

INSTITUTIONS INCREASE NETWORKING CAPACITY, GAP BETWEEN DOCTORATE- AND NONDOCTORATE- GRANTING INSTITUTIONS WIDENS

by Leslie Christovich¹

Cyberinfrastructure resources at doctorate-granting institutions are substantially greater than at institutions that do not grant doctorates, according to new data from the biennial Survey of Science and Engineering Research Facilities, sponsored by the National Science Foundation (NSF).² This is reflected both in networking capacity, reported here as network speed or bandwidth, and in computing capacity, characterized here by the number, type, and characteristics of the computing systems.

Over a decade ago, scientific and engineering research methods at U.S. academic institutions began experiencing fundamental change. Significant advancements in computing, communications, and information technology enabled traditional research methods to be expanded or even replaced.³ Some of these technological advances allow the collection and storage of enormous amounts of data, and others have led to exponential increases in the speed with which data can be analyzed. Research-performing academic institutions continue to expand these technologies, including networking and computing capacities.

Networking

Academic institutions can have multiple networking resources. In FY 2007 all institutions had direct or indirect connections to the commodity internet—the public, multiuse network often called the Internet—but the speed

of these connections varied across institutions.⁴ In addition, many institutions had direct or indirect connections to high-performance networks that support the development and use of advanced applications and technologies.⁵ Generally, in the academic community these high-performance networks are Internet2, the National LambdaRail (NLR), and connections to federal research networks.

Total institutional bandwidth may include access to both the commodity internet and to high-performance networks. In FY 2007 total bandwidth at 65% of all research-performing academic institutions (“institutions”) was faster than 100 megabits per second (mb) (table 1). Over 34% of institutions had total bandwidth of 1 gigabits per second (gb) or faster. In FY 2005, just 2 years earlier, 21% had bandwidth of at least 1 gb. Forty-two percent of all institutions estimated they would achieve bandwidth of 1 gb in FY 2008. (The FY 2007 survey asked respondents to estimate what their capacities would be in 2008.)

The number of institutions with bandwidth at the very fastest speeds is also increasing. By FY 2007 the total bandwidth at 16 institutions (4% of all institutions) was faster than 10 gb (table 2). Six percent anticipated their institutions’ bandwidth would be greater than 10 gb in FY 2008.



TABLE 1. Total bandwidth to commodity internet and Internet2 at academic institutions, by type of institution: FY 2005–08
(Percent distribution)

Speed	All		Highest degree granted		Control		Speed	All		Highest degree granted		Control	
	institutions		Doctorate	Nondoctorate	Public	Private		institutions		Doctorate	Nondoctorate	Public	Private
FY 2005						FY 2008 (estimated)							
No bandwidth	0	0	0	0	0	0	No bandwidth	0	0	0	0	0	0
< 1.6 mb	2	1	4	2	2	*	< 1.6 mb	*	0	1	*	*	0
1.6–9 mb	3	1	8	2	5	1	1.6–9 mb	1	1	1	1	1	2
10 mb	1	1	1	1	2	*	10 mb	*	0	1	1	1	0
11–45 mb	23	15	40	20	29	8	11–45 mb	8	5	15	6	12	
46–99 mb	16	16	15	17	14	10	46–99 mb	10	6	21	9	13	
100 mb	3	2	4	3	2	6	100 mb	6	5	8	4	10	
101–155 mb	9	9	9	11	5	8	101–155 mb	8	8	10	9	7	
156–622 mb	18	24	5	18	19	18	156–622 mb	18	19	15	15	23	
623–999 mb	3	4	0	4	2	4	623–999 mb	4	6	1	5	3	
1–2.5 gb	15	16	13	16	14	26	1–2.5 gb	26	27	23	31	16	
2.6–9 gb	4	5	1	5	2	5	2.6–9 gb	5	7	1	5	5	
10 gb	*	*	0	*	0	5	10 gb	5	6	1	6	3	
> 10 gb	2	3	0	2	3	6	> 10 gb	6	9	0	6	5	
Other	*	*	0	0	1	1	Other	1	1	1	2	0	
Institutions (number)	449	312	137	301	148	Institutions (number)	448	312	136	301	147		
FY 2007													
No bandwidth	0	0	0	0	0								
< 1.6 mb	*	0	1	*	0								
1.6–9 mb	2	1	3	1	2								
10 mb	*	0	1	*	0								
11–45 mb	13	9	23	11	18								
46–99 mb	14	10	21	11	18								
100 mb	7	4	13	6	9								
101–155 mb	10	9	10	11	7								
156–622 mb	17	21	8	17	18								
623–999 mb	4	6	1	5	3								
1–2.5 gb	24	26	19	27	16								
2.6–9 gb	4	6	1	5	3								
10 gb	2	2	0	2	1								
> 10 gb	4	5	0	3	4								
Other	0	0	0	0	0								
Institutions (number)	448	312	136	301	147								

* = greater than 0 but less than 0.5%.

gb = gigabits per second; mb = megabits per second.

NOTES: Detail may not sum to 100% due to rounding. Internet2 is a high-performance backbone network that enables development of advanced Internet applications and deployment of leading-edge network services to member colleges, universities, and research laboratories across the country.

SOURCE: National Science Foundation/Division of Science Resources Statistics, Survey of Science and Engineering Research Facilities.

TABLE 2. Academic institutions with total bandwidth greater than 10 gigabits: FY 2007

Institution	State
Claremont Graduate U.	California
U. California-Davis	California
U. California-Los Angeles	California
U. California-San Diego	California
Florida Atlantic U.	Florida
Florida International U.	Florida
Nova Southeastern U.	Florida
Emory U.	Georgia
Georgia Institute of Technology	Georgia
Northwestern U.	Illinois
U. Chicago	Illinois
Indiana U.	Indiana
Harvard U.	Massachusetts
U. Nebraska-Lincoln	Nebraska
U. Tennessee	Tennessee
U. Utah	Utah

SOURCE: National Science Foundation/Division of Science Resources Statistics, Survey of Science and Engineering Research Facilities, FY 2007.

In FY 2007 a majority of institutions (70%) had access to Internet2, one of the two major high-performance networks (table 3). A growing number had multiple high-performance network access. Seventy-two percent of institutions anticipated access to Internet2 by FY 2008. Although a smaller number of institutions had connections to the NLR, the percentage of institutions with NLR connections more than doubled from 10% in FY 2005 to 25% in FY 2007. Thirty-one percent estimated that their institutions would have NLR access in FY 2008.

Internal network speeds also increased at academic institutions (not shown in tables). The percentage of institutions with distribution speeds of at least 1 gb increased by 20 percentage points between FY 2005 and FY 2007, from 54% to 74% of all institutions. Seventy-eight percent anticipated having speeds of at least 1 gb by FY 2008.

The amount of dark fiber owned by institutions indicates an ability to expand existing network capabilities.⁶ Dark fiber may exist between campus buildings or from the campus to the institution's internet service provider. The percentage of academic institutions with dark fiber to their internet service providers increased significantly from 29% in FY 2005 to 37% in FY 2007.

The percentage of institutions with dark fiber between buildings remained stable.

Networking Capacity and Institutional Characteristics

Doctorate- and Nondoctorate-Granting Institutions

Bandwidth capacity at doctorate-granting institutions was significantly greater than at nondoctorate-granting institutions and was also more concentrated at the higher speeds (1 gb or more). In FY 2007 the percentage of doctorate-granting institutions with bandwidth of at least 1 gb (39%) was almost twice that of nondoctorate-granting institutions (20%) (table 1). Over half (62%) of the nondoctorate-granting institutions had bandwidth of 100 mb or less, compared with 24% of doctorate-granting institutions. In FY 2007 all of the 16 institutions with total bandwidth of at least 10 gb were doctorate-granting (table 2). The difference between doctorate and nondoctorate institutions likely increased in FY 2008, when almost half of doctorate-granting institutions estimated they would have bandwidth of 1 gb or faster, compared with 25% of nondoctorate-granting institutions.

The percentage of doctorate-granting institutions with access to Internet2 remained stable from FY 2005 to FY 2008: 82% in FY 2005 and 83% (anticipated) in FY 2008 (table 3). The percentage with NLR access increased significantly from 11% in FY 2005 to 32% in FY 2007, and 37% estimated having NLR access in FY 2008. The increase in NLR access was due to institutions adding multiple networking connections. Nondoctorate institutions increased their access to Internet2, from 38% in FY 2005 to 46% FY 2007. Access to NLR was expected to increase to 15% in FY 2008 for nondoctorate institutions.

Public and Private Institutions

The gap in bandwidth capabilities between public and private institutions was not as great as the gap between doctorate and nondoctorate institutions. Further, the differences between public and private institutions completely disappeared at the very highest bandwidth speeds. In FY 2007, 37% of public institutions had bandwidth of 1 gb or faster compared with 24% of private institutions (table 1). However, at the very fastest speeds of 10 gb or greater, the percentages were equal. Forty-eight percent of public institutions estimated their

TABLE 3. Institutions with high-performance network connections, by type of institution: FY 2005–08 (Percent)

Type of institution	Internet2	National LambdaRail	Federal government research network	State or regional high-performance network	Other
End of FY 2005					
All academic institutions	68	10	11	na	12
Doctorate granting	82	11	13	na	15
Nondoctorate granting	38	7	6	na	6
Public	73	11	12	na	14
Private	58	8	9	na	9
End of FY 2007					
All academic institutions	70	25	11	55	3
Doctorate granting	81	32	13	59	4
Nondoctorate granting	46	10	4	43	1
Public	75	29	12	61	3
Private	61	17	8	41	3
End of FY 2008 (estimated)					
All academic institutions	72	31	14	62	4
Doctorate granting	83	37	18	67	5
Nondoctorate granting	46	15	6	50	3
Public	75	37	16	71	4
Private	64	17	10	43	5

na = not applicable; data were not collected in FY 2005.

NOTES: Internet2 is a high-performance backbone network that enables development of advanced Internet applications and deployment of leading-edge network services to member colleges, universities, and research laboratories across the country. An institution may have a connection to more than one high-performance network. National LambdaRail is an advanced optical network infrastructure for research and education.

SOURCE: National Science Foundation/Division of Science Resources Statistics, Survey of Science and Engineering Research Facilities.

institutions would have bandwidth of 1 gb or faster in FY 2008, whereas 29% of private institutions did.

The percentages both of public and private institutions with access to Internet2 increased slightly from FY 2005 to FY 2007 (table 3). Another small increase was anticipated by private institutions for FY 2008. Both public and private institutions experienced significant increases in the percentage of institutions with NLR access from FY 2005 to FY 2007. Public institutions anticipated further significant increases in FY 2008, but private institutions did not.

High-Performance Computing

The survey requested that respondents provide information on their high-performance computing (HPC) sys-

tems, defined as computing systems operating at a speed of at least 1 teraflop. These systems could include a variety of computing architectures, such as computing clusters or symmetric multiprocessor systems.⁷ Historically, academic HPC capability has been available to recipients of research grants, departments of research grant recipients, or academic departments that make heavy use of HPC, but it has not been available campus-wide. In recent years a greater number of academic institutions have been expanding the availability of HPC to their entire campus community by centralizing the ownership and management of these resources.

In FY 2007 approximately 22% of research performing academic institutions made at least some of their high-performance computing generally available to their

campus community (table 4).⁸ Of the institutions with HPC, clusters were the most prevalent (94%), followed by symmetric multiprocessor systems (21%). Of all institutions with clusters, the most common peak performance was 3 teraflops or faster, with 39 institutions reporting clusters at this performance. Almost all HPC resources were located at doctorate-granting institutions. The seven nondoctorate-granting institutions with HPC capability all reported having clusters.

Colleges and universities tended to share their HPC resources with organizations outside their own institutions (table 5). In FY 2007 institutions with HPC capability were most likely to share with other colleges and universities (71%), with governments, industry, and nonprofit organizations constituting most of the other users (table 6). Doctorate granting institutions were more likely to have external users of their HPC than were nondoctorate granting institutions.

Data Sources

The data presented in this report were obtained from a census of colleges and universities that expended at least \$1 million in science and engineering research and developments funds. Each institution's level of expenditures was obtained from the NSF FY 2006 Survey of Research and Development Expenditures at Universities and Colleges.

The full set of detailed tables from the FY 2007 Survey of Science and Engineering Research Facilities will

TABLE 4. Centrally administered high-performance computing, by type of institution and computing architecture: FY 2007

Type of institution	Institutions		Type of architecture			
	Total	With HPC	Clusters	MPP	SMP	Other
All academic institutions	449	99	93	16	21	18
Doctorate granting	313	92	86	16	21	18
Nondoctorate granting	136	7	7	0	0	0
Public	302	73	69	9	18	14
Private	147	26	24	7	3	4

HPC = high-performance computing; MPP = massively parallel processors; SMP = symmetric multiprocessors.

NOTES: Each institution is counted only once in each architecture, but an individual institution may have multiple architectures. Centrally administered HPC is located within distinct organizational unit with staff and budget; unit has a stated mission that includes supporting HPC needs of faculty and researchers. Institutions may have HPC of more than one type of architecture.

SOURCE: National Science Foundation/Division of Science Resources Statistics, Survey of Science and Engineering Research Facilities, FY 2007.

be available at <http://www.nsf.gov/statistics/facilities/>. Individual detailed tables from the 2007 survey may be available in advance of publication of the full report. For further information, please contact the author. Current survey data for individual institutions are available from the Computer-Aided Science Policy Analysis and Research (WebCASPAR) database system, a Web tool for retrieval and analysis of statistical data on science and engineering resources (<https://webcaspar.nsf.gov/>).

TABLE 5. Academic institutions that share their high performance computing with other organizations, by type of organization: FY 2007

Institution	Type of organization					Institution	Type of organization				
	College/university	Government	Nonprofit organization	Industry	Other		College/university	Government	Nonprofit organization	Industry	Other
Arizona State U.	x	x		x		Tufts U.	x				
Baylor U.	x					Tulane U.	x				
Boston U.	x	x	x	x		U. of Alabama at Birmingham	x		x		
Bowling Green State U.	x					U. of Arkansas	x			x	
Brigham Young U.	x		x	x		U. of California-Irvine	x				
California Institute of Technology	x	x				U. of California-Los Angeles	x			x	
Clemson U.	x		x			U. of California-San Diego	x	x	x	x	
College of William and Mary	x					U. of Connecticut	x				
Cornell U.	x	x	x	x		U. of Dayton	x				
Duke U.	x					U. of Florida	x				
Emory U.	x					U. of Houston	x	x		x	
Florida State U.	x					U. of Illinois at Urbana-Champaign	x			x	
Georgia Institute of Technology	x			x		U. of Iowa	x				
Georgia State U.	x					U. of Kentucky	x				
Harvard U.	x					U. of Maryland Baltimore County	x	x			
Indiana State U.	x	x				U. of Minnesota	x	x			
Indiana U.	x					U. of Mississippi	x				
Louisiana State U.	x					U. of Missouri	x				
Louisiana Tech U.	x					U. of Nebraska-Lincoln	x				
Miami U.	x					U. of Nebraska-Omaha	x	x	x	x	
Mississippi State U.	x	x				U. of New Hampshire	x				
Missouri U. of Science and Technology	x					U. of New Mexico	x				
New York U.	x		x			U. of North Carolina at Charlotte	x				
North Carolina Agricultural and Technical State U.	x					U. of North Dakota	x	x			
North Carolina State U.	x	x				U. of Notre Dame	x	x	x	x	
Pennsylvania State U.	x	x	x	x		U. of Rochester	x				
Princeton U.	x	x				U. of Tennessee	x	x			
Purdue U.	x	x	x	x	x	U. of Texas at Austin	x	x	x	x	
Rensselaer Polytechnic Institute	x			x		U. of Utah	x	x			
Rice U.	x	x		x		U. of Washington				x	
Saint Louis U.	x					Vanderbilt U.	x				
Stanford U.	x					Virginia Commonwealth U.	x	x	x		
State U. of New York at Buffalo	x	x	x	x		Virginia Polytechnic Institute and State U.	x	x	x	x	
State U. of New York College at Geneseo	x			x		Washington U. in St. Louis	x				
Texas Tech U.	x					Wesleyan U.	x				
						Wichita State U.	x				

SOURCE: National Science Foundation/Division of Science Resources Statistics, Survey of Science and Engineering Research Facilities, FY 2007.

TABLE 6. Academic institutions with external users of their centrally administered high-performance computing, by type of user: FY 2007

Type of institution	Institutions with HPC		Type of user				
	All	With external users	Colleges/universities	Governments	Nonprofit organizations	Industry	Other
All academic institutions	99	71	70	24	15	19	4
Doctorate granting	92	70	69	24	15	18	4
Nondoctorate granting	7	1	1	0	0	1	0
Public	73	49	48	18	10	14	3
Private	26	22	22	6	5	5	1

HPC = high-performance computing.

NOTE: Institutions may provide HPC to more than one type of external user.

SOURCE: National Science Foundation/Division of Science Resources Statistics, Survey of Science and Engineering Research Facilities, FY 2007.

Notes

1. Leslie Christovich, Research and Development Statistics Program, Division of Science Resources Statistics, National Science Foundation, 4201 Wilson Boulevard, Suite 965, Arlington, VA 22230 (lchristo@nsf.gov; 703-292-7782).

2. The Survey of Science and Engineering Research Facilities collects data from academic institutions and nonprofit biomedical research institutions (hospitals and research organizations) receiving research funds from the National Institutes of Health. Networking and computing capabilities at biomedical institutions are not reported here.

3. National Science Foundation. 2003. *Revolutionizing Science and Engineering through Cyberinfrastructure: Report of the National Science Foundation Blue-Ribbon Advisory Panel on Cyberinfrastructure*. Arlington, VA. Available at <http://www.nsf.gov/od/oci/reports/atkins.pdf>.

4. *Fiscal year* in this report refers to each institution's fiscal year and thus varies across institutions. For example, for some it may be 1 January to 31 December, and for others it may be 1 July to 30 June.

5. A high-performance network is characterized by high bandwidth, low latency, and low rates of packet loss.

Additionally, a high-performance network is able to support delay-sensitive, bandwidth-intensive applications such as distributed computing, real-time access, and control of remote instrumentation.

6. Dark fiber is unused fiber within fiber optic cables that have already been laid; thus, it is available for future use.

7. Cluster architectures use multiple commodity systems with an Ethernet-based or high-performance interconnect network to perform as a single system. Symmetric multiprocessor systems use multiple processors sharing the same memory and operating system to simultaneously work on individual pieces of a program.

8. Survey respondents were asked to report only their centrally administered high-performance computing. Centrally administered was defined as high-performance computing that is located within a distinct organizational unit with a staff and a budget and is generally available to the campus community. The unit has a stated mission that includes supporting high-performance computing needs of faculty and researchers. This Info-Brief only reports on centrally administered HPC.

NATIONAL SCIENCE FOUNDATION

ARLINGTON, VA 22230

OFFICIAL BUSINESS

PRESORTED STANDARD
U.S. POSTAGE PAID
National Science Foundation

RETURN THIS COVER SHEET TO ROOM P35 IF YOU DO NOT WISH TO RECEIVE THIS MATERIAL , OR IF CHANGE OF ADDRESS IS NEEDED , INDICATE CHANGE INCLUDING ZIP CODE ON THE LABEL (DO NOT REMOVE LABEL).

NSF 10-328