UCSD's Research CyberInfrastructure (RCI) Program:

Create … Share … Discover

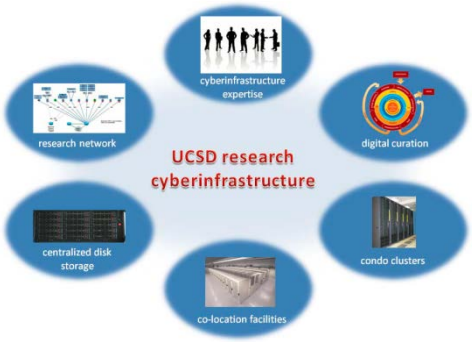

January 24, 2013

Richard Moore

RCI Project Manager

r3moore@ucsd.edu

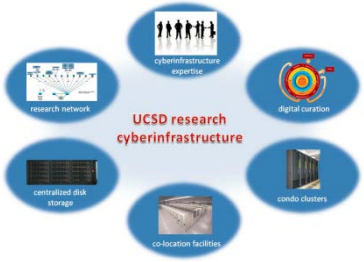

A brief history of UCSD's RCI Program …

- **2008-2009: Research Cyberinfrastructure Design Team (RCIDT)**
  - Chaired by Mike Norman and Phil Papadopoulos; broad campus participation
  - Campus-wide survey of research cyberinfrastructure needs (2008)
  - April 2009: RCIDT issued *Blueprint for the Digital University* (*http://rci.ucsd.edu/_files/Blueprint.pdf*)
- **2009-2010: *CyberInfrastructure Planning & Operations Committee (CIPOC)* developed a business plan with recommendations**
  - A principle of shared costs between PIs and campus RCI investments
  - April 2010: *CIPOC* issued their report
- **2011-Present: *RCI Oversight Committee* charged to implement RCI**
  - January 2011 - CIPOC business plan accepted, oversight committee charged
  - Chaired by Mike Norman and Mike Gilson; broad campus representation

Elements of UCSD's Integrated Research CyberInfrastructure

- **Data Center Colocation**
- **Networking**
- **Centralized Storage**
- **Data Curation**
- **Research Computing**
- **Technical Expertise**

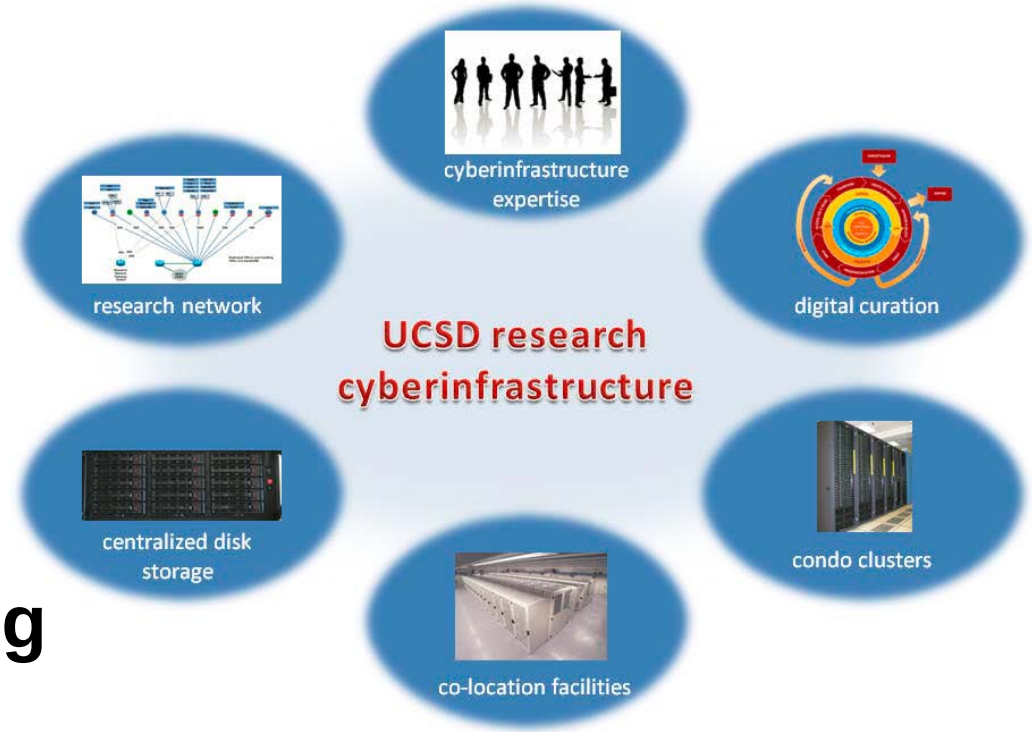

Why UCSD Is Investing in RCI Program

- **Increase competitiveness of UCSD researchers**
- **Realize cost efficiencies and improve service via economies of scale and shared services**
- **Preserve UCSD's digital intellectual property**
- **Save energy/$ and effectively use data center capital investments (colocation)**

RCI Program is "by campus, for campus"

- **RCI priorities driven by researcher requirements**
- **Oversight Committee represents all campus units and sets the strategic directions and implementation plans**
- **Implementation partners from across campus**
  - Administrative Computing & Telecommunications
  - Calit2
  - San Diego Supercomputer Center
  - UCSD Libraries

RCI Oversight Committee Charge & Membership

- Guide the development of an enhanced RCI
- Set policies and coordinate RCI elements
- Modify the CIPOC-proposed business plan as needed
- Continually seek opportunities for cost savings and recommendations for their implementation
- Provide accountability and document accomplishments on a regular basis

- Co-chairs
  - Michael Gilson, Professor, School of Pharmacy/Pharmaceutical Sciences
  - Michael Norman, Director, San Diego Supercomputer Center
- Members
  - Ed Babakanian, Chief Information Officer, Medical Center
  - Steven Briggs, Professor, Biological Sciences
  - Declan Fleming, Chief Technology Strategist, Library
  - Terry Gaasterland, Professor/Director Genomics, Marine Biology Research Jeffrey Gee, Professor, Geosciences Research Division (Alternate Jerry Wanetick, IT Director, SIO)
  - Jeff Henry, Director, Academic Computing/Media Services (Alternate Pedro Cruz, Assistant Director, Academic Computing/Media Services)
  - Bill Hodgkiss, Associate Vice Chancellor, Academic Planning & Resources, Academic Affairs
  - Andy McCammon, Professor, Chemistry and Biochemistry
  - Richard Moore, Deputy Director, San Diego Supercomputer Center
  - Bing Ren, Professor, Cellular and Molecular Medicine, Health Sciences
  - Tad Reynales, Chief Technology Officer, Calit2
  - Brian Schottlaender, The Audrey Geisel University Librarian, Library
  - Ed Spriggs, Assistant Vice Chancellor, Student Affairs
  - Shahrokh Yadegari, Chair, Committee on Academic Information Technology, Academic Senate, ex officio
  - Min Yao, Assistant Vice Chancellor, Administrative Computing and Telecommunications (Alternate: Charlotte Klock, Executive Director/Chf Info Sec-Prvcy Of)
- Consultants
  - Bobb Barile, Director, Organizational Development, Academic Affairs
  - Ron Espiritu, Assistant Vice Chancellor, Health Sciences, School of Medicine
  - Marianne Generales, Assistant Vice Chancellor, Research Affairs
  - Gene Hasegawa, Assistant Dean, School of Medicine (Alternate: Ron Spears, Assistant Dean, IT Enterprises, School of Medicine)
  - David Hutches, Director, Engineering Computing, Jacobs School of Engineering
  - Sylvia Lepe-Askari, Assistant Vice Chancellor, Campus Budget Office

RCI is rolling out production services for UCSD researchers

Service	Status	Lead/contact for service
Colocation	Production	Matt Campbell (SDSC) mattc@sdsc.edu
Networking	Production	Valerie Polichar (ACT) vpolichar@ucsd.edu
Research Computing	Production Feb 2013	Jim Hayes (SDSC) jhayes@sdsc.edu
Centralized Storage	Initial Production Spring 2013; Expanded Services thru 2013	Wilfred Li (SDSC) wilfred@sdsc.edu
Data Curation	Completing pilots; production FY13-14	Ardys Kozbial (Libraries) akozbial@ucsd.edu
Technical Expertise	Not planned as formal RCI service; expertise distributed across departments and projects	

Colocation – in production

- **Host IT equipment in energy-efficient, manned data center**
  - SDSC's 18kft2, 13MW datacenter
  - Standard rack space, secure facility, seismic protection
  - 24/7 operations staff provide facility oversight and emergency "remote hands" hardware assistance

- **RCI supplements rack rate: user pays $2500/rack/year**
- **NGN *may* cover basic networking costs; evaluated case-by-case**
- **Up to 10 Gb networking fabric connectivity available, both thru SDSC aggregation fabric and into CENIC**
- **UPS and generator capabilities available**
- **Cage and locked rack options available for security/compliance**

High-Performance Networking - in production

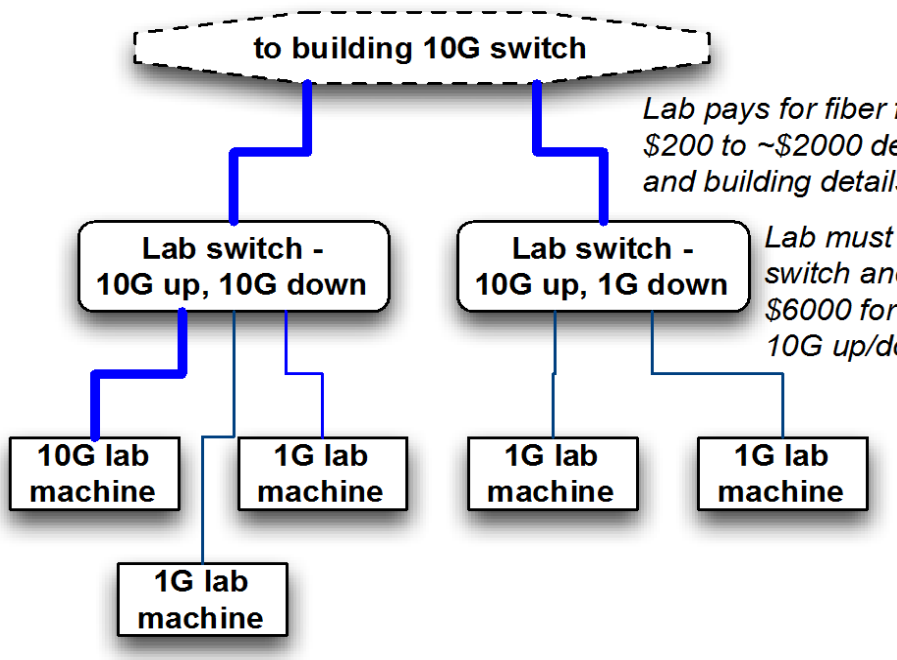

- **ACT is rolling out 10G network infrastructure to campus buildings under NGN3**
- **Department/lab pays for "last 100 feet" connectivity to use capability**

Sites requiring dedicated 1G or 10G pipes across campus (or to Internet2 or NLR) will incur specific additional costs, hard to estimate generally

List of 10Gbps connected buildings (as of 1/2013)

Black = already 10Gbps
Purple = upgraded to 10Gbps since June 2011
Green = projected by June 2013

Main Campus
- AP&M Bldg
- Bonner Hall
- CalIT2
- Center Hall
- Chancellor's Complex
- CMRR
- EBU1
- EBU2
- EBU3A (BioEngineering)
- EBU3B (CSE)
- Geisel Library
- High Energy Physics
- HSS Bldg
- IRPS Complex
- Mayer Hall
- Media Center
- Muir Biology
- Music Bldg
- Natural Sciences Bldg

Main Campus (Cont.)
- Pacific Hall
- Pepper Canyon Hall
- Police Dept
- Rady School of Management
- RIMAC Annex
- Sanford Consortium (SCRM)
- SERF Bldg
- Social Sciences Bldg
- Social Sciences Research Facility
- Structural and Materials
- Engineering Building (SME)
- Torrey Pines Center North
- Torrey Pines Center South
- Urey Hall
- Village West Residential Complex
- Visual Arts Complex
- York Hall

School of Medicine
- Basic Sciences Bldg
- BioMedical Library
- Black Boxes

School of Medicine (Cont.)
- CMG Building
- CMM East
- CMM West
- HSBRF2 Building
- Holly
- Laurel Complex
- Liechtag Building
- Keck Facility
- Skaggs School of Pharmacy
- Stein Building
- Telemedicine Building

SIO
- Birch Aquarium
- Deep Sea Drilling West
- Hubbs Hall (colo)
- IGPP Complex
- IGPP Munk Labs
- Keck/OAR/Nierenberg add'n complex
- Nierenberg Hall
- NORPAX/Isaacs Hall
- Ritter Hall
- Scholander Hall

SIO (Cont.)
- Scripps Satellite Facility
- SIO Library
- Sverdrup Hall
- Vaughan Hall

East Campus
- Moores Cancer Center
- Radiation Oncology Bldg
- Thornton Hospital
-

Remote Locations
- 10300 Campus Point Drive
- Convoy Court Data Center
- Dunhill Street
- Hillcrest Medical Center
- La Jolla Professional Center

Research Computing – in production Feb 2013!

- **RCI is evolving SDSC's Triton system to the "*Triton Shared Computing Cluster*" (TSCC)**
- **Condo model: Researchers purchase compute nodes which are operated as part of shared cluster for 3-4 years**
  - PI buys hardware & pays modest (RCI-supplemented) operations fee
  - Lower operations cost than local PI cluster; larger-scale resource available (core count and capacity); professionally-managed
- **Hotel: Purchase time by the core-hour; shared queue**
- **Provisional UCSD rates (accepting orders)**
  - Condo: ~$4K/node+ ~$1K/node infrastructure + ops fee ~$500/node/yr
  - Hotel: 2.5 c/core-hour (3.0 c/SU other UC campuses)

Centralized Storage (phased production thru 2013)

- Completing interviews of a broad sample of ~50 representative PIs across campus to understand their technical and cost requirements
- Identify common needs, and define sustainable RCI business model with likely adoption by campus users
- Report expected ~March 2013

- Anticipate production centralized storage services in CY13
  - RCI Data Services/Network-Attached Storage (RCI/NAS) available Spring CY13
  - Further services to be rolled out thru year, based on requirements analysis
- Note: SDSC already offers storage services for campus and beyond, but RCI is assessing campus' requirements and recommending appropriate service/business model

Data Curation – in pilot (production FY13-14)

- **Completing a two-year pilot phase**
  - How do lab personnel work with librarians to curate their data?
  - How much work is required to curate data and what are options?
  - What is a sustainable business model for curation within RCI project?
- **Five representative programs across UCSD selected as pilots**
  - The Brain Observatory (Annese)
  - Open Topography (Baru)
  - Levantine Archaeology Laboratory (Levy)
  - SIO Geological Collections (Norris)
  - Laboratory for Computational Astrophysics (Wagner)

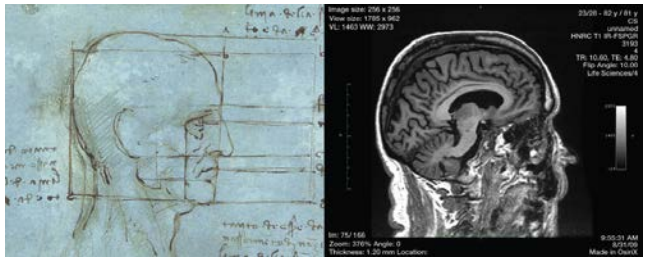

- **Using existing tools whenever possible**
  - Storage at SDSC, campus high-speed networking, Digital Asset Management System (DAMS) at UCSD Libraries, Chronopolis digital preservation network
- **Also, develop Data Management Plan tools and provide training**
- **Anticipate production curation services in FY13-14**

Data Management Plans

- **Resources and contacts available to UCSD researchers**
- **Examples from submitted proposals**
- **Guidance, tips and recommendations for DMP preparation**
- **UCSD-centered version of DMP Tool**

http://rci.ucsd.edu/dmp/index.html

How to get more info about RCI

- **Web site: rci.ucsd.edu**
- **For each production service, site includes**
  - Description of services
  - Cost summary
  - Approved verbiage for PIs to use in proposals
- **Email rci@ucsd.edu**
- **Call Richard Moore, RCI Proj Mgr, ext 25457**

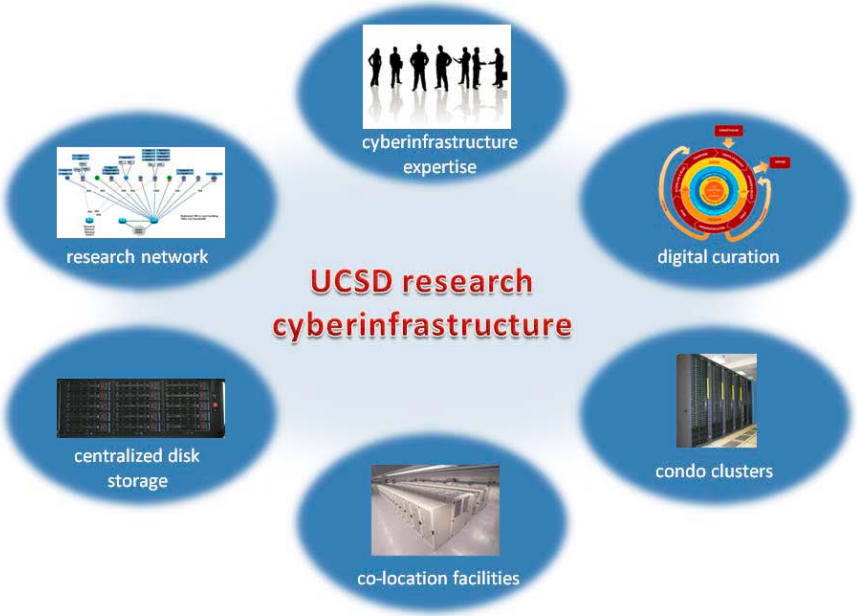

UCSD's RCI: Triton Shared Compute Cluster (TSCC)

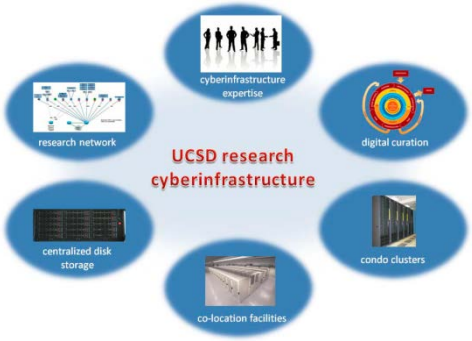

January 24, 2013

Jim Hayes

RCI Research Computing Lead

jhayes@sdsc.edu

TSCC Overview

- **TSCC is a shared computing resource designed to provide cost-effective computing for UCSD researchers**
  - Sharing provides economies of scale (costs and level of service) compared to multiple independent computing resources
- **RCI subsidizes fixed operating costs, increasing cost benefit to participants**
- **'Condo' model successful at many universities**
  - PIs buy hardware (nodes) and pay low (RCI-supplemented) operations cost for life of nodes (3-4 years)
- **TSCC 'hotel' allows pay-as-you-go access**

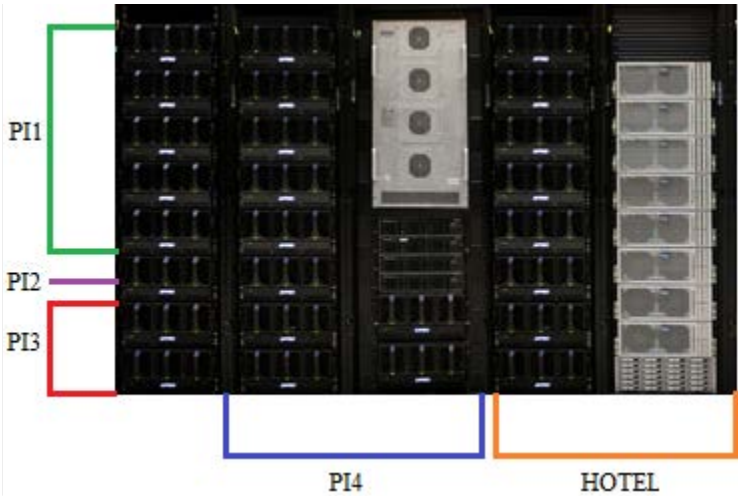

Who benefits from TSCC?

- **TSCC supports UCSD PIs, labs or departments with a wide range of research computing needs**
- **Condo is an attractive alternative to closet/lab clusters**
  - Good investment for compute cluster start-up funds
- **Hotel designed for researchers with modest/uncertain computing needs, or without upfront capital funds**
- **Can support UCSD parallel computing/computational science classes**
- **TSCC is not targeted to**
  - Administrative computing
  - Researchers needing national, petascale NSF or DOE systems

Lab clusters vs. TSCC

Lab Clusters

- **PI shops for and buys nodes, storage, interconnect, etc.**
- **PI hires staff or enlists student/ postdoc to administer system**
  - Significant $, diverts from science
- **Local control of schedule and priorities – within local cluster**
  - But often low-modest utilization
- **Uses office/department space or colo facility**

TSCC

- **PI provides index for nodes and pro rata infrastructure**
  - ~Same $, better infra, less work
- **Professionally-managed by RCI staff (~24*7)**
  - Modest ops $, NO admin work
- **Reasonable control of schedule/priorities + flexible usage for better utilization**
- **Energy-efficient colo facility**

Computing Options (CY 2013)

- **RCI competes/procures standard/GPU nodes each year**
  - Condo participants can buy any time during year @ RCI's cost
- **Standard compute nodes**
  - Dual-socket, 8 cores/socket, Intel Sandy Bridge, 2.6 GHz
  - 64 GB/node (4 GB/core); optional 128 GB/node (8 GB/core)
- **GPU nodes – still in selection**
- **Large-memory nodes (hotel only)**
  - 512 GB DRAM
  - 8-socket, 4 cores/socket
  - AMD Shanghai, 2.5 GHz

Supporting Infrastructure

- **RCI provides all supporting infrastructure**
  - Pro rata costs passed on to condo participants
- **Networking**
  - 1 GigE for management
  - 10 GigE for node-to-node, node-storage
  - Optional: InfiniBand (QDR)
    - Per-rack (1024-core) connections
    - Full bisection bandwidth up to 32 nodes
    - Hotel nodes are IB-connected

- **Home Directory storage**
  - 100 GB/login, replicated
- **Parallel File System (Data Oasis)**
  - Currently 700TB and ~25GB/s scratch space (~90 day purge)
- **Additional persistent storage available at additional cost**

Planned Software Stack

- **Rocks/CentOS (Initially v6.2)**
- **Compilers: Intel, PGI, GNU (4.6+), DDT, cilk, mono**
- **MPI: mvapich2, openmpi, fpmpi**
- **Applications/libraries to include:**

  abyss, amber, beast, emboss, fasta, glimmer, gmap, hmmer, mpiblast, mrbayes, ncbi, phylip, tcoffee, tigr, wgs, blat, bowtie, bwa, GATK, samtools, soapdenovo, velvet, apbs, cp2k, gamess, gromacs, lammps, namd, nwchem, columbus, cpmd, bbcp, bbftp, fftw, fsa, amos, mummer, gaussian, hadoop, hdf4, hdf5, idl, jags, atlas, lapack, metis, octave, parmetis, petsc, scalapack, sprng, superlu, trilinos, netcdf, nco, ipm, papi, pdt, tau, python3, R, rapidminer, scipy, siesta, tecplot, upc, vasp, visit, weka, matlab

Condo usage policies

- **Condo participants receive an annual amount of core-hours proportional to node purchase**
  - Can be used on own, hotel, and other condo nodes
  - Allocation expires/renewed annually
- **Guaranteed access to cores purchased within 8 hours**
  - Jobs within these bounds can run unlimited time
- **Can extend jobs to hotel and others' condo nodes**
  - Increases computing capability and flexibility
  - Hotel jobs compete with hotel users', have 72-hour time limit
  - Jobs running on others' condo nodes have 8-hour time limit
- **Gleaning jobs run on idle condo nodes without charge**
  - Allows opportunistic use of otherwise-wasted cycles
  - Gleaning jobs may be killed to make room for higher-priority runs
- **GPU/standard compute allocations cannot be mixed**

Hotel usage policy

- **Hotel (pay-as-you-go) users' jobs only run on hotel nodes**
  - Initially there will be 40 general computing nodes (640 cores)
- **Job limits: 72 hours / 128 cores (exceptions on request)**
- **General computing nodes will be allocated per-core, with up to 16 jobs on each node simultaneously**
- **For the large-memory PDAF nodes; a maximum of two jobs per node (i.e., 256GB or 512GB of memory); either 16 or 32 cores will be charged to each job**
- **No gleaning jobs by hotel users**

Provisional TSCC costs (UCSD users)

- **Note: all costs and projections of indirect costs (IDC) are provisional and reflect the recharge proposal currently under review by campus. Final costs and IDC will be charged when rates are approved.**
- **Condo**

Item (Prices for UCSD users)	Node Purchase	Infrastructure Fee	Annual Ops Fee
General Compute Node	$3,934	$920	$495/yr
128 GB memory	+ $575	+ $0	-
InfiniBand-connected	+ $0	+ $200	-
GPU Node	$TBD	$920	$495/yr
No options at this time			

  - Anticipate node/infrastructure fee to be IDC-free; annual ops fee will bear IDC
- **Hotel**
  - $0.025/SU (UCSD), $250 minimum (10,000 core-hours)
- **90-day trial accounts for UCSD researchers (250 SUs)**

When can I use TSCC?

- **TSCC will be in production Feb 2013**
  - GPU nodes expected March/April 2013
- **Condo users can place orders now or any time**
  - Generally, <1 month from time order is placed to using condo nodes; interim access possible
  - Current HW available thru CY2013; new options available late 2013 to keep up with technology
- **Hotel users can purchase and use time now**
  - Current Triton system available until TSCC in production
  - Current Triton users can migrate to TSCC

For more information

- **Web site rci.ucsd.edu/computing/**
  - Describes program and costs
  - User documentation will be available here as it's developed
  - Language that can be used directly in proposals
- **Technical questions: Jim Hayes (jhayes@sdsc.edu)**
- **Buy cycles/nodes or other business questions: Ron Hawkins (rhawkins@sdsc.edu)**
- **Join tscc-l@ucsd.edu mailing list for interested participants**