Introducing NQE

IN–2153 3.0

The UNICOS operating system is derived from UNIX® System V. The UNICOS operating system is also based in part on the Fourth Berkeley Software Distribution (BSD) under license from The Regents of the University of California.

New Features

The NQE 3.0 release has the following new features:

- The new `nqe` command invokes a new NQE graphical user interface (GUI) that lets you do the following:
  - Use the `Submit` window to open and edit a job script; to save changes made to a job script; to submit a request to NQE; to view, segment, delete, or reset your NQE GUI log; and to set or unset your password. A launching capability enables you to submit a request periodically at specific, even repeating, intervals. Also use the `Submit` window to set (configure) and to save your job-related options (job profiles).
  - Use the `Status` window to verify the status of your requests and FTA file transfers. Also use the `Status` window to delete a request, to send a specified signal to a request, to get a detailed status of a request, and to set or unset your password. Context-sensitive help is provided with the status windows; as you glide your mouse cursor over a menu or field name in a window, a brief description of the menu or field appears at the bottom of the display.
  - Use the `Load` window to view a continually updated display of the system load for machines in the execution server complex. You also can view data about a specific configured host, and you can view the same data that is provided on the main window but have the data grouped by host rather than by type of data.
  - Use the `Config` window to set (configure) the specific user preferences and to view how NQE-related variables are currently set.
- The `cqstat`(1) command was removed and replaced with the new NQE GUI `Status` window capability.
- Support was added for the NQE database and NQE scheduler. The following command changes also were made to support this feature:
  - The `-d` option was added to the `cqdel`, `cqstatl`, and `cqsub` user commands to specify the request's destination (NQE database or NQS). The default destination is NQS.
  - The `-u` option of the `cqdel`, `cqstatl`, and `cqsub` user commands was expanded to support specifying the database user name.

– The nqedbmgr(8) command was added for administration of the NQE database.

- Support was added for clusterwide rerun of requests submitted to the NQE database (instead of to NQS). A cluster also is known as an execution server complex or batch complex.
- The new NQE scheduler is written in Tcl. This new feature allows a site to create its own scheduler as opposed to using NLB destination selection and NQS.
- The nqeinfo file is now supported on NQE clients. NQS_SERVER now may be set in the nqeinfo file, which eliminates the need for client users to set the NQS_SERVER environment variable if they want to use the nqeinfo file default setting.
- Support was added for using the Distributed Computing Environment (DCE) processing capability when submitting a job to NQE.
- FTA has been enhanced to support byte-level recovery. You also can use the new NQE GUI Status window to verify the status of your FTA file transfers.
- The new ilb utility executes a command on a machine chosen by the NLB. To use the utility, enter the ilb command followed by the command you want to execute. After the login process is complete on the machine chosen by the NLB, the command is executed and I/O is connected to your terminal or pipeline. The ilb man page documents the ilb utility.
- Support was added to change job resource limits (qalter -l option).
- The PPermfile_limit option was added to the qmgr MODify Request subcommand so that you can modify the NQS per-process permanent file space limit for a request that was already submitted. The qmgr man page was revised to reflect this new option.
- Documentation was revised to support the new features described on this New Features page.

Record of Revision

The date of printing or software version number is indicated in the footer.

Version	*Description*
1.0	December 1993. Original printing. This manual introduces the Network Queuing Environment (NQE). This manual supplements complete usage and feature descriptions contained in the online guides available with NQE.
1.1	June 1994. Incorporates information for NQE release 1.1.
2.0	May 1995. Incorporates information for NQE release 2.0.
3.0	March 1996. Incorporates information for the NQE 3.0 release.

Preface

This manual describes the Cray Network Queuing Environment (NQE). NQE is a Cray Research software product that lets users submit, monitor, and control work across machines in a network.

NQE consists of the following components that provide a seamless environment for client users to do their work:

- The NQE graphical user interface (GUI) allows users to submit and control batch requests to a central storage database and to obtain status on their batch requests and file transfers.
- The Network Load Balancer (NLB) uses the system load information received from NQE execution servers to offer NQS an ordered list of servers to run a request; NQS uses the list to distribute the request.
- The NQE database provides an alternate mechanism for distributing work. Requests are submitted and stored centrally. The NQE scheduler examines each request and determines when and where the request is run.
- The File Transfer Agent (FTA) provides asynchronous and synchronous file transfer. You can queue your transfers so that they are retried if a network link fails.

This introduction describes basic user tasks, provides instructions on how to get help online, and provides troubleshooting information.

Related publications

The following documents contain additional information that may be helpful (this introductory guide is also provided online):

- *NQE User's Guide*, publication SG–2148 (also provided online)
- *NQE Administration*, publication SG–2150 (also provided online)

- *NQE Installation*, publication SG–5236
- *CrayDoc Reference Card*, publication SQ–6101 (also provided online)

For further information, see subsection 3.1, page 43.

Ordering publications

The *User Publications Catalog*, publication CP–0099, describes the availability and content of all Cray Research hardware and software documents that are available to customers. Cray Research customers who subscribe to the Cray Inform (CRInform) program can access this information on the CRInform system.

To order a document, either call the Distribution Center in Mendota Heights, Minnesota, at +1–612–683–5907 or send a facsimile of your request to fax number +1–612–452–0141. Cray Research employees may send electronic mail to `orderdsk` (UNIX system users).

Customers who subscribe to the CRInform program can order software release packages electronically by using the `Order Cray Software` option.

Customers outside of the United States and Canada should contact their local service organization for ordering and documentation information.

Licensing and ordering NQE product

To order additional copies of this product, or to increase the number of licensed users, contact the Cray Research, Inc. in one of the following ways:

1–800–289–2729

+1–612–683–3897

- Send mail to the following address:

  CrayDirect
  Cray Research, Inc.
  655A Lone Oak Drive
  Eagan, Minnesota 55121
  USA

- Send electronic mail from any system connected to Internet to the following address:

  `craysoft@cray.com`

- Send a facsimile (fax) to the following telephone number:

  +1–612–683–3559

- Contact Cray Research through the World Wide Web (WWW) at the following URL:

  `http://www.cray.com/craysoft`

Conventions

The following conventions are used throughout this manual:

Convention	Meaning
`command`	This fixed-space font denotes literal items such as commands, files, routines, path names, signals, messages, and programming language structures.
`manpage`(*x*)	Man page section identifiers appear in parentheses after man page names. The following list describes the identifiers: 1 User commands 1B User commands ported from BSD 2 System calls 3 Library routines, macros, and opdefs 4 Devices (special files) 4P Protocols 5 File formats 7 Miscellaneous topics 7D DWB-related information 8 Administrator commands Some internal routines (for example, the `ddcnt`() routine) do not have man pages associated with them.
variable	Italic typeface denotes variable entries and words or concepts being defined.
`user input`	This bold fixed-space font denotes literal items that the user enters in interactive sessions. Output is shown in nonbold, fixed-space font.

Convention	Meaning
`ABBreviation` or `abbreviation`	Capital letters of a `qmgr` command or underlining indicates the shortest possible abbreviation for a command; underlining also highlights a summary phrase or sentence.
[ ]	Brackets enclose optional portions of a command line.
...	Ellipses indicate that a preceding command-line element can be repeated.

The following machine naming conventions may be used throughout this manual:

Term	Definition
Cray PVP systems	All configurations of Cray parallel vector processing (PVP) systems, including the following: CRAY C90 series (CRAY C916, CRAY C92A, CRAY C94, CRAY C94A, and CRAY C98 systems) CRAY C90D series (CRAY C92AD, CRAY C94D, CRAY C98D systems) CRAY EL series (CRAY Y-MP EL, CRAY EL92, CRAY EL94, and CRAY EL98 systems) CRAY J90 series (CRAY J916 and CRAY J932 systems) CRAY T90 series (CRAY T94, CRAY T916, and CRAY T932 systems) CRAY Y-MP E series (CRAY Y-MP 2E, CRAY Y-MP 4E, CRAY Y-MP 8E, and CRAY Y-MP 8I systems) CRAY Y-MP M90 series (CRAY Y-MP M92, CRAY Y-MP M94, and CRAY Y-MP M98 systems)
Cray MPP systems	All configurations of Cray massively parallel processing systems, including the CRAY T3D series (CRAY T3D MC, CRAY T3D MCA, and CRAY T3D SC systems) and CRAY T3E series

Term	Definition
All Cray Research systems	All configurations of Cray PVP and Cray MPP systems that support this release
SPARC systems	All SPARC platforms that run the Solaris operating system version 2.3 or later

The default shell in the UNICOS operating system, referred to in Cray Research documentation as the standard shell, is a version of the Korn shell that conforms to the following standards:

- Institute of Electrical and Electronics Engineers (IEEE) Portable Operating System Interface (POSIX) Standard 1003.2–1992
- X/Open Portability Guide, Issue 4 (XPG4)

The UNICOS operating system also supports the optional use of the C shell.

Cray UNICOS Version 9.0 is an X/Open Base 95 branded product.

The POSIX standard uses *utilities* to refer to executable programs that Cray Research documentation usually refers to as *commands*. Both terms may appear in this document.

Readers comments

If you have comments about the technical accuracy, content, or organization of this manual, please tell us. You can contact us in any of the following ways:

- Send us electronic mail from a UNIX system, using the following UUCP address:

  `uunet!cray!publications`

- Send us electronic mail from any system connected to the Internet, using the following Internet address:

  `publications@timbuk.cray.com`

- Customers in the United States and Canada can call the Cray Research Customer Support Center at either of the following telephone numbers. These telephone numbers are also listed on the `cs_support`(7) man page and on the Service and Support card included in the product package.

  1–800–272–9349 (toll free from the United States and Canada)

  +1–404–631–2244

- Customers in other countries can call their local Cray Research support office. These telephone numbers are listed on the `cs_support`(7) man page and on the Service and Support card contained in the product package.

- Contact your Cray Research representative and ask that a Software Problem Report (SPR) be filed. Use `PUBLICATIONS` for the group name, `PUBS` for the command, and `NO-LICENSE` for the release name.

- Call our Software Publications Group in Eagan, Minnesota, through the Customer Service Call Center, using either of the following numbers:

  1–800–950–2729 (toll free from the United States and Canada)

  +1–612–683–5600

- Send a facsimile of your comments to the attention of “Software Publications” in Eagan, Minnesota, at the following fax number:

  1–612–683–5599

- Use the postage-paid Reader’s Comment Form at the back of the printed manual.

We value your comments and will respond to them promptly.

Contents

	Page

Figures

Overview [1]

The Cray Network Queuing Environment (NQE) lets you submit batch requests to execute on a server in the NQE complex that is running the Network Queuing System (NQS). *Batch requests* are shell scripts that are executed independently from an interactive terminal session.

You submit requests from NQE clients and they are executed at NQE servers. You also can log on to servers and submit requests. You can use the NQE graphical user interface (GUI) to submit, delete, and signal requests, monitor the status of your requests, and compare the machine load of servers in the NQE network.

This section describes the following:

- NQE components and NQE network components
- Tasks you can perform using NQE
- Where to look for additional information

NQE components and NQE network components

1.1

NQE includes the following components:

- The NQE client provides the client user interface to NQE. It provides the NQE GUI, which is accessed through the `nqe` command. It also provides a command-line interface, which is documented in the *NQE User's Guide*, publication SG–2148.

  For information about using the NQE GUI, see section 2, page 7, of this introductory guide.

- The Network Queuing System (NQS) initiates requests on *NQS servers*. An NQS server is the host on which NQS runs. Your default NQS server is designated by your system administrator and is specified in the NQE configuration file; you may submit your request to a specific NQS server by using the NQE GUI `Config` window or by setting the `NQS_SERVER` environment variable.

- The Network Load Balancer (NLB) gathers system load information from NQE execution servers. This information may then be used to load balance batch requests across NQS servers in the NQE network. The NLB offers NQS a list of servers, in order of preference, to run a request; NQS uses the list to route the request.

- The NQE database provides a central repository for batch requests in the NQE network. The NQE scheduler uses the NQE database and an alternate mechanism for distributing work. The NQE scheduler examines each request and determines when and where the request will run. The NQE database lightweight server (LWS) verifies validation, submits the copy of a request that is in the NQE database to NQS, and obtains exit status of completed requests from NQS.

- The File Transfer Agent (FTA) provides asynchronous and synchronous file transfer. You can queue your transfers so that they are retried if a network link fails.

The NQE network contains the following components:

- NQE *clients* contain software so users may submit, monitor, and control requests by using either the NQE graphical user interface (GUI) or the command-line interface. From clients, users also may monitor request status, delete or signal requests, monitor machine load, and receive request output using the FTA.

- NQE *execution servers* also contain NQE client software. In addition, they also contain the NQS, FTA, NLB collectors, and the NQE database LWS. These additional components have the following functions:
  - NQS schedules and initiates batch requests.
  - FTA transfers files.
  - NLB collectors collect and communicate system load information to the NLB server so that the NLB can provide NQS with a list of servers, in order of preference, to run a request and provide status information upon request.
  - The LWS obtains request information from the NQE database, verifies validation, submits the copy of a request to NQS, and obtains exit status of completed requests from NQS.

- An NQE *master server* contains all NQE components. To distribute work among all NQE execution servers, a master server runs an NLB server; it also runs the NQE database and the NQE scheduler. A master server also contains all of the software that clients and execution servers contain, and they can function as execution servers in your execution server complex. Only one NQE master server is required in an execution server complex, although there may be more.

Figure 1 shows the location of the NQE components within the NQE network:

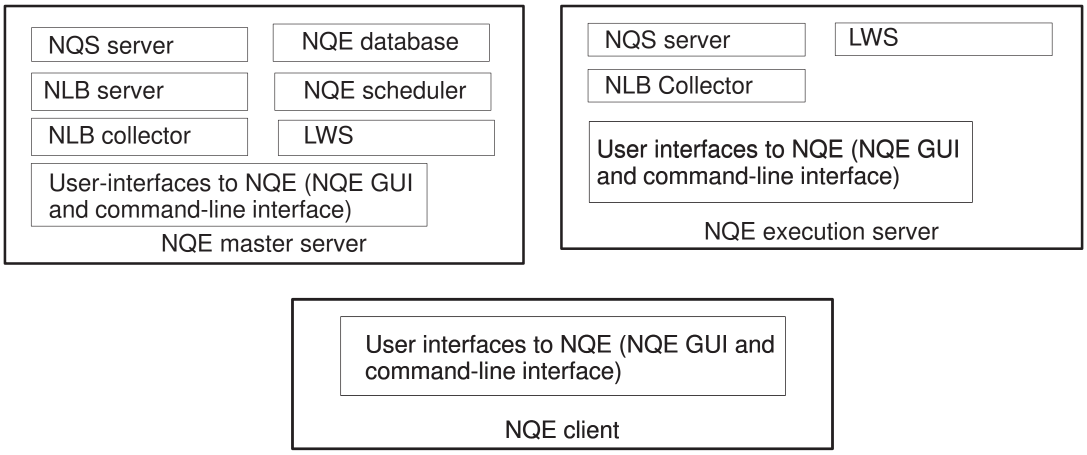

Figure 1. NQE component location

Tasks you can perform using NQE

1.2

NQE lets you do the following tasks:

- Submit batch requests to a server in the NQE complex that is running the Network Queuing System (NQS).

- Route requests automatically without knowing the names of queues or hosts.

- Monitor the status of requests across the NQE complex. The requests' status is refreshed at configurable intervals.
- Signal a request in the server complex.
- Delete a request in the server complex.
- Execute requests in interdependent sequences by using the `cevent` command to establish the sequence.
- Communicate with other UNIX systems running public domain versions of NQS.
- Monitor the machine load across the NQE server complex.
- Transfer files both interactively and from within batch requests without specifying a password.
- Queue file transfers so that you can go on to other work.
- Automatically retry file transfers that encounter network failures.
- Customize your environment.
- Read and search through online documentation.

Where to look for additional information

1.3

To learn more about the functions described in this introductory guide and about other user functions, see the *NQE User's Guide*, publication SG–2148 (for information about how to access the *NQE User's Guide* online, see subsection 3.1, page 43). Some of the functions described in the *NQE User's Guide* include the following:

- Designating either the NQE database and NQE scheduler or your NQS server to schedule and distribute your request for execution
- Using an alternate user name to submit your request for execution
- Customizing your request
- Customizing your environment

- Using the command-line interface
- Using job dependency, which lets you specify an interdependency among events in shell scripts or your requests (for example, you can ensure that one request runs after another so that the second request can use the output from the first request)
- Using the Distributed Computing Environment (DCE) and the DCE Distributed File System (DCE/DFS) environment within NQE

Getting Started [2]

This section describes the following information about NQE:

- Setting environment variables (subsection 2.1, page 7)
- Setting up authorization for the following:
  - NQE database authorization (subsection 2.2.1, page 10)
  - NQS validation to execute batch requests (subsection 2.2.2, page 11)
  - `.rhosts` validation files example (subsection 2.2.3, page 11)
- Using the NQE graphical user interface (GUI) for the following tasks:
  - Creating and submitting a batch request (subsection 2.3.1, page 16)
  - Monitoring a request (subsection 2.3.2, page 22)
  - Deleting a request (subsection 2.3.3, page 30)
  - Sending a signal to a request (subsection 2.3.4, page 31)
- Monitoring NQS queues (subsection 2.4, page 32)
- Monitoring machine load (subsection 2.5, page 35)
- Transferring files (subsection 2.6, page 39)

Setting environment variables

2.1

To use NQE, you must set the following environment variables:

- `DISPLAY` must be set to *local_workstation_name* `:0` for the NQE graphical user interface (GUI) to work.

  Note: If your site has access control in place for using X Windows applications, contact your system administrator to determine if you need additional settings.

- `PATH` must include the path name of the NQE commands. The default path name is as follows: `/usr/craysoft/nqe/bin`, except on Solaris systems where it is `/opt/craysoft/nqe/bin`. Optionally, it can contain the path name for the CrayDoc online documentation. The default path name is `/usr/craysoft/craydoc/v2.3fcs/bin`, except on Solaris systems where it is `/opt/craysoft/craydoc/v2.3fcs/bin`. System administrators also must include `/usr/craysoft/nqe/etc` or, for Solaris systems, `/opt/craysoft/nqe/etc` in their `PATH` environment variable to use certain NQE administrator commands.

- `MANPATH` must include the path name of the NQE man pages. The default name is `/usr/craysoft/nqe/man`, except on Solaris systems where it is `/opt/craysoft/nqe/man`.

- `EBTRC` can be set to the path name of a systemwide `.ebtrc` file, which is necessary to run the CrayDoc online documentation browser. (`EBTRC` is not necessary to run NQE.)

To verify whether your site's path names are the NQE system default, use one of the following commands:

```
cd /usr/craysoft/nqe/bin
```

For Solaris systems, use the following command:

```
cd /opt/craysoft/nqe/bin
```

If this command is not successful, ask your system administrator where the NQE software is located and add those directories to your `PATH` environment variables.

The commands that you use to set the environment variables depend on the shell that you use.

For a list of other NQE environment variables you may want to set to customize your environment, see section 9 of the *NQE User's Guide*, publication SG–2148.

The following example uses `sh` syntax to set and display NQE environment variables:

```
# PATH=$PATH:/usr/craysoft/nqe/bin:/usr/craysoft/nqe/etc:/usr/craysoft
/craydoc/v2.3fcs/bin; export PATH
# MANPATH=:/usr/craysoft/nqe/man; export MANPATH
# DISPLAY=snow32:0; export DISPLAY
# EBTRC=/usr/craysoft/craydoc/v2.3fcs/EBTRC; export EBTRC
# env
LOGNAME=you
MAIL=/var/mail/you
USER=you
SHELL=/bin/sh
PWD=/home/snow32/you
MANPATH=:/usr/craysoft/nqe/man
PATH=/usr/bin::/usr/craysoft/nqe/bin:/usr/craysoft/nqe/etc:/usr/crayso
ft/craydoc/v2.3fcs/bin
DISPLAY=snow32:0
EBTRC=/usr/craysoft/craydoc/v2.3fcs/EBTRC
```

The following example uses csh syntax to set and display the NQE environment variables:

```
% setenv PATH /usr/craysoft/nqe/bin:/usr/craysoft/nqe/etc:
/usr/craysoft/craydoc/v2.3fcs/bin:$PATH
% setenv MANPATH /usr/craysoft/nqe/man:$MANPATH
% setenv DISPLAY snow32:0
% setenv EBTRC /usr/craysoft/craydoc/craydoc/v2.3fcs/EBTRC
% env
HOME=/home/snow32/you
SHELL=/bin/csh
TERM=xterm
USER=you
LOGNAME=you
PWD=/home/snow32/you
MANPATH=/usr/man:/opt/local/man:/usr/craysoft/nqe/man
PATH=/usr/bin:/usr/craysoft/nqe/bin:/usr/craysoft/nqe/etc:/usr
/craysoft/craydoc/v2.3fcs/bin
DISPLAY=snow32:0
EBTRC=/usr/craysoft/craydoc/craydoc/v2.3fcs/EBTRC
```

Setting up authorization

2.2

To submit requests to the NQE database and to execute requests on the NQS server, you must have the proper authorization, as described in the following subsections.

NQE database authorization

2.2.1

To submit or control a request or to get a status of a request in the NQE database, you must have a database user (dbuser) account with the proper authorization (user privileges). This database user account name may be the same as or different from your login on the client host; this introductory guide assumes you are using the same login (user name) across the NQE network. Your NQE administrator controls who has access to the database and from which client host.

NQS validation to execute batch requests

2.2.2

By default, NQS uses file validation to authorize users. NQS also may be configured to use password validation or both file and password validation. Your NQE administrator specifies the validation method used in your NQS configuration.

If your site uses validation files, you must have an `.rhosts` or `.nqshosts` file in your home directory on each NQS server in the complex that might process your request. NQS uses these files to authorize your user name before it sends your request to a batch queue.

If your site uses aliases for host names, you also must include those names in the `.rhosts` or `.nqshosts` file entry. For example, you may have to include `ice.site.com` rather than simply `ice`.

Note: If your site uses password validation, you do not need validation files on the servers. If your site uses password validation, you must supply a password each time you submit, monitor, delete, or send a signal to a request.

`.rhosts` validation files example

2.2.3

An example follows of how user `jane` would set up her `.rhosts` files to use the NQE database in a multiple-server NQE network.

Note: For examples of setting up `.rhosts` files to submit requests directly to NQS or to use alternate user names, see the *NQE User's Guide* , publication SG–2148.

In this example, user `jane` is using the NQE GUI or NQE client commands on workstation `snow` and her NQE master server is `wind`. The NQE database resides on `wind`. User `jane` wants to submit requests to the NQE database on `wind`, where the NQE scheduler will select a target system to execute her request. To submit requests to the NQE database on `wind`, `jane` must have authentication on the NQE database (see subsection 2.2.1, page 10).

In this example, three other NQE servers exist in the NQE network that have the names `gust`, `storm`, and `rain`. This means that `jane` can potentially run requests on `rain`, `gust`, `storm`, and `wind`. To use all four servers, `jane` must have an `.rhosts` (or `.nqshosts`) file in her home directory on each of the four servers.

Since `jane` uses the same login ID on all of the NQE servers, the `.rhosts` file on `snow` would contain the following entries:

- On the NQE client workstation `snow`:
  `rain jane` (allows output to be returned to `snow`)
  `gust jane` (allows output to be returned to `snow`)
  `storm jane` (allows output to be returned to `snow`)
  `wind jane` (allows output to be returned to `snow`)

  Note: This is required only if `rcp` is the output mechanism; also, these entries need to be in the `.rhosts` file only; `rcp` does not use the `.nqshosts` file.

`jane`'s `.rhosts` file on each NQE server must contain the following entry (note that it is the same entry on each NQE server):

- On the NQE master server `wind`:
  `snow jane` (permits incoming requests from `jane` on `snow`)

- On the NQE execution server `gust`:
  `snow jane` (permits incoming requests from `jane` on `snow`)

- On the NQE execution server `rain`:
  `snow jane` (permits incoming requests from `jane` on `snow`)

- On the NQE execution server `storm`:
  `snow jane` (permits incoming requests from `jane` on `snow`)

Once she has her `.rhosts` files in place, `jane` may submit requests to `wind`.

Job output will be returned from the NQS server by default to the NQE client, `snow`, by either FTA or `rcp`. FTA will be attempted first.

FTA may require a password for the destination host (NQE client `snow`) and user (user `jane` at `snow`), which may be supplied by creating a `.netrc` file in the login directory of user `jane` on `wind` (`~jane/.netrc` on `wind`) (if it is executed there). (For further information, see the `netrc` man page.)

FTA may also be configured so that `jane` may not have to supply a password. This configuration in FTA is called network peer-to-peer authorization (NPPA). You may wish to check with your system administrator to learn if this option is available.

`rcp` is used only to return output if FTA fails immediately to return the output. `rcp` requires a `.rhosts` file at the destination host in the recipient's home directory (`~jane/.rhosts` on `snow`).

NQS will always try to use your `.nqshosts` file and will use the `.rhosts` file only if the `.nqshosts` file does not exist. You should use `.rhosts` files unless your system administrator tells you that you should not use them.

If you have validation files on all of the servers in the network on which your requests may run and you still receive authorization failure messages, you should read the subsections about authorization in the *NQE User's Guide*, publication SG–2148. To access the manual online, use the `cdoc` command (for more information on accessing documentation online, see subsection 3.1.3, page 45).

2.3 Using the NQE graphical user interface

The NQE graphical user interface (GUI) is similar to a Motif interface. To access the NQE GUI, execute the `nqe` command. Figure 2 shows the initial (top-level) NQE GUI button bar window that will appear:

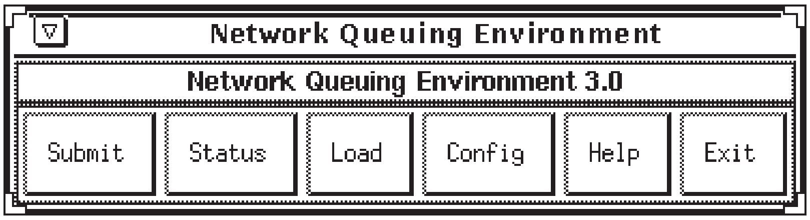

Figure 2. Initial, top-level NQE GUI button bar window

To access a window, position the mouse pointer over the desired window button and press ("click on") the left mouse button once. (The mouse button settings described in this guide are the default settings.)

You may use the NQE GUI for the following tasks:

- Use the `Submit` window to do the following:
  - Open and edit a job script
  - Save changes made to a job script
  - Submit a request to NQE
  - Launch a request on a periodic basis
  - From within the `Submit` window, reset your configuration preferences for the request you are submitting
  - View, segment, delete, or reset your NQE GUI log
  - Set or unset your password
  - Configure and save your job-related options (job profile)
- Use the `Status` window to do the following:
  - View updated status of your requests (the window is refreshed periodically)
  - View updated status of your FTA file transfers (the window is refreshed periodically)
  - Delete a request
  - Send a specified signal to a request
  - View the detailed status of a request
  - Set or unset your password

Context-sensitive help is displayed as you glide your mouse pointer over a menu or field name in the `Status` window; a brief description of the menu or field appears at the bottom of the display.

- Use the `Load` window to do the following:
  - Display continually updated system load information for machines in the execution server complex

- Display data about a specific host
- Display the same data that is provided on the main `Load` window, but have it grouped by host rather than by type of data

• Use the `Config` window to do the following:

- Set your preferences for the following: a specific NQS server; default job profile; temporary directory; job script; job output; job profile; and NQE GUI log directories.
- View your currently set preferences

• To access online help, use the left mouse button and click once on the `Help` button. For detailed information about using the NQE GUI and descriptions about menu options, select the `Help` button or see the *NQE User's Guide*, publication SG–2148. To read this manual online, use the `cdoc`(1) command (for more information on viewing documentation online, see subsection 3.1.3, page 45). For a summary of the NQE GUI displays and functions, see the `nqe`(1) man page.

• To exit the NQE GUI, use the left mouse button and click once on the `Exit` button.

When the mouse pointer is within a display area of a specific NQE GUI window, you may use the `ALT` key and the underscored letter from the menu bar to pop up submenus and to select more submenu options. An alternate way to do this is to use the `F10` key to activate the menu bar and then use the cursor movement keys to select submenus and options.

This introductory guide uses your default configuration (for information about changing your configuration or about using the `Config` window, see the *NQE User's Guide*, publication SG–2148).

Note: An NQE command-line interface also exists for most user functions. For additional information about using the NQE command-line interface, see the *NQE User's Guide*, publication SG–2148. Each command also is documented on a man page; man pages are provided only in online form.

Section 2

Creating and submitting a batch request
2.3.1

A batch request is a series of commands, usually contained in a script file that you create using a text editor (such as `vi`). A batch request can be one command, such as the following example:

```
ls    # list users
```

Usually , however, it contains several commands and can include directions to NQS, file transfer requests, and even your programs.

For example, the following batch request compiles and runs a C program:

```
date                         # Print date/time
cc loop.c    -o loop.out     # Compile loop.c
date                         # Print date/time
loop.out                     # Execute loop.out
date                         # Print date/time
echo job complete
```

You may submit a request to either the NQE database or to NQS (for information about submitting a request to run under the Distributed Computing Environment (DCE), see the *NQE User's Guide*, publication SG–2148).

To submit a request to NQE, access the NQE GUI by entering the `nqe` command. The initial NQE GUI button bar window will appear (as shown in Figure 2, page 13). Using the left mouse button, click once on the `Submit` button.

The `Submit` window shown in Figure 3 will appear.

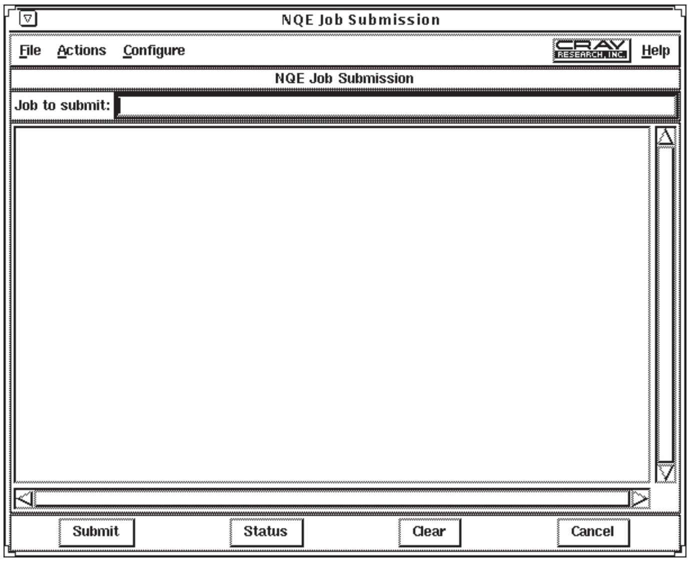

Figure 3. NQE GUI `Submit` window

The following subsections describe the components of the NQE GUI `Submit` window, and subsection 2.3.1.2, page 21, describes general steps to follow to submit an existing request.

Submit window description 2.3.1.1

The `Submit` window is composed of the menu bar, the file name (Job to submit) line, the job edit area, and the actions button bar; each of these four segments is described as follows.

> **Note:** For detailed information about the `Submit` window options or for information about how to configure mouse button settings differently, see the *NQE User's Guide*, publication SG–2148.

- *Menu bar.* Menu buttons are located within the menu bar, which is located at the top of the window. When you select a menu, it opens a submenu that contains lists of options. To open a menu window, place the pointer on the menu name and click once on the left mouse button.

– *File menu.* The `File` menu lets you open, edit, launch, save, and submit a job script (or request); view, segment, delete, or reset the NQE GUI log; and exit the `Submit` window. Figure 4 shows the `File` menu selections.

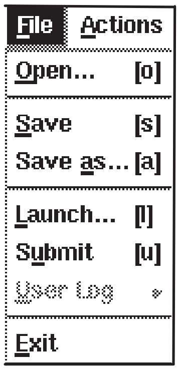

Figure 4. File pull-down menu of the `Submit` window

To open an existing job script, select `Open` from the `File` pull-down menu, which lets you select a file to open.

If you make changes to your job script and want to save the changes, select `Save` or `Save as` from the `File` pull-down menu.

To submit a job periodically at specific intervals or repeating intervals, select `Launch` from the `File` pull-down menu.

To submit the specified job script immediately (rather than submitting it by using the `Launch` selection, which lets you submit a job at a delayed or repeated interval), select `Submit` from the `File` pull-down menu. This selection provides the same action as using the `Submit` button located at the bottom of the window.

To view, segment, delete, or reset the NQE GUI log, select `User Log`.

To exit the `Submit` window, select `Exit`.

– *Actions menu*. The `Actions` menu lets you enter your password if your request will run on a host that uses password validation. This password is for the user name under which the request will execute. Figure 5 shows the `Actions` menu selections.

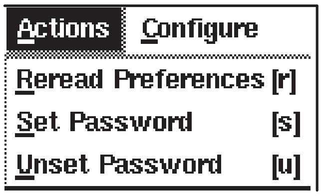

Figure 5. `Actions` pull-down menu of the `Submit` window

To reread your NQE GUI environment preferences before submitting a request, select `Reread Preferences`. This selection lets you refresh your preferences set through the `Config` window without closing and reopening your `Submit` window.

Figure 6 shows the `NQE Password Entry` window, which appears if you select either `Set Password` or `Unset Password`.

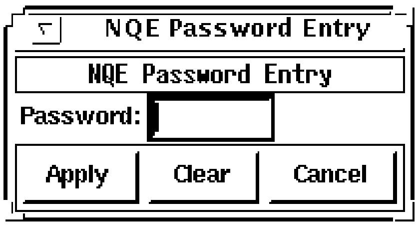

Figure 6. `NQE Password Entry` window

To set your password, enter your password, apply it, and cancel (close) the `NQE Password Entry` window.

To clear your password, select `Unset Password`, select `Clear` on the `NQE Password Entry` window, and then apply the action. If you entered the wrong password or if you receive a message that your password failed, you should use this selection to clear your password and then select `Set Password` and enter your correct password.

– *Configure menu*. The `Configure` menu of the `Submit` window lets you specify options for your request and then save them to be reused as needed. You do not have to include directives in your batch request file. Options you specify on the `Configure` menu take precedence over any options in the script file. Figure 7 shows the `Configure` pull-down menu of the `Submit` window.

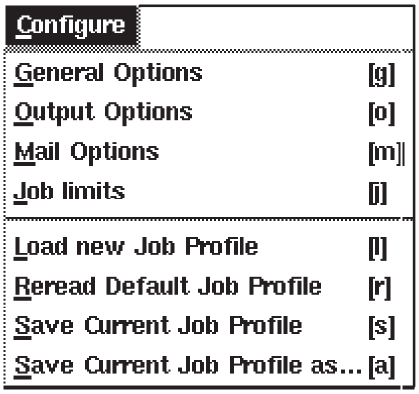

Figure 7. `Configure` pull-down menu of the `Submit` window

For example, you may set the destination of your request by selecting `Submit to NQE` or `Submit to NQS` on the `General Options` pull-down menu. To save loaded configuration options and limits so that you do not have to rekey them each time you submit the request, select `Save Current Job Profile` or `Save Current Job Profile as`.

– *Cray Research Logo*. This button opens a window that displays current NQE version and copyright information.

- *Help menu.* The `Help` menu lets you choose additional help information from either the `nqe`(1) man page or the *NQE User's Guide*, publication SG–2148.

- *Job to submit line*. The `Job to submit` line lets you specify the path name of a job to be edited or submitted to the NQS server. When you enter a valid job path name and press the `RETURN` key, the job text appears in the job edit area below the `Job to submit` line.

- *Job edit area*. You may make changes to the job displayed in the job edit area, and then save the changes by selecting `Save` or `Save as` from the `File` pull-down menu.

- *Actions bar.* The actions bar is located at the bottom of the `Submit` window (see Figure 3). When activated, these buttons perform the following actions (to activate a button, place the pointer on the button and click once on the left mouse button):

Button	Description
`Submit`	Submits the specified job script; the job script will use the loaded configuration options and limits.
`Status`	Provides the same summary status display as provided by the `Status` button on the top-level NQE GUI button bar.
`Clear`	Clears the `Job to submit` entry and the job edit area of the display .
`Cancel`	Cancels the `Submit` window.

General steps to follow to submit an existing request 2.3.1.2

To submit an existing request file, do the following:

- Either enter the path name of the request file on the `Job to submit` line and press the `RETURN` key, or select `Open` from the `File` menu (the `Open` option uses a standard Motif file selection interface) and select the request file you want to submit. The request file text appears in the job edit area below the `Job to submit` line.

Section 2

You may modify the content of your request file in the job edit area of the `Submit` window. To save changes for future use, select `Save` or `Save As` on the `File` menu.

- You may submit a request directly to NQS or to the NQE database. To set the destination for your request, select either the `Submit to NQE` or `Submit to NQS` on the `General Options` menu, and apply the change. If you do not select this option, the value of the `NQE_DEST_TYPE` environment variable is used, which you may set to be either `nqs` or `nqedb`; otherwise, the value of `NQE_DEST_TYPE` that is set in the `/etc/nqeinfo` file on your `NQS_SERVER` is used.

- If your request will run on a host that uses password validation, use the `Actions` menu to enter your password.

- To submit the request file, click once on the `Submit` button that is located at the bottom of the `Submit` window.

  If your request is submitted successfully, you will receive a message similar to one of the following messages:

  – For requests submitted to NQS, you will receive the following message:

    `Request` *number .host* `submitted to queue:`*queue* .

  – For requests submitted to the NQE database, you will receive the following message:

    `Task id t`*number* `inserted into database nqedb.`

  For additional information about messages received after submitting a request, see the *NQE User's Guide*, publication SG–2148.

To cancel the `Submit` window, click once on the `Cancel` button that is located at the bottom of the `Submit` window.

Monitoring a request
2.3.2

You may use the NQE GUI `Status` window to display information about your requests. The NQE GUI `Status` window provides a summary of request status that is refreshed periodically. By default, you see all requests in the execution server complex. However, it could be that your NQE administrator has disabled this function and that you will see only the requests you submit.

Using the NQE GUI `Status` window lets you do the following:

- Monitor status of all your requests. You do not have to know the location of your request before you request status on it. Request status is updated (refreshed) at configurable intervals.
- Tailor the display. You may specify how you want your display to look and what information is displayed.

To open NQE GUI `Status` window, access the NQE GUI by keying in the `nqe` command and, using the left mouse button, click once on the `Status` button of the initial NQE GUI button bar.

Figure 8 shows the `Status` window.

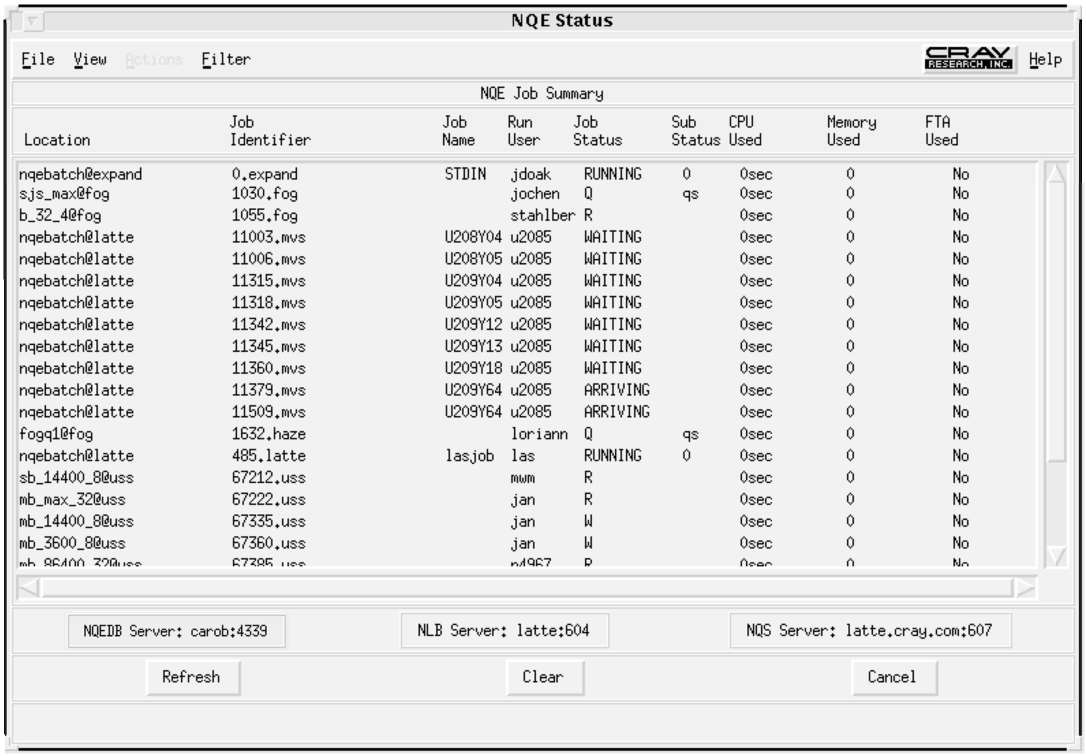

Figure 8. `Status` window

The `Status` window consists of the menu bar, the job summary area, the server identification area, the context-sensitive help area and the actions bar. Each of these segments is described as follows.

> **Note:** For detailed information about the `Status` window options, see the *NQE User's Guide,* publication SG–2148.

- *Menu bar.* Menu buttons are located within the menu bar, which is located at the top of the window. When you select a menu, it opens a submenu that contains lists of options. To open a menu, place the pointer on the menu name and click once on the left mouse button.
  - *File menu.* The `File` menu provides a selection for exiting the window.

– *View menu.* The `View` menu lets you display a summary of jobs (`Job Summary`) or FTA transfers (`FTA Summary`). The default `Status` window is the `Job Summary` view.

Figure 9 shows the `View` menu selections.

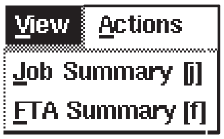

Figure 9. `View` pull-down menu of the `Status` window

– *Actions menu.* The `Actions` menu lets you delete a job, send a signal to a job, and view a detailed status of a job or its job log. This menu also lets you monitor status on hosts that use password validation.

Figure 10 shows the `Actions` menu selections.

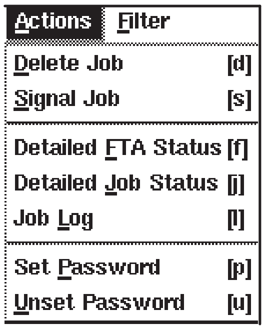

Figure 10. `Actions` pull-down menu of the `Status` window

To use the `Actions` menu, first position the pointer on the desired job line in the job summary area and select it by clicking the left mouse button once, then select the `Actions` menu selection.

To delete the currently selected job from the job summary area, select `Delete`. (For additional information about deleting a request, see subsection 2.3.3, page 30.)

To send a specified signal to the currently selected job from the job summary area, select `Signal Job`. (For additional information about signaling a request, see subsection 2.3.4, page 31.)

To display a new window that contains all details of the selected job, select `Detailed Status`. This action is the same as double clicking the left mouse button on a job in the job summary area.

To display information such as the date and time each job process runs, the status of your job as it is being executed, and the name of the queue in which the job is running, select `Job Log`.

To enter your password, select `Set Password`. The `NQE Password Entry` window appears. To set your password, enter your password, apply it, and cancel (close) the `NQE Password Entry` window.

To clear your password, select `Unset Password`, select `Clear` on the `NQE Password Entry` window, and then apply the action. If you entered the wrong password or if you receive a message that your password failed, you should use this selection to clear your password and then select `Set Password` to enter your correct password.

– *Filter menu*. The `Filter` menu lets you reduce the number of requests shown in the window. You may select specific hosts, user names, the originating user, the originating host, request IDs or task IDs, or locations.

Figure 11 shows the `Filter` menu selections.

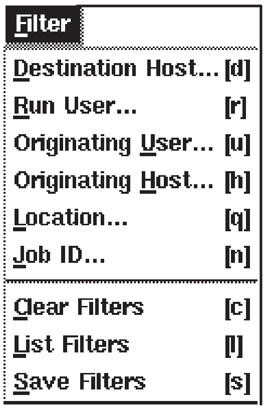

Figure 11. `Filter` pull-down menu of the `Status` window

After you select a `Filter` menu, to select an item, click once on the box next to it and then click once on the `Apply` button. The `Set All` button selects all available options in the display. The `Clear All` button eliminates all selected items. If you click on the `Clear All` button, you may then select other items you want to appear. After you have selected new items, you must apply your new selection.

– *Cray Research Logo*. This button opens a window that displays current NQE version and copyright information.

– *Help menu.* The `Help` menu provides general and administrative information related to each task and configuration option available to you through this window.

– *Job Summary area.* The job summary area shows a summary of information for the view option you specified by using the `View` menu (see Figure 9, page 25). The job summary view lists jobs according to filters that you have set by using the `Filter` menu selections (see Figure 11). The `FTA Summary` view lists FTA transfers that are in progress.

Section 2

The default view is `Job Summary`. The following data about requests is displayed by default:

Column name	Description
`Location`	The request's location, which may be either a queue or the NQE database
`Job Identifier`	The job identifier; possible identifiers are as follows: • NQS request ID (for example, `5703.fog`), as displayed when you submitted the request to NQS. • NQE database ID (known as the task ID) for example, `t1`), as displayed when you submitted the request to the NQE database. • NQE database ID with the NQS request ID. If you submitted a request to the NQE database, a copy or the request is executed while the original request remains in the NQE database. The request ID is the request identifier of the copy of the request executing in NQS and is displayed in parentheses after the `t`*id* (for example, `t4(61178.rain)`).
`Job Name`	Name of the request
`Run User`	User name with which the request was submitted.
`Job Status`	Status of the request.
`SubStatus`	Substatus of the request.
`CPU Used`	CPU usage (in seconds) for the request. On some platforms, a display of the amount of CPU that the request consumes is not available, and a 0 appears in this column.

Column name	Description
`Memory Used`	Memory usage (in words) for the request. On some platforms, a display of the amount of memory that the request consumes is not available, and a 0 appears in this column.
`FTA Used`	FTA usage for the request; usage setting may be `Yes` or `No`.

To select a request for use with the `Actions` menu selections, first position the pointer on the desired job line in the job summary area and select it by clicking the left mouse button once, then select the `Actions` menu selection.

To display a detailed status of a request, first position the pointer on the desired job line in the job summary area and then double-click the left mouse button on the desired summary line.

- The server identification area displays either the host name and port number for the NLB server or the host name and port number for the NQE database that is providing the information displayed by the `Job Summary` or `FTA Summary` selections of the `Status` window.
- The context-sensitive `Help` area is located at the bottom of the `Status` window. This area shows one-line informational messages about the area on the `Status` window that is directly under your mouse pointer.
- *Actions bar.* The actions bar is located at the bottom of the `Status` window. When activated, these buttons perform the following actions (to activate a button, place the pointer on the button and click once on the left mouse button):

Button	Description
`Refresh`	Refreshes the status window with new status information.
`Clear`	Clears the highlighted request line(s).
`Cancel`	Cancels the `Status` window.

Deleting a request
2.3.3

You may delete a request that you have submitted to NQE. The request may be executing or may be waiting to execute on an NQE server.

Note: If your site uses password validation, you must either set the `NQS_PASSWORD_NEEDED` environment variable or, in the NQE GUI `Submit` window, select `Set Password` on the `Actions` menu to ensure that you are prompted for a password; otherwise, your request will not execute. The password you supply is for the user name under which the request will execute.

After you submitted the request to be executed, you received a response similar to one of the following; the response includes the unique ID that is assigned to the request:

- If you submitted your request to NQS, you received a response similar to the following:

```
Request 46.latte submitted to queue: nqenlb
```

- If you submitted your request to the NQE database, you received a response similar to the following:

```
Task id t4 inserted into database nqedb
```

To delete a request that you have submitted but no longer want to execute, use the `Status` window of the NQE GUI.

Note: A request in the NQE database is known as a task and is assigned a task ID (for example, `t4`). When a copy of the request is executing under NQS, it also is assigned a request ID (for example, `288.latte`), which is displayed in parentheses after the task ID in the NQE GUI `Status` window `Job Identifier` field (for example, `t4(288.latte)`).

Select the request in the job summary area (its unique ID is included in the `Job Identifier` column). Then select `Delete Job` on the `Actions` menu (see Figure 10, page 25). This action deletes the request you selected. You will receive a response such as the following:

```
Request 46.latte has been deleted.
```

Using the NQE GUI to delete a request works whether or not the request is executing.

> **Note:** The original file is not deleted, just the request to execute the file is deleted.

Sending a signal to a request
2.3.4

You may signal a request that you have submitted to NQE. The request may be executing or may be waiting to execute on an NQE server. You may send one of several signals to a request; one of the most common is the `SIGKILL` signal, which aborts a running process.

> **Note:** If your site uses password validation, you must either set the `NQS_PASSWORD_NEEDED` environment variable or, in the NQE GUI `Submit` window, select `Set Password` on the `Actions` menu to ensure that you are prompted for a password; otherwise, your request will not execute. The password you supply is for the user name under which the request will execute.

After you submitted the request to be executed, you received a response similar to one of the following; the response includes the unique ID that is assigned to the request:

- If you submitted your request to NQS, you received a response similar to the following:

```
Request 46.latte submitted to queue: nqenlb
```

- If you submitted your request to the NQE database, you received a response similar to the following:

```
Task id t4 inserted into database nqedb
```

To send a signal to one or more of your requests, use the `Status` window of the NQE GUI.

Note: A request in the NQE database is known as a task and is assigned a task ID (for example, `t4`). When a copy of the request is executing under NQS, it also is assigned a request ID (for example, `288.latte`), which is displayed in parentheses after the task ID in the NQE GUI `Status` window `Job Identifier` field (for example, `t4(288.latte)`).

To signal the request, select the request in the job summary area (its unique ID is included in the `Job Identifier` column). Then select `Signal Job` on the `Actions` menu (see Figure 10, page 25).

A dialog box appears and asks you to select a signal to send. Select the signal you want to send by clicking once on one of the signal buttons. The dialog box disappears and the job is signaled.

A second dialog box will appear with a reply from the signal request. After you review the reply, click once on the `OK` button and the process is complete.

Standard output, standard error, and job log files are still produced for an executing request that is deleted by a signal. These files record the execution of the request up to the moment that the signal is received.

Monitoring NQS queues

2.4

Two types of NQS queues exist. Requests are executed from *batch queues.* Requests in *pipe queues* are routed to other queues, either batch or pipe. To display summary information on NQS queues, use the `cqstatl` command; this information is not available through the NQE GUI.

Note: The concept of queues does not exist in the NQE database. Only when a copy of an NQE database request is submitted to a given NQS does it enter a queue. If you have the `NQE_DEST_TYPE` environment variable set to be `nqedb` or if you use the `cqstatl` command `-d nqedb` option, you will not see queue status.

NQS uses both batch and pipe queues. To display a summary of a specific type of queue, use one of the following `cqstatl` command options:

Option	Description
`-b`	A summary of all batch queues
`-p`	A summary of all pipe queues

The following example uses the `-p` option to obtain information about pipe queues:

```
% cqstatl -p
----------------------------
NQS PIPE QUEUE SUMMARY
----------------------------
QUEUE NAME            LIM TOT ENA STS QUE ROU  WAI HLD ARR DEP  DESTINATIONS
--------------------- --- --- --- --- --- ---  --- --- --- ---  -------------
nqenlb                  1   0 yes  on   0   0    0   0   0   0
pendulum                1   0 yes  on   0   0    0   0   0   0  nqebatch
--------------------- --- --- --- --- --- ---  --- --- --- ---  -------------
pendulum               10   0           0   0    0   0   0   0
--------------------- --- --- --- --- --- ---  --- --- --- ---  -------------
```

The display shows that the NQS server `pendulum` has two NQS pipe queues. Both of the queues are enabled (shown in the `ENA` column). This means you may send requests to either queue. The queue `nqenlb` has no destinations because it is a load-balancing queue that sends your request to the most appropriate queue.

To display summary information about NQS batch queues, use the -b option of the cqstatl command, as shown in the following example:

```
% cqstatl -b
----------------------------
NQS BATCH QUEUE SUMMARY
-----------------------------
QUEUE NAME              LIM TOT ENA STS QUE RUN  WAI HLD ARR EXI
----------------------- --- --- --- --- --- ---  --- --- --- ---
nqebatch                  5   0 yes  on   0   2    0   0   0   0
batch123a                10  11 yes  on  11   0    0   0   0   0
----------------------- --- --- --- --- --- ---  --- --- --- ---
pendulum                 10  11          11   2    0   0   0   0
----------------------- --- --- --- --- --- ---  --- --- --- ---
```

This display contains information about the total NQS workload on the NQS server system. (The NQS server is defined by the NQS_SERVER variable, which you also may set as an environment variable.) This display helps you see why your requests may take a long time to begin executing. The TOT column tells you the total number of requests in each queue. The remaining columns tell you whether the queue is accepting requests (yes or no), starting requests (on or off), and how many requests in the queue are queued, running, waiting, held (by an NQS operator), arriving, and exiting.

To display summary information about all NQS queues at your NQS server, use the cqstatl command, as follows:

```
cqstatl
```

Often this is a large display and will scroll off your screen. To control scrolling, redirect the output from cqstatl to the more(1) command, as follows:

```
cqstatl | more
```

To display full details about NQS queues, use the `cqstatl -f` option, as follows:

```
cqstatl -f
```

You may use the `cqstatl` options `-b` and `-p` to restrict the display to a particular type of queue. To restrict the detailed display to a particular queue, use the following format:

`cqstatl -f` *queue*

To display a list of the limits that the NQE administrator defined for all NQS batch queues, use the `cqstatl -l` option (lowercase L), as follows:

```
cqstatl -l
```

For additional information about using `cqstatl` options, see the *NQE User's Guide*, publication SG–2148, or see the `cqstatl`(1) man page.

Monitoring machine load

2.5

The Network Load Balancer (NLB) displays information about the hosts in the execution server complex. This information is used for load balancing, display of machine load data, and display of request status in the network.

To access machine-load information, click once on the `Load` button of the NQE GUI interface. (To view this window from the command-line interface, execute the `cload` command.) The window provides the following:

- Continually updated status that allows easy comparison of the workload of servers
- Visual indication that a host is not providing new data
- Pop-up windows that provide information on a server

You may view data from the main window, by individual host, and through a miniature summary display. Figure 12 shows the main `Load` window.

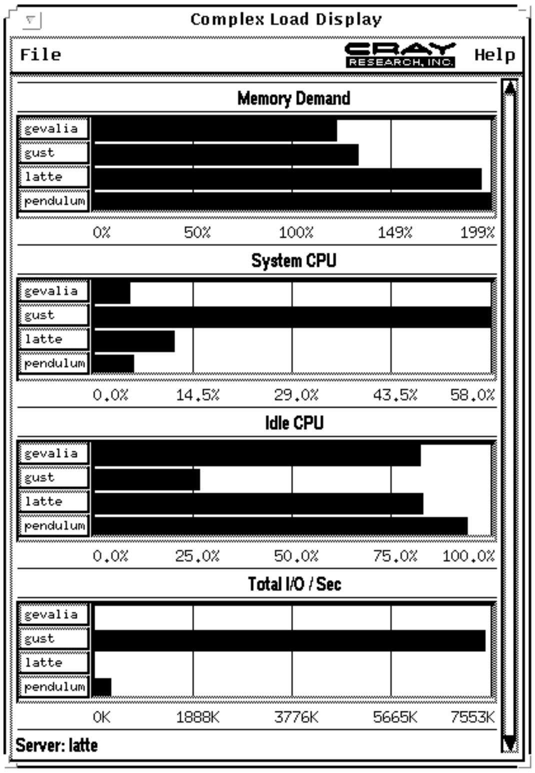

Figure 12. `Load` window

The `Load` window is composed of the menu bar, the complex load display, and the server name display area. Each of these segments is described as follows:

- *Menu bar.* The menu bar is located at the top of the `Load` window and displays buttons that open `Load` menus. To open a menu window, place the pointer on the menu name and click once on the left mouse button.
  - *File menu.* The `File` menu provides selections for hosts that you want to see in the window, for viewing the formula used to display the data, and to close the window.

- *Help menu.* The `Help` menu provides selections to view the `cload` man page or to display the name of the NQE administrator .

- *Complex load display.* The complex load display provides continually updated machine load data for participating machines in the complex. Each of the charts in the display has a title, a scale that automatically adjusts to the current data, and one button per host.

You may view data about a specific host as shown in Figure 13. You also may view the same data that is provided on the main `Load` window (memory demand, percentage of system CPU in use, idle CPU, and total I/O per second) grouped by host rather than by type of data (as shown in Figure 14, page 38).

To display information about a specific host, place the pointer over the button that has the name of the host on the main `Load` window and click once on the left mouse button. You will receive a display like the one shown in Figure 13. When the NLB receives new data, this display is updated.

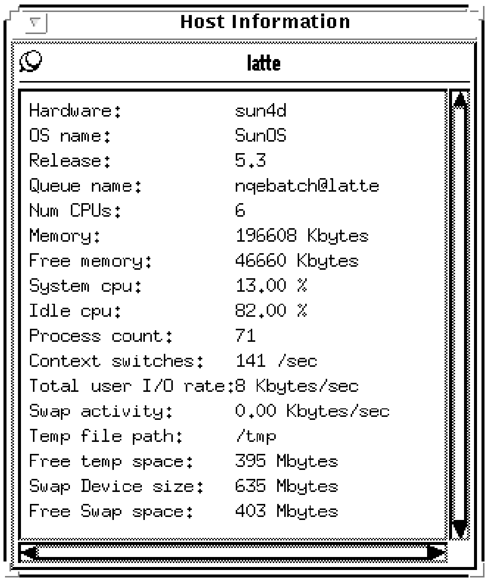

Figure 13. Load display for specific host

To cancel this window, place the pointer over the pin in the upper left corner of the window and click once on the left mouse button.

To display a window of data that is grouped by host rather than by type of data, position the pointer over the host name on the main `Load` window and click once on the middle mouse button. As you select additional hosts, the window will expand and include a display for each host you select. Figure 14 shows the load summary displayed for four hosts that were selected individually:

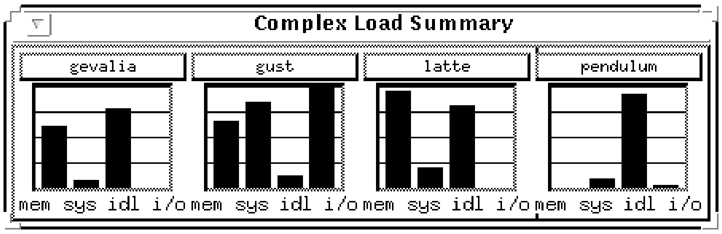

Figure 14. NLB `Load Summary` displayed by host

To view the same host information as is shown in Figure 13, page 37, position the pointer over the host name in the Complex Load Summary window and click once on the left mouse button. To cancel the host information window, position the pointer over the pin in the upper left corner and click once on the left mouse button.

To execute an `rlogin` command for a host, position the pointer over the host name and click once on the middle mouse button.

To cancel a host summary display, as shown in Figure 14, position the pointer over the host name and click once on the right mouse button. Repeat this action for each host summary you want to cancel.

You may customize the displays by editing the `.Xdefaults` file; see the *NQE User's Guide*, publication SG–2148, or the `nqe`(1) man page.

Transferring files

2.6

To transfer files from within a batch request or interactively, use the `ftua` and `rft` commands. The `ftua` interface to the File Transfer Agent (FTA) is similar to the `ftp` utility.

You might choose to use FTA for the following reasons:

- You may execute file transfers immediately or queue them for later execution. If the transfer is queued, it is executed after you leave the utility , allowing you to do other tasks.
- If the transfer fails for some transient reason (such as a network link failing), FTA automatically requeues the transfer. This is very useful in batch requests because the request can still be run even if a transfer cannot occur when it is first tried.
- FTA provides network peer-to-peer authorization, which lets you transfer files without specifying passwords either in batch request files or in `.netrc` files or by transmitting passwords over the network.

To view the status of an FTA transfer, use the NQE GUI `Status` button and select the `View` menu `FTA Summary` selection. The `FTA Summary` displays a one-line summary of all transfers. To get a detailed status display of a specific transfer, place the pointer over the transfer name and double-click on the left mouse button. To cancel the detailed status display, click once on the display's `Cancel` button by using the left mouse button.

The following example line in a batch request from user `jane` uses `rft` to copy a file named `remfile` from a remote host named `rain` to a file named `locfile` on the local host (the file names are relative to the login directory):

```
rft –user jane –host rain –domain nqe \
–nopassword –function get remfile locfile
```

In the following example, `jane` uses the `ftua` command to transfer a remote file from host `gust` to her local directory. The commands she uses perform the following steps:

1. Request a connection to the host `gust` by using domain `nqe`
2. Turn off interactive prompting

Section 2

3. Queue all transfers
4. Request the transfer (which is queued)
5. Display details about the request
6. Exit `ftua`

The commands in this example are described as follows:

Command	Description
`ftua gust nqe`	Opens a connection that will use the domain `nqe`
`prompt`	Turns off interactive prompting
`queue`	Queues all transfers
`get data.feb`	Requests a file transfer
`qdir`	Displays details about the queued transfer
`quit`	Exits `ftua`

```
% ftua gust nqe
Name (gust:eve):
Password:
ftua> prompt
ftua> queue
ftua> get data.feb
ftua> qdir
--QID-- ---Queued On --- --User-- --Group--
aa001t1 Wed Oct  6 10:39 eve      pubs
        State: Active
        Status: UA connected
        Priority: 0
        Last Attempt: Never
        Domain: inet
        Host: gust
        Username:eve
        Copy (get)
                RemoteFile: data.feb
                LocalFile:  /net/home/ice/eve/data.feb
250 QUEUE command successful.
ftua> quit
221 Goodbye.
```

For more information on using FTA, see the *NQE User's Guide*, publication SG–2148.

Getting Help [3]

This section describes the online documentation available in NQE. It also explains some possible problems and solutions.

Online documentation

3.1

Online documentation is available through the following sources:

- NQE GUI help function
- Online man pages
- CrayDoc online documentation reader, from which you can access the NQE documentation

The forms of online documentation are described in the following subsections.

Section 3

NQE GUI help function

3.1.1

The online documentation available within NQE through the NQE GUI help function gives you immediate access to information about the task you are performing. The *NQE User's Guide*, publication SG–2148, may be accessed through the NQE GUI help function.

Online man pages

3.1.2

The `man` command provides online help on all NQE commands. To view a man page online, type `man` *commandname*. You must ensure that your `MANPATH` is set properly, as described in subsection 2.1, page 7.

Section 3

The following user-level online man pages are provided with your NQE software:

User -level man page	Description
`cdoc`(1)	Brings up the CrayDoc online documentation reader , from which you can access the NQE documentation
`cevent`(1)	Posts, reads, and deletes job-dependency event information
`cload`(1)	Provides an X Window System display of machine load information
`cqdel`(1)	Deletes or signals to a specified batch request
`cqstatl`(1)	Provides a line-mode display of requests and queues on a specified host
`cqsub`(1)	Submits a batch request to NQE
`ftua`(1)	Transfers a file interactively (this command is issued on an NQE server only)
`ilb`(1)	Executes a load-balanced interactive command
`nqe`(1)	Provides a graphical user interface (GUI) to NQE functionality
`qalter`(1)	Alters the attributes of one or more NQS requests (this command is issued on an NQE server only)
`qchkpnt`(1)	Checkpoints an NQS request on a UNICOS system (this command is issued on an NQE server only)
`qdel`(1)	Deletes or signals NQS requests (this command is issued on an NQE server only)

User -level man page	Description
qlimit(1)	Displays NQS batch limits for the local host (this command is issued on an NQE server only)
qmsg(1)	Writes messages to stderr, stdout, or the job log file of an NQS batch request (this command is issued on an NQE server only)
qping(1)	Determines whether the local NQS daemon is running and responding to requests (this command is issued on an NQE server only)
qstat(1)	Displays the status of NQS queues, requests, and queue complexes (this command is issued on an NQE server only)
qsub(1)	Submits a batch request to NQS (this command is issued on an NQE server only)
rft(1)	Transfers a file in a batch request (this command is issued on an NQE server only)

CrayDoc online documentation
3.1.3

The following manuals are available through the CrayDoc online documentation reader:

- *Introducing NQE*, publication IN–2153 (this publication)
- *NQE User's Guide*, publication SG–2148
- *NQE Administration*, publication SG–2150

The cdoc command brings up the CrayDoc online documentation reader, from which you can access the NQE documentation. CrayDoc is based on Electronic Book Technologies' *Dyna* Text product.

CrayDoc must be able to access the .ebtrc file for operation. You can include this file in your home directory. The file may be in a systemwide start-up file that you can point to by setting the EBTRC environment variable. To determine whether EBTRC is set, use the following command:

```
echo $EBTRC
```

If this command returns a path name, you do not have to copy the file to your home directory or set the variable. If the command does not return a path name, create a file in your home directory by using the following commands:

```
% cd $HOME
% cp /usr/craysoft/craydoc/v2.3.fcs/EBTRC .ebtrc
```

Note: On Solaris systems, the file is named /opt/craysoft/craydoc/v2.3.fcs/EBTRC.

You also must ensure that the CrayDoc path is in your search path as described in subsection 2.1, page 7.

To start the reader, use the following command:

```
cdoc
```

To look at the documentation that describes how to use the reader, select the Reader Guide for UNIX book from the available books panel and click OPEN. For basic browser operations, see the *CrayDoc Reference Card*, publication SQ–6101.

In CrayDoc documents, you will see references to subsection numbers, tables, and figures marked as hot text. *Hot text* is a link to another place in the document. It is indicated by red, underscored text on a color monitor or by bold, underscored text on a monochrome monitor. The cursor changes from an arrow to a point finger when it is above hot text.

To use the cross-references, click on any hot text and the cross-reference will be displayed either in the window you are using or in a new window.

To search for a word or phrase, select `Search -> Show Search Panel` from the menu bar at the top of the main window. Enter the word or phrase you want matched in the `Search` field and click the `Search` button.

To close the book, from the menu bar at the top of the main window, select `File -> Close Book`. To exit CrayDoc, select `File -> Quit` from the Library window menu bar. When you exit CrayDoc, it closes any books that are open.

NQE troubleshooting

3.2

This subsection lists problems you may have in using NQE and possible actions you can take to solve the problems.

Commands do not execute

3.2.1

Before you can use the NQE commands, you must add the NQE `/bin` directory (and the `/etc` directory to use administrator commands) to your search path. Before you can use the man pages (which tell you about the NQE commands and command options) you must add the NQE `/man` directories to your search path. For a description of how to set these variables, see subsection 2.1, page 7.

Section 3

`Connection failed` messages

3.2.2

If you are using a client command or the NQE GUI to talk directly to an NQS server and a network error condition exists that can be retried, you will receive messages such as the following:

```
Retrying connection to NQS_SERVER on host ice (147.111.21.90) . . .
QUESRV:  ERROR:  Failed to connect to NQS_SERVER at ice [port 607]
NETWORK:  ERROR:  NQS network daemon not responding
```

The server may have just come up, it may not be listening, or the network may be busy. The client will retry the connection to the server for 30 seconds. You may also verify the value of the `NQS_SERVER` environment variable is set to a host running NQS.

If you are using a client command or the NQE GUI to talk to the NQE database, and the database server is not up, you will see a message like the following:

```
Connect: Connection refused
NETWORK: ERROR: NQE Database connection failure:
        Can't connect to MSQL server on latte
```

3.2.3 Authorization failure messages

The default validation type in NQS is file validation. File validation requires a `.rhosts` or `.nqshosts` file in the login directory of your account on all of the NQE server systems where your request may run. This applies even if your target NQS server is your local machine.

For summary information about validation files, see subsection 2.2, page 10. For a complete description, see the *NQE User's Guide*, publication SG–2148. To access the manual online, use the `cdoc`(1) command as described in subsection 3.1.3, page 45.

If you encounter authorization failure messages, consider these alternatives:

- A `.rhosts` or `.nqshosts` file may not exist for the user name under which the request will be run. This may be the case if the `$HOME` environment variable points to a location that does not have these files.
- The proper *hostname-username* pair is not contained in a validation file. You must have an entry in a validation file for every NQE server that the Network Load Balancer (NLB) or the NQE database may select to run your request.
- A password is required and was not supplied, or the wrong password was supplied.
- Depending on your local network configuration, host names of servers might have to be fully qualified (that is, you may have to include `ice.site.com` rather than simply `ice`).

- You may have created both `.rhosts` and `.nqshosts` files and only the `.rhosts` file is correct. When the `.nqshosts` file exists, NQS ignores the `.rhosts` file.

NQE database authorization failures

3.2.4

A client attempting to connect to the database without the proper validation will result in an error like the following:

```
latte$ cqstatl
NETWORK: ERROR: NQE Database connection failure:
         Connection disallowed.
latte$
```

Check with your system administrator to determine why you are unable to connect.

Request seems lost

3.2.5

If you cannot locate your request, try the following options:

- Check to see whether NQS sent you a mail message. This message may help you to determine the cause of the problem.
- In the NQE GUI, display all of your requests in the NQE database or in the NQS batch queues in the complex by using the left mouse button to click once on the `Status` button.
- Perhaps the request has completed. Check for the standard output and standard error files that the request produced. Unless you specified a different location for these files, they should be in the directory from which you submitted your request.

  If you used an alternative user name, try the home directory of the user name you used to submit the request.

  Check your home directory on the NQS server at which your request executed.

For more information about locating output, see subsection 3.2.8, page 51.

NQE scheduler not scheduling

3.2.6

If the NQE scheduler is not scheduling the request, an NQS server may not be available to accept your request.

Request seems to be stuck

3.2.7

If you find that a request remains in a queue for a long time, check the `Status` column of the NQE GUI `Status` window, which lists a status code that will help you determine the problem. The following are possible reasons:

- A letter `W` or the word `WAITING` can indicate the following:
  - Your request is waiting for a license (see subsection 3.2.9, page 52).
  - A pipe queue's destination queue is disabled.
  - A pipe queue's destination is at a remote system that is not available. NQS automatically retries sending the request at periodic intervals.
- A letter `Q` that has a substatus code `qs` indicates that queue is stopped. Only the NQE administrator can start and enable NQS queues.
- A letter `Q` that has a substatus code beginning with `c`, `g`, or `q` (except `qs`) indicates that a limit has been reached for the queue complex, globally on the NQS server, or for a queue. You can wait for resources to become available or delete the request.
- Your request may be accepted into an NQS queue and later be deleted. Check your electronic mail for a message such as the following:

```
The request could not be routed to any of the possible pipe
queue destinations because of the following reason(s)   :

No account authorization at transaction peer;
```

Section 3

If you receive this message, one of the following is true:

– Password checking is in use and either you did not request a password prompt (by selecting `Set Password` on the `Actions` menu of the NQE GUI `Submit` window, by using `cqsub -P`, or by setting the `NQS_PASSWORD_NEEDED` environment variable), or you supplied a password that is not valid.

– Validation file checking is in use and no correct entry exists in the validation file of the user name you used to submit the request. See section 2, page 7.

Subsection 3.2.3, page 48, includes more information about troubleshooting validation files.

Another possible reason that your request is not executing is that when NQS reads your `.cshrc` file or your request script file, it may encounter commands it cannot invoke and cause your request not to run. You will see these problems reflected in your standard error file. Also, NQS sets an environment variable called `ENVIRONMENT` to value `BATCH` for each NQS-initiated job. This variable can be checked within a `.profile`, `.login`, or `.cshrc` script and be used to differentiate between interactive and batch sessions; this action can be used to avoid performing terminal setup operations for a batch job. A benefit of NQS initiating the batch job as a login shell is that `.profile`, `.login`, or `.cshrc` scripts are run and your environment is set up as expected.

Section 3

Cannot locate output
3.2.8

If you do not specify another location, by default, NQS will place your output files in the directory you were in when you submitted the request. If the files are not where you expected them to be, you should have mail from NQS telling you where it sent your output.

Check your home directory on the NQS server at which your request executed. If you used an alternative user name, try the home directory of the user name you used to submit the request.

Note: For NQE database requests, the NQE database will send your output files to your home directory on your NQE client.

NQS uses the following methods to return your output to you:

1. NQS tries to send it by using NQS protocol. This method works if you submitted the request from an NQS server.
2. NQS tries to use `fta`. This method works if you have a `.netrc` file containing a password on the execution server or if FTA has been configured between client and server to use NPPA (no passwords).
3. NQS tries to use `rcp`. This method works if you have a `.rhosts` entry on your submitting system, and your `.rhosts` file has an entry for the NQE server that executed your request. Depending on your local network configuration, host names of servers might have to be fully qualified (that is, you may have to include `ice.site.com` rather than simply `ice`).
4. NQS sends you mail stating that it could not deliver your output to its first destination. It places your output in your home directory on the NQS execution server.
5. NQS delivers your output to an administrative directory. To retrieve the files, you must contact your system administrator.

Section 3

No licenses are available

3.2.9

When you try to submit a request and no NQE licenses are available, your request is accepted, but it will be held in a wait state in NQS. You will see the following status messages:

```
Wlm
WAITING
License Unavailable
```

NQE tries to obtain a license every 3 minutes.

Index

Index

Index

O

P

Q

R

S

Index

Index

Reader's Comment Form

Introducing NQE IN–2153 3.0

Your reactions to this manual will help us provide you with better documentation. Please take a moment to complete the following items, and use the blank space for additional comments.

List the operating systems and programming languages you have used and the years of experience with each.

Your experience with Cray Research computer systems: _____0-1 year _____1-5 year _____5+years

How did you use this manual: _____in a class _____as a tutorial or introduction _____as a procedural guide _____as a reference _____for troubleshooting _____other

Please rate this manual on the following criteria:

	Excellent			Poor
Accuracy	4	3	2	1
Appropriateness (correct technical level)	4	3	2	1
Accessibility (ease of finding information)	4	3	2	1
Physical qualities (binding, printing, illustrations)	4	3	2	1
Terminology (correct, consistent, and clear)	4	3	2	1
Number of examples	4	3	2	1
Quality of examples	4	3	2	1
Index	4	3	2	1

Please use the space below for your comments about this manual. Please include general comments about the usefulness of this manual. If you have discovered inaccuracies or omissions, please specify the number of the page on which the problem occurred.

Name ______________________
Title ______________________
Company ______________________
Telephone ______________________
Today's date ______________________

Address ______________________
City ______________________
State/Country ______________________
Zip code ______________________
Electronic mail address ______________________

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST CLASS MAIL PERMIT NO 6184 ST. PAUL, MN

POSTAGE WILL BE PAID BY ADDRESSEE

CRAY RESEARCH, INC.

ATTN: Software Publications Group
655 LONE OAK DR BLDG F
EAGAN MN 55121-9957

Fold

Cut along this line

Reader's Comment Form

Introducing NQE IN–2153 3.0

Your reactions to this manual will help us provide you with better documentation. Please take a moment to complete the following items, and use the blank space for additional comments.

List the operating systems and programming languages you have used and the years of experience with each.

Your experience with Cray Research computer systems: _____0-1 year _____1-5 year _____5+years

How did you use this manual: _____in a class _____as a tutorial or introduction _____as a procedural guide _____as a reference _____for troubleshooting _____other

Please rate this manual on the following criteria:

	Excellent			Poor
Accuracy	4	3	2	1
Appropriateness (correct technical level)	4	3	2	1
Accessibility (ease of finding information)	4	3	2	1
Physical qualities (binding, printing, illustrations)	4	3	2	1
Terminology (correct, consistent, and clear)	4	3	2	1
Number of examples	4	3	2	1
Quality of examples	4	3	2	1
Index	4	3	2	1

Please use the space below for your comments about this manual. Please include general comments about the usefulness of this manual. If you have discovered inaccuracies or omissions, please specify the number of the page on which the problem occurred.

Name ______________________________
Title ______________________________
Company ______________________________
Telephone ______________________________
Today's date ______________________________

Address ______________________________
City ______________________________
State/Country ______________________________
Zip code ______________________________
Electronic mail address ______________________________

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST CLASS MAIL PERMIT NO 6184 ST. PAUL, MN

POSTAGE WILL BE PAID BY ADDRESSEE

CRAY RESEARCH, INC.

ATTN: Software Publications Group
655 LONE OAK DR BLDG F
EAGAN MN 55121-9957

Fold

Cut along this line