

Survey of UC Computing Clusters

September, 2005

David Walker
Office of the President
University of California
David.Walker@ucop.edu

(Revised 10/17/2005))

The Survey

- Chartered by ITLC and Vice Chancellors for Research in September, 2005
- Intent was to *start* compiling an inventory of research computing resources within UC
- Survey developed by Bill Labate (UCLA), Tammy Welcome (LBNL), Frank Wessel (UCI), and David Walker (UCOP). Data collection by Connie Geraghty (UCOP).

Survey Questions - 1

- Cluster Name
- Contact Name
- Research Area
- # of Nodes
- # of Processors per Node
- Type of Processor and GHz
- Operating System

Survey Questions - 2

- Interconnect
- Transmission Speed of Internet Connection
- Source of System Support
- Notable Local Applications Developed
- Open Source Applications Used
- Commercial Applications Used

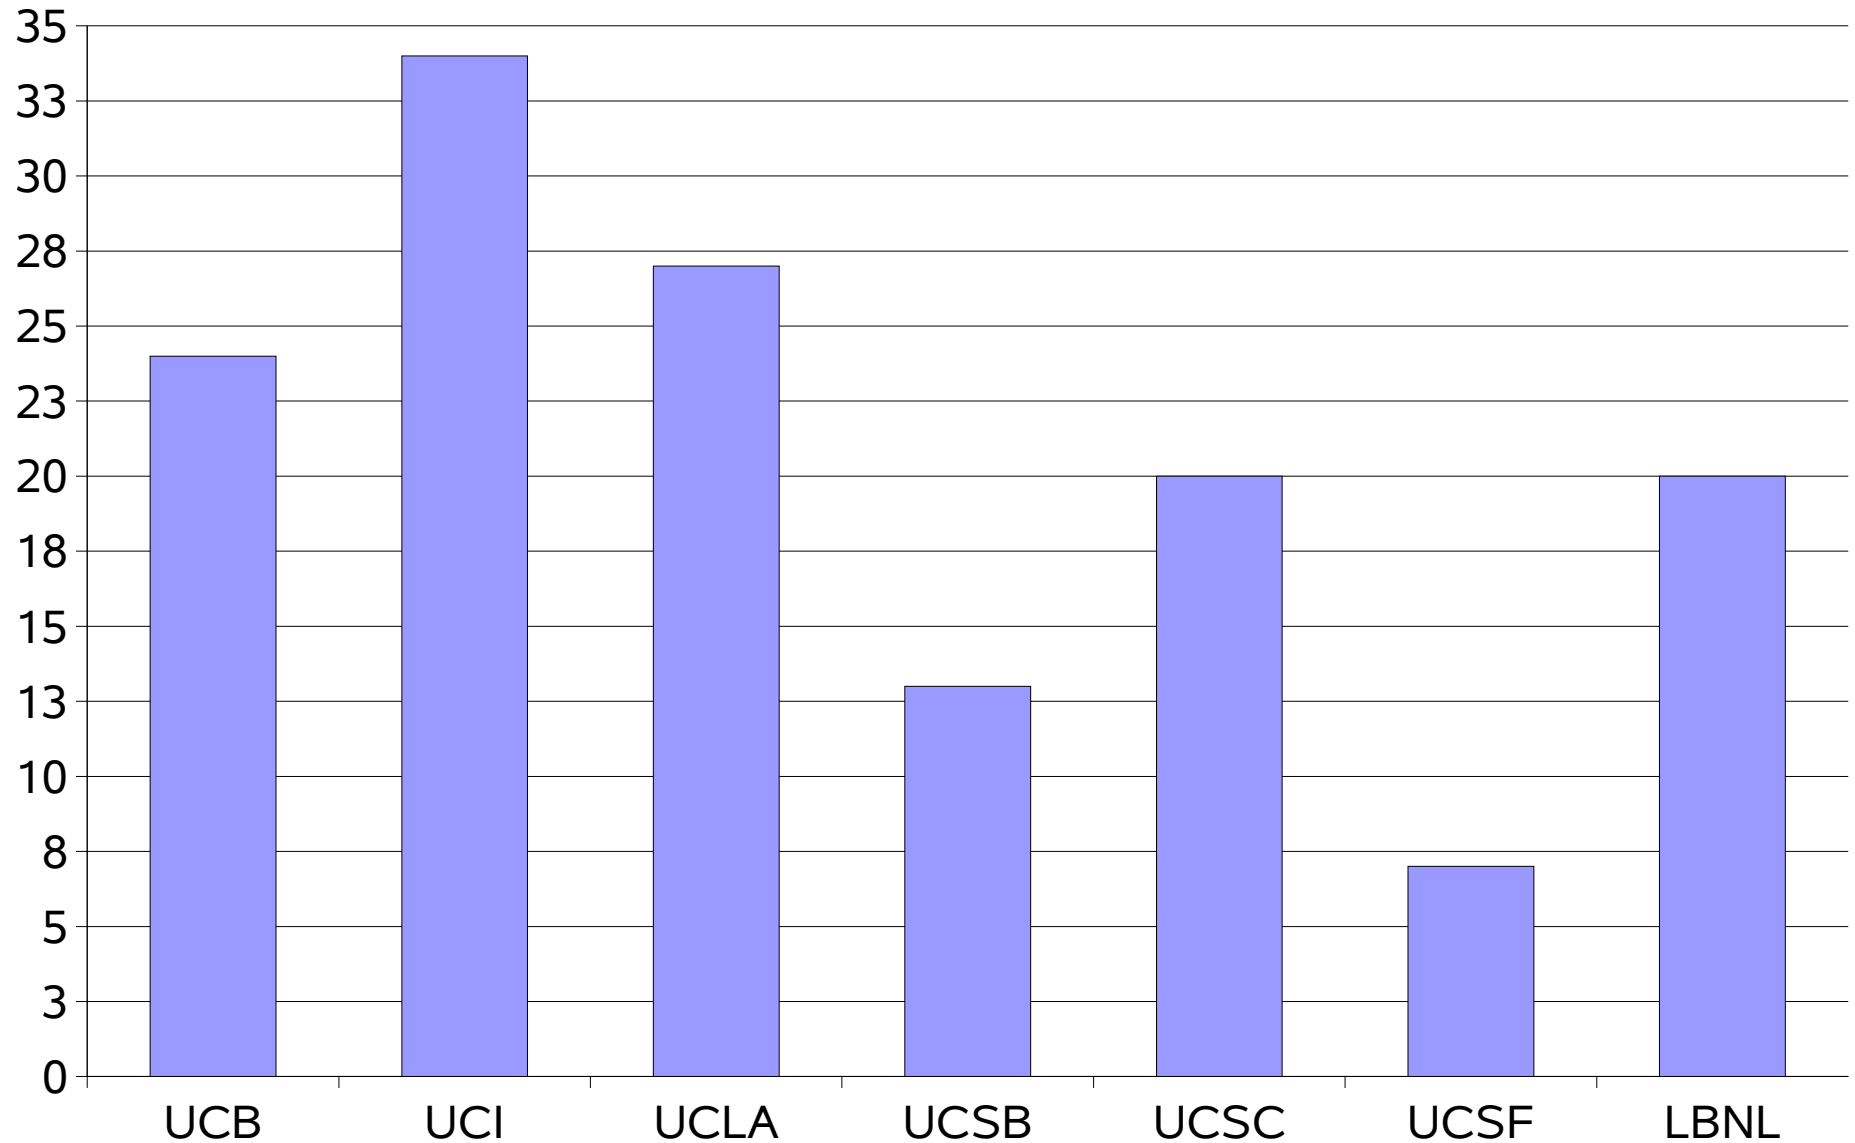
Disclaimers

- Survey designed and executed very quickly
 - Emphasis was on things that are (relatively) easy to collect.
 - Respondents did not always understand questions
- Significant amount of missing data
- Statistics related to cluster “size” are interesting but can be misleading.
 - *E.g.*, not all GHz are created equal

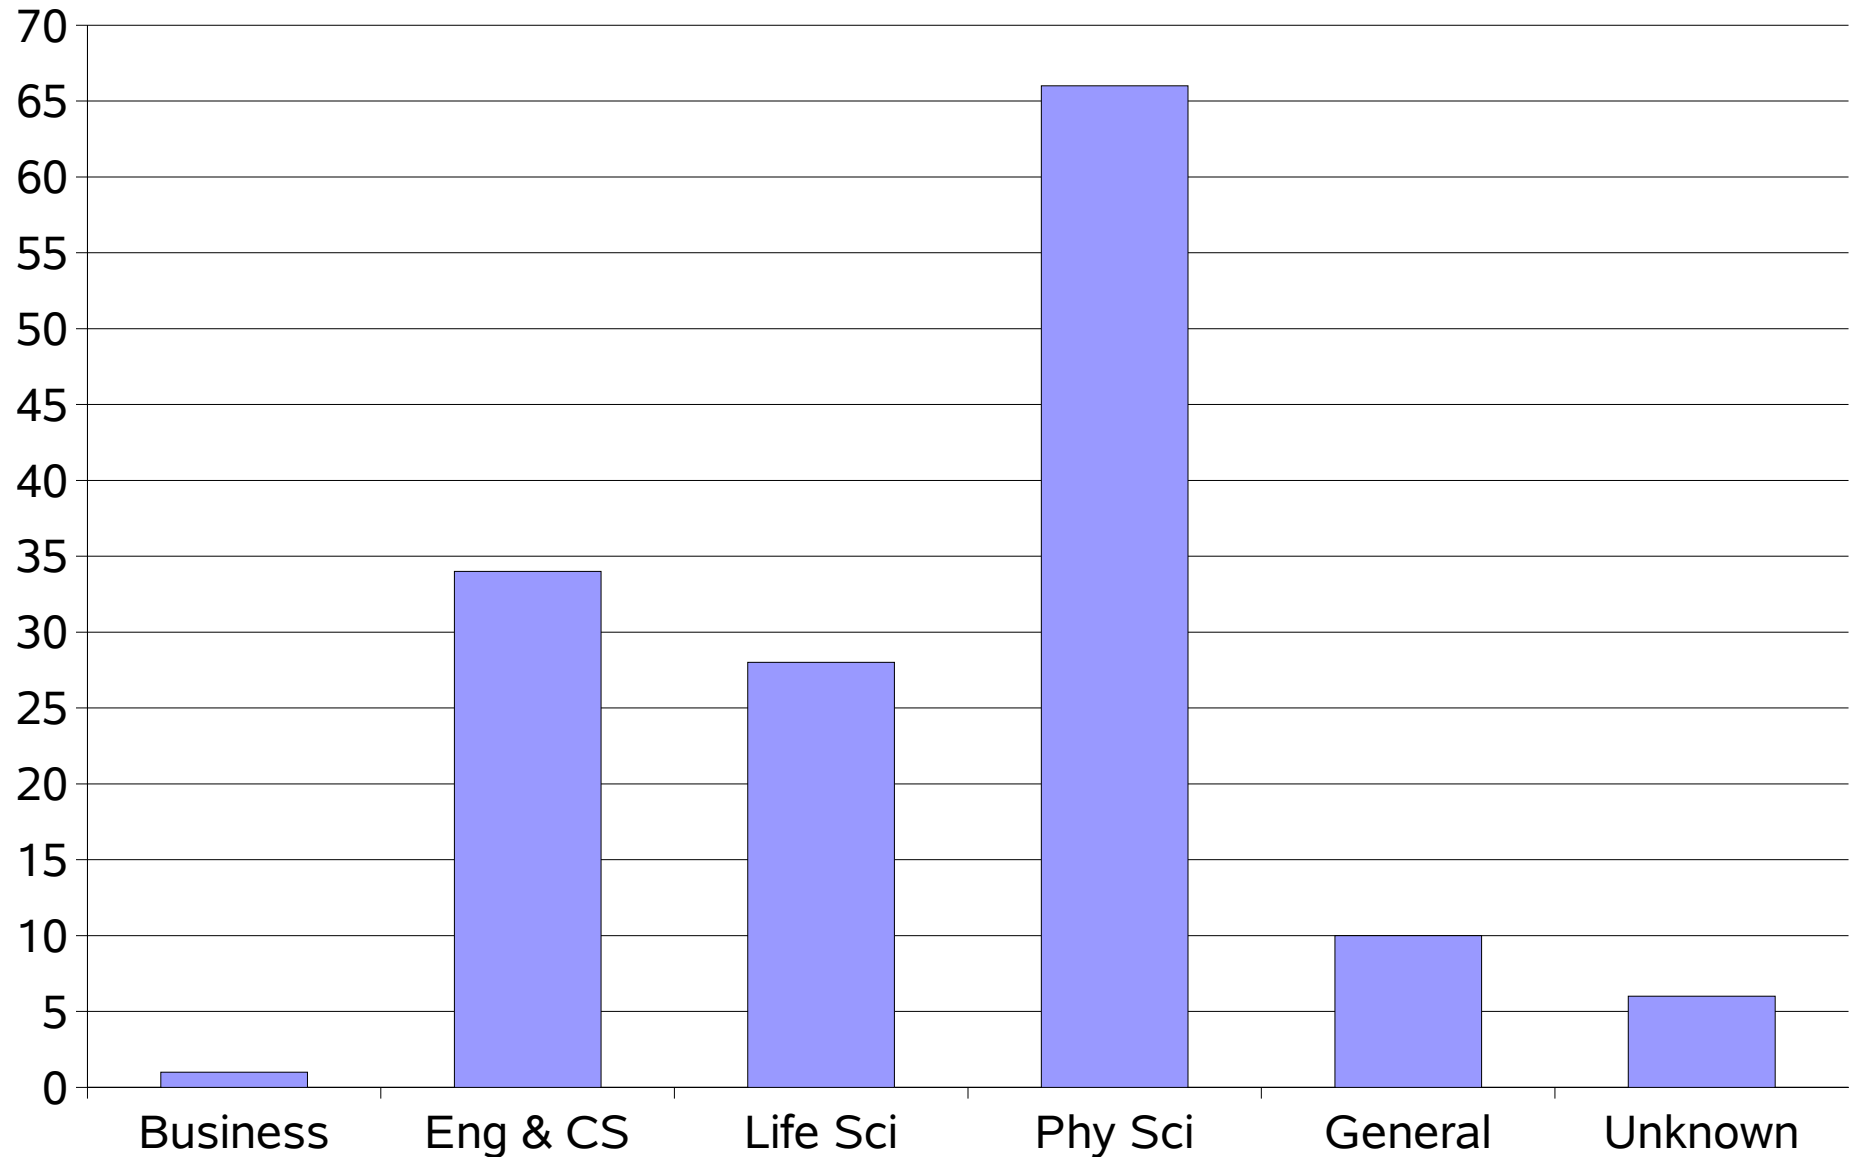
The Results

- Responses for 146 clusters at 6 campuses and 1 national lab (UCB, UCI, UCLA, UCSB, UCSC, UCSF, LBNL)
 - *Not* complete, even from responding locations
 - Results are probably more complete from locations with central support units

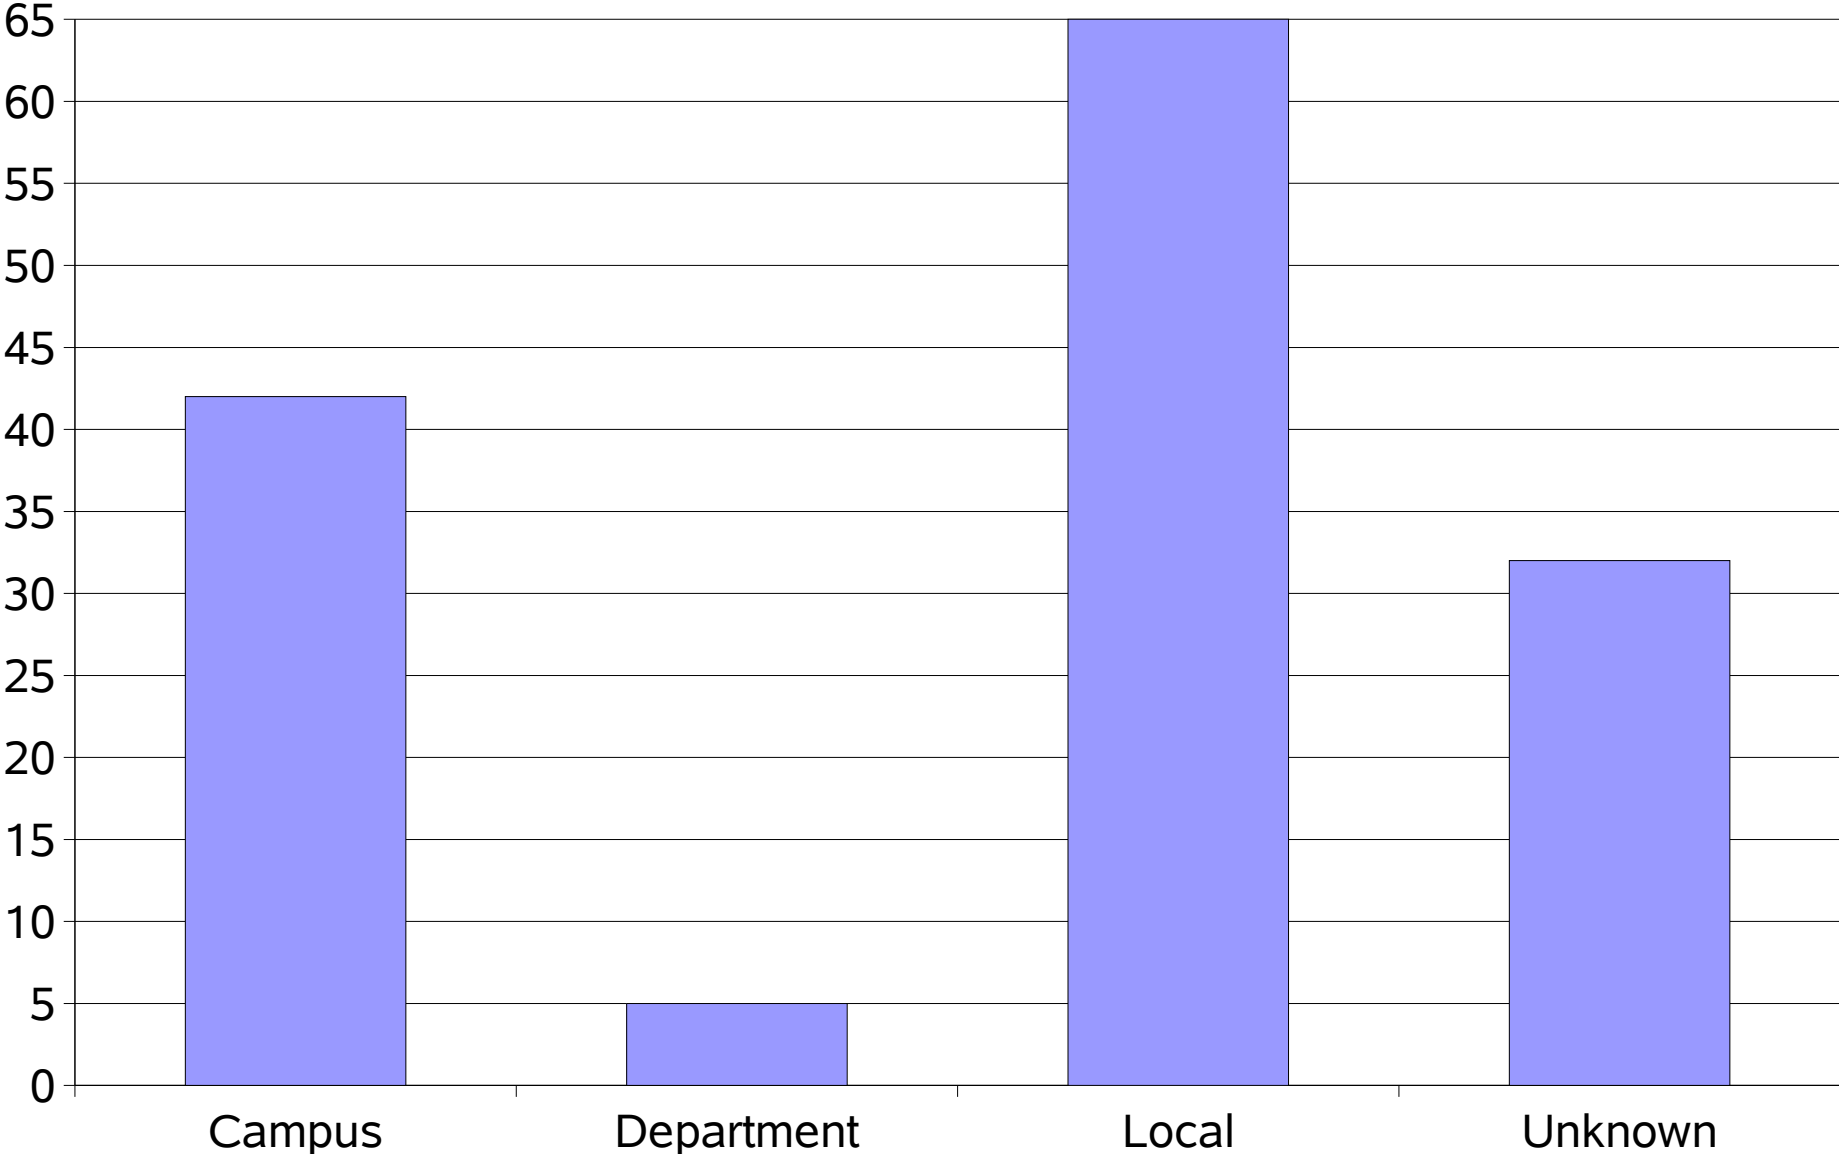
Clusters (That Responded) by Location



Clusters by Research Discipline



Clusters by System Support Organization



Operating Systems

- Linux is most common (135 of 146 clusters)
- Macintosh OS X is second with 10 clusters
- AIX, FreeBSD, Solaris, and Windows are also used
- Does Linux preponderance suggest possibility of a unified support framework for the widely distributed resource?

Interconnect and Internet Technologies

- 1 Gbps Ethernet is most common
 - 95 for interconnect, 73 for Internet
- 100 Mbps Ethernet is second
 - 18 for interconnect, 39 for Internet
- 15 clusters use parallel interconnects
 - Myrinet and Infiniband
- 3 clusters use 10 Gbps interconnect (WulfKit)
- Is Gigabit Ethernet a limiting factor?


Size Distributions

| | Median | 90th percentile | Maximum |
|--------------|--------|-----------------|---------|
| Nodes | 30 | 120 | 504 |
| CPUs | 50 | 250 | 1024 |
| Ghz | 86 | 320 | 1329 |

- Aggregate GHz (for responding clusters) is over 14,000 GHz. Assuming 2 GFlops/GHz, this is over 28 TFlops (28,000 GFlops).

Size Distribution (Example)

Clusters by # CPUs



All size metrics have similar distributions

Other Results

- 49 clusters have “notable” locally-developed applications
- Full responses available at <http://www.ucop.edu/irc/itlc/cyber.htm>