

TeraGrid Information Services: Building on Globus MDS4

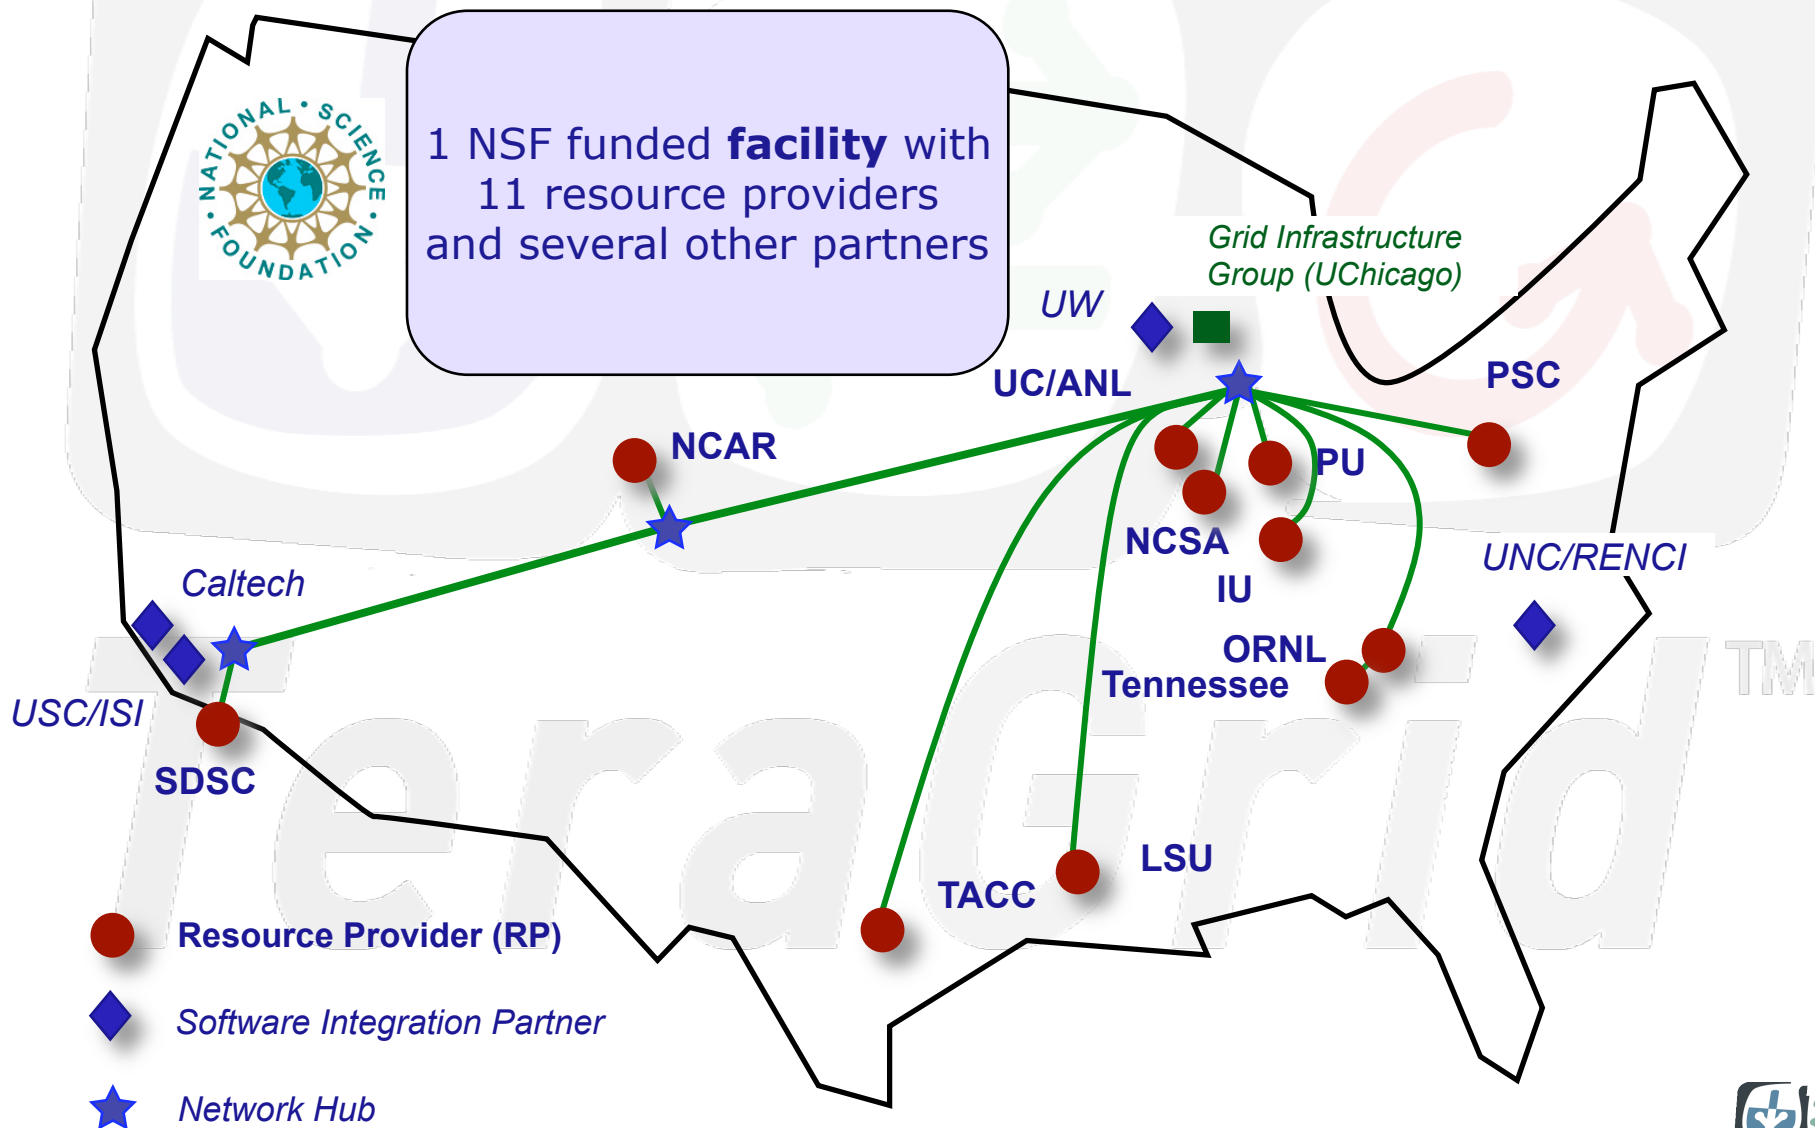
John-Paul "JP" Navarro

TeraGrid Grid Infrastructure Group "GIG"
Area Co-Director for Software Integration
University of Chicago, Argonne National Laboratory

GlobusWorld
May 13, 2008

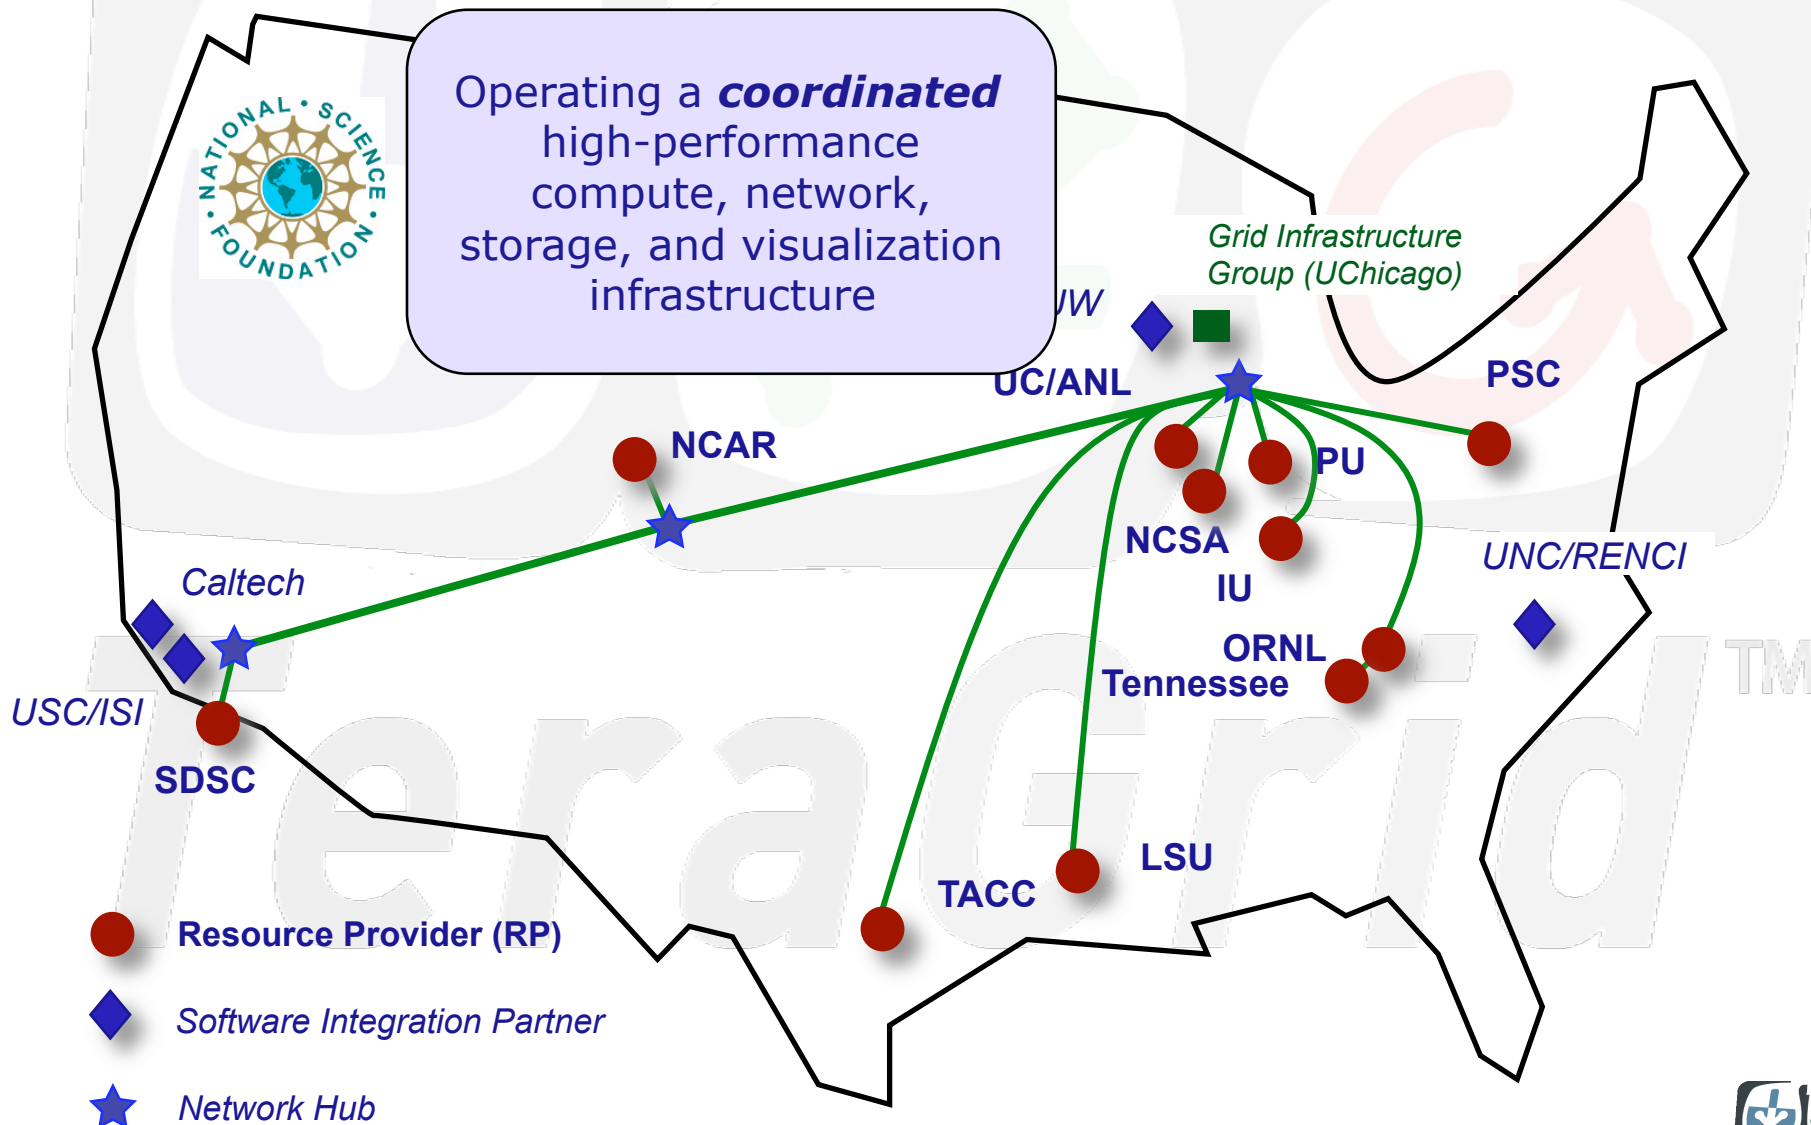


What is the TeraGrid?

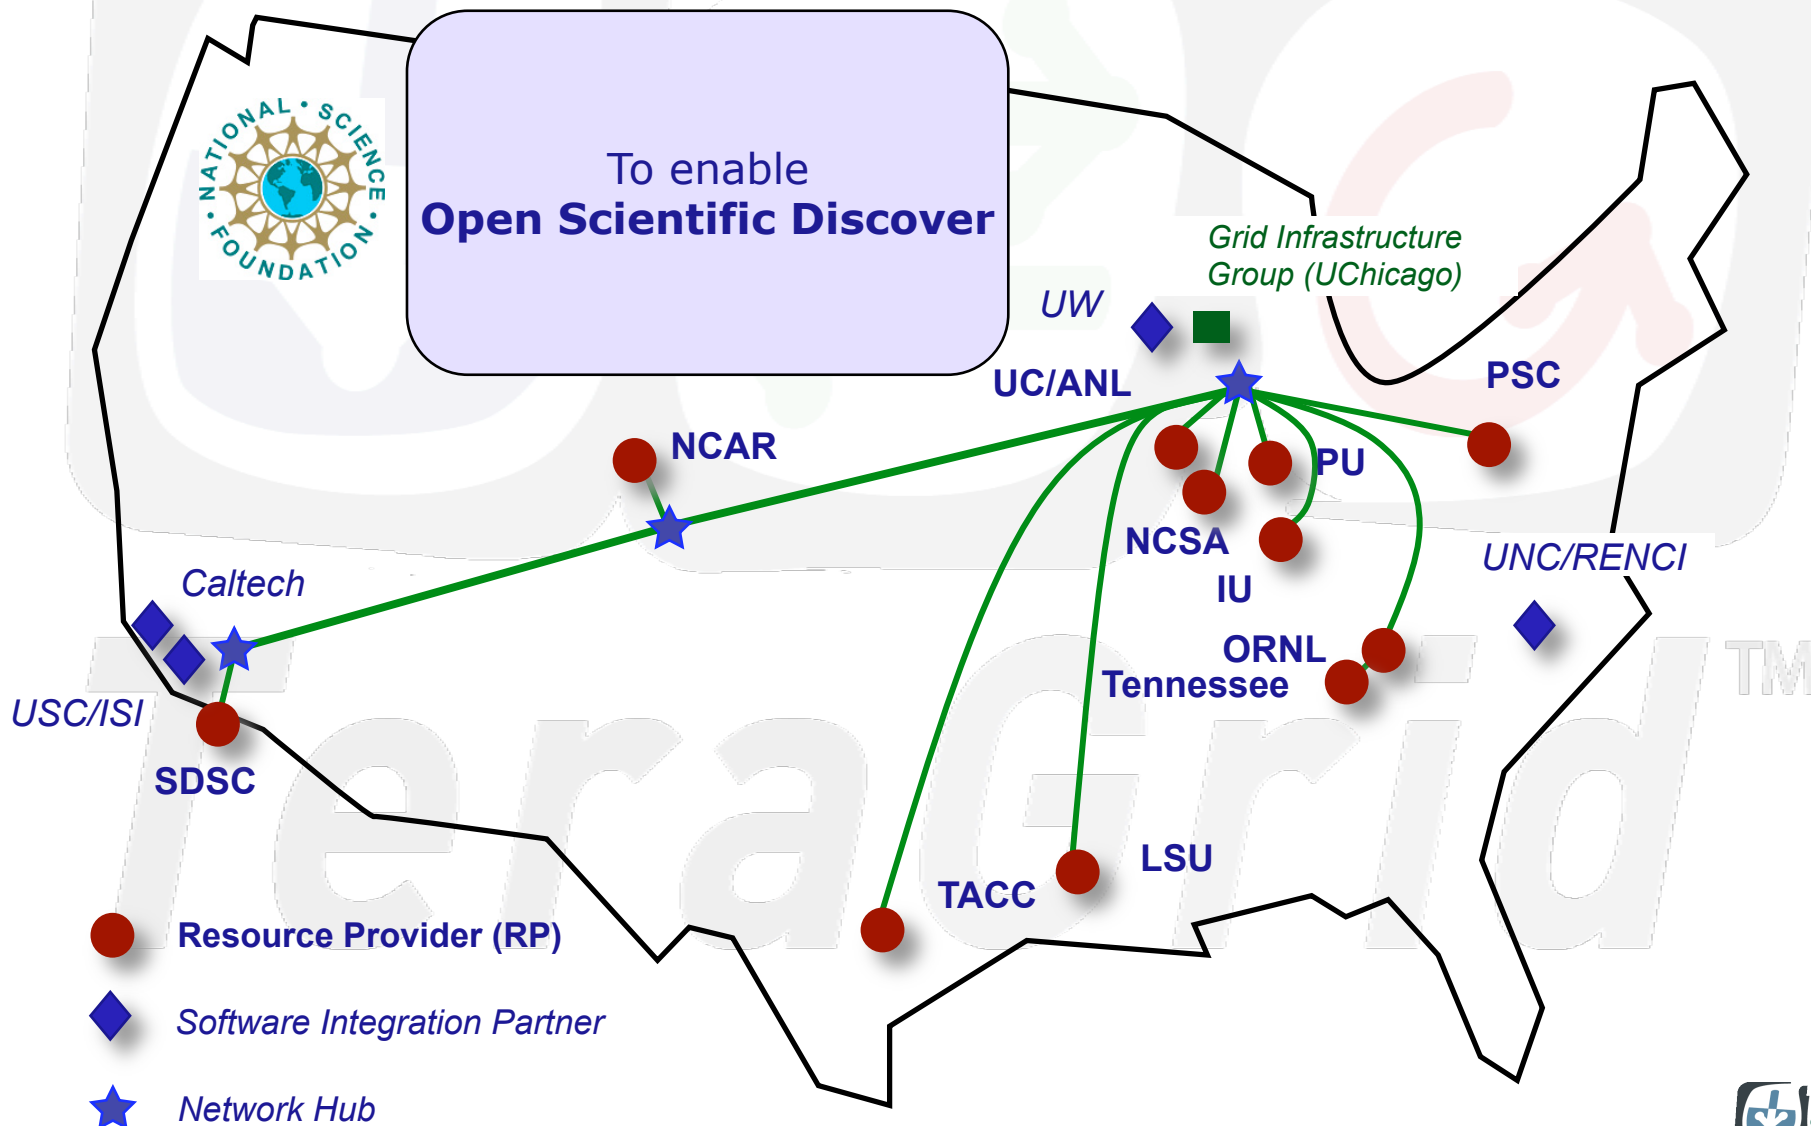


What is the TeraGrid?

Operating a **coordinated** high-performance compute, network, storage, and visualization infrastructure



What is the TeraGrid?



One facility in what sense?

Common People Interfaces, examples:


- Request allocations & accounts "POPS"
- Identify available resources
- Learn how to use resources
- Find resource status
- Ask for help
- Community events



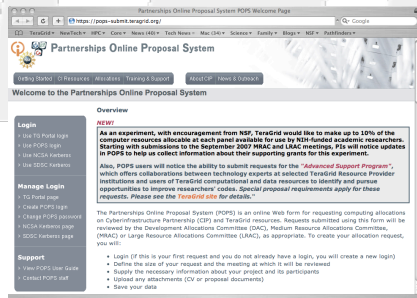
Events



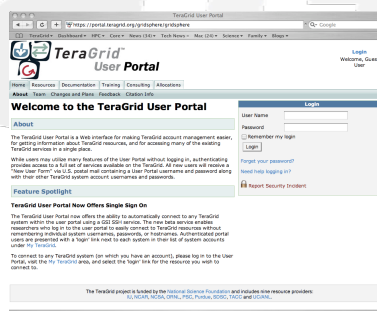
Helpdesk



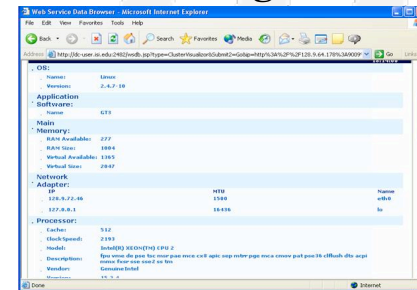
Proposal System "POPS"



User Portal



User Documentation & Knowledge Base



Challenges

Integrating many **more types of infrastructure resources**,
Including **more services providers**,
Providing **more types of services**,
With more resources, providers, and services diversity,
Supporting specialization.



How can users keep it all straight?

Meeting the Challenge 1

Re-architect CTSS

Took Coordinated TeraGrid Software and Services "CTSS" v3:

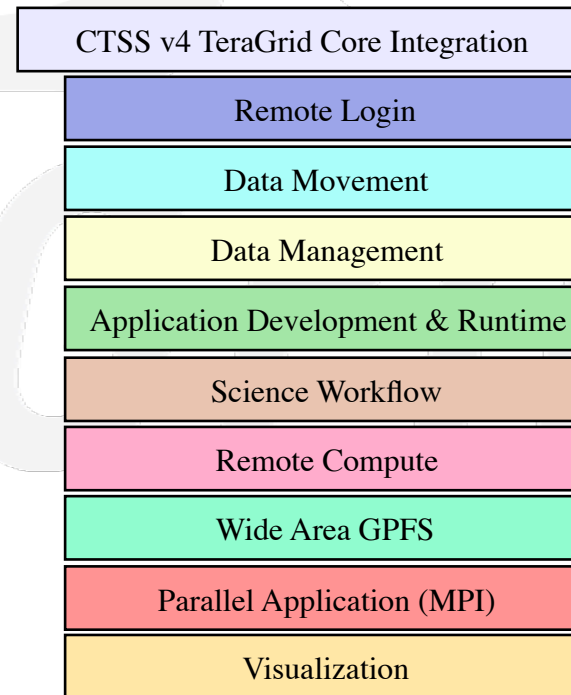
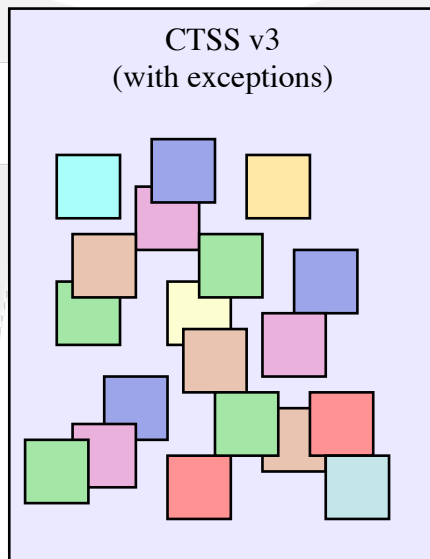
Monolithic coordinated capabilities

In production since mid 2006

Re-architected into CTSS v4 "capability kits":

Small required **core integration kit** with minimum needed to integrate resources

Plus ~10 optional **user capability kits** (at least 1 required)



Meeting the Challenge 2

Information Services

Information Services Vision:

- 1) Create a **coordinated way for TeraGrid participants to publish** about the services they offer,
- 2) Create a **way for the TeraGrid to aggregate and index the** information from TeraGrid participants, and
- 3) to **publish this information to the public in a form that can easily be used** by user software, user interfaces, and TeraGrid service providers themselves to

discover capabilities and how to access them

Information Services Design Goals

Applies Grid concepts to *original* information publishing

- ① Publishing is the responsibility of the information owner
- ② Publishing is done using standard (content) schemas
- ③ Publishing thru standard interfaces regardless of content and where the data comes from
- ④ Publishing services should be available globally (subject to authentication/authorization)
- ⑤ **Information owners publish to EVERYONE, not just the TeraGrid**
- ⑥ Publishing is a grid service

Applies Grid concepts to *aggregated* information publishing

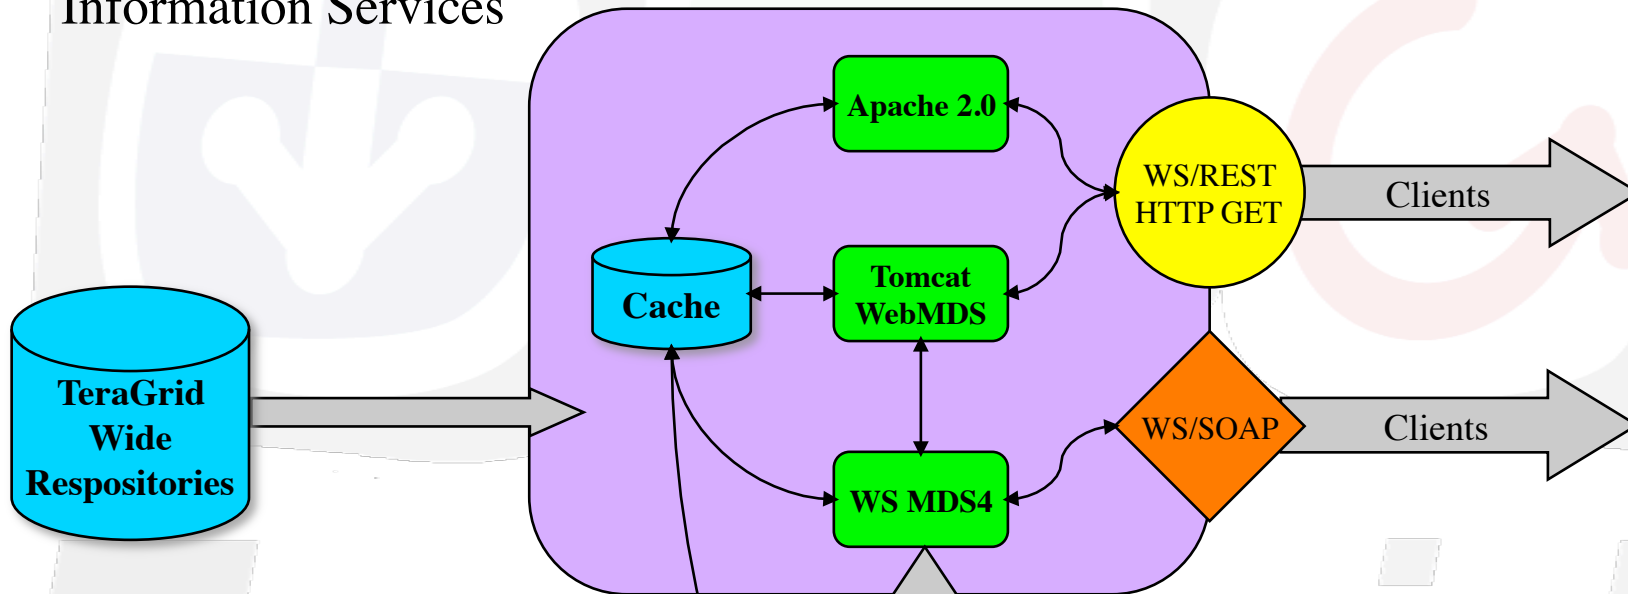
- ① Aggregation uses standard information services interfaces to retrieve information
- ② Publishing aggregated information is done exactly like original information publishing
- ③ This is how a collaboration, such as the TeraGrid, aggregates participant information

Applies Grid concepts to querying information


- ① Querying can use standard interfaces regardless of content
- ② Querying can use standard interfaces for original or aggregated information

High-Level Architecture

TeraGrid Wide
Information Services



Service Provider
Information Services



Information Services Tooling

WS/SOAP (Globus 4.0.x MDS4)

Benefits

- Indexing, Trigger
- Registration, Publish, Subscribe
- Security/Authorization
- Robust WSRF interface

Content

- XML

WS/* (Tomcat 5.0, Apache 2.0)

Benefits

- Very common web services platform
- Supports several web service interfaces (including simple)
- Supports multiple styles like REST, Web 2.0
- Can be highly scalable

Content

- Many formats: HTML, XHTML/XML, XML, RSS/Atom, ...

WebMDS (Globus 4.0.x)

Benefits

- Live MDS4 content access
- XPath support
- XSLT transforms

Content

- Many formats: HTML, XHTML/XML, XML, RSS/Atom



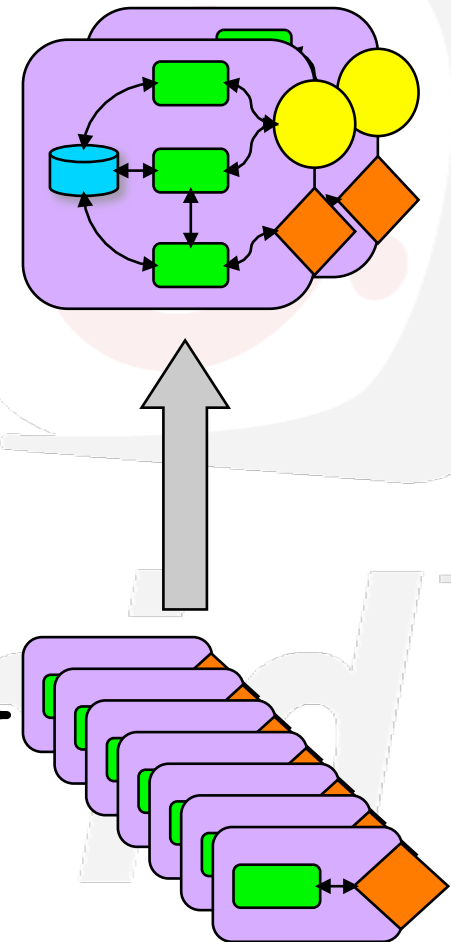
Service Provider vs TG Wide Services

Content:

Locally owned and published information
Can come from existing local systems

Services:

1 general purpose MDS service
1 remote execution MDS services



Service Provider vs TG Wide Services

Content:

Aggregate/index service provider information
Plus central information (TeraGrid databases)
Cached
Authenticated registrations

Services:

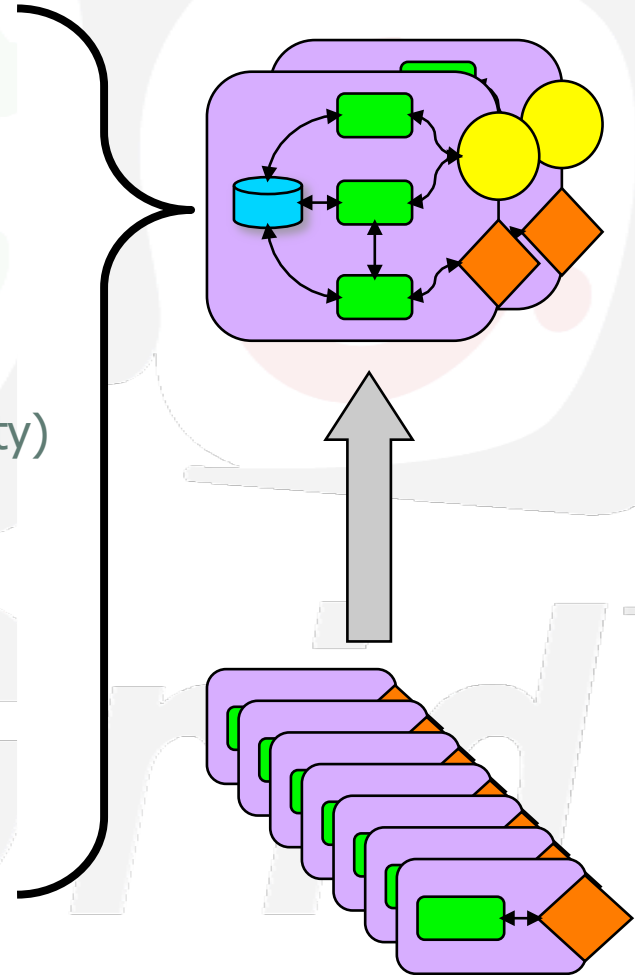
Several redundant servers (>99.5% availability)

Each server:

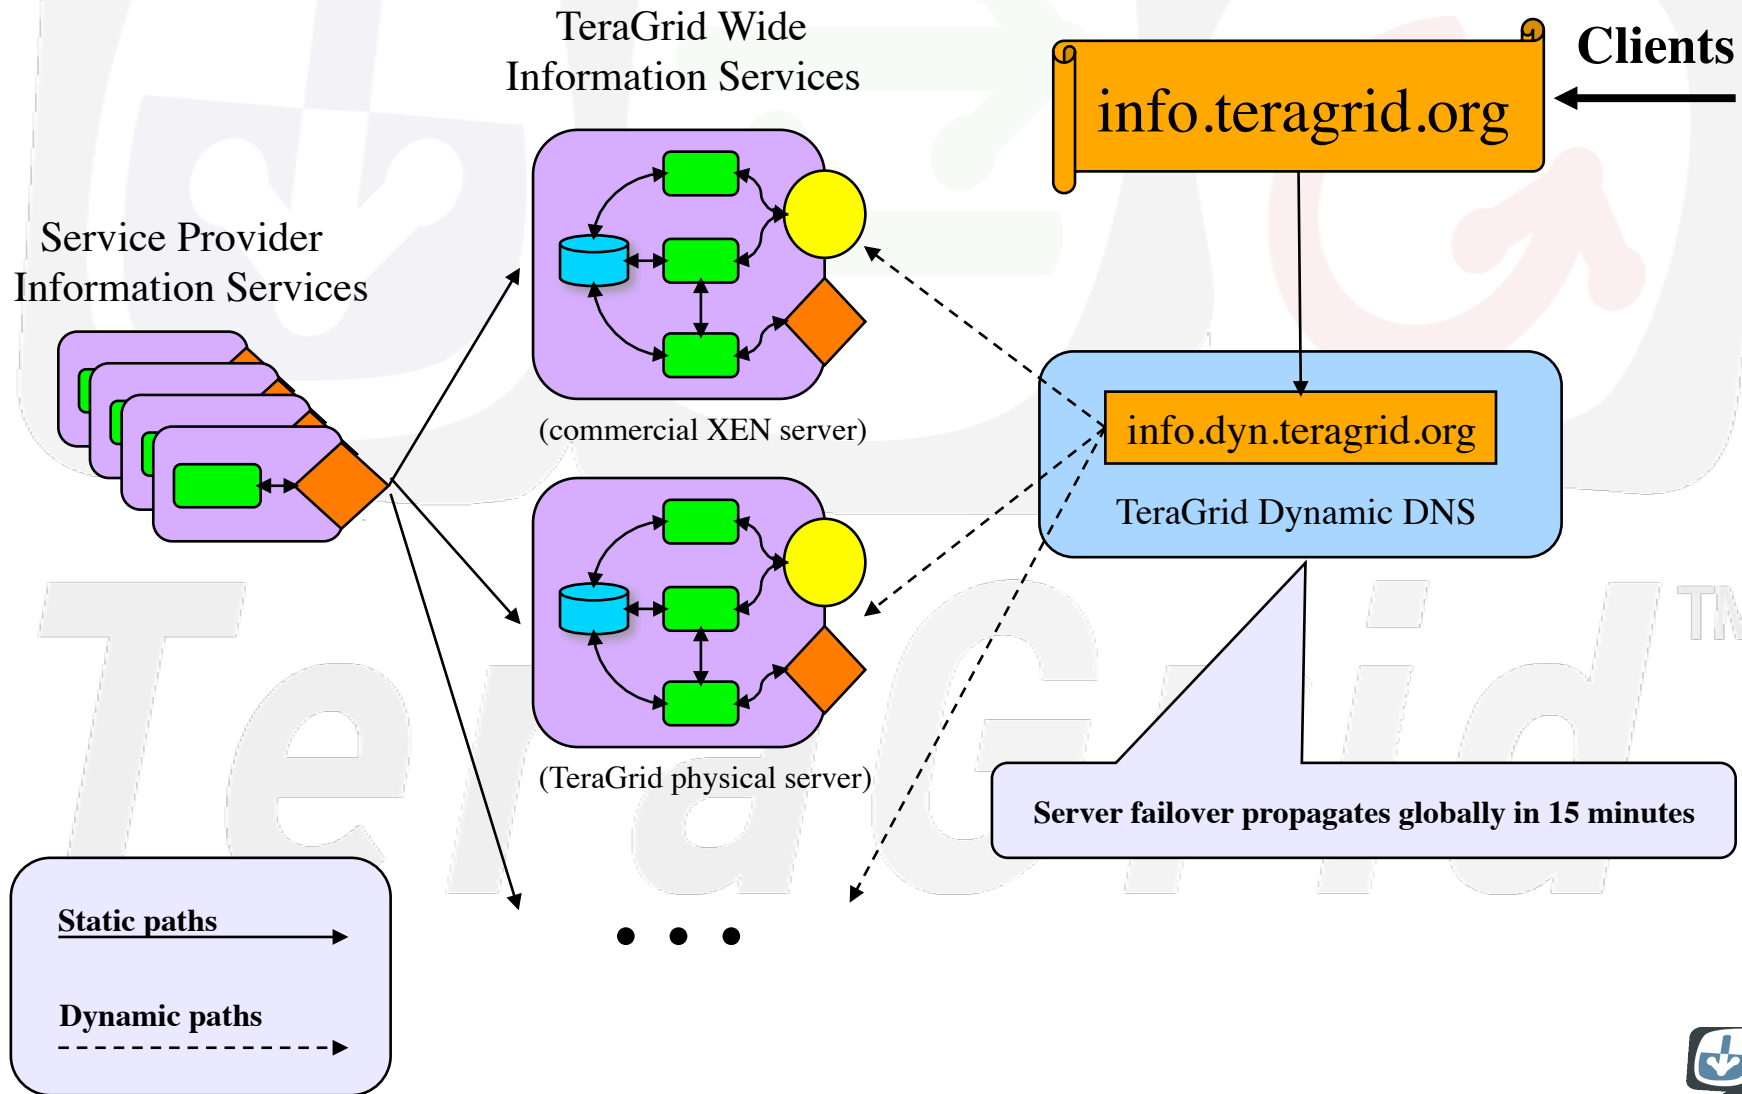
- Information caching programs
- MDS4 index services (WS/SOAP)
- WebMDS/Tomcat, Apache 2.0, ... services (WS/REST)

Services publish in:

- HTML
- XML
- CSV
- (Atom, JSON, RSS being prototyped)



High-Availability Design



Caching Design

Goal

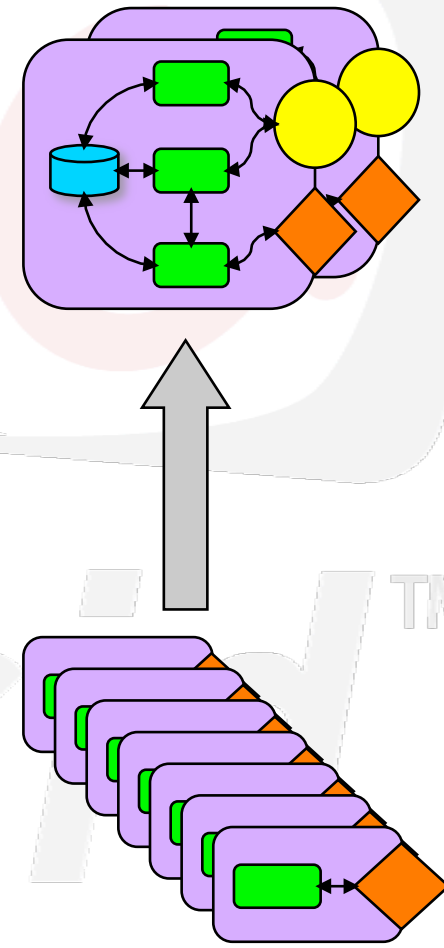
Information persistence when service providers are down
and when central information services go down

MDS configuration:

Service provider register their existence
Central MDS do not aggregate automatically
Central MDS publishes from local file-system cache

Caching programs:

External and asynchronous to MDS
Query registered
Service provider MDSs, and
Other configured MDSs
Store query results in file-system
Replaces previously cached data only with new data
Not if downstream is down
Or if downstream returns an error
Or if downstream doesn't return valid data



Primary MDS4 features used

Registration:

Parallel upstream registration to multiple central servers
Authenticated/authorized upstream registration

Aggregation:

Some information is automatically aggregated
(WS services registration and default GLUE)

Information providers:


Useful RP custom information providers
Scheduling load and queue contents (User Portal)
CTSS 4 Capability kit registration
TeraGrid ID and description cross-reference (TGCDDB)

Secure MDS:

Authenticated/authorized queries (non-anonymous)
Multiple index services in same container
DefaultIndexService AND SecureIndexService

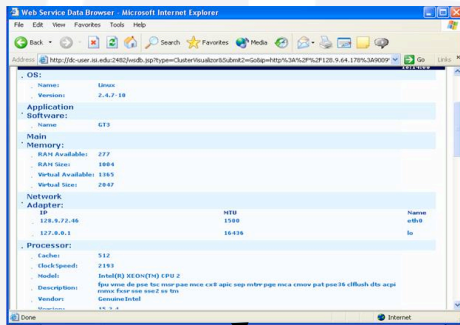
WebMDS

XSLT, Xpath

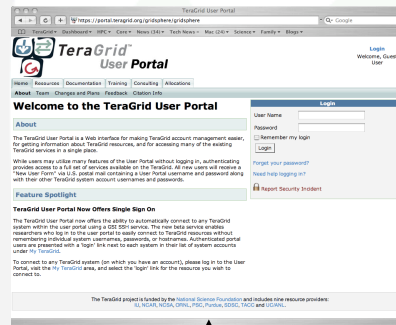


Information Services Users

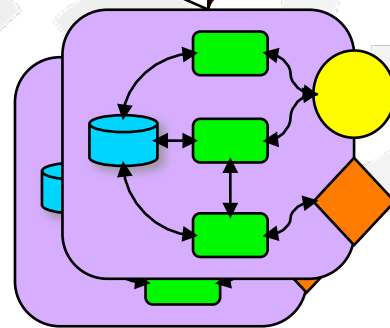
User Documentation
<http://www.teragrid.org/>



User Portal
<http://portal.teragrid.org/>



Inca Testing Harness



Gateways

Peer Grids

User Applications

info.teragrid.org



TeraGrid™

CTSS Capability Kit Availability

TeraGrid [User Support: View CTSS Results]

http://www.teragrid.org/userinfo/software/ctss_results.php

View CTSS Kits Available on TeraGrid

Home > User Info > Software > CTSS > Results

CTSS Resources and Available Kits [\(Read the Kit Definitions\)](#).

Site: Resource	TeraGrid Core Integration Kit	Remote Login Capability Kit	TeraGrid Remote Compute Kit	Application Development and Runtime Support Capability Kit	TeraGrid Data Movement Kit	Parallel Application Capability Kit	Science Workflow Support Kit	TeraGrid Data Management Capability Kit	Data Visualization Support Kit
NCAR: Frost	currently available	currently available	currently available	currently available	currently available	currently available	currently available	Not Installed	Not Installed
PSC: Big Ben	currently available	currently available	currently available	currently available	currently available	currently available	Not Installed	currently available	Not Installed
PSC: Rachel	currently available	currently available	currently available	currently available	currently available	currently available	Not Installed	Not Installed	Not Installed
Purdue: Condor	currently available	currently available	currently available	currently available	currently available	Not Installed	currently available	currently available	Not Installed
Purdue: Lear	currently available	currently available	currently available	currently available	currently available	currently available	currently available	currently available	Not Installed
TACC: Lonestar	currently available	currently available	currently available	currently available	currently available	currently available	currently available	Not Installed	Not Installed
TACC: Maverick	currently available	currently available	currently available	currently available	currently available	currently available	currently available	Not Installed	currently available
NCSA: Cobalt	currently available	currently available	currently available	currently available	currently available	currently available	currently available	Not Installed	Not Installed
NCSA: Abe	currently available	currently available	currently available	currently available	currently available	currently available	currently available	Not Installed	Not Installed
NCSA: IA-64 Cluster	currently available	currently available	currently available	currently available	currently available	currently available	currently available	Not Installed	Not Installed



Where are the GridFTP services?

TeraGrid [User Support: CTSS Software Search Results]

http://www.teragrid.org/userinfo/software/ctss_soft.php

TeraGrid Dashboard HPC Core News (33) Tech News Mac (17) Science Family Blogs

Details of your selection(s)

Service selection: globus-gridftp-server

NCAR Blue Gene/L (Frost)

Type : gridftp
Version : 4.0.3
Endpoint : gsiftp://gridftp.frost.ncar.teragrid.org:2811/
Name : globus-gridftp-server

PSC XT3 (BigBen)

Type : gridftp
Version : 4.0.3
Endpoint : gsiftp://gridftp.bigben.psc.teragrid.org:2811/
Name : globus-gridftp-server

PSC Alpha EV7 (Rachel)

Type : gridftp
Version : 4.0.3
Endpoint : gsiftp://gridftp.psc.teragrid.org:2811/
Name : globus-gridftp-server

Purdue Condor Pool

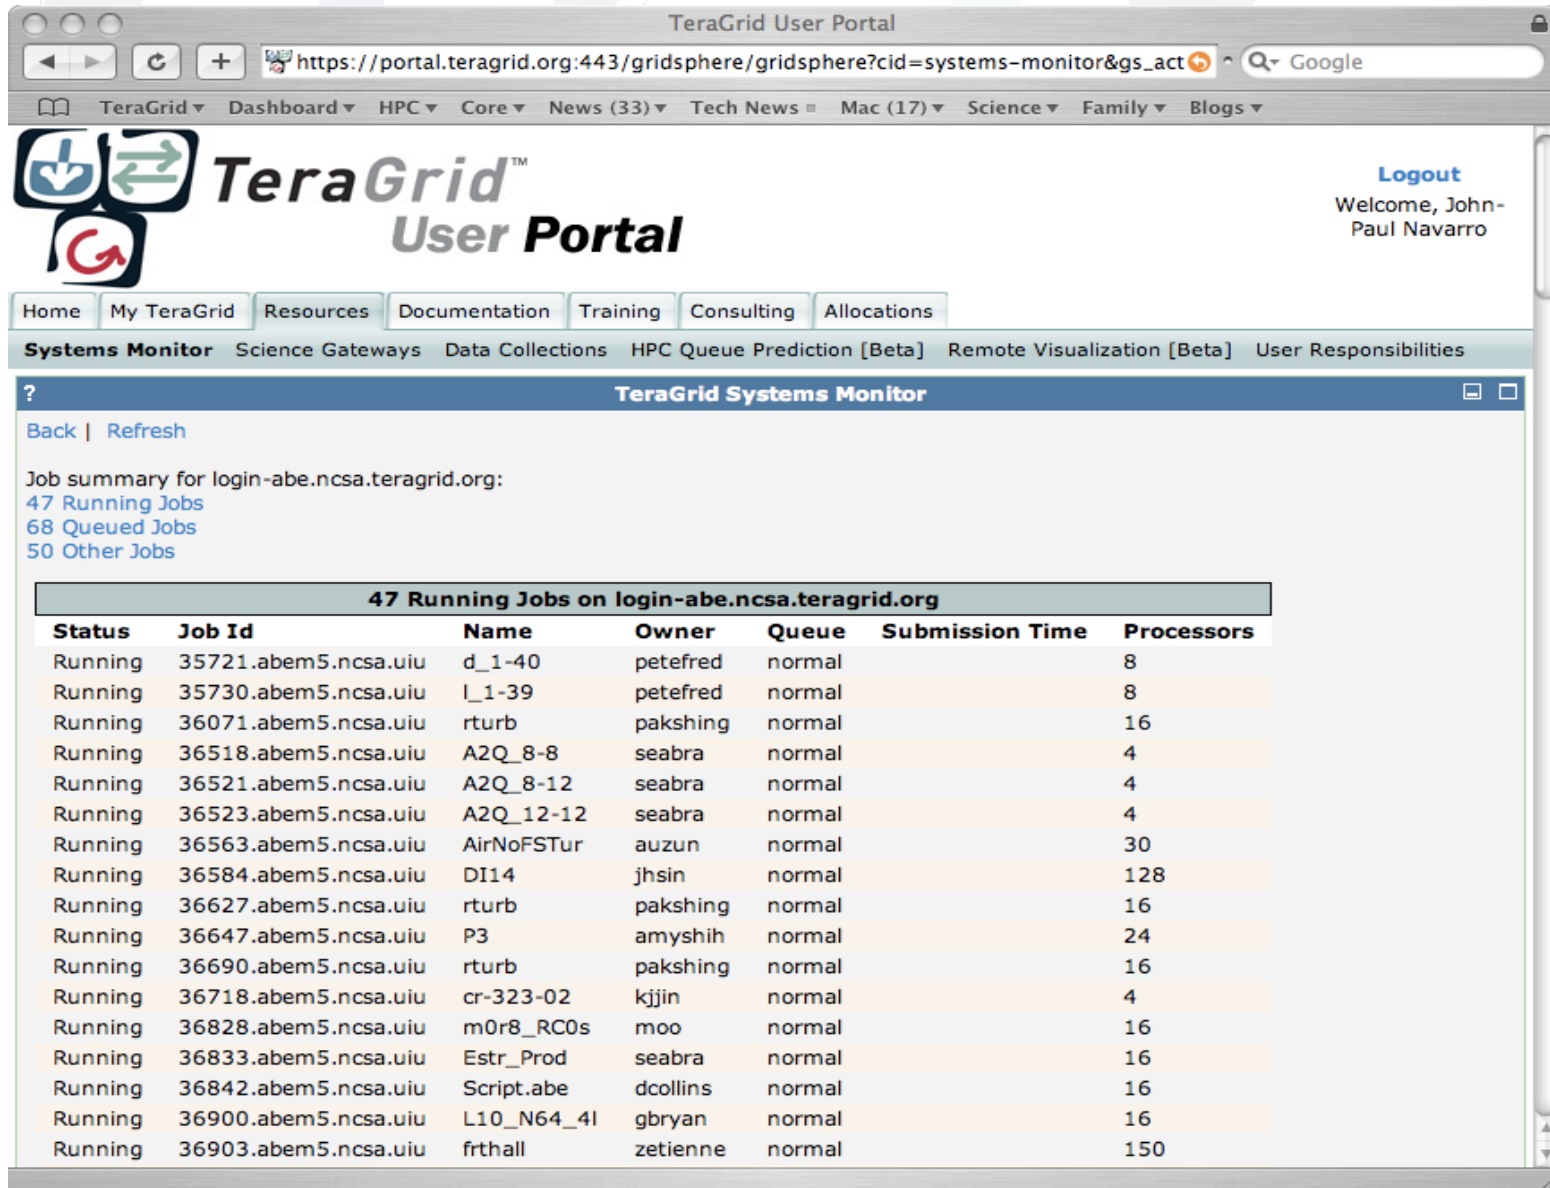
Type : gridftp
Version : 4.0.3
Endpoint : gsiftp://tg-data.purdue.teragrid.org:2811/
Name : globus-gridftp-server

Purdue EM64T Cluster (Lear)

Type : gridftp
Version : 4.0.3
Endpoint : gsiftp://tg-data.purdue.teragrid.org:2811/



Queue Contents in User Portal



The screenshot shows the TeraGrid User Portal interface. The browser address bar displays the URL: https://portal.teragrid.org:443/gridsphere/gridsphere?cid=systems-monitor&gs_act. The page title is "TeraGrid User Portal". The navigation menu includes: Home, My TeraGrid, Resources, Documentation, Training, Consulting, Allocations, Systems Monitor, Science Gateways, Data Collections, HPC Queue Prediction [Beta], Remote Visualization [Beta], and User Responsibilities. The main content area is titled "TeraGrid Systems Monitor" and shows a job summary for "login-abe.ncsa.teragrid.org":

- 47 Running Jobs
- 68 Queued Jobs
- 50 Other Jobs

The table below displays the details for the 47 running jobs:

Status	Job Id	Name	Owner	Queue	Submission Time	Processors
Running	35721.abem5.ncsa.uiu	d_1-40	petefred	normal		8
Running	35730.abem5.ncsa.uiu	l_1-39	petefred	normal		8
Running	36071.abem5.ncsa.uiu	rturb	pakshing	normal		16
Running	36518.abem5.ncsa.uiu	A2Q_8-8	seabra	normal		4
Running	36521.abem5.ncsa.uiu	A2Q_8-12	seabra	normal		4
Running	36523.abem5.ncsa.uiu	A2Q_12-12	seabra	normal		4
Running	36563.abem5.ncsa.uiu	AirNoFSTur	auzun	normal		30
Running	36584.abem5.ncsa.uiu	DI14	jhsin	normal		128
Running	36627.abem5.ncsa.uiu	rturb	pakshing	normal		16
Running	36647.abem5.ncsa.uiu	P3	amyshih	normal		24
Running	36690.abem5.ncsa.uiu	rturb	pakshing	normal		16
Running	36718.abem5.ncsa.uiu	cr-323-02	kjjin	normal		4
Running	36828.abem5.ncsa.uiu	m0r8_RC0s	moo	normal		16
Running	36833.abem5.ncsa.uiu	Estr_Prod	seabra	normal		16
Running	36842.abem5.ncsa.uiu	Script.abe	dcollins	normal		16
Running	36900.abem5.ncsa.uiu	L10_N64_4l	gbryan	normal		16
Running	36903.abem5.ncsa.uiu	frthall	zeticne	normal		150



Inca Capability Kit Testing

Inca Status Pages

http://sapa.sdsc.edu:8080/inca/jsp/status.jsp?suiteNames=remote-compute.teragrid.org-3.0.0&resource=select&xml=remote-cor

TeraGrid Inca Status Pages

CURRENT DATA | HISTORICAL DATA | REPORTS | INFO

Page loaded: 05-13-2008 02:38 PM (PDT)

down err	resource is down	pkgWait	test under development	tk-#	TeraGrid ticket number
n/a	does not apply to resource	incaWait	waiting for package delivery	timeOut	reporter timed out
	missing (not yet executed)	incaErr	waiting for inca test change	*	result is stale
pass	passed		inca framework error		
error	error				
noSoftenv	missing softenv				

KIT: RESOURCE:

Remote Compute 3.0.0

- [ctss-remote-compute-registration](#)
- [wsgram](#)
- [globus-wsrf](#)
- [prews-gram](#)
- [tg-audit-config](#)

	anl-grid	anl-ia64	anl-viz	indiana-bigred	ncar-frost	ncsa-cobalt	ncsa-grid-abe	ncsa-grid-hg	ncsa-tungsten	ornl-login	psc-bigben	psc-rachel	purdue-condor	purdue-grid	purdue-steele	sdsc-bg	sdsc-datastar	sdsc-ia64	tacc-lonestar	tacc-ranger	tacc-viz
ctss-remote-compute-registration	anl-grid	anl-ia64	anl-viz	indiana-bigred	ncar-frost	ncsa-cobalt	ncsa-grid-abe	ncsa-grid-hg	ncsa-tungsten	ornl-login	psc-bigben	psc-rachel	purdue-condor	purdue-grid	purdue-steele	sdsc-bg	sdsc-datastar	sdsc-ia64	tacc-lonestar	tacc-ranger	tacc-viz
* need test *	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
globus-wsrf	anl-grid	anl-ia64	anl-viz	indiana-bigred	ncar-frost	ncsa-cobalt	ncsa-grid-abe	ncsa-grid-hg	ncsa-tungsten	ornl-login	psc-bigben	psc-rachel	purdue-condor	purdue-grid	purdue-steele	sdsc-bg	sdsc-datastar	sdsc-ia64	tacc-lonestar	tacc-ranger	tacc-viz
wsrf-query-mds	n/a	pass	pass	pass	pass	pass	pass	pass	error	pass	pass	n/a	pass	pass	error	error	pass	pass	pass	pass	pass
wsrf-query-mds-auth-JobsRP	n/a	pass	pass	pass	pass	pass	error	error	pass	pass	pass	pass	error	error	pass	error	pass	pass	pass	pass	n/a
wsrf-query-mds-auth-LoadRP	n/a	pass	pass	pass	pass	pass	error	error	pass	pass	pass	pass	error	error	pass	error	pass	pass	pass	pass	n/a
prews-gram	anl-grid	anl-ia64	anl-viz	indiana-bigred	ncar-frost	ncsa-cobalt	ncsa-grid-abe	ncsa-grid-hg	ncsa-tungsten	ornl-login	psc-bigben	psc-rachel	purdue-condor	purdue-grid	purdue-steele	sdsc-bg	sdsc-datastar	sdsc-ia64	tacc-lonestar	tacc-ranger	tacc-viz
prews-gram-netstat	noSoftenv	n/a	n/a	pass	pass	pass	pass	pass	pass	pass	n/a	n/a	pass	pass	pass	pass	pass	pass	pass	pass	pass
prews-gram-gatekeeper-ping	incaErr	n/a	n/a	pass	pass	pass	pass	pass	pass	pass	n/a	n/a	pass	pass	pass	pass	pass	pass	pass	pass	pass
prews-gram-fork	incaErr	n/a	n/a	pass	pass	pass	pass	pass	pass	pass	n/a	n/a	pass	pass	pass	pass	pass	pass	pass	pass	pass
prews-gram-batch	n/a	pass	pass	pass	pass	pass	pass	pass	pass	pass	pass	pass	pass	n/a	pass	pass	pass	pass	pass	pass	pass
tg-audit-config	anl-grid	anl-ia64	anl-viz	indiana-bigred	ncar-frost	ncsa-cobalt	ncsa-grid-abe	ncsa-grid-hg	ncsa-tungsten	ornl-login	psc-bigben	psc-rachel	purdue-condor	purdue-grid	purdue-steele	sdsc-bg	sdsc-datastar	sdsc-ia64	tacc-lonestar	tacc-ranger	tacc-viz
* need test *	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a

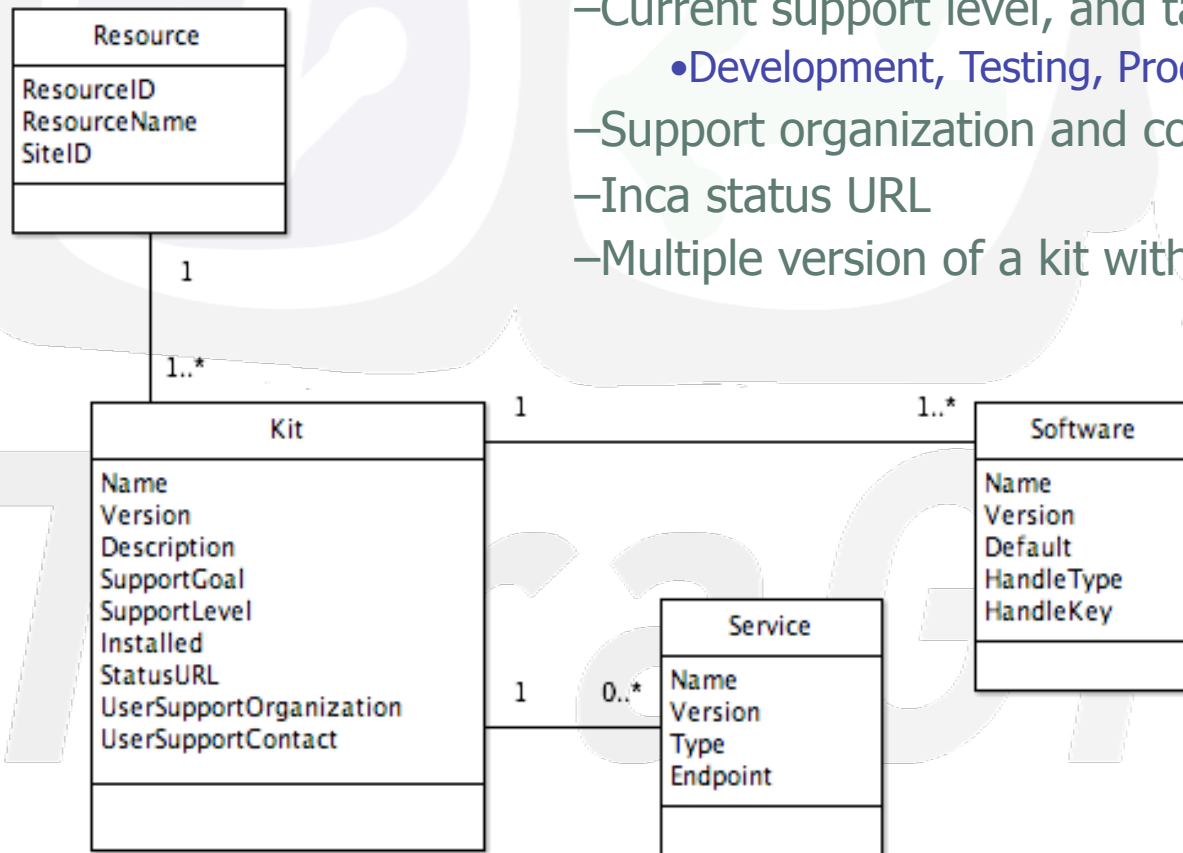
TM



CTSS 4 Capability Kits

For each capability kit on each resource

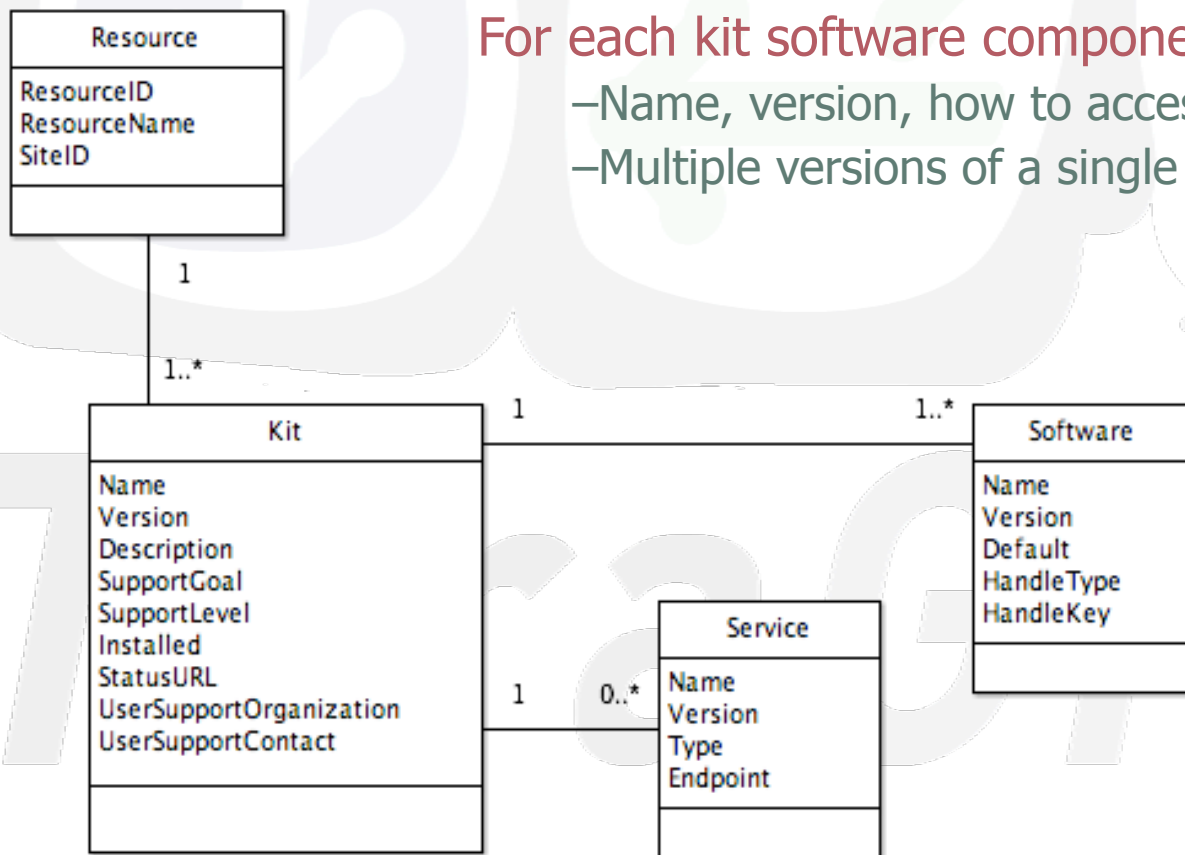
- Current support level, and target support level
 - Development, Testing, Production
- Support organization and contact
- Inca status URL
- Multiple version of a kit with different support levels



CTSS 4 Capability Kit Software

For each kit software component on each resource

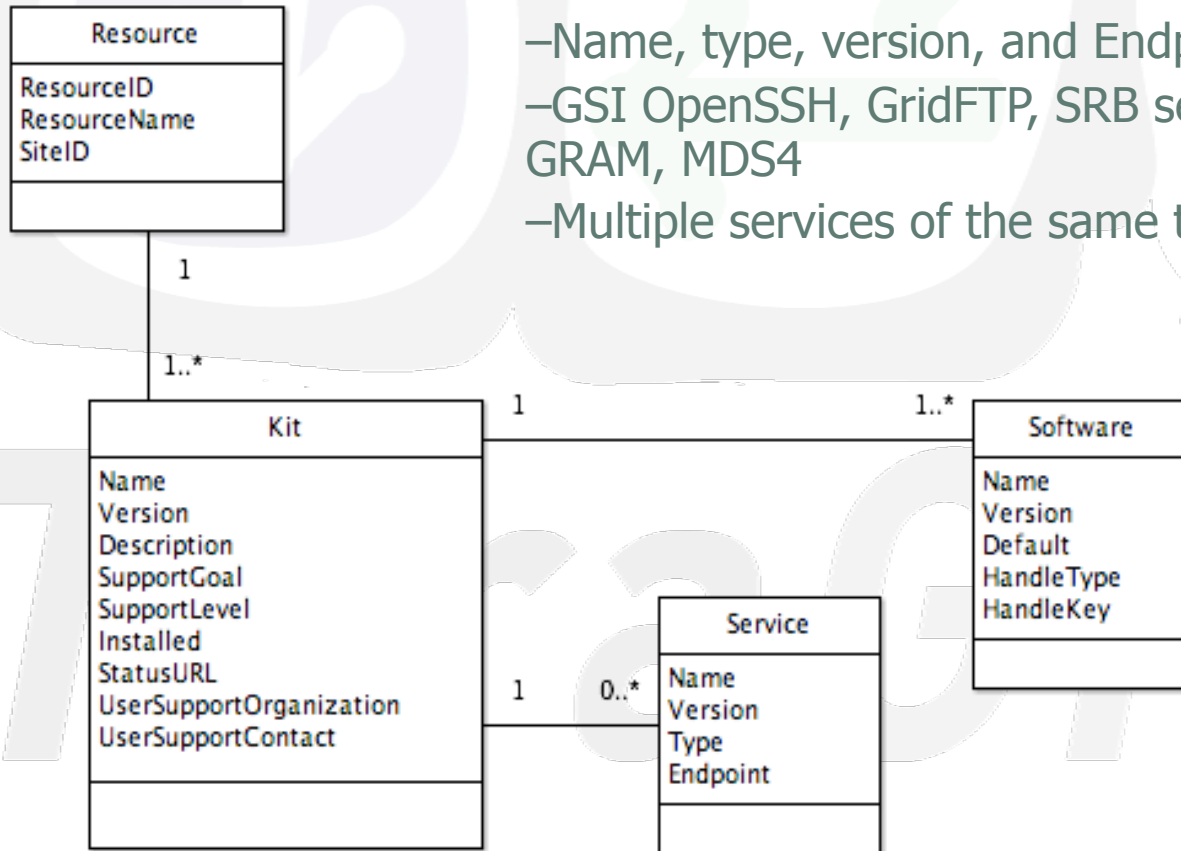
- Name, version, how to access it
- Multiple versions of a single component



CTSS 4 Capability Kit Services

For each kit service on each resource

- Name, type, version, and Endpoint (contact location)
- GSI OpenSSH, GridFTP, SRB servers, PreWS & WS
- GRAM, MDS4
- Multiple services of the same type



What's in Development?

Expanded content

- TeraGrid Gateways list
- Extended GridFTP service information (striping, bandwidth, etc)
- Local HPC Software
- (Meta)Scheduling support information
- Information Services Metadata

Enhanced caching, aggregation, and storage

Improved information access

- tginfo, universal command line query tool
- WS/REST, Web 2.0 style information access
- Multiple formats: CSV TEXT, XML, JSON, RSS/Atom, ...
- GLUE 2.0
- Usage metrics



The Team

Information Services design, development, and support team

Eric Blau
Jason Brechin
Lee Liming
JP Navarro
Laura Pearlman

Information Services downstream developers and consumers

End-to-end data transfer optimization team @ PSC (Derek Simmel, others)
INCA (Kate Ericson)
Operations (Jason Brechin, Tony Rimovsky)
User Documentation (Mike Dwyer, Diana Diehl)
User Portal (Maytal Dahan)
Other lurkers...

More Information

Information Services main page:

<http://info.teragrid.org/>

(links to content and documentation)

User Documentation (CTSS 4 kits, software, services):

<http://www.teragrid.org/userinfo/software/ctss.php>

User Portal (scheduler load & queue contents):

<https://portal.teragrid.org:443/gridsphere/gridsphere?cid=resources>

Inca Monitoring Framework

<http://www.teragrid.org/userinfo/software/ctss.php>

E-mail:

<navarro@mcs.anl.gov>