

Data Center Observatory: A Living Lab for Data Center Study

A Utility and a Research Vehicle

- A computing and storage utility
 - for CMU researchers and services
- A source of measurement data
 - on operational costs
 - on administrator time
- A testbed for cost-saving technologies
 - automation of administrative functions
 - reducing energy usage

Physical Design Parameters

- ~2000 sq.ft. machine room space
 - plus adjacent administrator space
- Very visible showcase
 - main hallway, main floor of CIC building
 - windowed wall plus monitor with live stats
- Accommodates 40 compute/storage/networking racks
 - each rack at up to 2000 pounds and 25 kW
 - ~775 kW total, 425 W/sq.ft.
- Cooling via APC high density enclosures
 - contains hot air for greater cooling efficiency
 - simplifies staged deployment

Progress to Date

- DCO started operating in April 2006
- As of October 2013, 20 racks fully populated:
 - 750+ computers including blade servers, computing nodes, storage servers
 - 80 GbE link between zones
 - Minimum Gigabit Ethernet to each server
 - Network switches, file servers and monitoring gear
 - Susitna deployed with 40GbE to core network
- Deployment of OpenCloud and OpenCirrus clusters
 - 80 Gb Ethernet backbone

Related Projects

- | | |
|---|---------------------------------------|
| • Failure Data Analysis | • Open Cirrus |
| • Problem Analysis | • Open Cloud |
| • Self-* Storage | • Power Monitoring and Load Balancing |
| • CyLab's Biometric Research | • Tashi |
| • The Datapositionary | • vCloud |
| • Fingerprinting: automated problem diagnosis | • PRObE |

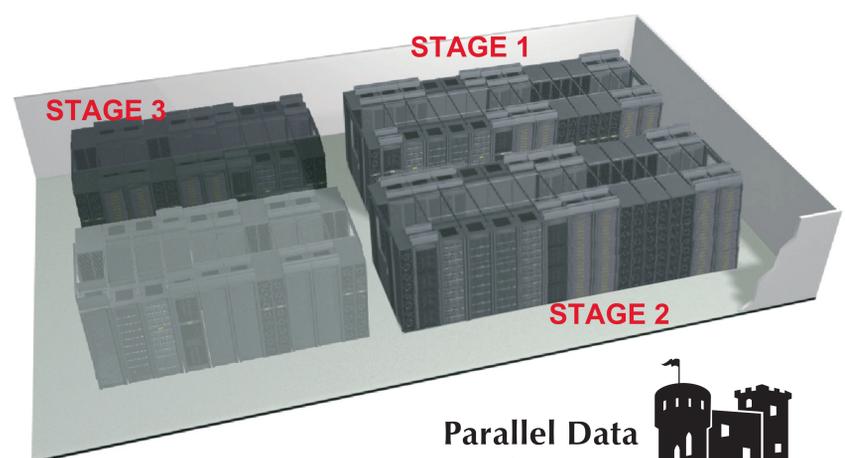
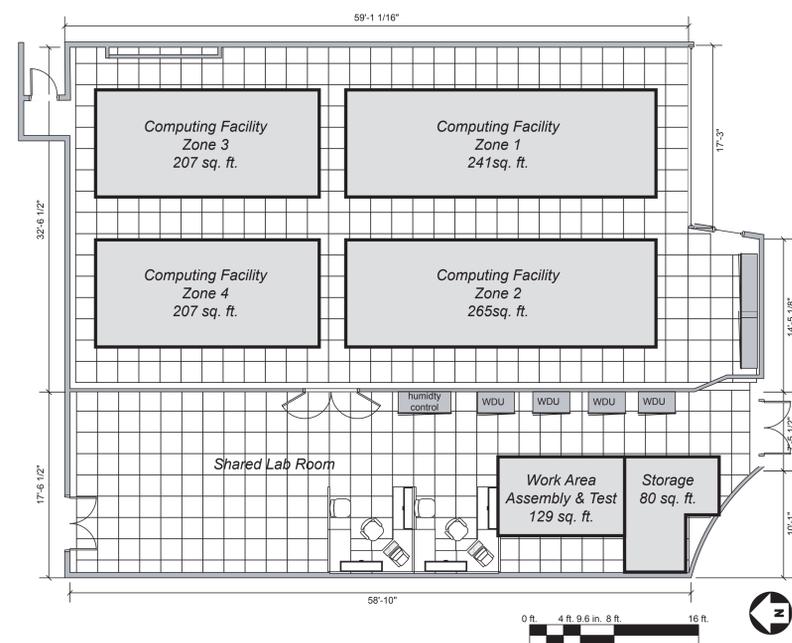
Carnegie Mellon University
CyLab

Deeply Instrumented

- Physical sensors
 - server to room level power metering
 - air temperature and chilled water flow
 - fan speeds, humidity, leaks, etc.
- Operational records
 - administrator tasks and times
 - component, system, application failures
 - per-customer resource utilization
- Activity tracing
 - I/O, network, etc.

Deployment Schedule

- Staged deployment plan as follows:
 - Engineering plans complete
 - Construction began November 2005
 - Room construction completed March 2006
 - Zone 1 came online April 2006
 - Zone 2 came online October 2008
 - Zone 3 came online November 2012
 - Zone 4 planned for 2015



Parallel Data
Laboratory  PDLR13

