

Grids at Indiana University

Scott McCaulay

IU TeraGrid Site Lead

INDIANA UNIVERSITY



Grid Activities at IU

- TeraGrid Resource Partner
- Open Science Grid – Atlas Tier 2 Site
- Hydra Portal/Condor Pool
- IVDGL Grid Operations Center (iGOC)
- Indianapolis/Bloomington intercampus grid
- Environmental, political factors that shaped the development of grid activities at IU?

One IT Department Serving Two Research Communities

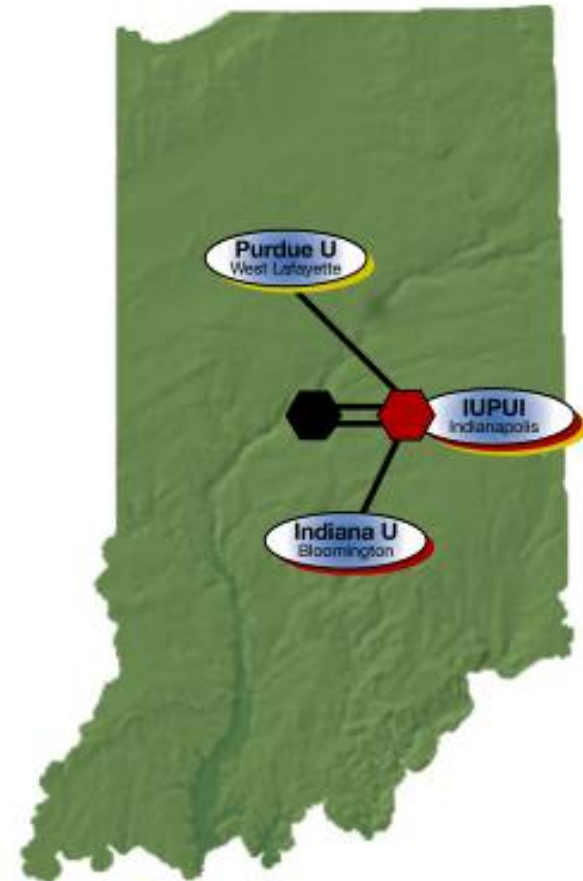
- Bloomington campus
 - Astronomy, Biology, Computer Science, Chemistry, Physics...



- Indianapolis campus
 - IU School of Medicine
 - Indiana University-Purdue University Indianapolis (IUPUI)

The I-Light Network

- \$5.3 million appropriation from the state of Indiana in 1999, operational in 2001
- Indiana became the first state to deploy its own high-performance data network
- Connects IU Bloomington, IUPUI Indianapolis and Purdue West Lafayette campuses
- Opened up a host of new opportunities...



AVIDD Cluster

(Analysis and Visualization of Instrument-Driven Data)



- First geographically distributed Linux cluster to top 1 TFLOPS on Linpack
- Nodes communicate through a Force 10 gigabit Ethernet network
- Indianapolis and Bloomington clusters connected via dual 10-gigabit pipeline through I-Light
- Funded in part by a \$1.8 million grant from the NSF
- Aggregate peak computing power of 2.2 TFLOPS

IBM Research SP

- IBM Scalable POWERparallel system geographically distributed between Bloomington and Indianapolis
- Provided a computing capacity of approximately 1 teraFLOP
- Funding for the SP system and its expansions came from Lilly Endowment, Inc. and from IBM SUR grants
- Supported over 900 researchers on both campuses, partially retired as of September 30,2005



Massive Data Storage System


- HPSS (High Performance Storage System)
- Automatic data replication between Indianapolis and Bloomington, via I-Light
- 720 TB current tape capacity, total capacity of 2.2 PB
- >250 TB currently in use;
>5 TB for biomedical data
- Will be available to TeraGrid users via HSI, eventually via gridFTP



Photo by Tyagen Miller, © Trustees of Indiana University

Uses of Linked Computational Resources

- The I-Light network has been used in a Purdue-led simulation of the Sept 11 crash of a jet into the Pentagon
- Global Grid - HPC Challenge award at the SuperComputing 2003 conference
- Support for Center for Computational Homeland Security
- Ongoing cycle sharing agreement, IU-Purdue

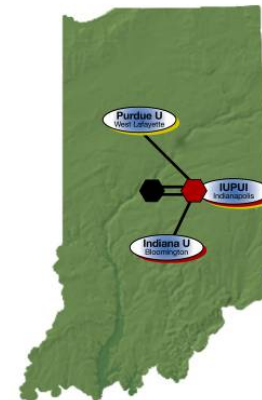


Grids at Indiana

Distribution of Systems Across Two Campuses

- A potential weakness becomes a strength for grid computing at Indiana University
- Disaster resiliency, AVIDD & SP users on both campuses can still be supported if one machine room goes out
- Automated data replication between Indianapolis and Bloomington
- Spread machine room load, some relief from power and cooling headaches
- Additional opportunities at regional campuses with I-Light 2

Chicago



Gary



Purdue West Lafayette



Indianapolis



Richmond

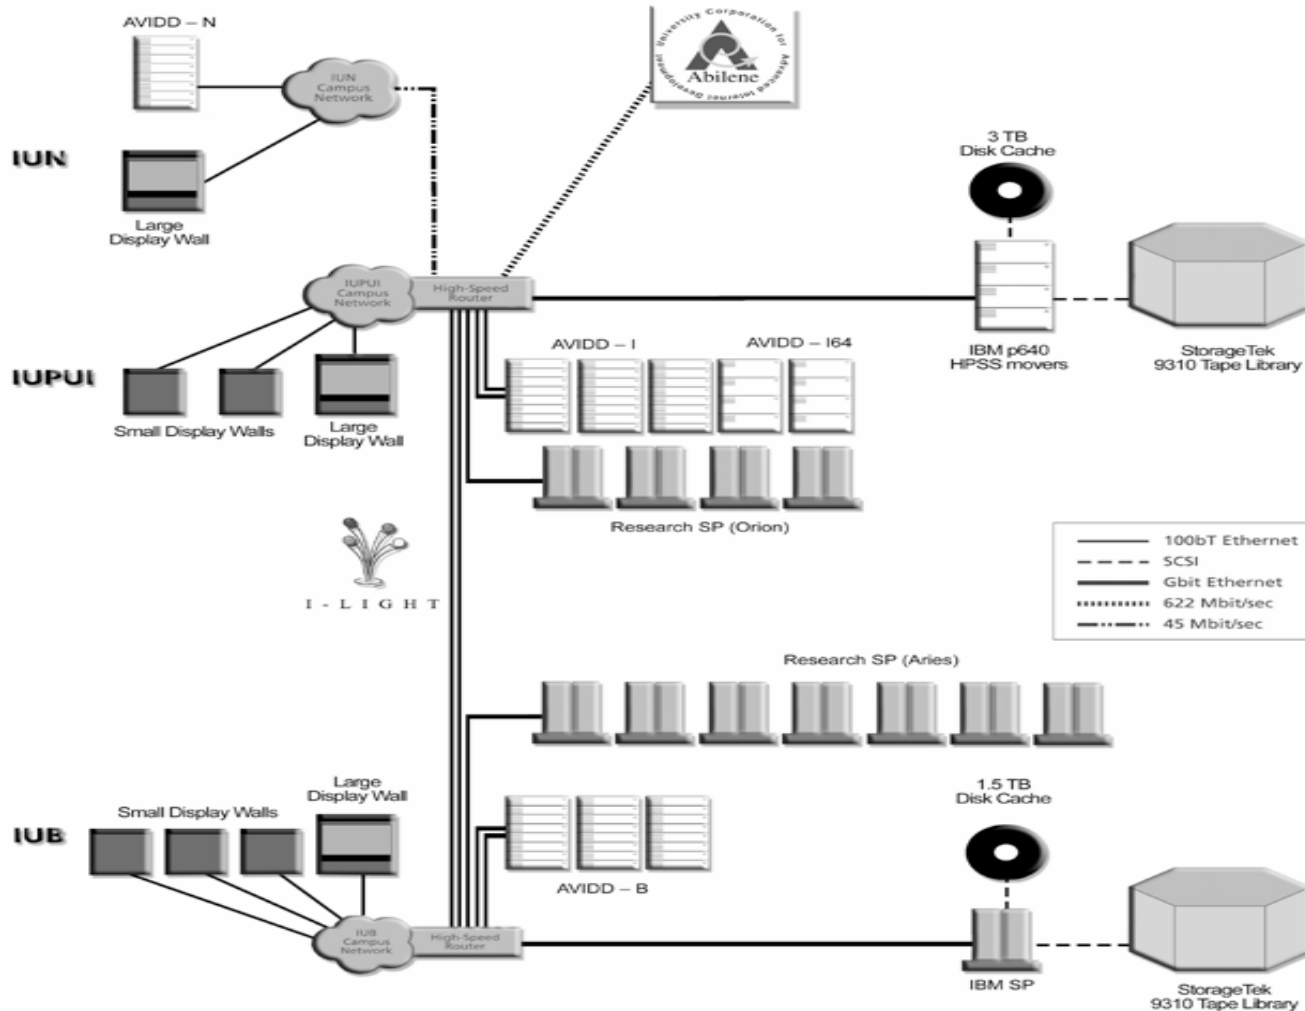


1 TFLOPS SP – INGEN, IBM
2 TFLOPS AVIDD – NSF, IBM
2.2 PB Storage – STK, NSF, INGEN

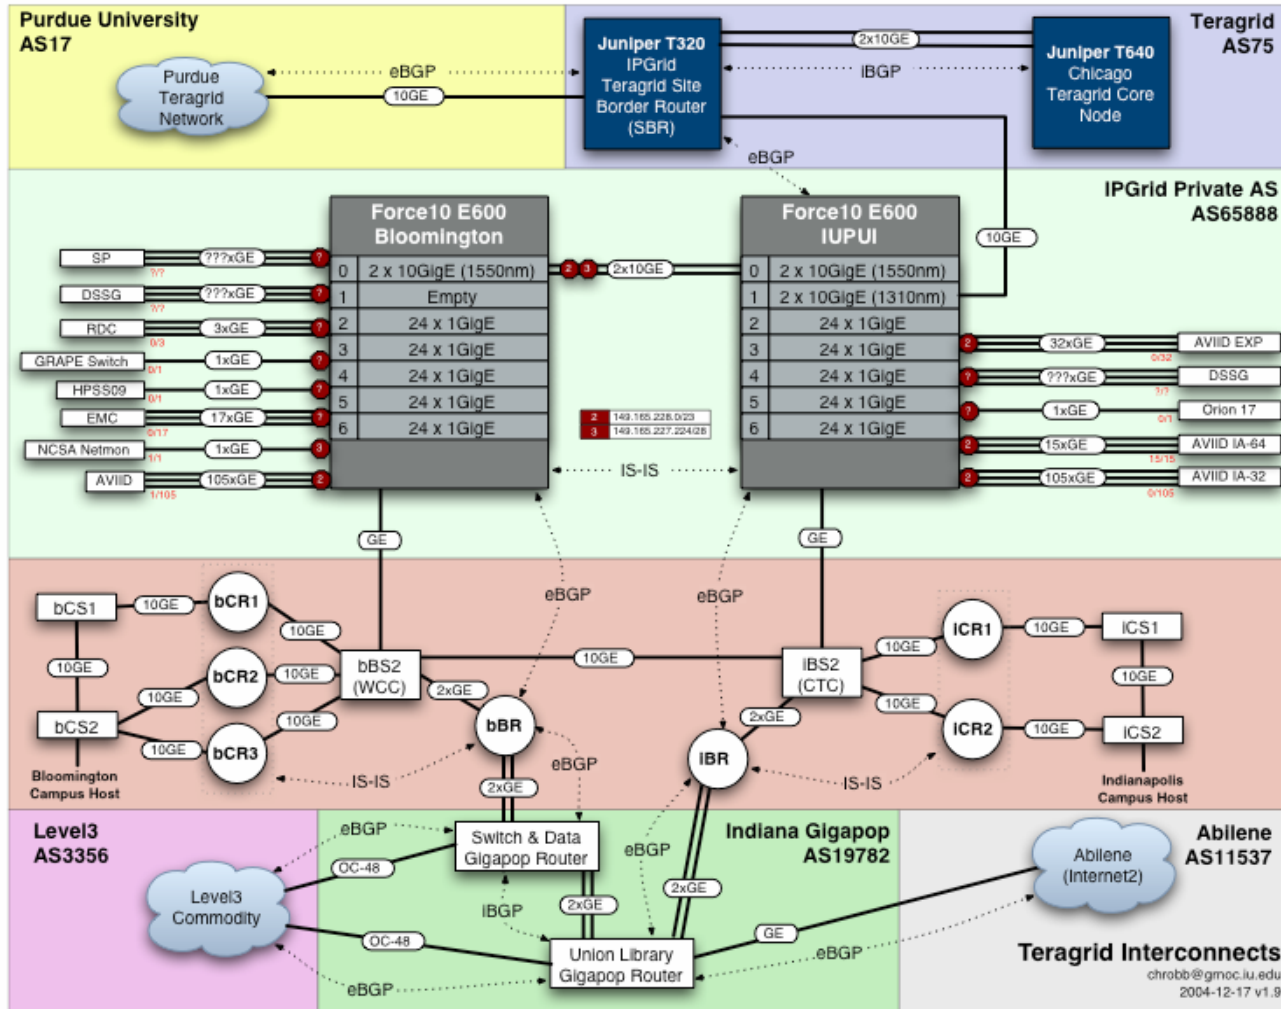
Bloomington



One view of the IU System



Network



Grids at Indiana University



Open Science Grid

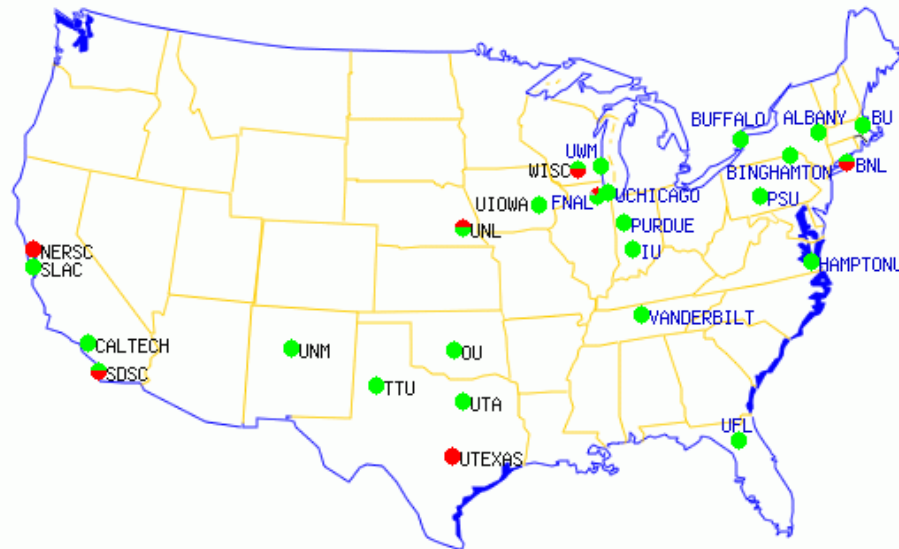
IVDGL Tier 2 Operations for Atlas



Open Science Grid

X509

User Proxy



Wed Sep 28 20:16:52 GMT 2005



Grids at Indiana University



TeraGrid

- Received \$4.4 million NSF grant to participate as a resource partner in the TeraGrid
- Plan to provide computing resources, storage, data sources, “unique” resources

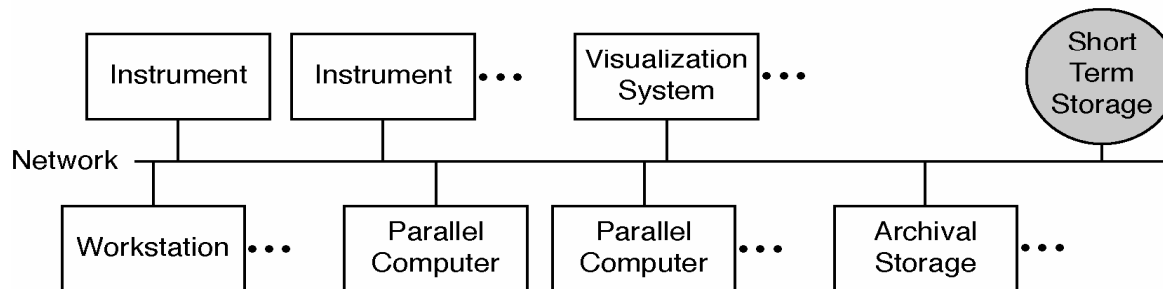


Grids at Indiana University



Data Capacitor

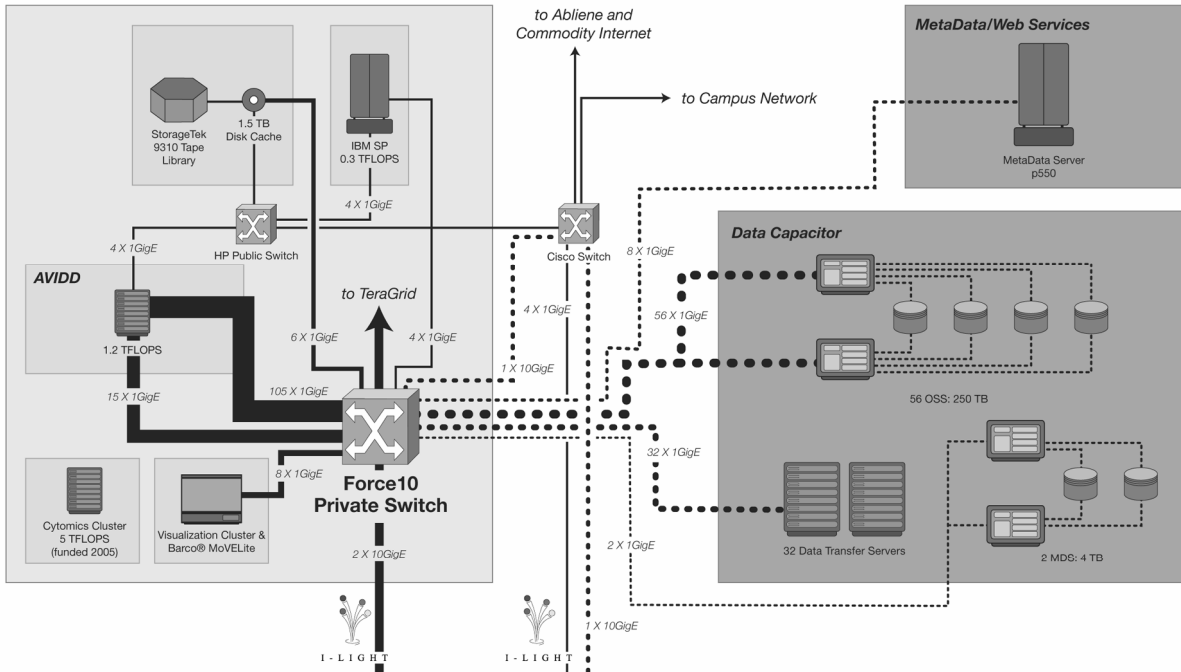
- Announced September 27, 2005, \$1.72 million NSF grant to build a massive short-term data storage system
- Data acts (sort of) as an incompressible fluid
- Catching the data deluge
- Parallel high speed I/O of files transported in serial
- Temporary storage between different parts of a work flow



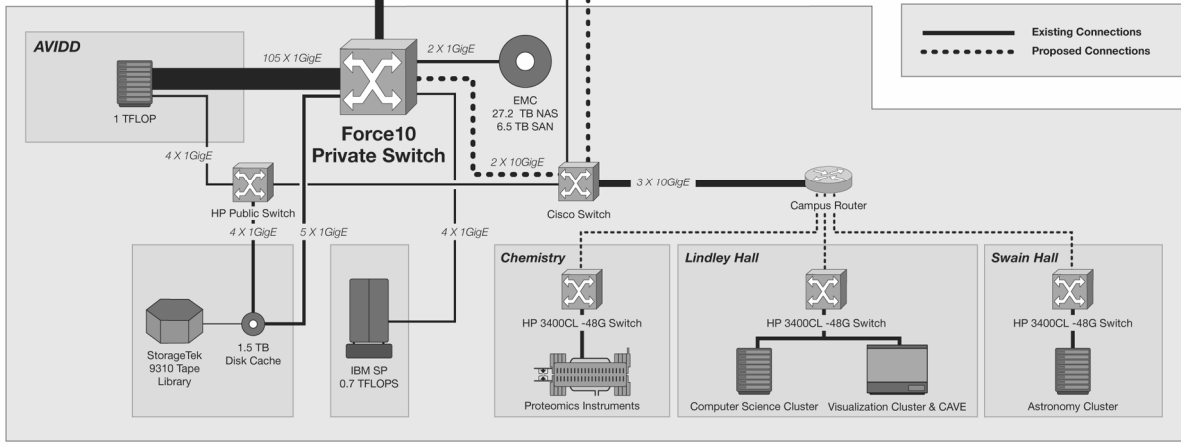
Project	Production rate (TB/year)		Disk use per job (TB)		Local archive or remote data source (TB)		Published to Web (TB)	
	Current	2008	Current	2008	Current	2008	Current	2008
Chemistry	3.5	12,900	1	100				
Biology								
Protein structure searching			1	10	0.26	10	1	10-100
Text mining for protein sites			1	10	0.1	1		
Functional genomics	2	6	4	4-10	3*	12*	4	10
Genome assembly			10	100	2.5	10		
Genome comparisons			10	100s	0.26	10		
Bio knowledge warehousing	0.1	0.3	0.2	0.5	1-2	10s to 100s	0.5	~1
Astronomy								
WIYN 3.5M & 0.9M	3.5	3.5	1-2	≥2	7*	10-20	2	
WIYN 3.5-m+ODI		120	-	120	-	50		
High energy physics		1.5	20	100	2.5*	**		
Informatics								
Complex network analysis			10	100	1	1000		10
Linguistic analysis					1			
Computer Science								
NEXRAD Level II data management	1	5	5	5	68*	88*	4	100s
Network traffic analysis		13	1	13	12.8*	100s*		
TOTALS	10.1	13049	-	-	≥99	≥1300	>10	≥100



Indiana University – Purdue University Indianapolis



Indiana University Bloomington



Grids at Indiana University



Acknowledgments

- Funding for projects described in this talk has come from the National Science Foundation, National Institutes of Health, Lilly Endowment, Inc., State of Indiana (particularly through support of I-light Initiative and the 21st Century Fund)
- This material is based upon work supported by the National Science Foundation under Grant No. 0116050. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation (NSF)."
- "The project described was supported by grant number 1U24AA014818-01 from NIAAA/NIH. Its contents are solely the responsibility of the authors and do not necessarily represent the official views of the NIAAA/NIH."
- "This research was supported in part by the Indiana Genomics Initiative. The Indiana Genomics Initiative of Indiana University is supported in part by Lilly Endowment, Inc."
- "This work was supported in part by Shared University Research grants from IBM, Inc. to Indiana University."
- The work described here was made possible by the faculty, students, and staff of Indiana University. Thanks especially to the staff of RAC, CPO, Telecommunications, PTL, UITS generally, the participants in the Indiana Genomics Initiative, and the participants in the METACyt Initiative.



For additional info

- www.purdue.teragrid.org/
- www.itap.purdue.edu/
- www.rcac.purdue.edu/
- rac.uits.indiana.edu/
- www.iu.teragrid.org/
- uits.iu.edu/
- email: smccaula@indiana.edu
- Send email to Scott McCaulay if you are interested in getting a starter (DAC) account on the TeraGrid

