

# Information for New Users of the NIH Helix Systems

---

The High Performance Scientific Computing Section of the Computing Facilities Branch currently supports four systems: a Silicon Graphics Inc (SGI) Challenge multiprocessor system (*helix.nih.gov*), a SGI Power Challenge (*churn.nih.gov*), a SGI Origin 2000 (*galaxy.nih.gov*) and an IBM Scalable Powerparallel (SP) system (*spica.nih.gov*). *Helix* and *churn* are both open system platforms which run the UNIX operating system and utilities, *galaxy* is a 32-processor parallel computer that runs the UNIX operating system and utilities and *spica* is a 104-node parallel system that runs the AIX operating system and utilities. Software and applications are tailored to the needs of the NIH scientific community. Machines are configured as a heterogeneous processing environment in which *helix* is a general purpose system and *churn* is a computation server which runs applications and programs optimally configured for its hardware capabilities. *Spica* and *galaxy* are designed for the development and execution of parallel programs. For more information on *spica*, see the web page: <http://helix.nih.gov/documents/sp-guide.html>. For more information on *galaxy*, see the web page: <http://helix.nih.gov/platforms/galaxy.html>.

Most users will have their computing needs met by the SGI Challenge system (*helix*) where the following services are available: **mail**, **POP**, **news**, anonymous **ftp**, **web browsers**, **ph**, and **archie**. However, some users will need to use *churn*, *galaxy* or *spica* in addition to *helix*. The table at the end of this document (“*Summary of Network Services and Scientific Applications*”) should help you decide which system meets your needs. If you have a *helix* account, you can login to *churn* and *galaxy*. If you need an SP account, visit the following web page: <http://helix.nih.gov/register.html>.

## Connecting to the Helix Systems

### Network Access

The SGI Challenge system has the Internet network name *helix.nih.gov*, the SGI PowerChallenge system has the name *churn.nih.gov*, the SGI Origin 2000 has the network name *galaxy.nih.gov* and the IBM SP system has the name *spica.nih.gov*. Users connecting to the Helix Systems should use the above hostnames with the **telnet** program.

### Dial-in Access

Dial-in service is available to *helix*. Call 301-402-2222 or, for long distance access, 800-358-2022.

## How to Log in and Log out

After you are connected, the system will type out “login:”. Enter your username (using only lowercase letters), and press the “return” key. The system will then prompt you for your password. Enter it, and press the “return” key. Your password will not appear on your terminal when you type it in. Remember that UNIX is case-sensitive, so you must enter your username and password exactly. At this point the system will prompt you to choose a new password. Your new

password must:

- be at least 6 characters
- include 2 letters (a-z) and 2 special characters (punctuation, digits)
- not have a special character as the first or last character
- differ from your old password by at least 3 characters.

When you are finished with your session you can log out by typing the **logout** command.

## Documentation

Important announcements regarding such things as new software, system availability and changes in policy are made using messages printed out when you login to *helix*. It is your responsibility to read and review this information. To read current messages from the *helix%* prompt, type **msgs**. The most recent messages are also posted on our web pages at <http://helix.nih.gov/whatsnew.html>.

Printed documentation is available to registered users through the Technical Assistance and Support Center (TASC), Building 12A, Room 1011. Users can order documentation by issuing the **pubware** command and following the instructions or visiting the URL <http://publications.cit.nih.gov/>.

For more information on topics mentioned in this write-up, and UNIX in general, users should request the “*Learning the UNIX Operating System*” book. More in-depth descriptions are found in “*A Student’s Guide to Unix*.”

## Getting On-line Help

The UNIX **man** command gives on-screen information from the UNIX “man” pages. These pages contain reference material designed to help you find information about a specific topic. The command **apropos** lets you search for “man” page entries based on a topic or keyword. Given a keyword, **apropos** lists all commands that are related to it.

Here is a list of a few useful man commands:

<b>man man</b>	get help about the man command
<b>man software</b>	get information about software available on the NIH Helix systems
<b>man docs</b>	get information about available documentation

You can also access the man pages at the following URL: <http://helix.nih.gov/manpage.html>

## Printing Files

Text and PostScript files can be printed to a networked printer or via the central IBM System/370 printers by using the **print370** command. Type **man lpr** to find out how to print to a

networked printer and **man print370** to find out how to use the **print370** command.

## Startup Files

When you log in, the operating system creates a login shell for you. This shell looks in your home directory for your `.cshrc` file and follows the instructions it finds in this file. Next, the shell looks for a file called `.login` and follows the instructions in it. You can list your startup files using the command **dots**. The `.cshrc` file is the startup file most commonly edited by users, and includes instructions on changing default settings.

## Mail

Electronic mail can be sent and received on helix using the **pine** command. An example of sending mail to user smith at college.edu follows:

```
helix% pine smith@college.edu RET
To       :smith@college.edu RET
CC:      : RET
Attchmnt: RET
Subject  : Seminar on Friday RET
Your seminar is scheduled here at NIH for next Friday. See you then.
CTRL-X
Send message? [y] : RET
helix%
```

You can read mail by typing **pine** at the helix% prompt. Other mail programs available on helix are UNIX **mail** and GNU Emacs **rmail**.

## Editors

Several line- and screen-oriented editors are available on both systems. Because screen editors let you see and access many lines at a time, they are usually preferred over line editors. The screen editors available are **GNU Emacs**, **vi**, **pico** and **edt**. The line editors available are **ed** and **ex**.

## Transferring Files to the Helix Systems

From the helix systems, you can use **ftp** to transfer files to and from other Internet hosts. Users with personal computers can use **kermit** for transferring files.

## Backups

The helix systems are backed up regularly by the Operations staff. Full backups are done weekly, and incremental backups are done daily. If you need to have files restored, send e-mail to the operators at the address `operators@helix.nih.gov`.

## Assistance and/or Problems

Please send comments or problem reports by typing the **ptr** command. **ptr** prompts for information and sends the completed report to the Helix Systems Staff. You can also call CIT's TASC at (301) 594-6248.

## Keeping You Informed

Messages displayed at login will keep you informed of new developments. Call us if you are interested in particular facilities. News and events are also posted on our web site at: <http://helix.nih.gov/>.

# Hardware Configurations

## The NIH Challenge System

The NIH Challenge, known as *helix*, is a 6-processor SGI Challenge XL system. The processors are based on the MIPS R4400 150 MHz chip. Each processor has a 32 KB primary and a 1 MB secondary cache, with peak performance of 75 MFLOPS (64-bit words). Helix has a total capacity of 4.5 Gbytes of virtual memory; 1.5 Gbytes is 8-way interleaved memory and 3 Gbytes is used as swap space and has 80 GB total disk capacity.

## The NIH Origin 2000

The NIH Origin 2000, known as *galaxy*, is a 32-processor SGI Origin 2000. Each processor is a 250MHz MIPS R10000 (IP31) processor with 4 MBytes of secondary cache. The peak floating point performance is 16 GFLOPS. Galaxy has a total of 8 GBytes system memory and appears as global shared memory to the programmer.

**Technical note:** while the system appears as a large shared memory computer, the memory, in fact, is distributed across 16 "nodes". The technology used to make this distributed memory appear equally accessible from any processor is called **ccNUMA** (*cache coherent non-uniform memory access*).

## The NIH SP System

The NIH SP, known as *spica*, is a 104-node system, each of which is a fully contained RS/6000 processor with disk and memory. There are eight wide nodes, each with 512 MB of memory and 96 thin nodes, each with 128 MB of memory. The nodes are interconnected by a low-latency, high-performance switch network.

# Summary of Network Services and Scientific Applications

	<i>helix.nih.gov</i> (SGI Challenge)	<i>churn.nih.gov</i> (SGI Power Challenge)	<i>galaxy.nih.gov</i> (SGI Origin 2000)	<i>spica.nih.gov</i> (IBM SP)
Mail (mail, pine, POP)	x			
POP Server	x			
News (rn, xrn, tin)	x			
mosaic, lynx, netscape	x			
archie, xarchie	x			
ph	x			
Anonymous ftp	x			
ftp	x	x	x	x
telnet	x	x	x	x
AVS	x			
Mathematica	x			
MATLAB	x	x		
MATLAB Toolboxes (Signal Processing, Optimi- zation, Wavelet, PDE, Fre- quency Domain, Higher- Order Spectral Analysis, Image Processing, Neural Network, Statistics and Sys- tem Identification)		x		
GCG, Blast	x			
Quest	x			
S-PLUS	x			
AMPAC		x		
Gaussian	x			
CHARMm	x			
IDL	x			
XPLOR	x			x
LAP, Fastlink, Vitesse	x			
C, C++	x	x	x	x
Perl	x	x	x	x
Fortran 77, 90, High Per- formance	x	x	x	x
Pascal	x			
Javac			x	
IMSL		x		
LSF (Batch)	x	x	x	