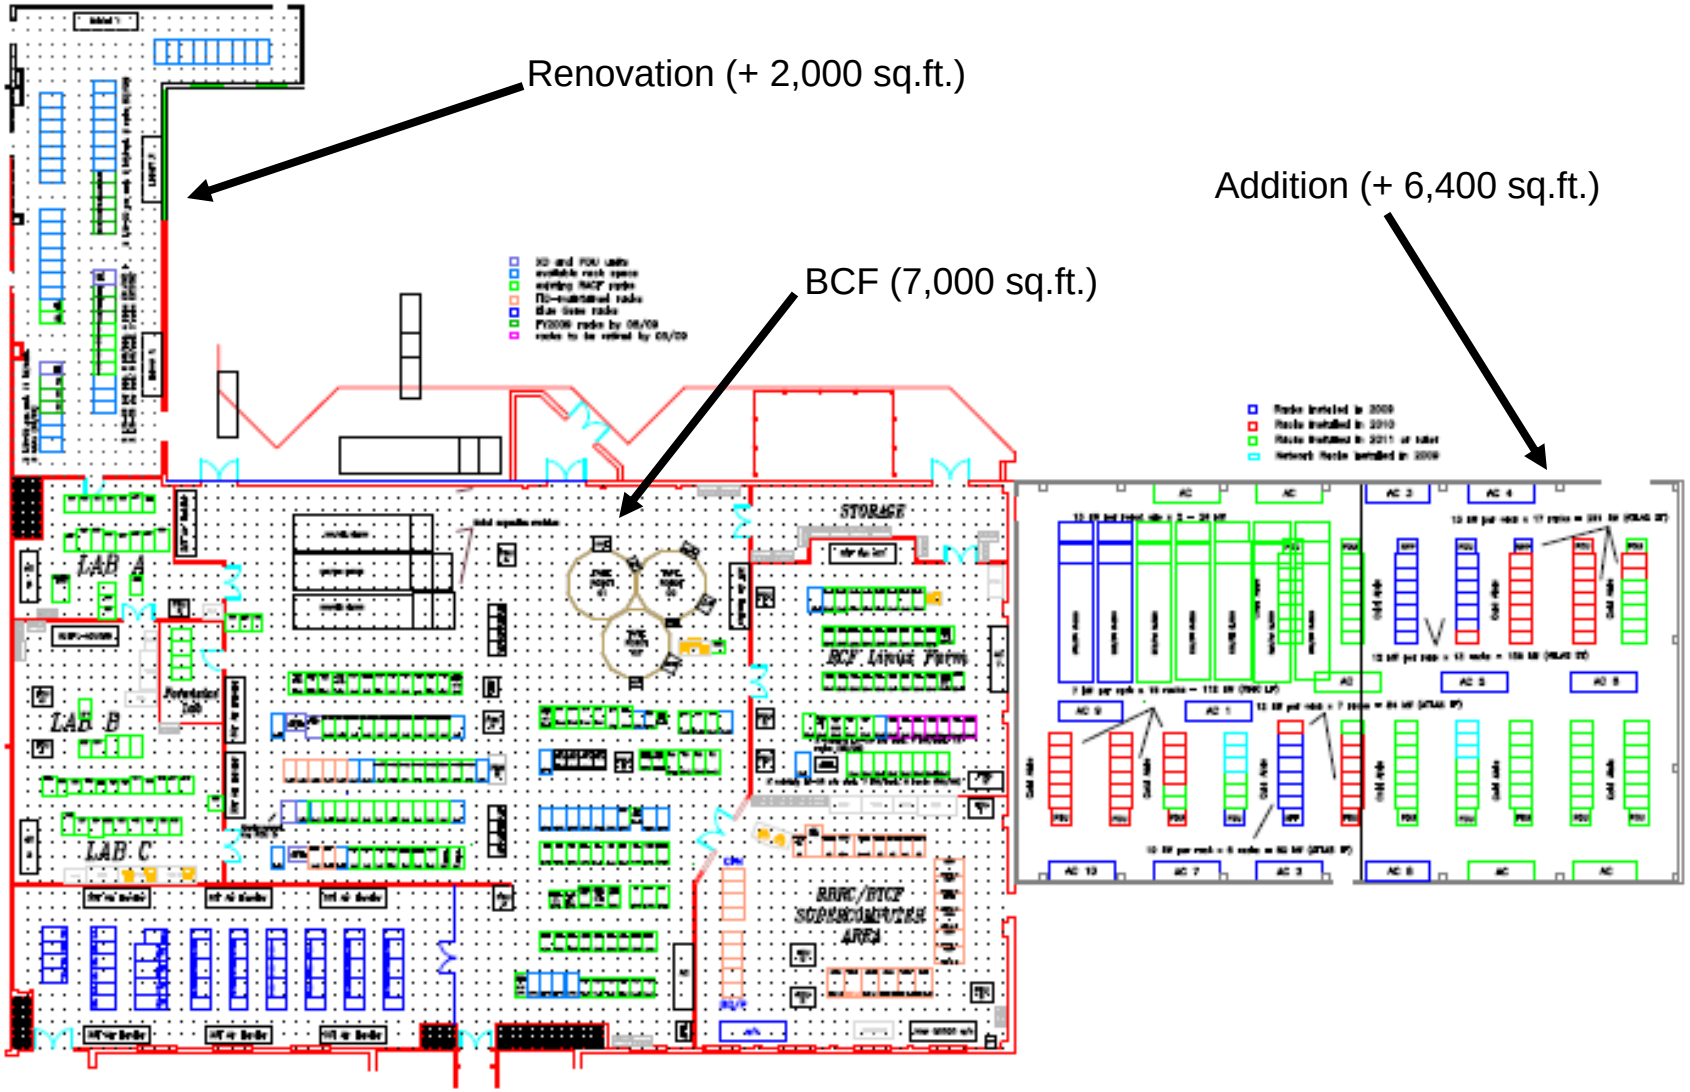

New Data Center (available 08/2009)

RACF Equipment Space Projection

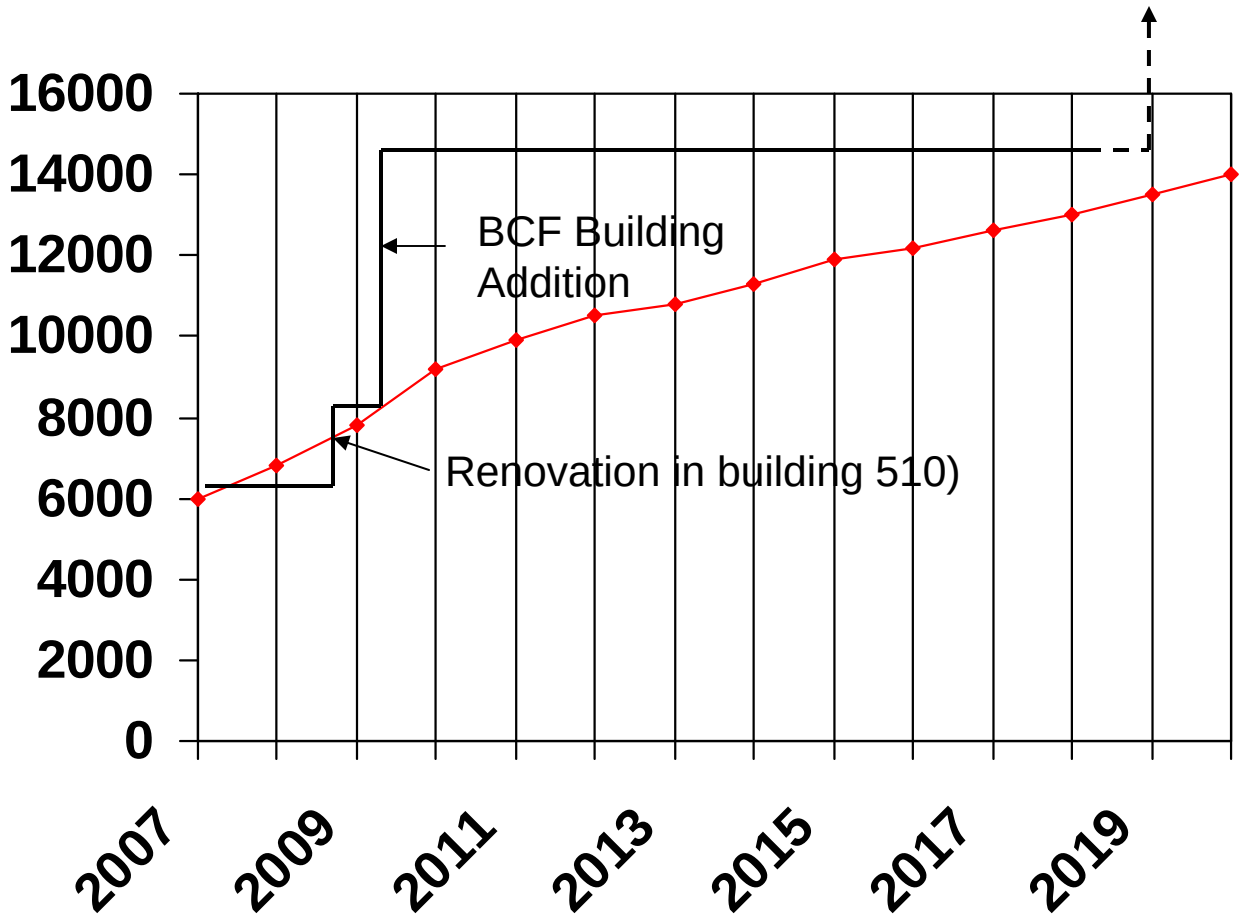

Space available to RACF in BCF (7k sq.ft.)
Renovation of 2000 sq.ft. in 510 and CDCE in 2009 (6400 sq.ft.)

Space (sq.ft.)

This will suffice RACF requirements until ~2018 *provided*

- 15k sq. ft. remain available to RACF
- There is no other community at BNL that requires space in the Data Center

… Provided we can increase the Power Density in BCF to ~300W/sq.ft.

RACF Equipment Power Usage Estimate

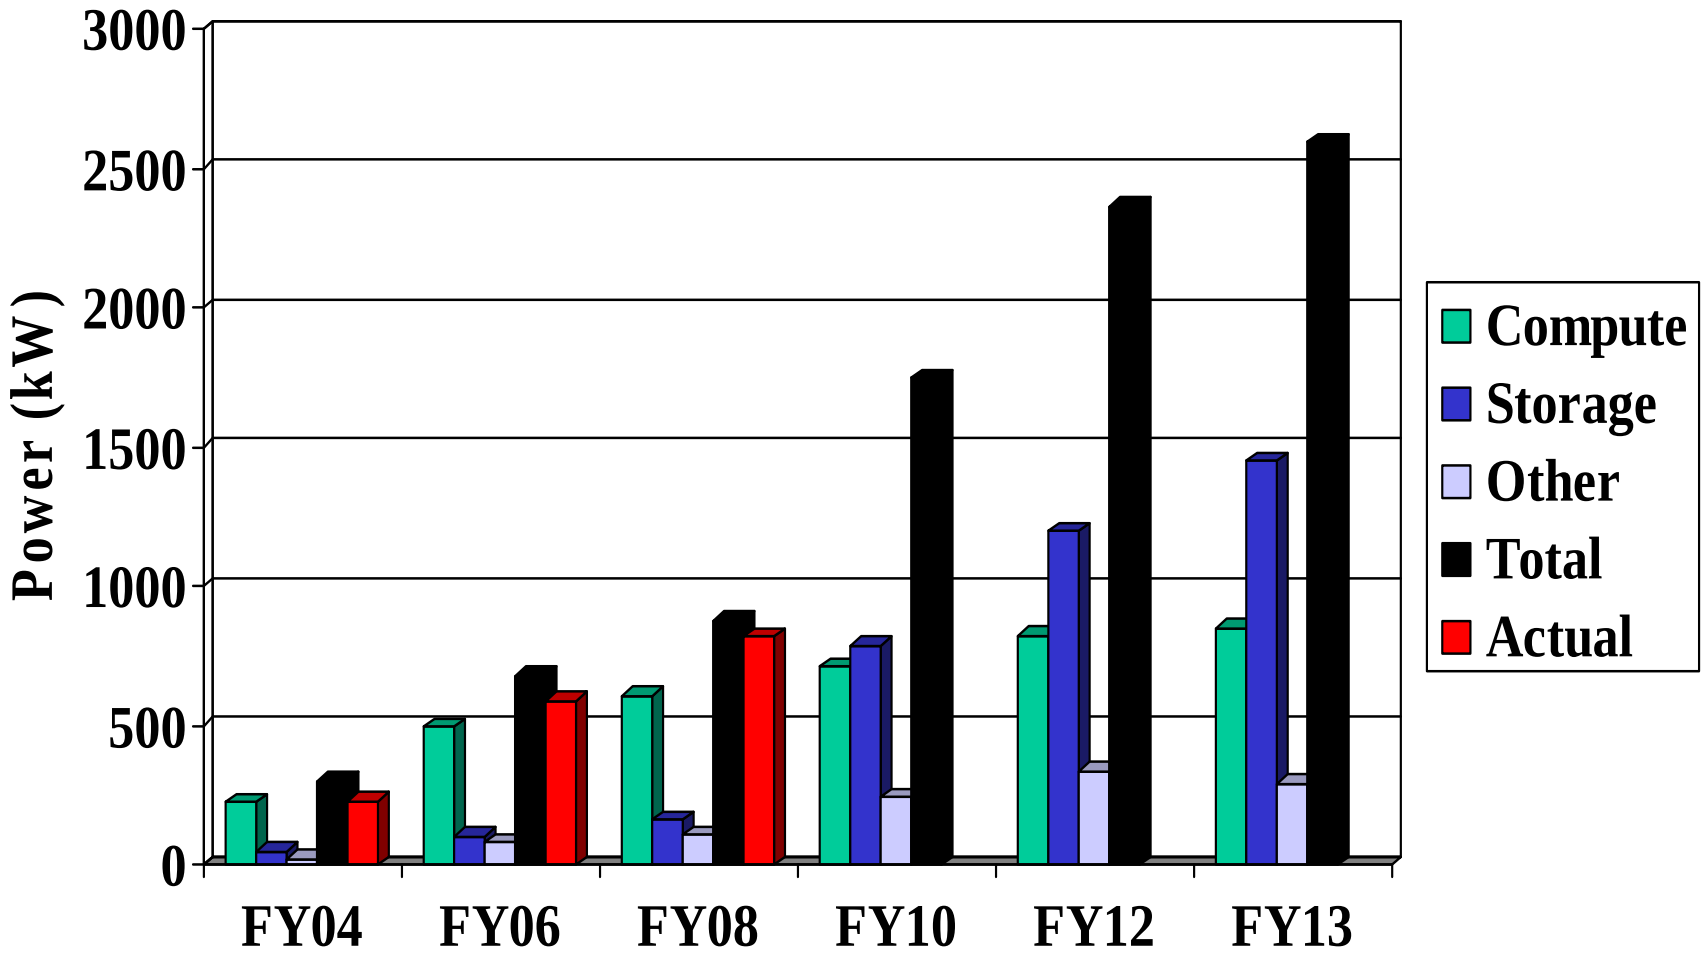

NB: This is Equipment Power only – Cooling is not included

RACF Long Term Power Estimates

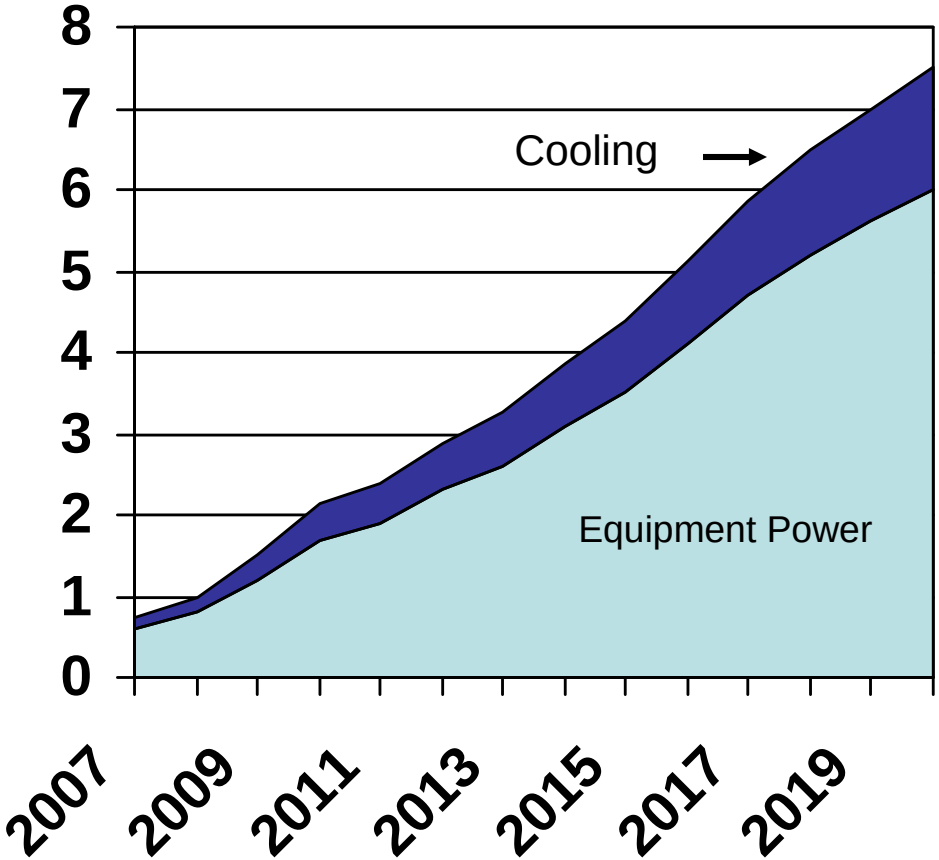

Computer Equipment and Cooling Power Requirements (MW)

NB: The development beyond 2013 implies substantial uncertainties and are to be considered rough estimates, The numbers include a contingency of 30%.

Is UPS Power Required ?

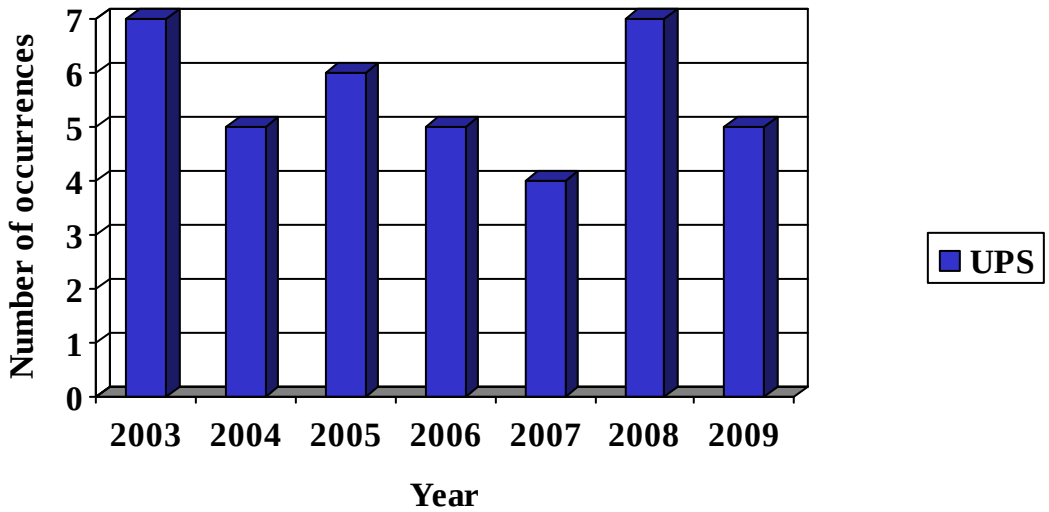

Number of instances the RACF power was switched over to the UPS system since the UPS monitoring system went into operation in March 2003. Data for 2009 is through June

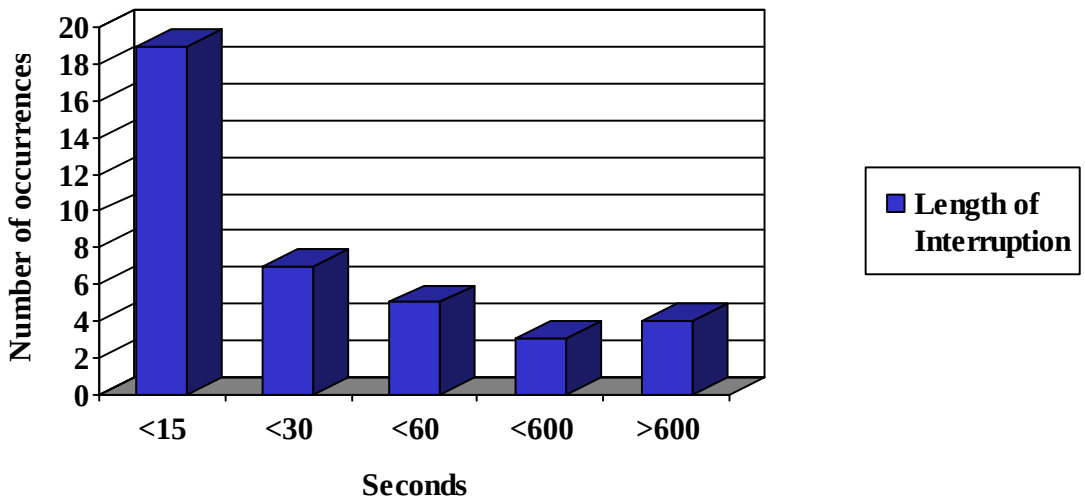

Length of utility power interruption (in seconds) when power supply was transferred to UPS, as measured by the UPS monitoring system

- With ~6 power failures/yr causing disks to fail w/o UPS this would cause a loss of ~16-19 days or 4%-5% /yr of analysis resources

Summary and Outlook

- In support of PHENIX and STAR Data Taking and Data Analysis the RHIC Computing Facility has provided stable and high performance processing and storage services
- Plans to evolve and expand facility services to meet expected needs
  - Are based on successful adjustments of technical directions
  - Remain within the mainstream of NP and HEP computing
  - Requires agreed and planned for increases (capital and operating) in 2010 and beyond
  - Drive down cost for operating and computing support
    - Improved instrumentation & monitoring and problem resolution efficiency
- Funding tight in FY09 and FY10
  - Request to re-program savings made on operating funds to address shortfall with equipment funds
  - Resources external to BNL at collaborating Institutions and at institutions open for opportunistic usage via the Grid are vital to accomplish the scientific mission
  - Grid/Cloud computing technology is likely to change future RHIC computing (at least for STAR)
    - Building on OSG Middleware and support
- Physical infrastructure expansions and improvements
  - Projects to expand space by 8,400 sq.ft. were completed in time and on budget
  - Funding for Power Infrastructure needed in 2011 - 2013 an issue

Backup Slides

Archived RAW & DST 2003 - 2009

4 PB

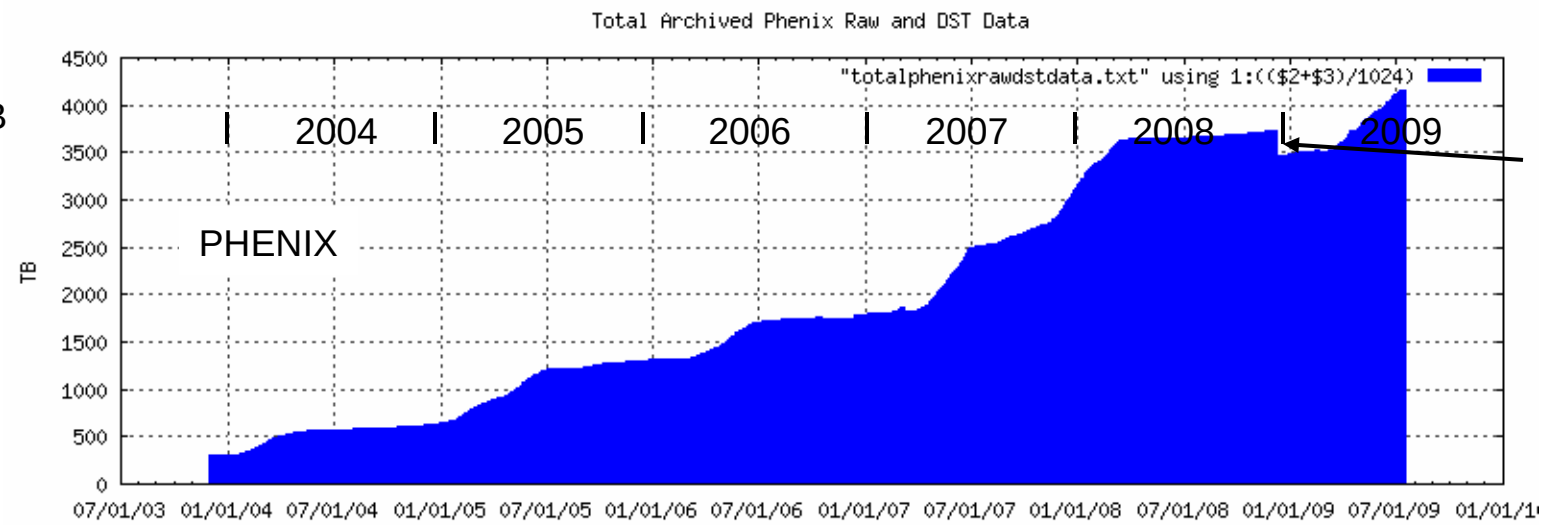

Deletion of Derived Data

2.5 PB

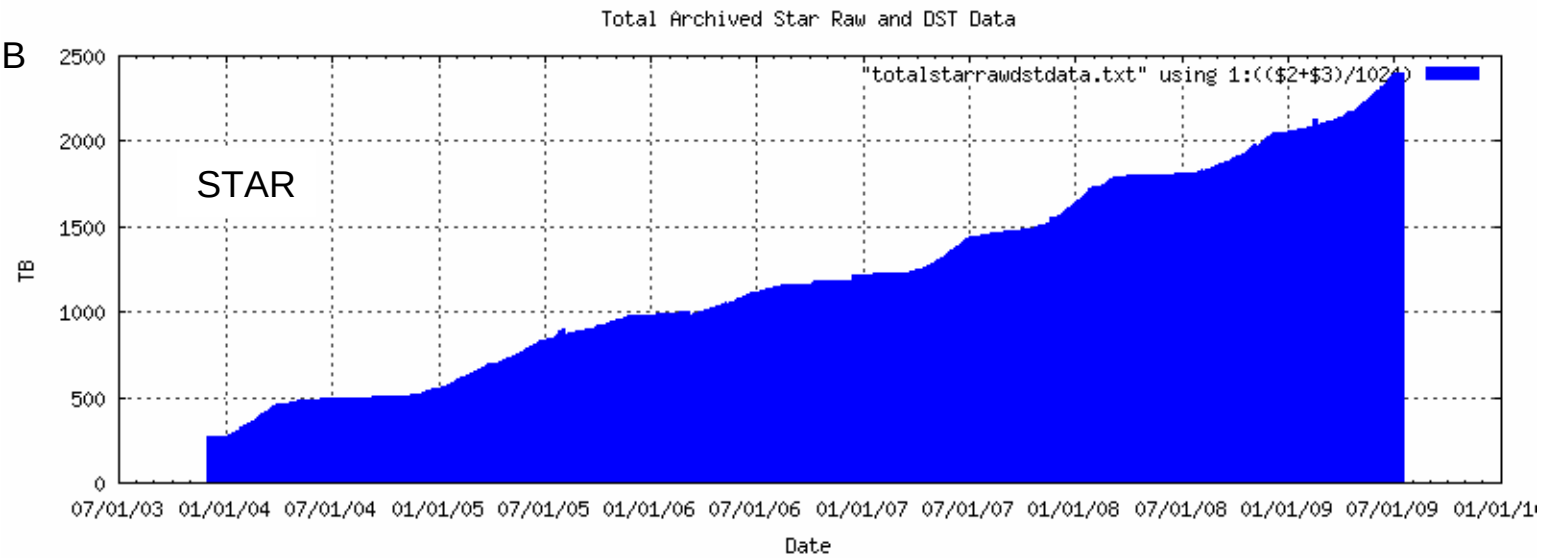

STAR Input and outlook

J. Lauret

- Run plan as input
  - Strawman for projections
  - Integrated, probably fine (devil may be in the details)
- Some good news?
  - Number of passes increased to 2.2 minimum
  - **Always assume scaling of analysis demands with production demand (1/1)**

Run	Species	Purpose	Dataset
2009	p+p 200 GeV	$\Delta G(x)$ with dijets	900 M events, 50 pb^{-1} sampled
2010	Au+Au 200 GeV	Precision Au	600 M events, 2 nb^{-1} sampled
	p+p 500 GeV	First W measurements	250 M events, 10 pb^{-1} sampled
2011	Au+Au 5-40 GeV	Energy Scan	50 M events
	U+U 200 GeV	Highly elliptical zone at high density	650 M events
2012	p+p 500 GeV	Precision W	550 M events, 150 pb^{-1} sampled
	Au+Au 200 GeV	Heavy Flavor, RHIC II	750 M events, 5 nb^{-1} sampled
2013	p+p 500 GeV	Precision W	550 M events, 150 pb^{-1} sampled
	Au+Au low E	Energy Scan	50 M events
2014	p+p 200 GeV	Au reference with HFT	2100 M events
	Au+Au 200 GeV	Heavy Flavor, RHIC II	1200 M events, 10 nb^{-1} sampled

Assumed 2015 similar to 2014 within uncertainties

2009	2010	2011	2012	2013	2014	2015
2.2	2.2	2.5	2.2	2.2	2.5	2.5

STAR: Storage outlook - MSS

- Storage needs are large [DAQ+Reco]
  - Costs prohibitive
  - Have to keep RAW
  - Potential for savings
    - **All passes, all files**
    - **One pass but all files**
    - **One pass + calib, 1/10th DST**
    - **All passes + calib, 1/10th DST**

- **Approach will have to be a combination of green and blue**
  - STAR cannot keep all DST
  - STAR will keep cost under control by deleting un-used production data

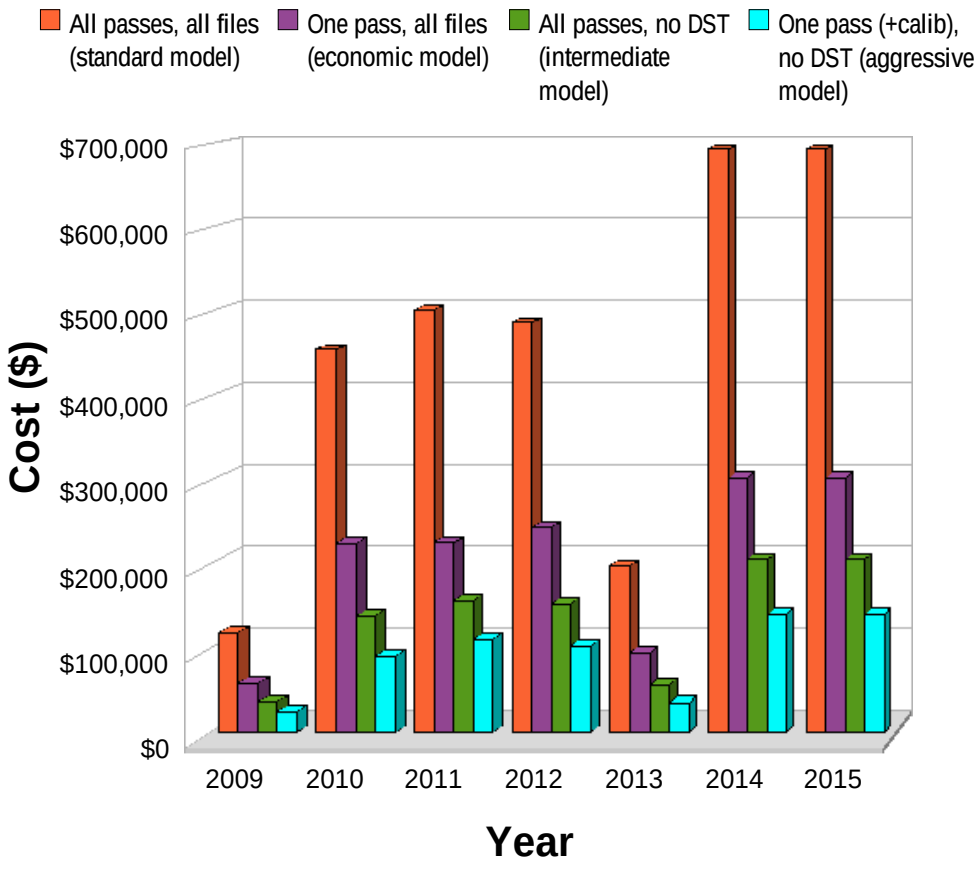

J. Lauret

STAR Networking needs

Within run plan, LAN / WAN in and out of the counting house is in between 2-3 Gb/sec

NB: transfer out ⇔ also bring produced data back.

LAN planned for 5 Gb/sec

WAN planned for 1.5 Gb/sec by 2011

WAN plan secured via ESNet long term planning (SN0466)

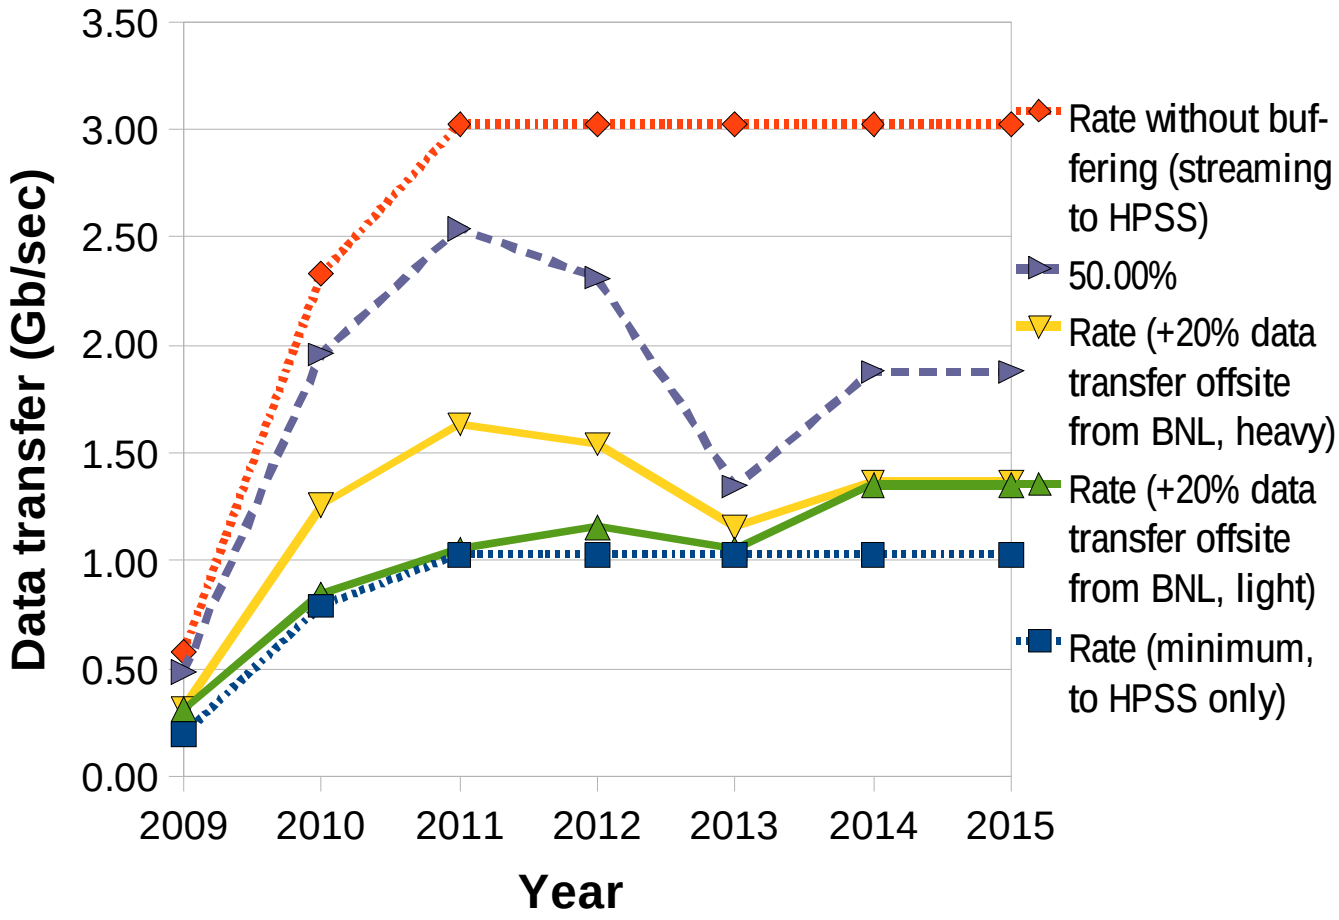

J. Lauret

STAR: Other concerns

- Noted in past years, drop of support from the collaboration trickles down to the core-team
  - Dedicated TPC software support inexistent – absorbed by reconstruction leader and calibration coordinator, essentially spending ½ their time on this area …
  - Dropped support for analysis module (code is typically dropped as core-team has no room for expansion)
  - R&D continues to draw effort

- Need to address code obsolescence, external support decrease for an increasing collaboration (number of members, number of authors)
  - A STAR paradox??
  - Institutional support direly needed – funding need to ACCEPT reasoning based on “service” for long term survival

J. Lauret

Failed Supplementary Cooling Unit

- 3 cooling units failed within 24 hours (equivalent to 30% of BCF cooling capacity)
  - Corrosion of the heat exchanger and probably over pressure has caused leaks
  - Temporary chiller (rental) was ordered, delivered and installed within 24 hours
  - RCF sustained running at full capacity
  - Took 3 months to replace the failed units