

# Campus Research Computing: On Ramp to National Supercomputing Resources

---

**Jan E. Odegard**

Executive Director

Computer & Information Technology Institute

Rice University

odegard@rice.edu; 713 348 3128

<http://www.citi.rice.edu>

Campus Research Computing Cyberinfrastructure

April 27, 2006



Computer and Information Technology Institute



# People not here that deserve credit




- Members of the Rice BOG
- David Leebron, President
- Eugene Levy, Provost
- Moshe Vardi, Director C I T I
- Ken Kennedy, Director HiPerSoft
- Sallie Keller-McNulty, Dean of Engineering
- Kathleen Matthews, Dean of Natural Sciences
- Kamran Khan, VP for I T
  - William Deigaard, Director I T
  - Richard Peterson, Director I T
  - Barry R. Ribbeck, Director I T
  - Kim Andrews, Manager R CSG
- many others



To build a community of scholars that engages in collaborative research and education covering virtually every aspect of information technology and computing

## Directors:

Ken Kennedy (1986-1992)



Sidney Burrus (1992-1998)



Willy Zwaenepoel (1998-2001)



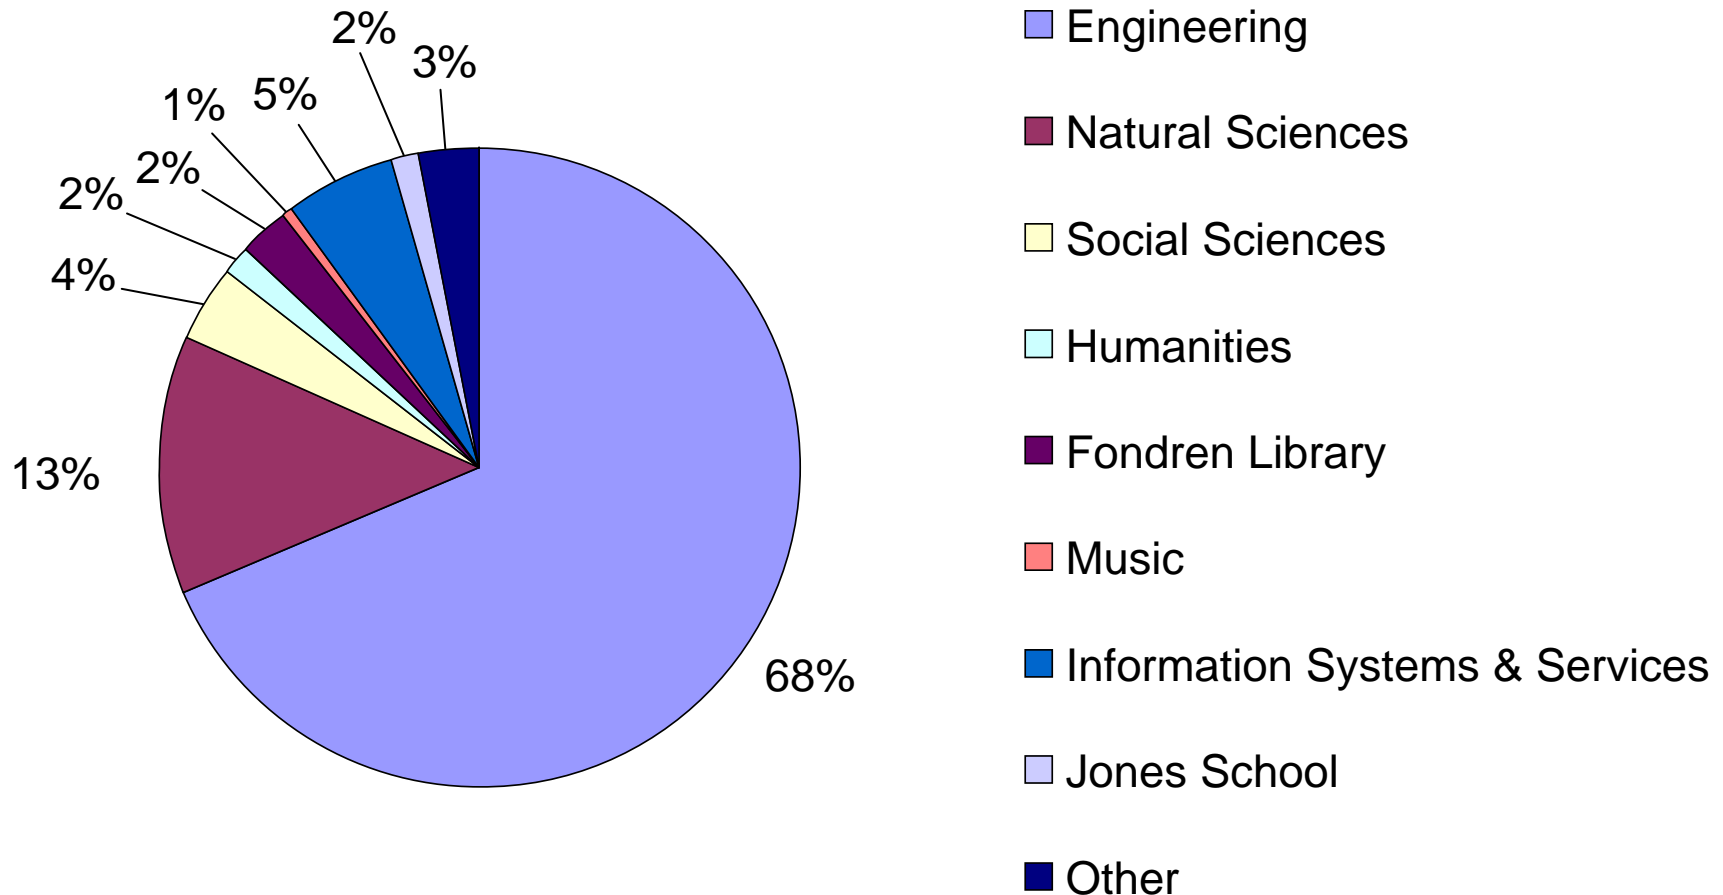
Moshe Vardi (2001-...)



# CITI Members



6 schools  $\leftrightarrow$  20 departments  $\leftrightarrow$  ~135 members  
7 centers  $\leftrightarrow$  ~15 ad hoc research groups



# Research Centers




- Center for High Performance Software (HiPerSoft)
  - Director: Ken Kennedy, CS
- Center for Multimedia Communication (CMC)
  - Director: Behnaam Aazhang, ECE
- Center for Computational Geophysics (CCG)
  - Co-directors: B. Symes, CAAM / A. Levander, ES
- Center for Computational Finance & Economic Systems (CoFES)
  - Director: Kathy Ensor, STAT
- Laboratory for NanoPhotonics (LANP)
  - Director: Naomi Halas, ECE
- Center for Technology in Teaching and Learning (CTTL)
  - Director: Tony Gorry, CS
- Center for Excellence and Equity in Education (CEEE)
  - Director: Richard Tapia, CAAM

# Research Groups & Labs



- Dynamical Systems Group
- Statistical Consulting Lab
- Bioinformatics Group
- Robotics Group
- Rice Networking Group
- Sensor Nets Group
- Theoretical and Computational Neuroscience
- Rice Computer Architecture Group
- Complex Flow of Complex Fluids Group
- Advanced Research Initiative on the Emerging Library
- Connexions: Open content education repository
- Digital Library
- ...



Computer Modeling of a Cerebral Aneurysm. Blood-flow patterns at an instant during the systolic cycle

- C I T I set out to:
  - Coordinate and develop a model for supporting large scale computer resources for the research community @ Rice
- Why?
  - Growing need for (larger) resources
  - Perceived barriers to access national resources
  - Cluster explosion (“Linux – it must be free”)
  - Faculty need a supported system that permit development of code that would be ready for consuming large national resources

- From the archives:



- and many other systems from past and current...
- January 2002:
  - CITI pulled together a team of ~30 investigators and wrote a successful MRI proposal
- January 2004:
  - CITI pulled together another team of ~35 investigators and wrote a second successful MRI proposal

### "Rice Terascale Cluster"

- ~1 TeraFLOP HP Linux cluster\*
  - 286 Intel® Itanium® 2 processors
    - 900MHz, 1.5MB
    - HP zx1 chipset

Support about 230 users  
plus support several classes

- 640GB memory
- 11TB Disk
  - 6.5TB on node
  - 1TB scratch back-end
  - 3.5TB shared front-end



NSF MRI, Rice, Intel and HP



# Cray XD1 System

Dual-Core AMD Opteron™




## "Rice Computational Research Cluster"

- ~3 TeraFLOP Cray Linux cluster\*
  - 336 Dual-Core AMD Opteron™ 275
    - 2.2GHz, 1MB / Core
    - 168 dual socket nodes (4 way SMP)

100 users and growing fast  
will support several classes

- 1.4 TB DDR2 400
- 12 TB Local Disk
- 6 TB Lustre parallel file system
- 10 TB NFS file system



NSF MRI , Rice, AMD and Cray

# ADA: Cray XD1



- <http://rcsg.rice.edu/ada/>
- "Friendly Users" January 1, 2006
- "Ribbon Cutting" February 10, 2006
- Pre-Production queues opened March 8, 2006
  - March 8-31, 2006 - 32% utilization
- Announced limited availability April 1, 2006
  - April 1-12, 2006 - 52% utilization
- Target general availability May 1, 2006

# Rewind 2002: What's the problem?



- Faculty voiced need for increased IT support for research
- S&E IT committee formed August 2002
  - Reviewed current and past
  - Interviewed colleagues
  - Interviewed Rice faculty
- Major findings:
  - Most of the institutions we interviewed had cost recovery to support some or all of research computing support
  - Rice was the exception (we had no service center providing IT support for research)
  - Rice faculty was surprised to learn that their colleagues at other institutions in many cases paid for “better support”
- Challenge:
  - Converting findings into funded actions
- Solution:
  - Forge partnership with IT Division (this is not CITI)

- Rice hired Vice Provost for I T
  - Kamran Khan
- Re-aligning I T organization
- Faculty now have strong advocate
- Created I TAC
  - Information Technology Advisory Committee
- Challenge
  - Sell and build a new campus network
  - Sell to renovate or build a new datacenter
  - Develop a sustainable model for research support
  - We needed this yesterday



# Information Technology Support Structure

Kamran Khan, Vice Provost for IT

Donna Shorey, Executive Assistant

## Academic & Research Computing

Rick Peterson, Director

Staff: 44

## Enterprise Applications

Andrea Martin, Director

Staff: 18.3

## Networking, Telecommunications, & Operations

William Deigaard, Director

Staff: 23

## Systems, Architecture, & Infrastructure

Barry Ribbeck, Director

Staff: 21

## Administration, Planning, & Finance

Yemeen Rahman, Director

Staff: 9

### CLIENT SERVICES

#### HELP DESK

- Diane Butler
- Rick Russell
- Arnold Chee
- Bill Klemm
- Jeff Koffler
- Carolynne White

#### SCIENCE AND ENGINEERING

- Michael Dewey
- Eric Aune
- Jackie Blyden
- Fernando Gonzalez
- Daniel Jackson
- Harry Le
- Michael Lightfoot
- Keith Poli
- Niki Serakiotou
- Jason Scott †
- Tim Stafford
- Ahmed Syed
- Clinton Heider
- TBN

#### ADMINISTRATION, ARCHITECTURE,

#### HUMANITIES, MUSIC,

#### SOCIAL SCIENCES

- Al Grazis
- Eric Bailey
- David Bergman
- Scott Blackard †
- John Croft
- Hans Krause
- Stan Kwan
- Tom Murray
- Chesson Olawaiye
- Tom Oster
- Rick Roberts
- Rav Jones

### RESEARCH COMPUTING GROUP

- Kim Andrews
- Franco Bladilo
- TBN
- Roger Moye
- Keith Schincke
- Kiran Thyagaraja
- TBN

### INSTRUCTIONAL TECH. SUPPORT

- Carlos Solis
- Karen Ross
- Sheldon Cohn
- David del Pino
- Terry Graham
- Hector Piñeda
- Angela Rabuck

### MEDIA SERVICES / EVENTS

- Doug Killgore
- Tom Lytle

### LIBRARY/DIGITAL SERVICES

- Diane Butler, Assistant Director
- Denis Galvin
- Karen Oster
- Chris Windham
- Ali Salman

### DATABASE ARCHITECTURE

- Mir Mirhashimali
- Omer Piperdi

### APPLICATIONS & DATABASE SERVICES

- Linda Humphreys
- Ali Aijaz
- Jimmy Rannik

### COMMUNICATIONS & MARKETING

- Liz Brigman (.5)
- Carlyn Chatfield
- Iris Chang (.5)

### WEB SERVICES

- Jeff Frey
- Krystal Bivens
- Iris Chang (.5)
- David Dean
- Michael Chiu
- Katy McKinin (.6)
- Chris Pound
- Trey Rouse

### NETWORKING

- ◊ Long Pham

### CAMPUS NETWORK

- Bao Dam
- Danny Eaton
- Ryan Moore

### TECHNICAL SERVICES

- Jeremy Reichert
- Paul Williams
- John Thompson

### TELECOMMUNICATIONS

- Reggie Clarkson
- Delia Diaz
- Sylvia Hillsman
- Mary Perez
- Sharon Scoppa
- Stan Shearer
- Ruben Tamez
- Mike Thorn

### DATA CENTER

- Shelby Sims
- Robert Castillo
- James Chien
- Edwin Martinez
- Prabesh Vijayananda
- Dora Lomeli

### SYSTEMS MANAGEMENT

- Amado Aguilar

### CLIENT SYSTEMS

- Michael Lucas

### CONSULTANTS

- Omar Abul-Aziz

### EDUCATIONAL SYSTEMS

- Mayfield Harris
- Lance Ogletree

### SYSTEMS

- Art Gorski
- David Hagan
- Phillip Hsieh
- Daniel Stephens
- Charles Wright

### DIGITAL MEDIA TECHNOLOGY

- Hubert Daugherty

### MIDDLEWARE DEVELOPMENT & INTEGRATION

- Kenneth Marshall
- Karl Burkett
- Paul Engle
- Dean Lane
- Susan McClure
- Al Ramos
- Tod Sandman
- TBN

### FACILITIES & PROJECT MANAGEMENT

- Derek Rabuck

### PROCUREMENT & VENDOR RELATIONS

- Rick Nelson
- Carl Furra
- Henry Halliburton

### ADMINISTRATION & FINANCE

- Tim Pitts
- Alice Butler
- Elizabeth Ford
- Jin Ung

## Information Technology Security Office

Marc Scarborough, Information Security Officer

Staff: 2

Levi Ball

- ◊ Assistant Director

- Manager

- † On military leave

# Network Projects - Convergence

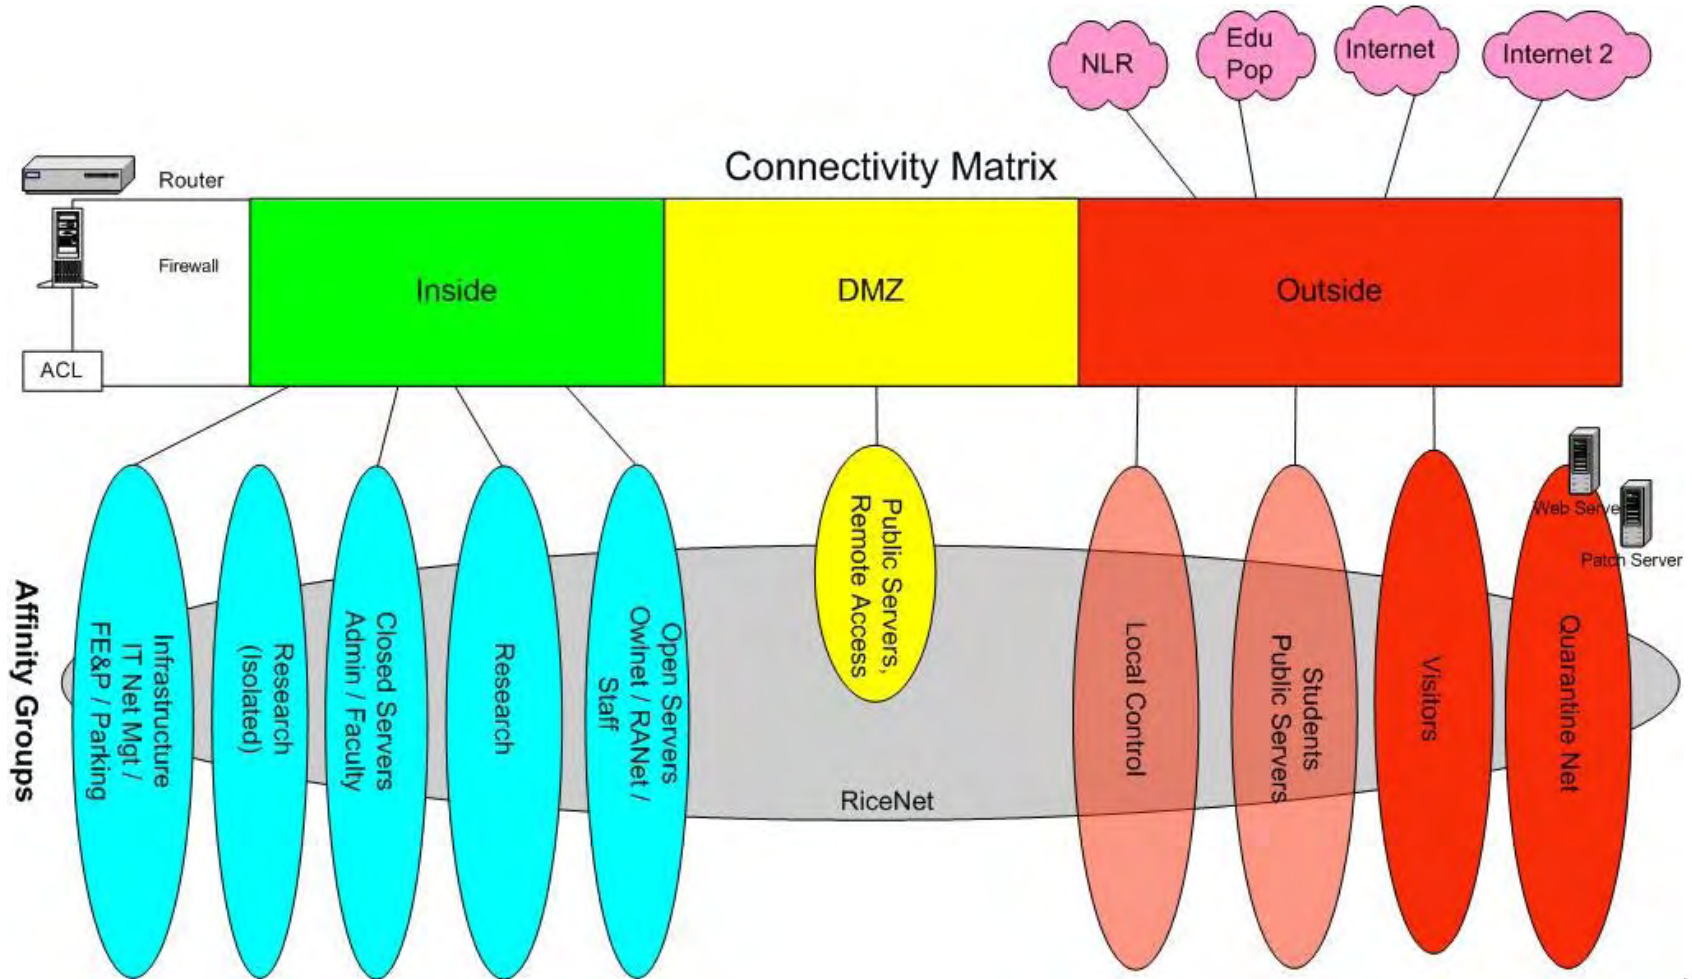


- New campus network
  - Network=Storage
  - 44,000 ports needed
  - 150+ wiring closets
  - 18 months to do it
- Southeast Texas GigaPOP
- Research and Education Network of Houston
- Lonestar Education and Research Network
- National Lambda Rail

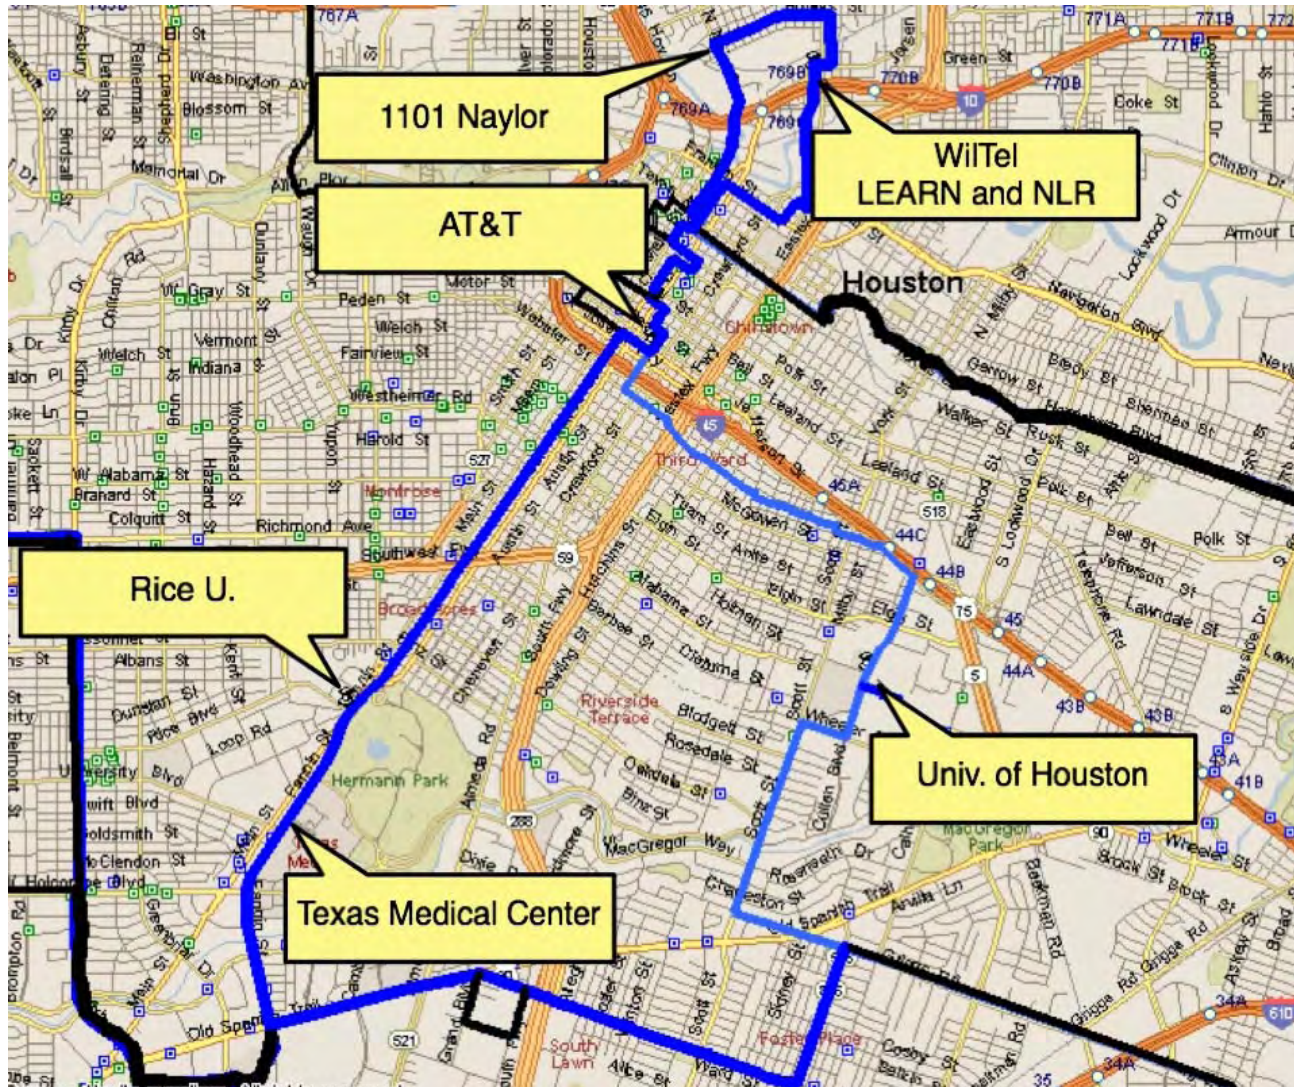


- Concept: Storage is part of the network!
- Rolling out “today”
  - Enterprise Storage (~20TB scale to ...)
  - Commodity Storage (~10TB scale to ...)
  - Near Line Storage (~85TB scale to 225TB)
  - Archive Storage (tape/DVD/...)
  
  - Support connectivity to grandfather “private storage” for backup and archiving (lifecycle management, leverage & partnerships)


# Campus Network Affinity Groups



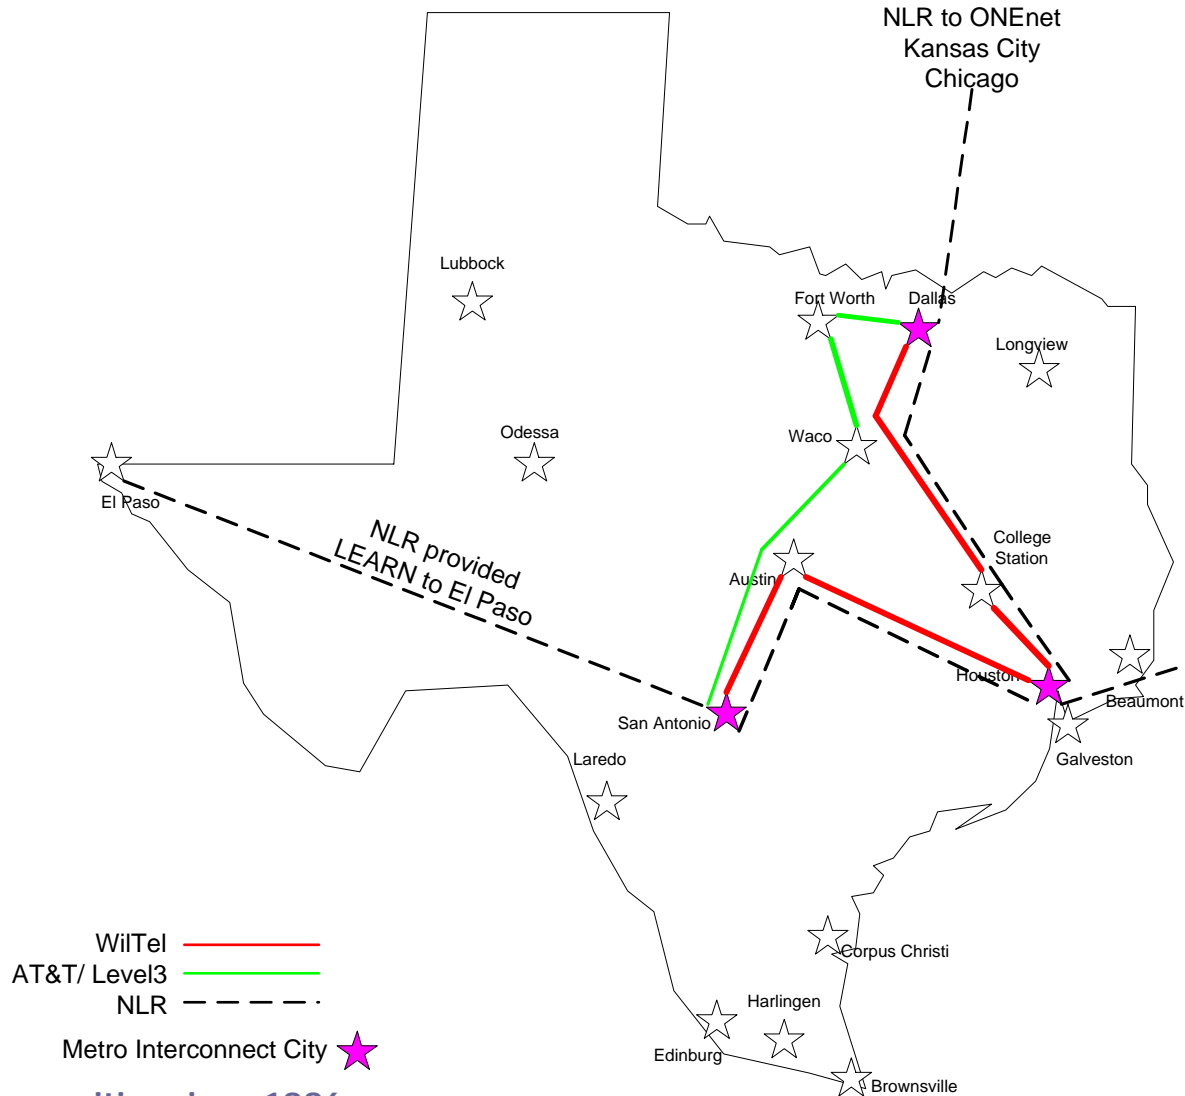
# RENoH Physical Layer



# Rice & Texas Medical Center



# LEARN: Aggregation in Texas!



# HiPCAT - TIGRE



- HiPCAT - High Performance Computing Across Texas
  - Share knowledge and high performance computing (HPC) resources
  - Pursue the development of a competitive Texas Grid
- TIGRE - Texas Internet Grid for Research and Education
  - Promotion of research, education and competitiveness of Texas
  - Computation, information, and collaboration
- Participants:
  - Rice University
  - Texas A&M University
  - Texas Tech University
  - University of Houston
  - University of Texas
  - Baylor College of Medicine
  - University of Texas at Arlington
  - University of Texas at El Paso
  - University of Texas Southwestern Medical Center at Dallas
  - University of Texas Health Science Center at San Antonio



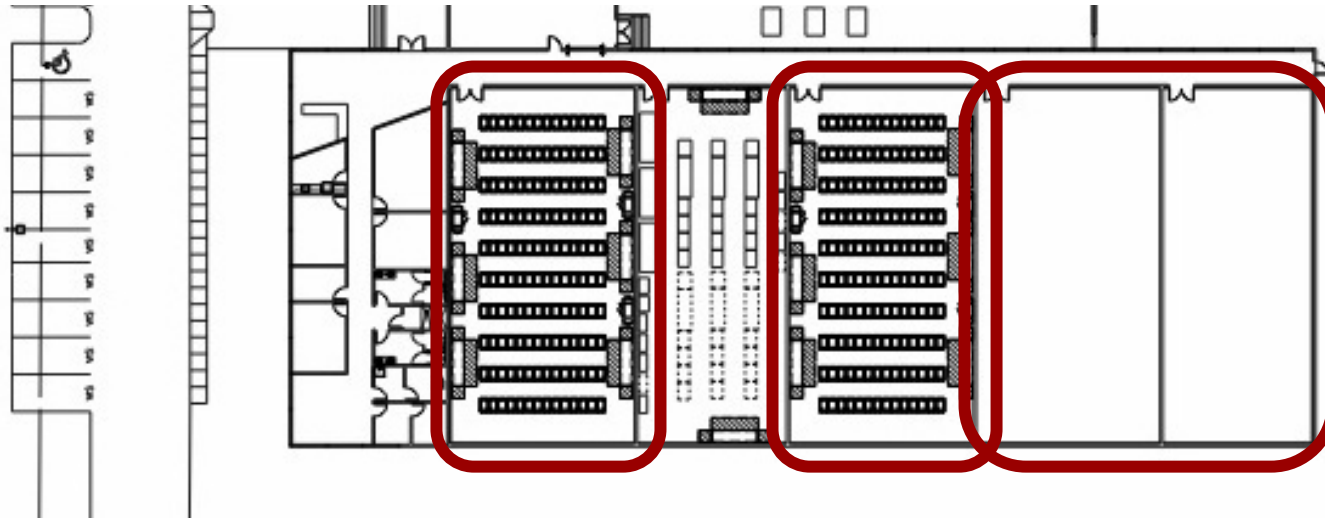
# Internet2





# New Datacenter

- 20,000 sq.ft.
- 6 MWatt inbound power
- Three pod configuration (3-15+ kW/rack)
- Three separate electrical systems (A, B & C)
- Mechanical: N+1 or better



- Buying & managing computing resources large enough for peak demands is generally not cost efficient or feasible for individuals or small groups – offer condo
- Framework for sharing with an eye towards:
  - Leveraging investments
  - Leveraging support
  - Leveraging knowledge & training
  - Permit semi-private partitions “condo”
    - Managed by Research Computing Support Group
    - Procurement is an “addition” to a major resource
    - Significant portion of private resource can be shared
    - Include “management team” of large resource in procurement
    - Discuss and sign MOU prior to procurement starts
- Act as an on-ramp for National Resources!

# Challenges for IT Research Support



- OMB A21 (principles for determining costs applicable to federal grants, contracts, and other agreements with educational institutions)
  - Core Services SLA
  - Education SLA
  - Research SLA
- Security
  - Front and center
- Sustainability
  - Infrastructure
  - Infrastructure support
  - Total cost of ownership
  - Lifecycle (it's short) management
- Incentives for “on ramps”
  - Campus centers
  - National centers
- Success metrics
  - is it when a user successfully “move” from campus to national

