Urgent Computing, Sharing Grid Resources, and Elastic Computing

Pete Beckman
Argonne National Laboratory
University of Chicago

http://www.mcs.anl.gov/~beckman


VISSIM 3.60 - ..._vision_data\munic_abds_stadium\vissim\testphase2\samstag.inp
File Edit Network Editor Signal Control Options Simulation Animation Test
A92/E53
A99
New exit
Stadium
A9
Legend
Density [veh/km/l]
Speed[km/h]
Macroscopic value per lane segment
1.00 10.00
2.00 25.00
4.00 40.00
8.00 50.00
16.00 60.00
32.00 70.00
64.00 80.00
96.00 90.00
128.00 100.00
999.00 200.00
www.ptv.de
4464496.0:5343419.6 17:48:39.5 39 2071 16106

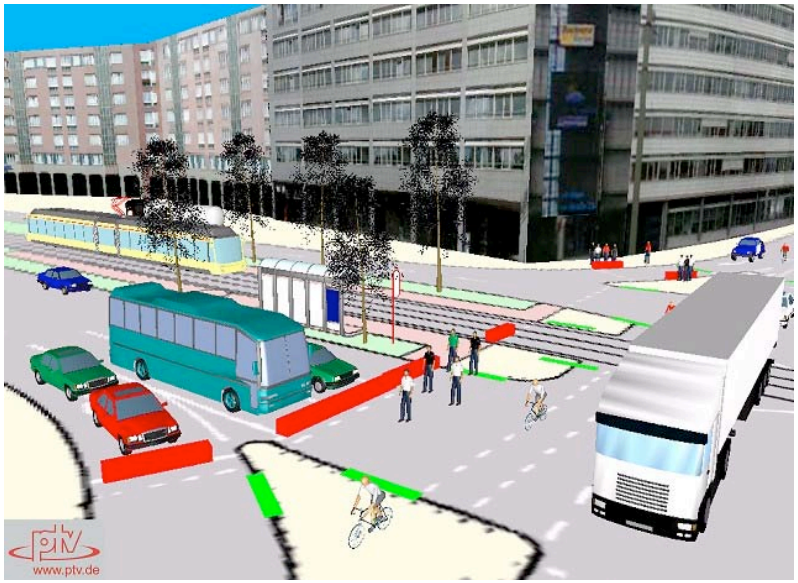
www.ptv.de

Urgent Computing: I Need it Now!

- Applications with dynamic data and ***result deadlines*** are being deployed
- Late results are useless
  - Wildfire path prediction
  - Storm/Flood prediction
  - Influenza modeling
- Some jobs need priority access "**Right-of-Way Token**"

How can we get cycles?

- Build supercomputers for the app
  - *Pros*: Resource is ALWAYS available
  - *Cons*: Incredibly costly (99% idle)
  - *Example*: Coast Guard rescue boats
- Share public infrastructure
  - *Pros*: low cost
  - *Cons*: Requires complex system for authorization, resource mgmt, and control
  - *Examples*: school buses for evacuation, cruise ships for temporary housing

EBS AUTHENTICATOR LIST D

OCTOBER 1973

ACTIVATION	TERMINATION
1 GLORY	1 AFTERPIECE
2 CHINAMAN	2 ORGY
3 FANFOLD	3 HEMPSEED
4 RATIONAL	4 FORTRESS

Introducing SPRUCE

- The Vision:
  - Build cohesive infrastructure that can provide urgent computing cycles
- Technical Challenges:
  - Provide high degree of reliability
  - Elevated priority mechanisms
  - Resource selection, data movement
- Social Challenges:
  - Who? When? What?
  - How will emergency use impact regular use?
  - Decision-making, workflow, and interpretation

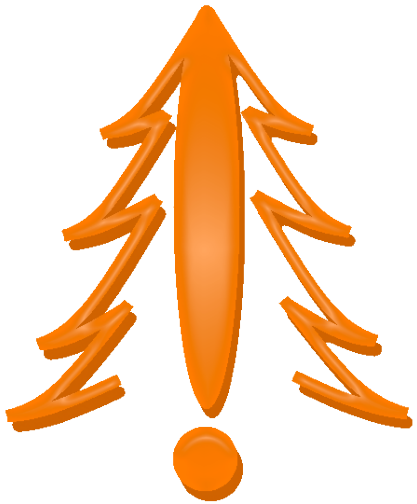

Existing "Digital Right-of-Way" Emergency Phone System

Calling cards are in widespread use and easily understood by the NS/EP User, simplifying GETS usage

GETS priority is invoked "call-by-call"

GETS is a "ubiquitous" service in the Public Switched Telephone Network…if you can get a DIAL TONE, you can make a GETS call

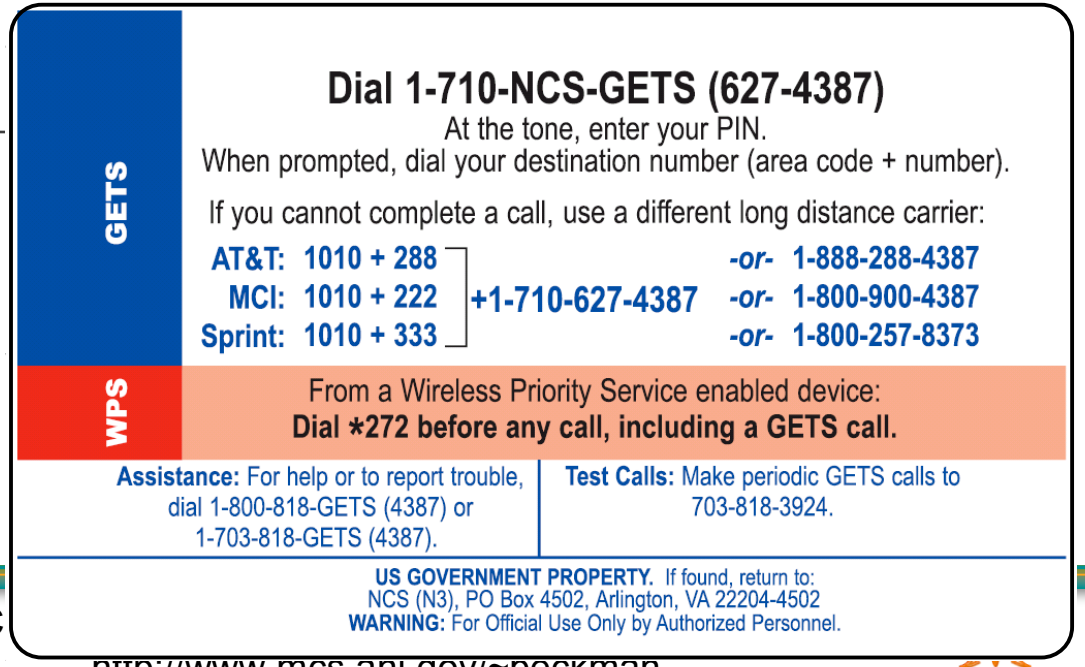

SPRUCE Architecture Overview (1/3)
Right-of-Way Tokens

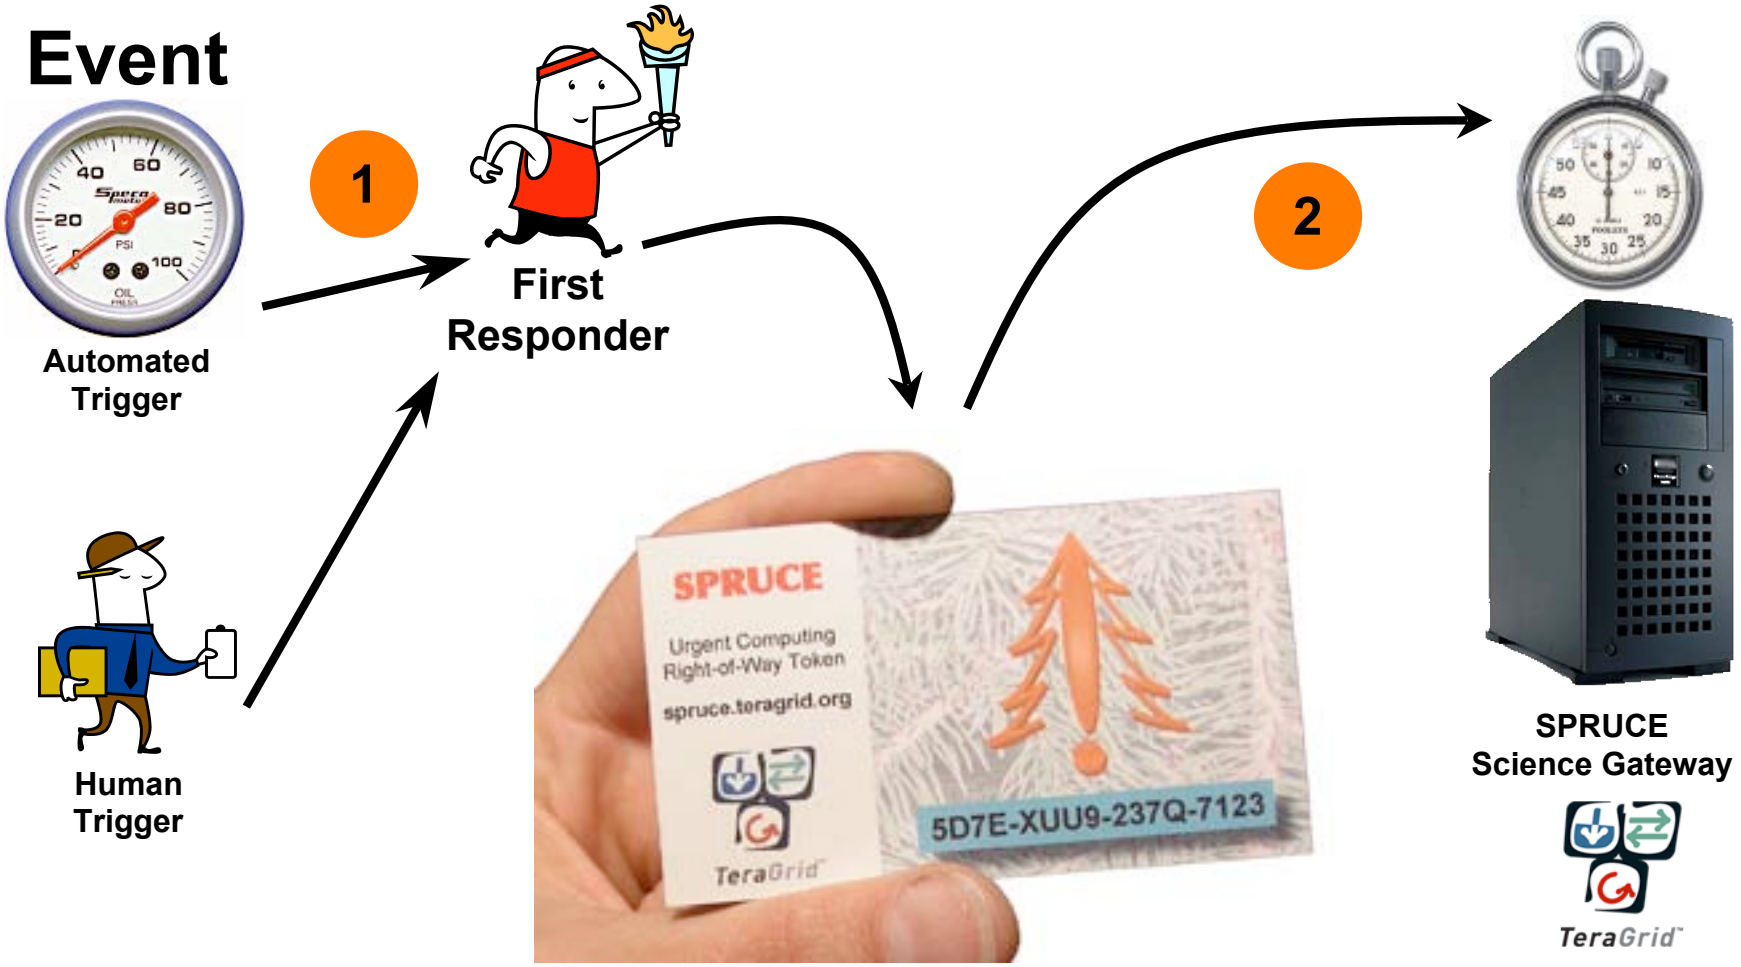

SPRUCE Architecture Overview (2/3)
Submitting Urgent Jobs

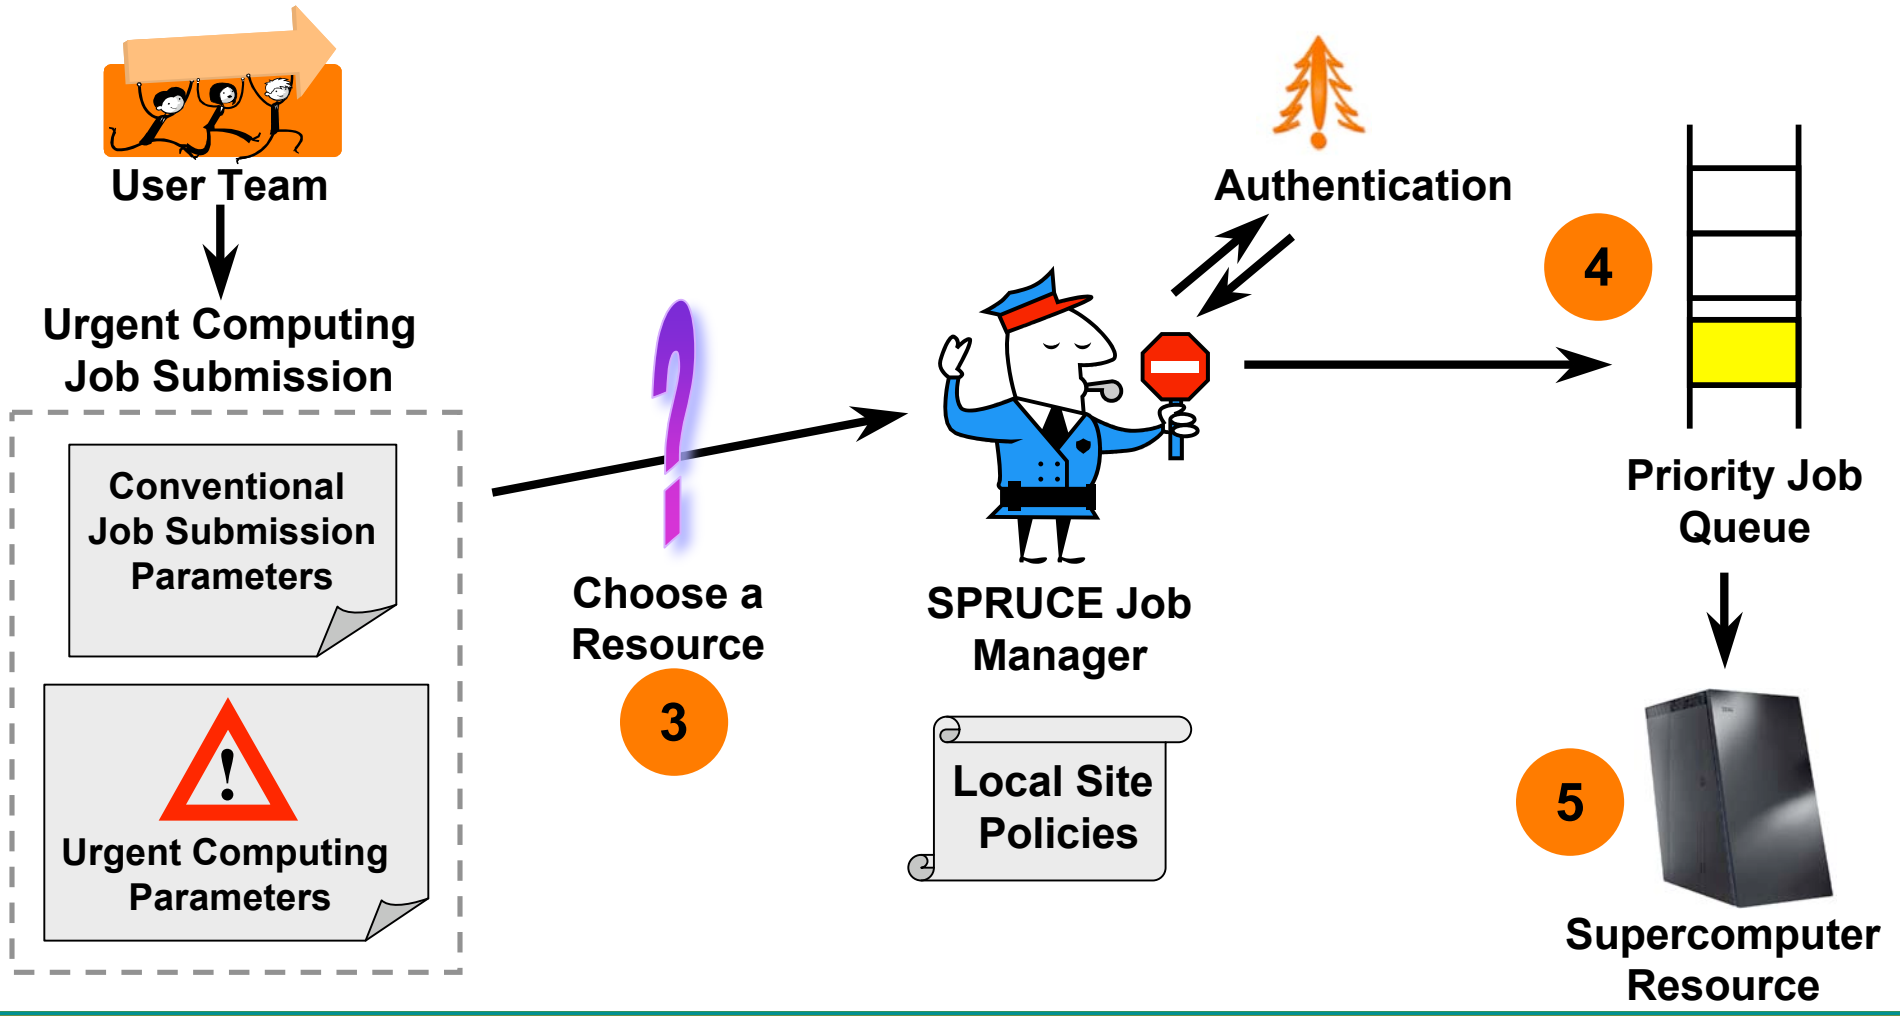

SPRUCE Architecture Overview (3/3)
Analyzing Urgent Jobs

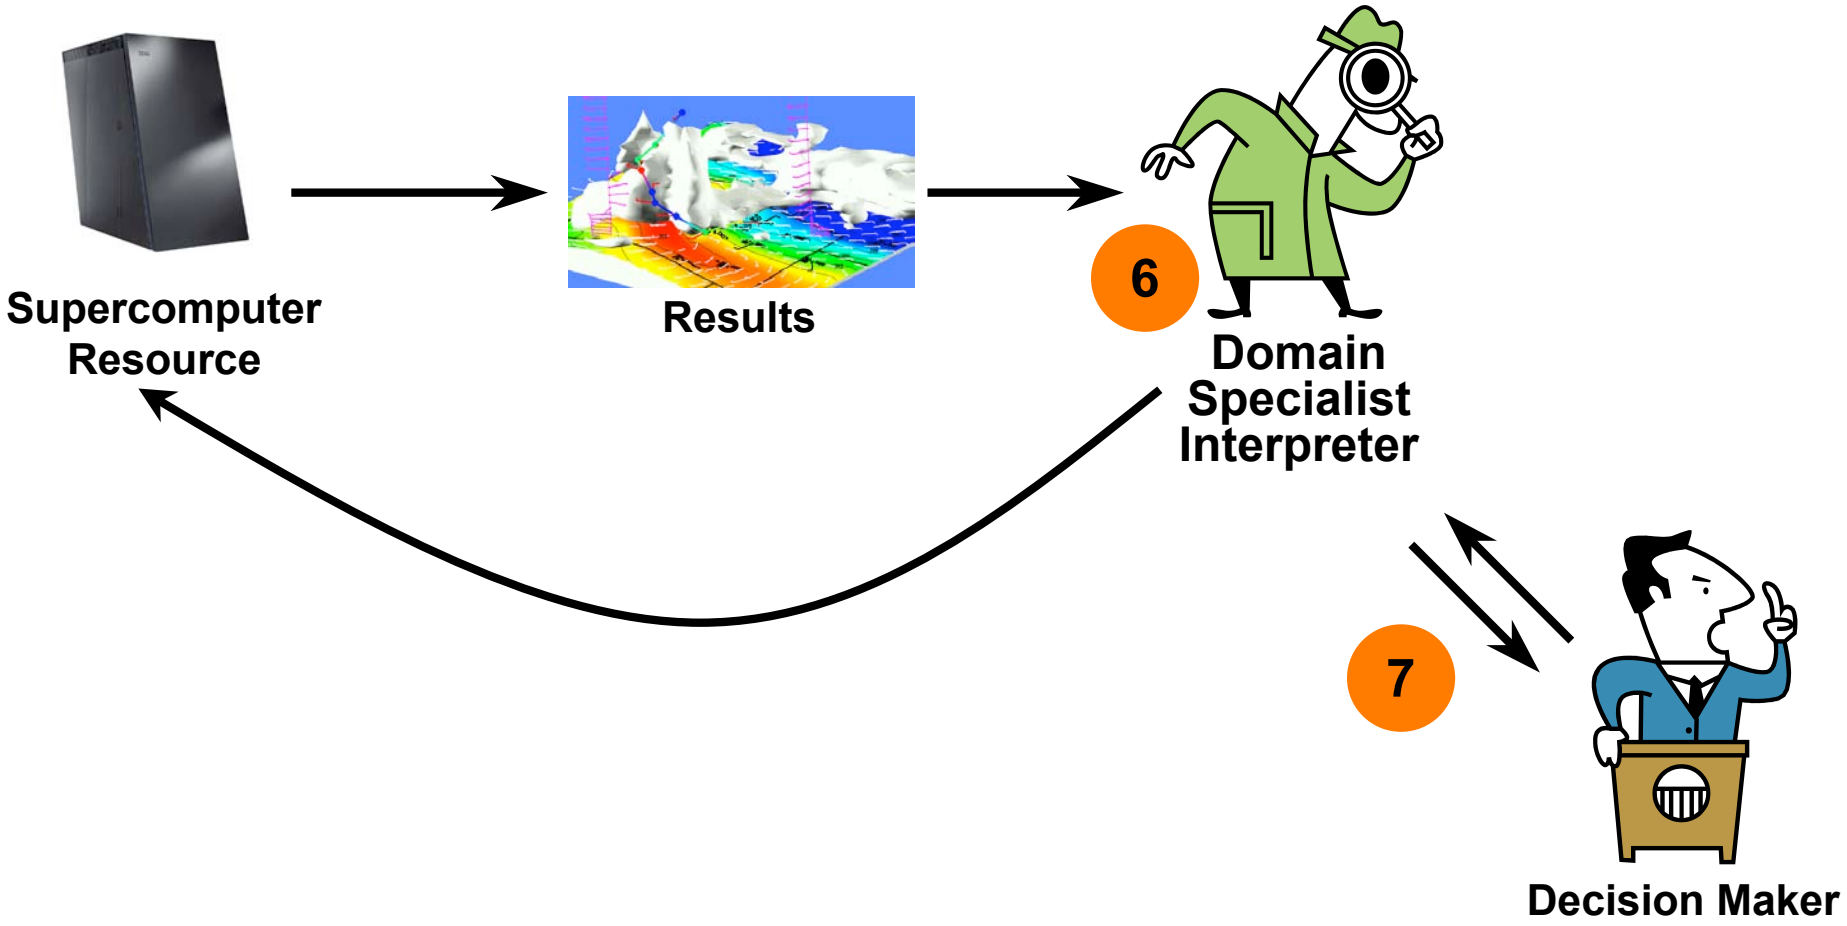

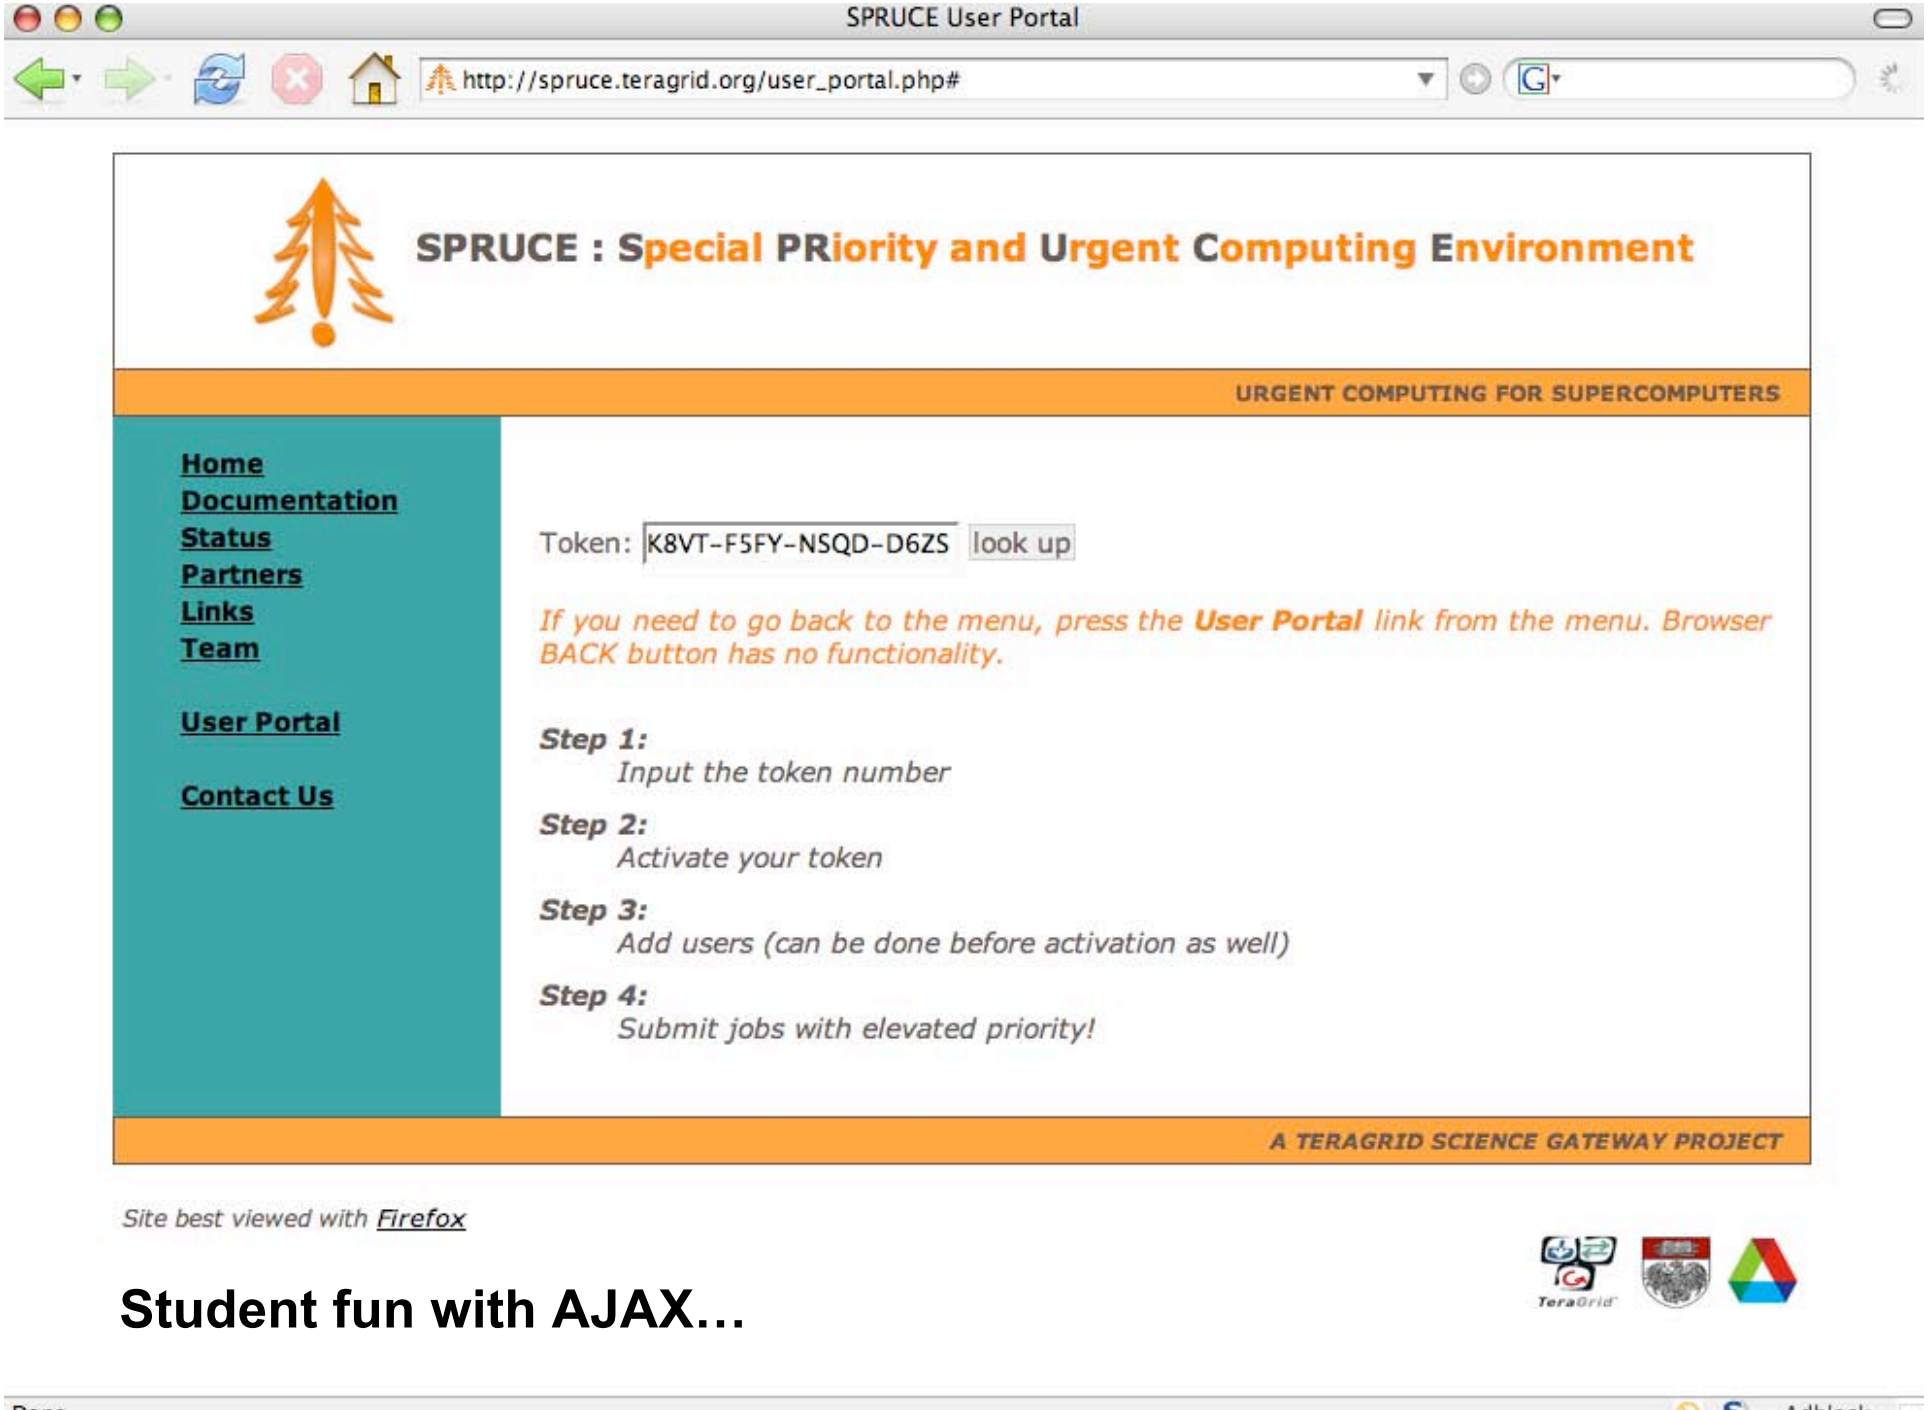

Student fun with AJAX…

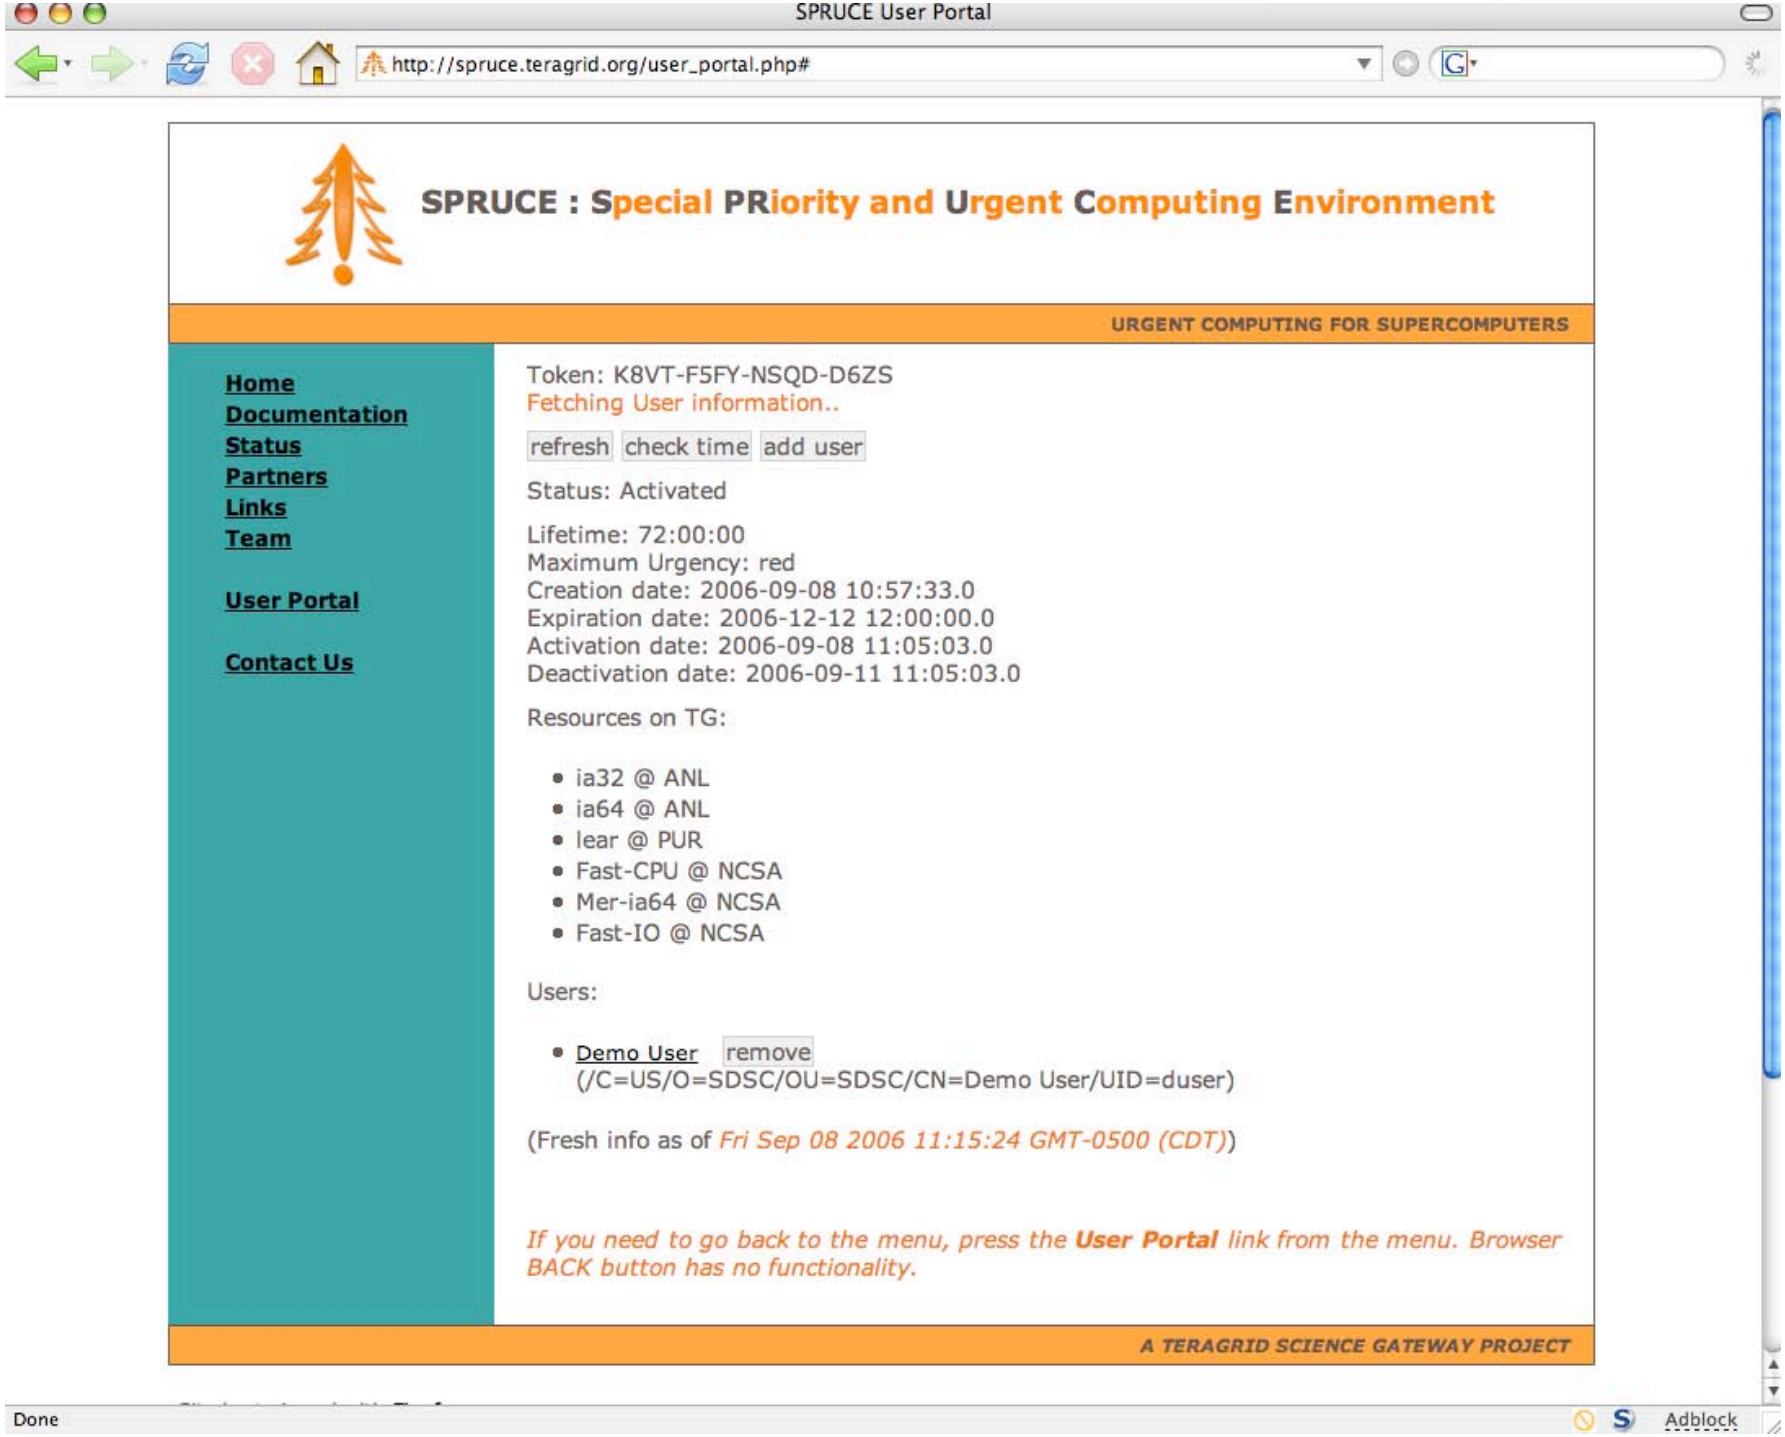
SPRUCE User Portal
http://spruce.teragrid.org/user_portal.php#
SPRUCE : Special PRiority and Urgent Computing Environment
URGENT COMPUTING FOR SUPERCOMPUTERS
Home
Documentation
Status
Partners
Links
Team
User Portal
Contact Us
Token: K8VT-F5FY-NSQD-D6ZS
Fetching User information..
refresh check time add user
Status: Activated
Lifetime: 72:00:00
Maximum Urgency: red
Creation date: 2006-09-08 10:57:33.0
Expiration date: 2006-12-12 12:00:00.0
Activation date: 2006-09-08 11:05:03.0
Deactivation date: 2006-09-11 11:05:03.0
Resources on TG:
ia32 @ ANL
ia64 @ ANL
lear @ PUR
Fast-CPU @ NCSA
Mer-ia64 @ NCSA
Fast-IO @ NCSA
Users:
Demo User remove
(/C=US/O=SDSC/OU=SDSC/CN=Demo User/UID=duser)
(Fresh info as of Fri Sep 08 2006 11:15:24 GMT-0500 (CDT))
If you need to go back to the menu, press the User Portal link from the menu. Browser BACK button has no functionality.
A TERAGRID SCIENCE GATEWAY PROJECT
Done
Adblock

Site-Local Response Policies: How will Urgent Computing be treated?

- "Next-to-run" status for priority queue; wait for running jobs to complete
- Force checkpoint of existing jobs; run urgent job
- Suspend current job in memory (kill -STOP); run urgent job
- Kill all jobs immediately; run urgent job

- Provide differentiated CPU accounting
  - "jobs that can be killed because they maintain their own checkpoints will be charged 20% less"
- Other incentives

Emergency Preparedness Testing: "Warm Standby"

- In urgent computing situation, there is no time to port applications
  - Applications must be in "warm standby"
  - Verification and validation runs test readiness periodically (Inca)
  - Only verified apps participate in urgent computing
- Grid-wide Information Catalog
  - Application was last tested & validated on <date>
  - Also provides key success/failure history logs

Choosing a Resource

An Advisor

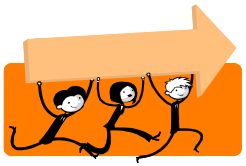

User Team

Urgent Computation Request

Deadline
Urgency Level

Live Job/Queue Data

Platform	Next Available Job (Policy Based)	…
NCSA::Cobalt	Immediate	…
SDSC::Datastar	(5.3 hrs, 1024 nodes)	
PSC::Rachel	Immediate	…

the globus toolkit

MDS4 Service

Site Policies

Platform	Policy	…
NCSA::Cobalt	Human-in-the-loop, immediate access, kill existing jobs, 15 min. turnaround	…
SDSC::Elimidata	Automated, next job	…
SDSC::Datastar	Normal priority, no SPRUCE support	
PSC::Rachel	Automated, immediate access, kill existing jobs, 10 min turnaround	…

SPRUCE Data

Warm Standby Validation History

Platform	App.	Validated	Reliability	…
NCSA::Cobalt	Tornado	8 days ago	95%	…
NCSA::Cobalt	City Airflow	14 days ago	98%	…
SDSC::Elimidata	City Airflow	45 days ago	78%	…
SDSC::Elimidata	Influenza	30 days ago	59%	…

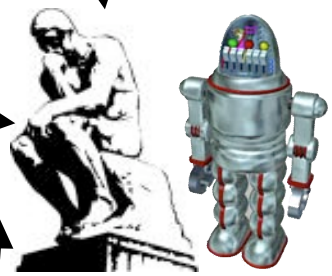

Advisor

Best HPC Resource

Deployment Status

- Deployed and available:
  - UC/ANL
  - Purdue
  - TACC
  - SDSC
- Very close:
  - Indiana
  - LSU
- Ready to integrate LEAD into SPRUCE
  - First user-customer
  - Warm standby apps

What About "Capacity" Computing?

- SPRUCE works well with "capability" computing:
  - Interface to small set of large resources
- Imagine a larger set of smaller resources?
  - Condor management?
  - Real on-demand servers?
- Amazon S3 & EC2

Amazon S3 & EC2
It's a Web Services World

- S3: Simple Storage Service
  - Cost: $0.20/GB transfer, $.15/GB-month
- EC2: Elastic Compute Cloud
  - Cost: $0.10/cpu-hr, $0.20/GB transfer
  - **No cost for internal bandwidth**
- Cost is extraordinarily good
- Commoditization is good!!
- The the real keys are reliability and dynamic behavior

amazon web services™

Account Activity

Learn About Amazon Web Services

AWS Home
Why Use AWS?
What's New in AWS?
Upcoming Events
Success Stories
Solutions Catalog
Create an Account
FAQs

Browse Web Services

Amazon E-Commerce Service
Amazon Elastic Compute Cloud (Beta)
Amazon Historical Pricing
Amazon Mechanical Turk (Beta)
Amazon Simple Storage Service
Amazon Simple Queue Service
Alexa Web Services

Browse All Web Services

Developer Connection

Resource Center
Forums
Blog
Newsletter

Account Number: [redacted]

Summary of This Month's Activity as of August 30, 2006

Billing Cycle for this Report: August 1 - August 31, 2006

Usage Charges	Rate	Usage	Totals
Amazon Simple Storage Service			
View/Edit Service	$0.20 per GB of data transferred (details)	0.007 GB	0.01
	$0.15 per GB-Month of storage used (details)	0.000 GB-Mo	0.01
			0.02
		Usage Report	
Amazon Elastic Compute Cloud			
View/Edit Service	$0.10 per instance hour consumed (or part of an hour consumed)	1 BoxUsage	0.10
	$0.20 per GB of data transferred outside of Amazon (i.e, Internet traffic)	8.61212611198425e-05 WANUsage	0.01
			0.11
		Usage Report	
Subtotal			**$ 0.13**
Taxes			
Estimated Taxes due on September 1, 2006			**$ 0.00**
Charges due on September 1, 2006*			**$ 0.13**

* All charges for this billing cycle will be charged to your credit card on your next billing date, September 1, 2006. These charges include 1) next billing cycle's subscription charges due on the next billing date and 2) usage charges from the current billing cycle. Not included in the charges displayed here are any additional usage charges you will accrue this billing cycle. The current billing cycle ends August 31, 2006 GMT. Visit the Amazon Web Services FAQs to learn more about web services pricing models and billing.

All web services are sold by Amazon Digital Services, Inc.

Imagine…

- Other companies catching up…
- Commoditization (like web email)
- A standardized interface to web-service "request vm"
- Dynamic capacity provides availability of 250K "node instances"
- urgent computing resources available immediately
- Missing bisection bandwidth, but great for capacity computing

The Future

- Web services interfaces to all the portal functions
- Extended submission schema
- Flexible tokens – aggregation, extension
- Encode local site policies
- Warm standby integration
- Automated ‘advisor’
- Data movement
- Redundancy to avoid downtime of portal

Questions?
Ready to Join?

spruce@ci.uchicago.edu
beckman@mcs.anl.gov

http://spruce.teragrid.org