
INDUSTRIAL ENGAGEMENT AT SDSC



Ron Hawkins

Director of Industry Relations



September 18, 2012

Allan E. Snavely (1962-2012)



Chapters in SDSC's History

- **1985-1997 (“Centers program”):** General Atomics established SDSC with one of initial NSF supercomputer center awards
 - Strong engagement with industry selling cycles on highest-end Cray systems
- **1997-2006 (“PACI era”):** SDSC transitioned to UCSD and won the ‘NPACI’ program as a center leading a national partnership
 - Decreasing industry engagement as HPC became democratized with commodity cluster computing
- **2006-future (Present):** Transition to a multi-constituency cyberinfrastructure center with increased engagement across UCSD, U California and industry
 - Growing services and technical collaborations




Tech Transfer / Spin-Off


- **Nirvana Storage (Division of General Atomics Corporation)**
 - Distributed data management (“data grid”) software based on the open source Storage Resource Broker
 - Establish global namespace & metadata for large-scale, distributed data systems
- **StackIQ**
 - Automated large-scale software deployment solution (Hadoop, Cloud, HPC cluster) based on SDSC’s open source ‘Rocks’ system
 - Venture backed: \$3M Series A financing in 2011 by Avalon Ventures & Anthem Venture Partners
- **High Performance Wireless Research Network**
 - Fixed wireless microwave WAN throughout San Diego County
 - Backhaul for seismic sensor arrays, Palomar Observatory, public safety, and other uses
 - Funded by NSF for 10+ years
 - In 2011, SDSC transitioned HPWREN to a sustainable business model funded entirely by user fees

REVENUE SHARE BY CATEGORY (FY10-13)

FY10 REVENUES




FY11 REVENUES




= 2X FY10

FY12 REVENUES



= 1.5X FY11

FY13 REVENUES (FORECAST)




= 1.7X FY12

Industry Program Thrusts

- **Sponsored Research**
- **Providing Services**
 - Computing, Storage, Networking, Consulting
 - Mechanism for engagement with industry
 - Spread fixed costs to broader user base
- **Focused ‘Centers of Excellence’ in Areas of Interest to Industry**
 - ‘Big Data’
 - Data Analytics

Exemplary Research & Consulting Projects

- Evaluation of emerging hardware architectures for near real-time visualization of multi-TB 3D data sets
- Performance characterization of hybrid computing (GPGPU/FPGA)
- Benchmarking of HPC codes on different architectures
- Development of a large-scale HIPAA-compliant storage & analytics system
 - “Big data” handling for health IT



Cyberinfrastructure Services

- **HPC**
 - Triton Cluster
 - NSF Clusters (Gordon, Trestles)
- **Storage**
 - SDSC Cloud Storage (OpenStack Swift)
 - 'Project' Storage (file system storage services)
 - Data preservation services (partnership with DuraSpace)
- **Data Center Co-location**
- **HPWREN**

Exemplary Resource Services

- **Compute & storage for next generation sequencing and analysis (multiple projects)**
- **Finite element analysis for large civil engineering project (BART earthquake retrofit)**
- **Simulation & data analysis for an explosives detection system**
- **Image processing in support of a DoD-funded SBIR project**
- **Running specialized version of WRF for localized weather predictions**



HPC ‘Uber-Cloud’ Experiment

- **Currently participating in the ‘Uber-Cloud’ experiment in delivering HPC as a service**
- **Working with a small, San Diego-based product designer & manufacturer to run CFD modeling on SDSC’s Triton cluster**
- **Looking to use SDSC’s Gordon to run ISV codes on Windows over vSMP**
- **Also looking at Abaqus over vSMP for large-memory simulations**
- **More information: www.hpccexperiment.com**

Center for Large-Scale Data Systems (CLDS) Research

- Industry & NSF-funded center to research issues related to growth, value, and use of big data in enterprises
- Vision: Create a joint industry and research forum that can visualize future trends, provide management and technical insights, conduct system tests at scale and offer solutions
- Technical track: Big Data Benchmarking
- Management track: Data Growth Index
- Workshop held in San Jose, CA in May 2012
- Next workshop planned for Pune, India Dec 2012
- More information: clds.sdsc.edu



Predictive Analytics Center of Excellence (PACE)

- **Focus on research & education in predictive analytics**
- **Leverage strong analytics tech cluster in San Diego**
- **Big data analytics research on Gordon**
- **Major focus on analytics for smart grids & renewable energy systems**
- **Will hold first data mining 'boot camp' in October**
- **More information: pace.sdsc.edu**

Observations on Industry Engagement

- **Applied R&D and technical services will be a large part of the mix for industrial engagements**
 - Can be an end in themselves or a steppingstone to broader partnerships
- **Need a ‘customer’ focus and ‘solution’ mindset**
- **Smooth, repeatable processes are important**
- **Address intersection of HPC & Internet-scale computing**
- **Challenges to consider for university centers:**
 - Expectations of deliverables vs. grant-funded research
 - Reacting to industry timelines – “need it yesterday”
 - Technology Transfer & Intellectual Property issues
 - Proprietary information
 - Security
 - Commercial SW licenses – can be tough to get short-term or pay-as-you-go
- **Can be rewarding, especially with startups and small businesses**
 - Often, “HPC users that don’t know they’re HPC users”

Contact:

Ron Hawkins

Tel (858) 534-5045

Email rhawkins@sdsc.edu

THANK YOU!