



# GridX1 : A Canadian Particle Physics Grid

**A. Agarwal**, M. Ahmed, B.L. Caron, A. Dimopoulos, L.S. Groer, R. Haria, R. Impey, L. Klektau, C. Lindsay, G. Mateescu, Q. Matthews, A. Norton, D. Quesnel, R. Simmonds, R.J. Sobie, B. St. Arnaud, D.C. Vanderster, M. Vetterli, R. Walker, M. Yuen

CANARIE Inc., Ottawa, Ontario, Canada  
Institute of Particle Physics of Canada  
National Research Council, Ottawa, Ontario, Canada  
TRIUMF, Vancouver, British Columbia, Canada  
University of Alberta, Edmonton, Canada  
University of Calgary, Calgary, Canada  
Simon Fraser University, Burnaby, British Columbia, Canada  
University of Toronto, Toronto, Ontario, Canada  
University of Victoria, Victoria, British Columbia, Canada

# Motivation

- GridX1 is driven by the scientific need for a Grid
  - the ATLAS particle physics experiment at CERN
  - Linked to the Large Hadron Collider (LHC) Grid Project
  - BaBar particle physics application at SLAC
- Particle physics (HEP) simulations are “embarrassingly parallel”; multiple instances of serial (integer) jobs
- We want to exploit the unused cycles at non-HEP sites
  - Minimal software demands on sites
- Open to other applications (serial, integer)
  - Grid-enabling application is as complicated as making the Grid
- A number of facilities are dedicated to particle physics groups but most are shared with researchers in other fields
- Each shared facility may have unique configuration requirements



# Overview

GridX1 currently has 8 clusters:

Alberta(2), NRC Ottawa, WestGrid, Victoria(2), McGill, Toronto

Total resources >> (2500 CPUs, 100 TB disk, 400 TB tape)



## Resources:

OS: Red Hat Enterprise Linux 3, Scientific Linux, Suse Linux

VDT Globus software + PBS or Condor batch system for local resource management

External network access needed for worker nodes

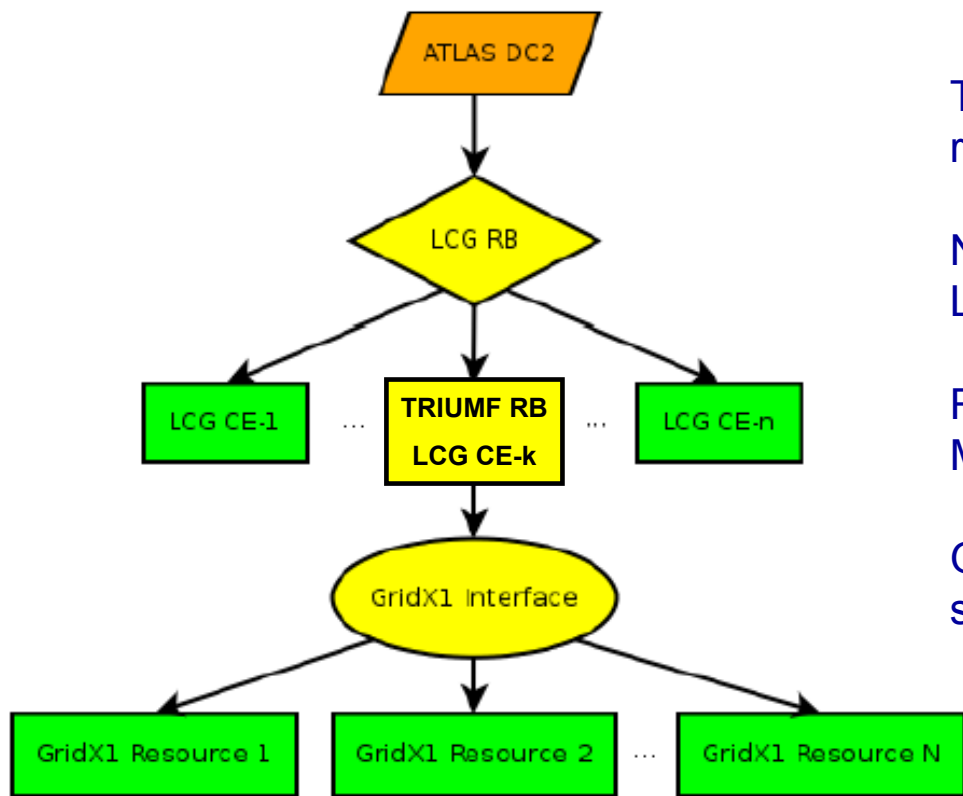
Most sites have 1Gbit/s network connectivity

# Infrastructure

- Security and User Management
  - Hosts and users require an X.509 certificate issued by the Grid Canada Certificate Authority (CA)
  - User certificates are mapped to the local accounts following a standard naming scheme gcprodx, where x is an account number
- Resource Management
  - Resources periodically publishes their state to the collector
    - Free/total CPUs; Num of running and waiting jobs; estimated queue waiting time
  - Job ClassAds contain a resource requirements expression.
    - CPU requirements, OS, application software

# Interface

GridX1 resources federated into the LCG, with all GridX1 resources appearing as a single LCG compute element



TRIUMF RB is a gateway to the GridX1 resources

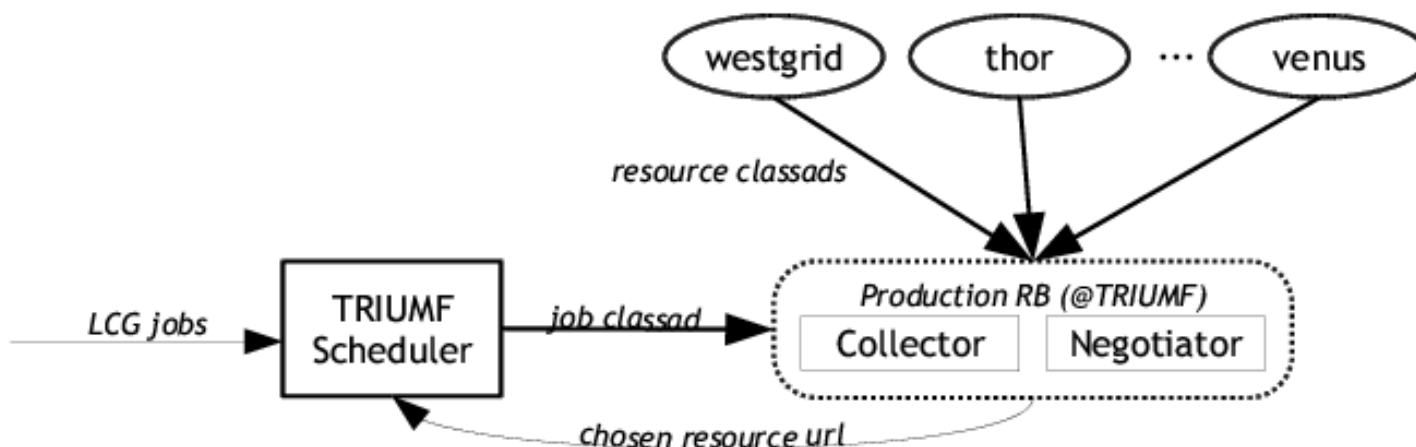
Non-LCG resources are integrated into LCG without LCG middleware

Requires full user proxy obtained from the MyProxy server

Greatly simplifies the management of shared resources

# Job management

- Each site specifies the maximum number of grid jobs, maxJobs. (100 at UVictoria)
- Job is sent to site with lowest wait time.
- Sites are selected on a round-robin basis.
- RB submits jobs to a site until number of jobs pending at a site is 10% of maxJobs



# Monitoring

home monitor publications documentation contact

gridx1logger condorg

you are here: home > grid monitor > grid x1 condor-g monitoring

## Grid Monitor Condor-G Monitoring

### Job Status

job id	user id	owner	command	resource	status	run time	time submitted
3879.0	atlas005	Guido Negri	data 'UI= ...'	mercury.uvic.ca	ACTIVE	27:17:01	9 Mar, 08:29
3899.0	atlas005	Guido Negri	data 'UI= ...'	thuner-grid.nic.ualberta.ca	ACTIVE	26:59:56	9 Mar, 09:09
3900.0	atlas005	Guido Negri	data 'UI= ...'	thuner-grid.nic.ualberta.ca	ACTIVE	26:59:56	9 Mar, 09:09
3911.0	atlas005	Guido Negri	data 'UI= ...'	thuner-grid.nic.ualberta.ca	ACTIVE	26:57:56	9 Mar, 09:11
3919.0	atlas005	Guido Negri	data 'UI= ...'	thuner-grid.nic.ualberta.ca	ACTIVE	26:55:22	9 Mar, 09:15
3924.0	atlas005	Guido Negri	data 'UI= ...'	thuner-grid.nic.ualberta.ca	ACTIVE	26:47:27	9 Mar, 09:22
3931.0	atlas005	Guido Negri	data 'UI= ...'	thuner-grid.nic.ualberta.ca	ACTIVE	26:45:20	9 Mar, 09:24
3933.0	atlas005	Guido Negri	data 'UI= ...'	thuner-grid.nic.ualberta.ca	ACTIVE	26:38:28	9 Mar, 09:25
3941.0	atlas005	Guido Negri	data 'UI= ...'	thuner-grid.nic.ualberta.ca	ACTIVE	26:41:28	9 Mar, 09:28
4090.0	atlas005	Guido Negri	data 'UI= ...'	mercury.uvic.ca	ACTIVE	6:35:07	10 Mar, 05:35
4091.0	atlas005	Guido Negri	data 'UI= ...'	thuner-grid.nic.ualberta.ca	ACTIVE	6:33:22	10 Mar, 05:36
4093.0	atlas005	Guido Negri	data 'UI= ...'	hep.westgrid.ca	ACTIVE	6:01:21	10 Mar, 05:37
4094.0	atlas005	Guido Negri	data 'UI= ...'	mercury.uvic.ca	ACTIVE	5:39:24	10 Mar, 06:31
4095.0	atlas005	Guido Negri	data 'UI= ...'	mercury.uvic.ca	ACTIVE	5:30:24	10 Mar, 06:40
4096.0	atlas005	Guido Negri	data 'UI= ...'	thuner-grid.nic.ualberta.ca	ACTIVE	5:28:51	10 Mar, 06:40
4097.0	atlas005	Guido Negri	data 'UI= ...'	hep.westgrid.ca	ACTIVE	5:25:19	10 Mar, 06:40
4100.0	atlas005	Guido Negri	data 'UI= ...'	hep.westgrid.ca	ACTIVE	5:16:46	10 Mar, 06:42
4101.0	atlas005	Guido Negri	data 'UI= ...'	hep.westgrid.ca	ACTIVE	5:16:51	10 Mar, 06:42
4102.0	atlas005	Guido Negri	data 'UI= ...'	hep.westgrid.ca	ACTIVE	5:15:46	10 Mar, 06:42
4104.0	atlas005	Guido Negri	data 'UI= ...'	hep.westgrid.ca	ACTIVE	5:15:51	10 Mar, 06:43
4105.0	atlas005	Guido Negri	data 'UI= ...'	mercury.uvic.ca	ACTIVE	5:04:23	10 Mar, 07:05
4106.0	atlas005	Guido Negri	data 'UI= ...'	thuner-grid.nic.ualberta.ca	ACTIVE	5:03:49	10 Mar, 07:06
4107.0	atlas005	Guido Negri	data 'UI= ...'	hep.westgrid.ca	ACTIVE	5:00:45	10 Mar, 07:06
4108.0	atlas005	Guido Negri	data 'UI= ...'	hep.westgrid.ca	ACTIVE	5:00:45	10 Mar, 07:06
4109.0	atlas005	Guido Negri	data 'UI= ...'	hep.westgrid.ca	ACTIVE	5:00:45	10 Mar, 07:07
4110.0	atlas005	Guido Negri	data 'UI= ...'	hep.westgrid.ca	ACTIVE	5:00:45	10 Mar, 07:07

grid monitor

- Victoria-Mercury
  - Details
  - May 28, 2005 21:16:05
  - Gatekeeper **Up**
  - Job Submission **Up**
  - GridFTP **Up**
- Victoria-Muse
- Alberta-Thor-grid
- NRC-Venus
- NRC-Mercury
- Vancouver-Westgrid
- Alberta-Thor
- Toronto-BigMac

GridX1 monitoring website:  
[www.gridx1.ca](http://www.gridx1.ca)

### Condorg Resource

Site	Active + Pending	Waiting Time
NRC-Venus	0+0	12:04:44
Victoria-Mercury	66+0	00:00:00
NRC	0+0	00:00:00
Alberta-Thor	11+0	00:00:00
Westgrid	0+0	24:00:00
Alberta-Thor-grid	26+1	00:00:00
Unsubmitted:	1	
Total:	103+1	


As of: 2005-04-05 00:13:11

### GridX1 - mercury.uvic.ca

Site info and current status for mercury.uvic.ca.

Current Status: **Up** Apr 11, 2005 13:31:04

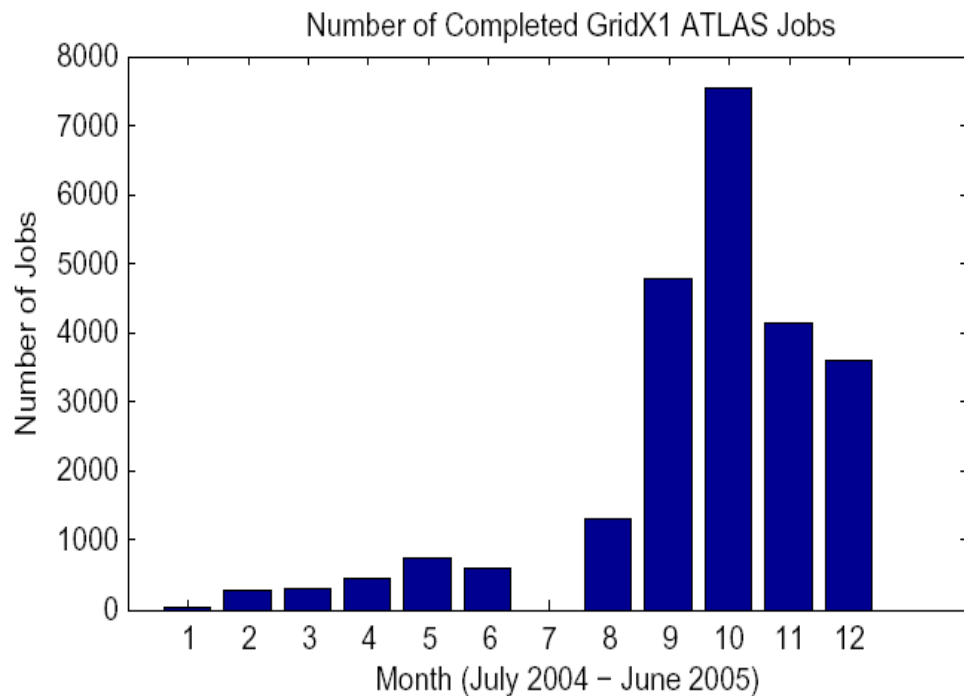
Active CPUs: 110 IP Address: 142.104.21.1  
 Max Grid CPUs: 100 Administrator: Drew Leske: dleske@uvic.ca  
 Grid Jobs: 99 Local Jobs: 1



Generated: Mon Apr 11 13:32:04 PT

View past monitoring test results for mercury.uvic.ca.  
 The current status of jobs and resources can be viewed on the Condor-G Monitoring page.  
 View the current resource ClassAd.

# Status



GridX1 used by the ATLAS experiment via the LCG-TRIUMF gateway

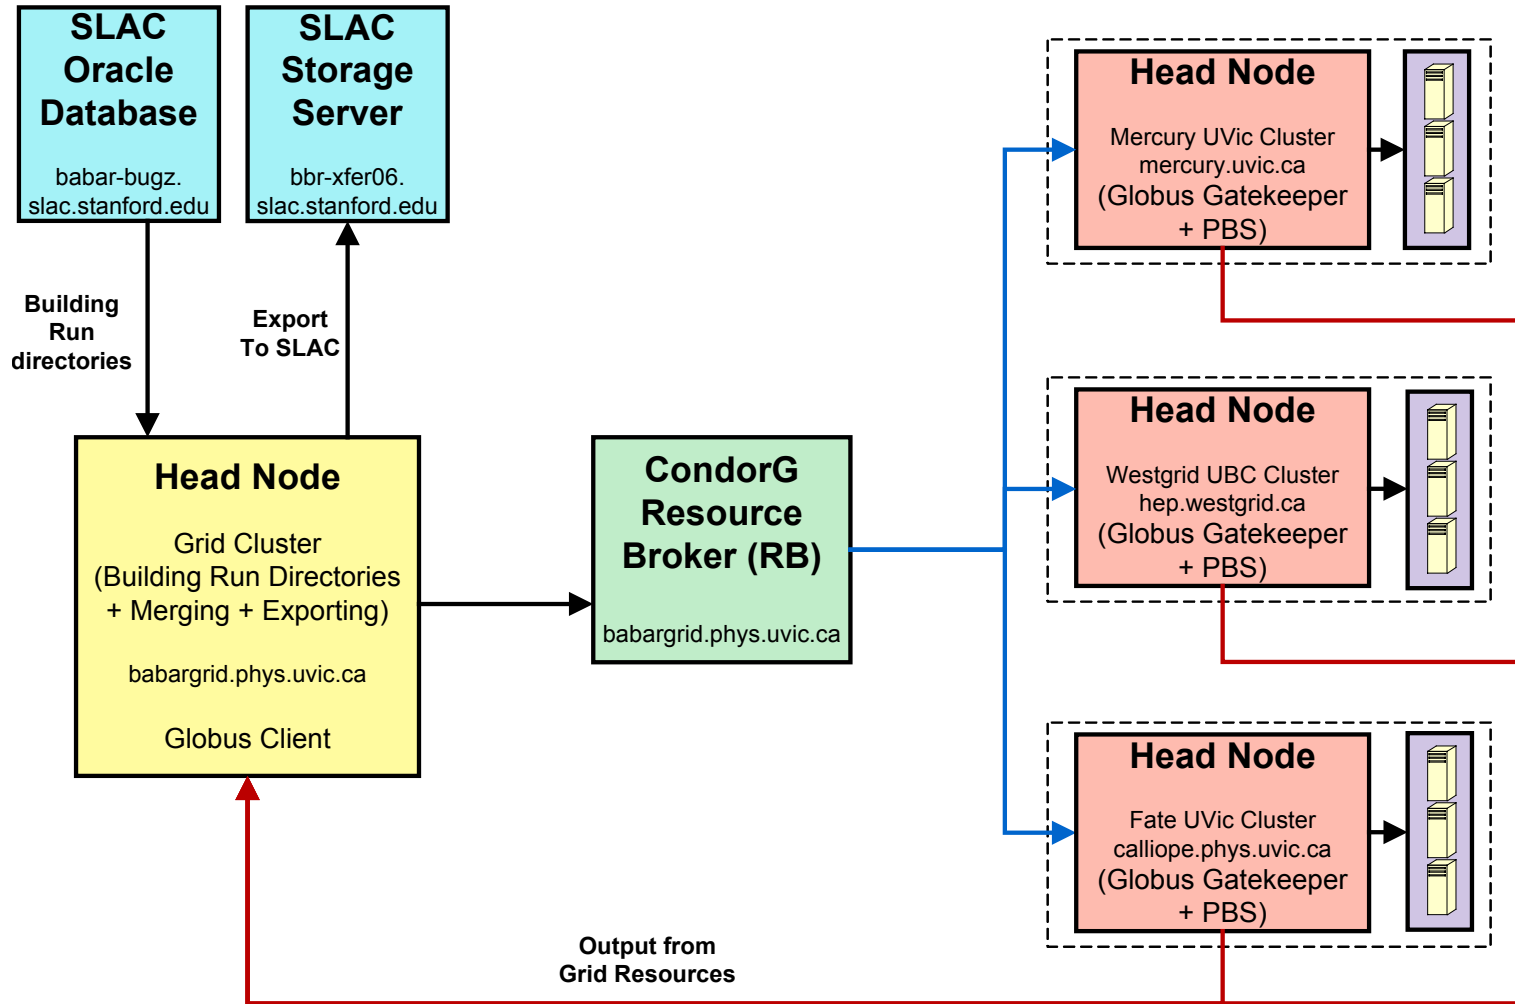
Over 20,000 ATLAS jobs successfully completed

Success rate of jobs is similar to LCG (50%)

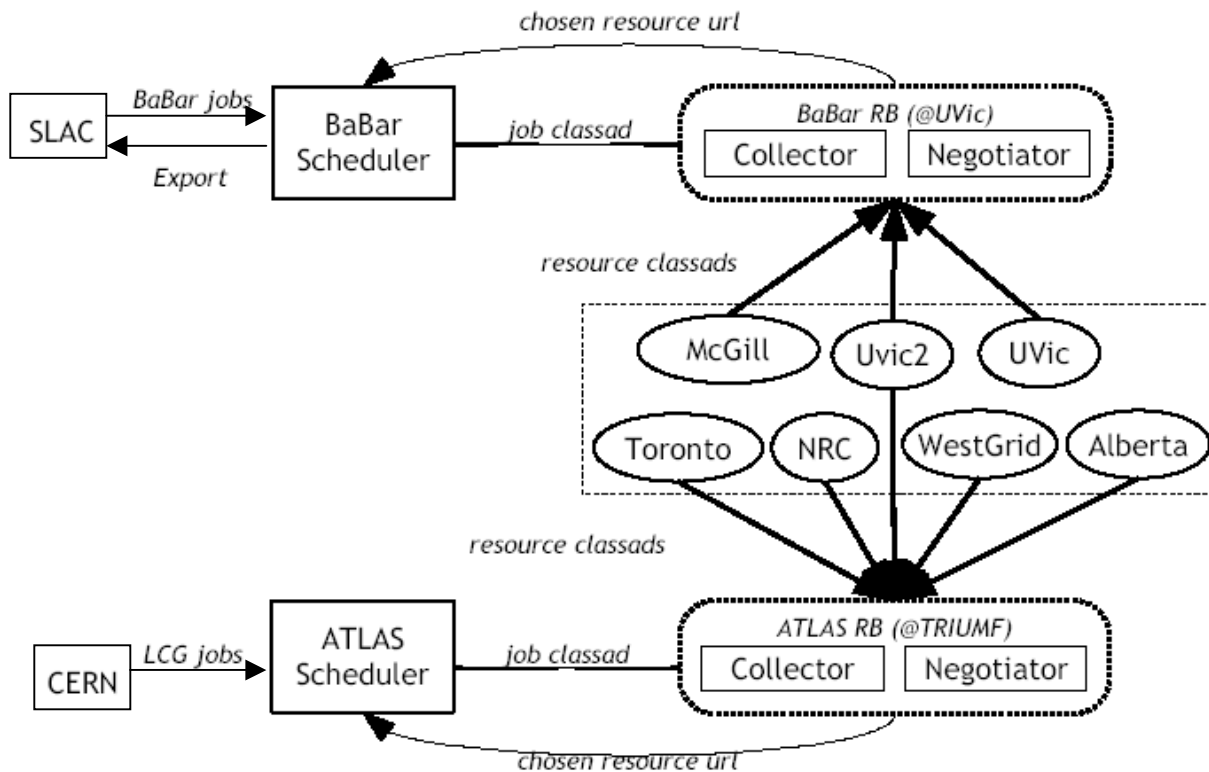
Large failure rate is due to unavailability of SE



# BaBar MC Production Setup Using Resource Broker



# Metascheduling Architecture



Condor-G system for BaBar

Condor-G system for ATLAS

# Condor-G Monitoring

job id	user id	owner	command	resource	status	run time	time submitted
8763.0	mcgrid	Ashok Agarwal	run.csh 6845 ...	caliope.phys.uvic.ca	ACTIVE	2:22:22	2 Nov, 12:15
8764.0	mcgrid	Ashok Agarwal	run.csh 6845 ...	caliope.phys.uvic.ca	ACTIVE	4:08:48	2 Nov, 12:15
8770.0	mcgrid	Ashok Agarwal	run.csh 6845 ...	caliope.phys.uvic.ca	PENDING		2 Nov, 12:15
8771.0	mcgrid	Ashok Agarwal	run.csh 6845 ...	caliope.phys.uvic.ca	ACTIVE	0:25:14	2 Nov, 12:15
8772.0	mcgrid	Ashok Agarwal	run.csh 6845 ...	caliope.phys.uvic.ca	ACTIVE	2:08:22	2 Nov, 12:16
8773.0	mcgrid	Ashok Agarwal	run.csh 6845 ...	caliope.phys.uvic.ca	ACTIVE	4:23:23	2 Nov, 12:16
8780.0	mcgrid	Ashok Agarwal	run.csh 6845 ...	caliope.phys.uvic.ca	PENDING		2 Nov, 12:16
8781.0	mcgrid	Ashok Agarwal	run.csh 6845 ...	caliope.phys.uvic.ca	ACTIVE	2:07:22	2 Nov, 12:16
8782.0	mcgrid	Ashok Agarwal	run.csh 6845 ...	caliope.phys.uvic.ca	ACTIVE	3:57:52	2 Nov, 12:16
8790.0	mcgrid	Ashok Agarwal	run.csh 6845 ...	caliope.phys.uvic.ca	ACTIVE	2:04:47	2 Nov, 12:16
8791.0	mcgrid	Ashok Agarwal	run.csh 6845 ...	caliope.phys.uvic.ca	ACTIVE	3:50:48	2 Nov, 12:16
8800.0	mcgrid	Ashok Agarwal	run.csh 6846 ...	caliope.phys.uvic.ca	ACTIVE	0:30:20	2 Nov, 12:16
8801.0	mcgrid	Ashok Agarwal	run.csh 6846 ...	caliope.phys.uvic.ca	PENDING		2 Nov, 12:16
8802.0	mcgrid	Ashok Agarwal	run.csh 6846 ...	caliope.phys.uvic.ca	PENDING		2 Nov, 12:16
8803.0	mcgrid	Ashok Agarwal	run.csh 6846 ...	caliope.phys.uvic.ca	PENDING		2 Nov, 12:16
8804.0	mcgrid	Ashok Agarwal	run.csh 6846 ...	caliope.phys.uvic.ca	PENDING		2 Nov, 12:16
8805.0	mcgrid	Ashok Agarwal	run.csh 6846 ...	caliope.phys.uvic.ca	PENDING		2 Nov, 12:16
8806.0	mcgrid	Ashok Agarwal	run.csh 6846 ...	caliope.phys.uvic.ca	PENDING		2 Nov, 12:16
8807.0	mcgrid	Ashok Agarwal	run.csh 6846 ...	caliope.phys.uvic.ca	ACTIVE	0:15:19	2 Nov, 12:16
8808.0	mcgrid	Ashok Agarwal	run.csh 6846 ...	caliope.phys.uvic.ca	ACTIVE	1:20:51	2 Nov, 12:16
8810.0	mcgrid	Ashok Agarwal	run.csh 6846 ...	caliope.phys.uvic.ca	PENDING		2 Nov, 12:16
8811.0	mcgrid	Ashok Agarwal	run.csh 6846 ...	caliope.phys.uvic.ca	PENDING		2 Nov, 12:16
8812.0	mcgrid	Ashok Agarwal	run.csh 6846 ...	caliope.phys.uvic.ca	PENDING		2 Nov, 12:16
8813.0	mcgrid	Ashok Agarwal	run.csh 6846 ...	caliope.phys.uvic.ca	PENDING		2 Nov, 12:16

job id	user id	owner	command	resource	status	run time	time submitted
8763.0	mcgrid	Ashok Agarwal	run.csh 6845 ...	caliope.phys.uvic.ca	ACTIVE	2:22:22	2 Nov, 12:15
8764.0	mcgrid	Ashok Agarwal	run.csh 6845 ...	caliope.phys.uvic.ca	ACTIVE	4:08:48	2 Nov, 12:15
8770.0	mcgrid	Ashok Agarwal	run.csh 6845 ...	caliope.phys.uvic.ca	PENDING		2 Nov, 12:15
8771.0	mcgrid	Ashok Agarwal	run.csh 6845 ...	caliope.phys.uvic.ca	ACTIVE	0:25:14	2 Nov, 12:15
8772.0	mcgrid	Ashok Agarwal	run.csh 6845 ...	caliope.phys.uvic.ca	ACTIVE	2:08:22	2 Nov, 12:16
8773.0	mcgrid	Ashok Agarwal	run.csh 6845 ...	caliope.phys.uvic.ca	ACTIVE	4:23:23	2 Nov, 12:16
8780.0	mcgrid	Ashok Agarwal	run.csh 6845 ...	caliope.phys.uvic.ca	PENDING		2 Nov, 12:16
8781.0	mcgrid	Ashok Agarwal	run.csh 6845 ...	caliope.phys.uvic.ca	ACTIVE	2:07:22	2 Nov, 12:16
8782.0	mcgrid	Ashok Agarwal	run.csh 6845 ...	caliope.phys.uvic.ca	ACTIVE	3:57:52	2 Nov, 12:16
8790.0	mcgrid	Ashok Agarwal	run.csh 6845 ...	caliope.phys.uvic.ca	ACTIVE	2:04:47	2 Nov, 12:16
8791.0	mcgrid	Ashok Agarwal	run.csh 6845 ...	caliope.phys.uvic.ca	ACTIVE	3:50:48	2 Nov, 12:16
8800.0	mcgrid	Ashok Agarwal	run.csh 6846 ...	caliope.phys.uvic.ca	ACTIVE	0:30:20	2 Nov, 12:16
8801.0	mcgrid	Ashok Agarwal	run.csh 6846 ...	caliope.phys.uvic.ca	PENDING		2 Nov, 12:16
8802.0	mcgrid	Ashok Agarwal	run.csh 6846 ...	caliope.phys.uvic.ca	PENDING		2 Nov, 12:16
8803.0	mcgrid	Ashok Agarwal	run.csh 6846 ...	caliope.phys.uvic.ca	PENDING		2 Nov, 12:16
8804.0	mcgrid	Ashok Agarwal	run.csh 6846 ...	caliope.phys.uvic.ca	PENDING		2 Nov, 12:16
8805.0	mcgrid	Ashok Agarwal	run.csh 6846 ...	caliope.phys.uvic.ca	PENDING		2 Nov, 12:16
8806.0	mcgrid	Ashok Agarwal	run.csh 6846 ...	caliope.phys.uvic.ca	PENDING		2 Nov, 12:16
8807.0	mcgrid	Ashok Agarwal	run.csh 6846 ...	caliope.phys.uvic.ca	ACTIVE	0:15:19	2 Nov, 12:16
8808.0	mcgrid	Ashok Agarwal	run.csh 6846 ...	caliope.phys.uvic.ca	ACTIVE	1:20:51	2 Nov, 12:16
8810.0	mcgrid	Ashok Agarwal	run.csh 6846 ...	caliope.phys.uvic.ca	PENDING		2 Nov, 12:16
8811.0	mcgrid	Ashok Agarwal	run.csh 6846 ...	caliope.phys.uvic.ca	PENDING		2 Nov, 12:16
8812.0	mcgrid	Ashok Agarwal	run.csh 6846 ...	caliope.phys.uvic.ca	PENDING		2 Nov, 12:16
8813.0	mcgrid	Ashok Agarwal	run.csh 6846 ...	caliope.phys.uvic.ca	PENDING		2 Nov, 12:16

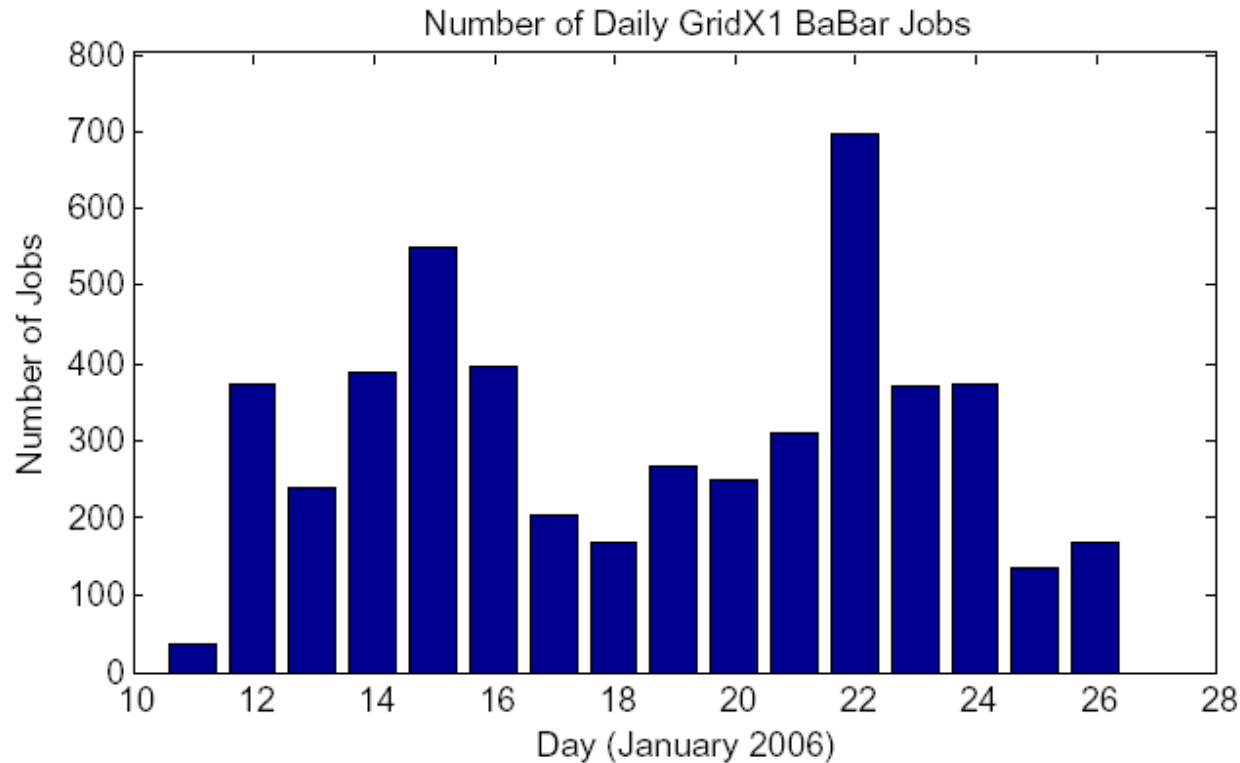
### babargrid.phys.uvic.ca resource

Site	Active + Pending	Waiting Time
caliope.phys.uvic.ca:2119/jobmanager-pbs-gc-production	16+260	00:00:00
sl-gw.physics.mcgill.ca:2119/jobmanager-pbs-workq	13+66	00:00:00
mercury2.uvic.ca:2119/jobmanager-pbs-gc-production	6+0	00:00:00
Unsubmitted:	0	
Total:	35+326	

As of: 2005-11-04 15:21:25

GridX1 monitoring website:  
[www.gridx1.ca](http://www.gridx1.ca)

# BaBar Jobs



Plot of BaBar jobs successfully executed on GridX1 over the period of January 11-26, 2006.

# Challenges

- GridX1 is equivalent to a moderate-sized computing facility
  - It requires a “grid” system administrator to keep system operational
- We need a more automated way to install applications
- Monitoring is in good shape but further improvements are needed
  - Improve reliability and scalability
- Error recovery has not been an issue with LCG and BaBar jobs

# Summary

- GridX1 setup for both ATLAS and BaBar applications proved to be highly efficient and successful
- Over 20,000 ATLAS jobs for July 2004 – June 2005
- Typically 200 or more BaBar jobs run on the Uvic and McGill clusters
- We want to add more sites to enhance the activity of HEP applications
- Other applications could be run on the Grid
- Please visit our poster on using Xen in a Grid environment titled “[Evaluation of Virtual Machines for HEP Grids](#)”