Building the TeraGrid

Applied Grid Infrastructure

Charlie Catlett

Pete Beckman

Italy Grid School 2003, July 23-24

Argonne National Laboratory
University of Chicago

Acknowledgments: The TeraGrid is Big… Really Big…

- Director: Rick Stevens, Argonne Nat'l Lab/Univ Chicago
- Executive Director: Charlie Catlett, Argonne Nat'l Lab/Univ Chicago
- Director of Eng.: Pete Beckman, Argonne Nat'l Lab/Univ Chicago
- Chief Architect: Dan Reed, National Center for Supercomputing Applications (NCSA)
- Exec Committee Chair: Fran Berman, San Diego Supercomputer Center (SDSC)
- Jim Pool, Caltech
- Ian Foster, Argonne Nat'l Lab / Univ Chicago
- Mike Levine, Pittsburgh Supercomputer Center (PSC)
- Ralph Roskies, Pittsburgh Supercomputer Center (PSC)
- Site leads: JP Navarro (ANL), Mark Bartelt (CIT), J. Ray Scott (PSC), Rob Pennington (NCSA), Phil Andrews (SDSC)
- + 40 or 50 more….

Goals For The 4 Hour Presentation

- Fun
  - Our presentation will be the best
- Practical
  - No over-hyped claims (starting now)
  - No vaporware or claims how Grid computing will reinvent the Internet
- Comprehensive
  - The practical design, construction and operation of large-scale Grids for scientific computation
  - Not about technology (which service to run on what port, or why to choose Myrinet over Quadrics)

Outline of Presentation

Day 1:		Min:
Chapter 1:	Introduction	30
Chapter 2:	Creating a blueprint	60
Chapter 3a:	Design examples (part 1)	30
Day 2:		
Chapter 3b:	Design examples (part 2)	60
Chapter 4:	Construction phase	20
Chapter 5:	Operating the grid	20
Chapter 6:	Summary / questions	20

Chapter 1
Introduction

30 Minutes

Yes, Grid Hype has made Everyone's Job Harder

- **NYT**: "Grid computing, a concept that originated in supercomputing centers, is taking a step toward the mainstream: Sony will announce today that it will use the technology to accelerate its push into the emerging market for online games with thousands of players at a time."
- **Financial Times**: "Grid computing involves yoking together many cheap low-power computers to create a system with the high processing power typical of a large supercomputer, at a fraction of the price."

Hype Hype Hype

- **Technology Review**: "The coming explosion of activity could create a world of interlinked computer Grids – a development dwarfing the Internet boom of the 1990s."
- **Dan Farber@ZDnet**: "On my hype meter, grid computing receives a rating of 6.5 on a scale of 10"
- **Gartner Report**: "Biometrics, Web services and grid computing are the most hyped technologies"

Forget The Hype: Applied Grid Infrastructure

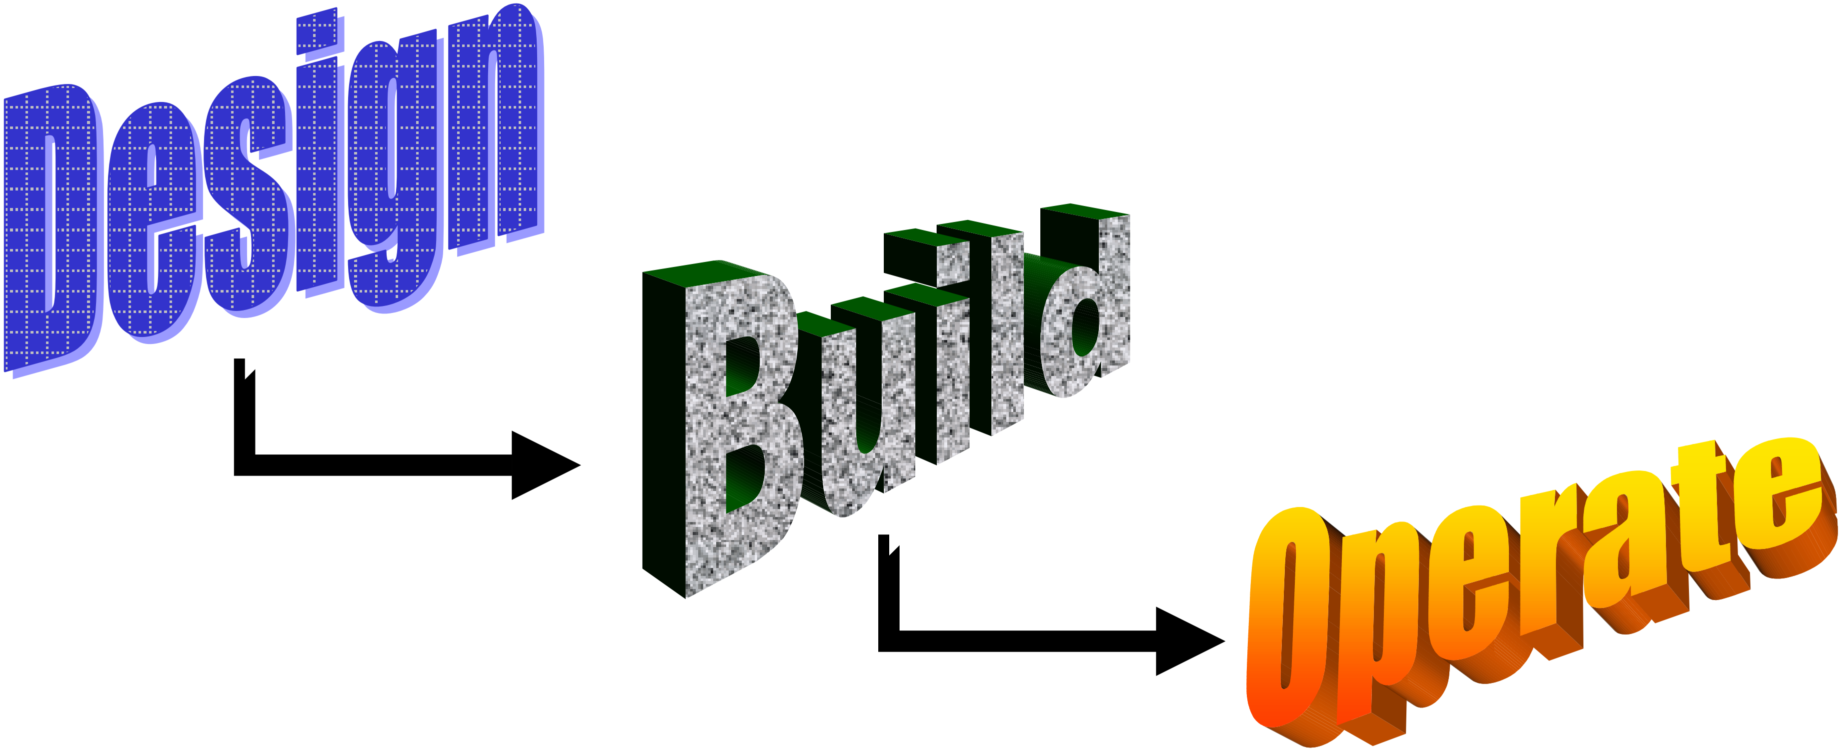

- Architects, Builders, Operators
  - Passing from role to role over time

The TeraGrid in a Nutshell (1/2)

- Transform US supercomputing R&D centers
  - Prototype NSF vision for Cyberinfrastructure
- Fun Science: Invent and Discover
  - Design new “Grid Hosting” SW infrastructure
    - Are you building a instance, or a hosting environment?
  - Design new optical network as national backplane for Grid computing
  - Create new processes and techniques for construction, management, and operations
  - Integrate new hardware computing platforms
  - Enable new Grid applications (100s)

Applications: The Real Reason for Cyberinfrastructure & the TeraGrid

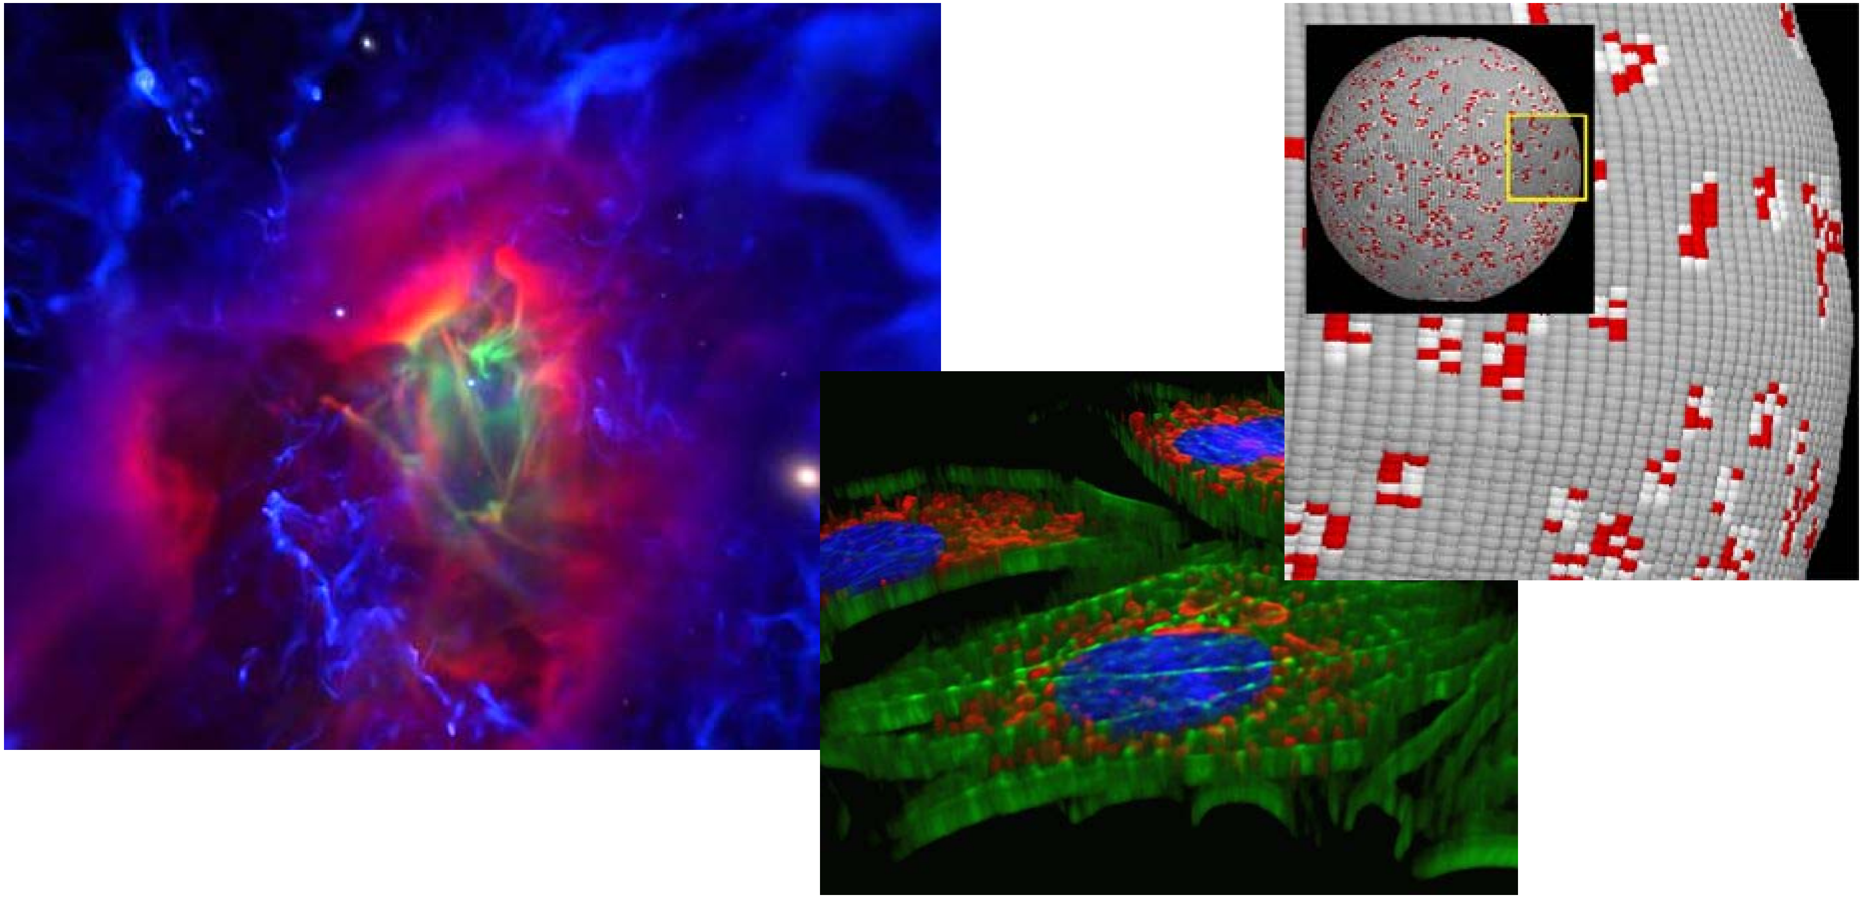

The TeraGrid in a Nutshell (2/2)

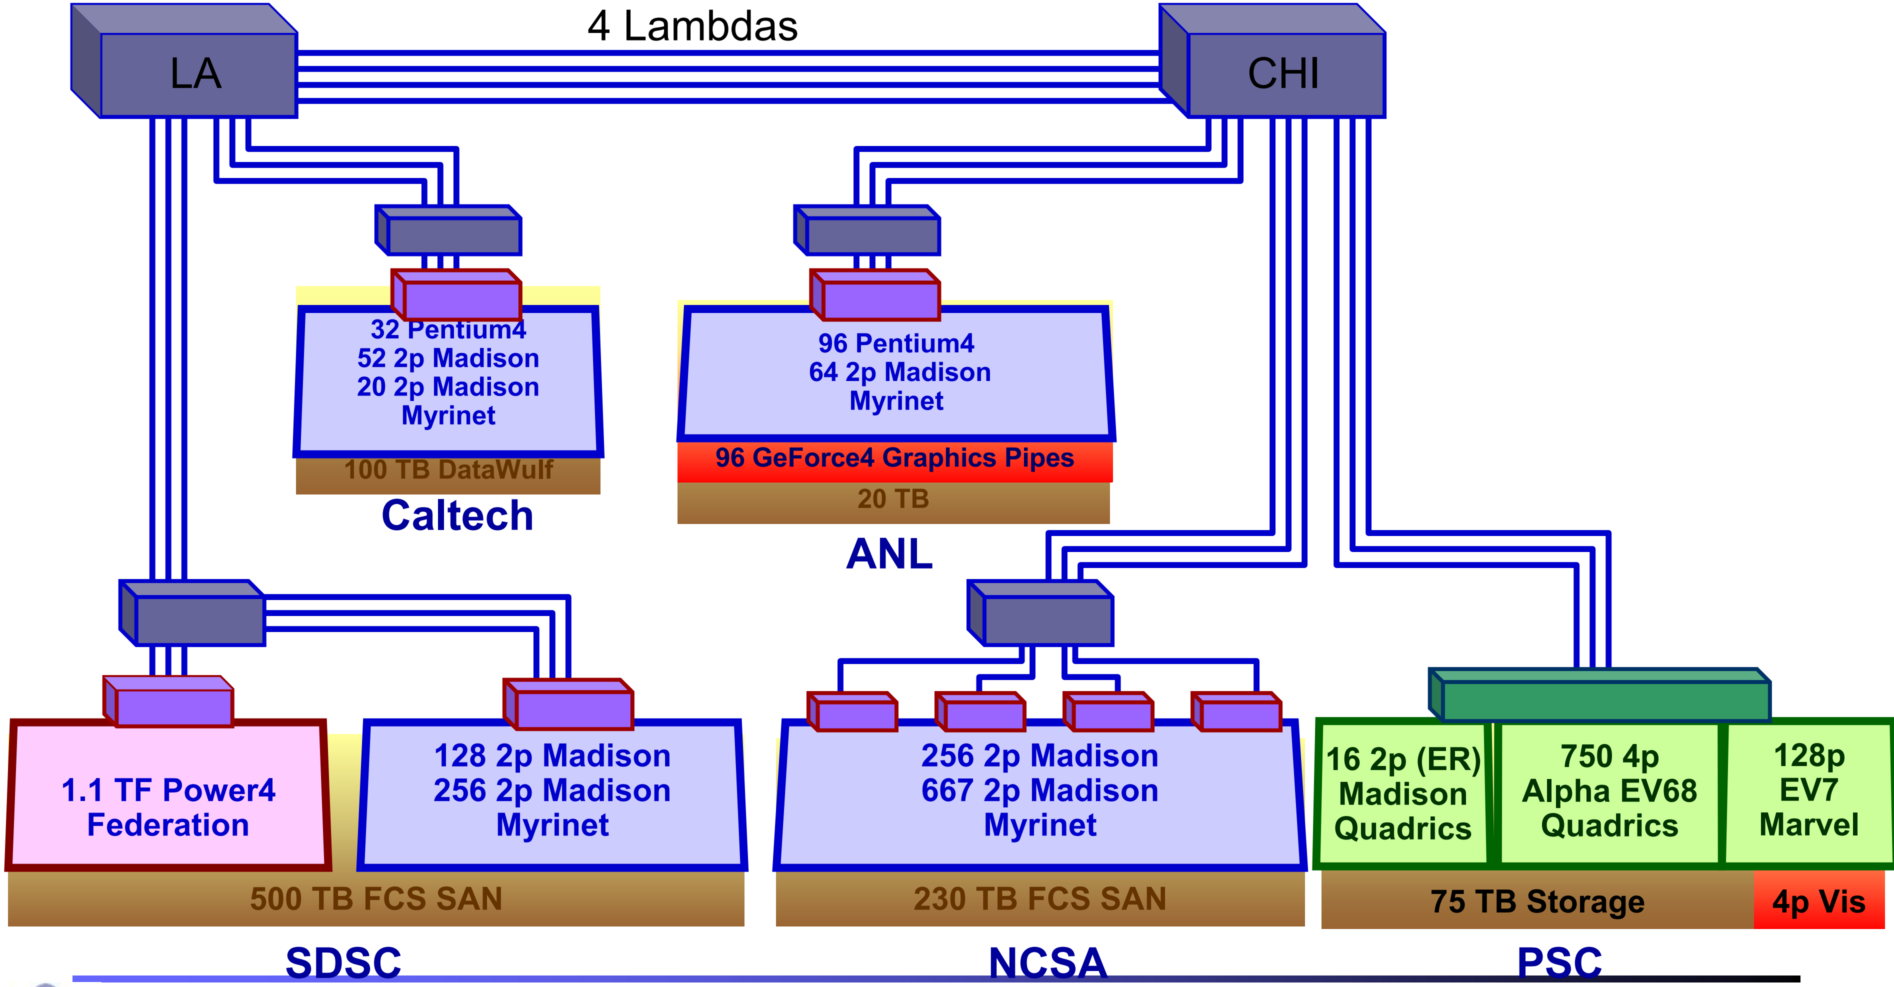

Ballpark Prices (actual values are NDA)

Software Team	$4 Million
Operations	$6 Million
Routers	$8 Million
Optical Networks	$10 Million
Hardware	$60 Million
Class with Pete & Charlie	Priceless

History of the TeraGrid

How to reproduce our failures or our successes

- Proposal Submitted April 2001, Funded October 2001
  - (artist’s conception)
- Fall 2001
  - “All-Hands” meeting
  - WG assignment: develop action plans, identify resource needs
  - Principles: exploit homogeneity wherever possible
- Early 2002
  - WG’s floundering with only an architecture, no blueprint
  - Address “Centers” versus “Grid” approach and “Scalability”
  - Applications Survey
- Spring 2002
  - BLUEPRINT in April 2002
  - Applications/Operations Workshop in April 2002

History of the TeraGrid

How to reproduce our failures or our successes

- Summer 2002
  - Straw Man User Guide
  - “Services” approach, initial service definitions
  - Initial network, initial test grid
  - Add 5th site
- Fall 2002
  - TeraGrid Primer
  - Project Plan Workshop
  - Final hardware configuration & ordering
  - Network operational
- Early 2003
  - Initial Software Stack (service) definition
  - Detailed Project Plan
  - Deployment of initial hardware

History of the TeraGrid

How to reproduce our failures or our successes

- Spring 2003
  - Comprehensive “work breakdown structure” plan
  - Selection of “Flagship” applications
  - Initial deployment of TeraGrid hardware + software
- Today
  - Focus on training, allocations, user services roll out
  - “Hardening” software stack
  - Completing distributed accounting systems
  - Completing test infrastructure
- Tomorrow
  - Finish adding 5th site, beginning discussions with 6, 7, 8

Why the Scientific Community Cares About the TeraGrid

- Science is bigger
  - Fewer garage labs, more centralized global resources
- More teams
  - Large, geographically distributed research teams
- More diverse interactions
  - Global collaborations spanning organizations

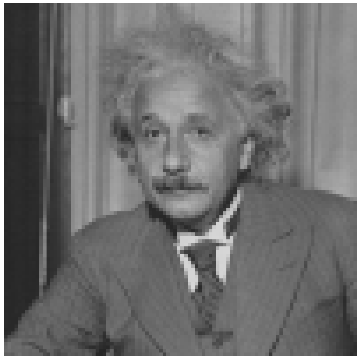

- **Result: More Sharing:**
  - Computers, Instruments, Databases, Storage
  - Today's scientists share more than arcane journal articles

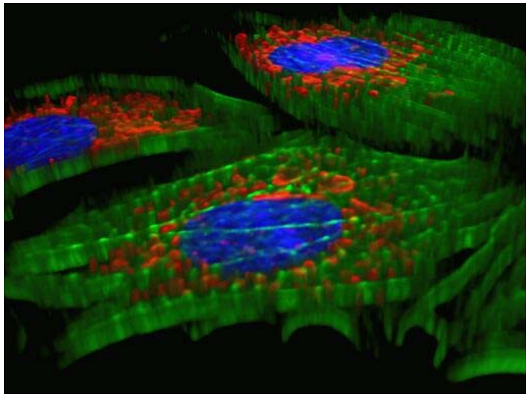

Why Building a Production Grid is Hard

- We perpetuate a great big fat lie:
  - Technology solves problems
    - “Grid Technology builds collaboration…”
- Truth: Grid technology is raw, it only enables new procedures and operational methods
  - Installing “Globus technology” does not create a Grid, just like installing Apache does not create an e-commerce site
- An enormous amount of architecture, design, and process management is required to convert technology into a solution for users

For The Consumer: It Must Be Easy^3

- Even the smallest barrier to sharing prevents collaboration
- "Technologists call it revolutionary… users call it unusable"
- Benefits to collaboration less tangible
- Examples:
  - Dialing a few extra telephone digits
  - Sharing between laptops (using email)
  - Gopher / FTP / Archive verses the Web

“TeraGrid Roaming”

It must be easy, it must be easy, it must be easy
Nearly eliminate the barrier to entry

- Develop application at ANL, run at NCSA
- Run at CalTech with data from SDSC
- Run large job across all sites
- Unified accounting and billing
- Predictable levels of service

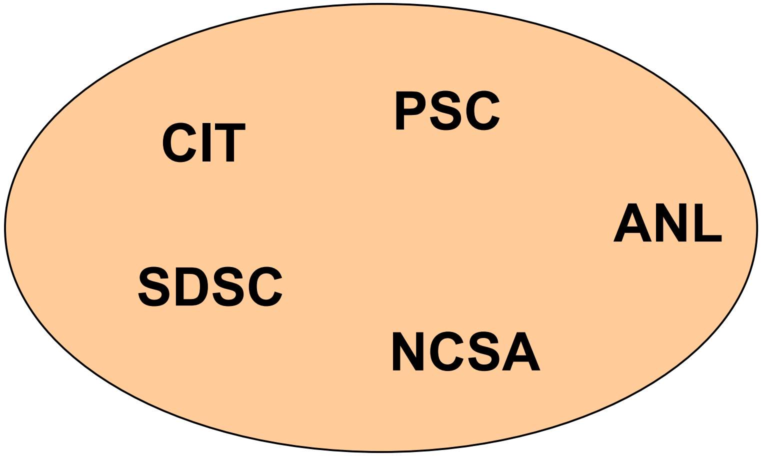

Launching a new service is hard

- **Enormous investment**
- **Ubiquitous**
- **Easy**
- **Paying Customers**

(remember Iridium?)

A Grid Hosting Environment

A SLA for a Virt Org that hosts other Virt Orgs and Grid Applications

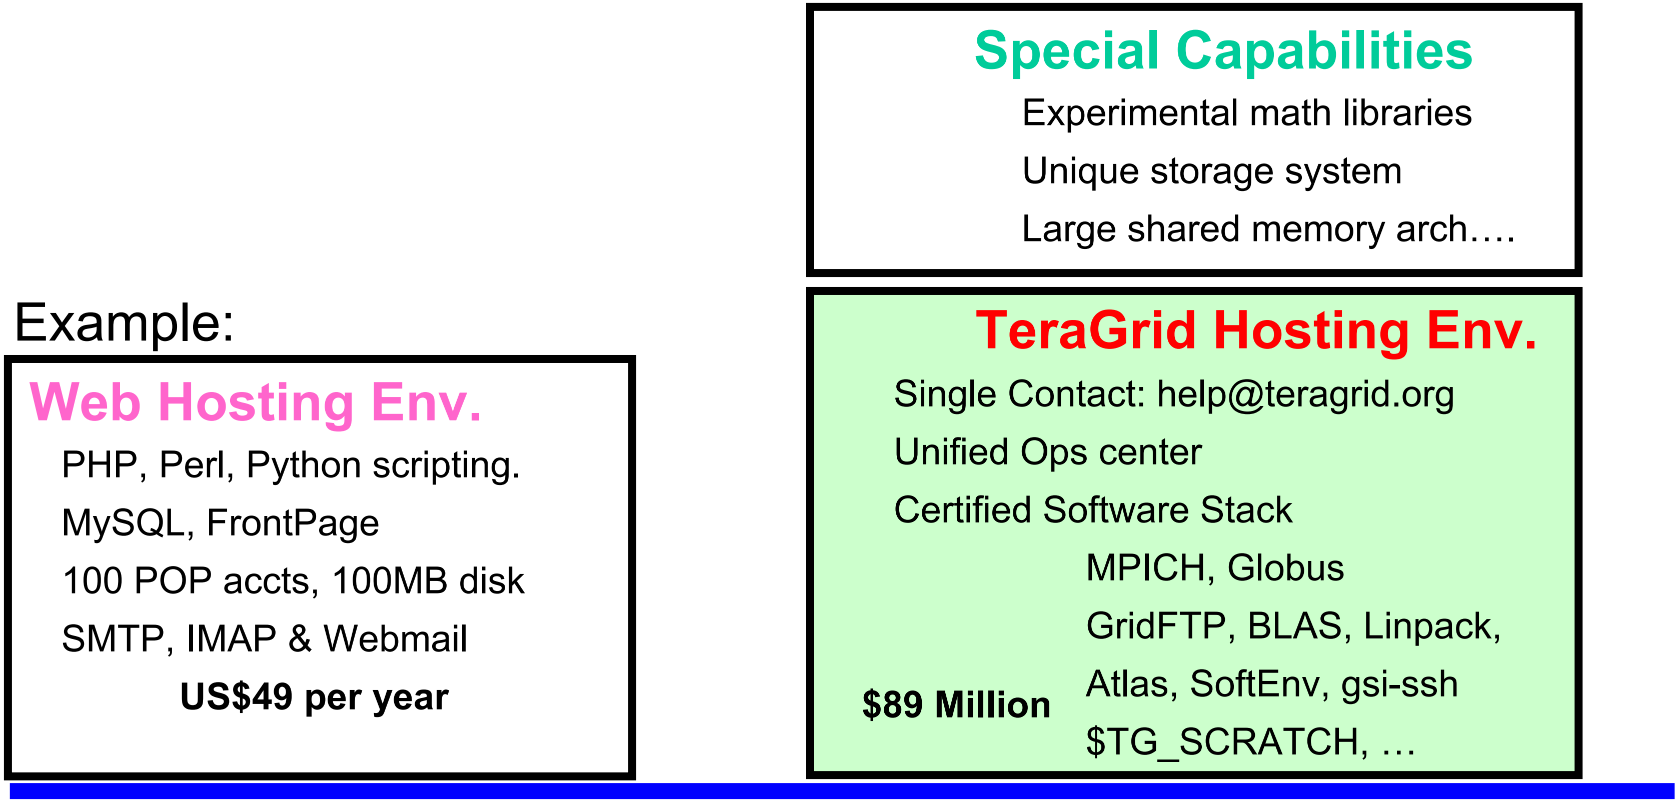

Special Capabilities

Experimental math libraries

Unique storage system

Large shared memory arch….

Example:

Web Hosting Env.

PHP, Perl, Python scripting.

MySQL, FrontPage

100 POP accts, 100MB disk

SMTP, IMAP & Webmail

US$49 per year

TeraGrid Hosting Env.

Single Contact: help@teragrid.org

Unified Ops center

Certified Software Stack

MPICH, Globus

GridFTP, BLAS, Linpack,

Atlas, SoftEnv, gsi-ssh

$TG_SCRATCH, …

$89 Million

Classic Unix-like Environment

- /bin/sh, bin/cp, /bin/ls, Unix file system & tools, dev tools (make, compilers) etc

Chapter 2
Creating a Blueprint

60 Minutes

The Art and Science of Design

- Who are the consumers?
  - HPC Users
  - Funding agencies
- What are their expectations?
  - Grid hype, overselling, “and then a miracle occurs”
  - The art of negotiation…
    - Not covered in this talk, but we can provide professional consulting services on an hourly basis
- What is your budget?
- What technologies will you choose?
  - Free software, $$$ SW/support, build it from scratch

Job #1:
Create Virtual Organization for Participants

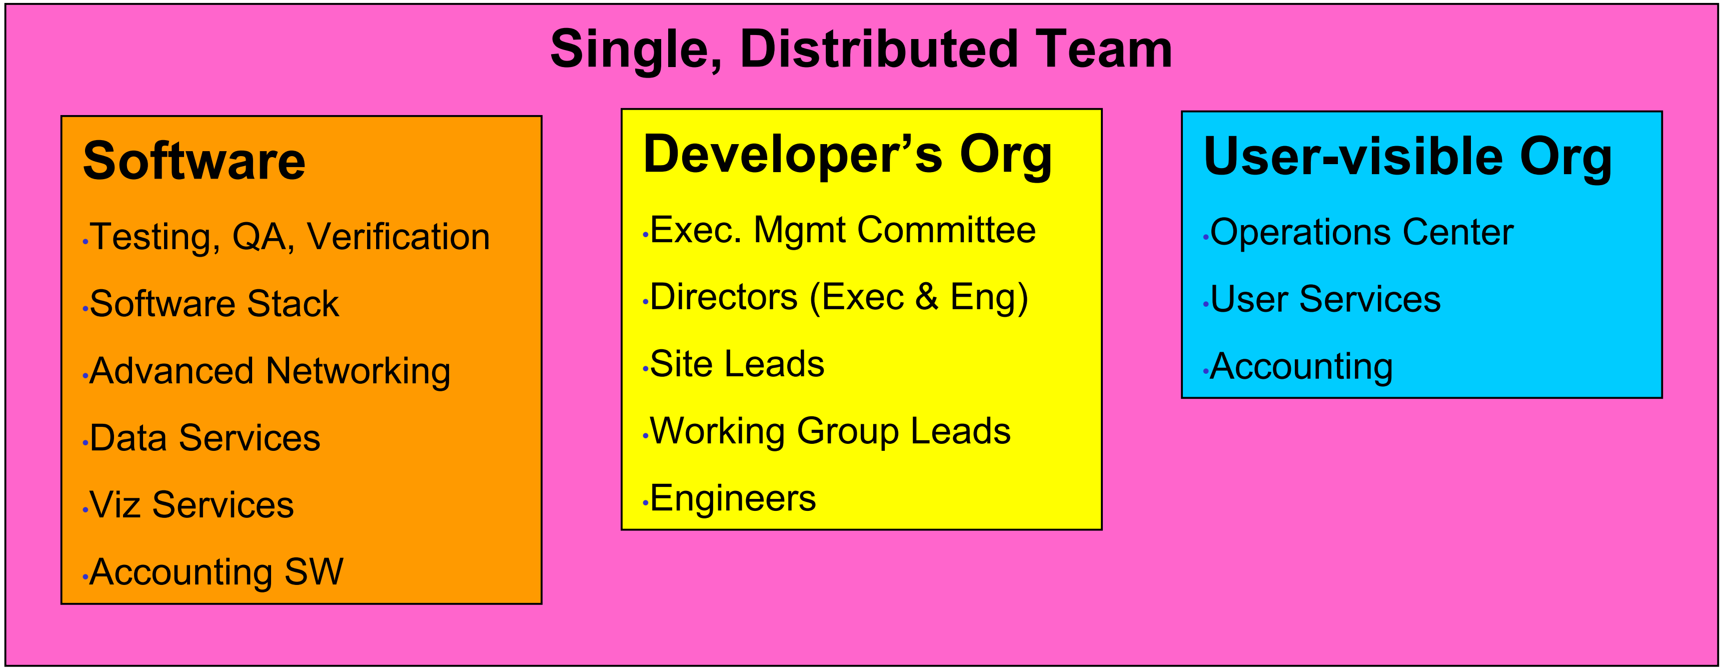

Single, Distributed Team

Software

- Testing, QA, Verification
- Software Stack
- Advanced Networking
- Data Services
- Viz Services
- Accounting SW

Developer's Org

- Exec. Mgmt Committee
- Directors (Exec & Eng)
- Site Leads
- Working Group Leads
- Engineers

User-visible Org

- Operations Center
- User Services
- Accounting

It is not your code, it is our code

It is not your doc, it is our doc

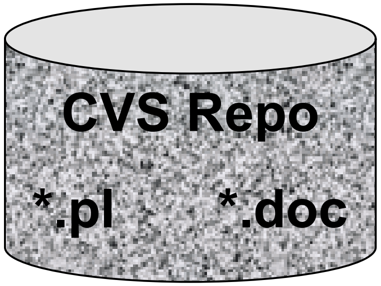

One repo to rule them all,
one repo to find them
one repo to bring them all
and in the grid-world bind them

Collaboration Structures: (strange: most are not Grid based…)

Repository and version control:	CVS
Problem tracking:	Bugzilla
Mailing lists:	Mailman
Web page authoring:	Wiki
Document Library:	homemade
Project Plan:	MS Project
Real-time conferencing:	Access Grid
Software Repository:	CVS

- Integrated Solutions?

Design & Operation Document Guide

- Cooperative Agreement
- Mgmt & Org Chart
- User Survey
- Strawman user guide
- Flagship apps
- Project Plan
- Security Policies
- Elevator Overview
- QA / Test / Accept Plan

- Primer
- Risk Mgmt Plan
- Accounting practices & policy
- Data practices & policy
- Customer service: Help desk, trouble tickets, SLAs
- Work and task list
- Real user guide

Organizing From the Top

- Cooperative Agreement:
  - The “Virt Org Bylaws”
    - How will decisions be made?
    - How will money be spent?
    - What are the approval processes?
    - How will resources be shared and billed?
    - How are disputes resolved?
    - What is Virt Org ***not*** responsible for?
  - Critical for real collaboration
- Mgmt & Org Chart
  - Executive Director, Director of Eng
  - Site Leads, etc

Organize Users (manage expectations)

- User survey: What ***features*** do the users want?
- Representatives from user community use survey results and write *Strawman User Guide*: detail each service, interfaces, user expectations
  - Computation
    - Any user can develop and run anywhere (transparently)
      - Implies identical environments: same paths/directories, same default tool versions (gcc, mpi), etc.
    - Single job can run across multiple clusters
      - Implies advanced scheduling capabilities, reservations?
  - Data
    - Intra-site High-Performance parallel filesystem
      - Implies GPFS, PVFS, etc
    - Consistent way to reach any HSM from any cluster
      - Implies ubiquitous HSM clients or GridFTP to remote HSM
    - Move data between clusters
      - Implies parallel GridFTP (GridFTP)
  - Support Services
    - Common system to track and respond to problems
      - Implies single interface used by both distributed operations and users
    - Single "transaction" for applying for a account
- Choose flagship apps

Users Are Practical What Services Are *Required*?

- Develop Applications (mostly, not grid tech)
  - compilers, viz, mpi, etc
- Submit jobs, transfer data (GRID!)
- Unified accounting (minimal grid tech)
- Parallel filesystem (not grid tech)
- Simplified portal
- Unified support (not grid tech)
- Testing and monitoring support (minimal grid tech)

Develop a Project Plan

- With *Strawman User Guide* and budget, pick technologies
  - Free software, $$$ SW/support, build it from scratch
  - How much of the strawman user guide can be implemented?
  - Which features must be dropped?
- Iterate with users & management
- Develop project plan
  - Staffing
  - Timelines
  - Contingency plans
  - Expansion plans early… what is non scalable from day 1?

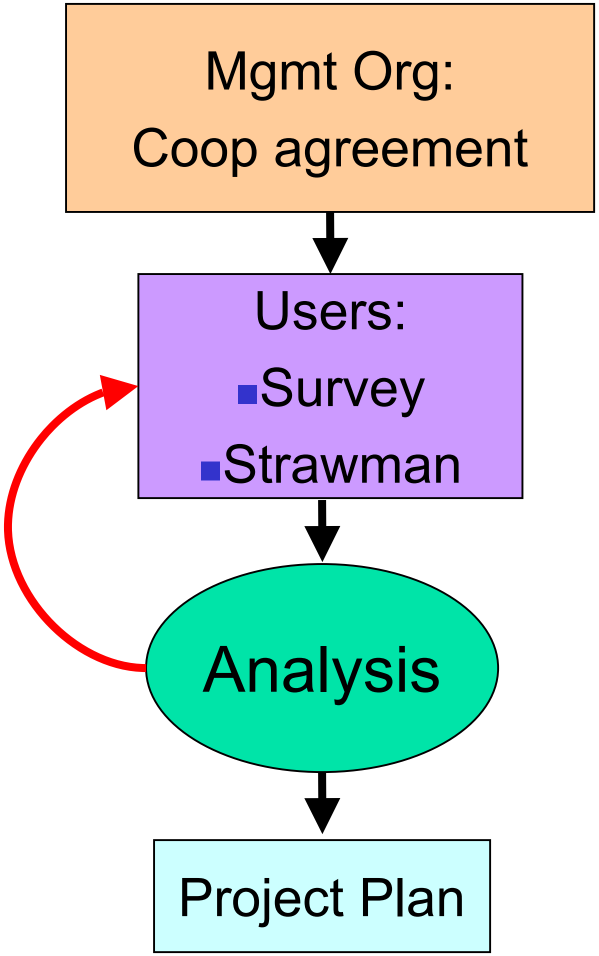

Work Breakdown Structure

- Seven Project Areas
  - Resources (Hardware Acquiring and Deployment)
  - Software & Services Environment
  - Network
  - Operations and System Support
  - Applications and User Services
  - Project Management and Service Policies
  - Expansion
- Integrates current status and progress
  - A "reset" injecting lessons learned
  - Useful to plan additional sites

Project Management Methodology

Detailed task-level planning
by working group

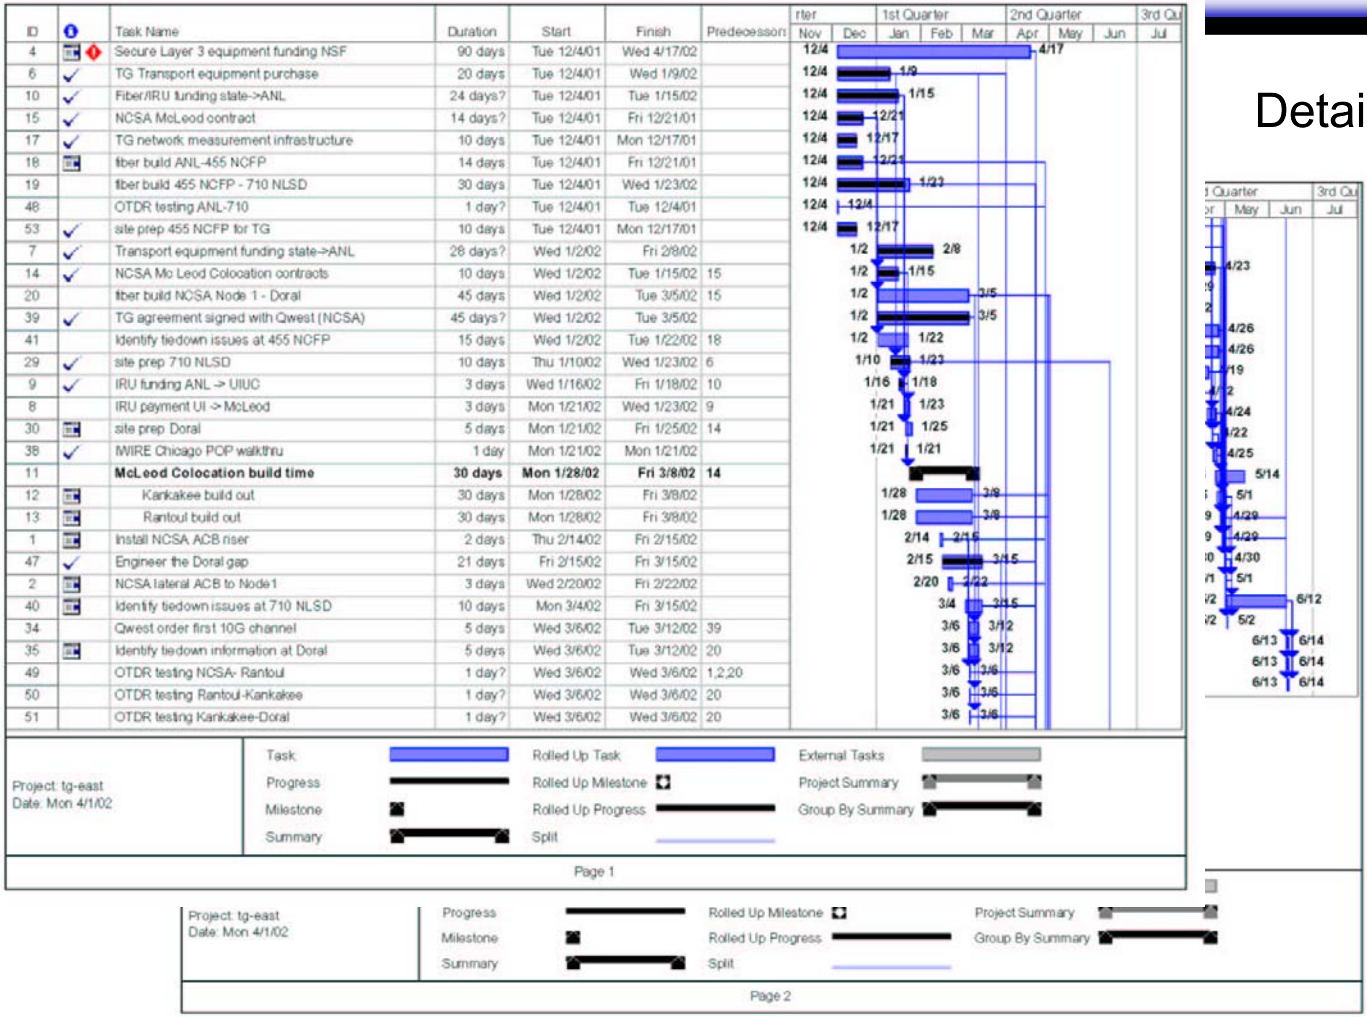

ID	ⓘ	Task Name	Duration	Start	Finish	Predecessors
4	▦ ◆	Secure Layer 3 equipment funding NSF	90 days	Tue 12/4/01	Wed 4/17/02	
6	✓	TG Transport equipment purchase	20 days	Tue 12/4/01	Wed 1/9/02	
10	✓	Fiber/IRU funding state->ANL	24 days?	Tue 12/4/01	Tue 1/15/02	
15	✓	NCSA McLeod contract	14 days?	Tue 12/4/01	Fri 12/21/01	
17	✓	TG network measurement infrastructure	10 days	Tue 12/4/01	Mon 12/17/01	
18	▦	fiber build ANL-455 NCFP	14 days	Tue 12/4/01	Fri 12/21/01	
19		fiber build 455 NCFP - 710 NLSD	30 days	Tue 12/4/01	Wed 1/23/02	
48		OTDR testing ANL-710	1 day?	Tue 12/4/01	Tue 12/4/01	
53	✓	site prep 455 NCFP for TG	10 days	Tue 12/4/01	Mon 12/17/01	
7	✓	Transport equipment funding state->ANL	28 days?	Wed 1/2/02	Fri 2/8/02	
14	✓	NCSA Mc Leod Colocation contracts	10 days	Wed 1/2/02	Tue 1/15/02	15
20		fiber build NCSA Node 1 - Doral	45 days	Wed 1/2/02	Tue 3/5/02	15
39	✓	TG agreement signed with Qwest (NCSA)	45 days?	Wed 1/2/02	Tue 3/5/02	
41		Identify tiedown issues at 455 NCFP	15 days	Wed 1/2/02	Tue 1/22/02	18
29	✓	site prep 710 NLSD	10 days	Thu 1/10/02	Wed 1/23/02	6
9	✓	IRU funding ANL -> UIUC	3 days	Wed 1/16/02	Fri 1/18/02	10
8		IRU payment UI -> McLeod	3 days	Mon 1/21/02	Wed 1/23/02	9
30	▦	site prep Doral	5 days	Mon 1/21/02	Fri 1/25/02	14
38	✓	MWIRE Chicago POP walkthru	1 day	Mon 1/21/02	Mon 1/21/02	
11		**McLeod Colocation build time**	**30 days**	**Mon 1/28/02**	**Fri 3/8/02**	**14**
12	▦	Kankakee build out	30 days	Mon 1/28/02	Fri 3/8/02	
13	▦	Rantoul build out	30 days	Mon 1/28/02	Fri 3/8/02	
1	▦	Install NCSA ACB riser	2 days	Thu 2/14/02	Fri 2/15/02	
47	✓	Engineer the Doral gap	21 days	Fri 2/15/02	Fri 3/15/02	
2	▦	NCSA lateral ACB to Node1	3 days	Wed 2/20/02	Fri 2/22/02	
40	▦	Identify tiedown issues at 710 NLSD	10 days	Mon 3/4/02	Fri 3/15/02	
34		Qwest order first 10G channel	5 days	Wed 3/6/02	Tue 3/12/02	39
35	▦	Identify tiedown information at Doral	5 days	Wed 3/6/02	Tue 3/12/02	20
49		OTDR testing NCSA- Rantoul	1 day?	Wed 3/6/02	Wed 3/6/02	1,2,20
50		OTDR testing Rantoul-Kankakee	1 day?	Wed 3/6/02	Wed 3/6/02	20
51		OTDR testing Kankakee-Doral	1 day?	Wed 3/6/02	Wed 3/6/02	20

Project: tg-east
Date: Mon 4/1/02

Task	🟦	Rolled Up Task	🟦	External Tasks	▭
Progress	▬	Rolled Up Milestone	◇	Project Summary	▬
Milestone	◆	Rolled Up Progress	▬	Group By Summary	▬
Summary	▬	Split	┈		

Page 1

Project: tg-east
Date: Mon 4/1/02

Progress	▬	Rolled Up Milestone	◇	Project Summary	▬
Milestone	◆	Rolled Up Progress	▬	Group By Summary	▬
Summary	▬	Split	┈		

Page 2

Other Useful Documents

- Elevator pitch overview (5 slides)
  - **Everyone** in the project should be able to describe and represent the project clearly
- QA / Test / Acceptance Plans
  - Ongoing, **independent** QA and testing is vital
- Risk management
  - Go / No-go decision points
  - Alternative technologies
  - Contingency budgets & plans

Design and Implementation Responsibilities:
The Working Groups

Working Groups

- User Services
- Applications
- Accounting
- Networking
- Visualization
- Data
- Testing / InfoServ
- Perf Eval
- External Relations
- Grid
- Security
- Operations
- Software / Clusters
- Interoperability

User Services

- Responsibilities
  - Direct customer (user) interaction
  - Maintain user-based focus during D, B, & O
  - Ombudsman
  - Allocations process
- Deliverables
  - Training and documentation
  - Allocations announcement & mgmt
- Key Issues
  - Match expectations

Applications

- Responsibilities
  - Represent users
  - Help negotiate expectations
  - Ombudsman
  - Porting and flagship app trials
- Deliverables
  - User survey
  - Strawman
- Key Issues
  - Rallying users to try new infrastructure

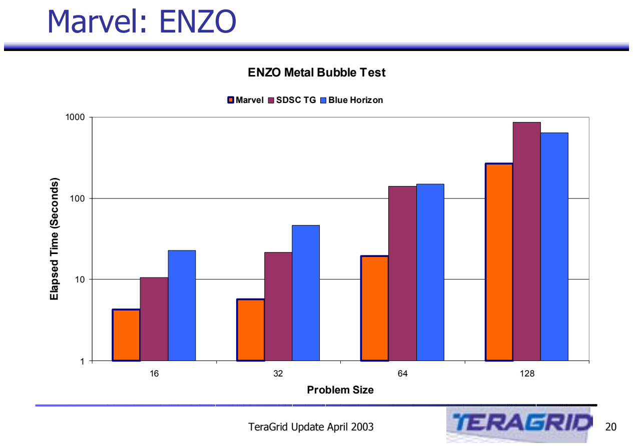

Marvel: ENZO

TeraGrid Update April 2003

Accounting

- Responsibilities
  - Develop and operate unified accounting system
  - Work with User Services on allocations
- Deliverables
  - Accounting system
  - Reporting system
- Key Issues
  - Do not reinvent wheel
  - Fairness
  - Reservation billing

Networking

- Responsibilities
  - Design, cost, and build network
  - Operate network
- Deliverables
  - A working network
  - Vendor quotes & purchasing
  - Help develop security plans
  - Network monitoring
- Key Issues
  - Reliability & Performance

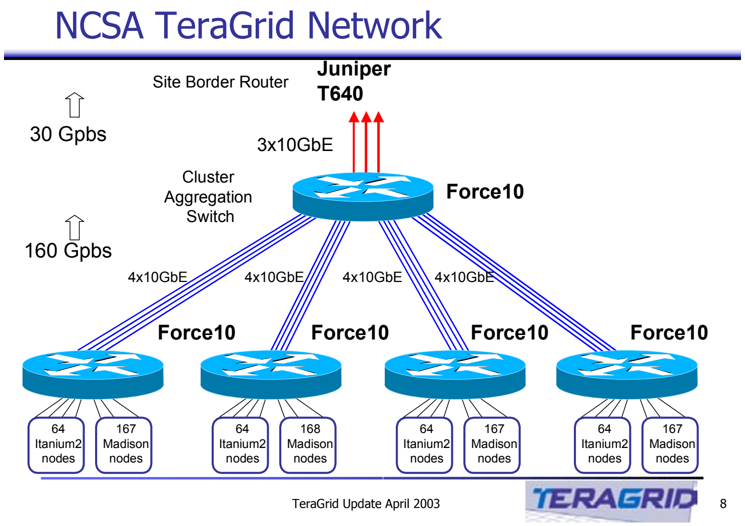

Visualization

- Responsibilities
  - Define viz services
- Deliverables
  - Software stack for viz
  - Operation models
- Key Issues
  - Scheduling (on demand, reservations)

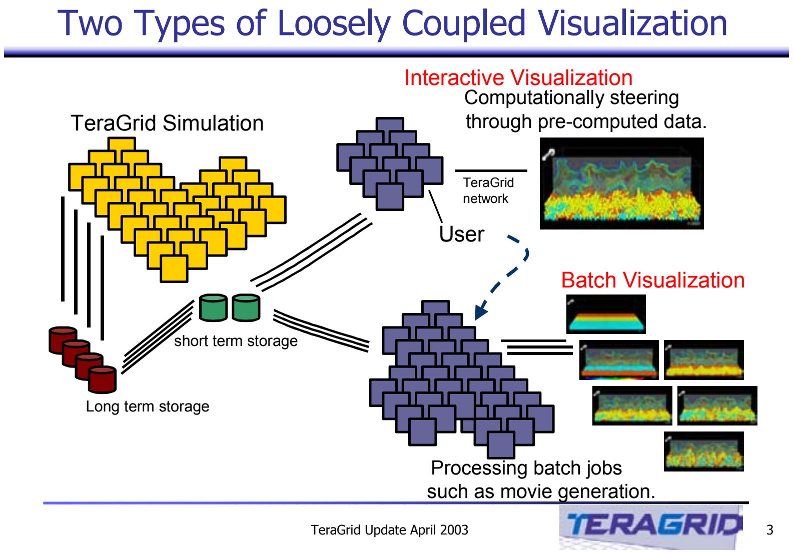

Two Types of Loosely Coupled Visualization

Interactive Visualization

Computationally steering through pre-computed data.

TeraGrid Simulation

TeraGrid network

User

Batch Visualization

short term storage

Long term storage

Processing batch jobs such as movie generation.

TeraGrid Update April 2003

TERAGRID

3

Data

- Responsibilities
  - Storage design
  - Access design
- Deliverables
  - Blueprints & custom softw
  - Implementation
- Key Issues
  - Practical tools and models
  - Close collaboration with users

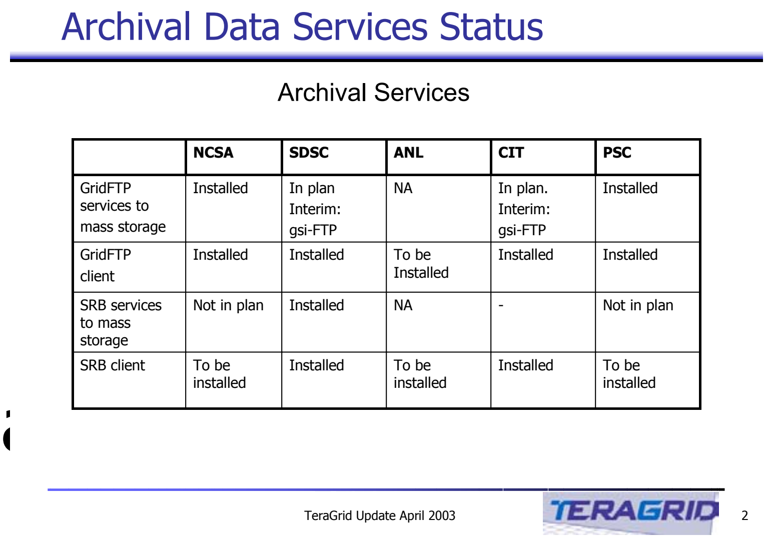

Archival Data Services Status

Archival Services

	NCSA	SDSC	ANL	CIT	PSC
GridFTP services to mass storage	Installed	In plan Interim: gsi-FTP	NA	In plan. Interim: gsi-FTP	Installed
GridFTP client	Installed	Installed	To be Installed	Installed	Installed
SRB services to mass storage	Not in plan	Installed	NA	-	Not in plan
SRB client	To be installed	Installed	To be installed	Installed	To be installed

TeraGrid Update April 2003 2

Testing / Info Serv

- Responsibilities
  - Design for testing & monitoring
  - Testing Grid components
  - Collecting monitoring data
- Deliverables
  - Testing and data collection system
- Key Issues
  - Scalability
  - Broad participation

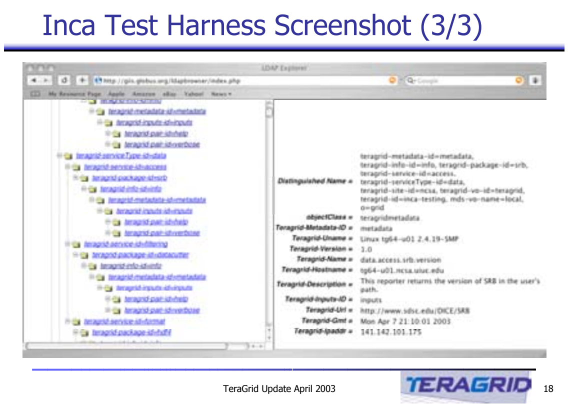

Inca Test Harness Screenshot (3/3)

LDAP Explorer
http://gis.globus.org/ldapbrowser/index.php
Google
My Resource Page
Apple
Amazon
eBay
Yahoo!
News
teragrid-metadata-id=metadata
teragrid-inputs-id=inputs
teragrid-pair-id=help
teragrid-pair-id=verbose
teragrid-serviceType-id=data
teragrid-service-id=access
teragrid-package-id=srb
teragrid-info-id=info
teragrid-metadata-id=metadata
teragrid-inputs-id=inputs
teragrid-pair-id=help
teragrid-pair-id=verbose
teragrid-service-id=filtering
teragrid-package-id=datacutter
teragrid-info-id=info
teragrid-metadata-id=metadata
teragrid-inputs-id=inputs
teragrid-pair-id=help
teragrid-pair-id=verbose
teragrid-service-id=format
teragrid-package-id=hdf4

Distinguished Name = teragrid-metadata-id=metadata, teragrid-info-id=info, teragrid-package-id=srb, teragrid-service-id=access, teragrid-serviceType-id=data, teragrid-site-id=ncsa, teragrid-vo-id=teragrid, teragrid-id=inca-testing, mds-vo-name=local, o=grid
objectClass = teragridmetadata
Teragrid-Metadata-ID = metadata
Teragrid-Uname = Linux tg64-u01 2.4.19-SMP
Teragrid-Version = 1.0
Teragrid-Name = data.access.srb.version
Teragrid-Hostname = tg64-u01.ncsa.uiuc.edu
Teragrid-Description = This reporter returns the version of SRB in the user's path.
Teragrid-Inputs-ID = inputs
Teragrid-Url = http://www.sdsc.edu/DICE/SRB
Teragrid-Gmt = Mon Apr 7 21:10:01 2003
Teragrid-Ipaddr = 141.142.101.175

TeraGrid Update April 2003 TERAGRID 18

Perf Eval

- Responsibilities
  - Assist tuning the machines
  - Understand performance
  - Predict performance
- Deliverables
  - Benchmark suite for regression
- Key Issues
  - User interaction
  - Regression testing for performance

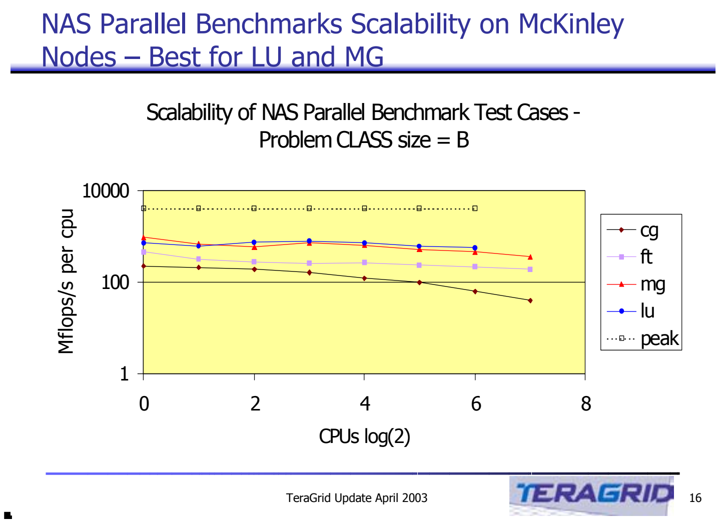

NAS Parallel Benchmarks Scalability on McKinley Nodes – Best for LU and MG

Scalability of NAS Parallel Benchmark Test Cases - Problem CLASS size = B

TeraGrid Update April 2003

16

External Relations

- Responsibilities
  - Organize unified external msg
  - Forum for discussing balanced PR message
- Deliverables
  - Communication plan
  - PR schedule
- Key Issues
  - Balance
  - Coherent msg

Grid

- Responsibilities
  - Design practical service model
  - Explore production-ready technologies
- Deliverables
  - Grid implementation plan
  - Train operations
- Key Issues
  - Research software is often not production ready
  - Keep feet firmly planted by working closely with user services and applications (user proxies)
  - Grid technology is new, Ops team may have no experience supporting it

Security

- Responsibilities
  - Examine vulnerability scenarios
  - Explore certificate acceptance policies
    - Where are the mismatches?
  - Depending on project, create a CA
- Deliverables
  - Security plan
  - Certificate policy
  - Certificate management plan (Gridmap tools, etc)
- Key Issues
  - Trust between organizations

Operations

- Responsibilities
  - Running production systems
- Deliverables
  - Operations plan, including:
    - Trouble tix system
    - Account generation & accounting mgmt
    - Ops-based monitoring
  - Workplan for working Ops (virtual) center
- Key Issues
  - Reliability, blame, new models, interacting services

Software / Clusters

- Responsibilities
  - Software stack
  - Build & administer resources
- Deliverables
  - Implementation plan for software stack
  - Working compute & data resources
- Key Issues
  - Not Grid weenies
  - Change control
  - Matching functionality expected by users

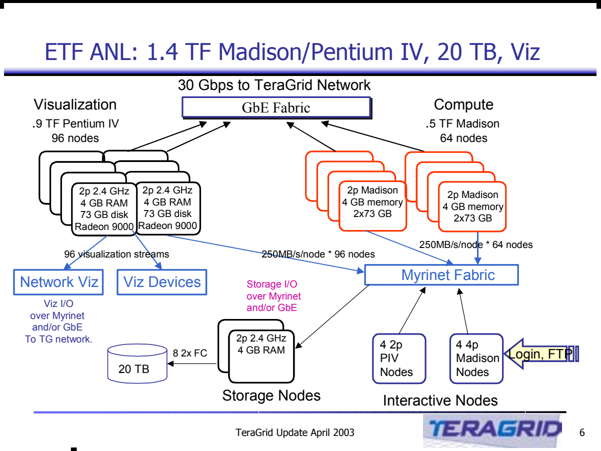

ETF ANL: 1.4 TF Madison/Pentium IV, 20 TB, Viz

TeraGrid Update April 2003

TERAGRID 6

Interoperability

- Responsibilities
  - Explore interactions to diverse systems
    - E.g: MPI across Myrinet and Quadrics together
- Deliverables
  - Enumeration of components that can be shared, and where interoperable interfaces are required
- Key Issues
  - End to end (a to b to c)
  - Complexity and resources required to add and maintain abstract interfaces

Chapter 3a
Design Examples (part 1)

Current HPC Cluster Environments are Painful for Users

- Every machine is completely different, essentially we tell users:

  "Please learn this site-specific information before you can use the cluster"

  - My scheduler
  - My installed math libraries
  - My default paths
  - My parallel IO system
  - My help/ticket system
  - My documentation
- They may be **functional**, but not **interoperable**

Common TeraGrid Software Stack

- A social contract with the user:
  - LORA: Learn Once, Run Anywhere
- Precise definitions:
  - Services
  - Software
  - User Environment
- **Reproducibility**
  - Standard configure, build, and install
  - Single CVS repository for software

CTSS Example

- SuSE SLES 8
- X-cat 1.2 beta4
- OpenPBS v2.3.16
- Maui 3.2.5p2
- MPICH v1.2.5 (G2, many builds)
- MPICH v1.1.2 (VMI build)
- gm v1.6.3
- VMI/CRM
- Globus 2.2.4 +patches
- Condor-G
- gsi-ssh 2.1
- GPT Wizard
- GPT 2.2.8b1
- SoftEnv 1.4.2
- MyProxy v1.2
- Intel compilers 7.0
- Gnu compilers
- HDF4/5
- SRB client

Testing, Verification, Monitoring

New software, built to create the Grid Hosting Environment

- Goal: Use ‘unit test’, ‘version reporter’, and ‘integration tests’ to assure each the quality of each component in the system

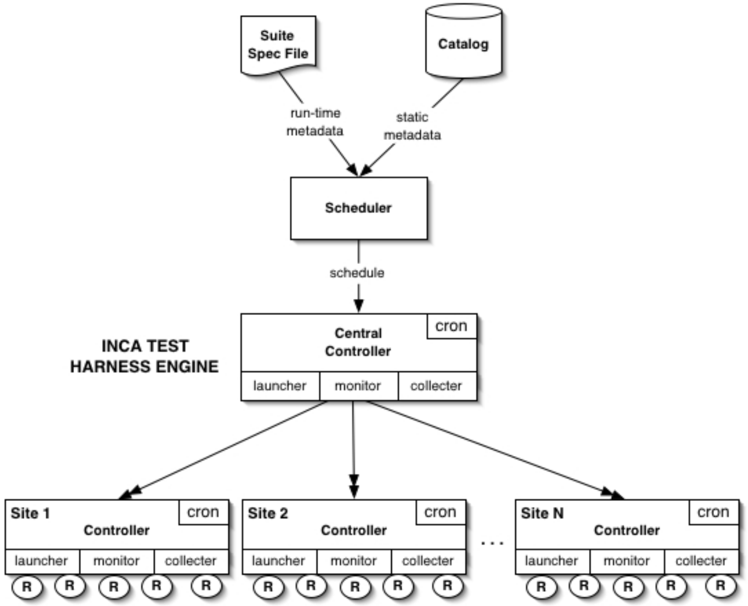

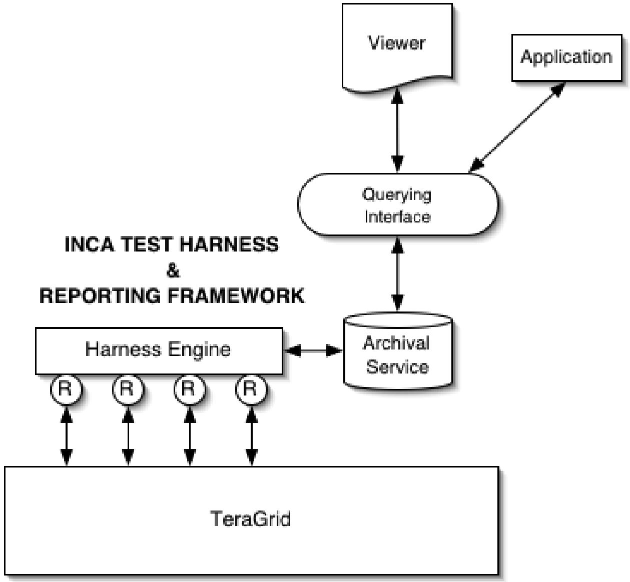

Software Development Principles For Building The TeraGrid

- Drive development with applications
- TG Software Stack and environment is homogeneous across all sites **except where a difference is clearly justified to the users and driven by their requirements**
- Every package is versioned and the build/install/config parameters reproducible
- CTSS packages have specific versions and tests; change is carefully managed
- A Test Harness is constantly working to insure stability and conformance to the TG Hosting Environment
- After successful deployment on the TestGrid, new components are tested on the ProdGrid

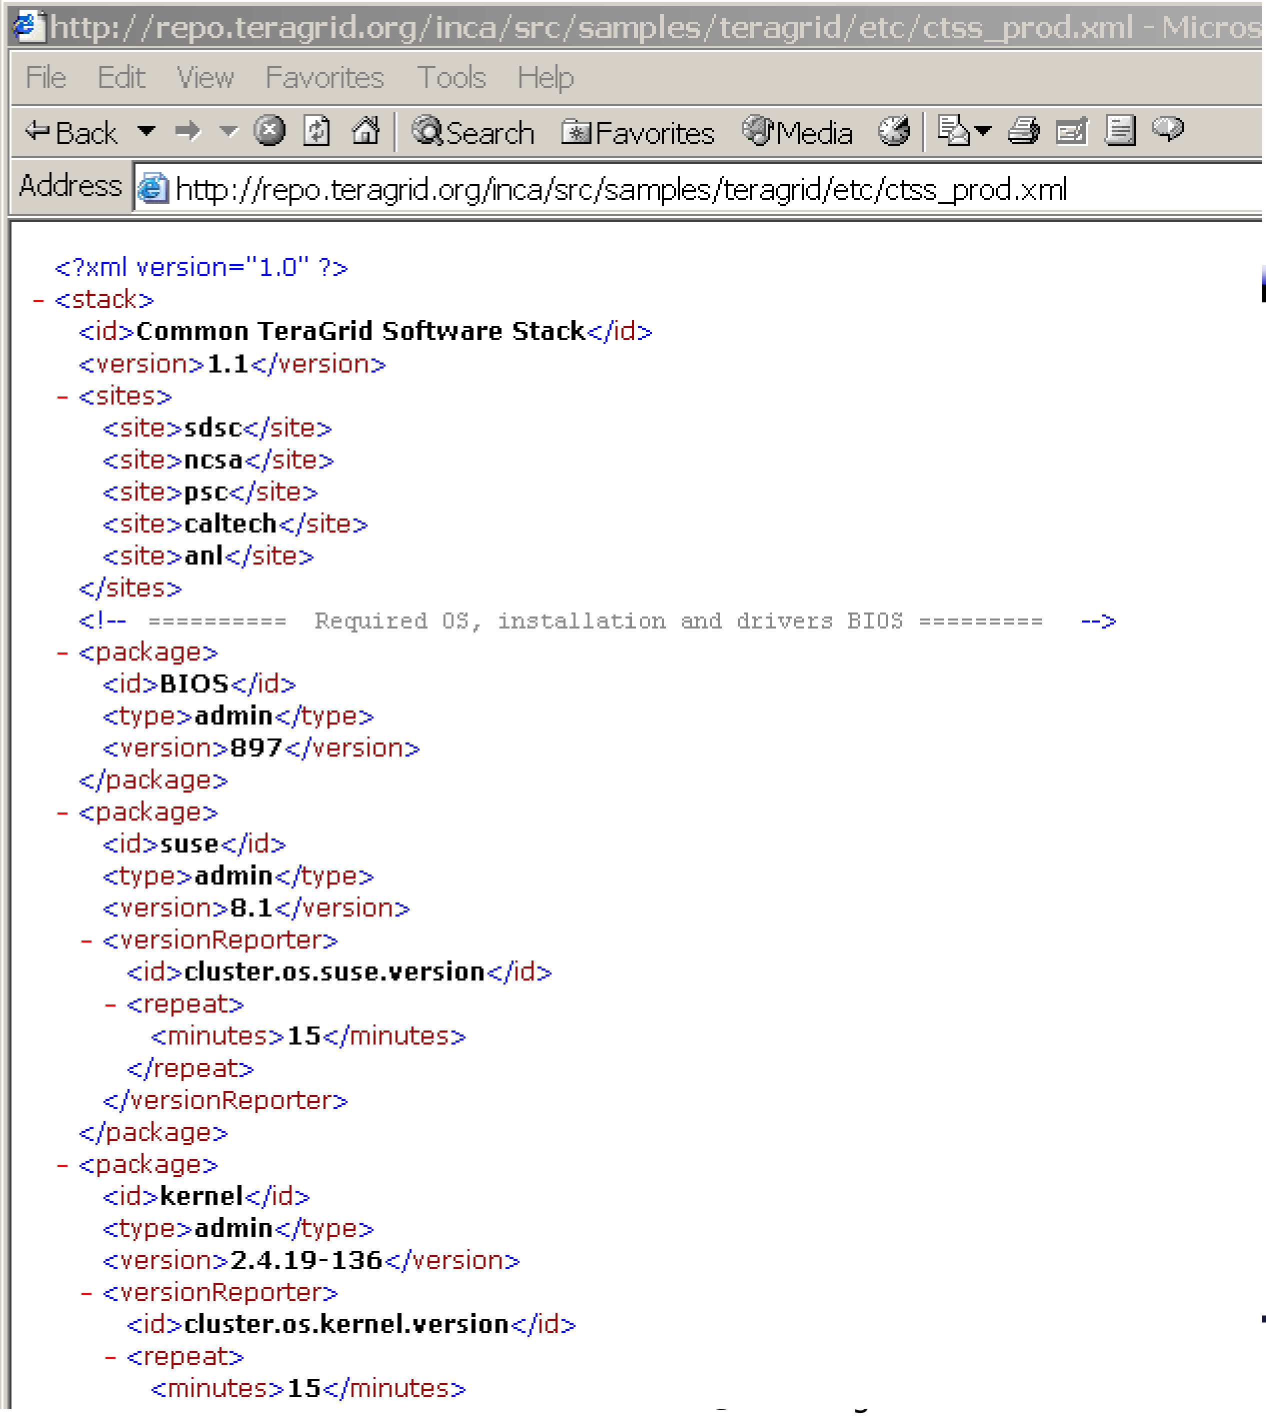
http://repo.teragrid.org/inca/src/samples/teragrid/etc/ctss_prod.xml - Micros
File Edit View Favorites Tools Help
Back
Search
Favorites
Media
Address
http://repo.teragrid.org/inca/src/samples/teragrid/etc/ctss_prod.xml

```
  <?xml version="1.0" ?>
- <stack>
    <id>Common TeraGrid Software Stack</id>
    <version>1.1</version>
  - <sites>
      <site>sdsc</site>
      <site>ncsa</site>
      <site>psc</site>
      <site>caltech</site>
      <site>anl</site>
    </sites>
    <!-- ========== Required OS, installation and drivers BIOS =========  -->
  - <package>
      <id>BIOS</id>
      <type>admin</type>
      <version>897</version>
    </package>
  - <package>
      <id>suse</id>
      <type>admin</type>
      <version>8.1</version>
    - <versionReporter>
        <id>cluster.os.suse.version</id>
      - <repeat>
          <minutes>15</minutes>
        </repeat>
      </versionReporter>
    </package>
  - <package>
      <id>kernel</id>
      <type>admin</type>
      <version>2.4.19-136</version>
    - <versionReporter>
        <id>cluster.os.kernel.version</id>
      - <repeat>
          <minutes>15</minutes>
```

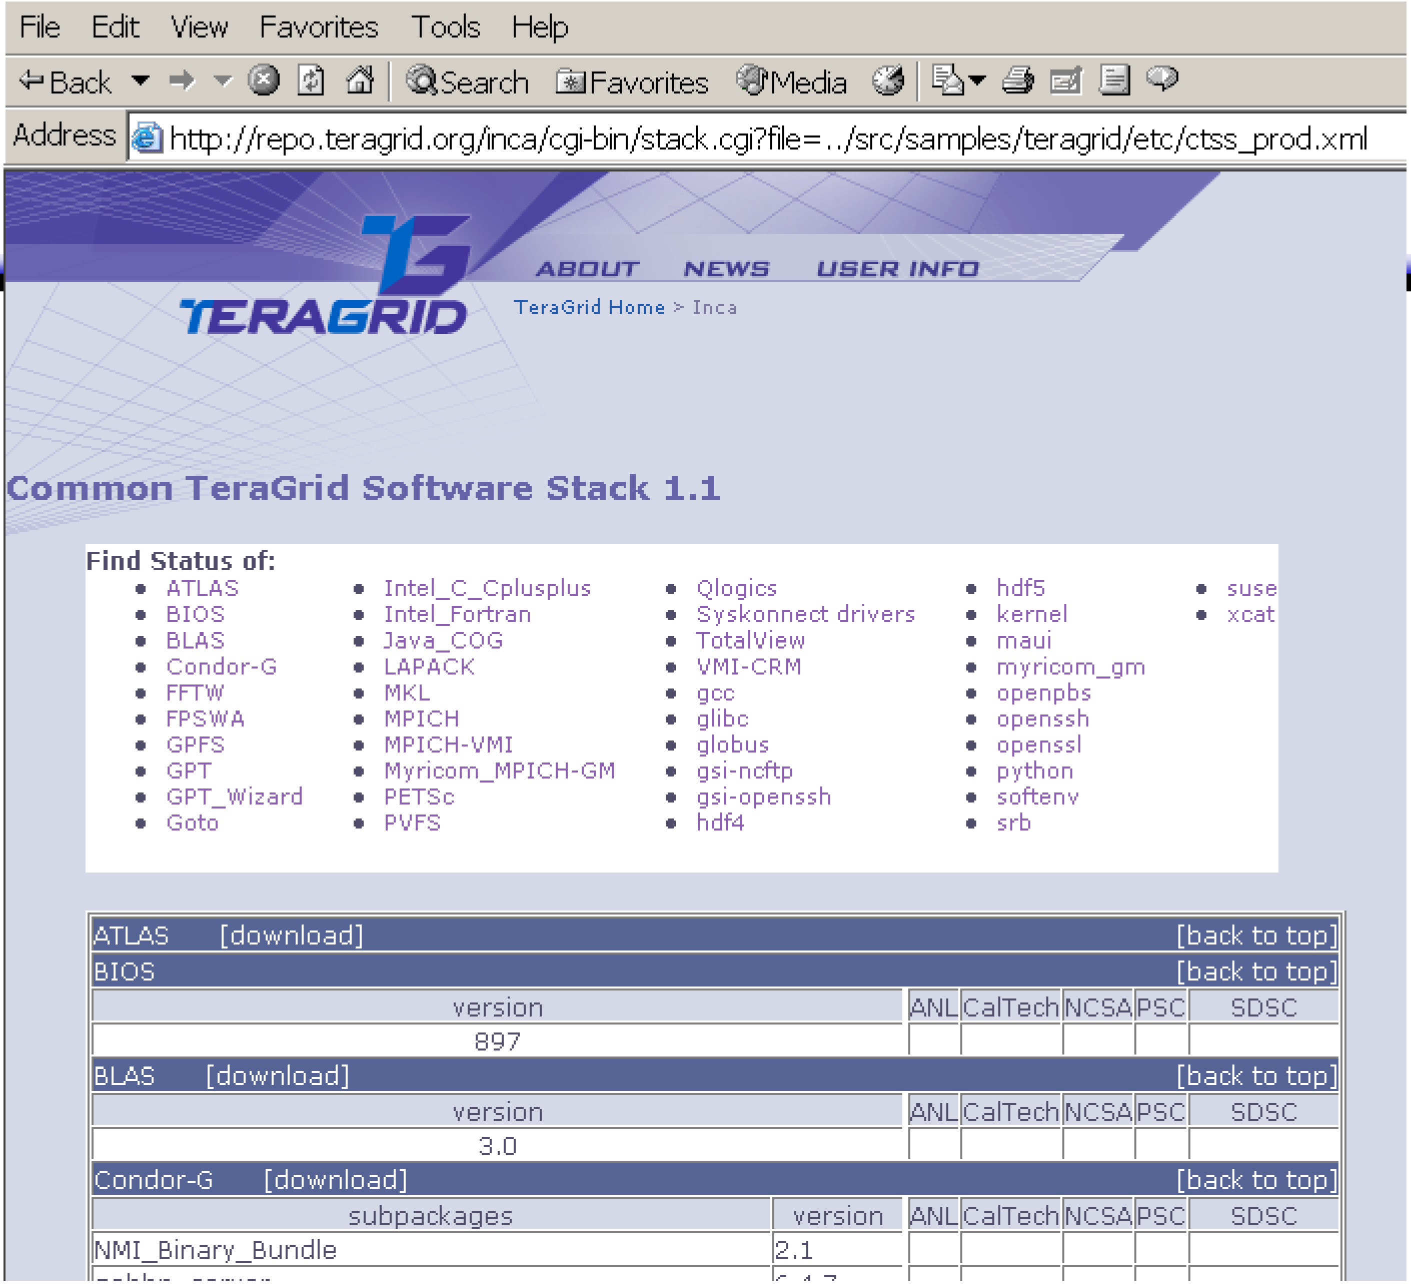

http://repo.teragrid.org/inca/cgi-bin/stack.cgi?file=../src/samples/teragrid/etc/ctss_prod.xml

TeraGrid Home > Inca

Common TeraGrid Software Stack 1.1

Find Status of:

- ATLAS
- BIOS
- BLAS
- Condor-G
- FFTW
- FPSWA
- GPFS
- GPT
- GPT_Wizard
- Goto
- Intel_C_Cplusplus
- Intel_Fortran
- Java_COG
- LAPACK
- MKL
- MPICH
- MPICH-VMI
- Myricom_MPICH-GM
- PETSc
- PVFS
- Qlogics
- Syskonnect drivers
- TotalView
- VMI-CRM
- gcc
- glibc
- globus
- gsi-ncftp
- gsi-openssh
- hdf4
- hdf5
- kernel
- maui
- myricom_gm
- openpbs
- openssh
- openssl
- python
- softenv
- srb
- suse
- xcat

ATLAS [download] [back to top]

BIOS [back to top]

version	ANL	CalTech	NCSA	PSC	SDSC
897					

BLAS [download] [back to top]

version	ANL	CalTech	NCSA	PSC	SDSC
3.0					

Condor-G [download] [back to top]

subpackages	version	ANL	CalTech	NCSA	PSC	SDSC
NMI_Binary_Bundle	2.1					

Two Types of Loosely Coupled Visualization

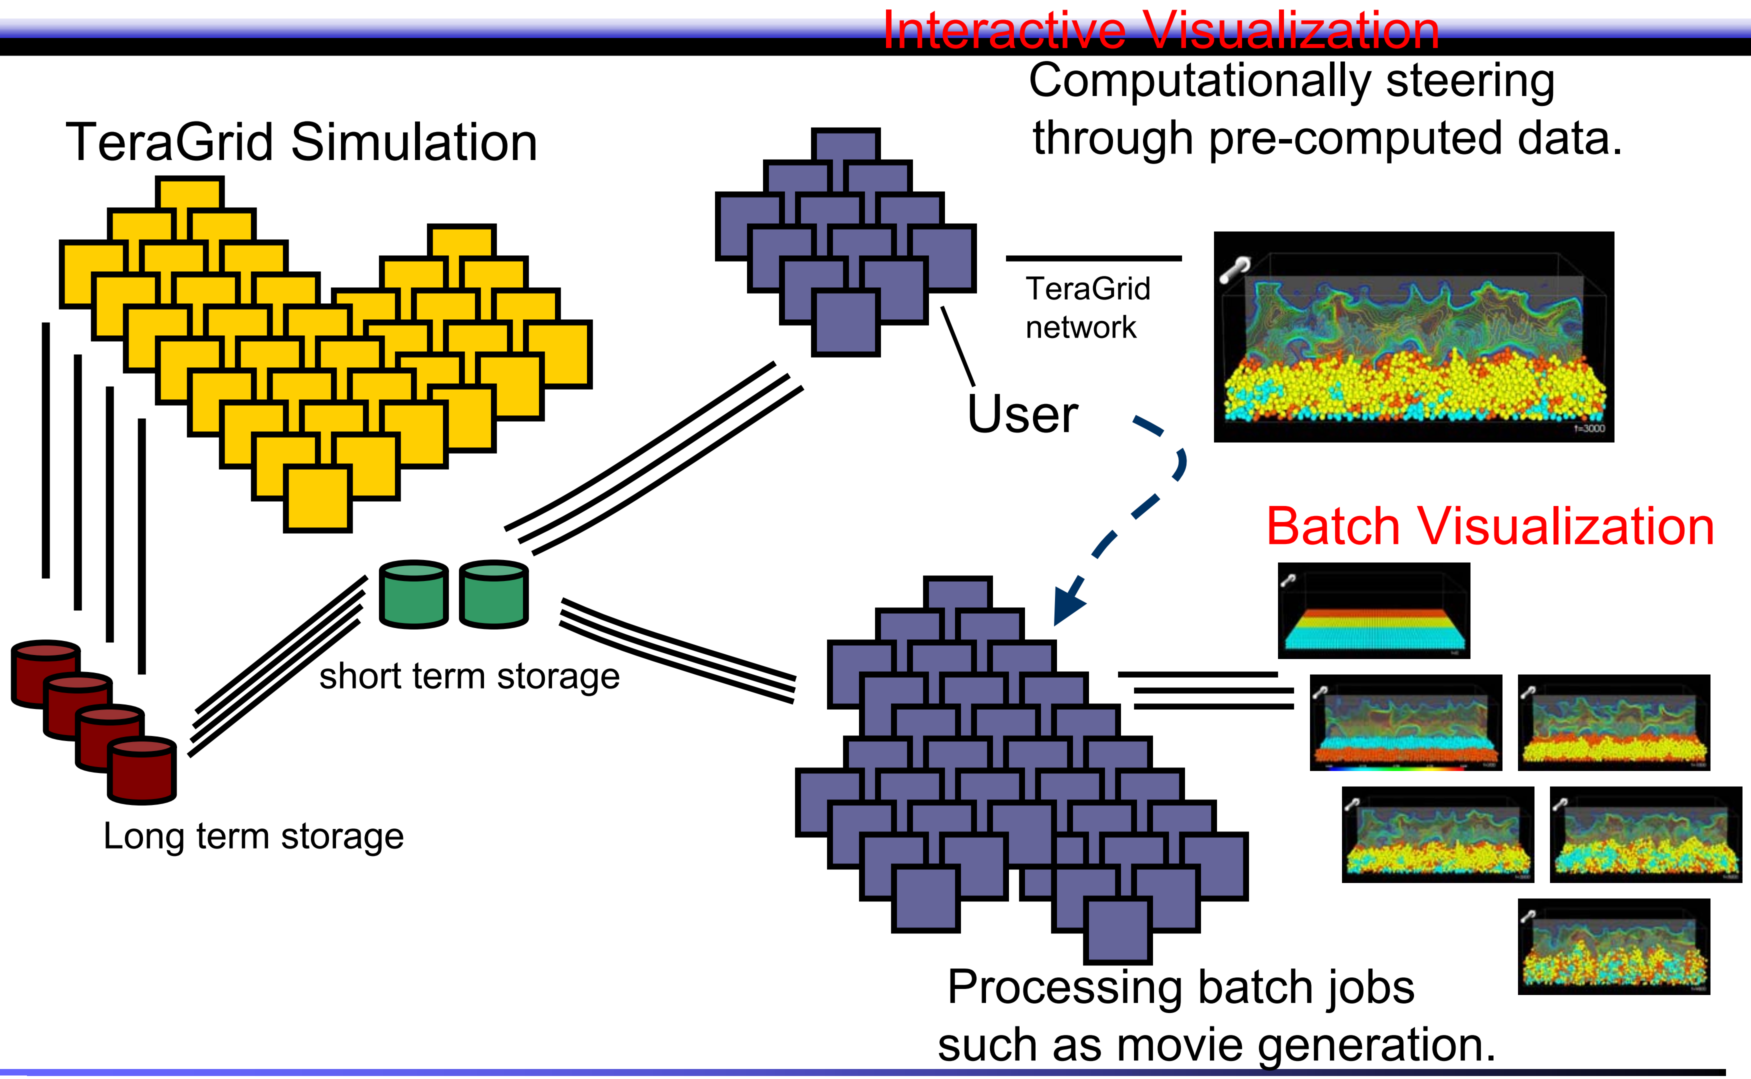

Interactive Visualization

Computationally steering through pre-computed data.

TeraGrid Simulation

TeraGrid network

User

Batch Visualization

short term storage

Long term storage

Processing batch jobs such as movie generation.

On-Demand and Collaborative Visualization

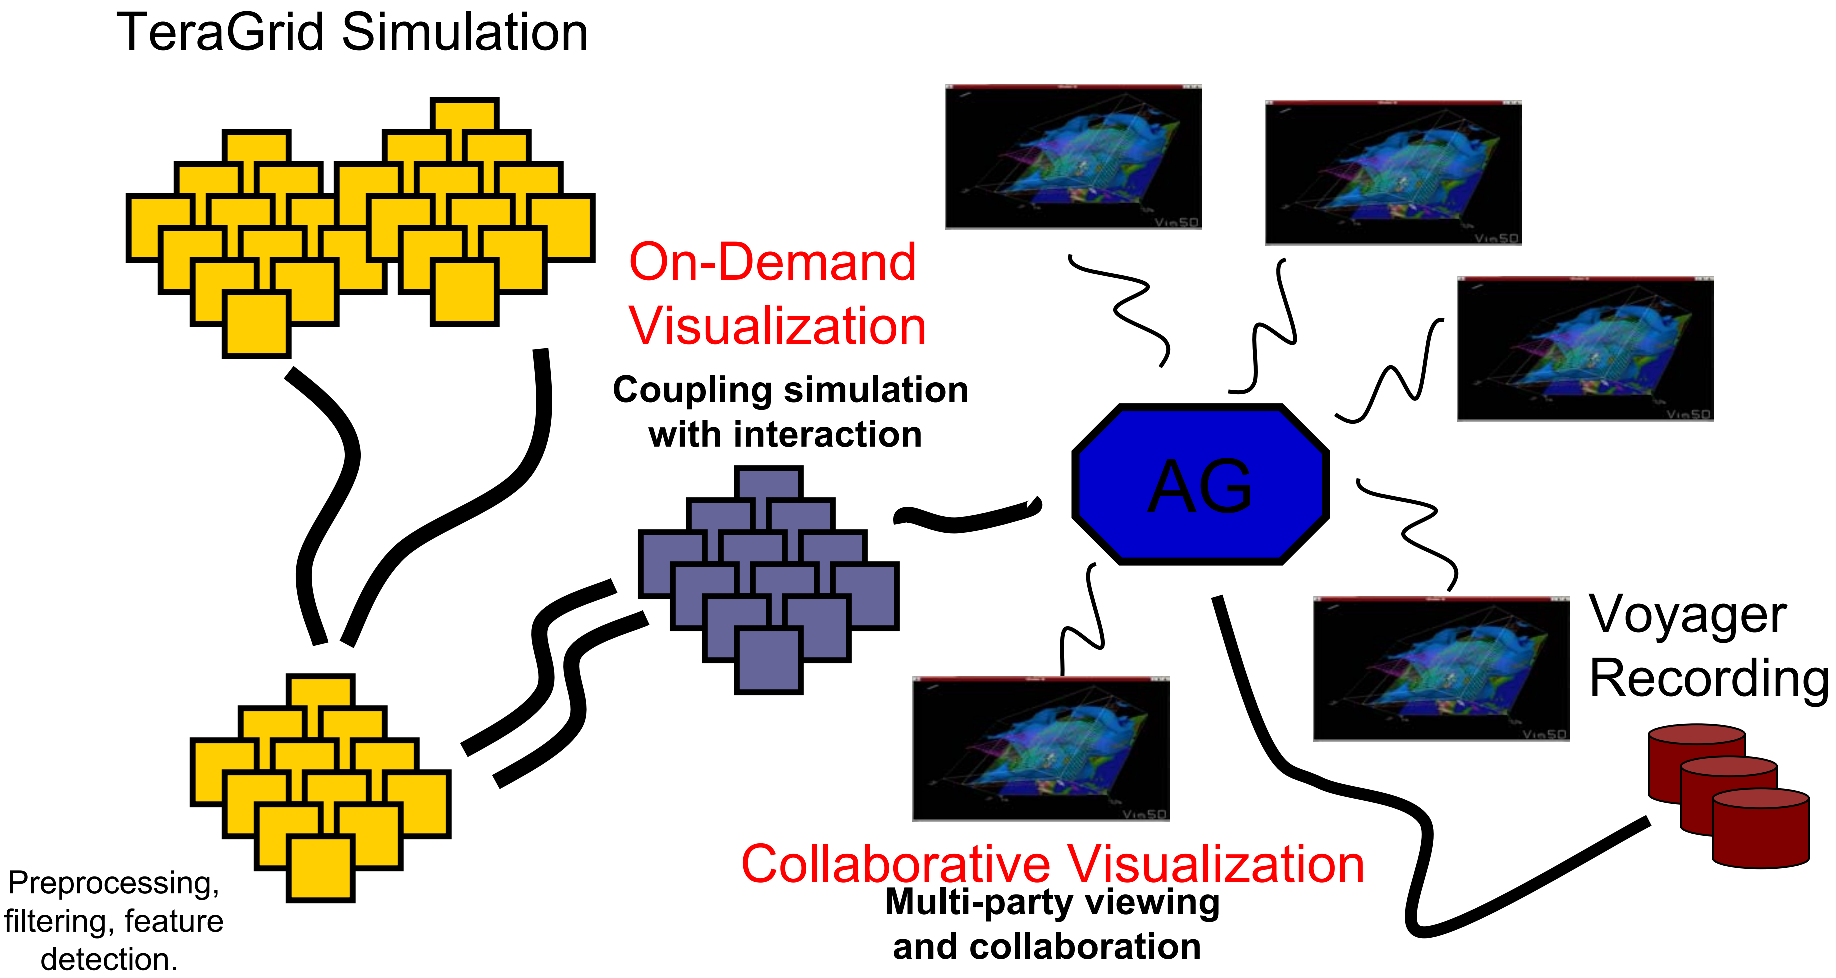

TeraGrid Simulation
On-Demand Visualization
Coupling simulation with interaction
AG
Voyager Recording
Collaborative Visualization
Multi-party viewing and collaboration
Preprocessing, filtering, feature detection.

Visualization – Sample Use Cases

App	Modality	Tools	Data
Collaborative Analysis of Atmospheric Simulations	Collaborative, Interactive Analysis.	AG 2.0 Grid VTK MPICHG2	GridFTP in. AG distribution
Interactive Visualization of Time-Dependent CFD Data	Collaborative, Interactive	AG 2.0 Grid VTK VisBench	GridFTP in. AG distribution
Volume Rendering for Production Visualization	Batch Rendering of Movies	NPACI Scalable Visualization Alpha Project Volume Render. Radiosity renderer. APST, PBS.	Dynamic, on-demand use of SRB.

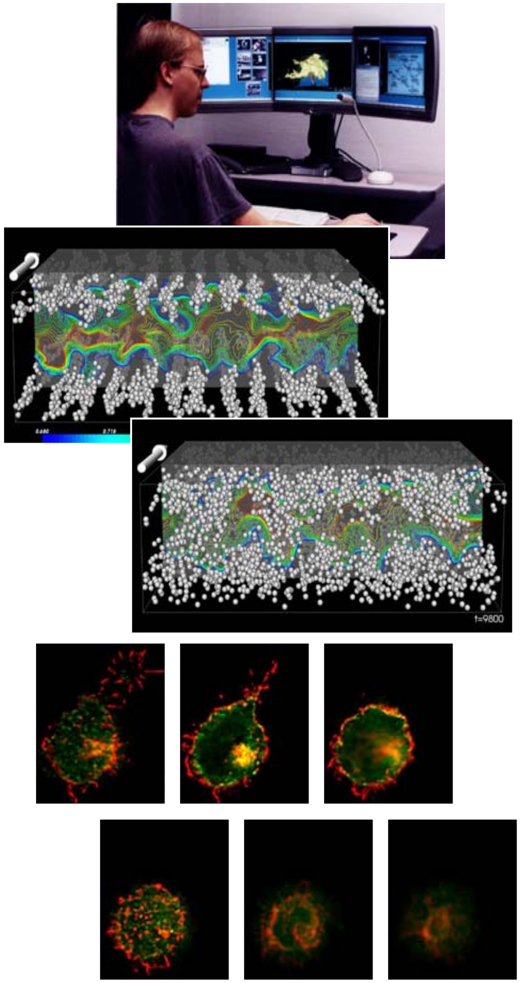

Summary: Day 1

- Virtual Organization & Cooperative Agreement
- Management and Decision Making
  - Exec Dir, Dir of Eng, Team Leaders, etc
- Design -> Build -> Operate
- Eliminate hype, be practical, get user commitment early
- Collaboration technology required from day 0
- Installing and understanding Globus is the least of your concerns
- Integration and engineering is costly, budget for it

Chapter 3b
Design Examples (part 2)

Outline of Presentation

- Day 1: Min:
  - Chapter 1: Introduction 30
  - Chapter 2: Creating a blueprint 60
  - Chapter 3a: Design examples (part 1) 30
- Day 2:
  - Chapter 3b: Design examples (part 2) 60
  - Chapter 4: Construction phase 20
  - Chapter 5: Operating the grid 20
  - Chapter 6: Summary / questions 20

TeraGrid Networking

- Clusters and Sites (what to connect)
- TeraGrid Cluster Architecture
- WAN Architecture
  - Original Design: Optical Mesh
  - Scalable Backplane
- Optical Networks
- Expansion and Management

TeraGrid Circa July 2003

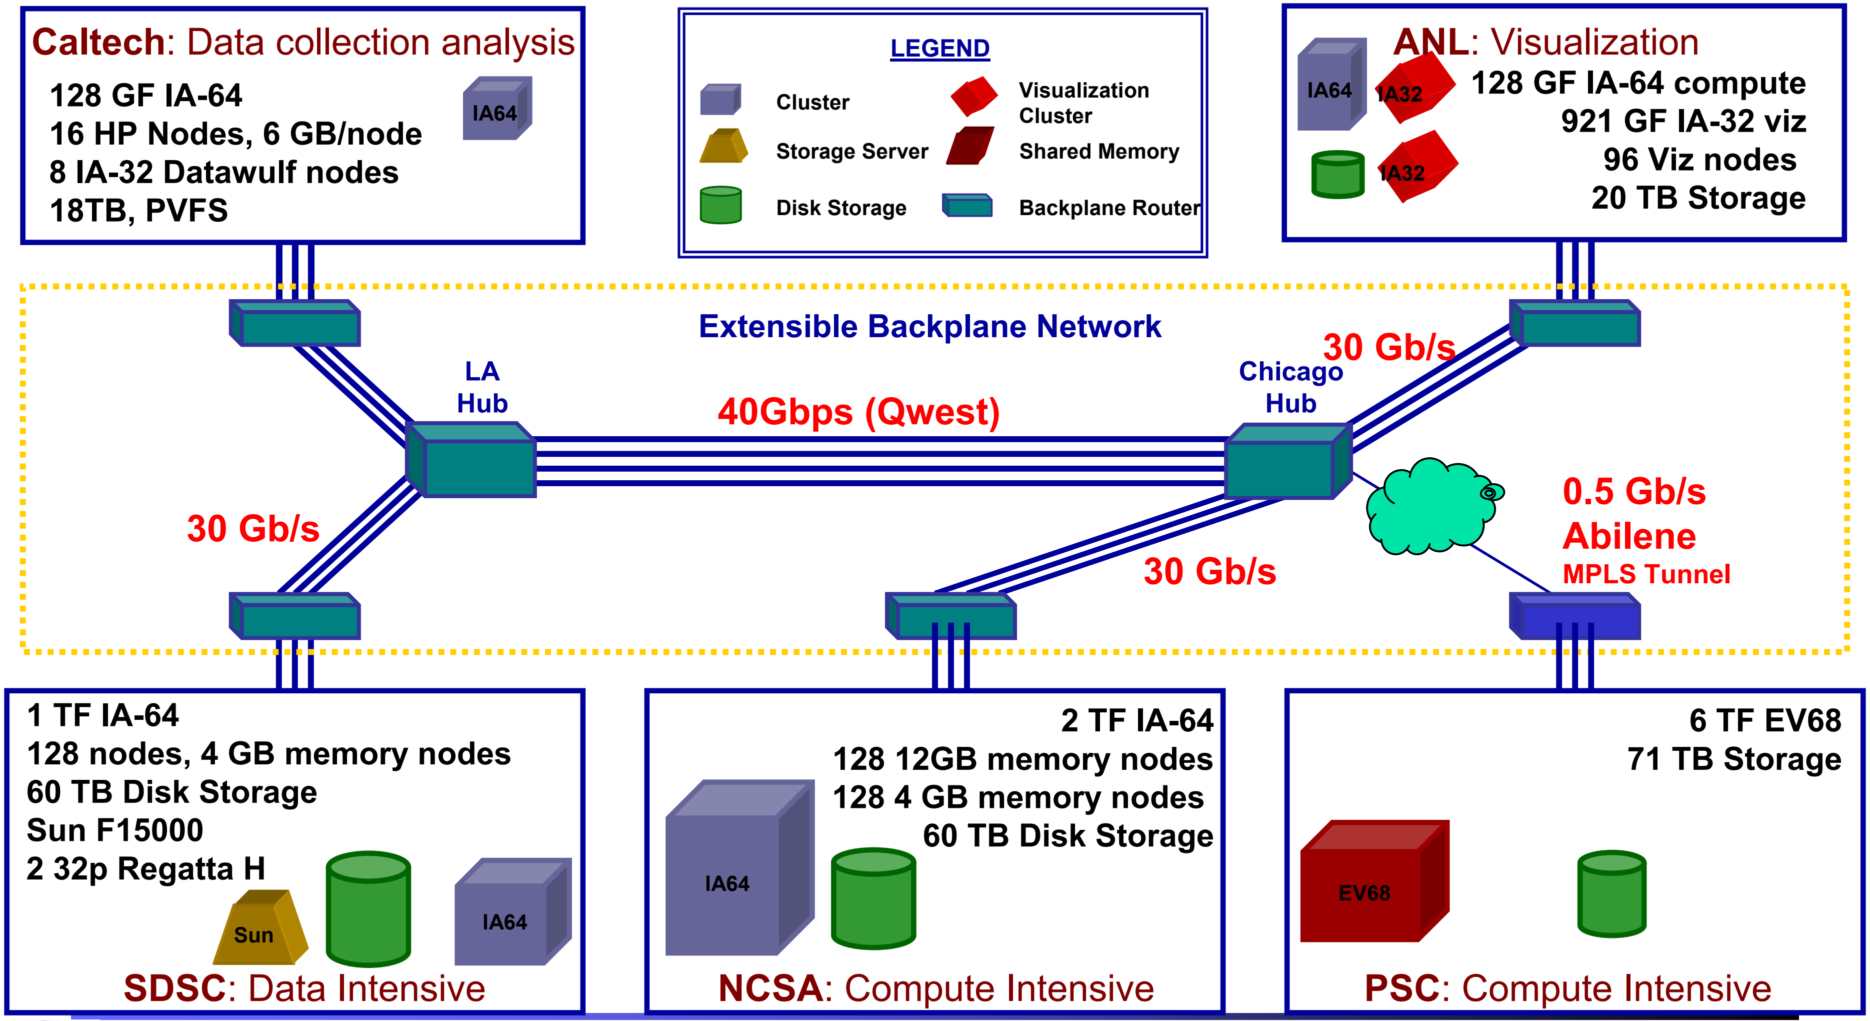

Caltech: Data collection analysis

128 GF IA-64
16 HP Nodes, 6 GB/node
8 IA-32 Datawulf nodes
18TB, PVFS

LEGEND

Cluster
Storage Server
Disk Storage
Visualization Cluster
Shared Memory
Backplane Router

ANL: Visualization

128 GF IA-64 compute
921 GF IA-32 viz
96 Viz nodes
20 TB Storage

Extensible Backplane Network

LA Hub

Chicago Hub

40Gbps (Qwest)

30 Gb/s

30 Gb/s

30 Gb/s

0.5 Gb/s
Abilene
MPLS Tunnel

1 TF IA-64
128 nodes, 4 GB memory nodes
60 TB Disk Storage
Sun F15000
2 32p Regatta H

SDSC: Data Intensive

2 TF IA-64
128 12GB memory nodes
128 4 GB memory nodes
60 TB Disk Storage

NCSA: Compute Intensive

6 TF EV68
71 TB Storage

PSC: Compute Intensive

Terascale Cluster Architecture

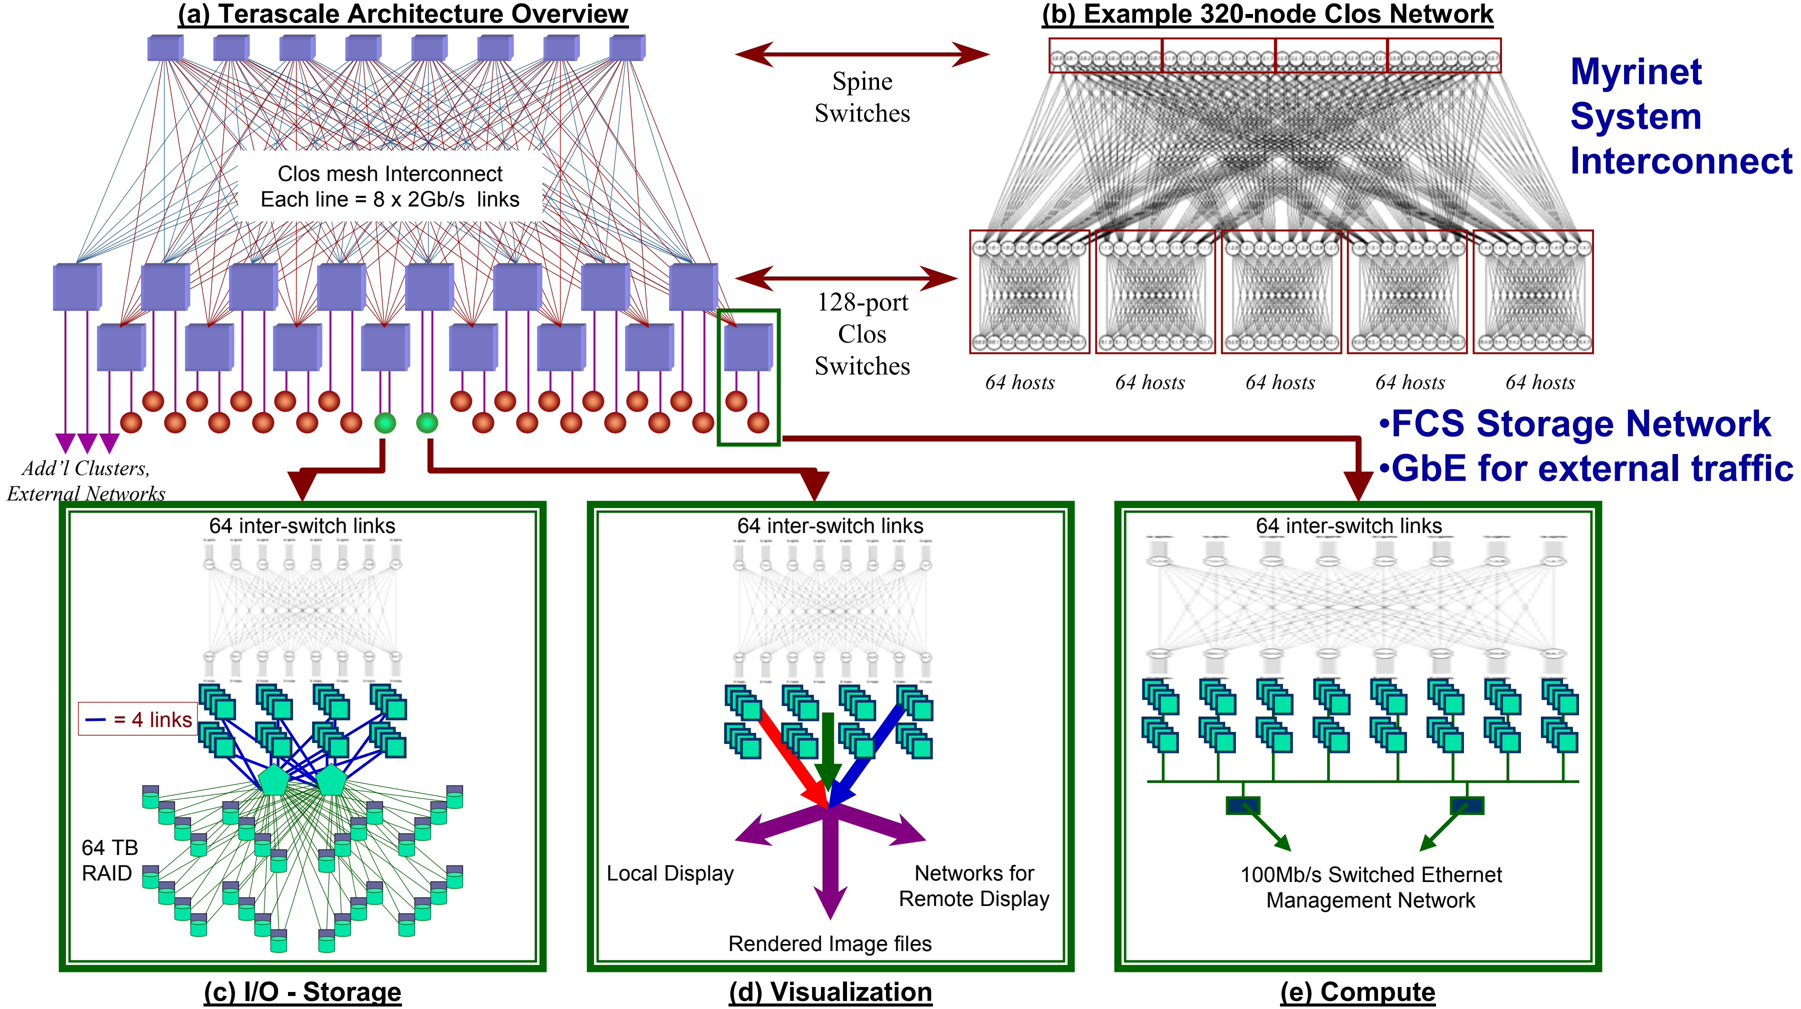

(a) Terascale Architecture Overview

(b) Example 320-node Clos Network

Myrinet System Interconnect

- FCS Storage Network
- GbE for external traffic

(c) I/O - Storage

(d) Visualization

(e) Compute

Charlie Catlett <catlett@mcs.anl.gov>
Pete Beckman <beckman@mcs.anl.gov>

TeraGrid Cluster Architecture

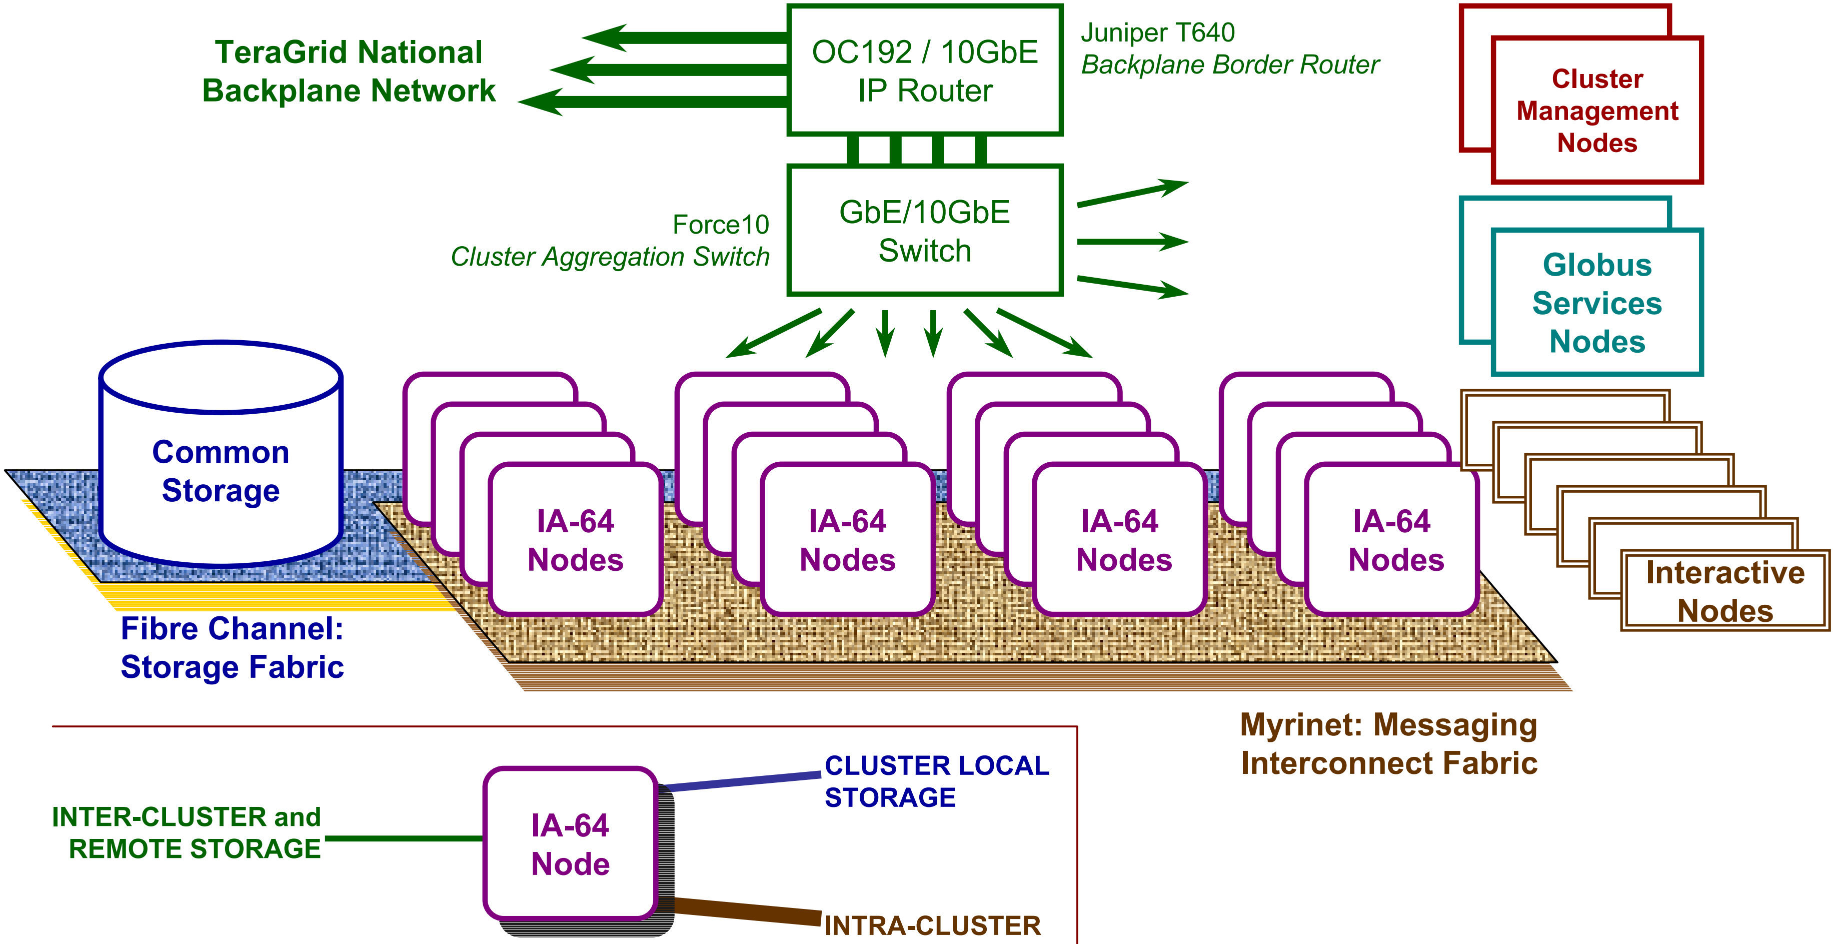

Each Node connects to 3 networks

Typical Cluster-WAN Architecture

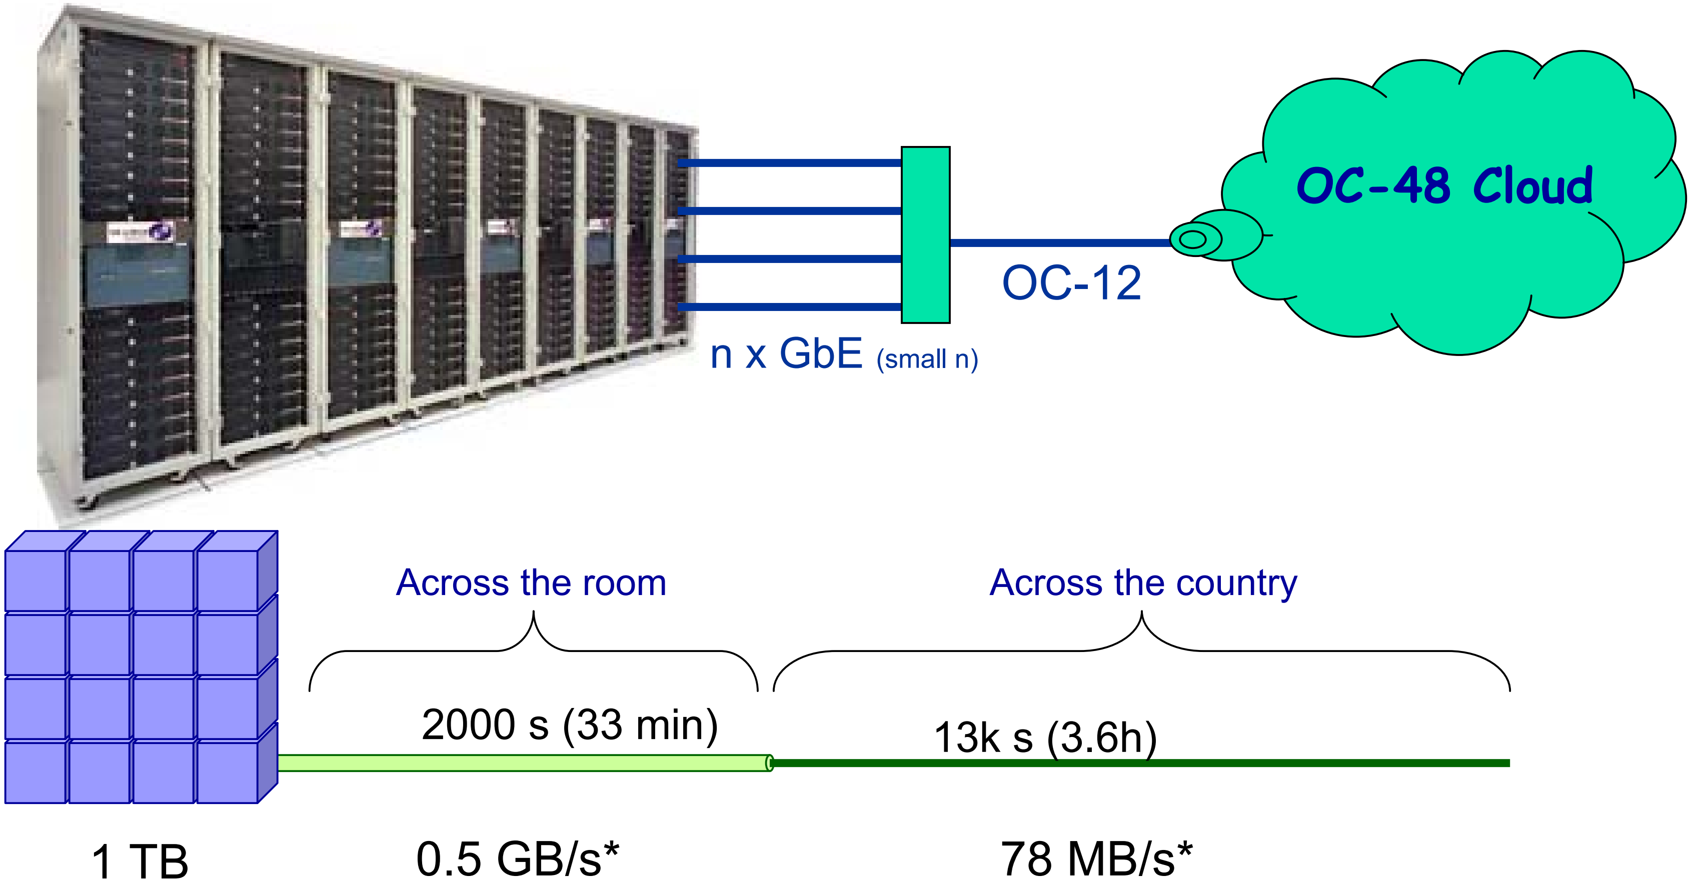

*of course in practice the performance is some percentage of this peak number

TeraGrid Cluster WAN

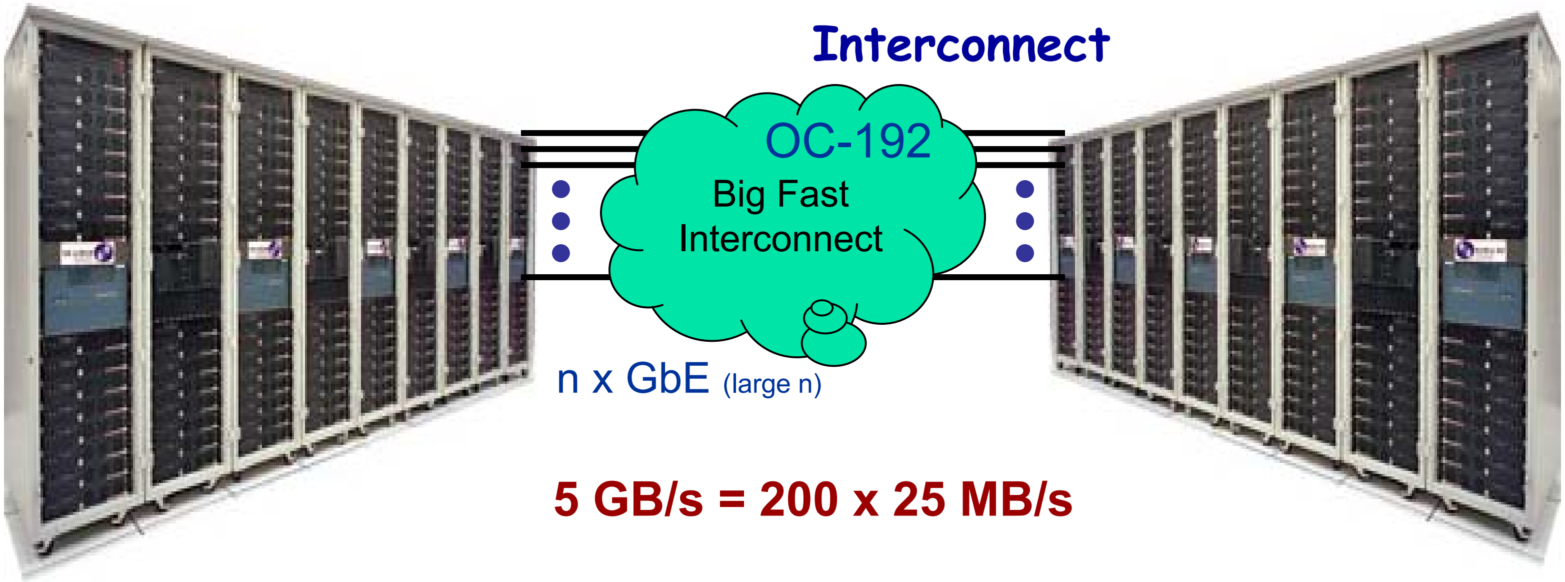

5 GB/s = 200 x 25 MB/s

25 MB/s = 200 Mb/s or 20% of GbE

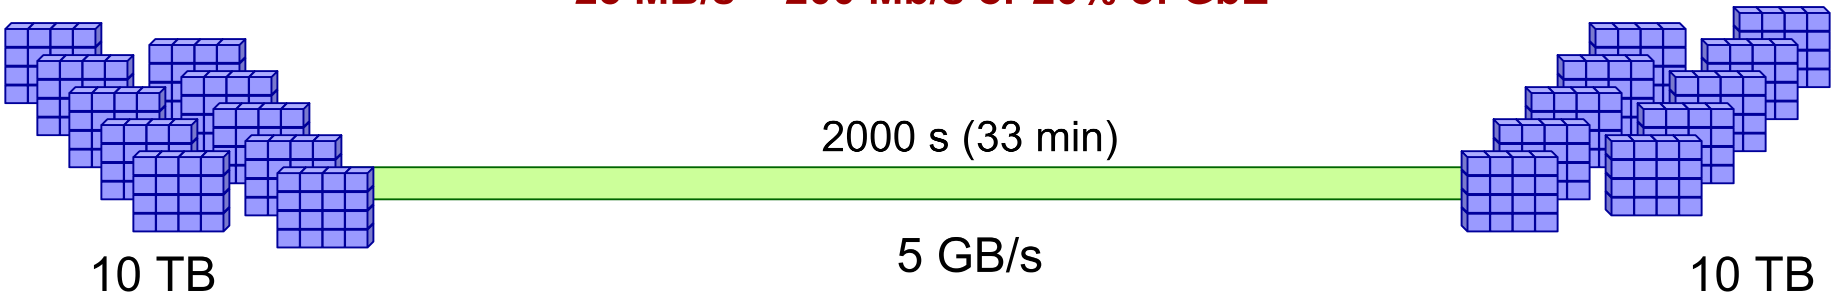

Initially Proposed Architecture

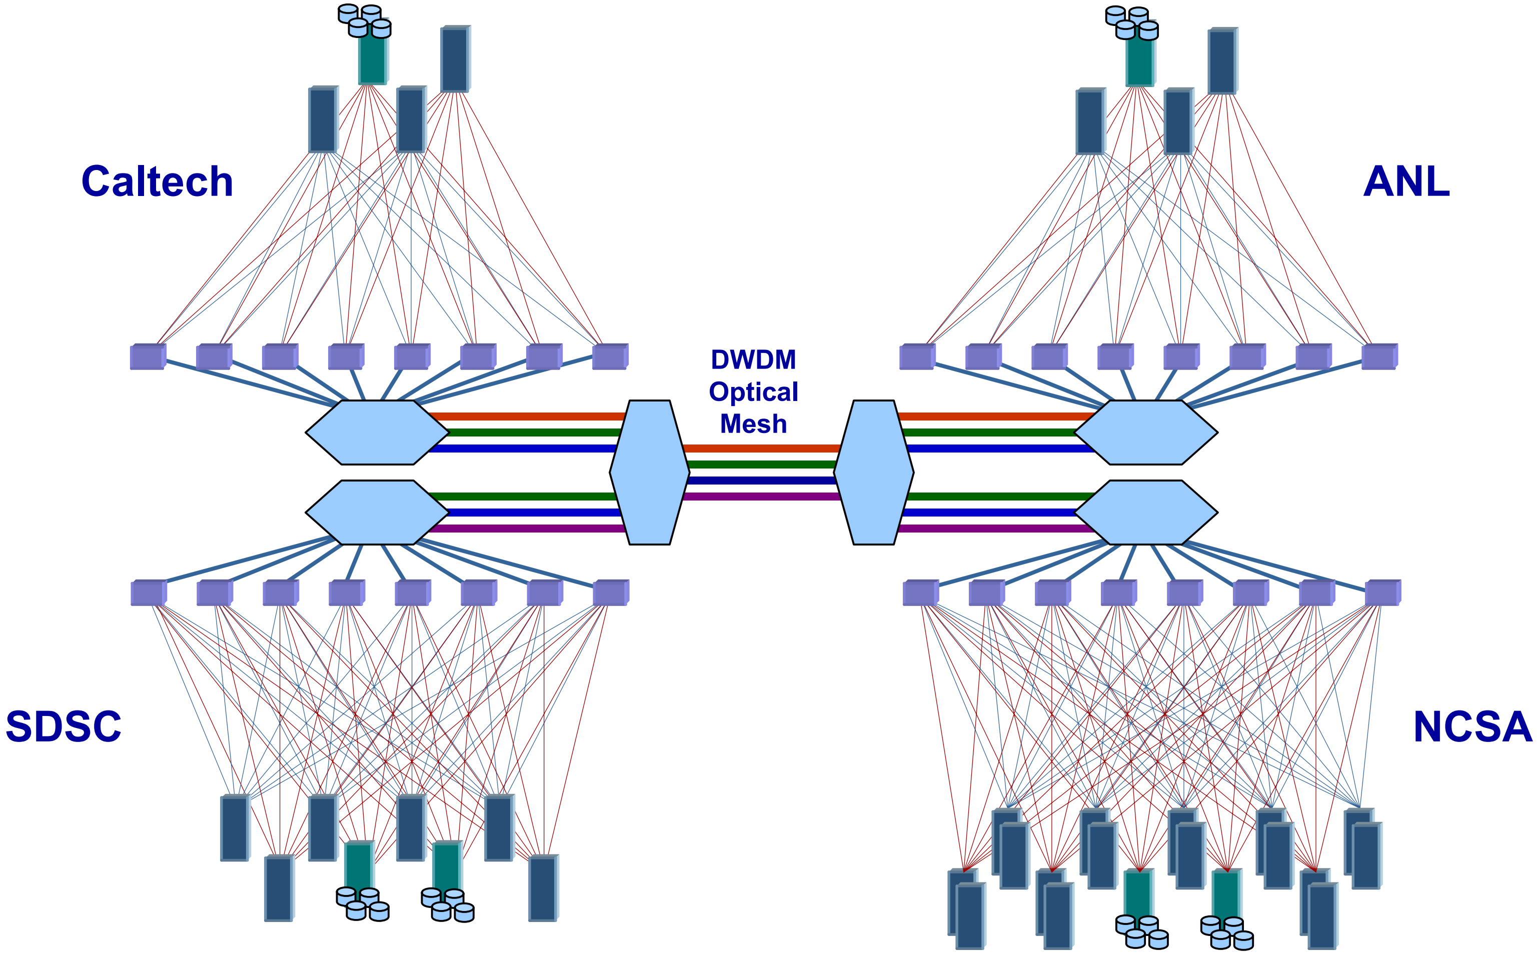

Initial Site Configuration

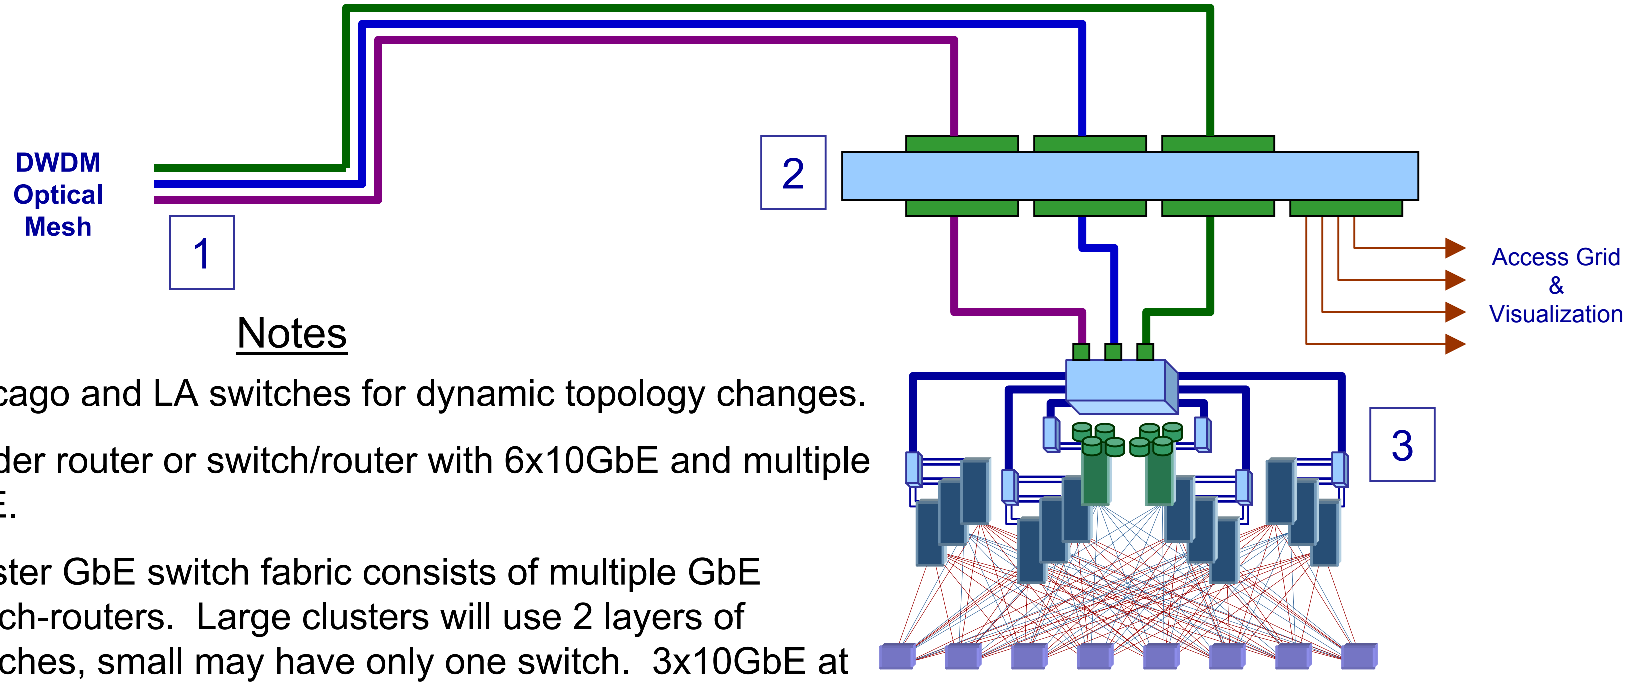

Notes

(1)Chicago and LA switches for dynamic topology changes.

(2)Border router or switch/router with 6x10GbE and multiple GbE.

(3)Cluster GbE switch fabric consists of multiple GbE switch-routers. Large clusters will use 2 layers of switches, small may have only one switch. 3x10GbE at the top, nx10GbE at the bottom going out to n bottom layer switch/routers. Bottom layer switch/routers have 10GbE at the top and nxGbE at the bottom, connecting to individual cluster servers.

Cluster-Backplane Interface

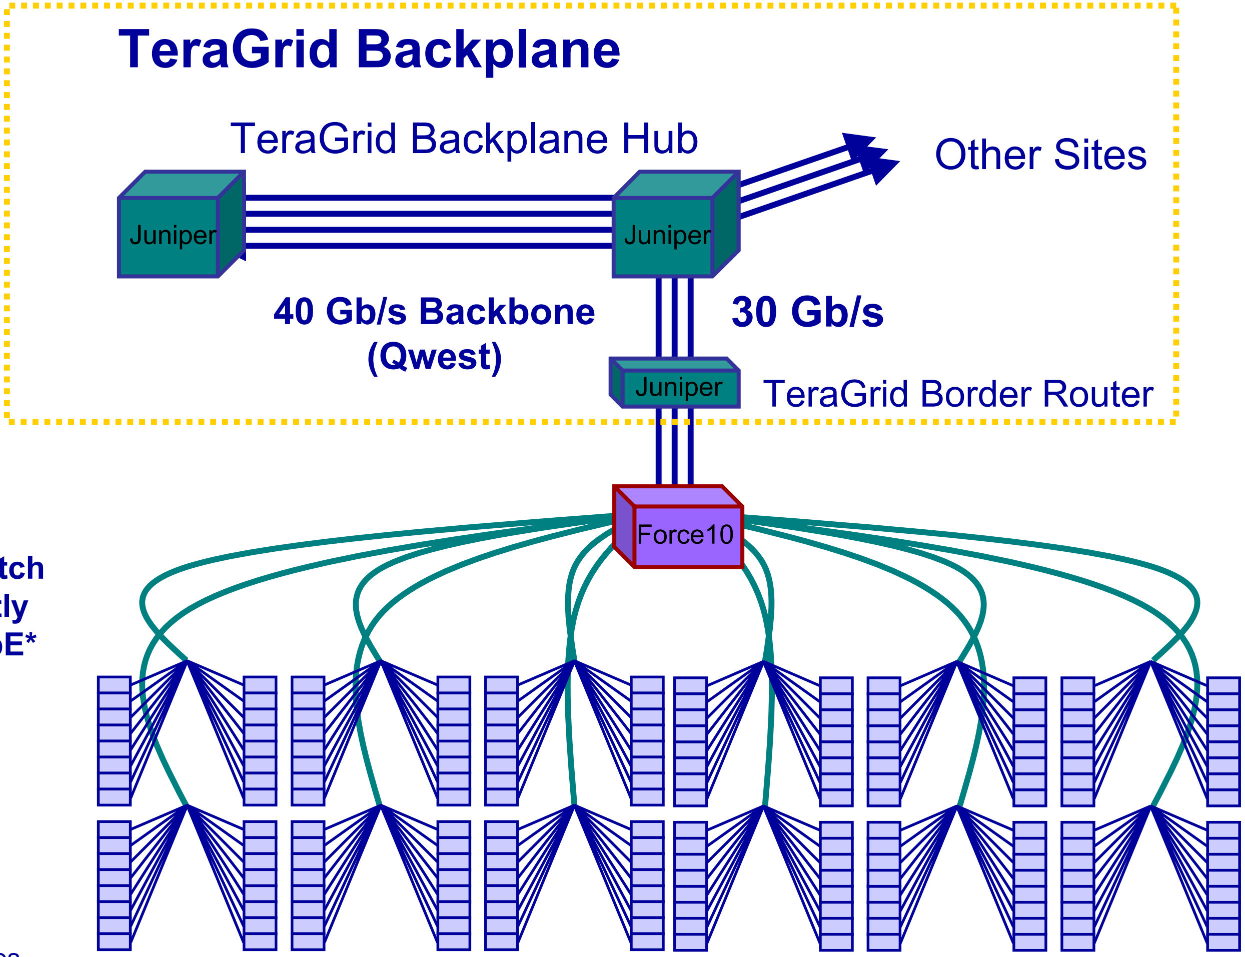

GbE/10GbE Aggregation Switch All nodes directly connected at GbE*

*NCSA Cluster uses 2 layers of Force10 switches

Intra-Cluster I/O via Myrinet (not shown)

Wide Area Network Topology

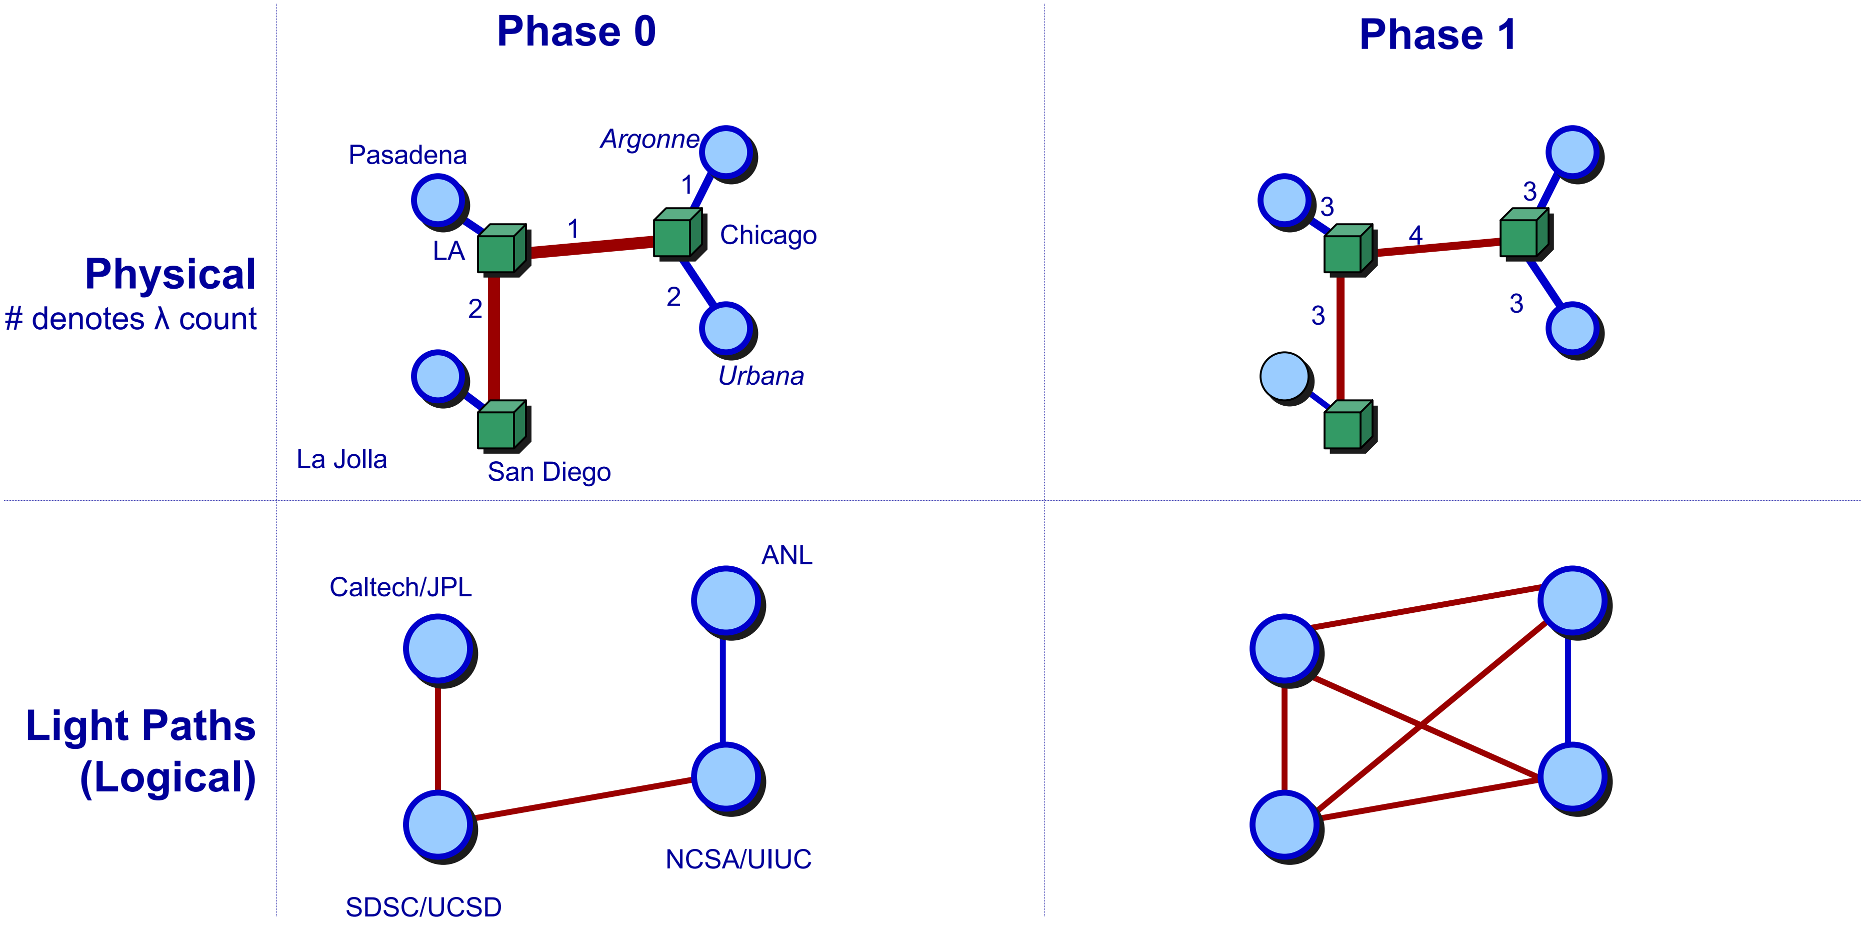

Phase 0

Phase 1

Physical

# denotes λ count

Light Paths (Logical)

TeraGrid Optical Network

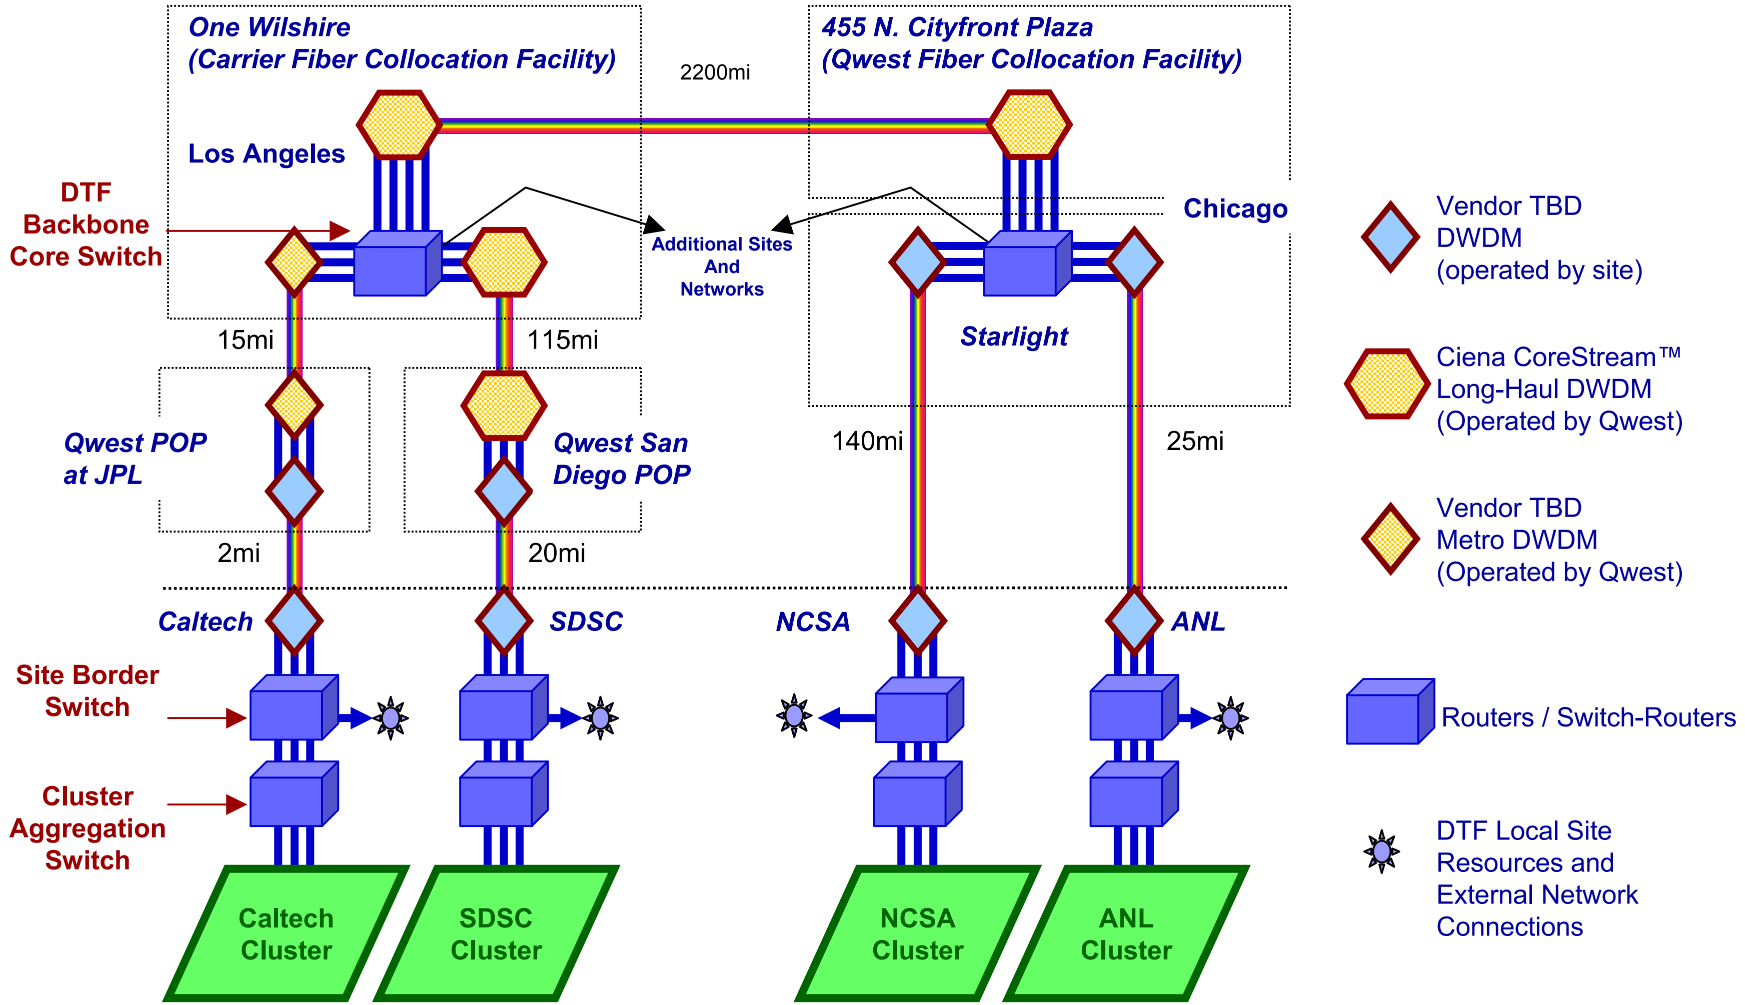
One Wilshire
(Carrier Fiber Collocation Facility)
455 N. Cityfront Plaza
(Qwest Fiber Collocation Facility)
2200mi
Los Angeles
Chicago
DTF Backbone Core Switch
Additional Sites And Networks
Starlight
15mi
115mi
140mi
25mi
Qwest POP at JPL
Qwest San Diego POP
2mi
20mi
Caltech
SDSC
NCSA
ANL
Site Border Switch
Cluster Aggregation Switch
Caltech Cluster
SDSC Cluster
NCSA Cluster
ANL Cluster
Vendor TBD DWDM (operated by site)
Ciena CoreStream™ Long-Haul DWDM (Operated by Qwest)
Vendor TBD Metro DWDM (Operated by Qwest)
Routers / Switch-Routers
DTF Local Site Resources and External Network Connections

Optical Network Cost Components

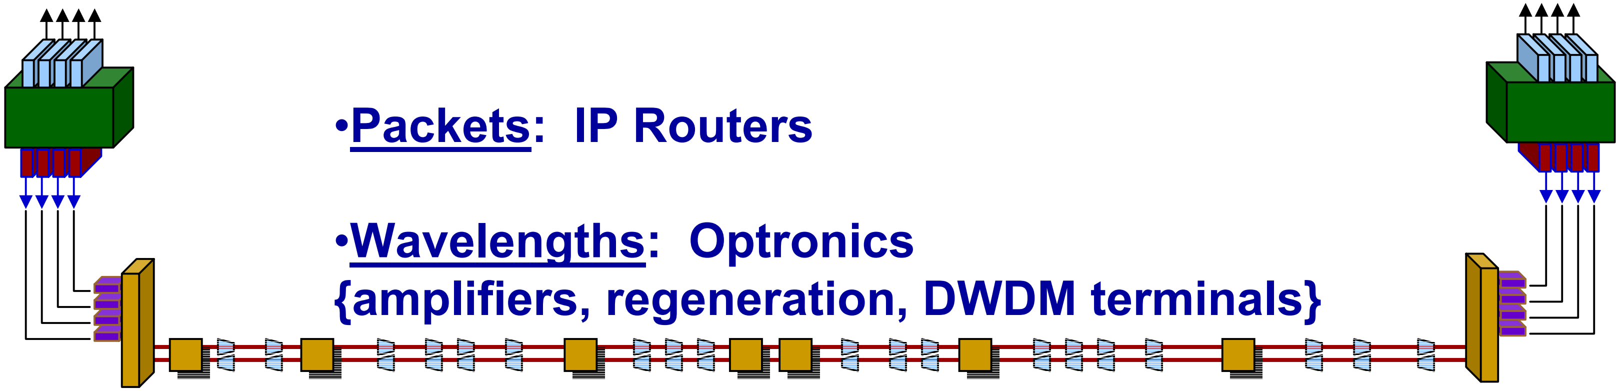

•**Packets**: IP Routers

•**Wavelengths**: Optronics {amplifiers, regeneration, DWDM terminals}

•**Medium**: Fiber and facilities

Optical Transport Costs

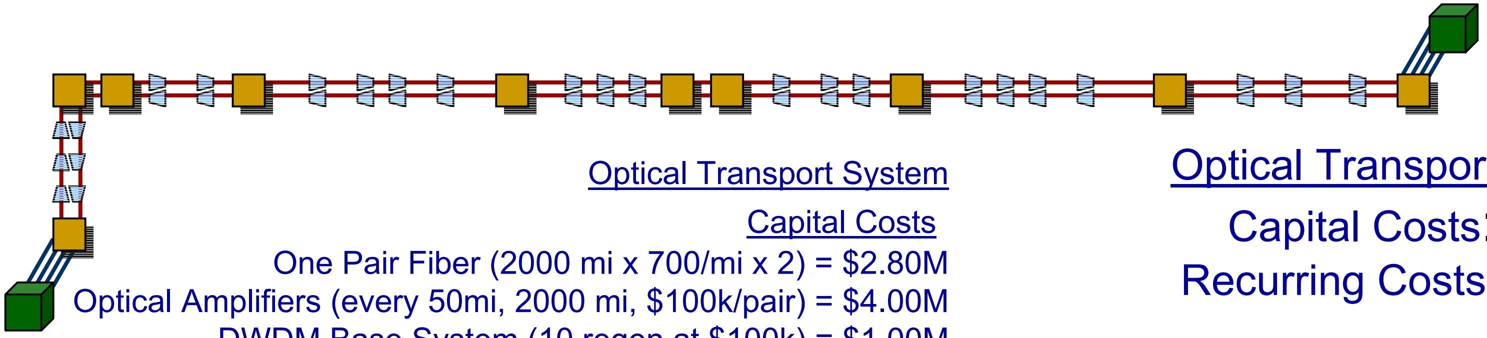

Optical Transport System

Capital Costs

One Pair Fiber (2000 mi x 700/mi x 2) = $2.80M

Optical Amplifiers (every 50mi, 2000 mi, $100k/pair) = $4.00M

DWDM Base System (10 regen at $100k) = $1.00M

Four 10Gb/s Lambda ($75k/transponder x 18 x 4) = $5.40M

Total for 4-lambda system = $13.2M

Recurring Costs

Fiber Maintenance (@$200/mi) = $0.20M

Equipment space & power ($10k, every 50 mi) = $0.40M

Amplifier and DWDM Maintenance (15%) = $1.15M

Total for 4-lambda system = $1.75M

Integration with LANs

Capital Costs

IP Routers (2, $400k) = $0.80M

WAN interfaces (8, $200k) = $1.60M

LAN interfaces (8, $200k) = $1.60M

Total for 4-lambda system = $4.00M

Recurring Costs

15% of capital = $0.60M

Optical Transport System

Capital Costs: $13.2M

Recurring Costs: $1.75M

Integration with LANs

Capital Costs: $4.00M

Recurring Costs: $0.60M

Total 4-Lambda Network

$17.2M, $2.35M/yr

🟦🟧 Full “3R” Regeneration (optical → electrical → optical)

▷ Optical amplifiers.

— LAN circuit.

— Strand of fiber

🟩 IP Router

I-WIRE

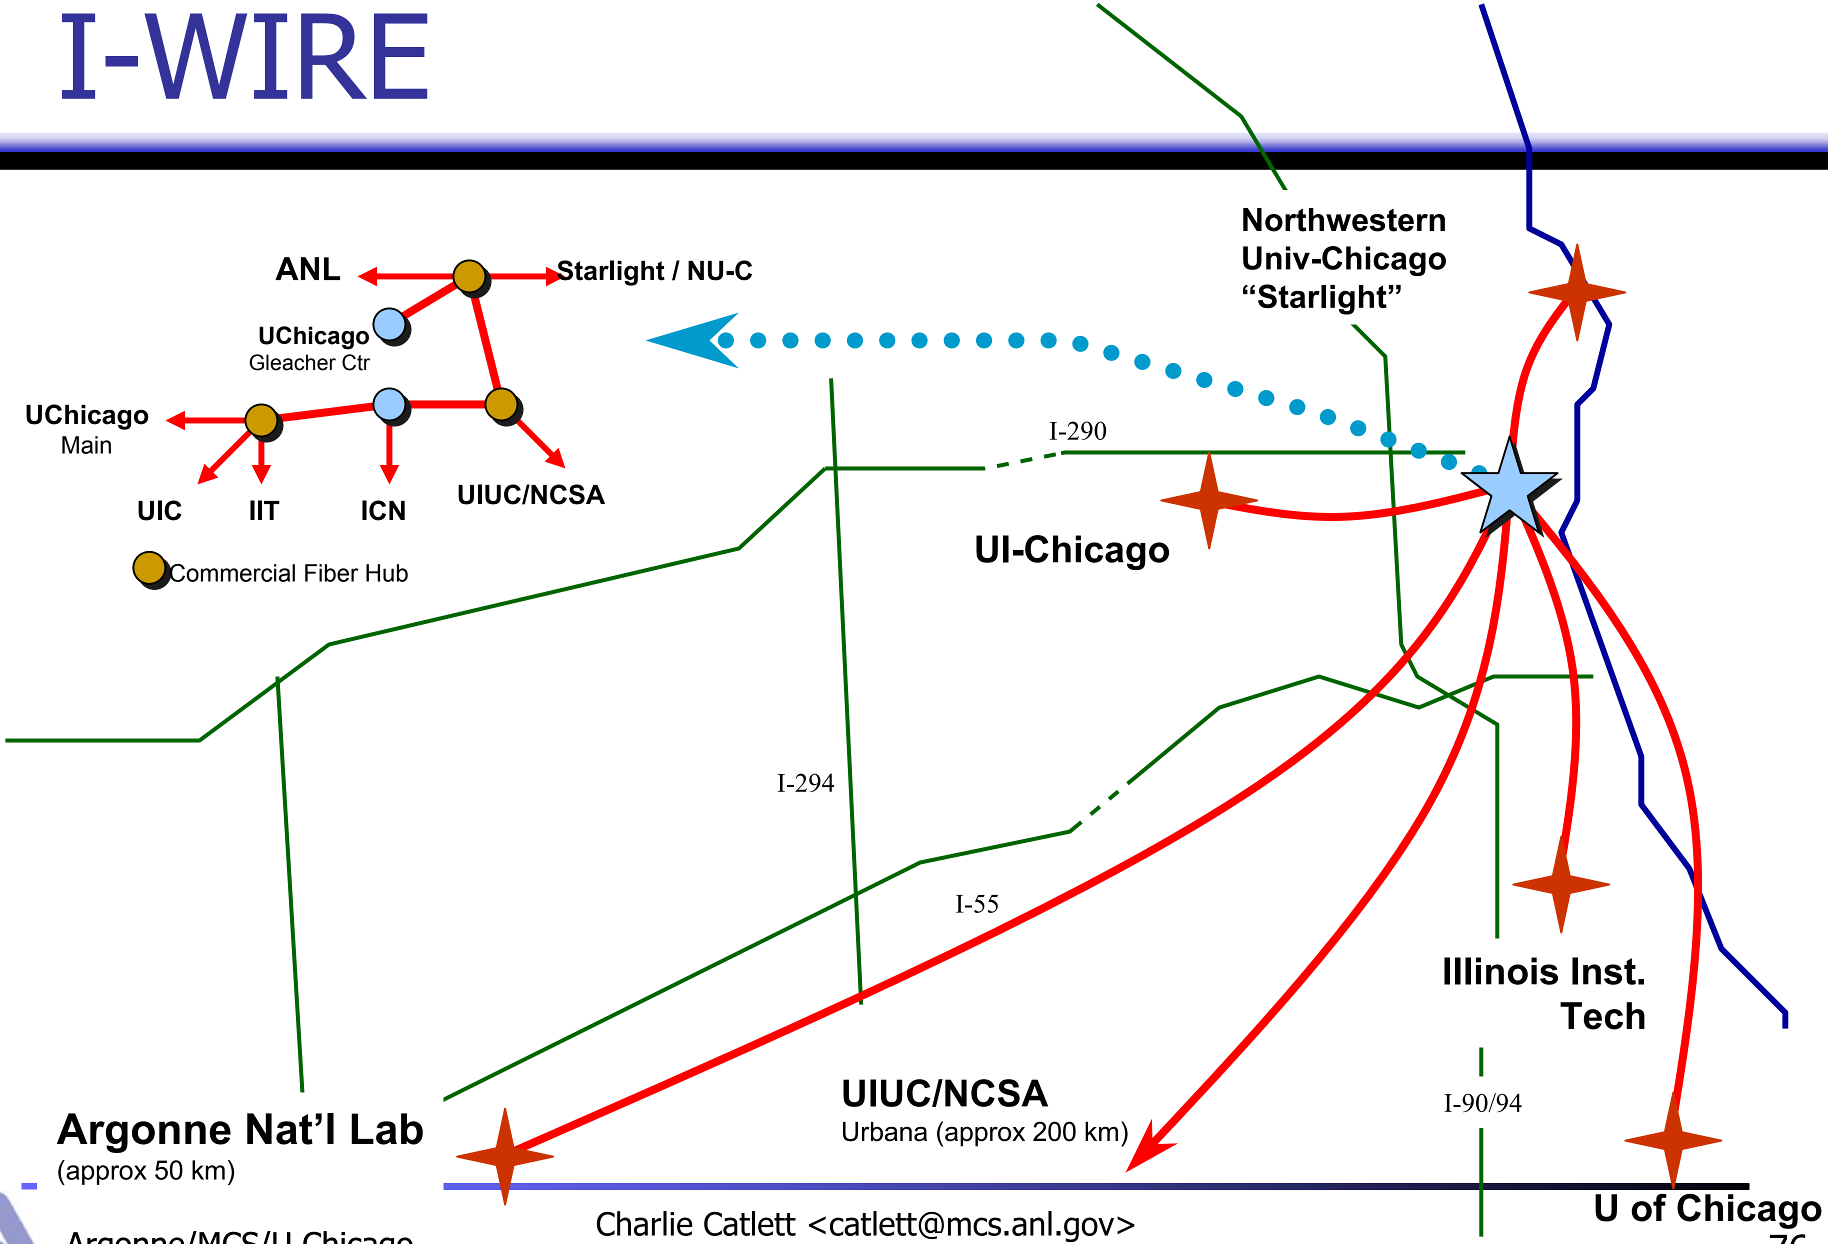

Charlie Catlett <catlett@mcs.anl.gov>
Pete Beckman <beckman@mcs.anl.gov>

I-Wire Fiber Topology

🟩 Qwest

🟧 McLeodUSA

🟪 Level(3)

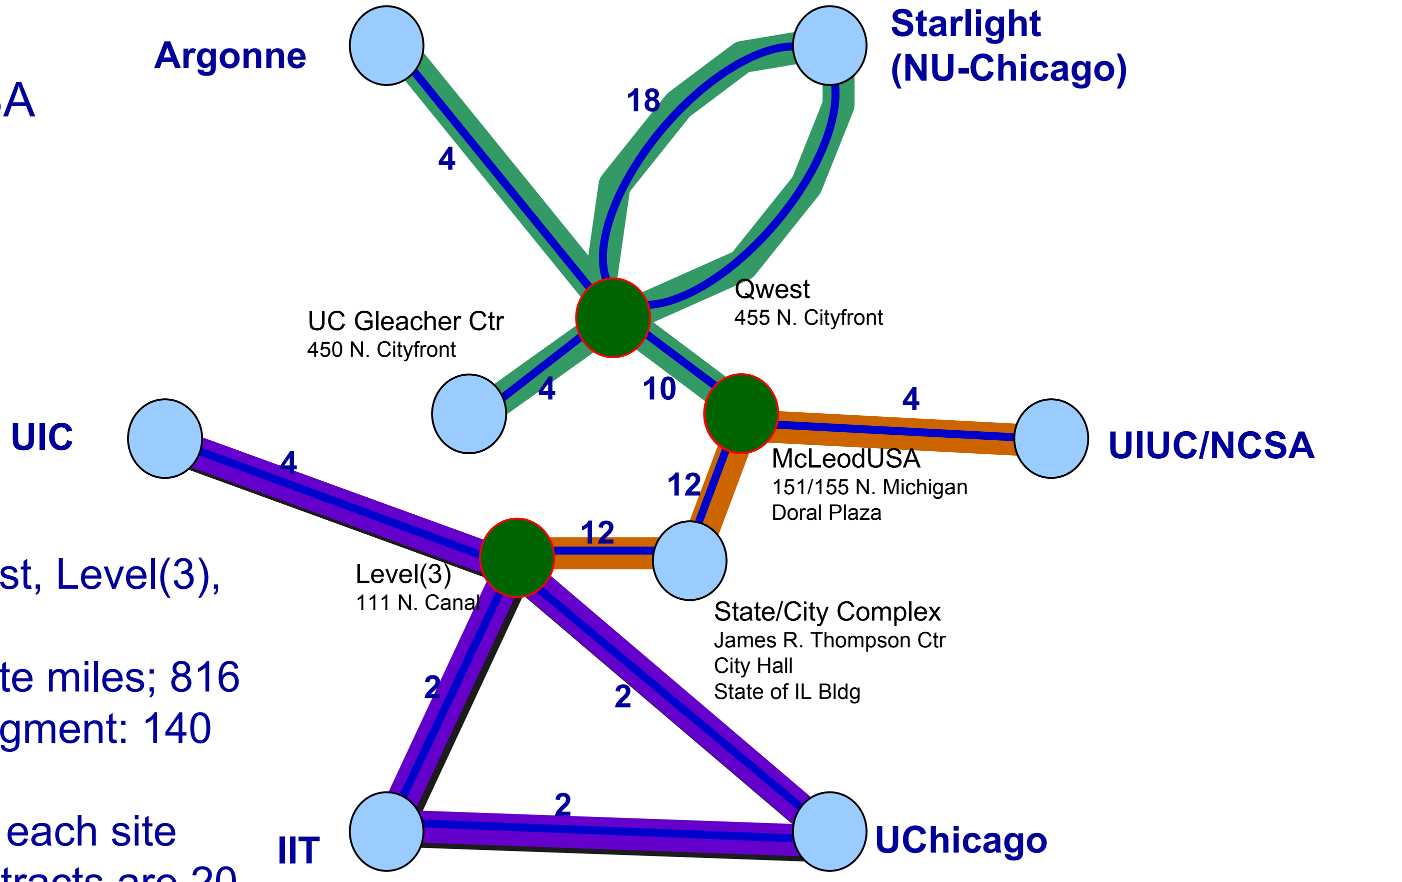

- Fiber Providers: Qwest, Level(3), McLeodUSA
- 10 segments, 190 route miles; 816 fiber miles, longest segment: 140 miles
- 4 strands minimum to each site
- ~$4M for fiber- all contracts are 20 year “IRU”

Numbers indicate fiber count (strands)

I-WIRE Transport

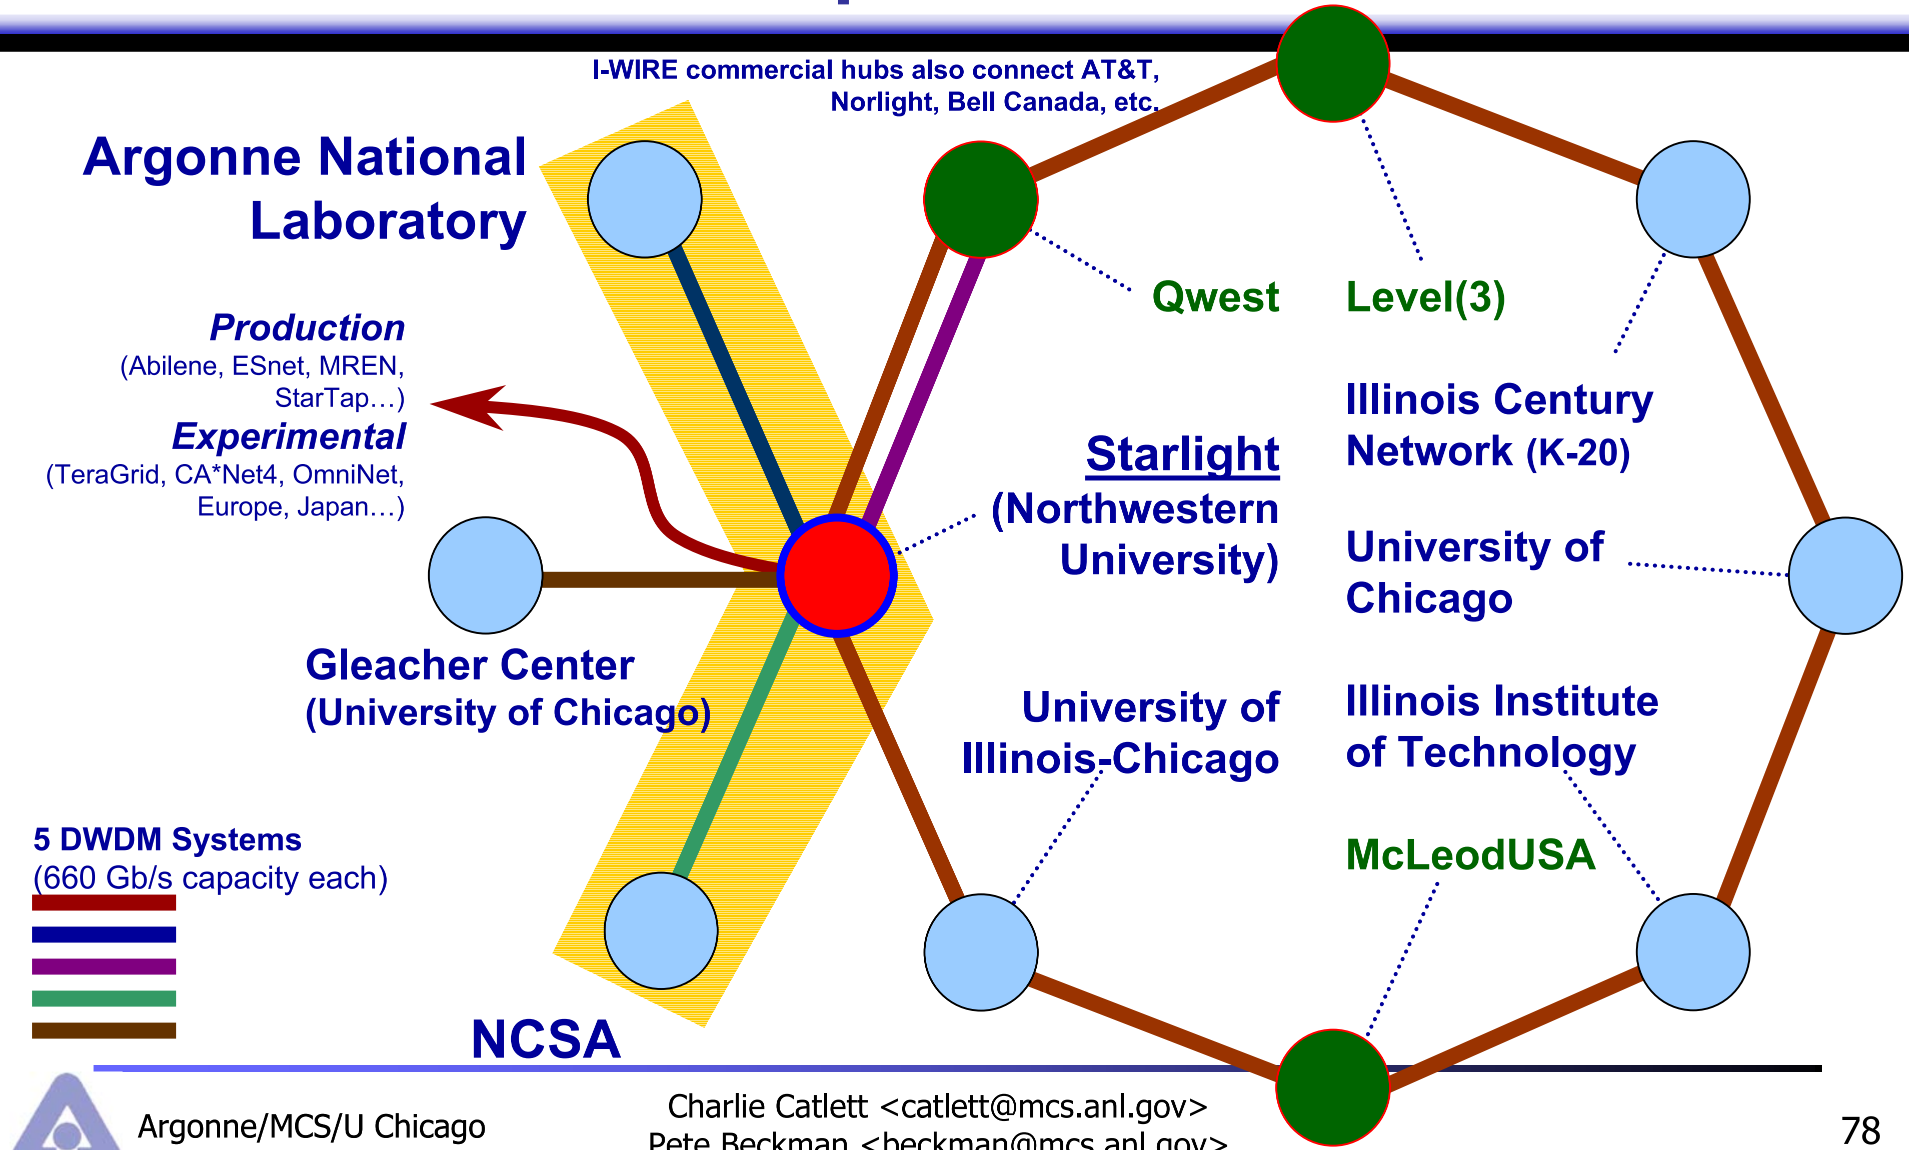

I-WIRE commercial hubs also connect AT&T, Norlight, Bell Canada, etc.

Argonne National Laboratory
Production
(Abilene, ESnet, MREN, StarTap…)
Experimental
(TeraGrid, CA*Net4, OmniNet, Europe, Japan…)
Gleacher Center (University of Chicago)
NCSA
Qwest
Level(3)
Starlight (Northwestern University)
Illinois Century Network (K-20)
University of Chicago
University of Illinois-Chicago
Illinois Institute of Technology
McLeodUSA

5 DWDM Systems
(660 Gb/s capacity each)

I-WIRE Costs

- Initial Costs
  - $4M for fiber
    - 20-year IRU Fiber (existing fiber) ($700 to $5,000/strand-mile)
    - New construction (for 'last mile') ($30 to $100/foot ($160-530k/mile))
  - $2.5M for equipment
    - Dual GbE channels over OC-48, one per campus
    - Three OC-192 ANL←→Starlight and NCSA←→Starlight
- Annual costs
  - $100k for fiber maintenance, equipment space/power
  - $300k for equipment maintenance, engineering staff
- Example of Potential Savings to the State of Illinois
  - NCSA: 622 Mb/s Urbana to Chicago, $50,000 per month
    - Replace with 2.5 Gb/s channel (4x capacity) using I-WIRE
      - $35k equipment, ~$k/mo maintenance

Original Plan: 10 GbE MESH

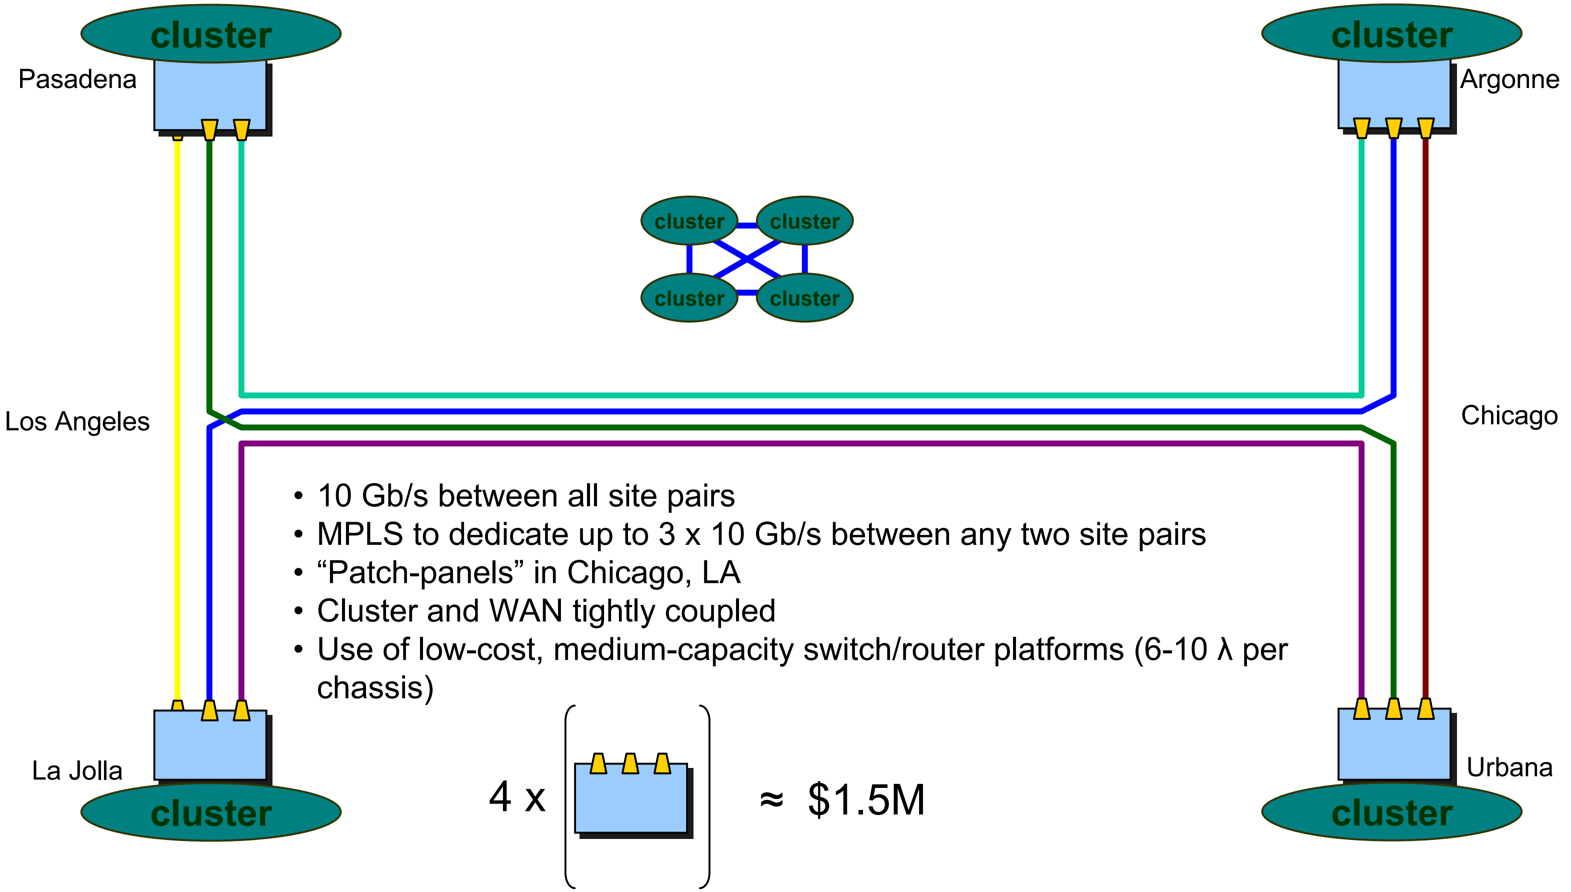

- 10 Gb/s between all site pairs
- MPLS to dedicate up to 3 x 10 Gb/s between any two site pairs
- “Patch-panels” in Chicago, LA
- Cluster and WAN tightly coupled
- Use of low-cost, medium-capacity switch/router platforms (6-10 λ per chassis)

Enhanced Scalable Network

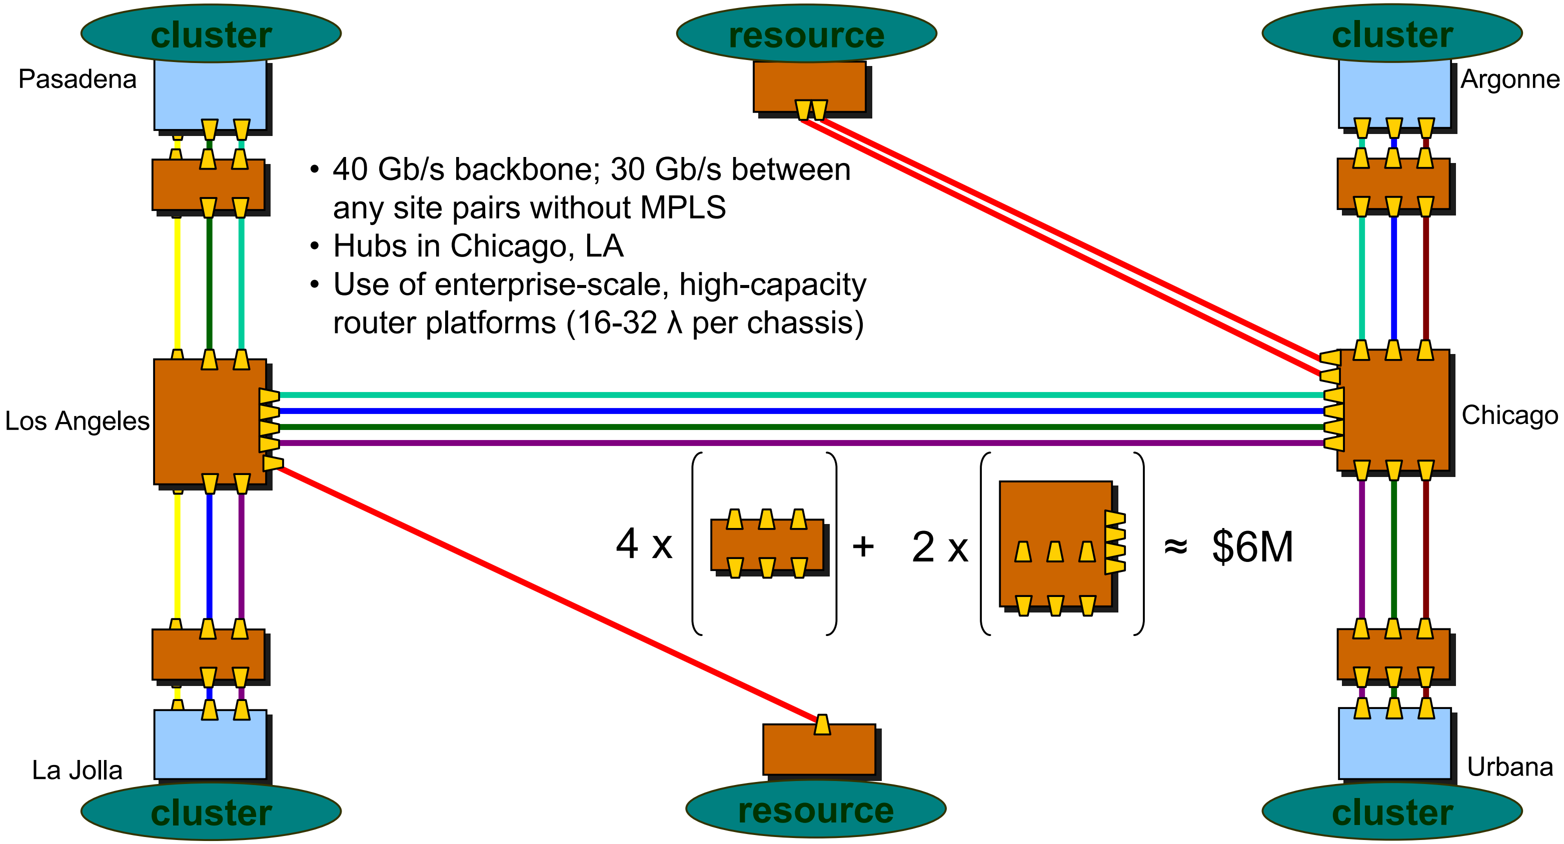

- 40 Gb/s backbone; 30 Gb/s between any site pairs without MPLS
- Hubs in Chicago, LA
- Use of enterprise-scale, high-capacity router platforms (16-32 λ per chassis)

Solution 1: Create connection opportunities at major fiber hubs, where lambda markets drive prices down.
Solution 2: Create hierarchical topology with core (backbone), hub, and spokes.
Solution 3: Create traditional site border router separate from internal cluster interconnect.

Economics

Architecture	Sites		Hubs		Total
	Chassis	Interfaces	Chassis	Interfaces	
Full Mesh	Part of cluster	3x $125k = $375k/site	n/a	n/a	$1.5M
Backplane	1x $100k = $100k/site	6x $125k = $750k/site	1x $100k = $100k/hub	10x $125k = $1.25M/hub	$6.1M

Scaling: Mesh → Backplane

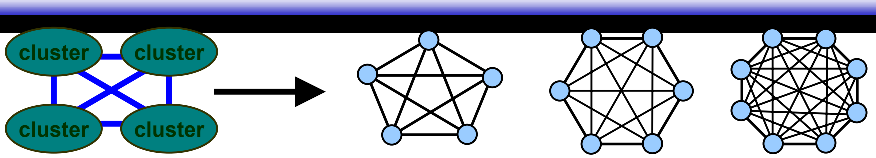

- Connection Points
  - Connection points to the mesh are only at the sites, not major fiber hubs.
    - No cost-effective way to provide a wavelength to sites.
- Scaling a Mesh
  - Core site n+1 adds n lambdas to maintain full mesh. 5th site adds 4, 6th site adds 5, 7th site adds 6…
- WAN/Cluster Coupling
  - Tightly coupled WAN/Cluster relies on tightly coupled hardware/software management and homogeneity. Additional sites with various resources loosely coupled will require decoupling WAN from Clusters.

Tracking Network Deployment

Site	Regional Fiber	Bandwidth to Hub	Border Router	Backplane to LAN	Cluster Connection
ANL	√	√ 3 x OC-192 (to Chicago)	√	√ 3 * 10GbE	√
Caltech	√	DWDM install 4/18 (to LA)	√	Tested	√
NCSA	√	√ 3 * OC-192 (to Chicago)	√	√ 3 * 10GbE	√
SDSC	√	√ 3 * OC-192 (to LA)	√	√ 3 * 10GbE	√
PSC	√	**500 Mb/s Interim** 1 x OC-192 (6/03) 3 x OC-192 (10/03)	Hardware Selected	Hardware Selected	Evaluating Prototypes

Status as of April 2003 (Caltech completed May 2003)

Hub Operations and Policy

- Juniper Connections
  - Available Hub Router Capacity
    - Chicago: 17 x 10 Gb/s
    - Los Angeles: 20 x 10 Gb/s
- Hub Policies
  - Minimum 10 Gb/s connections
    - Otherwise Abilene likely to be sufficient
    - <10 Gb/s wastes limited slot capacity at hubs
- Establishing 10 Gb/s Abilene interconnect
  - Connecting sites pay for hub access & hardware
  - Interim PSC connection via Abilene provides proof of concept

Backplane Policies

- Resources directly connected (i.e. at layer2) to border routers
  - No intervening LANs, firewalls, etc.
  - Implies one site per connection to hub
- New Hubs should extend backplane at full rate
  - 40 Gb/s minimum for new hubs to preserve backplane uniformity

Backplane Operations Policies

- Initial (ramp-up) Policies
  - Use of individual lambdas for experiments as warranted
    - E.g. recent storage over SONET experiments
  - Experimenting with MPLS for reserved bandwidth
- Steady State Policies
  - Backplane typically will operate at 40 Gb/s
  - Option to use 10 Gb/s (25% resource) for experiments
    - Available for scheduling, time granularity tbd based on experience
    - Additional hardware and funding will be required for equipment capable of interconnecting to lambda

TeraGrid Security

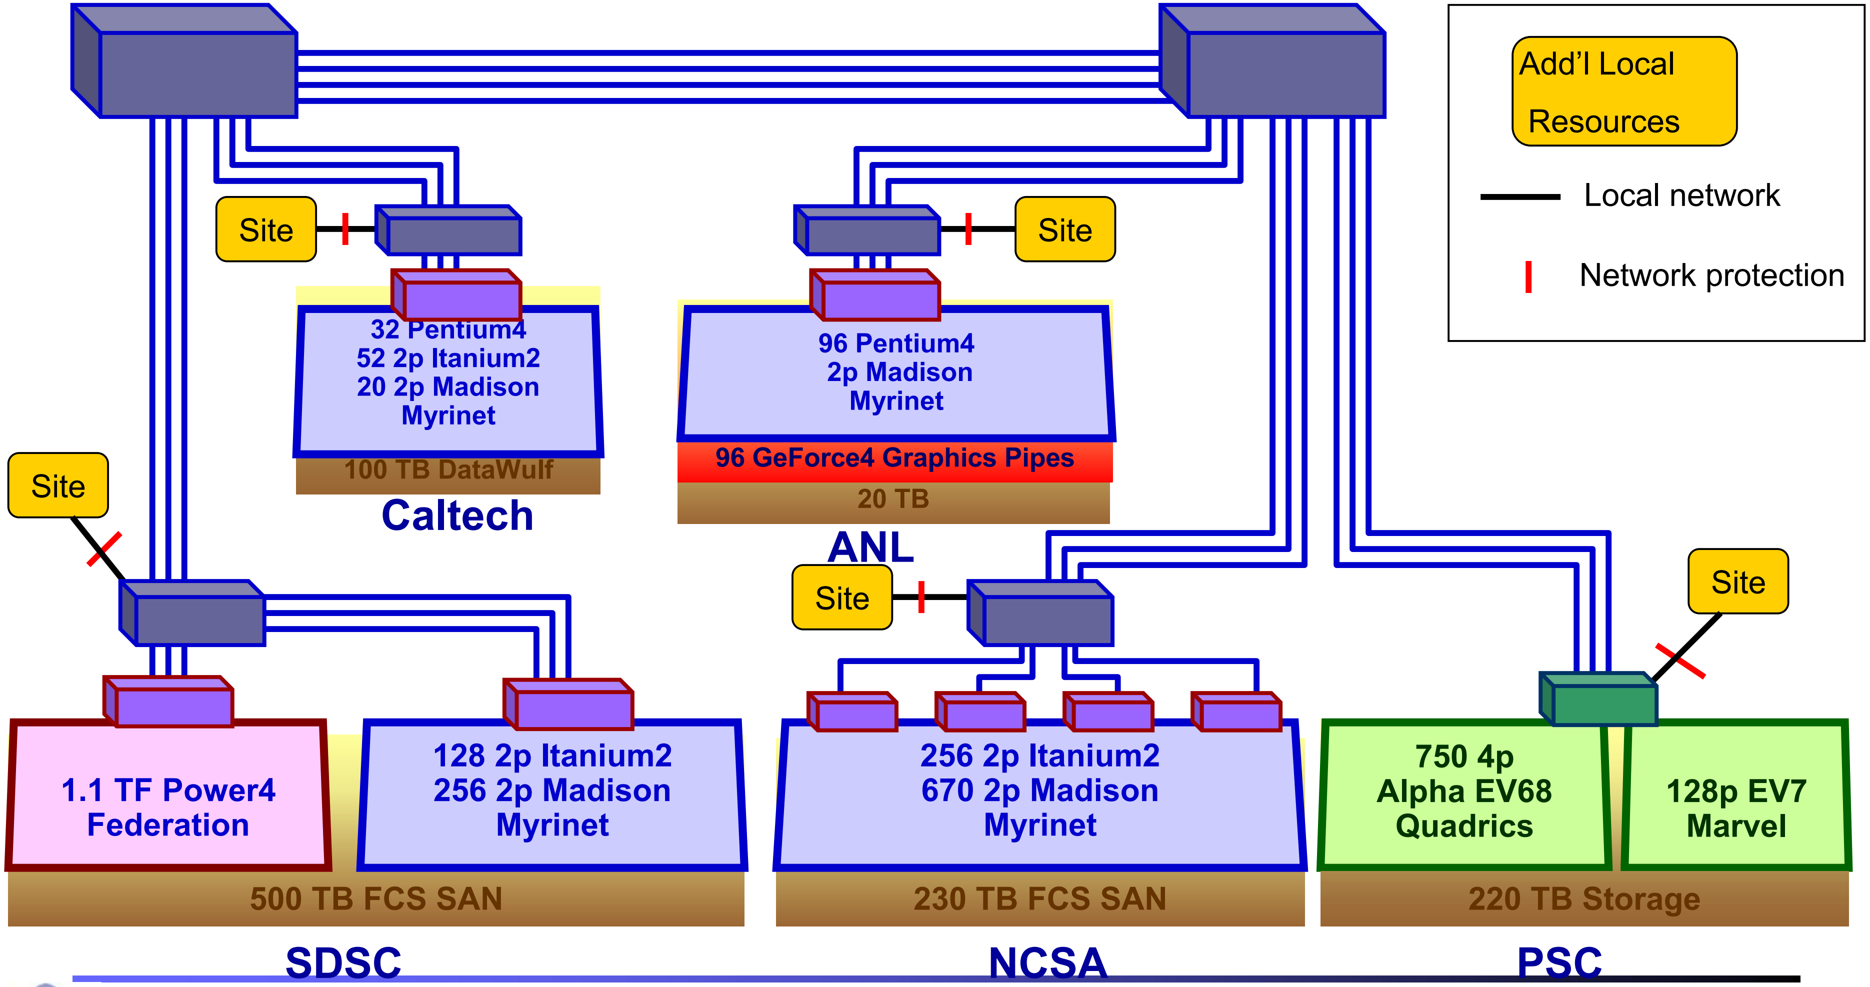

Chapter 4
Construction Phase

Begin Software Development

- During the design phase, you should have determined:
  - Freely available software
  - $$$ SW/support
  - Software that requires development
- Two implementation teams, two kinds of people, two different cultures
  - “Classic sysadmin”
    - Install freely available or $$ SW
  - “Programmer”
    - Comfortable with CVS, design specifications, design reviews, user interface design

Culture Clash Warning…

- Sysadmins:
  - Totally demand driven: “Which emergency do I handle first today?”
  - “I don’t have time for software engineering”
  - “I wrote this really ugly, un-maintainable, 500 line Perl program to temporarily solve our problem, isn’t it so cool!”
  - “Version control? I can remember all the details, I don’t need a version system, and besides, I don’t have time to learn that and maintain it”
  - Underestimate time: “A weekend of hacking should do it”
- Programmers:
  - Comfortable with timelines and budget
  - Understand project plans, technical reviews, documentation, etc
  - Live and die by versioning and build systems
  - Think sysadmining is easy, even though they have never done it
  - Overestimate time: “32.9 man-years to create the specification”

REPRODUCABILITY!!!!

- To maintain the Service Level Agreement and consistent software stack (CTSS) requires versioning and reproducible results
- **The TeraGrid continues to loose this battle**
- SysAdmins:
  - “I know secret magic to setup and configure the machines, and it is so complicated it would be nearly impossible to automate, and besides, I don’t have time”
  - Therefore, SysAdmins will spend most of their time painfully duplicating, by hand, and with human errors, their work

Reproducibility: What We Need

- ALL source code, patches, and build scripts go into repository
    - Examples:
        - 1.2.5..9-gcc, 1.2.5..9-intel, mpich-1.2.5..10-gcc, mpich-1.2.5..10-intel
- Automated (non-human) steps for:
    - Configuring source (patches, config), building source, building package, installation, configuration, post-installation
- Automated process for rebuilding all packages
- (similar to how Linux companies build a distro)
- For many (most?) sysadmins, they have never used CVS (or the discipline of versioning) on a daily basis

Make the Repository Part of the Culture

- If it is not in the repository, it does not get included in status reports or press releases
- Code cannot be installed on machines unless it is in the Repo first
- PPT slides go in the repository
- Web pages are checked into the repository
- TortoiseCVS is used on Windows desktops

Testing: Part of the Culture

- All programs that are part of the software stack must have:
  - Version Reporter (for automated collection)
  - Unit Test (the Self Test for packages)
  - Integrated Test (test that confirms proper integration with the other system components)
- Automated test harness
- Historical archive of test results for mining
  - Comparing performance
  - Finding components that fail often

Grid-wide Change Management

- Set up TestGrid / development platforms
- Changes to production machines must be handled by Change Management process
  - Functional on TestGrid
  - Rationale for upgrading production machine
  - Test on production machine by sysadmins
  - Trial by limited number of users
  - Move to full production status

Project Plans and Task Lists

- Create long-range plans and milestones for entire project
- Create detailed task lists for 8 – 10 week horizon
- Working groups (teams) must be responsible for weekly reporting on their milestones
- Weekly "architecture" meetings used for:
    - Discussing the 8-10 week detailed tasks
    - Synchronization and dependencies between groups
    - Failures in the system

Technical Reviews and All Hands Meetings

- Plan for frequent mini-reviews of architecture, status, and technologies (8 – 10 week horizon)
- Quarterly careful review of technical achievements and designs.
  - Results used for quarterly reports
    - Weekly reports very helpful
  - Risk assessment and go/no-go decisions
- Semi-annual "All Hands Meeting" to synchronize all the Grid participants

Chapter 5
Operating the Grid

Unified Support

- Single 1-800 number that moves between sites
  - 12 hrs at NCSA
  - 12 hrs at SDSC
- Single page for requesting allocation
- Unified trouble tix system
- Working groups that can work directly with users:
  - Performance Evaluation
  - Interop
  - Networking

Portals

- Simplified mechanisms for single-point status
  - Current accounting
  - System status (as needed by user)
    - Network
    - Queue
    - Projected downtime
- Domain-specific portals for job submission, mass storage, etc, can be built by their respective communities
  - We are a Grid Hosting Environment, we don't want nor need to build user-level portals… they can do it!

Operations Monitoring

- There are many types of status info users:
  - Developers
  - Users
  - Operators
- Operations must construct set of status monitors from basic tools to meet their needs:
  - Clumon / Ganglia
  - Inca
- Set of actions and procedures to take based on status information
- Planning for uptime / downtime
  - Users are already frustrated with TeraGrid…

User Interactions

- Project-wide training team
  - Creating SC03 Tutorial
- Early user workshops
- Support for Flagship Applications

Chapter 6
Summary / Questions

Summary: Day 1

- Virtual Organization & Cooperative Agreement
- Management and Decision Making
  - Exec Dir, Dir of Eng, Team Leaders, etc
- Design -> Build -> Operate
- Eliminate hype, be practical, get user commitment early
- Collaboration technology required from day 0
- Installing and understanding Globus is the least of your concerns
- Integration and engineering is costly, budget for it

Job #1:
Create Virtual Organization for Participants

Single, Distributed Team

Software

- Testing, QA, Verification
- Software Stack
- Advanced Networking
- Data Services
- Viz Services
- Accounting SW

Developer’s Org

- Exec. Mgmt Committee
- Directors (Exec & Eng)
- Site Leads
- Working Group Leads
- Engineers

User-visible Org

- Operations Center
- User Services
- Accounting

It is not your code, it is our code

It is not your doc, it is our doc

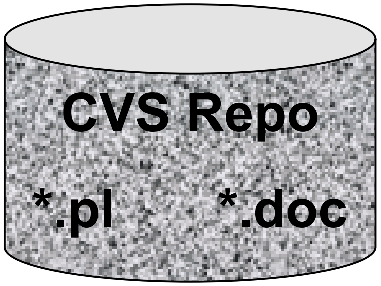

One repo to rule them all,
one repo to find them
one repo to bring them all
and in the grid-world bind them

Collaboration Structures:
(strange: most are not Grid based…)

Repository and version control:	CVS
Problem tracking:	Bugzilla
Mailing lists:	Mailman
Web page authoring:	Wiki
Document Library:	homemade
Project Plan:	MS Project
Real-time conferencing:	Access Grid
Software Repository:	CVS

- Integrated Solutions?

Working Groups

- User Services
- Applications
- Accounting
- Networking
- Visualization
- Data
- Testing / InfoServ
- Perf Eval
- External Relations
- Grid
- Security
- Operations
- Software / Clusters
- Interoperability

Design & Operation Document Guide

- Cooperative Agreement
- Mgmt & Org Chart
- User Survey
- Strawman user guide
- Flagship apps
- Project Plan
- Security Policies
- Elevator Overview
- QA / Test / Accept Plan
- Primer
- Risk Mgmt Plan
- Accounting practices & policy
- Data practices & policy
- Customer service: Help desk, trouble tickets, SLAs
- Work and task list
- Real user guide

Summary: Day 2

- Networking:
  - Dedicated backplane
  - Snowplow utilization
  - Policies: Security, ownership & mgmt of routers
  - Fiber required for 10Gb/s+
    - dark fiber, lease, IRU
- Construction
  - Reproducibility
  - Sysadmin & Soft Eng cultures
- Operations
  - Training, running, support (it must be easy!)

The End