



Computing Operations Roadmap

Kaushik De

University of Texas At Arlington

US ATLAS Tier 2/3 Meeting, UM

May 27, 2008

Introduction




- ❑ Expect first data to arrive end of summer
- ❑ After 6+ years of preparing computing facilities and software systems – are we really ready?
- ❑ In this talk, I will review what is ready, but more importantly what remains to be done
- ❑ I will review summer plans (alas, not vacation plans!)

Some Recent Achievements




- ❑ Excellent MC production capacity and availability
 - ❑ Good completion rates recently (when jobs available!) – typically 0.5M fully simulated events/day (~5k CPU's)
 - ❑ Many Tier 3's ready – more joining soon
- ❑ Impressive throughput during CCRC08 tests
 - ❑ Tier 1 and Tier 2's demonstrated prompt transfers
- ❑ Analysis sites ready
 - ❑ But not much data availability at Tier 2's
 - ❑ Need more users!
- ❑ Overall, U.S. facilities in good shape
 - ❑ Production hardened, analysis ready, SRM v2.2 (in record time)
 - ❑ But need to complete 2008 procurement


PanDA production (few weeks ago)




Production – This Year




CCRC Data Transfers – on May 21




Last 24 errors



Last 24h CCRC throughput



Last 24h throughput (should be 1GB/s)



From Simone Campana

Throughput Details

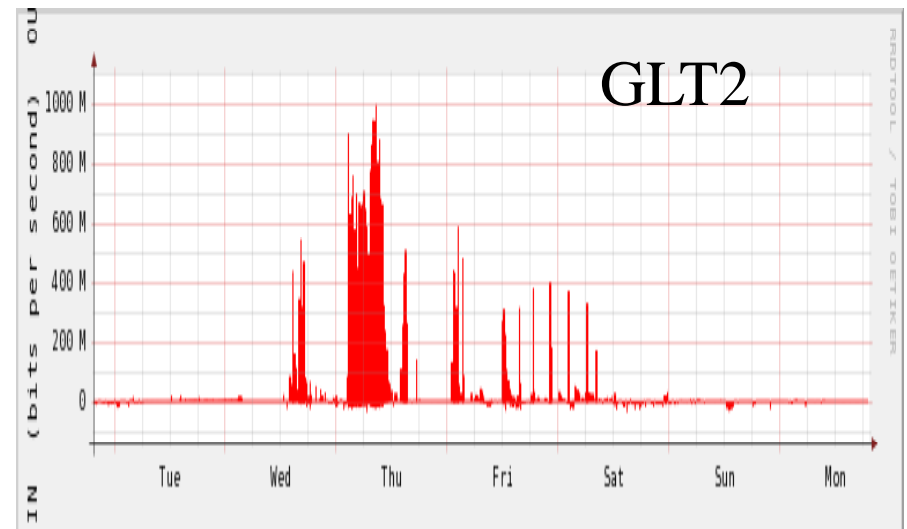
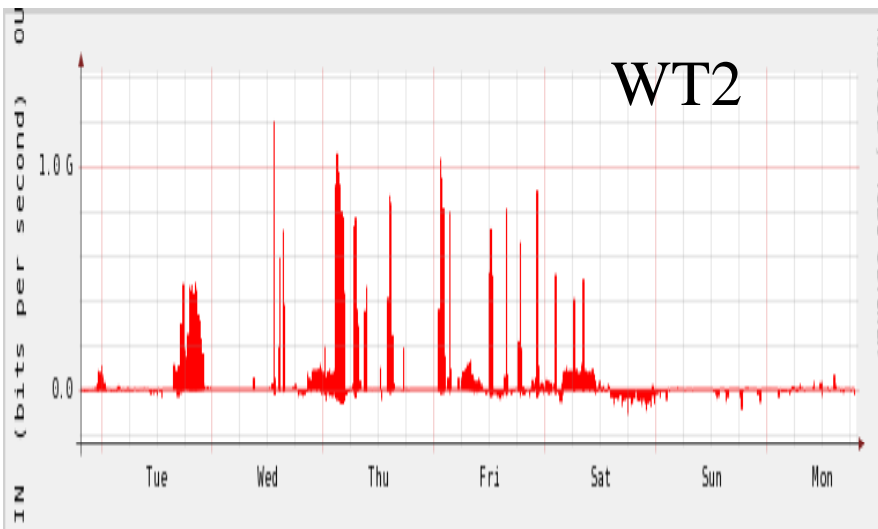




□ The rate is almost what one would expect

Cloud	Efficiency	Transfers	Rates(MB/s)	TAPE	DISK	Total
		Throughput				
ASGC	94%	51 MB/s	BNL	79.63	218.98	298.61
BNL	64%	319 MB/s	IN2P3	47.78	79.63	127.41
CERN	0%	0 MB/s	SARA	47.78	79.63	127.41
CNAF	55%	77 MB/s	RAL	31.85	59.72	91.57
FZK	85%	118 MB/s	FZK	31.85	59.72	91.57
LYON	71%	120 MB/s	CNAF	15.93	39.81	55.74
NDGF	63%	67 MB/s	ASGC	15.93	39.81	55.74
PIC	75%	60 MB/s	PIC	15.93	39.81	55.74
RAL	84%	92 MB/s	NDGF	15.93	39.81	55.74
SARA	43%	106 MB/s	Triumf	15.93	39.81	55.74
TRIUMF	79%	48 MB/s	Sum	318.5	696.76	1015.28

From Simone Campana

CCRC08 Tier 2 Netmon – Past Week



CCRC Plans – This Week



- ❑ Plan for realistic data taking exercise
 - ❑ Data transfers T0->T1->T2, assuming 14h run
 - ❑ Simultaneous MC production – all week
 - ❑ Simultaneous reprocessing from tape (M5, M6, FDR data)
 - ❑ Also involve T1->T1 transfers end of week
 - ❑ Software validation & FDR-2 mixing will continue during this period!
- ❑ Finally, full scale exercise
 - ❑ Except for distributed analysis load – chaotic but low so far
- ❑ Need help from everyone!

Operating Facilities



- ❑ Mixture of Global and Local approach
- ❑ Global Operations
 - ❑ Production shifts – many years of experience
 - ❑ Works well – now integrated ATLAS-wide
 - ❑ Primarily monitor and escalate
 - ❑ Distributed Analysis support – being organized
 - ❑ Task support – being organized at global ATLAS level (Nevski, Yu)
- ❑ Local Operations
 - ❑ Tier 1 facilities support – primarily through global support
 - ❑ Tier 1 support for Tier 2's
 - ❑ Tier 2 support – proactive site monitoring
 - ❑ Communication works well within U.S. among various teams

Are We Ready?



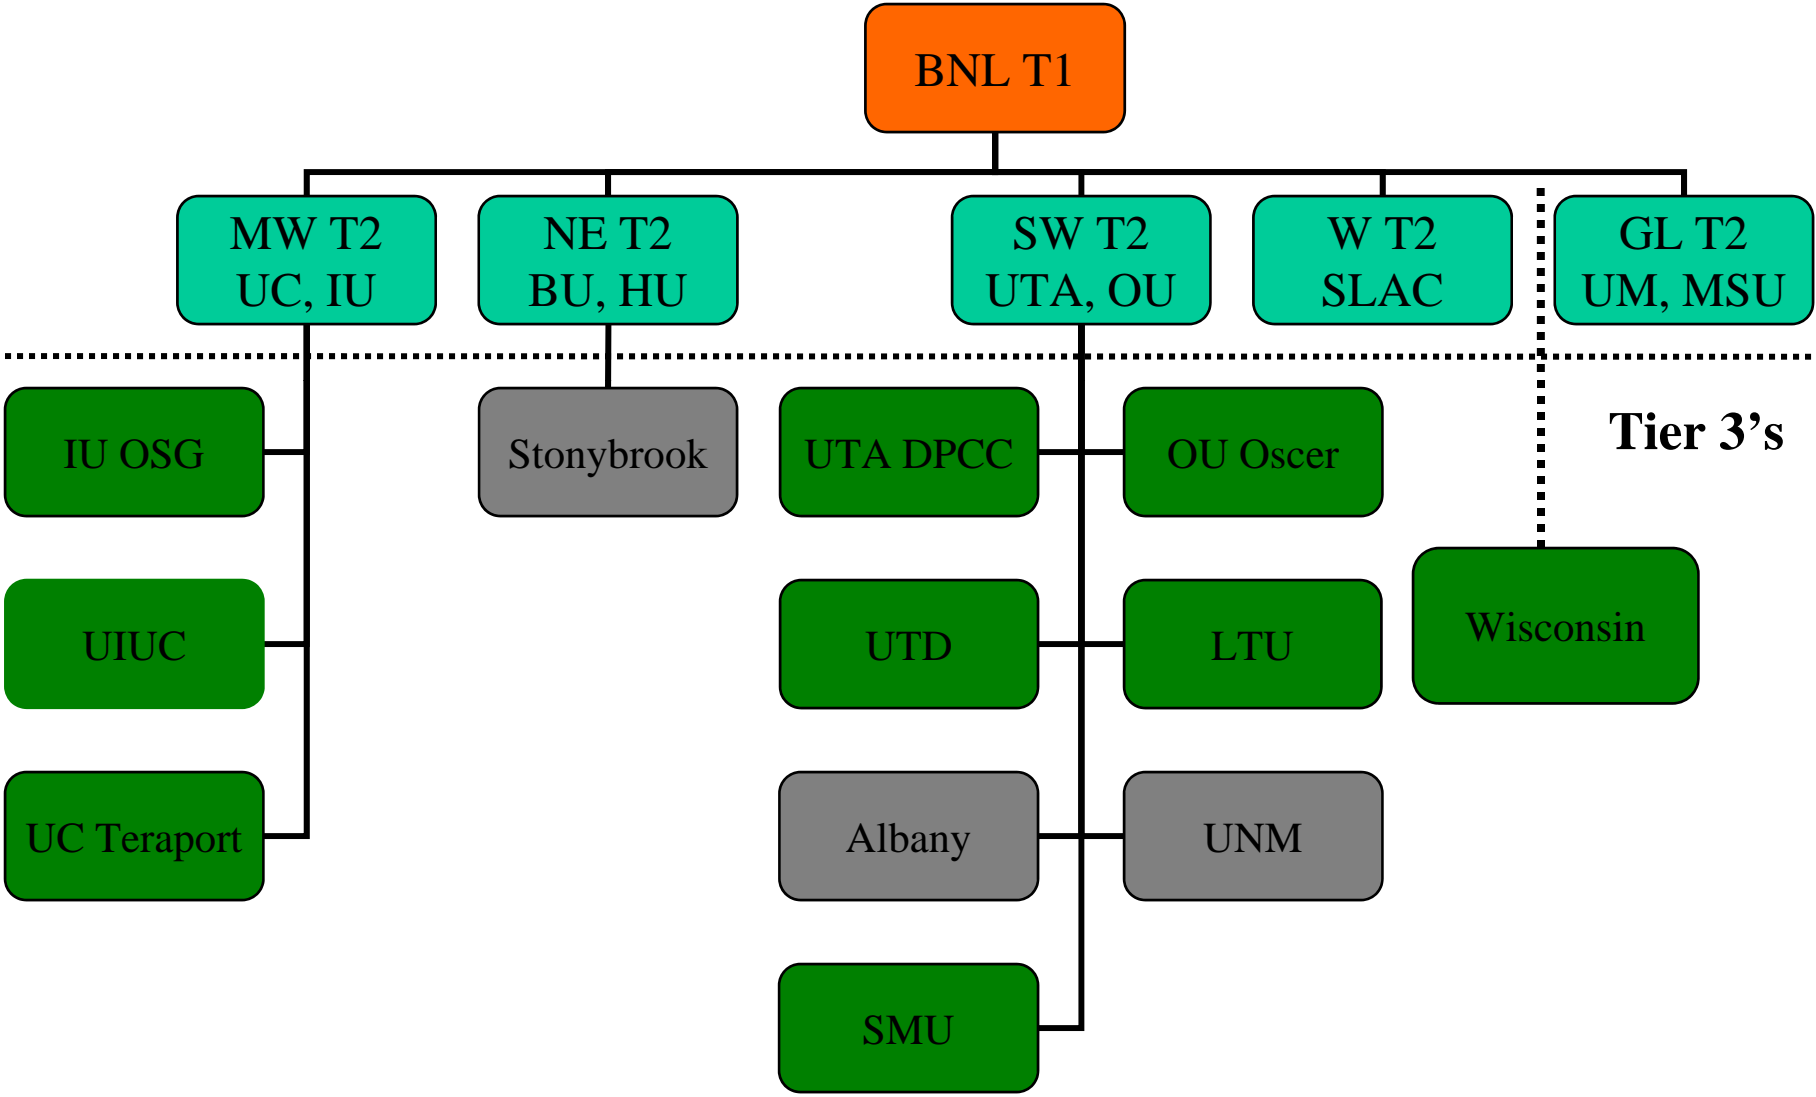
- ❑ Everything looks good – thanks to hard work by many of you
- ❑ But are we ready – not quite!
 - ❑ Distributed analysis – never exercised at scale
 - ❑ Distributed Data Management – still on critical path (more details in next talk)
 - ❑ Reprocessing – critical test this week
 - ❑ Core software – delayed and understaffed
 - ❑ Several other areas are still understaffed – transformations, job definition, Panda core ...
 - ❑ LRC vs LFC – need resolution quickly

Distributed Analysis



- ❑ All US sites are ready – but user demand not very high
- ❑ Need to implement load based brokering – now most jobs go to BNL (in computing model, most jobs go to Tier 2)
- ❑ Data management issues not fully resolved
 - ❑ Storage quotas/data archiving/data lifetime issues
 - ❑ Storage allocation for DA
- ❑ Have no experience (actually, no clue) what the scale will be – urgently need to plan for larger exercise
- ❑ Need work this summer

Tier 3 Integration



Distributed Data Management



- ❑ We can get data efficiently from T0 (mostly!)
- ❑ We can get data from some Tier 1's – and send data
- ❑ Tools to manage data hierarchy within US not ready
 - ❑ We recently decided not to rely on central operations
 - ❑ U.S. cloud will be managed by us
 - ❑ ATLAS-wide tools not fully ready – need new U.S. development
 - ❑ This will be major activity during summer
- ❑ Many storage related issues
 - ❑ Data distribution policies
 - ❑ Data deletion strategies
 - ❑ Allocation among various activities
 - ❑ Proper mixture between central vs local management ...

Reprocessing



- ❑ Stand-alone tests at BNL & Triumpf worked well
 - ❑ But many steps manual – not final configuration
 - ❑ Used pandamover for prestaging
- ❑ New tests past two weeks somewhat successful
 - ❑ At all Tier 1's – but few test jobs
 - ❑ Trying to use DQ2 (pandamover still at BNL)
- ❑ Big tests planned this week
 - ❑ New conditionDB tarball
 - ❑ Schema changes in Panda
 - ❑ Much higher rate planned
 - ❑ Tier 2's will also participate
- ❑ Still on critical path

Organization Issues



- ❑ I started as Operations Coordinator in April
 - ❑ Main goal – better integration of applications and operations
 - ❑ Hope to spend more time on planning, this summer!
 - ❑ Need input from all of you!
- ❑ Deputy Operations Coordinator
 - ❑ Armen Vartapetian – recently returned from CERN (after >2 years setting up ATLAS control room operations)
 - ❑ Will start in new role, as my deputy, next week!
 - ❑ Will be stationed at BNL
- ❑ I also plan to spend 5 weeks at BNL this summer
 - ❑ Review overall operations organization and staffing

Summary



- ❑ The glass is more than half full
- ❑ But we are not fully ready
- ❑ DDM, DA, software ... some of the unknowns
- ❑ Expect busy summer
- ❑ And then finally data!