



Open Science Grid

Open Science Grid

High Throughput Computing On A National Scale


Alain Roy





OSG: HTC at National Scale

- OSG provides high-throughput computing across the United States.
 - 70 or so sites
 - For 28-Nov-2008:
 - 131,261 jobs for 393,312 hours
 - Used 54 sites
 - Jobs by 30 different virtual organizations
 - 86% of jobs succeeded
 - Underestimate: 64% of sites reported statistics



Who Uses OSG?

- About 30 virtual organizations
 - High-energy physics uses a large chunk of OSG
 - But several other sciences are actively using OSG.
 - nanoHUB: nanotechnology simulations
 - LIGO: detecting gravitational waves
 - CHARMM: molecular dynamics
 - Football pool: mathematical coding theory


More at:

http://www.opensciencegrid.org/Science_on_the_OSG/



Focus on DZero

- High-energy physics experiment
- Based at Fermilab, near Chicago, US
- Searching for new particles by smashing together protons and antiprotons at nearly the speed of light.
- I'm not a physicist: this is as deep as my understanding goes






DZero & HTC

- DZero exemplifies HTC
- They do local HTC
- They do HTC on OSG




Colliders are Big





Physics Collaborations Are Big



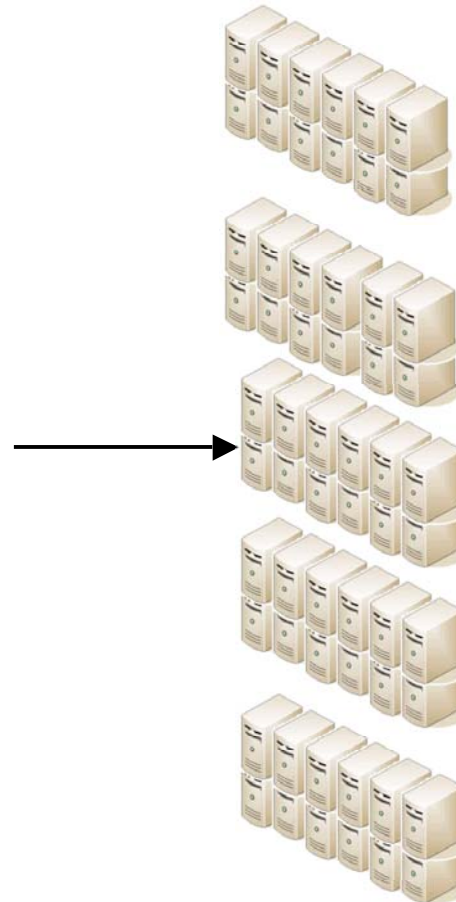


Dzero's Computing is Big

Detector
(three stories tall)



Process/Analyze
(~1000 CPUs)





Dzero's Problem

- As data comes off the detector, it is:
 - Processed once on everyone's behalf
 - Analyzed many times by many scientists
- Recently, they wanted to re-process all the data from the detector in time for scientists to analyze for summer conferences
- They needed ~ 500, 1GHz computers for one year
- They only had 1000 CPUs for a few months
 - And they were also doing new processing, not just reprocessing



DZero's Core Problem

- DZero needed a peak capacity that was beyond their local capacity



DZero's Solution

- Expand HTC onto OSG, and other grids
 - Requested 1500 CPUs from OSG for four months



How Did it Go? (1/2)

- Used about 12 OSG sites
 - Number fluctuated over time
 - Ramped up: Added sites one at a time
 - Certified that each site produced correct answers
 - Only three of these were “DZero” sites
 - Kept roughly 1500 CPUs busy, after ramp up
- Reprocessed 445 million events
 - 286 million on OSG



How Did it Go? (2/2)

- 90 TB of input data
- 250 TB of application
 - The application is 1GB
 - It was transferred many times
 - Easier than pre-installing on all nodes



Open Science Grid

Beyond DZero

OSG's goal is to provide for many scientists what we provided for for DZero.



The OSG Vision

Transform

processing and data intensive science

through

a cross-domain,
self-managed,
national,
distributed cyber-infrastructure

that brings together

campus and
community infrastructure

and facilitates

the needs of Virtual Organizations (VO)

at all scales



The OSG Vision

Transform

processing and data intensive science

through

a cross-domain,
self-managed,
national,
distributed cyber-infrastructure

Implies:

Autonomy

Heterogeneity

Large-Scale

organizations (VO)



Autonomy & Heterogeneity

- Autonomy & heterogeneity are a pain
- But they are also a fact of life
- If we accept it, we have access to more resources

- This is opportunistic computing



The OSG Vision

Transform

processing and data intensive science

through

a cross-domain,
self-managed,
national,
distributed cyber-infrastructure

that brings together

campus and
community infrastructure

and facilitates

the needs of Virtual Organizations (VO)

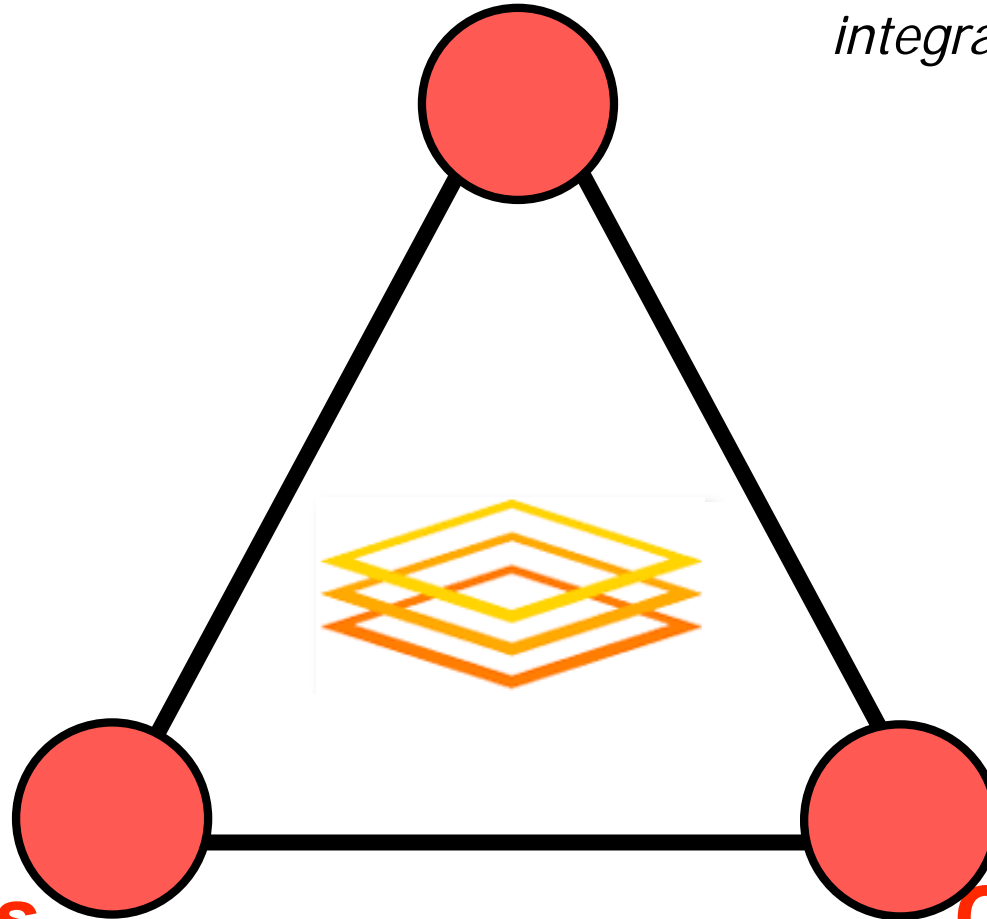
at all scales



The Three Cornerstones

National

Needs to be harmonized into a well integrated whole.



Campus

Community



OSG Needs

- OSG needs many things to be successful:
 - Good people
 - Good software ← **My Focus**
 - Good security
 - Good policies
 - Good communication
 - Good testing
 - ...



VDT: OSG Software Stack

- Virtual Data Toolkit (VDT)
 - A software distribution for Grid computing
 - A packaging of other software
(Like a Linux distribution, but different)
 - No software development
 - We get Condor, Globus, and other software from other groups
 - We “glue” it together



Why Have The VDT?


- Everyone could download the software from the providers...
- But the VDT:
 - Figures out dependencies between software
 - Works with providers for bug fixes
 - Provides automatic configuration
 - Builds it (we provide binaries)
 - Packages it
 - Tests everything on a dozen or so platforms (and growing)



Example: VOMS

- VOMS can authorize people in a VO
- VOMS has a web interface
- We:
 - Install Tomcat
 - Install Apache
 - Built with Globus SSL
 - Patched so GSI pass through to Apache works
 - Install VOMS
 - Install VOMS Admin
 - Install Perl modules needed by VOMS Admin
 - Install MySQL and set up database (with command-line tool)
 - Configure all software
 - Configure rotation of log files

We pre-build binaries for each of these



Example: Security Update

- Last year, a security update to Globus software:
 - We decided to patch four versions of the VDT
 - We built updated binaries three times on about six platforms
 - We coordinated creation of patch for unsupported version of Globus
 - We patched the Globus updates with our patches
 - We took subset of Globus updates
 - We packaged an update that was reversible, if there were problems.





What's in the VDT?

- Job management:
 - Globus GRAM
 - Condor
- Data management
 - Globus GridFTP
 - dCache
 - Bestman
- Security
 - VOMS
 - GUMS
 - PRIMA
 - MyProxy
- Information/Monitoring
 - CEMon
 - Generic Info Provider
 - Site Validation
- Infrastructure:
 - Apache
 - Tomcat
 - Python
 - Perl modules
 - ...
- Miscellaneous
 - Squid
 - Wget

And More !




Supported Platforms

- RedHat Enterprise Linux
 - 3, 4, 5
 - x86, x86-64, ia64
- Scientific Linux
 - 3, 4, 5,
 - x86, x86-64, ia64
- Fedora Core 4
- Debian 3 (soon)
- SLES 9
- Mac OS X

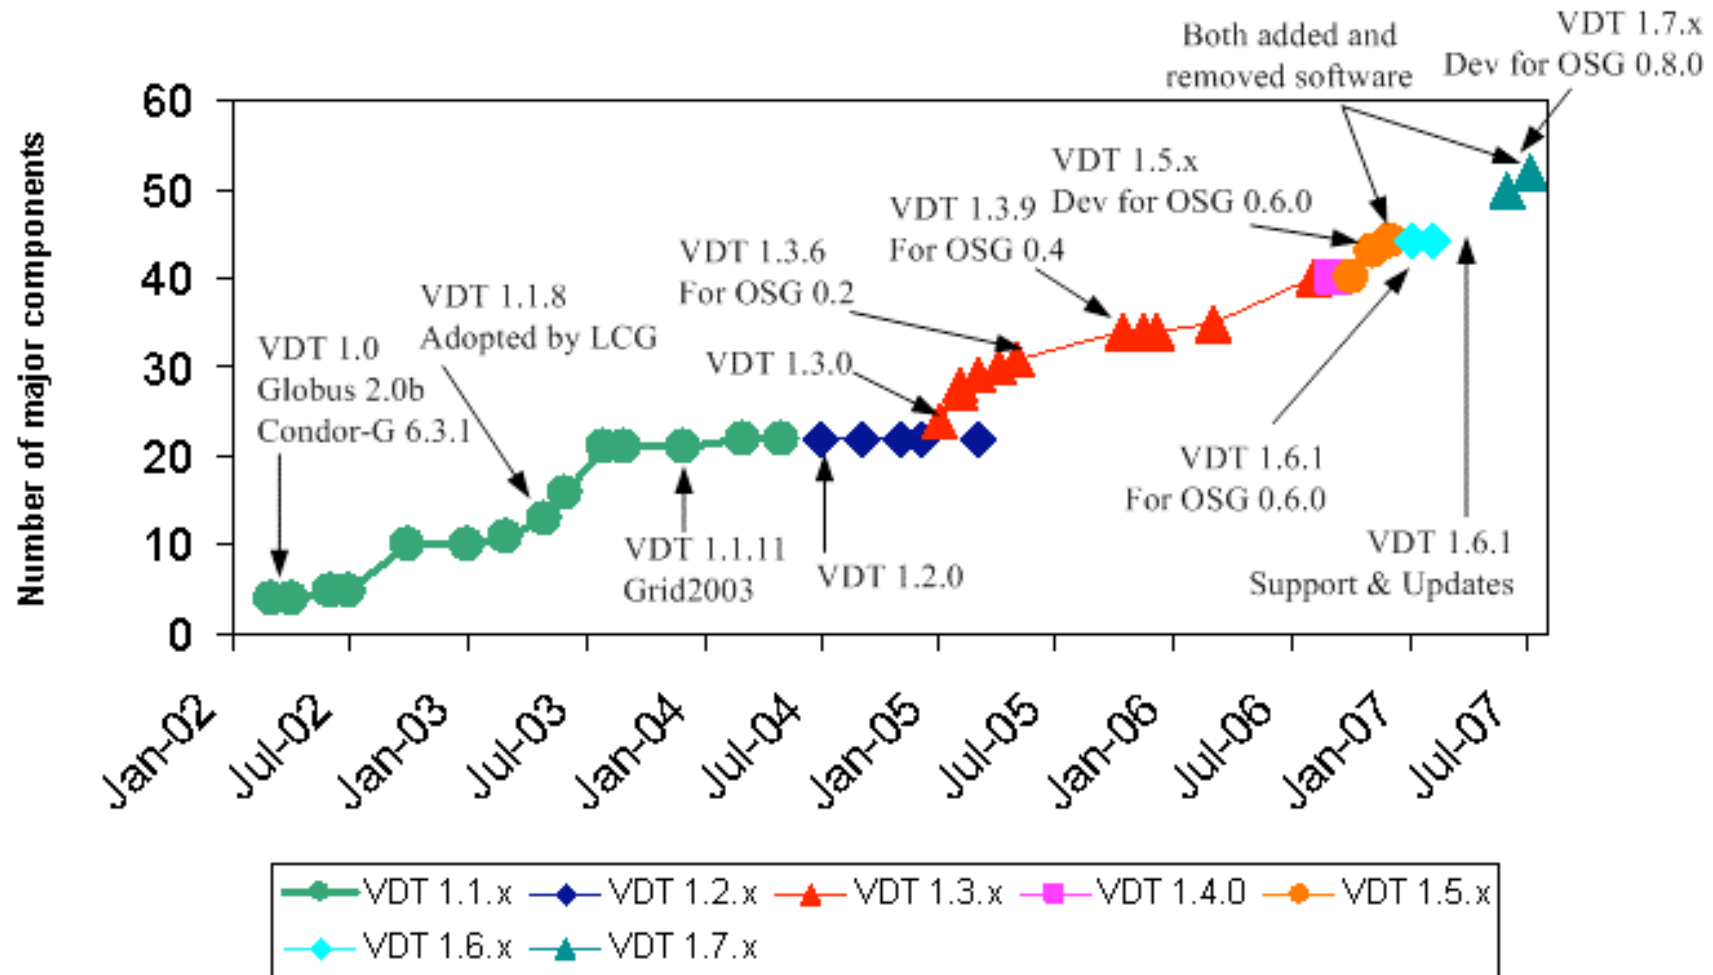


Keeping up with Linux distributions is like sprinting a marathon

But to support autonomous, heterogeneous sites, it's a necessity



VDT Growth





The VDT's Challenge

- Keep software up to date
- Add new software
- Support latest OS versions (and old ones!)
- Keep it secure
- Make it easy to update
- Make it easier to install
- Create better documentation



Questions?

Alain Roy

roy@cs.wisc.edu

vdt-support@opensciencegrid.org

