Grid Operations

Doug Pearson
Senior Manager Global NOC and Director REN-ISAC
Indiana University
dodpears@iu.edu

LCG Workshop on Operational Issues

November 2, 2004

Agenda

- Introduction to iVDGL, Grid3, and the iGOC
- Efforts, Accomplishments and Lessons Learned
- Future Directions

http://www.ivdgl.org/grid3

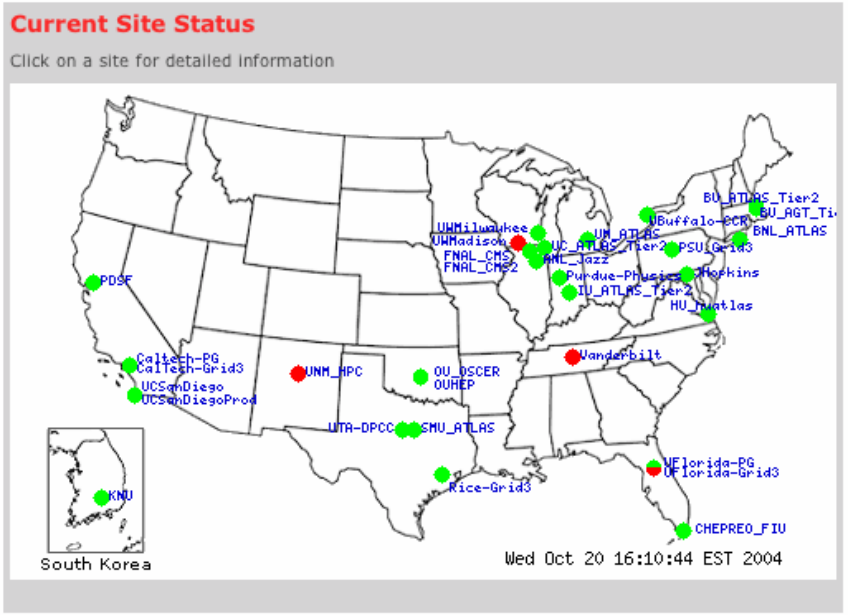

iVDGL Project Goals

- Deploy a Grid laboratory
  - Support research mission of data intensive experiments
  - Provide computing and personnel resources at university sites
  - Provide platform for computer science technology development
  - Prototype and deploy a Grid Operations Center (iGOC)
- Integrate Grid software tools
  - Into computing infrastructures of the experiments
- Support delivery of Grid technologies
  - Hardening of the Virtual Data Toolkit (VDT) and other middleware technologies developed by GriPhyN and other Grid projects
- Education and Outreach

iVDGL and Grid3 Structural Overview

- iVDGL
  - Steering committee
  - Core software
  - Operations
  - Applications
  - Outreach
- Grid3
  - Taskforce
  - iGOC (operations)
  - Admins
  - Grid3dev (development environment)

Grid3

- The first multi-VO grid environment shared across U.S. laboratories and universities to run applications from a variety of scientific disciplines.
- A collaboration of U.S. Physics Grid projects GriPhyN, iVDGL and PPDG, and U.S. ATLAS, U.S. CMS, plus participation from the Condor and Globus teams.
- Initial goal was to demonstrate specific metrics (demonstrated at the SC2003 conference), and then
- To continue to operate to the benefit of the collaborating organizations; continue to contribute to the U.S.CMS and U.S. ATLAS community data challenges, and others.
- Over 30 sites with over 3000 processors.

Grid3 Architecture

- Simple approach:
  - Sites consisting of
    - Computing element (CE)
    - Storage element (SE)
    - Information and monitoring services
  - VO level, and multi-VO
    - VO information services
    - Operations (iGOC)

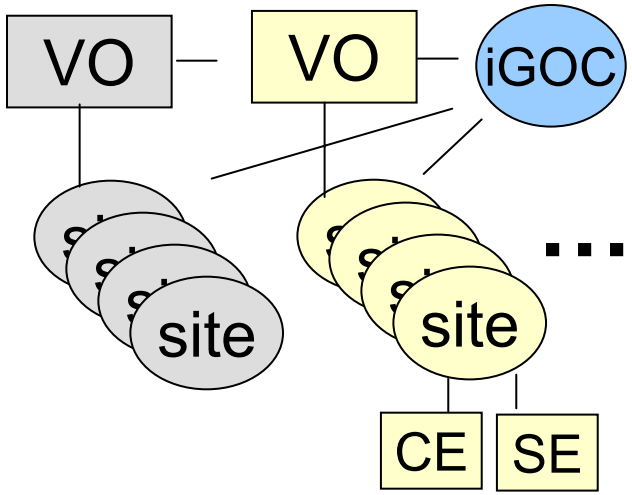

- Minimal use of grid-wide systems
  - No centralized workload manager, production replica or data management catalogs, or command line interface
    - higher level services are provided by individual VO's
      - E.g. VDS (Virtual Data Service), controls flow of data to and from site, Atlas Capone executes Atlas code on grid sites; CMS has MOP similar. I.E. creating own resource brokers for job submission and data movement.

Grid3 Architecture

- A few distinguishing characteristics:
  - No interactive logins to remote sites, i.e. use globus-job-submit or globus-job-run; fork job manager enabled
  - No centralized scheduler – users determine where to submit their work (user checks Grid3catalog, ACDC job monitor, and Ganglia real-time load monitoring)
    - Some VO's, e.g Atlas use a job submitting person/team

Grid3 Services

- Authorization Services
  - Virtual Organization Management System (VOMS) consists of VO owned and operated user Distinguished Name (DN) repositories, exposed to the grid via dedicated LDAP servers.
- Software Packaging Service
  - VDT + Grid3 specifics via Pacman
- Monitoring Services
  - GridCat, ACDC Job Monitor, MonALISA, Ganglia, and MetricsData Viewer:
- Grid3 Operations
  - The iGOC .

Grid3: an application grid laboratory

Continuous Operations since deployment in September 2003

Successfully upgraded VDT during June 2004

CERN LHC: US testbeds

C: US ATLAS challenges

app

earch

ual data grid laboratory

iVDGL iGOC

- Mission
  - Deploy, maintain, and operate Grid3 as a NOC manages a inter-network, providing a single point of operations for configuration support, monitoring of status and usage (current and historical), problem management, support for users, developers and systems administrators, provision of grid services, security incident response, and maintenance of the Grid3 information repository.
- Staffing:
  - 2 FTE at Indiana University, plus effort from University of Chicago (monitoring development), University Florida at Gainsville (Grid3catalog, web site, site verify script, etc.), and leveraged resources of the 24x7 NOC at Indiana University

iVDGL iGOC

- Proposed Areas of Research:
  - Access control and policy - Security
  - Trouble Ticket System - Problem coordination
  - Configuration and Information Services
  - Health and Status Monitoring
  - Experiment Scheduling

iVDGL/Grid3 Operations Approach

- The iVDGL Grid3 Operations group
  - Sets up and maintains a cooperative grid community
  - Facilitates work to and among responsible agents
  - Has no direct control: uses notification with follow-ups
  - Tunes services to the capabilities of the sites
- Cooperative and mentoring principles are employed:
  - Identifies community vision – i.e. the Project Plan (anchor)
  - Utilizes a participatory decision making process -- Taskforce
  - Makes clear agreements -- Service Descriptions and MOUs
  - Makes clear communication and conflict resolution a priority
    - Weekly operations (problem solving) and management teleconferences.

Agenda

- Introduction to iVDGL, Grid3, and the iGOC
- Efforts, Accomplishments and Lessons Learned
- Future Directions

http://www.ivdgl.org/grid3

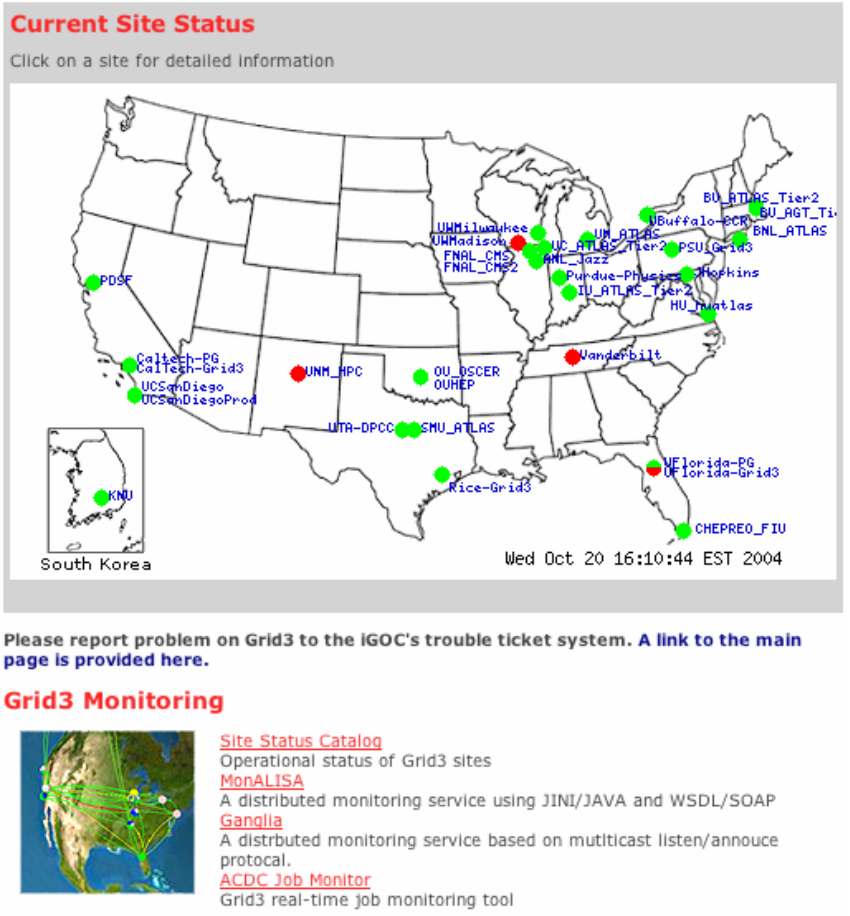

iGOC – Division of Work

- Service Desk
- Engineering
- Experts collective
- Web development
- Management

iGOC – Service Desk

- Activities
  - A common face to collaboratively-provided support
  - Facilitate and support communications:
    - Direct email with site administrators and Grid users
    - Web page resources
    - Status reporting to mailing list
  - Monitor status of Grid resources
  - Coordinate and track:
    - Problems
    - Changes (software updates, resource additions)
    - Security incidents
    - Requests for assistance

iGOC – Service Desk

- Activities (continued)
  - Provide reports
    - Problem summaries, service desk activity
  - Maintain the repository of support and process information
  - User support, such as:
    - How to join a VO
    - How to get and maintain a cert
    - How to run an application
    - How to use monitoring tools
    - Troubleshooting application failures
    - Information about policies, etc.

iGOC – Engineering

- Activities
  - Maintain the grid-controlled software packages and cache
  - Provide site software not supported through VDT
  - Verify software compatibility
  - Provide ease-of-installation tools
  - Develop instructions on how to plug things together
  - Provide site installation and configuration support
  - End-to-end troubleshooting for resources
  - Provide and maintain common Grid services such as VOMS, GIIS, RLS, archives, and monitoring systems

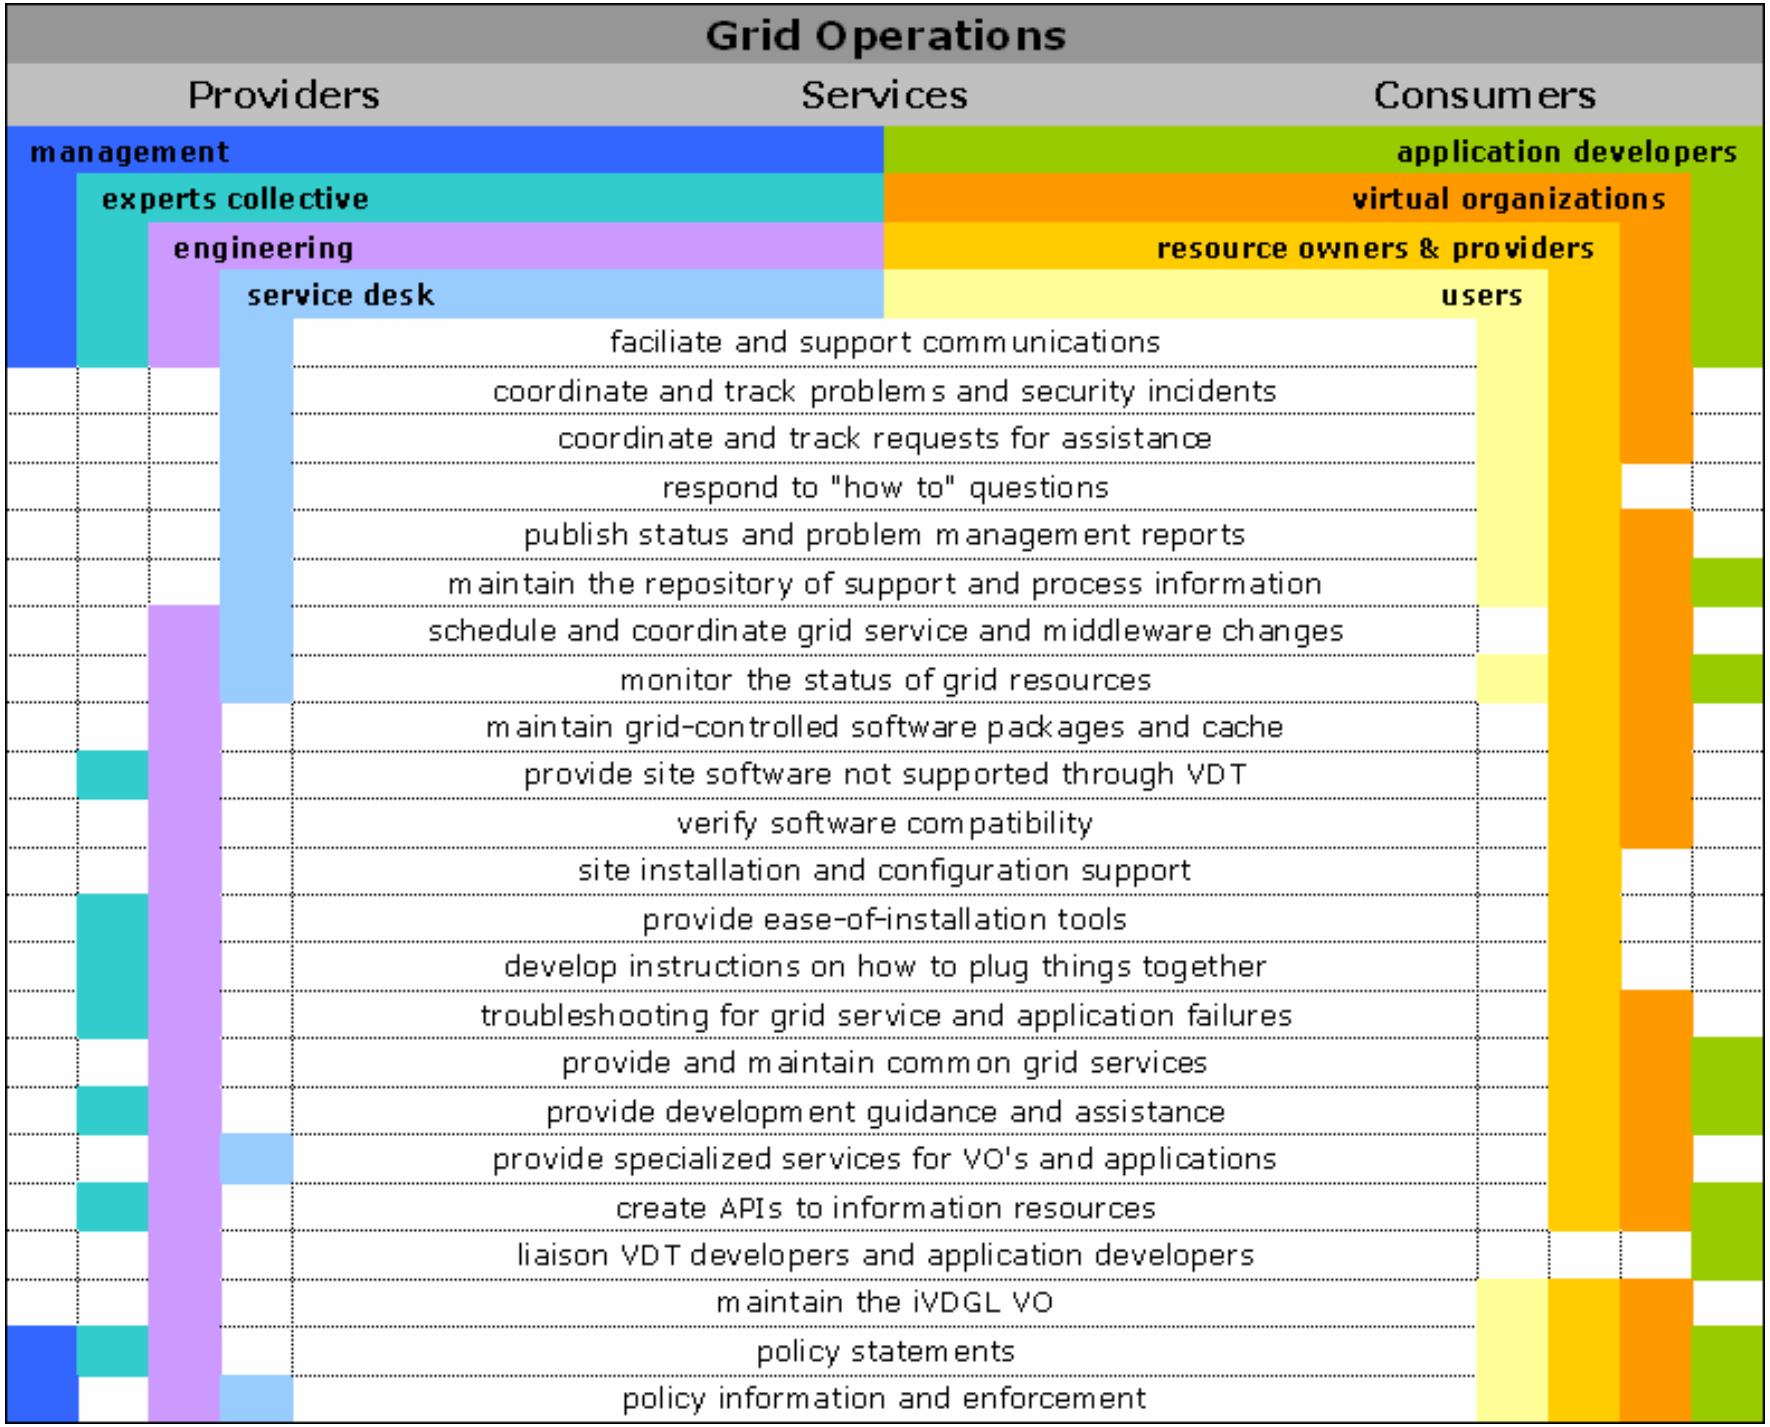
Grid Operations
Providers
Services
Consumers
management
application developers
experts collective
virtual organizations
engineering
resource owners & providers
service desk
users
faciliate and support communications
coordinate and track problems and security incidents
coordinate and track requests for assistance
respond to "how to" questions
publish status and problem management reports
maintain the repository of support and process information
schedule and coordinate grid service and middleware changes
monitor the status of grid resources
maintain grid-controlled software packages and cache
provide site software not supported through VDT
verify software compatibility
site installation and configuration support
provide ease-of-installation tools
develop instructions on how to plug things together
troubleshooting for grid service and application failures
provide and maintain common grid services
provide development guidance and assistance
provide specialized services for VO's and applications
create APIs to information resources
liaison VDT developers and application developers
maintain the iVDGL VO
policy statements
policy information and enforcement

Operations Enables Applications

- Provide operational services that provide Applications with the “instruments” to:
  - Publish site policies and environment
  - Know the status of grid middleware on sites
  - Know the job queue for compute resources
  - Know the status and load of grid resources
  - Access monitoring archives
  - Manage VO services
  - Keep apprised of security incidents in the collaborative

Resource Monitoring

- Ganglia: Open source tool to collect cluster monitoring information such as CPU and network load, memory and disk usage
- Mona LISA: Monitoring and Archiving tool to support resource discovery, access to information and gateway to other information gathering systems
- ACDC Job Monitoring System: Application using grid submitted jobs to query the job managers and collect information about jobs. This information is stored in a DB and available for aggregated queries and browsing.
- Metrics Data Viewer (MDViewer): analyzes and plots information collected by the different monitoring tools, such as the DBs at iGOC.
- Globus MDS: Grid3 Schema for Information Services and Index Services for Information services
- GridCat: Graphical display of middleware testing results, provides Site database repository also include extended functions for storage, retrievable configuration and human contacts.

Leveraging the NOC

- Global NOC at Indiana University
  - The Global NOC provides 24x7 network engineering and operations services for research and education networks and international interconnections, including Internet2 Abilene, National LambdaRail, TransPAC and AMPATH networks, the STAR TAP and MANLAN layer 3 international exchange points, and the STAR LIGHT optical exchange. In addition, the Global NOC supports activities of the iVDGL Grid Operations Center and the REN-ISAC cybersecurity Watch Desk. By virtue of the R&E network, grid, and cybersecurity activities, the Global NOC possesses a unique and embracing view of R&E cyberinfrastructure.

Leveraging the NOC

- 24x7 front line
- Monitoring (watch for red indicators)
- Problem management
- Management overhead

Analysis of Effort by Area

Issues relating to resource owners and providers	60%
Special issues for Virtual Organizations (VO's)	20%
Issues relating to developers of applications and workflow environments (portals)	10%
Support to individuals using Grid resources	10%

iGOC ATLAS Data Challenge 2 Service Support

iGOC contact information: 24x7, igoc@ivdgl.org, 317-278-9699

BNL Operation Center: M-F 9AM-12AM EDT, 631-344-5480

Contents

Introduction

The Indiana University based Grid Operation Center (iGOC) provides operations services for participating sites of the US ATLAS Data Challenge 2 (DC2). Services include monitoring, problem notification, tracking, and reporting, covering hours when sites are not staffed, thereby providing DC2 with 24x7 support for critical production hardware and services.

Problem Reporting and Discovery

Problems are typically discovered by iGOC technicians via observation of monitoring systems. Additionally, problems may be reported to the iGOC from end-users or others via the trouble ticket submission webform, e-mail (igoc@ivdgl.org), web-based or phone (317-278-9699). The hours of iGOC service provided to each site is according to the Site Coverage Schedule, however, the iGOC will promptly handle all matters reported to it regardless of time of day.

- Provided 24x7 monitoring and problem discovery during Atlas DC2
- Successfully interoperated with BNL Tier1 Support Center
- Provided research advancements toward Grid to VO operations coordination

iGOC Daily Use Case

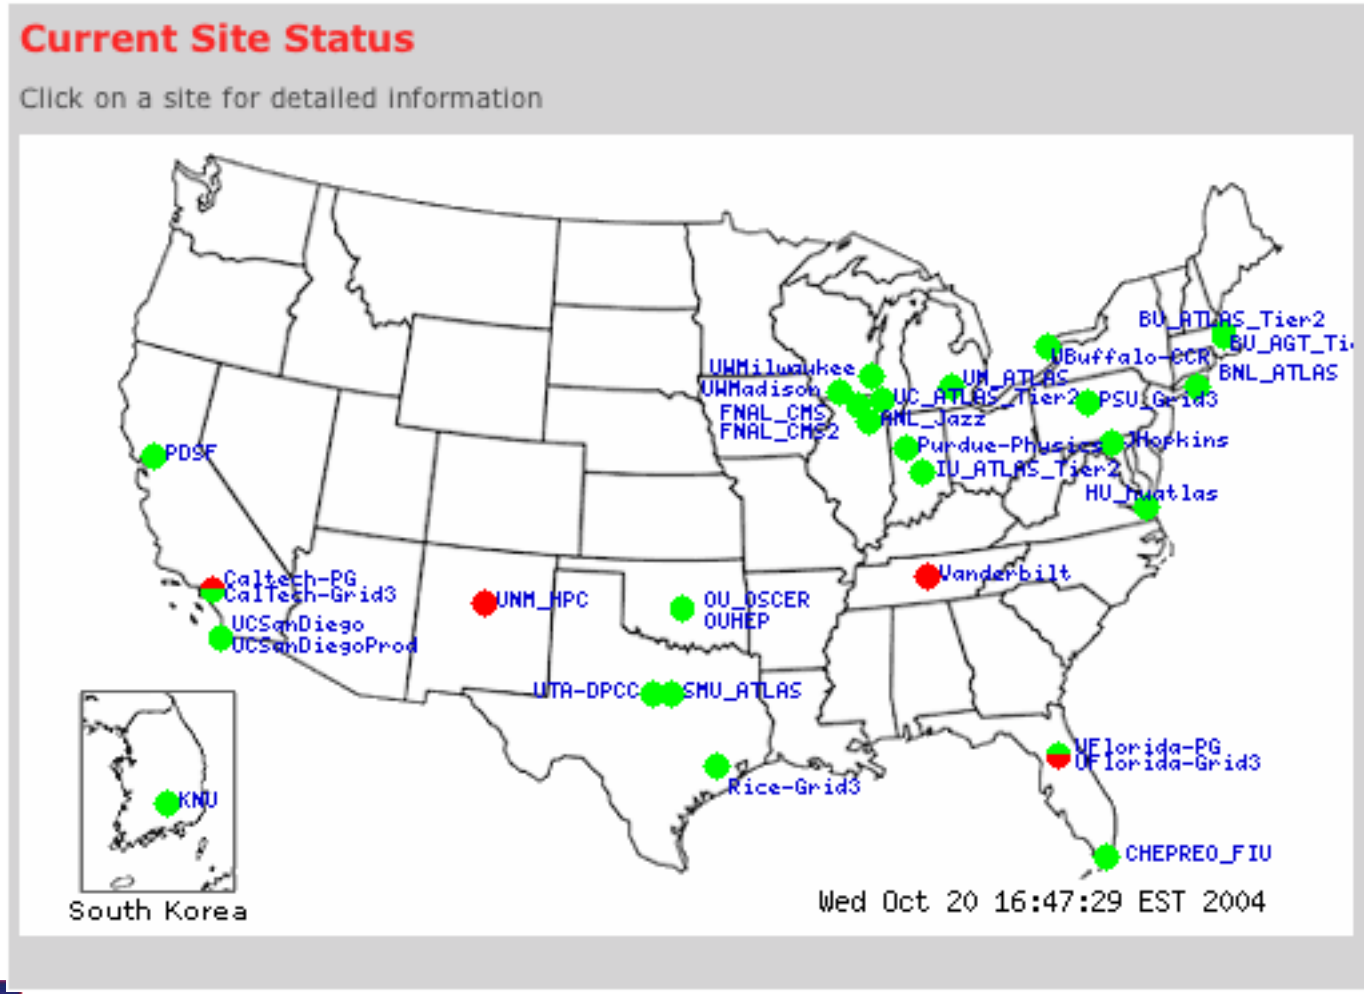

Gridcat Tests

- Tests are run every 5 hours
  - authentication (globusrun) (insures that site is in grid map file, equivalent of doing a ping)
  - helloworld, via globus-job-run (through the fork job manager).
  - GITS; submit a long job; see if the submit works; if yes then query for that job in the batch queuing system; then cancel job
  - gsiftp data transfer to and from
- Test results are world viewable

Following up on a "Red" Status

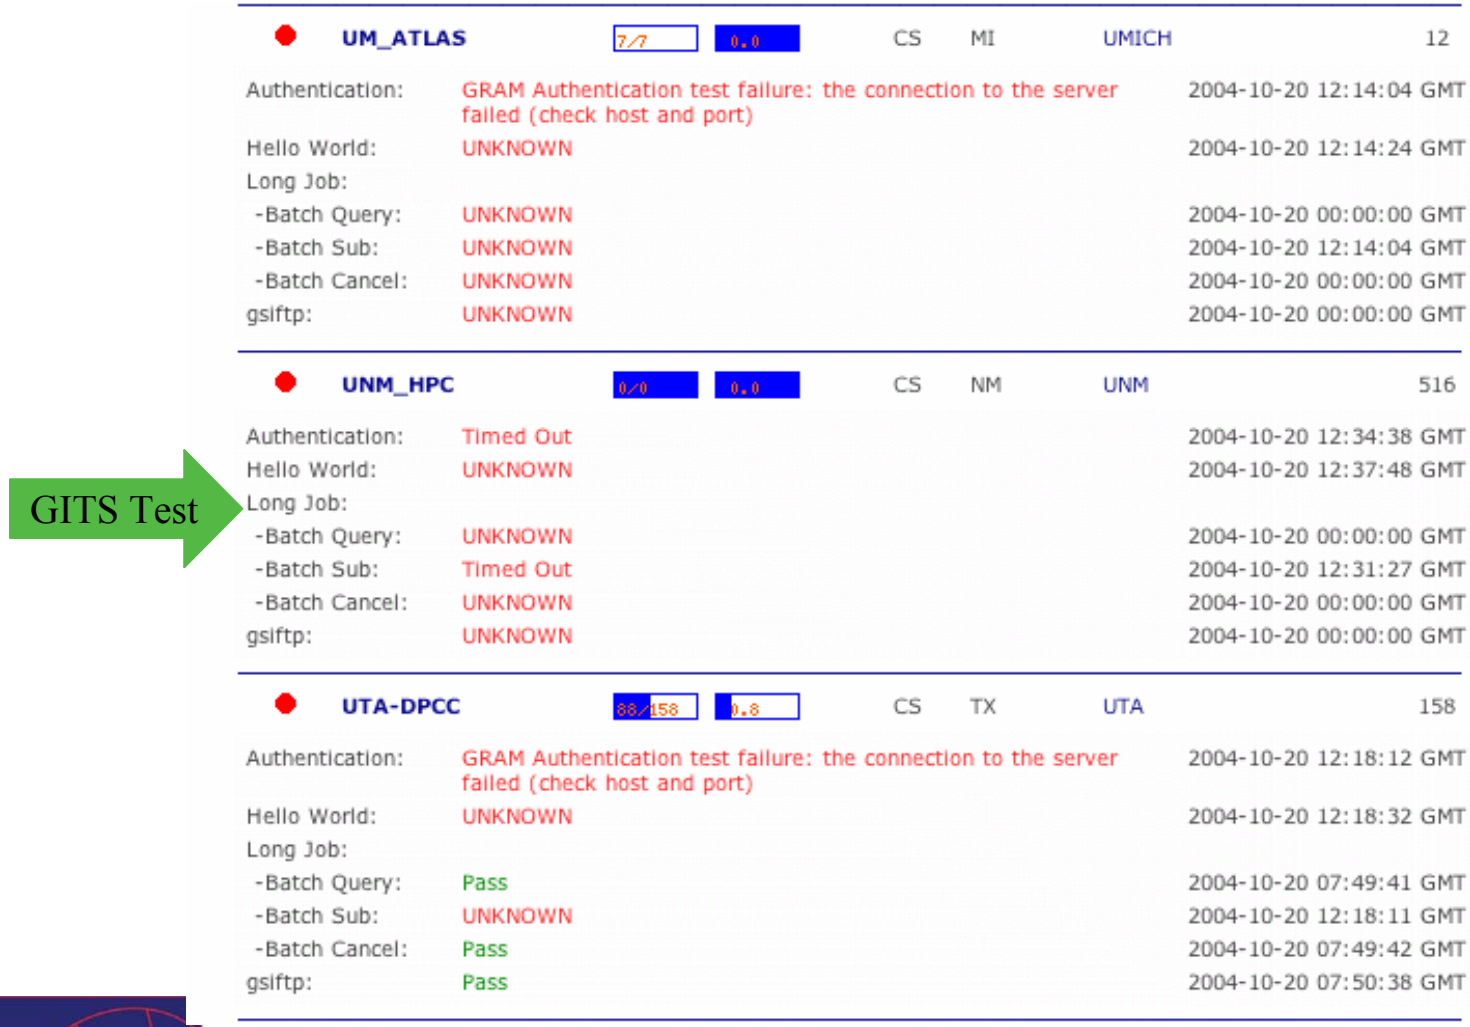

●	**UM_ATLAS**	7/7	0.0	CS	MI	UMICH	12
Authentication:	GRAM Authentication test failure: the connection to the server failed (check host and port)					2004-10-20 12:14:04 GMT
Hello World:	UNKNOWN					2004-10-20 12:14:24 GMT
Long Job:						
-Batch Query:	UNKNOWN					2004-10-20 00:00:00 GMT
-Batch Sub:	UNKNOWN					2004-10-20 12:14:04 GMT
-Batch Cancel:	UNKNOWN					2004-10-20 00:00:00 GMT
gsiftp:	UNKNOWN					2004-10-20 00:00:00 GMT
●	**UNM_HPC**	0/0	0.0	CS	NM	UNM	516
Authentication:	Timed Out					2004-10-20 12:34:38 GMT
Hello World:	UNKNOWN					2004-10-20 12:37:48 GMT
Long Job:						
-Batch Query:	UNKNOWN					2004-10-20 00:00:00 GMT
-Batch Sub:	Timed Out					2004-10-20 12:31:27 GMT
-Batch Cancel:	UNKNOWN					2004-10-20 00:00:00 GMT
gsiftp:	UNKNOWN					2004-10-20 00:00:00 GMT
●	**UTA-DPCC**	88/158	0.8	CS	TX	UTA	158
Authentication:	GRAM Authentication test failure: the connection to the server failed (check host and port)					2004-10-20 12:18:12 GMT
Hello World:	UNKNOWN					2004-10-20 12:18:32 GMT
Long Job:						
-Batch Query:	Pass					2004-10-20 07:49:41 GMT
-Batch Sub:	UNKNOWN					2004-10-20 12:18:11 GMT
-Batch Cancel:	Pass					2004-10-20 07:49:42 GMT
gsiftp:	Pass					2004-10-20 07:50:38 GMT

GITS Test

Test Time

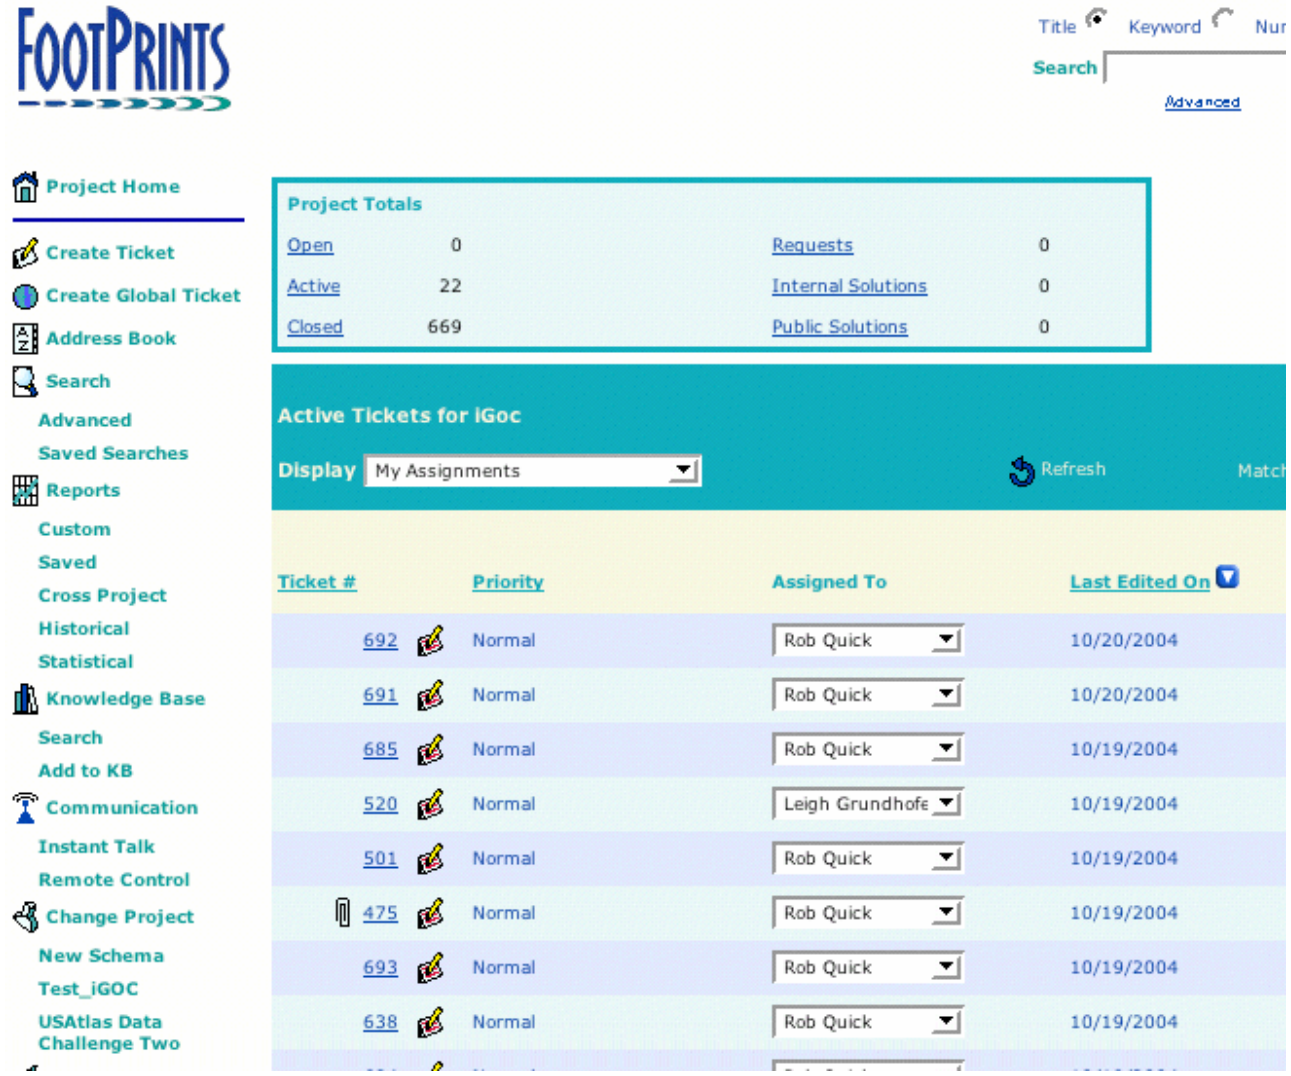

Nearly 700 tickets created since Jan 2004

22 open tickets

Ticket Creation since Nov. 2003

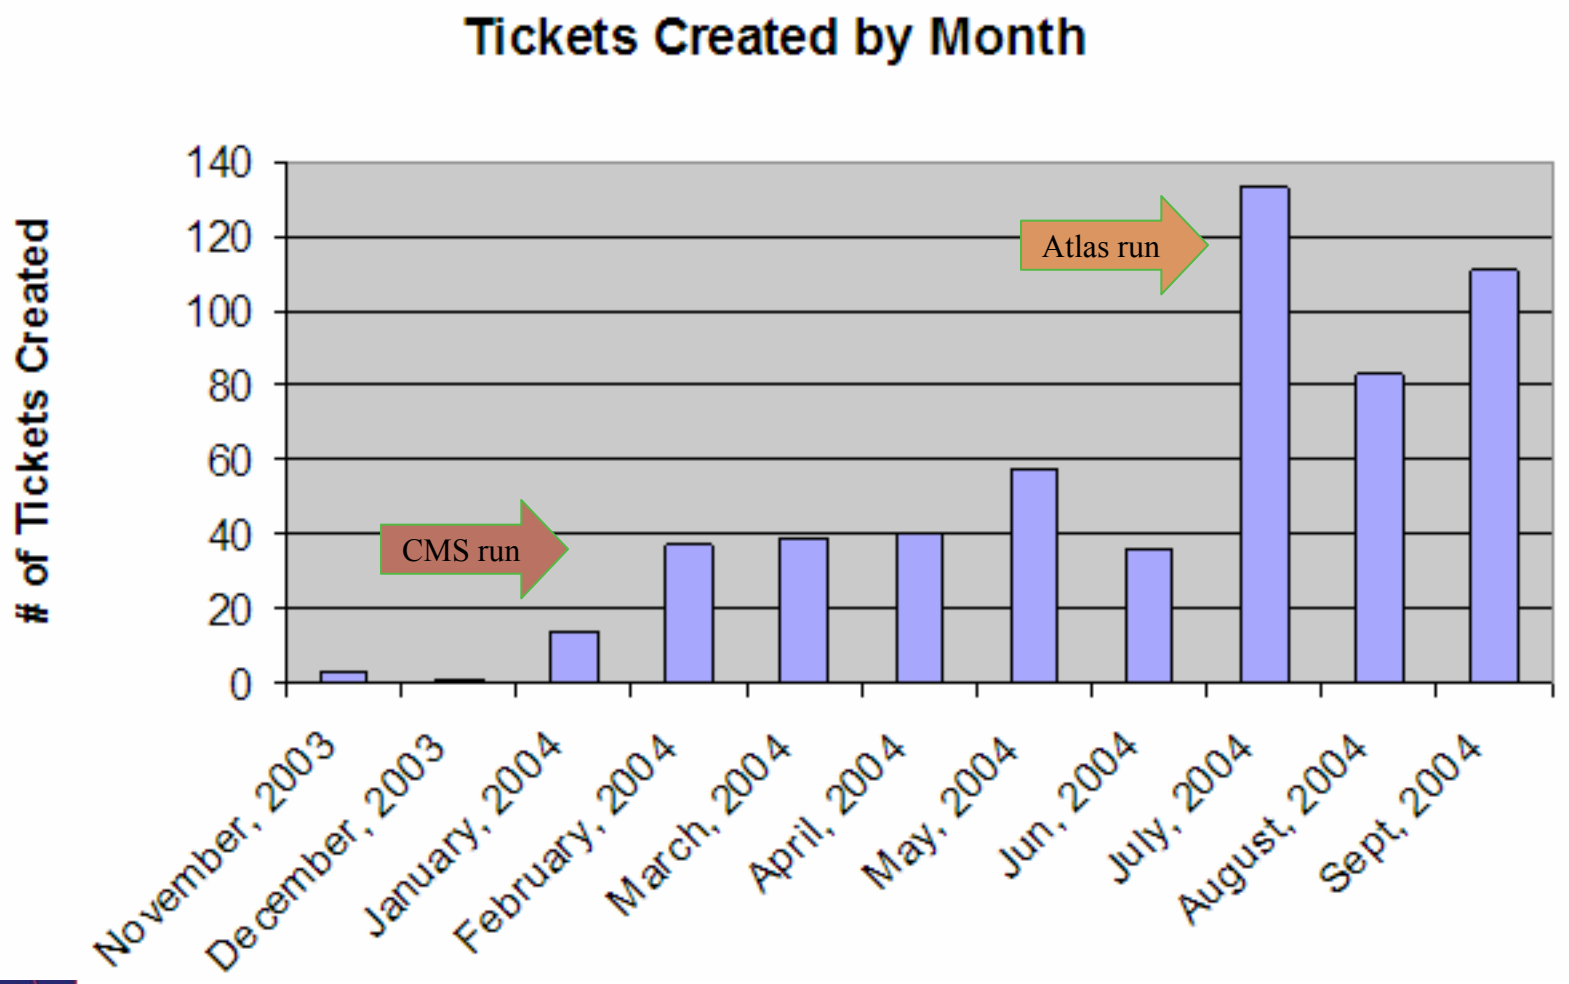

Grid3 TT Handling by Type

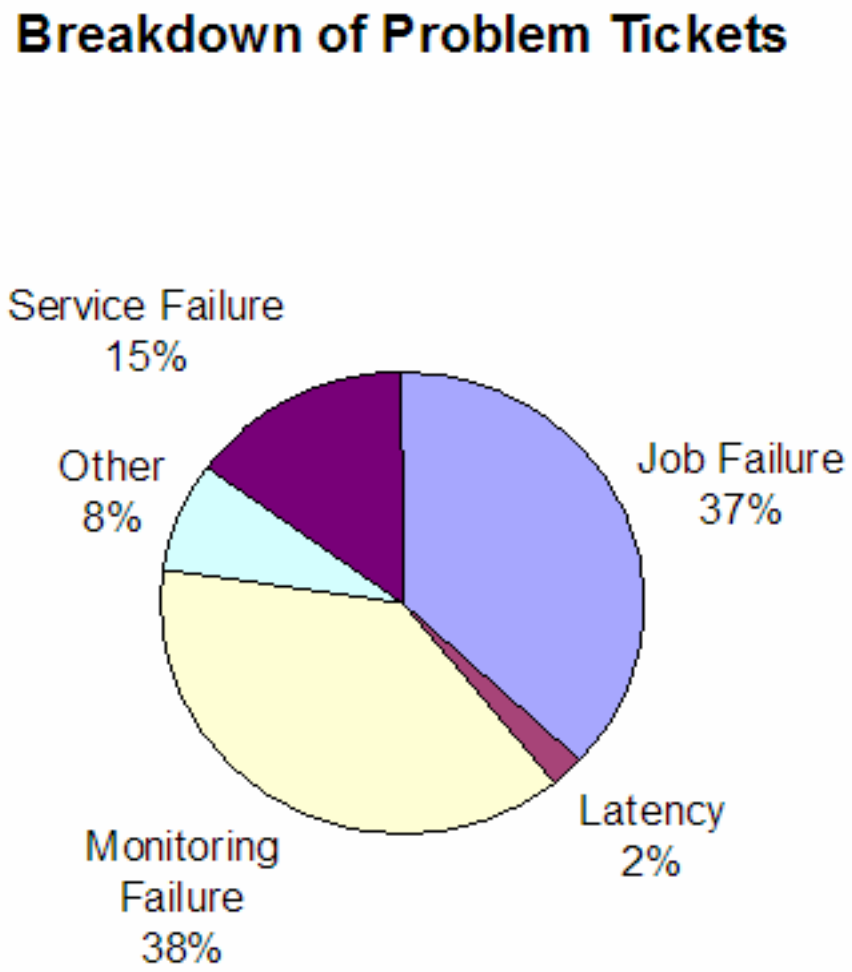

Atlas DC2 TT Handling by Type

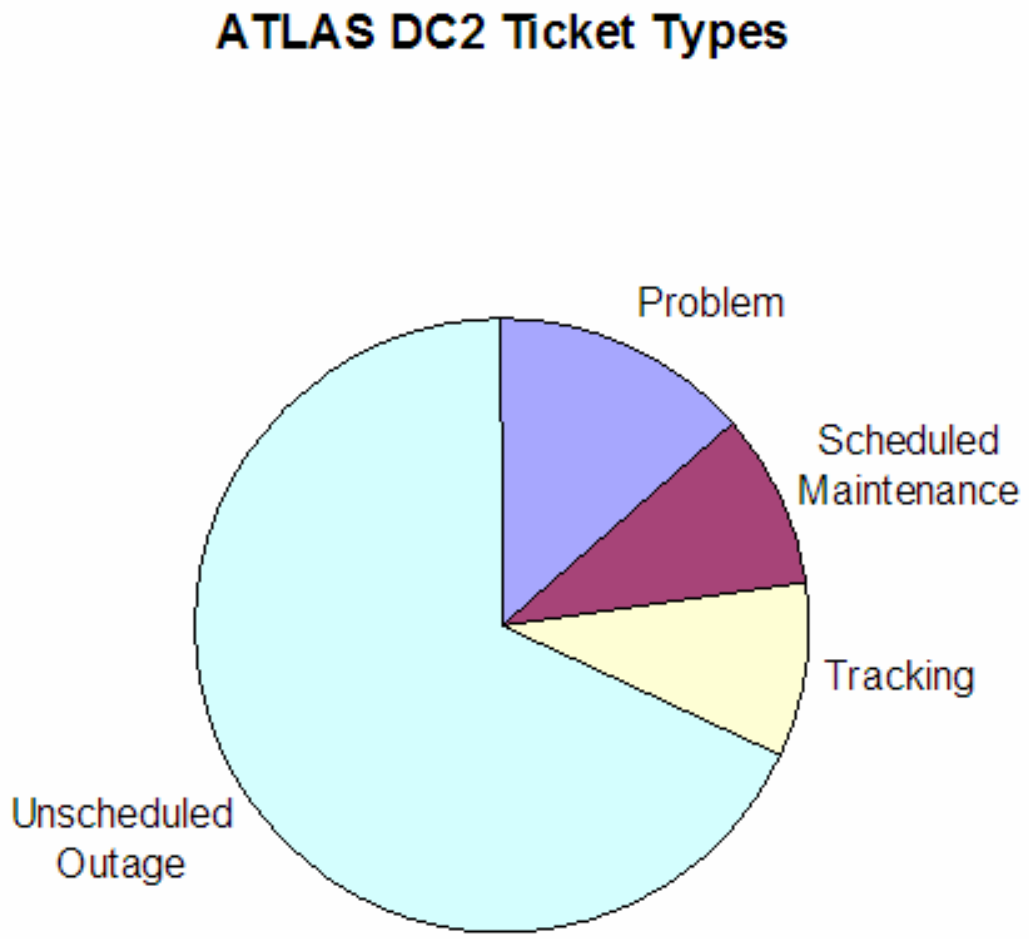

Catalog Site History Analysis

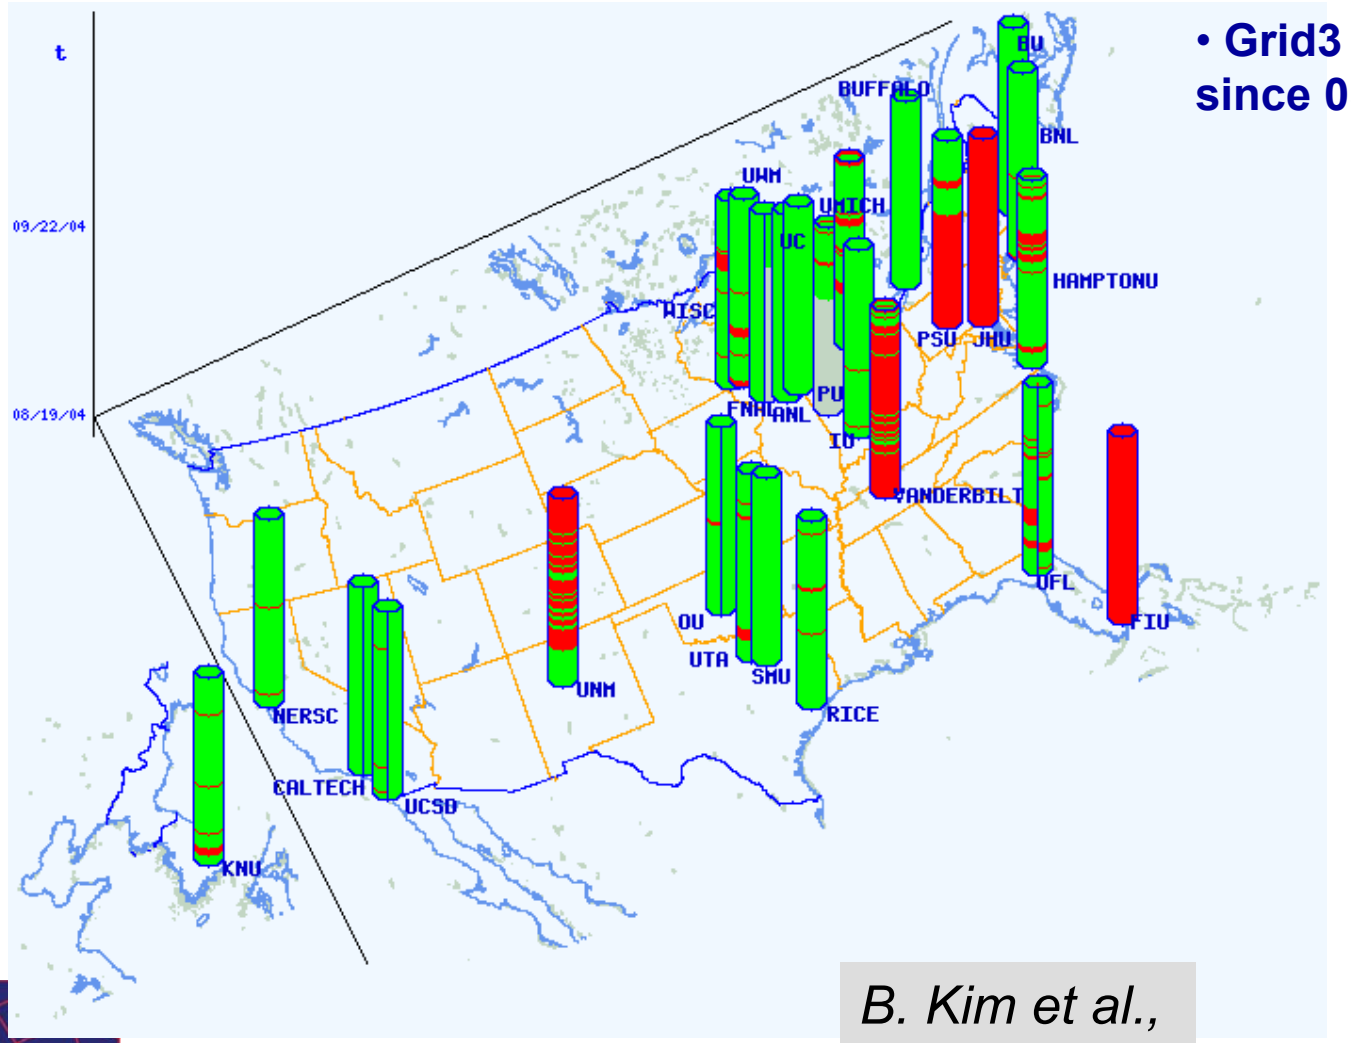

- **Grid3 status collected since 08/19/04**

B. Kim et al.,

Use of Grid3 – led by US LHC

- 7 Scientific applications and 3 CS demonstrators
  - A third HEP and two biology experiments also participated
- Over 100 users authorized to run on Grid3
  - Application execution performed by dedicated individuals
  - Typically ~few users ran the applications from a particular experiment

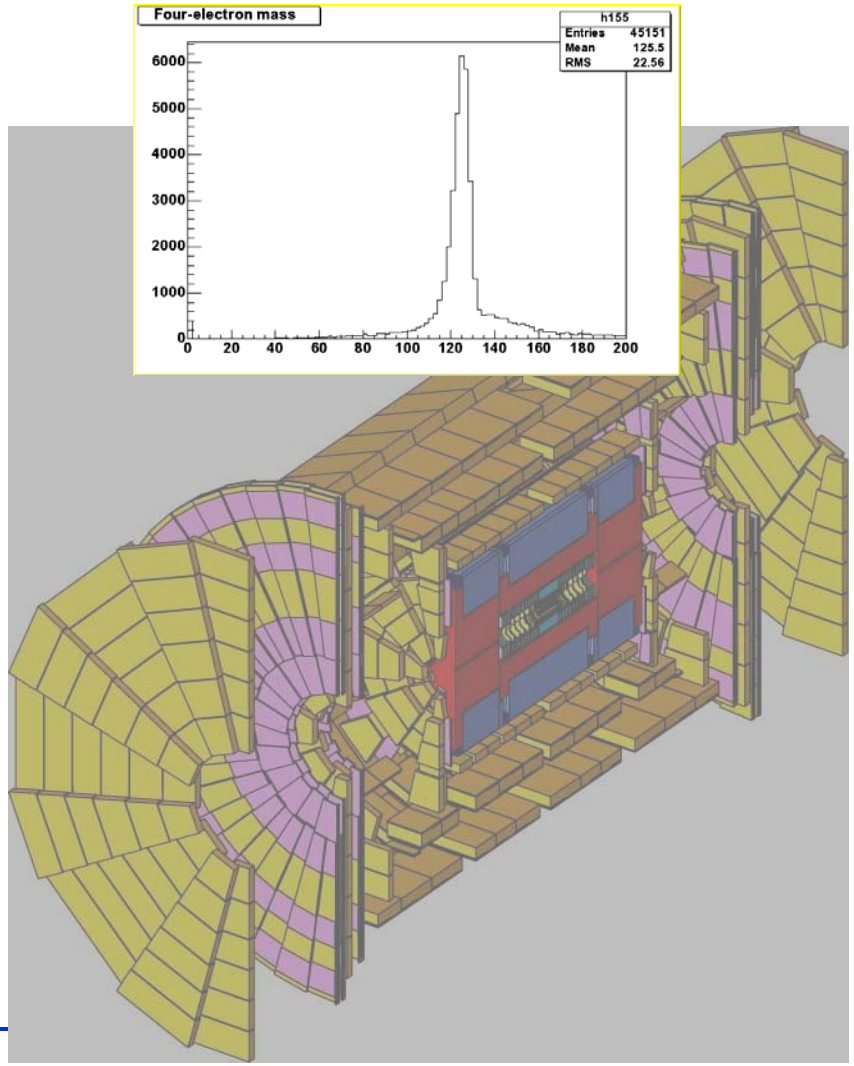

Usage of the Grid3 (6 months)

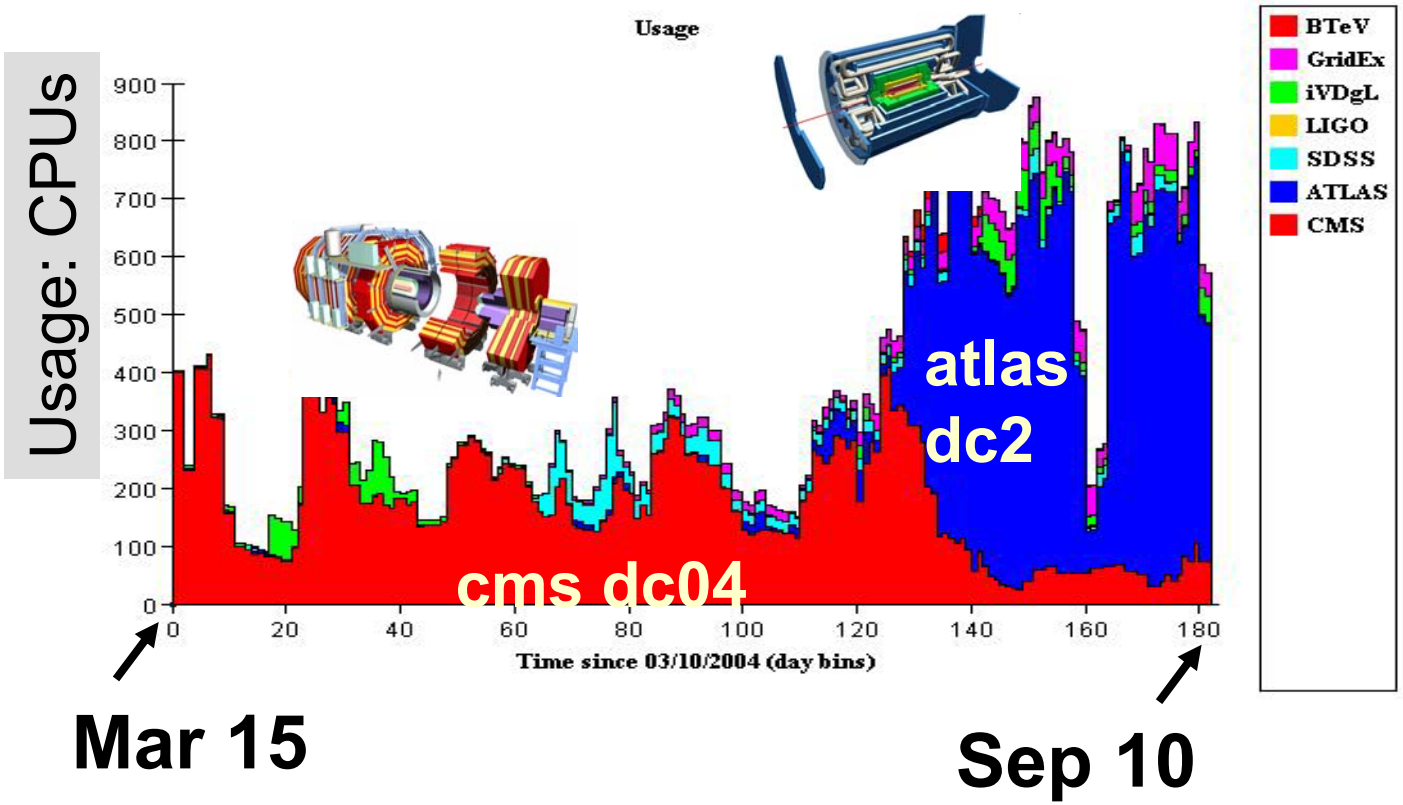

Lessons Learned

- Configuration management and assistance efforts in development and deployment are rewarded many times over during production.
- Middleware updates can be painless.
- Certificates are a hassle (just like all security)
- Not all resource information should be public
- A production monitoring infrastructure including people provides a significant problem solving advantage, esp. redundant monitoring.
- Resource providers and owners are more responsive and comfortable working with a central operations center.
- The GOC provides more than operations – it provides focus, continuity of effort, and community.

Agenda

- Introduction to iVDGL, Grid3, and the iGOC
- Efforts, Accomplishments and Lessons Learned
- Future Directions

http://www.ivdgl.org/grid3

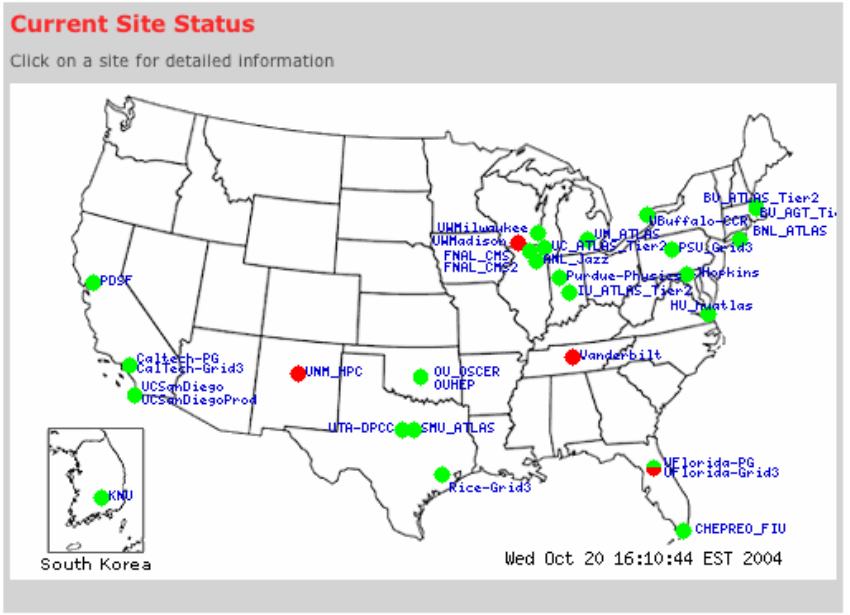

Please report problem on Grid3 to the iGOC's trouble ticket system. A link to the main page is provided here.

Grid3 Monitoring

Site Status Catalog
Operational status of Grid3 sites
MonALISA
A distributed monitoring service using JINI/JAVA and WSDL/SOAP
Ganglia
A distrbuted monitoring service based on mutlticast listen/annouce protocal.
ACDC Job Monitor
Grid3 real-time job monitoring tool

Grid3 is evolving into OSG

- Main features/enhancements
  - Storage Resource Management
  - Improve authorization service
  - Add data management capabilities
  - Improve monitoring and information services
  - Service challenges ***and interoperability with other Grids***
- Timeline
  - Current Grid3 remains stable through 2004
  - Service development continues
  - Grid3dev platform

Open Science Grid

c.f. R. Pordes

OSG deployment landscape

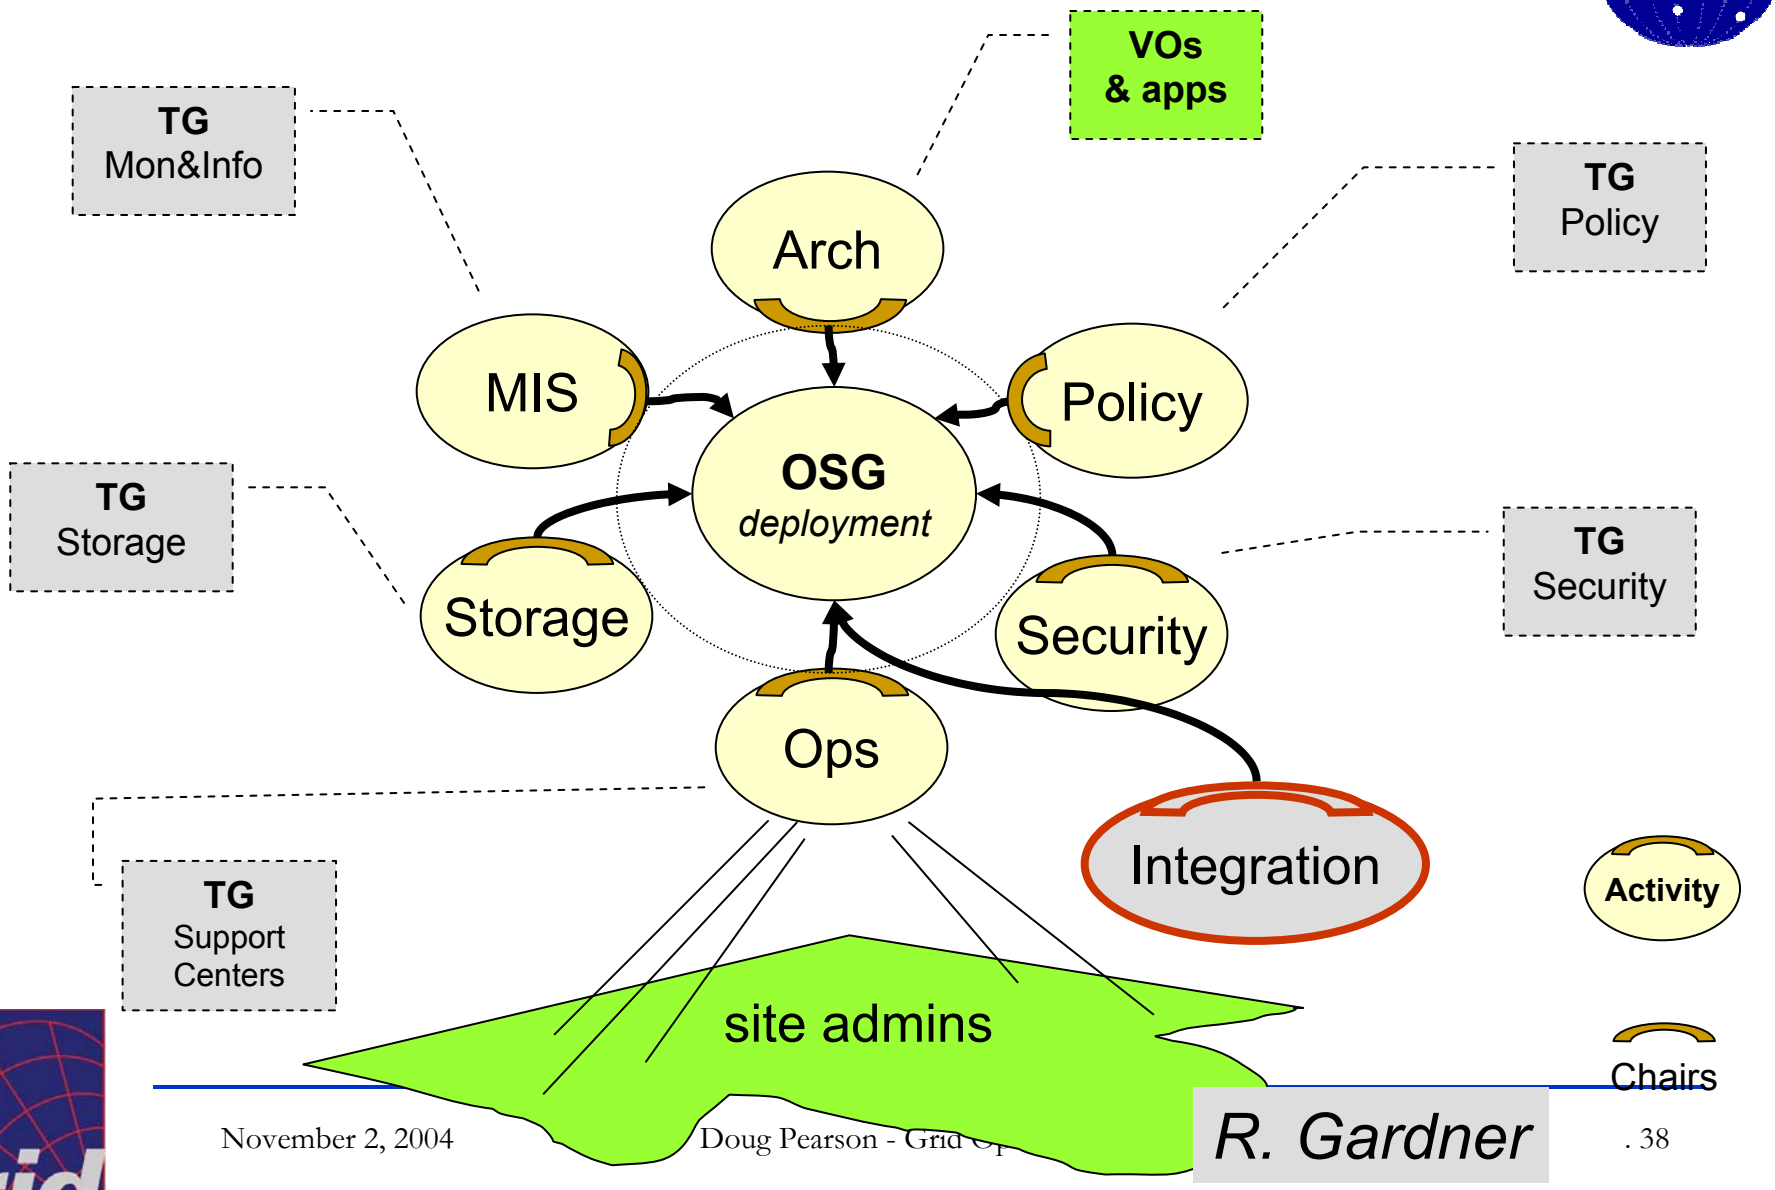

Support Centers Technical Group

- is responsible for discussing and coordinating the OSG activities that relate to support centers and services. These services include:
  - definition of the support model for user, infrastructure, service and technology support.
  - communication and publication of information for support helpdesk and trouble ticket infrastructures.
  - communication and interoperation with other grid infrastructures, in particular the LCG/EGEE.

Challenges

"OSG is a project with little central control or resources – almost everything has to be done by the sites or the VOs"

The GOC is demonstrated as a valuable central entity, minimally to facilitate, coordinate, establish software caches, monitor, assist in site installation, etc.

How to bring these two facts together?

THE END

EXTRA SLIDES

US CMS Data Challenge DC04

Opportunistic use of Grid3

non-CMS (blue)

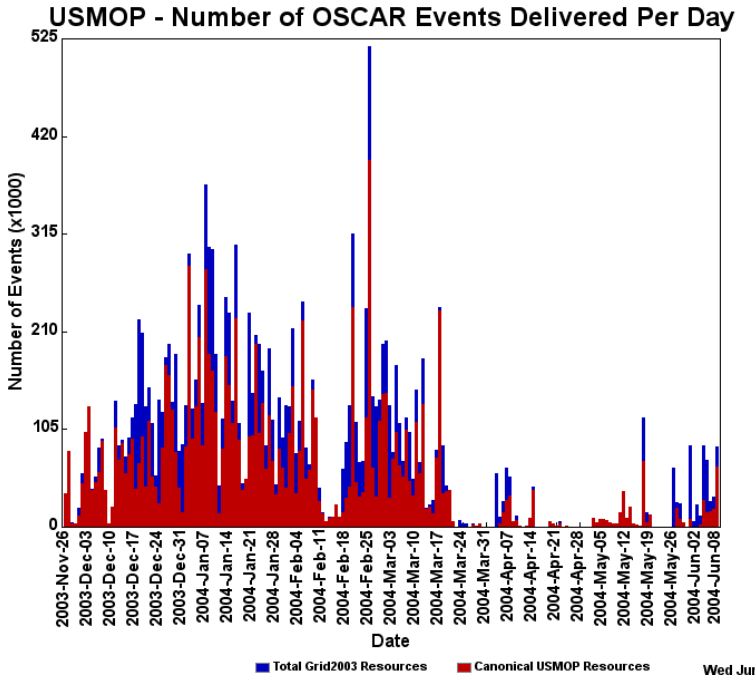

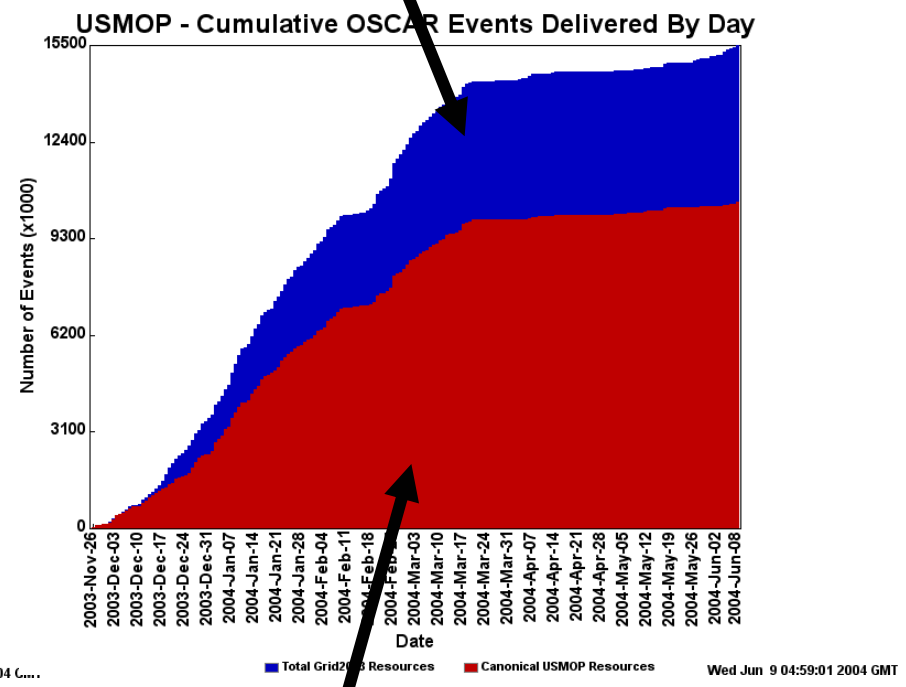

Events produced vs. day

CMS dedicated (red)

c.f. A. Fanfani, #497

Ramp up ATLAS DC2

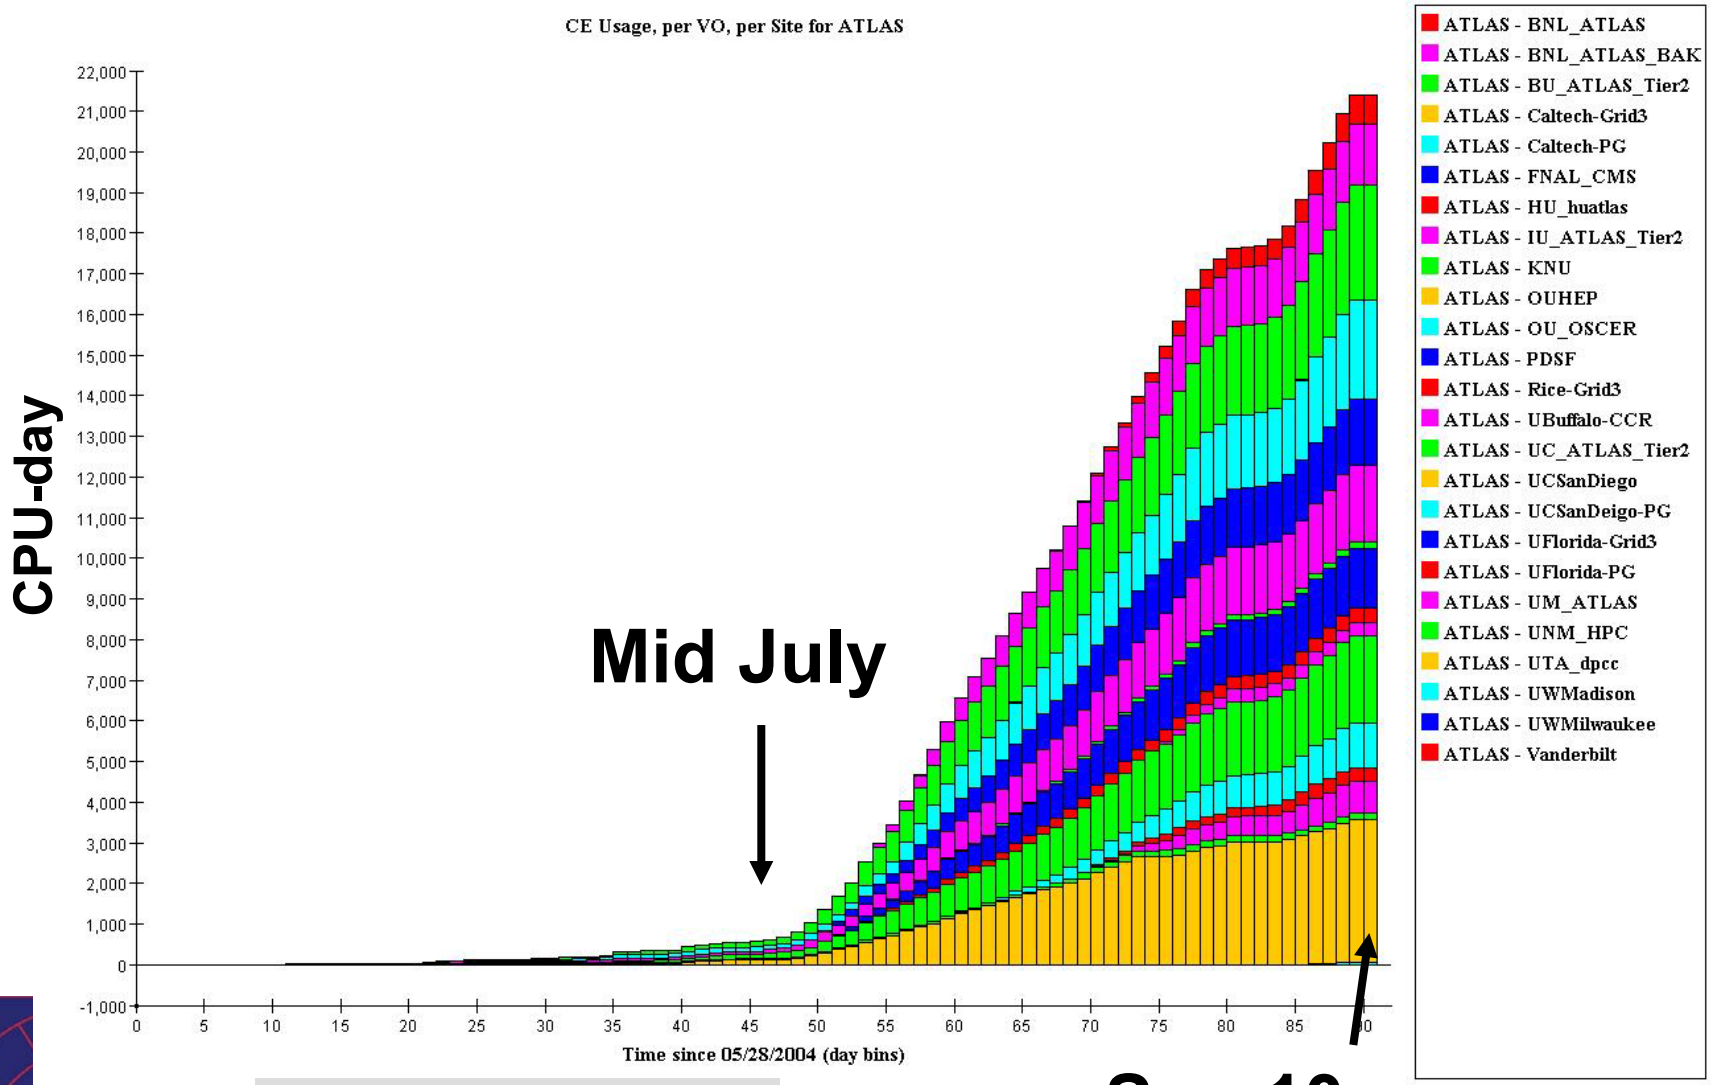

c.f. R. Gardner, et al.,

Resource Owners and Providers

- Pre-installation configuration (cluster configuration review/ batch queuing/ distributed file system)
- Software installation
- Configuration management
- Outages
- Capacity planning for storage
- Policy statement for CPU resources
- Policy enforcement
- Network and system performance

Virtual Organization Issues

- Adjustable trouble handling procedures
- 24x7 monitoring of specialized services
- Ratings for response levels of services “Critical”, “Elevated”, etc.
- Monitoring of VO’s grid services such as VOMS

Application Developers

- Provide specialized services for applications
- Create APIs to obtain published information from Site Resources
- Provide a liaison between VDT developers and Application developers
- Grid3 schema to publish file system location information for dynamic application installation

Support of Individual Grid users

- How to get and maintain a cert
- How to run an application
- What site policies are in place
- How to use monitoring tools
- Troubleshooting application failures
- Managing datasets
- Joining a VO

Catalog -- Site Status

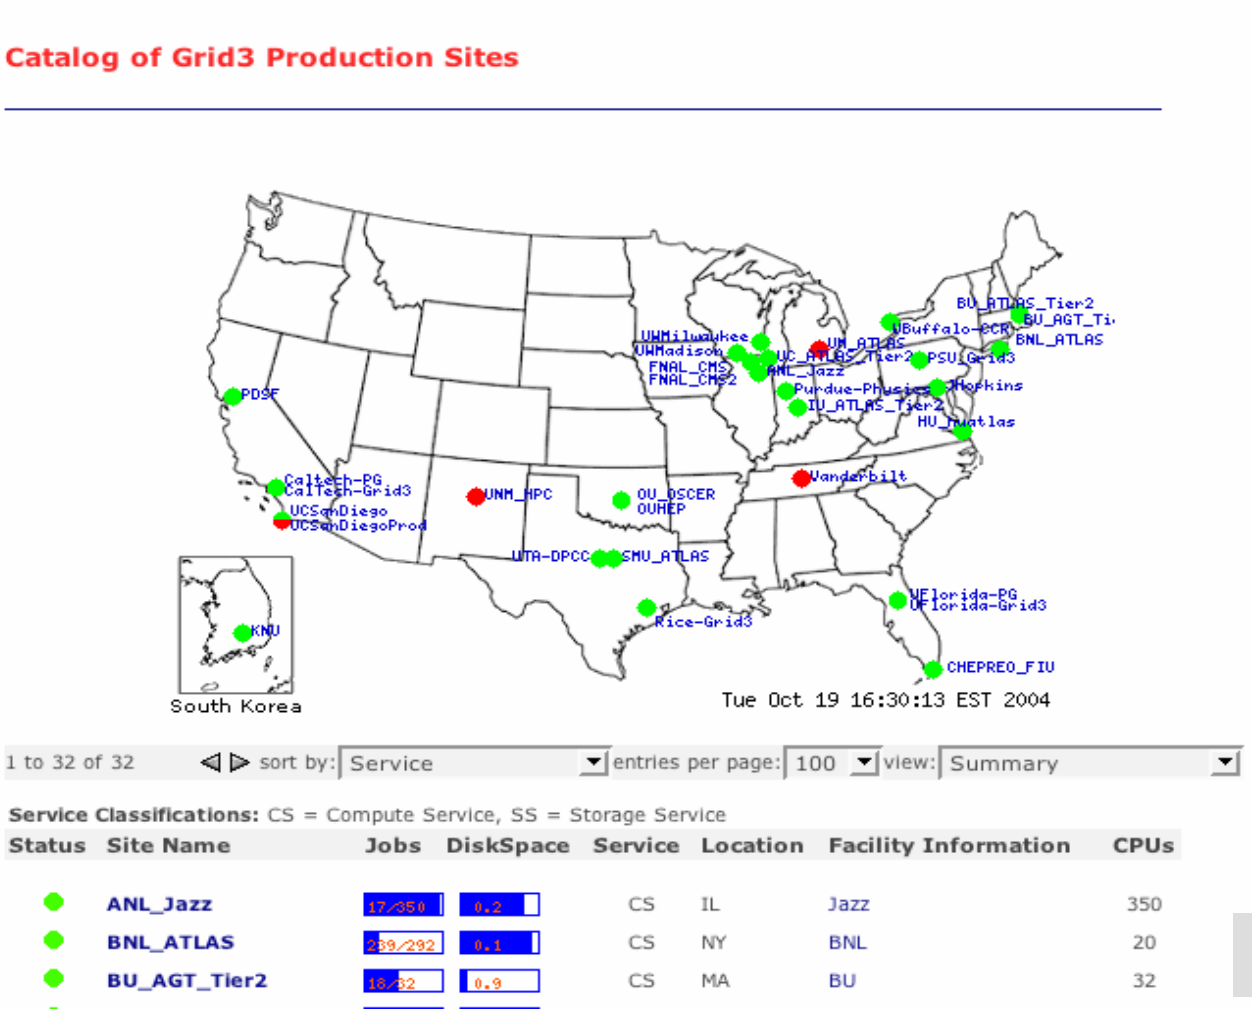

- **Status map on**
- **Facility--> Sites**
- **Grid test results clickables**
- **Dynamic CPU/Disk info**
- **Optional views: different information**
- **Map : US-Korea map or Worldmap <- New release**

c.f. B. Kim et al.,

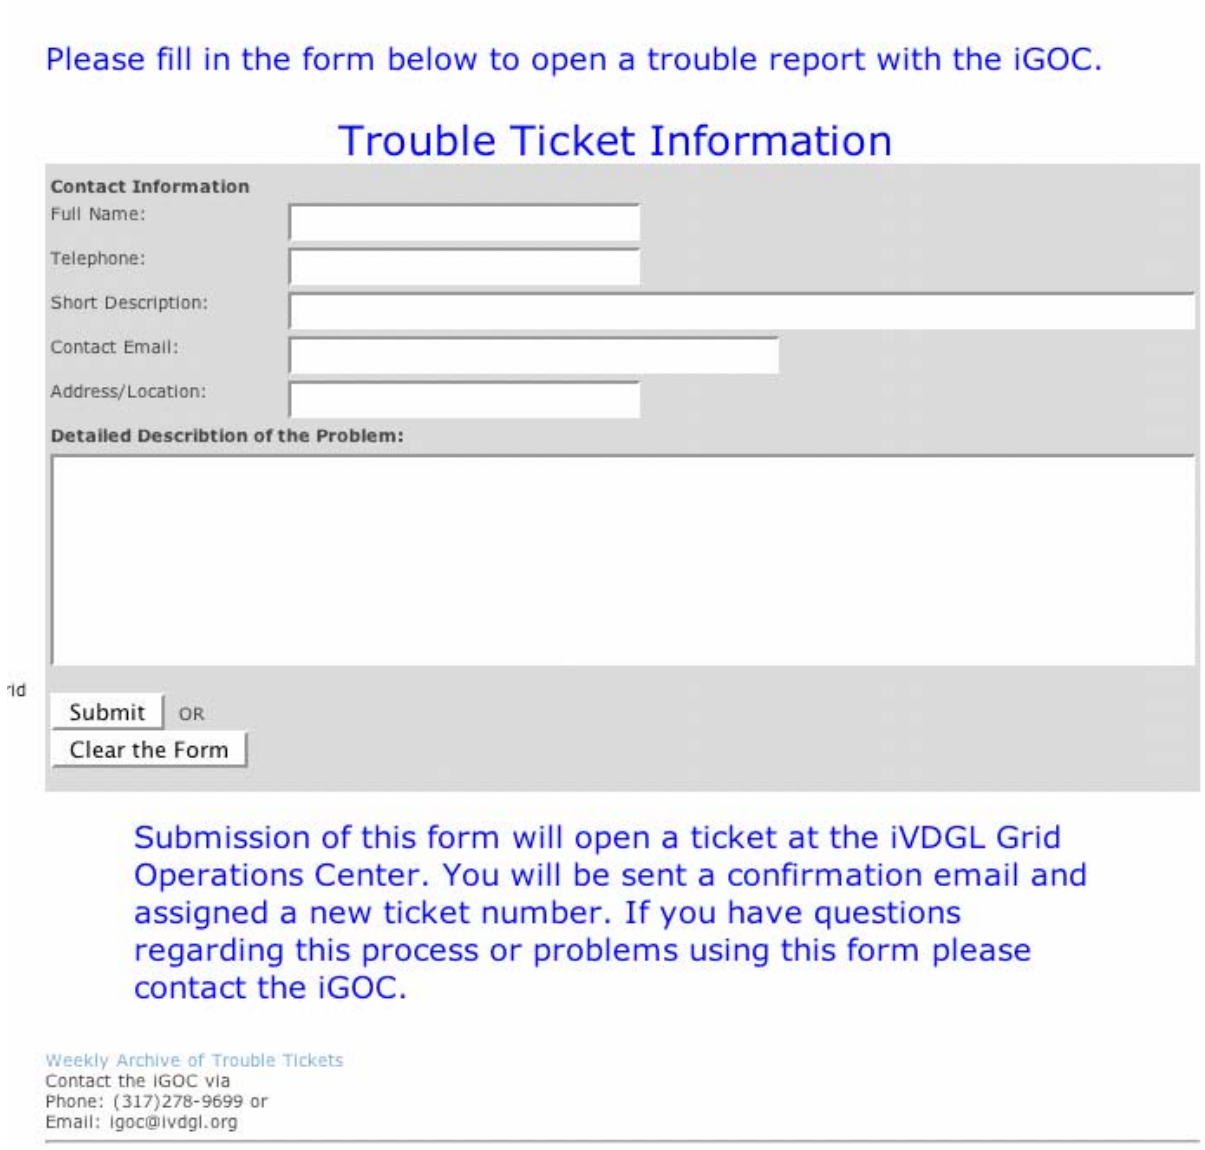

Web form to open a Trouble Ticket

Or send email to igoc@ivdgl.org