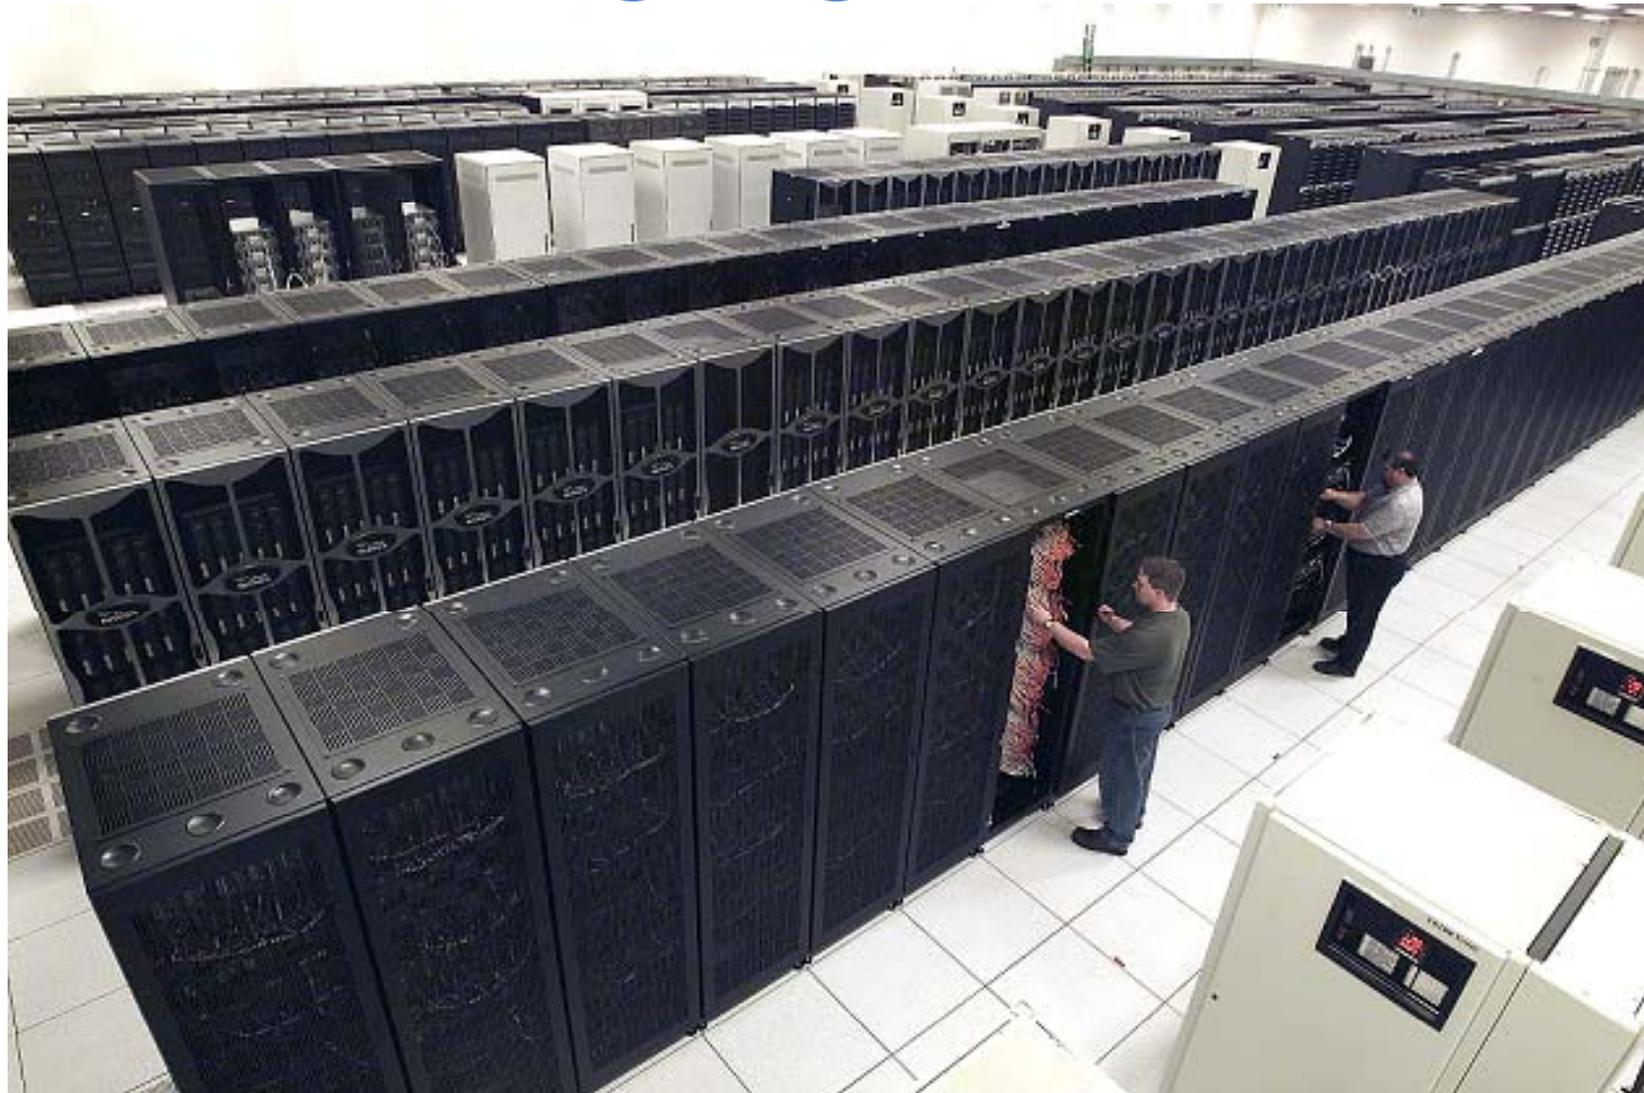


# PRObE

An NSF large scale systems research center  
in Los Alamos, New Mexico

<http://www.newmexicoconsortium.org/probe>

LANL was going to trash this!



# PRObE to the rescue!

- NSF Funds the New Mexico Consortium (NMC) to bring LANL supercomputers back to life
- PRObE –

Parallel Reconfigurable

Observational Environment

# Motivation

- Systems research community lacks *very large* dedicated resource for experiments, fault injection, and hardware control.
- Research on large compute resources often constrained by imposed software stack
- Large systems are hurried through testing phase into production. Inhibits systems research at scale.
- And...

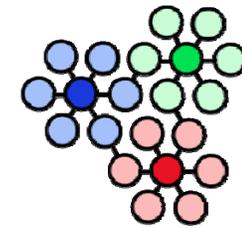
# What is PRObE?

- Low level systems research center
- Days to weeks of dedicated usage of a large computer resource for projects
  - **Physical** and remote access
- Complete control of hardware and software
- Enables fault injection and failure statistics collection
- End-of-life destructive testing
- Supports parallel and data intensive workloads

Brought to you by:



**Carnegie  
Mellon  
University**

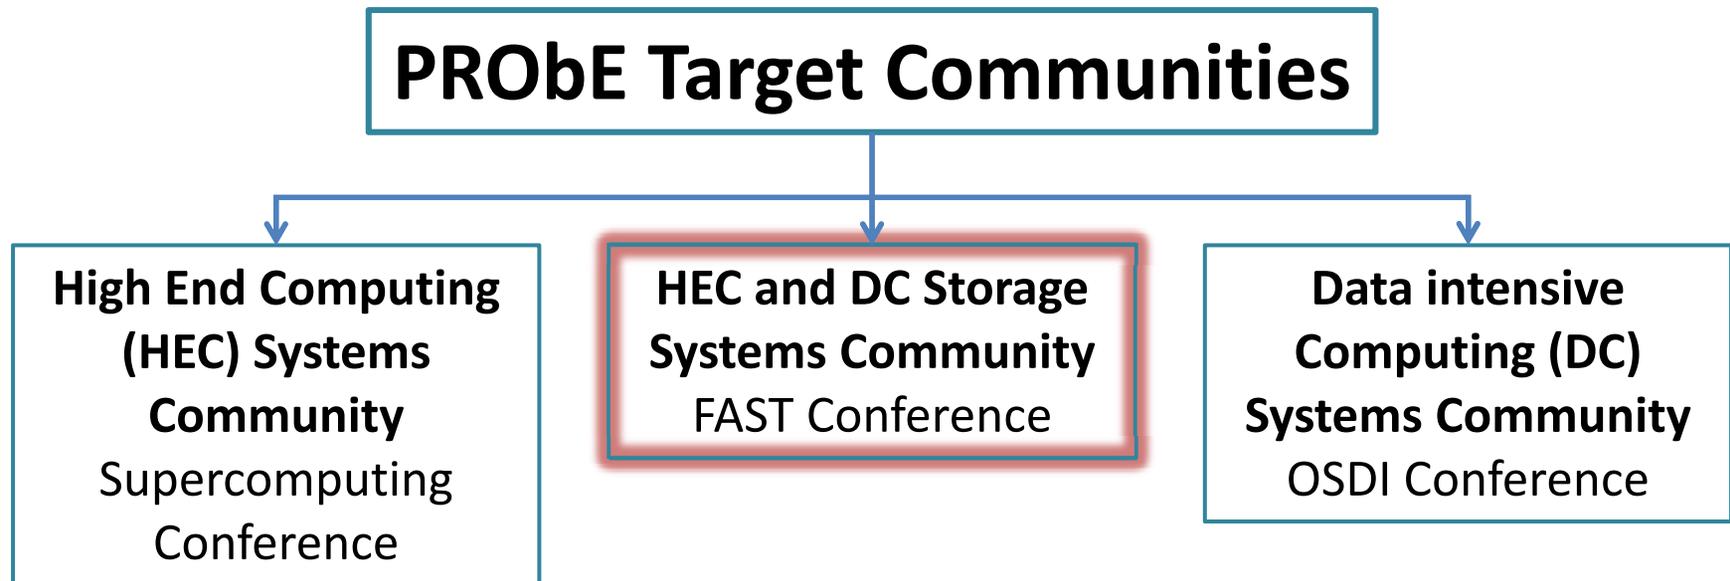


**emulab**



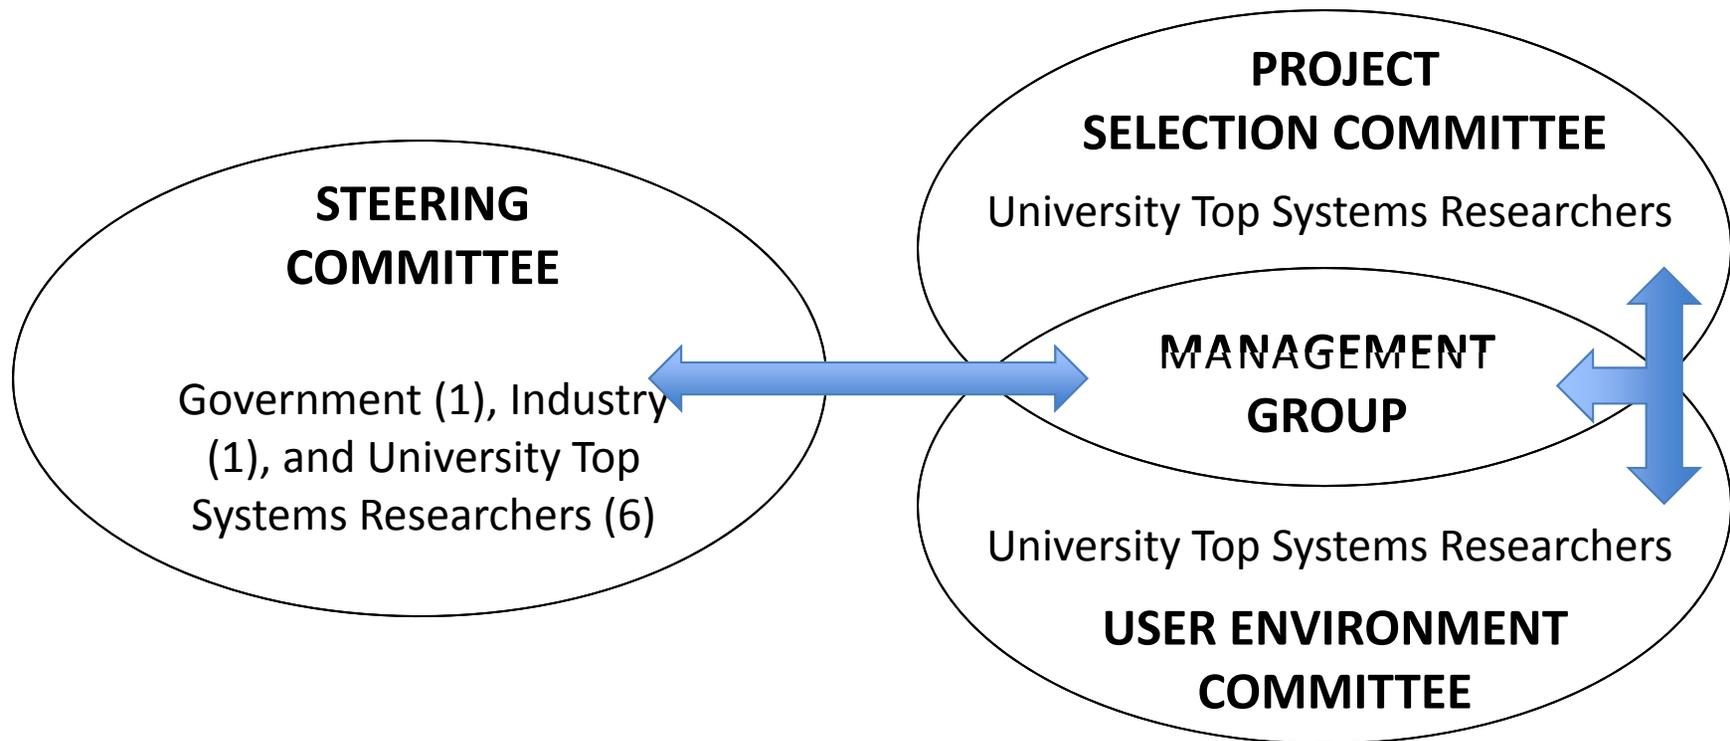
# For Systems Research Users

- NFS's "who can apply" rules
  - Includes international and corporate research projects (partnership with US university preferred)

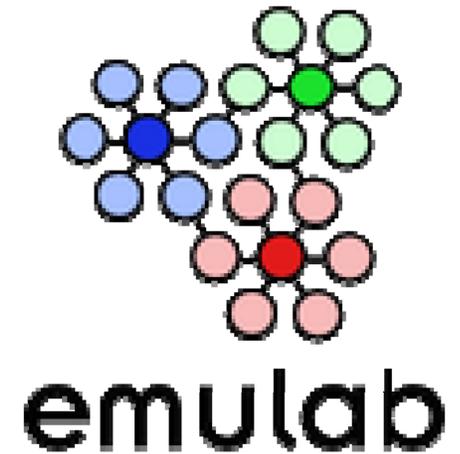


# PRObE Decision Making

3 Committees, members selected from community



# Software



- First, “none” is allowed
  - Researchers can put any software they want onto the clusters
  - Full OpenCirrus stack possible
- Second, a well known tool managing clusters of hardware for research
  - Emulab ([www.emulab.org](http://www.emulab.org)), Flux Group, U. Utah
  - Widely used in academic systems community
  - Enhanced for PRObE hardware, scale, networks, resource partitioning policies, remote power and console, failure injection, deep instrumentation

# Cluster Installation Timeline

When	Nodes	Cores	Purpose / Type	Where	Name
Q1 CY2011	128	256	Front end test cluster (IB)	CMU	Marmot
Q3 CY2011	128	256	Front end (Myri)	NMC	Denali
Q3 CY2011	36	1728	High core count cluster (IB)	CMU	Susitna
Q4 CY2011	1024	2048	High node count cluster (Myri)	NMC	Sitka
Q1 CY2012	1024	2048	High node count cluster (IB)	NMC	Kodiak
Q3 CY2013	16	128	Front end (IB)	NMC	Yakutat
Q3 CY2013	200	1600	High node count cluster (IB)	NMC	Nome
Q4 CY2013	36	3456	High core count cluster (100GigE)	CMU	Matanuska
Q2 CY2014	Next high node count cluster identified and...				

..first 1024 node cluster decommissioned to make room for next large cluster. Research contest to see how best to torture the machine on its way out will be conducted.

# Contacts

- Website
  - <http://www.newmexicoconsortium.org/probe>
    - Will soon house: Wiki's, Published data  
Committee Nomination & Election pages
- Email
  - [probe@newmexicoconsortium.org](mailto:probe@newmexicoconsortium.org)