

# CONVEX CXbatch User's Guide

*Third Edition*




CONVEX

CONVEX COMPUTER CORPORATION



**CONVEX Computer Corporation**  
3000 Waterview Parkway  
P.O. Box 833851  
Richardson, TX 75083-3851  
United States of America  
(214)497-4000



---

# CONVEX CXbatch User's Guide



---

Order No. DSW-183

Third Edition  
January 1992

---

**CONVEX**

**CXbatch User's Guide**

Order No. DSW-183

©1992 CONVEX Computer Corporation  
All rights reserved.

This document is copyrighted. This document may not, in whole or part, be copied, duplicated, reproduced, translated, electronically stored, or reduced to machine readable form without prior written consent from CONVEX Computer Corporation.

Although the material contained herein has been carefully reviewed, CONVEX Computer Corporation does not warrant it to be free of errors or omissions. CONVEX reserves the right to make corrections, updates, revisions or changes to the information contained herein. CONVEX does not warrant the material described herein to be free of patent infringement.

UNLESS PROVIDED OTHERWISE IN WRITING WITH CONVEX COMPUTER CORPORATION (CONVEX), THE PROGRAM DESCRIBED HEREIN IS PROVIDED AS IS WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. SOME STATES DO NOT ALLOW THE EXCLUSION OF IMPLIED WARRANTIES. THE ABOVE EXCLUSION MAY NOT BE APPLICABLE TO ALL PURCHASERS BECAUSE WARRANTY RIGHTS CAN VARY FROM STATE TO STATE. IN NO EVENT WILL CONVEX BE LIABLE TO ANYONE FOR SPECIAL, COLLATERAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES, INCLUDING ANY LOST PROFITS OR LOST SAVINGS, ARISING OUT OF THE USE OR INABILITY TO USE THIS PROGRAM. CONVEX WILL NOT BE LIABLE EVEN IF IT HAS BEEN NOTIFIED OF THE POSSIBILITY OF SUCH DAMAGE BY THE PURCHASER OF ANY THIRD PARTY.

CONVEX and the CONVEX logo ("C") are registered trademarks of CONVEX Computer Corporation.

Printed in the United States of America

---

## Revision information for

# CONVEX CXbatch User's Guide

---

Edition	Document No.	Description
Third	710-002730-206	Released with CONVEX CXbatch, V2.1, January 1992. Added conceptual information and explanation of fields shown in output.
Second	710-002730-205	Released with CONVEX CXbatch, V2.0, September 1990, contains information about integration with CONVEX Share Scheduler and Checkpoint Restart, new qsub options, and enhanced accounting features.
First, Rev.1	710-002730-203	Released with CONVEX CXbatch V1.1, April 1990.
First	710-002730-200	Released with CONVEX CXbatch V1.0, February 1989.



---

# Contents

---

<b>Using this guide .....</b>	<b>xiii</b>
Purpose and audience .....	xiii
How to use this guide .....	xiii
Notational conventions .....	xiv
General conventions .....	xiv
Command syntax .....	xiv
Associated documents .....	xv
Ordering documentation .....	xv
Technical assistance .....	xvi

---

<b>1 CXbatch processing .....</b>	<b>1</b>
Managing queues and queue requests .....	2
General users .....	2
CXbatch operators .....	4
Managing queue requests .....	4
Managing queues .....	4
Managing CXbatch .....	5
CXbatch managers .....	5
Incompatibilities between NQS and CXbatch .....	6
Integration with CONVEX Share Scheduler .....	7
COVUEbatch considerations .....	7
Integration with Checkpoint Restart .....	8

---

<b>2 Displaying information .....</b>	<b>9</b>
Displaying status of queues .....	10
Displaying standard output .....	10
Displaying additional request information .....	14
Displaying date and time requests will run .....	16
Displaying additional queue information .....	17
Displaying queue status interactively .....	21
Displaying queue limits .....	26
Displaying status of related processes .....	28

---

---

**3 Submitting a request..... 31**

Three methods for submitting batch requests ..... 32

- Using interactive commands ..... 32
- Using a script file ..... 32
- Using a compiled program ..... 33

qsub command options ..... 34

- Controlling run requests ..... 34
  - Specifying a future run time .....35
  - Specifying a billing account ..... 36
  - Putting a request on hold at submission ..... 36
  - Controlling importation of the current directory ..... 36
  - Specifying use of login shell ..... 38
  - Setting priority on batch jobs ..... 40
  - Submitting to a specific queue ..... 40
  - Naming a request ..... 41
  - Exporting all environment variables ..... 42
  - Submitting a request silently ..... 43
- Redirecting output files and error messages ..... 44
  - Redirecting standard error output ..... 44
  - Redirecting error messages to the standard output file ..... 45
  - Redirecting error output on the executing machine .. 46
  - Redirecting standard output on executing machine . 46
  - Redirecting standard output ..... 47
  - Appending accounting information to stdout output file ..... 47
- Specifying resource limits ..... 48
- Sending notification of request status ..... 50
  - Sending mail when request begins execution ..... 50
  - Sending mail when request completes execution ..... 51
  - Sending mail to specified user ..... 52
- Signalling processes when a request completes executing ..... 53

Embedded options ..... 54

---

**4 Controlling queue requests..... 57**

Deleting queue requests ..... 58

- Using the qdel command ..... 58
- Using the delete request command ..... 59

Preventing queue requests from executing ..... 60

- Placing a queue request on hold ..... 60
- Removing the hold on a queue request ..... 60

Moving a request to another queue ..... 61

Changing the queue request priority ..... 62

---

<b>5 Checkpoint Restart features.....</b>	<b>63</b>
Controlling checkpointing when request is submitted .....	64
Overriding not-checkpointable default .....	64
Overriding checkpointable default .....	65
Periodically checkpointing a request .....	65
Preventing a checkpointed request from being restarted .....	65
Two methods of checkpointing after request is submitted ..	66
Using the qchkpnt command .....	66
Using the chkpnt request command .....	68
Suspending executing requests .....	69
Suspending an executing request .....	69
Resuming a suspended request .....	69
Restarting checkpointed requests .....	70

---

<b>A Transaction completion messages.....</b>	<b>71</b>
---	-----------

---

<b>B Request completion messages .....</b>	<b>99</b>
--	-----------

---

<b>C Man pages .....</b>	<b>113</b>
--------------------------	------------

---



---

# Figures

Figure 1	Standard <code>qstat</code> output .....	11
Figure 2	Additional queue request information .....	14
Figure 3	Future run time output .....	16
Figure 4	Additional queue information .....	17
Figure 5	Additional queue information .....	18
Figure 6	Interactive queue status output .....	22
Figure 7	Interactive queue request output .....	24
Figure 8	Queue limit output .....	26
Figure 9	Status of <code>CXbatch</code> -related processes .....	29
Figure 10	Contents of a shell script .....	30
Figure 11	Submitting a batch job interactively .....	32
Figure 12	Submitting batch job using a script file .....	32
Figure 13	Submitting a batch job using a compiled program .....	33
Figure 14	Submitting a compiled program using a script file .....	33
Figure 15	Incorrect way to submit compiled program .....	33
Figure 16	Default environment .....	42
Figure 17	All environment variables exportation .....	42
Figure 18	Mail received when job starts executing .....	50
Figure 19	Mail received when job completes executing .....	51
Figure 20	Mail sent to another user .....	52
Figure 21	Embedded options in a script file .....	55



---

# Tables

Table 1	Packet numbers recognized by CXbatch .....	6
Table 2	qsub options that control run requests .....	34
Table 3	qsub options that redirect output and error messages .....	44
Table 4	Per-process and per-request limits .....	48
Table 5	Invalid per-request limits for CONVEX machine .....	49
Table 6	qsub options that allow sending notification .....	50
Table 7	qsub options related to Checkpoint Restart .....	64
Table 8	Exit statuses .....	67



---

# Using this guide

---

## Purpose and audience

The *CONVEX CXbatch User's Guide* describes basic CXbatch concepts and how to use the features available to the general user.

---

## How to use this guide

If you have never used CXbatch, read Chapter 1 for basic CXbatch concepts before attempting to perform any of the tasks described in this book.

For information on viewing queue and queue request information, read Chapter 2.

For information on how to submit a job to the batch network, read Chapter 3.

For information on manipulating your own jobs submitted to a batch queue, read Chapter 4.

For information on how to checkpoint and restart batch requests, read Chapter 5.

For information on messages put out by CXbatch, read Appendix A and Appendix B.

---

## Notational conventions

This section discusses notational conventions used in this book.

---

### General conventions

The following conventions are used in this guide:

- *Italics*
  - Designate user-supplied variables in a command-line example
  - Indicate document titles
- Constant-width font designates input that must be typed exactly as it appears and output displayed on the terminal screen. This includes:
  - Command names and options
  - Directives, program statements, display examples, printout examples, and error messages returned.
- Horizontal ellipsis (...) shows repetition of the preceding item(s).
- Words and abbreviations that indicate keyboard keys you press are identified in a distinctive bold type. For example, **RETURN** refers to the carriage return key. Words separated by a hyphen indicate two keys that you must press simultaneously. For example, **CTRL-x** indicates that you must press and hold down the **CTRL** key and then press the **x** key.
- The word “enter” in a phrase such as “enter **ls**” means that you type the command and then press **RETURN**.
- References to the *ConvexOS Programmer’s Reference* appear in the form `adb(1)`, where the name of the man page is followed by its section number enclosed in parentheses.

---

### Command syntax

In order to use the commands in this document, you must understand the conventions used when describing command syntax. Consider this example:

```
Command input_file [input_file ...] {a|b} [output_file]
```

①                    ②                    ③                    ④                    ⑤

1. Constant-width font indicates that you must type the characters exactly as they appear (uppercase and lowercase are identical). The letters that appear in uppercase indicate

the minimum amount you must type to make the command unique. However, they do not have to be typed in uppercase.

2. *Italics* indicate a variable that must be supplied by the user. In this case, the user must supply the name of an *input\_file*.
3. Square brackets [ ] indicate optional data. Horizontal ellipsis (...) shows repetition of the preceding item(s). In this case, the user can optionally specify more than one *input\_file* on the command line.
4. Curly brackets { } indicate a choice. The choices available are shown inside the curly brackets and separated by the pipe (|) sign. In this case, the user can enter either a or b.
5. [*output\_file*] indicates the user can optionally specify an output file name with the command.

---

## Associated documents

Using this software may require information not specific to the tasks described in this document.

For more information on the ConvexOS operating system, you can order the following books from CONVEX Computer Corporation:

- *CONVEX COVUEbatch Guide* (DSW-151). This book describes how to use and maintain the COVUEbatch batch processing software.
- *ConvexOS Extensions User's Guide* (DSW-053). This book describes ConvexOS utilities from the user's perspective.

---

## Ordering documentation

To order the current edition of this or any other CONVEX document, send requests to:

CONVEX Computer Corporation  
Customer Service  
P.O. Box 833851  
Richardson TX 75083-3851 USA

Include the order number or the exact title, as listed on the front cover.

---

## **Technical assistance**

If you have questions that are not answered in this book, contact the CONVEX Technical Assistance Center (TAC).

- Within the continental U.S., call 1(800)952-0379.
- From Canada, call 1(800)345-2384
- Outside continental U.S., contact local CONVEX office.

CONVEX CXbatch permits users to submit jobs to a queue for batch execution. A queue is a list of batch requests that are ready and waiting to execute. A batch request is one or more commands submitted by a user or a user program to a batch queue. These commands are usually executed after a certain time or event has passed and do not require further interaction with the user. A typical batch request is a program containing commands that perform large-scale, detailed computations on a static data file.

Users can submit batch jobs for execution on either the local machine or a remote machine configured with CXbatch or another version of Network Queueing System (NQS). There is one CXbatch daemon per machine and each machine has its own set of queues.

When a user submits a request to CXbatch, CXbatch assigns it to a queue and assigns it a priority. The request waits in the queue until it is selected for execution. Requests are selected according to their priority. Once selected, CXbatch executes it and routes output to the specified recipient.

There are two types of CXbatch queues:

- **Batch**—Batch queues are used only to execute CXbatch batch requests. They hold requests for scheduled, perhaps delayed, processing by subsystems within CXbatch.
- **Pipe**—Pipe queues are routing queues that do not directly process requests but, instead, transmit requests to other queues on either the same or a remote machine. Each pipe queue has a set of destination queues that are possible recipient queues for requests submitted to that pipe queue. A destination queue can be either a batch queue or another pipe queue.

---

## Managing queues and queue requests

Commands available through the CXbatch `qmgr` utility allow a user to control submitted requests, track requests through the system, and with the right privileges, control CXbatch queues. The `qmgr` commands available to each user depend on their user type. CXbatch distinguishes three user types:

- General users
- CXbatch operators
- CXbatch managers

The following sections list and describe the commands available with each user type. The rest of this document describes the syntax and usage of each command available to the general user.

---

### General users

General CXbatch users can track and control their own batch requests using a small subset of `qmgr` commands:

<code>chkpnt request</code>	Checkpoint a request. The state of the batch request is saved into a set of checkpoint files stored in the checkpoint directory. The request continues to run.
<code>delete request</code>	Delete a request. The request can either be running or not running.
<code>exit</code>	Exit <code>qmgr</code> .
<code>help</code>	Get help on the commands available to that user.
<code>hold request</code>	Place a request on hold, preventing its execution. The request must be in the queued state in order to place it on hold.
<code>modify request</code>	Modify the priority of a request. Users can only decrease the priority of their requests. CXbatch operators can increase the priority of any user's request.
<code>move my_request</code>	Move a request from one queue to another. The request is not moved if any queue limits, access restrictions, or attributes prevent the request from being submitted to the queue.
<code>release request</code>	Release a request previously placed on hold, allowing the request to execute.

<code>resume request</code>	Resume execution of suspended request. A resumed request starts out in the queued state. Once it is about to enter the running state, it is restarted from its checkpointed state.
<code>show</code>	View the status of requests (all <code>show</code> commands). Refer to the <code>qmgr(8)</code> man page for a complete list of <code>show</code> commands.
<code>suspend request</code>	Temporarily suspend execution of a request. The request is checkpointed and execution is terminated. However, the request remains in the queue so it can be resumed later. If a request fails to checkpoint, it continues executing. Only checkpointable requests can be suspended.

---

## CXbatch operators

Users with CXbatch operator privileges have access to commands that allow them to manipulate batch requests for other users, control queues, and set run limits. The CXbatch manager assigns operator privileges to a user in order for them to act as a CXbatch operator.

### Managing queue requests

The qmgr commands listed under "General users" are available to users with CXbatch operator privileges to track and control any user's batch requests. In addition, the following commands are available to CXbatch operators:

move request	Move a request from one queue to another queue. The request is moved regardless of any queue limit violations, access restrictions, or attribute violations.
run request	Force a request to begin executing immediately. If running the request exceeds the current run limit of the queue, the queue's run limit is increased by one until the request completes.

### Managing queues

The following qmgr commands are available to CXbatch operators to manage queues:

abort queue	Abort all currently running requests in the queue.
disable queue	Prevent queue from accepting requests.
enable queue	Allow queue to accept requests.
move queue	Move all requests in one queue to another queue.
purge queue	Delete all queued requests from the queue.
start queue	Put queue into operation to allow it to execute jobs assigned to it.
stop queue	Prevent queue from executing any other requests; the currently executing request is allowed to complete.

## Managing CXbatch

The following qmgr commands are available to CXbatch operators to manage CXbatch:

<code>set copy_open_files</code>	Preserve the state of open files within a process when the request is checkpointed.
<code>set share policy</code>	Specify where CPU usage is charged for jobs running in the queue.
<code>set per-user run limit</code>	Establish the per-user run limit on the queue.
<code>set run limit</code>	Establish the run limit of an existing queue.
<code>shutdown</code>	Checkpoint requests that are checkpointable and shut down CXbatch on the local machine. You must execute <code>shutdown</code> to shut down CXbatch before executing <code>/etc/shutdown</code> to shut down the system. Otherwise, requests are not checkpointed and therefore, not recoverable.
<code>start Cxbatch</code>	Start CXbatch on the local machine.

---

## CXbatch managers

Users with CXbatch manager privileges have access to all qmgr commands. Refer to the `qmgr(8)` man page for a complete list. By default, root is a CXbatch manager and can assign this privilege to other users.

## Incompatibilities between NQS and CXbatch

CXbatch is based on NASA's Network Queueing System (NQS) but has many enhancements, including checkpoint restart, load balancing, fair share scheduling for batch jobs, and batch accounting. There are six major differences between the way you can use CONVEX CXbatch and other NQS systems:

- `move my_request` does not exist in other NQS systems and can only be used within CXbatch.
- CXbatch can mount remote files using NFS; other systems cannot.
- CXbatch's `maximum_request_priority` command, if used when submitting a request between CXbatch and another NQS system, causes the priority of previously-submitted requests to be moved to a lower priority. It does not delete previously-submitted requests.
- You cannot directly submit batch requests between CONVEX machines and other machines. You must use CXbatch to submit them.
- Several CXbatch `qsub` options are not applicable if the destination machine is a CONVEX machine. These options are listed in Chapter 3, "Submitting a request," of the *CXbatch User's Guide*.
- Several CXbatch packet numbers are not recognized by other NQS systems. Table 1 lists packet numbers recognized only by CXbatch.

**Table 1** Packet numbers recognized by CXbatch

Packet Number	Type	Use
206	nqs	Queue request from remote <code>qsub</code>
207	nqs	Get sequence number
208	nqs	Hold request
209	nqs	Release request
216	nqs	Checkpoint a request
218	nqs	Force request to run
219	nqs	Suspend request
220	nqs	Resume request
223	nqs	Force run request
205	net	Get remote queue and request information

---

## Integration with CONVEX Share Scheduler

CXbatch is integrated with the CONVEX Share Scheduler, an optional product. Share Scheduler is a per-user process scheduler that operates with the standard process scheduler. It provides equitable allocation of machine resources among users and groups of users according to their allocation of shares.

---

## COVUEbatch considerations

CONVEX COVUEbatch is an optional product that allows a user to submit batch jobs from VAX/VMS systems to remote CONVEX systems. A user submits a command procedure to COVUEbatch. COVUEbatch then uses COVUENet (an optional network management product) to transmit the job to CXbatch on the CONVEX system and to transmit the results of the batch job to the VMS user's home directory.

There are several items that you must consider if COVUEbatch is installed on your system.

- The `covue show queue` command does not display full CXbatch queue information, such as queue type (batch or pipe), queue limits, request limit, and attributes such as import, type of pipe queue, and accounting.
- If COVUEbatch is installed and running on only one machine in the CXbatch network and a user wants to submit a job to a remote queue, there must be a pipe queue on the local machine that has the remote queue as a destination. In this situation, the same performance degradation occurs as mentioned above. If COVUEbatch is installed and running on all machines that have CXbatch, however, remote queues can be accessed without use of pipe queues.
- The `covue show queue` command displays the batch queues on the local machine and pipe queue destinations. The `covue delete/entry` command deletes only those entries from CXbatch queues on the local machine.
- Because COVUEbatch uses lowercase queue names, references to COVUEbatch queue names must also be lowercase.

Please refer to the *CONVEX COVUEbatch Guide* for further information.

---

## Integration with Checkpoint Restart

Checkpoint Restart is a standard feature of ConvexOS. It permits the state of selected processes or process hierarchies to be saved to disk files and later to be restarted from the saved state.

Checkpoint Restart is useful for application programs that must run for long lengths of time and that—once halted—cannot be started over from the beginning without wasting significant time and resources. Such applications can be saved to files (checkpointed) either by operator intervention or by CXbatch periodically checkpointing them. If the application is then halted, either by a system crash, by a scheduled shutdown, or by an operator, it can be restarted as it was when it was last checkpointed. If desired, the process can be restarted under the control of a debugger.

Users and operators can checkpoint and restart processes that are running under CXbatch by using Checkpoint Restart features built into CXbatch. See Chapter 5, "Checkpoint Restart features," for details on these features.

This chapter describes how to display information about queues and requests in queues using the following CXbatch commands:

- `qstat`—Displays status of CXbatch queues and requests in queues.
- `qwatch`—Displays status of CXbatch queues and requests in queues interactively.
- `qlimit`—Displays queue limits.
- `qps`—Displays information about CXbatch-related processes.
- `qjlist`—Displays the contents of a shell script.

How to use each of these commands is described in the following sections.

---

## Displaying status of queues

You can display information about CXbatch queues and requests in the queues using the `qstat` command. The format for this command is

```
qstat [option...] [queueName[@hostname]...]
```

where

*option* controls the type and amount of information displayed. If no options are specified, `qstat` shows only those requests belonging to the user issuing the command. *option* can be one or more of the following:

- a Displays status for all requests in the queue.
- l Displays additional information about queues and queue requests. See the “Displaying additional request information” section in this chapter for a description of the output for this option.
- m Displays the date and time requests will run. See the “Displaying date and time requests will run” section in this chapter for a description of the output for this option.
- u *username* Displays only those requests belonging to the specified *username*.
- x Displays additional information about queues. See the “Displaying additional queue information” section in this chapter for a description of the output for this option.

*queueName* is the name of the queue for which you wish status information. If you do not specify a queue, information for all queues on the requested host is displayed.

*hostname* is the name of the machine that receives the request. If *hostname* is omitted, the local host is assumed.

---

## Displaying standard output

There are two sections to the standard output for the `qstat` command: information on the queues and information on each individual request in the queue. Figure 1 illustrates the standard output for the `qstat` command.

Figure 1 Standard qstat output

```
Queue information | % qstat v
                  | verylong@hostC; type=BATCH; [ENABLED, RUNNING]; pri=16
                  | aliases: v, verylong_queue, V, VERYLONG
                  | 0 exit; 1 run; 0 stage; 0 queued; 0 wait ;0 hold; 0 arrive;
Request information | REQUEST NAME   REQUEST ID   USER      PRI  STATE      PGRP
                  | 1:myjob       47.mach2    test      31  RUNNING    12103
```

The first line of the output describes information about the queue:

```
verylong@hostC; type=BATCH; [ENABLED, RUNNING]; pri=16
 ①                ②                ③                ④                ⑤
```

- ① Name of the queue and what host it is configured on. In this case, the attributes shown apply to the *s* queue on *hostC*.
- ② Type of queue. This can be:
  - BATCH Executes CXbatch requests.
  - PIPE Routes CXbatch requests to queues that can execute them.
- ③ Indicates whether or not the queue can accept requests. This can be:
  - ENABLED CXbatch is running on the local machine and the queue is accepting requests.
  - DISABLED CXbatch is running on the local machine but the queue is not accepting requests.
  - CLOSED CXbatch is not running on the local machine.
- ④ Indicates whether or not the queue can execute requests. This can be:
  - INACTIVE Requests in the queue are permitted to run; none are running.
  - RUNNING Requests in the queue are permitted to run; some are running.
  - STOPPING New requests sent to the queue are not permitted to run, but requests that are currently running are allowed to complete.
  - STOPPED Requests in the queue are not permitted to run; none are running.
  - SHUTDOWN CXbatch is not running on the local machine.
- ⑤ Inter-queue priority. This priority affects which queue is looked at first for the next job to run. Priority can be any number from 0 to 63; 0 is the lowest priority and 63 the highest.

The second line of the output tallies the number of requests in specific states:

0 exit; 0 run; 0 stage; 0 queued; 0 wait; 0 hold; 0 arrive;  
 ①            ②            ③            ④            ⑤            ⑥            ⑦

- ① Number of requests in this queue in a state of exiting.
- ② Number of requests in this queue in the state of running.
- ③ Number of requests in this queue in the staged state, which means the request has completed executing and is moving the stdout and stderr files to the appropriate destination directory.
- ④ Number of requests in this queue in a state of being queued.
- ⑤ Number of requests in this queue in a state of waiting.
- ⑥ Number of requests in this queue in a state of holding.
- ⑦ Number of requests in this queue in a state of arriving.

The rest of the output displays status information about each of the requests in the queue:

REQUEST NAME	REQUEST ID	USER	PRI	STATE	PGRP
1:myjob	47.mach2	test	31	RUNNING	12103
①	②	③	④	⑤	⑥

- ① Name assigned to the request.
- ② Unique identifier assigned to request when it is submitted to the queue.
- ③ User submitting the request.
- ④ Intra-queue priority assigned to request. This priority affects which job in a queue is executed next. This number can be from 0 to 63; 0 is the lowest priority and 63 the highest.
- ⑤ State of the request. This can be:

ARRIVING	Request is arriving at the queue.
CHECKPOINTED	A failed attempt was made to restart the checkpointed request. You can attempt to manually restart the request using the <code>qrestart</code> command. Refer to Chapter 5, "Checkpoint Restart features," for details on using this command.
DEPARTING	Request is departing from the queue but has not yet been received by the destination queue.
EXITING	Batch request has completed executing and will exit from the system after the required

output files are returned to their intended destinations.

HOLDING	A hold has been placed on the request preventing it from entering any other state.
QUEUED	Request is queued and eligible for running or routing. This is the most common state.
ROUTING	Request has reached the head of a pipe queue and is being routed to another queue.
RUNNING	Request has reached its final destination batch queue and is executing.
WAITING	Request is waiting for a specified amount of time to pass before attempting to execute. This could be because it was submitted with a future date and time specified for running, or because a pipe queue could not route the request and will attempt to route it later.
SUSPENDED	Request is suspended from running. It will remain suspended until manually resumed.

- ⑥ Process group of the request, if available to the local CXbatch daemon. This information is displayed only for processes that are running.

For more information on queue or request properties, refer to the `qstat(1)` man page.

---

## Displaying additional request information

You can use the `-l` option to display additional information on queue requests. Figure 2 illustrates the output for this option.

Figure 2 Additional queue request information

```
% qstat -l long
Queue for long jobs.
  long@mach2; type=BATCH; [ENABLED, RUNNING]; pri=32
aliases: l, long_queue, L, LONG
  0 exit; 1 run; 0 stage; 0 queued; 0 wait; 0 hold; 0 arrive;

Request 1: Name=STDIN Id=25.mach2
  Owner=test Priority=31 RUNNING Pgrp=4918
  Created at Mon Jan 23 16:30:03 CST 1989
  Mail = [NONE]
  Mail address = test@mach2
  Owner user name at originating machine = test
  Import directory: Yes
  Per-proc. core file size limit= UNLIMITED <DEFAULT>
  Per-proc. data size limit= UNLIMITED <DEFAULT>
  Per-proc. permanent file size limit= UNLIMITED <DEFAULT>
  Per-proc. execution nice priority = 0 <DEFAULT>
  Per-proc. stack size limit= UNLIMITED <DEFAULT>
  Per-proc. CPU time limit= [600.0, 600.0]<DEFAULT>
  Per-proc. working set limit= UNLIMITED <DEFAULT>
  Standard-error access mode = SPOOL
  Standard-error name = mach2:/mnt/test/STDIN.e272
  Standard-output access mode = SPOOL
  Standard-output name = mach2:/mnt/test/STDIN.o272
  Shell = DEFAULT
  Umask = 2
  Restartable = Yes
  Checkpointable = Yes
```

Additional  
request  
information



- ① Created at Mon Jan 23 16:30:03 CST 1989
- ② Mail = [NONE]
- ③ Mail address = test@mach2
- ④ Owner user name at originating machine = test
- ⑤ Import directory: Yes
- ⑥ Per-proc. core file size limit= UNLIMITED <DEFAULT>
- ⑦ Per-proc. data size limit= UNLIMITED <DEFAULT>
- ⑧ Per-proc. permanent file size limit= UNLIMITED <DEFAULT>
- ⑨ Per-proc. execution nice priority = 0 <DEFAULT>
- ⑩ Per-proc. stack size limit= UNLIMITED <DEFAULT>
- ① Per-proc. CPU time limit= [600.0, 600.0]<DEFAULT>
- ② Per-proc. working set limit= UNLIMITED <DEFAULT>
- ③ Standard-error access mode = SPOOL
- ④ Standard-error name = mach2:/mnt/test/STDIN.e272
- ⑤ Standard-output access mode = SPOOL
- ⑥ Standard-output name = mach2:/mnt/test/STDIN.o272
- ⑦ Shell = DEFAULT
- ⑧ Umask = 2
- ⑨ Restartable = Yes
- ⑩ Checkpointable = Yes

- ① Created at specifies the day, date, and time the request was created.
- ② Mail specifies whether or not mail is sent when the job starts and when it finishes.
- ③ Mail address specifies where mail generated for the request is sent.
- ④ Owner user name at originating machine defines the user submitting the request.
- ⑤ Import directory describes whether or not the current working directory is imported for this request.
- ⑥ Per-process core file size limit is the maximum size for this request that a core file for a process can be. If a

process attempts to create a core file exceeding this maximum size, the core file is not written.

- ⑦ Per-process data size limit is the maximum size for this request that a data segment for a process can be.
- ⑧ Per-process permanent file size limit is the maximum size for this request that a file created by a process can be. If a process attempts to write to a file larger than this maximum, CXbatch sends a signal to the process. If nothing is defined in the process to handle that signal, the process is killed.
- ⑨ Per-process execution nice value is the maximum nice value for this request that a process can have. The nice value determines the proportion of CPU time allocated to a process relative to all other processes in the system. The lower the nice value assigned to a process, the higher the proportion of CPU time allocated to that process. A practical range is from -20 to 20.
- ⑩ Per-process stack size limit is the maximum size for this request that a stack segment for a process can be.
- ① Per-process CPU time limit is the maximum total CPU time for this request that a process can use. If this limit is exceeded, CXbatch sends a signal to the process. If this signal is not caught or ignored, the process exits.
- ② Per-process working set limit is the maximum amount of physical memory for this request that a process can use. If this limit is reached, the job is targeted for paging.
- ③ Standard-error access mode indicates that the stderr file is written to spooling directories during job execution and copied to a destination at job termination.
- ④ Standard-error name is the name of the file where error output is sent.
- ⑤ Standard-output access mode indicates that the stdout file is written to spooling directories during job execution and copied to a destination at job termination.
- ⑥ Standard-output name is the name of the file where standard output is sent.
- ⑦ Shell is the shell specified for this request to interpret shell scripts.
- ⑧ Umask defines the default umask for this request.
- ⑨ Restartable specifies whether or not this request can be restarted.
- ⑩ Checkpointable specifies whether or not this request is checkpointable.

---

## Displaying date and time requests will run

You can use the `-m` option to display information about the date and time a request is scheduled to run. Only those requests submitted with a specific date and time have this information available. Figure 3 illustrates the output for this option when a request in the queue is submitted to run at a future date and time.

**Figure 3** Future run time output

```
% qstat -m long
Queue for long jobs.
long@mach2; type=BATCH; [ENABLED, INACTIVE]; pri=32
aliases: l, long_queue, L, LONG
0 exit;0 run; 0 stage; 0 queued;1 wait; 0 hold; 0arrive;

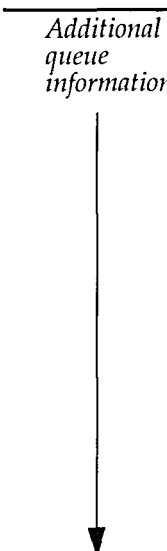
Request 1: Name=myjob Id=297.mach2
Owner=test Priority=31 WAITING Wed Jan 25 00:00:00 CST 1989
```

The request is shown as `WAITING` and includes the day, date, and time it is scheduled to run.

## Displaying additional queue information

You can use the `-x` option to display additional information about the queue without receiving additional information about queue requests. Figure 4 illustrates the output from this option.

Figure 4 Additional queue information

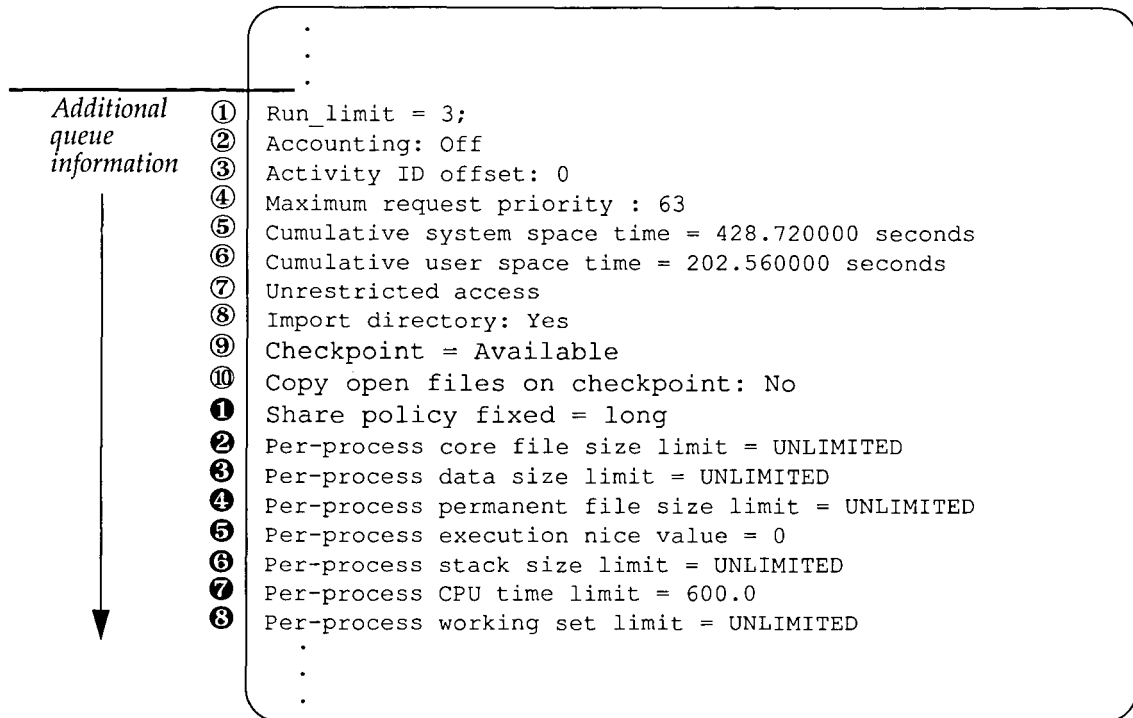


```
% qstat -x long
Queue for long jobs.
long@mach2; type=BATCH; [ENABLED, RUNNING]; pri=32
aliases: 1, long_queue, L, LONG
0 exit; 1 run; 0 stage; 0 queued; 0 wait; 0 hold; 0 arrive;
① Run_limit = 3;
② Accounting: Off
③ Activity ID offset: 0
④ Maximum request priority : 63
⑤ Cumulative system space time = 428.720000 seconds
⑥ Cumulative user space time = 202.560000 seconds
⑦ Unrestricted access
⑧ Import directory: Yes
⑨ Checkpoint = Available
⑩ Copy open files on checkpoint: No
❶ Share policy fixed = long
❷ Per-process core file size limit = UNLIMITED
❸ Per-process data size limit = UNLIMITED
❹ Per-process permanent file size limit = UNLIMITED
❺ Per-process execution nice value = 0
❻ Per-process stack size limit = UNLIMITED
❼ Per-process CPU time limit = 600.0
❽ Per-process working set limit = UNLIMITED

REQUEST NAME   REQUEST ID   USER   PRI   STATE   PGRP
1:myjob        47.mach2    test   31   RUNNING 12103
```

- ❶ Run\_limit limits the number of requests that can run in a queue at any one time. When the limit is exceeded, requests are queued until the number of jobs running is less than the limit.
- ❷ Accounting indicates whether or not batch accounting is activated.
- ❸ Activity ID offset is the number added by CXbatch to a job's activity identification (ID) before the request is executed. The activity ID offset is typically an integer from 1 to 9, with 0 reserved for jobs not submitted to a batch queue. The activity ID is used for accounting purposes.
- ❹ Maximum request priority is the maximum priority at which a request can be submitted to the queue.

Figure 5 Additional queue information



- ⑤ Cumulative system space time for batch queues is the total amount of system time used by batch jobs since the queue was created. For pipe queues, the status indicates the total time used to route jobs.
- ⑥ Cumulative user space time total is the amount of user time used by completed batch jobs since the queue was created.
- ⑦ Unrestricted access specifies the access restrictions placed on the queue. A request submitted by the superuser is an exception to these limitations; superuser requests are always queued. Access restrictions can be:
  - Unrestricted Queue can receive any request from any submitter.
  - Restricted Queue can receive only those requests submitted by a specified group(s) or user(s).
  - Pipeonly Queue accepts requests only from a pipe queue.
- ⑧ Import directory describes whether or not the current working directory for a request is imported (mounted on the

machine processing the request) before the request is executed. This can be:

**Yes** Queue automatically imports the current working directory for any request executing in the queue. The user can override this setting for individual requests using the `-ni` option of `qsub`.

**Available** Allows the user to specify importation of the current working directory for any request submitted to the queue using the `-i` option of `qsub`.

**No** Queue does not allow the current working directory to be imported for requests submitted to the queue. Requests requiring imported directories are rejected when submitted.

⑨ **Checkpointable** specifies whether or not jobs are automatically checkpointed in this queue in the event CXbatch shuts down. This can be:

**Yes** Requests running in the queue are automatically checkpointed when CXbatch shuts down. The user can override this setting for individual requests using the `-nc` option of `qsub`. Requests submitted to this queue can also be manually checkpointed at any time by the owner of the request, a CXbatch operator, or a CXbatch manager, unless they were submitted with the `-nc` option. Requests submitted with the `-nc` option are not automatically checkpointed and cannot be manually checkpointed.

**Available** Requests submitted with the `-c` option to `qsub` are automatically checkpointed when CXbatch shuts down. Requests not submitted with the `-c` option are not automatically checkpointed when CXbatch shuts down. Requests can be manually checkpointed at any time by the owner of the request, a CXbatch operator, or a CXbatch manager, unless they were submitted with the `-nc` option. Requests submitted with the `-nc` option are not automatically checkpointed and cannot be manually checkpointed.

- No Requests are not automatically checkpointed if CXbatch shuts down and cannot be manually checkpointed.
- ⑩ Copy open files defines whether or not CXbatch preserves the state of open files within a process when the request is checkpointed.
- ① Share policy describes the queue share policy. This can be:
- User Charges CPU usage to the user submitting the request.
- Fixed Charges CPU usage to a specified account.
- ② Per-process core file size limit is the maximum size a core file for a process can be. If a process attempts to create a core file exceeding this maximum size, the core file is not written.
- ③ Per-process data size limit is the maximum size the data segment for a process can be.
- ④ Per-process permanent file size limit is the maximum size a file created by a process can be. If a process attempts to write to a file larger than this maximum, CXbatch sends a signal to the process. If nothing is defined in the process to handle that signal, the process is killed.
- ⑤ Per-process execution nice value is the maximum nice value a process can have. The nice value determines the proportion of CPU time allocated to a process relative to all other processes in the system. The lower the nice value assigned to a process, the higher the proportion of CPU time allocated to that process. A practical range is from -20 to 20.
- ⑥ Per-process stack size limit is the maximum size a stack segment for a process can be.
- ⑦ Per-process CPU time limit is the maximum total CPU time that a process can use. If this limit is exceeded, CXbatch sends a signal to the process. If this signal is not caught or ignored, the process exits.
- ⑧ Per-process working set limit is the maximum amount of physical memory that a process can use. If this limit is reached, the job is targeted for paging.

---

## Displaying queue status interactively

You can view the status of queues and requests in queues interactively using the `qwatch` command. This allows you to monitor one or more queues without making repetitive calls to `qstat`.

`qwatch` is a tool designed to be used by system managers in fine-tuning CXbatch queue loads. While it does not interfere with CXbatch daemons, it does use a small amount of CPU resources and should not be used to check on the status of one request. Use `qstat` for checking the status of a request.

The format for the `qwatch` command is

```
qwatch [-i interval] [-c count]
```

where

- `-i interval` specifies how many seconds to wait between each screen update. The default is five seconds.
- `-c count` specifies how many screen updates must occur before updating terminates. If this number is 0, updating continues until manually stopped by pressing **CTRL-c**. The default is 0.

You can display status information about queues either with or without status information about queue requests. When you first invoke `qwatch`, you only receive status information about queues. The first page displays status information on the first five queues on the system. If there are more than five queues, you can display additional pages of queue information by typing the number of the additional page. If there are more than 25 queues, `qwatch` ignores those above 25. Figure 6 on the next page illustrates the initial page output for the `qwatch` command.

**Figure 6** Interactive queue status output

```
% qwatch
① mach1 CXbatch Activity                               Wed Jun 6 15:59:54 1991
②                                     3 queues
      ③          ④          ⑤          ⑥          ⑦
a short@mach1;type=BATCH; [ENABLED,RUNNING];pri=48
⑧  aliases: s
    0 exit;1 run;0 stage;0 queued;0 wait;0 hold;0 arrive;
      ⑨          ⑩          ①          ②          ③          ④          ⑤
b long@mach1; type=BATCH; [ENABLED,RUNNING];pri=32
  aliases: 1
    0 exit;1 run; 0 stage; 0 queued;0 wait;0 hold;0 arrive;

c verylong@mach1; type=BATCH; [ENABLED,INACTIVE];pri=16
  aliases: 1
    0 exit;0 run;0 stage;0 queued;0 wait;0 hold;0 arrive;

Page 1 of 1 Type 'a' - 'c' for queue.
```

- ① Name of the host containing the queue and the day, date, and time the status was taken.
- ② Number of queues on the host.
- ③ Name of the queue and what host it is configured on.
- ④ Type of queue. This can be:

BATCH	Executes CXbatch requests.
PIPE	Routes CXbatch requests to queues that can execute them.
- ⑤ Indicates whether or not the queue can accept requests. This can be:

ENABLED	CXbatch is running on the local machine and the queue is accepting requests.
DISABLED	CXbatch is running on the local machine but the queue is not accepting requests.
CLOSED	CXbatch is not running on the local machine.
- ⑥ Indicates whether or not the queue can execute requests. This can be:

INACTIVE	Requests in the queue are permitted to run; none are running.
RUNNING	Requests in the queue are permitted to run; some are running.

**STOPPING** New requests sent to the queue are not permitted to run, but requests that are currently running are allowed to complete.

**STOPPED** Requests in the queue are not permitted to run; none are running.

**SHUTDOWN** CXbatch is not running on the local machine.

- ⑦ Specifies the inter-queue priority. This priority affects which queue is looked at first for the next job to run. Priority can be any number from 0 to 63; 0 is the lowest priority and 63 the highest.
- ⑧ Alternate names by which the queue can be referenced.
- ⑨ Number of requests in this queue in a state of exiting.
- ⑩ Number of requests in this queue in the state of running.
- ❶ Number of requests in this queue in the staged state, which means the request has completed executing and is moving the stdout and stderr files to the appropriate destination directory.
- ❷ Number of requests in this queue in a state of being queued.
- ❸ Number of requests in this queue in a state of waiting.
- ❹ Number of requests in this queue in a state of holding.
- ❺ Number of requests in this queue in a state of arriving.

If you wish to display status information about the requests in queues, enter the letter assigned to the queue. The first page displays status information on the first 16 requests in the queue. If there are more than 16 requests, you can display additional pages by typing the number of the additional page. `qwatch` can display up to nine pages or 144 requests. Figure 7 on the next page illustrates the output for selected queues.

**Figure 7 Interactive queue request output**

```
mach1 CXbatch Activity Thu Aug 23 16:00:44 1990
long@mach1; type=BATCH; [ENABLED,RUNNING];pri=32
aliases: 1
0 exit;1 run;0 stage;2 queued;0 wait;0 hold;0 arrive;
REQUEST NAME    REQUEST ID    USER    PRI    STATE    PGRP
1:STDIN        346.mach1    johndoe 31    RUNNING  16257
1:STDIN        346.mach1    johndoe 31    QUEUED
1:STDIN        346.mach1    johndoe 31    QUEUED
```

Page 1 of 1 Type '-' for queues.

Return to the initial screen by typing a dash (-). Exit `qwatch` by pressing **CTRL-c**.

The information provided about each queue request is:

REQUEST NAME	REQUEST ID	USER	PRI	STATE	PGRP
1:STDIN	346.mach1	johndoe	31	RUNNING	16257
①	②	③	④	⑤	⑥

- ① Name assigned to the request.
- ② Unique identifier assigned to request when it is submitted to the queue.
- ③ User submitting the request.
- ④ Intra-queue priority assigned to request. This priority affects which job in a queue is executed next. This number can be from 0 to 63; 0 is the lowest priority and 63 the highest.
- ⑤ State of the request. This can be:

ARRIVING	Request is arriving at the queue.
CHECKPOINTED	A failed attempt was made to restart the checkpointed request. You can attempt to manually restart the request using the <code>qrestart</code> command. Refer to Chapter 5, "Checkpoint Restart features," for details on using this command.
DEPARTING	Request is departing from the queue but has not yet been received by the destination queue.
EXITING	Batch request has completed executing and will exit from the system after the required output files are returned to their intended destinations.

HOLDING	A hold has been placed on the request preventing it from entering any other state.
QUEUED	Request is queued and eligible for running or routing. This is the most common state.
ROUTING	Request has reached the head of a pipe queue and is being routed to another queue.
RUNNING	Request has reached its final destination batch queue and is executing.
WAITING	Request is waiting for a specified amount of time to pass before attempting to execute. This could be because it was submitted with a future date and time specified for running, or because a pipe queue could not route the request and will attempt to route it later.
SUSPENDED	Request is suspended from running. It will remain suspended until manually resumed.

- ⑥ Process group of the request, if available to the local CXbatch daemon. This information is displayed only for processes that are running.

For more information on queue or request properties, refer to the `qstat(1)` man page.

## Displaying queue limits

You can display the resource limits of a queue using the `qlimit` command. Queue limits are set by the system manager when the queue is created.

CXbatch supports many types of resource limits, but not all operating systems support all of them. The `qlimit` command displays those limits that CXbatch can directly support under the operating system on the indicated host machine. The format is

```
qlimit [hostname ...]
```

where *hostname* is the desired host machine. If *hostname* is omitted, the local host is assumed.

Figure 8 illustrates use of the `qlimit` command to display the local host limits and shell strategy.

**Figure 8** Queue limit output

```
% qlimit
① Core file size limit (-lc)
② Data segment size limit (-ld)
③ Per-process permanent file size limit (-lf)
④ Nice value (-ln)
⑤ Stack segment size limit (-ls)
⑥ Per-process cpu time limit (-lt)
⑦ Working set limit (-lw)
⑧ Shell strategy = FREE
```

The information displayed is:

- ① Per-process core file size limit is the maximum size a core file for a process can be. If a process attempts to create a core file exceeding this maximum size, the core file is not written.
- ② Per-process data size limit is the maximum size the data segment for a process can be.
- ③ Per-process permanent file size limit is the maximum size a file created by a process can be. If a process attempts to write to a file larger than this maximum, CXbatch sends a signal to the process. If nothing is defined in the process to handle that signal, the process is killed.
- ④ Per-process execution nice value is the maximum nice value a process can have. The nice value determines the proportion of CPU time allocated to a process relative to all other processes in the system. The lower the nice value

assigned to a process, the higher the proportion of CPU time allocated to that process. A practical range is from -20 to 20.

- ⑤ Per-process stack size limit is the maximum size a stack segment for a process can be.
- ⑥ Per-process CPU time limit is the maximum total CPU time that a process can use. If this limit is exceeded, CXbatch sends a signal to the process. If this signal is not caught or ignored, the process exits.
- ⑦ Per-process working set limit is the maximum amount of physical memory that a process can use. If this limit is reached, the job is targeted for paging.
- ⑧ Shell strategy specifies the ConvexOS shell that is used to interpret the script commands submitted through CXbatch if one is not specified when the request is submitted using `qsub -s`. The shell strategy can be:

FIXED	The system manager specifies the shell that interprets script file commands.
LOGIN	CXbatch uses the shell specified in the <code>/etc/passwd</code> file for the user submitting the job. CXbatch does not read the script file for a shell specification.
FREE	CXbatch uses the shell specified in the <code>/etc/passwd</code> file for the user submitting the job unless a shell is specified in the script file. When a user submits a batch request, CXbatch supplies the name of the script file to the login shell as standard input. The login shell reads the first line of the script file. If the first line specifies a shell, that shell is used to interpret the script file commands instead. If there is no shell specified in either of these places, CXbatch uses <code>sh</code> .

For further information, refer to the `qlimit(1)` man page.

---

## Displaying status of related processes

You can display the status of CXbatch-related processes from batch queues on the local system using the `qps` command. Pipe queues and remote queues are not significant because they do not have processes that affect the local machine.

The format for the `qps` command is

```
qps [option][process-id]
```

where

*option* controls which CXbatch-related processes are reported. If no options are given, `qps` displays information about all CXbatch-related processes, including the daemon processes. *option* can be one of the following:

`-r` displays information on a specific request. Only the processes related to the specified request are displayed.

`-p` determines if a specific process is running under CXbatch. If it is, the following message displays showing the queue and request to which the process is related:

```
Process ID 508 is being executed by
Request ID 3782 in the short queue.
```

If it is not running, the following message displays:

```
Process ID 508 is not being executed
by CXbatch.
```

With `qps`, the top-level daemons are not considered to be running under the CXbatch system, but CXbatch shepherd processes are.

`-q` is a silent version of the `-p` option. Nothing is printed to the screen, but the exit status of `qps` is set appropriately.

*process-id* is the identifying number of the process for which you wish to see the status.

Figure 9 illustrates the output for the `qps` command.

**Figure 9** Status of CXbatch-related processes

①	②	③	④	⑤	⑥
QUEUE	REQ	PID	STAT	TIME	COMMAND
<daemon>		21052 S		0:00	CXbatch logdaemon
<daemon>		21053 S		0:00	CXbatch nqsdaemon
<daemon>		21054 S		0:00	CXbatch netdaemon
<daemon>		210571 S		0:00	CXbatch queued
long	508.mach	112535 S		0:00	CXbatch shepherd
long	508.mach	112536 S	R	0:00	/bin/make -k ...
long	508.mach	112539 S	T	0:00	/bin/make -k ...
long	508.mach	112535 S	D	0:00	tsch
long	508.mach	112535 S	R	0:00	tsch

The information provided is:

- ① QUEUE is the queue that contains the request related to the process.
- ② REQ is the identification number assigned to request initiating process and machine where it originated.
- ③ PID is the unique identification number assigned to the process.
- ④ STAT is the status of the process. This can be:
  - R Process is runnable.
  - T Process is stopped.
  - P Process is in page wait.
  - D Process is in disk wait or other short term wait.
  - S Process is active (sleeping for less than 20 seconds).
  - I Process is idle (sleeping for more than 20 seconds).
  - Z Process is trying to exit but the parent of the process is not waiting for it (zombie).
- ⑤ TIME is the amount of CPU time used so far.
- ⑥ COMMAND is the command that started the process.

---

## Displaying contents of shell scripts

You can view the contents of a shell script used in a batch request using the `qjlist` command. To display a shell script, you must either:

- Own the request
- Have superuser privileges
- Be designated as a batch operator or manager

The format for the `qjlist` command is

```
qjlist request_id.[@hostname]
```

where

*request\_id* is the number assigned to the request when it is submitted to CXbatch. By default, CXbatch displays the *request\_id* of the request at the submitter's terminal when a batch request is successfully submitted. You can also get this number using the `qstat` command. Refer to the "Displaying status of queues" section in this chapter for details on using `qstat`.

*hostname* is the host machine that receives the request. If *hostname* is omitted, the local host is assumed.

Figure 10 illustrates the contents of the `52.mach2` shell script using the `qjlist` command.

**Figure 10** Contents of a shell script

```
% qjlist 52.mach2
sleep 20
echo "It is time to go home"
CTRL-d
```

For further information, refer to the `qjlist(1)` man page.

This chapter describes how to submit a request to a CXbatch queue using the `qsub` command. It describes:

- Three methods for submitting batch requests
- Format of the `qsub` command
- Controlling how, when, and where a request is submitted and executed
- How to redirect output and error information
- How to specify per-process and per-request resource limits
- How to send notification that a request is starting or ending
- How to embed `qsub` command options in a shell script

---

## Three methods for submitting batch requests

A batch request consists of one or more commands submitted by a user or a user program to CXbatch. You can submit a request for batch execution in one of three ways:

- With interactive commands
- In a script file
- In a compiled program

Each of these methods are described in the following sections. See the “qsub command options” section in this chapter for details on using qsub command options.

---

### Using interactive commands

You can submit batch requests interactively by issuing the qsub command and following it with the commands you want to execute. You complete the transaction by pressing **CTRL-d**. Figure 11 illustrates an interactive session.

Figure 11 Submitting a batch job interactively

```
% qsub
sleep 20
echo "It is time to go home"
CTRL-d
```

---

### Using a script file

You can submit a batch request using a script file. To do this, create a file that contains all the commands you want to execute. Then, issue the qsub command and include the name of the script file on the command line. Figure 12 illustrates how to create a script file named *my.script* and submit it to CXbatch. A script file is run by a shell, so the script file must contain commands the shell can recognize.

Figure 12 Submitting batch job using a script file

```
Create script { % cat > my.script
                  echo "hello world"
                  CTRL-d
Submit script to CXbatch → % qsub my.script
                             Request 271.mach2 submitted to queue: long.
```

When you submit a script file, you must include the name of the file on the command line. Make sure that script files submitted to CXbatch are completely self-contained so that when the script

exits, no background child processes remain. If a request sends mail to a local user, the mail program runs a background process to deliver the mail and does not wait for it to complete.

If you send mail at the end of a request's script, have the request sleep for a few seconds just before exiting to allow the delivery process to complete. If the request exits before the mail delivery process completes, the shepherd process aborts the delivery and the mail disappears. This does not occur if you send mail to a remote machine or if you pass the `-v` (verbose) option to the mail program.

---

## Using a compiled program

You can submit a batch request using a compiled program (binary code). To do this, include the name of the program in a script file or enter the name as standard input following the `qsub` command line. Figure 13 illustrates how to submit the compiled program `test` by entering the program name as standard input.

**Figure 13** Submitting a batch job using a compiled program

```
% qsub
test
CTRL-d
```

Figure 14 illustrates how to submit the compiled program `test` by creating a shell script file and submitting that script file using the `qsub` command.

**Figure 14** Submitting a compiled program using a script file

```
% cat > myjob
test
CTRL-d
% qsub myjob
```

Figure 15 illustrates an incorrect way to submit the compiled program `test` as a batch request.

**Figure 15** Incorrect way to submit compiled program

```
% qsub test
Line length exceeds 1025 characters.
Script file line 2.
```

This command is incorrect because `qsub` is expecting either the name of a shell script or input from the terminal.

---

## qsub command options

Use the `qsub` command to submit a request to a batch queue. The format for the `qsub` command is

```
qsub option [option ...]
```

where *option* can be one or more values that allow you to:

- Control how, when, and where the request is run
- Redirect output files and error messages
- Specify resource limits
- Request status updates

Unlike ConvexOS command format, which allows you to string several options together after a single dash, CXbatch command format requires that you list each option separately. For example, to execute the `qsub` command with two options, you must enter

```
qsub -option1 -option2
```

The following sections describe the options available with the `qsub` command, except for those related to checkpointing. Refer to Chapter 5, “Checkpoint Restart features” for information on options related to checkpointing.

---

## Controlling run requests

The `qsub` command includes options to control how, when, and where a request is run. Table 2 lists each option and its meaning.

**Table 2** `qsub` options that control run requests

Option	Definition
-a [date] [time]	Run request after a stated time
-b bill_account	Run request under specified billing account
-h	Place request on hold
-i	Import current working directory
-l	Run request under login shell
-ni	Do not import current directory
-p priority	Set intra-queue priority
-q queue	Specify queue to submit request
-r name	Assign name to request
-s shell	Specify shell to interpret script files
-x	Export value of all environment variables
-z	Submit request silently

The following sections explain these options in detail and give examples of their use.

## Specifying a future run time

Unless you specify otherwise, a batch request is eligible for execution immediately after you issue the `qsub` command. You can use the `-a` option to specify a future run time. The format for this option is

```
qsub -a [date] [time]
```

where

*date* specifies the date the job should run. If no date is indicated, the current month, day, and year are assumed.

If you specify a date, the year is optional. If it is omitted, the current year is assumed. You can abbreviate the names of the month and day by using the first three (or more) characters of the word. For example, you can indicate a day of the week (like, Mon, Tue). A period following the abbreviation is optional. Uppercase and lowercase letters are equivalent, so that "WeD" is treated the same as "WED" or "wed."

You can use the words "today" or "tomorrow" rather than specifying a month, day, and year.

*time* specifies the time of day the job should run. A 24-hour clock is the default, so if 1:00 is specified without "am" or "pm," CXbatch interprets it as 01:00 or 1:00 a.m.

You can use a 12-hour clock as the basis for the time by including an "am" or "pm" designation with the time. 1:00 p.m. is shown as 1:00pm or 1pm. If 1:00 p.m. is specified, it is interpreted as 13:00 on the 24-hour clock.

"12am" refers to 0:00:00, "12n" refers to noon, and "12 pm" refers to 24:00:00. You can optionally use the words "midnight" or "noon."

You can include a time-zone designation, as in "13:01-PDT." CXbatch understands European time zones. If a time-zone designation is omitted, the local time zone is assumed, with daylight savings time included when appropriate.

The order of *date* and *time* is irrelevant. The following formats are equivalent:

```
qsub -a [date] [time]
```

```
qsub -a [time] [date]
```

If you use *date* and *time*, enclose the entire string in double quotes. You must also enclose either *date* or *time* in double quotes if you include spaces within either option.

```
"January 1, 1989, 12:31"
```

shows the correct way to indicate combined *date* and *time* as well as the correct way to use spaces. This rule applies when this option is included on the command line or in the shell script. Refer to the last section in this chapter, "Embedded Options," for information on embedding options in shell scripts.

Some valid examples of *date* and *time* are:

```
"01-Jan-1988 13:00"  
tomorrow  
"today 5pm"
```

For example, the following command executes the script file *test* on July 4, 2026, at 12:31 Eastern Daylight Time:

```
qsub -a "July 4, 2026 12:31-EDT" test
```

### Specifying a billing account

If accounting is set up for the queue, you can use the `-b` option to specify that a request run under a specific billing account. For example, the following command bills the script file *my\_job* to the *activity1* billing account:

```
qsub -b activity1 my_job
```

### Putting a request on hold at submission

You can use the `-h` option to place a request on hold at the time you submit the job. For example, the following command places the script file *my\_job* on hold when it is submitted:

```
qsub -h my_job
```

### Controlling importation of the current directory

Often, batch requests require access to all files in all subdirectories of your current working directory. Because of this, you may need to make the files in your current directory available to batch requests. Granting this access is called *importing* the current working directory.

The system manager defines the ability to import directories on a per-queue basis. You can use the `qstat` command to view whether or not the queue you are using allows importation of current directories. See Chapter 2, "Displaying information," for details on using the `qstat` command.

If you need to import files and the queue you are using does not allow this, see your system manager. If you do not want importing or it has been denied by the system manager, CXbatch expects to find any files required to run the job in your home directory.

You can use the `-i` option to import the current working directory before executing the batch request. If the execution queue is on the same machine as the current working directory, CXbatch changes directories before it starts. If the execution queue is on another machine, CXbatch imports the directories and subdirectories with NFS by making temporary mount points in the `/tmp` directory on the originating machine. Be aware of any local automatic cleanup facilities of `/tmp` that might affect these NFS mount points.

For example, the following command imports the current working directory for the script file *my\_job*:

```
qsub -i my_job
```

CXbatch cannot import a current working directory that is already within an NFS file system. For example, if you are on a machine called *sleepy* in a file system imported from *happy*, and you submit a job with the `-i` option to a queue on *grumpy*, the file system cannot be imported from *happy* to *grumpy*.

You can use the `-ni` option to prevent the importation of the current working directory.

For example, the following command does not import the current working directory for the script file *my\_job*:

```
qsub -ni my_job
```

For further information on the import option, refer to the *CONVEX CXbatch System Manager's Guide* and the `qsub (1)` man page.

## Specifying use of login shell

During configuration of the CXbatch system, the system manager defines a shell strategy for each queue. A shell strategy specifies the ConvexOS shell that is used to interpret the script commands submitted through CXbatch if one is not specified by the user when the request is submitted. The shell strategy can be:

- **Fixed**—The system manager specifies the shell that interprets script file commands.
- **Login**—CXbatch uses the shell specified in the `/etc/passwd` file for the user submitting the job. CXbatch does not read the script file for a shell specification.
- **Free**—CXbatch uses the shell specified in the `/etc/passwd` file for the user submitting the job unless a shell is specified in the script file. When a user submits a batch request, CXbatch supplies the name of the script file to the login shell as standard input. The login shell reads the first line of the script file. If the first line specifies a shell, that shell is used to interpret the script file commands instead. If there is no shell specified in either of these places, CXbatch uses `sh`.

You can use the `qlimit` command to discover which strategy is in effect for your system. See Chapter 2, “Displaying information” for details on using the `qlimit` command.

Unless otherwise specified, CXbatch runs jobs using a noninteractive shell. You can use the `-l` option to specify using an interactive login shell.

If you specify use of a login shell, your default login shell (determined by the shell strategy specified for this queue) is used unless you specify otherwise. You can use the `-s` option to specify an alternate shell. The format for this option is

```
qsub -s shell-name
```

where *shell-name* indicates the absolute path name on the executing host of the shell to be used.

If you use a C shell, CXbatch gets environment variables from `/etc/login` and the `.login` file in your home directory. If you use the Bourne or Korn shell, CXbatch gets environment variables from `/etc/profile` and the `.profile` file in your home directory.

To prevent `stty`, `tset`, and `msgs` from running during batch jobs, you can add the string `ENVIRONMENT=BATCH` to your `.login`, `.profile`, or `.cshrc` file so shell scripts can test for batch request execution. Then place an `if` statement in the `.login`, `.profile`, or `.cshrc` file.

If your login shell is a C shell, the following `if` statement in your `.login` file prevents `stty`, `tset`, and `msgs` from running during batch jobs. C shell always reads your `.cshrc` file regardless of whether or not it is running as a login shell.

```
if (! $?ENVIRONMENT) then
  stty crt erase ^H kill ^U
  tset -Q
  msgs -q
endif
```

If your login shell is the Bourne shell or Korn shell, the following `if` statement in your `.profile` file has the same effect of preventing `stty`, `tset`, and `msgs` from running during batch jobs.

```
if test "$ENVIRONMENT" != "BATCH"
then
  stty crt erase ^H kill ^U
  tset -Q
  msgs -q
fi
```

As an alternative, you might use

```
if ($?ENVIRONMENT) exit
```

## Setting priority on batch jobs

You can use the `-p` option to assign an intra-queue priority to a job request. This priority does not determine execution priority of the request; it is used to determine only relative ordering of requests within a queue. The ConvexOS nice value determines execution priority of requests. Refer to the `nice(1)` man page for details.

When CXbatch adds a request to a queue, it compares the intra-queue priority of the request with the intra-queue priority of all existing requests. It then places the request in the queue ahead of those requests with a lower priority and behind existing requests with a higher priority. When the priority of the new request is equal to the priority of an existing request, the existing request takes precedence over the new request.

If you do not specify an intra-queue priority, CXbatch assigns the default intra-queue priority to each request.

The specified priority can be any number from 0 and 63; 0 is the lowest priority and 63 the highest. You cannot assign a priority to a request that is higher than the maximum request priority assigned to the queue.

For example, the following command sets a priority of 48 on the script file `test`:

```
qsub -p 48 test
```

## Submitting to a specific queue

Unless otherwise specified, CXbatch determines which queue should receive a request by checking the value of the `QSUB_QUEUE` environment variable in the `.login` file of the user submitting the job. If this environment variable is not defined, the request is submitted to the default batch queue defined by the system manager. If no variable is found and there is no default queue defined, an error message is displayed and the request is not accepted.

You can use the `-q` option to specify a queue to run the request. If you do not specify a queue, the job is submitted to the default queue. The format for this option is

```
qsub -q queuename [@hostname]
```

where

*queuename* is the name of the queue.

*hostname* is the name of the host machine for that queue. If *hostname* is omitted, CXbatch submits the request to the named queue on the local machine. Your system manager can tell you which host machines are running CXbatch.

For example, the following command submits the script file, *test*, to the short queue:

```
qsub -q short
```

### Naming a request

Unless otherwise specified, CXbatch assigns a request the same name as the script file (less the path name) given on the command line. If there is no script file, CXbatch assigns the default name STDIN. You can use the `-r` option to assign a name to the request. The format for this option is

```
qsub -r request-name
```

where *request-name* is the name to be assigned. The *request-name* option can have any number of characters, but its display is truncated to fifteen characters. In the actual output file, *request-name* is truncated to seven letters, plus regular extensions. If it begins with a digit, CXbatch prefixes the name with the letter R.

For example, the following command sequence assigns the name *myjob* to the interactive batch session:

```
qsub -r myjob  
echo "hello world"  
CTRL-d
```

## Exporting all environment variables

Unless otherwise specified, CXbatch uses the current value of the following ConvexOS environment variables when a batch request is executed: HOME, SHELL, PATH, USER, MAIL. Figure 16 illustrates the default environment.

**Figure 16** Default environment

```
ENVIRONMENT=BATCH
HOME=/mnt/test
MAIL=/usr/spool/mail/test
PATH=:/usr/ucb:/bin:/usr/bin
QSUB_AID=0
QSUB_HOME=/mnt/test
QSUB_HOST=mach2
QSUB_PATH=./mnt/test/bin:/usr/convex:/bin:/usr/ucb:/usr/bin:/usr/local/bin
QSUB_REQID=278.mach2
QSUB_REQNAME=STDIN
QSUB_SHELL=/bin/csh
QSUB_USER=test
```

You can use the `-x` option to export the value of all other ConvexOS environment variables with a batch request. Figure 17 illustrates sample environment variables exported when the `-x` option is used to export all environment variables.

**Figure 17** All environment variables exportation

```
EDITOR=gnumacs
ENVIRONMENT=BATCH
HOME=/mnt/test
MAIL=/usr/spool/mail/test
MCSDIR=/test/work/lib/MCS
NFED=/usr/local/bin/gnumacs
PATH=:/usr/ucb:/bin:/usr/bin
PRINTER=vaxip
QSUB_AID=0
QSUB_HOME=/mnt/test
QSUB_HOST=mach2
QSUB_PATH=./mnt/test/bin:/usr/convex:/bin:/usr/ucb:/usr/bin:/usr/local/bin
QSUB_REQID=279.mach2
QSUB_REQNAME=STDIN
QSUB_SHELL=/bin/csh
QSUB_USER=test
QSUB_WORKDIR=/mnt/test
SHELL=/bin/csh
TERM=sun
USER=test
```

### Submitting a request silently

When a batch request is successfully submitted, CXbatch displays the *request\_id* assigned to the request at the user's terminal who submitted the request. You can stop the display of this message by using the `-z` option. This option has no effect on error messages; they are always displayed.

---

## Redirecting output files and error messages

If you submit a request without supplying a request name, output goes to a file named `STDIN.oseq#` on the machine that originated the request where `seq#` is the sequence number assigned to the request in the queue. Errors go to a file named `STDIN.eseq#` on the machine that originated the request, where `seq#` is the sequence number assigned to the request in the queue. The sequence number assigned to the request is displayed on the user's terminal after the request is submitted.

When submitting a request from a machine running an automounting daemon, you must explicitly specify output and error files with full path names. Otherwise, output and error files may not be placed where expected.

If you are near your quota limit on your home file system when you submit a job, you will lose the information normally stored in the standard output file, `STDIN.oseq#`. If the request runs on the same machine as it was submitted, you will receive notification that the output could not be saved. If the request ran on a remote machine, you will not receive notification.

The `qsub` command includes options to control direction of output and error messages. Table 3 lists each option and its meaning.

**Table 3** `qsub` options that redirect output and error messages

Option	Definition
<code>-e</code>	Redirect standard error output
<code>-eo</code>	Send error output to standard output file
<code>-ke</code>	Leave standard error output file on executing machine
<code>-ko</code>	Leave standard output file on executing machine

The following sections explain each option in detail and give examples of its use.

### Redirecting standard error output

You can use the `-e` option to redirect error messages to a specified file name instead of the default file. The format for this option is

```
qsub -e [hostname:] [[/] [pathname/] filename
```

where

*hostname* is the host machine to contain the error file. If you omit *hostname*, CXbatch sends error messages to the machine that originated the request. If you also use the `-ke` option, CXbatch keeps the error messages on the machine that executes the request.

If you omit *hostname* and the `/` preceding the path name, CXbatch sends the error messages to the specified file name in the current working directory, provided that `-ke` is not used. Otherwise, any portion of the path name included on the command line is interpreted in relation to the user's home directory on the standard error destination machine.

*pathname* is the path name of the designated file.

*filename* is the file name to contain the error output.

For example, the following command sends error messages for the *test* script file to the file *qsub07.redirected\_err* on the CONVEX machine:

```
qsub -e mach2:qsub07.redirected_err test
```

You cannot use the `-e` option if you use the `-eo` option. Refer to the "Redirecting error messages to the standard output file" section in this chapter for details on this option.

### Redirecting error messages to the standard output file

You can use the `-eo` option to redirect error messages to the standard output file instead of the standard error file. For example, the following command redirects error messages for the script file *test* to the standard output file:

```
qsub -eo test
```

Errors are then directed to the file *STDIN.oseq#*. For example, if the sequence number assigned to this request is 282 on host named *mach2*, errors are directed to a file named *STDIN.o282*.

You cannot use the `-eo` option if you use the `-e` or the `-ke` options. Refer to the "Redirecting standard error output" section in this chapter for details on the `-e` option. Refer to the "Redirecting error output on the executing machine" section in this chapter for details on the `-ke` option.

## Redirecting error output on the executing machine

You can use the `-ke` option to have the standard error output file created on the machine executing the request rather than on the machine originating the request. For example, the script file *test* is submitted from the host named *mach2* but executes on the host named *liberty*. The following command creates the standard error output file for the script file on the host named *liberty*, the host executing the request:

```
qsub -ke test
```

You cannot use the `-ke` option if you use the `-eo` option.

You should use the `-e` option to specify the path or file name. If you do not specify a path or file name, the file is placed in your home directory. Do not specify a hostname with the `-e` option, or the `-ke` option is ignored.

## Redirecting standard output on executing machine

You can use the `-ko` option to have the standard output file created on the machine executing the request rather than on the machine originating the request. For example, the script file *test* is submitted from the host named *mach2* but executes on the host named *liberty*. The following command creates the standard output file for the script file on the host named *liberty*, the host executing the request:

```
qsub -ko test
```

You cannot use the `-ko` option if you specify a specific *hostname* in conjunction with the `-o` option. The `-ko` option is also meaningless if the current directory has been imported by the executing machine.

You should use the `-e` option to specify the path or file name. If you do not specify a path or file name, the file is placed in your home directory.

## Redirecting standard output

You can use the `-o` option to send standard output to a specified file name instead of the default file name. The format for this option is

```
qsub -o [hostname:][[/][pathname/]filename
```

where

*hostname* is the host machine where the output file will be sent. If you omit *hostname*, CXbatch sends the output to the machine that originated the request. If you also use the `-ko` option, CXbatch keeps the output on the machine that executes the request.

If you omit *hostname* and the `/` preceding the path name, CXbatch sends the output to the specified file name in the current working directory, provided that `-ko` is not used. Otherwise, any portion of the path name included on the command line is interpreted in relation to the current working directory at the time of submission.

*pathname* is the path name of the designated file.

*filename* is the file name where the standard output will be sent.

For example, the following command sends standard output for the *test* script file to the file *myjob.redirect\_out* on the local machine:

```
qsub -o myjob.redirect_out test
```

## Appending accounting information to stdout output file

You can use the `-y` option to append accounting information to your request's standard output file. Accounting must be enabled for the queue in which the job runs to use this option.

This information saved includes queue, host, sequence number, remote host, submission time, completion time, time spent executing in user mode, and time spent executing in the system.

For more information on accounting, please refer to the *CONVEX CXbatch System Manager's Guide*

---

## Specifying resource limits

qsub has options that allow you to set limits on resources consumed by each process within a batch request (per-process limits) and on resources that all processes within a batch request together can consume (per-request limits). If you do not specify limits, the limits set for the queue are used. If no limits are set for the queue, limits do not apply.

The format for these options is

```
qsub option size-limit [, warning-limit]
```

where

*option* is an option that controls a resource limit. See Table 4 for per-process limit options and for per-request limit options.

*size-limit* is the desired maximum limit.

*, warning-limit* is the limit at which a warning signal is delivered to the executing process. If no *warning-limit* is set, no warning signal is delivered.

Table 4 lists each option that controls a per-process or per-request resource limit and the limit it defines.

**Table 4** Per-process and per-request limits

Option	Limit defined	Value
-lc	Core file size	A number representing less than 100,000,000 bytes.
-ld	Data segment	A number representing less than 100,000,000 bytes.
-lf	Permanent file	A number representing less than 100,000,000 bytes.
-lm	Memory	A number representing less than 100,000,000 bytes.
-ln	Nice value	A number between -64 and +64.
-ls	Stack segment	A number representing less than 100,000,000 bytes.
-lt	CPU time	hours:minutes:seconds:milleseconds
-lv	Temporary file	A number representing less than 100,000,000 bytes.
-lw	Working set	A number representing less than 100,000,000 bytes.

With the exception of the nice value, not all operating systems support these per-process limits. If you submit a request specifying a limit and the machine that runs the batch request does not enforce the specified limit, the limit is ignored.

Other implementations of NQS may support features that CXbatch does not. For example, CXbatch does not support the per-process memory or temporary file limits. However, if the batch system on the execution machine supports these limits, they are enforced.

For further information on resource limits, refer to *CXbatch System Manager's Guide* and `qsub(1)` and `getrlimit(2)` man pages.

CXbatch ignores all per-request limits for batch requests submitted to a CONVEX machine for execution on a CONVEX machine. Table 5 lists each per-request option that is not applicable if the destination machine is a CONVEX machine.

**Table 5** Invalid per-request limits for CONVEX machine

Option	Limit defined
-lf	Permanent file
-lm	Memory
-lt	CPU time
-lv	Temporary file

---

## Sending notification of request status

CXbatch sends mail to the user submitting the request if it cannot execute the submitted shell script. Unless otherwise specified, it does not notify the user if the request is accepted and executed. The `qsub` command includes options that allow you to send mail when a request has started or finished executing. Although mail originates on the specified machine, it is routed according to standard mail-forwarding mechanisms. Table 6 lists the options that allow sending of notification and their meanings.

**Table 6** `qsub` options that allow sending notification

Option	Definition
<code>-mb</code>	Send mail when request begins execution
<code>-me</code>	Send mail when request completes execution
<code>-mu</code>	Send mail to specific user
<code>-t</code>	Signal a process when request completes execution

The following sections explain these options and give examples of their uses.

### Sending mail when request begins execution

You can use the `-mb` option to have CXbatch send you mail on the originating machine when the request begins execution. For example, the following command sends mail to the user submitting the request named *myjob* when the job begins execution:

```
qsub -mb myjob
```

Figure 18 illustrates the mail received as a result of using the `-mb` option when submitting the job.

#### Figure 18 Mail received when job starts executing

```
Request name: myjob  
Request owner: johndoe  
Mail sent at: Fri Aug 24 10:43:07 CDT 1990
```

## Sending mail when request completes execution

You can use the `-me` option to have CXbatch send you mail on the originating machine when the request finishes execution. For example, the following command sends mail to the user submitting the request named *myjob* when the job finishes execution.

```
qsub -me myjob
```

Mail sent as a result of using the `-me` option contains a summary of CPU time used by the request. Figure 19 illustrates an excerpt of a mail message received as a result of using `-me`.

**Figure 19** Mail received when job completes executing

```
Request name: STDIN
Request owner: johndoe
Mail sent at: Fri Aug 24 10:47:36 CDT 1990
Request exited normally.
_Exit() value was: 0.
```

## Sending mail to specified user

You can use the `-mu` option to have CXbatch send mail to a particular user or host instead of to yourself. The format for this option is

```
qsub -mu username [@hostname]
```

where

*username* is the user to receive the mail.

*hostname* is the machine where *username* is located.

For example, the following command sends mail to user *smith* on host *mach2* when the request begins and completes execution.

```
qsub -mu smith mach2
```

Figure 20 illustrates the mail received by the user *smith* as a result.

### Figure 20 Mail sent to another user

```
>From daemon Tue Jan 24 15:39:55 1990
Received: by mach2 (5.51/4.7)
 id AA18211; Tue, 24 Jan90 15:39:53 CST
Date: Tue, 24 Jan90 15:39:53 CST
From: root (Superuser)
Subject: CXbatch request: 2172.mach2 beginning.
Apparently-To: smith
Status: R

Request name: STDIN
Request owner: johndoe
Mail sent at: Tue Jan 24 15:39:52 CST 1990

>From daemon Tue Jan 24 15:40:03 1990
Received: by mach2 (5.51/4.7)
 id AA18224; Tue, 24 Jan90 15:40:01 CST
Date: Tue, 24 Jan90 15:40:01 CST
From: root (Superuser)
Subject: CXbatch request: 2172.mach2 ended.
Apparently-To: smith
Status: R

Request name: STDIN
Request owner: johndoe
Mail sent at: Tue Jan 24 15:40:01 CST 1990

Request exited normally.
_Exit() value was: 1.
```

---

## Signalling processes when a request completes executing

You can use the `-t` option to have CXbatch signal a particular process when the request has executed. The format for this option is

```
qsub -t process-id
```

where *process-id* is the process to be signalled.

One of the following signals is sent, depending on how the request completes:

- **SIGTERM**—Request completed normally.
- **SIGUSR1**—Request was aborted while running.
- **SIGUSR2**—Request was deleted before it ran.

---

## Embedded options

You can include `qsub` options in a script file and submit it as a batch request. The options must be included in the first comment block of the script file for CXbatch to recognize them.

Begin comment lines with either `#` (ConvexOS) or `#!` (COVUEshell). CXbatch interprets both the ConvexOS (`#`) and COVUEshell (`#!`) comment indicators. (COVUEshell is an optional CONVEX interface to ConvexOS that emulates the VAX/VMS Digital Command Language.) If the next nonblank characters are `@$`, CXbatch treats the line as an embedded option. Signify the end of the option with either the end of the line or unquoted `#` or `#!` characters. Indicate the end of the comment block with any character other than `#` or `#!` as the first character on a line.

The options embedded in the script file set the default characteristics for the batch request. Command line options override embedded values in the script file. For example, when you enter the command:

```
qsub -a 10:00am EDT testjob
```

and the option

```
-a "11:30pm EDT"
```

is embedded in the script file *testjob*, the request is run at 10:00 a.m. EDT because the value shown on the command line takes precedence over the embedded option.

Figure 21 illustrates use of embedded options in a script file.

**Figure 21** Embedded options in a script file

```
#
# Batch request script example:
#
# @$-a "11:30pm EDT" -lt "21:10, 20:00"
# # Run request after 11:30 EDT by default,
# # and set a maximum per-process CPU time
# # limit of 21 minutes and ten seconds.
# # Send a warning signal when any process
# # of the running batch request consumes
# # more than 20 minutes of CPU time.
# @$-lt 1:45:00
# # Set a maximum per-process CPU time limit
# # of one hour and 45 minutes. (The
# # implementation of CPU time limits is
# # completely dependent upon the ConvexOS
# # implementation at the execution
# # machine.)
# @$-mb -me
# Send mail at beginning and end of
# # request execution.
# @$-q batch1
# Queue request to queue: batch1 by
# # default.
# @$ # No more embedded flags.
```



Once a batch job is submitted to a queue, it becomes a queue request. Once in a queue, the owner can:

- Delete a request
- Place a request on hold
- Take a request off hold
- Move a request
- Change a request's priority

CXbatch operators and CXbatch managers can perform these same actions on any user's queue requests. This chapter describes how to perform these actions.

---

## Deleting queue requests

There are two commands to delete a batch request in a queue: the CXbatch `qdel` command and the `qmgr` utility `delete request` command. Each of these commands are described in the following sections.

---

### Using the `qdel` command

The `qdel` command is a CXbatch command that is executed from a shell command line. The format is

```
qdel [option] request_id [@hostname] [request_id [@hostname]...]
```

where

option can be one of the following:

- `-u username` Allows you to delete requests other than your own, where *username* is the name of the user who owns the request.  
By default, only the user who submitted the request can delete it from a queue. This option allows the superuser, CXbatch manager, or CXbatch operator to delete someone else's request from a queue.
- `-k` Sends a SIGKILL (-9) signal to a request. The request then exits and is deleted. If a request is running, all processes of the request are sent a SIGKILL signal.
- `sig` Sends the specified signal to an executing request where *sig* can either be the signal number or the signal name found in the `/usr/include/signal.h` file.

*request\_id* is the number assigned to the request when it is submitted to CXbatch. This number can be for any request, whether or not it is executing. You can specify more than one request.

If you are using the `-u` option, the *request\_id* must belong to the user defined in *username* of that option.

By default, when a batch request is successfully submitted, CXbatch displays the identification number of the request on the terminal of the user submitting the request. You can also use the `qstat` command to find the *request\_id*. Refer to Chapter 2, "Displaying information" for details on using `qstat`.

*hostname* is the name of the machine where the queue containing the request resides. If *hostname* is omitted, the local host is assumed.

For example, a user with batch operator privileges issues the following command to delete the request identified as 291 on the local machine submitted by user *smith*.

```
qdel -u smith 291
```

---

## Using the delete request command

The `delete request` command is a `qmgr` utility command. You must start `qmgr` before you can use this command. To start `qmgr`, enter

```
qmgr
```

The format for the `delete request` command is

```
delete request request_id [request_id ...]
```

where *request\_id* is the number assigned to the request when it is submitted to CXbatch. This number can be for any request, whether or not it is executing. You can specify more than one request on the command line. If a request is running, all processes of the request are sent a SIGKILL signal. Deleted requests are removed from the queue and discarded.

By default, when a batch request is successfully submitted, CXbatch displays the identification number of the request on the terminal of the user submitting the request. You can also use the `qstat` command to find the *request\_id*. Refer to Chapter 2, “Displaying information,” for details on using `qstat`.

---

## Preventing queue requests from executing

It may some times be desirable to prevent a specific queue request from executing for a period of time after it has been submitted to a queue and then later allowing the queue request to execute. The `hold request` and `release request` commands allow you to do this.

The `hold request` and `release request` commands are `qmgr` utility commands. You must start `qmgr` before you can use these command. To start `qmgr`, enter

```
qmgr
```

The format for each of these requests is described in the following sections.

---

### Placing a queue request on hold

You can use the `hold request` command to place a queue request on hold preventing it from executing. The format is

```
hold request request_id [request_id ...]
```

where *request\_id* is the number assigned to the request when it is submitted to CXbatch. You can specify more than one request on the command line. The request must be in the queued state to place it on hold.

By default, when a batch request is successfully submitted, CXbatch displays the identification number of the request on the terminal of the user submitting the request. You can also use the `qstat` command to find the *request\_id*. Refer to Chapter 2, "Displaying information," for details on using `qstat`.

---

### Removing the hold on a queue request

You can use the `release request` command to remove the hold on a queue request, making it eligible for execution. The format is

```
release request request_id [request_id ...]
```

where *request\_id* is the number assigned to the request when it is submitted to CXbatch. You can specify more than one request on the command line. The request must be in the holding state in order to be released.

By default, when a batch request is successfully submitted, CXbatch displays the identification number of the request on the terminal of the user submitting the request. You can also use the `qstat` command to find the *request\_id*. Refer to Chapter 2, "Displaying information," for details on using `qstat`.

---

---

## Moving a request to another queue

You can use the `move my_request` command to move a queue request from one queue to another. The request cannot be running. The `move my_request` command is a `qmgr` utility command. You must start `qmgr` before you can use this command. To start `qmgr`, enter

```
qmgr
```

The format for the `move my_request` command is

```
move my_request request_id [request_id ...] queue
```

where

*request\_id* is the number assigned to the request when it is submitted to CXbatch. You can specify more than one request on the command line. If any queue limits, access restrictions, or attributes prevent the request from being placed in the receiving queue, the queue request is not moved.

By default, when a batch request is successfully submitted, CXbatch displays the identification number of the request on the terminal of the user submitting the request. You can also use the `qstat` command to find the *request\_id*. Refer to Chapter 2, "Displaying information," for details on using `qstat`.

*queue* is the name of the queue where you wish the queue request to be moved.

---

## Changing the queue request priority

You can use the `modify request` command to change the priority of a queue request. The request cannot be running. The `modify request` is a `qmgr` utility command. You must start `qmgr` before you can use this command. To start `qmgr`, enter

```
qmgr
```

The format for the `modify request` command is

```
modify request priority=value request_id [request_id ...]
```

where

`priority = value` specifies the new priority of the queue request, where *value* is a number between 0 and 63; 0 is the lowest priority and 63 is the highest. A user can only decrease the priority of a request. A CXbatch manager or operator can raise a request's priority.

*request\_id* is the number assigned to the request when it is submitted to CXbatch. You can specify more than one request on the command line.

By default, when a batch request is successfully submitted, CXbatch displays the identification number of the request on the terminal of the user submitting the request. You can also use the `qstat` command to find the *request\_id*. Refer to Chapter 2, "Displaying information," for details on using `qstat`.

Checkpoint Restart is a standard feature of ConvexOS useful for application programs that must run for long lengths of time and that—once halted—cannot be started from the beginning without wasting significant time and resources. Such applications can be saved to files (*checkpointed*) either by operator intervention or by a periodically-executed script that includes Checkpoint Restart commands. If the application is then halted—either by a system crash, by a scheduled shutdown, or by an operator—it can be restarted as it was when it was last checkpointed.

You can checkpoint and restart processes, or groups of processes, running under CXbatch using Checkpoint Restart features built into CXbatch. These are:

- `qsub`—Allows you to specify checkpoint capabilities when you submit the request to CXbatch.
- `qchkpnt`—Allows you to checkpoint any checkpointable request currently running.
- `qrestart`—Allows you to restart any request that meets restart requirements.

This chapter describes how to use these commands. For information on other Checkpoint Restart commands, refer to the *ConvexOS Extensions User's Guide*.

---

## Controlling checkpointing when request is submitted

When a queue is created, the system manager specifies whether or not the queue is checkpointable. If the queue is specified as checkpointable, all jobs that are running are checkpointed if the system shuts down for any reason. If the queue is specified as not checkpointable, all jobs that are running are killed if the system shuts down for any reason. You can use the `qstat` command to find out whether or not the queue you wish to submit a job to is checkpointable or not. Refer to Chapter 2, "Displaying information," for details on using the `qstat` command.

You can override the checkpoint specification on a per job basis using options available with the `qsub` command. These options and their descriptions are listed in Table 7.

**Table 7** `qsub` options related to Checkpoint Restart

Option	Effect
<code>-c</code>	Request is checkpointable
<code>-nc</code>	Request is not checkpointable
<code>-cp <i>period</i></code>	Checkpoints request every <i>period</i> where <i>period</i> is the interval of time that must elapse between checkpoints
<code>-nr</code>	Request is not restartable

Each of these options are described in the following sections.

---

### Overriding not-checkpointable default

You can use the `-c` option when you submit a job to CXbatch to checkpoint the request if it is running when CXbatch shuts down. You only need to use this option if the queue you are using is specified as `checkpoint=available` and you want the job automatically checkpointed if CXbatch shuts down or if you want to manually checkpoint it while it is running.

For example, the following command checkpoints the request named `my_job` if it is running when CXbatch shuts down:

```
qsub -c my_job
```

---

## Overriding checkpointable default

You can use the `-nc` option when you submit a job to CXbatch to prevent a request from being checkpointed if it is running when CXbatch shuts down. You only need to use this option if the queue you are using is specified as `checkpoint=yes` and you do not want the job automatically checkpointed when CXbatch shuts down.

For example, the following command does not checkpoint the request named `my_job` if it is running when CXbatch shuts down.

```
qsub -nc my_job
```

## Periodically checkpointing a request

You can use the `-cp` option to periodically checkpoint a request. The format is

```
qsub -cp [number]unit
```

where

*number* is a positive integer specifying the number of *units* that must pass before checkpointing. If you do not specify a number, `qsub` assumes 1 unit.

*unit* can be *minutes*, *hours*, *days*, or *weeks*.

For example, the following command checkpoints the request named `my_job` every 20 minutes.

```
qsub -cp 20 minutes my_job
```

For more information, please refer to the `qsub(1)` man page.

## Preventing a checkpointed request from being restarted

You can use the `-nr` option to prevent a checkpointed request from being automatically restarted when CXbatch is restarted after it has been shut down. For example, the following command prevents the request `my_job` from being restarted if CXbatch shuts down.

```
qsub -nr my_job
```

---

## Two methods of checkpointing after request is submitted

There are two methods to checkpoint a request after it is running: using the `qchkpnt` command (a CXbatch command) and using the `chkpnt request` command (a CXbatch qmgr utility command). Each command is described in the following sections.

CXbatch uses `nqsdemon` running as root to checkpoint requests. Checkpoint files are owned by the owner of the request. These conditions must be met before you can checkpoint a request:

- User requesting the checkpoint must be the owner of the request or must have CXbatch manager or operator privileges.
- Request must be running in a batch queue.
- Request must be checkpointable. This means that the queue is set to `checkpoint=available` and the job was submitted with the `-c` option to `qsub`, or `checkpoint=yes` and the job was not submitted with the `-nc` option to `qsub`. Refer to the “Controlling checkpointing when request is submitted” section in this chapter for more details.

---

## Using the `qchkpnt` command

The `qchkpnt` command is a CXbatch command and can be executed from the shell command line. The format is

```
qchkpnt [option] request_id [request_id ...]
```

where

*option* can be one of the following:

`-e number unit`

specifies to checkpoint the request every *number unit*, where *number* is a positive number specifying the number of *units* that must pass before checkpointing. *unit* can be *minutes*, *hours*, *days*, or *weeks*. Checkpointing begins when the request begins running and ends when the request is terminated.

If you do not specify the `-e` option, the request is checkpointed immediately. You can cancel checkpointing by specifying 0 for *number*.

-f forces a request to be checkpointed, even if conditions exist that inhibit checkpointing.

*request\_id* is the number assigned to the request in the queue. You can specify more than one request on the command line.

By default, when a batch request is successfully submitted, CXbatch displays the identification number of the request on the terminal of the user submitting the request. You can also use the `qstat` command to find the *request\_id*. Refer to Chapter 2, "Displaying information," for details on using `qstat`.

For example, the following command checkpoints the request identified as 162 every hour:

```
qchkpnt -e 1 hour 162
```

When a request is checkpointed, an exit status describing results of the checkpoint is returned to the user's terminal submitting the checkpoint request:

- If no errors are encountered, the exit status is zero.
- If one or more of the requests are not checkpointed, the exit status is the number of requests that did not get checkpointed.
- If a fatal error occurs and none of the requests are checkpointed, the exit status is one of the codes listed in Table 8.

**Table 8** Exit statuses

Exit status	Reason
EX_USAGE	Syntax error.
EX_OSFILE	Batch file system does not exist, cannot be opened, or has an error.
EX_TEMPFAIL	Temporary failure; retry the command at a later time.
EX_NOPERM	User does not have permission to use command.

Refer to the `qchkpnt(1)` man page for more information.

---

## Using the `chkpnt` request command

The `chkpnt` command is a `qmgr` utility command. You must start `qmgr` before you can use this command. To start `qmgr`, enter

```
qmgr
```

The format for the `chkpnt` command is

```
chkpnt request request_id [ request_id ...]
```

where *request\_id* is the number assigned to the request when it is submitted to CXbatch. You can specify more than one request on the command line.

By default, when a batch request is successfully submitted, CXbatch displays the identification number of the request on the terminal of the user submitting the request. You can also use the `qstat` command to find the *request\_id*. Refer to Chapter 2, "Displaying information," for details on using `qstat`.

---

## Suspending executing requests

It may some times be necessary to suspend a queue request that is currently executing and then later restarting the queue request. The `suspend request` and `resume request` commands allow you to do this.

The `suspend request` and `resume request` commands are `qmgr` utility commands. You must start `qmgr` before you can use this command. To start `qmgr`, enter

```
qmgr
```

The format for each of these requests are described in the following sections.

---

### Suspending an executing request

You can use the `suspend request` command to temporarily “freeze” execution of a queue request. The request is checkpointed and then execution is terminated. However, the queue request remains in the queue.

Only checkpointable requests can be suspended. If a request fails to checkpoint, the request continues to execute. The format is

```
suspend request request_id [request_id ...]
```

where `request_id` is the number assigned to the request when it is submitted to CXbatch. You can specify more than one request on the command line.

By default, when a batch request is successfully submitted, CXbatch displays the identification number of the request on the terminal of the user submitting the request. You can also use the `qstat` command to find the `request_id`. Refer to Chapter 2, “Displaying information,” for details on using `qstat`.

---

### Resuming a suspended request

You can use the `resume request` command to resume execution of a suspended request. Resumed requests start in the queued state. Once the resumed request is about to enter the running state, it is restarted from its checkpointed state instead of being run in its entirety. The format is

```
resume request request_id [request_id ...]
```

where `request_id` is the number assigned to the request when it is submitted to CXbatch. You can specify more than one request on the command line.

By default, when a batch request is successfully submitted, CXbatch displays the identification number of the request on the

terminal of the user submitting the request. You can also use the `qstat` command to find the `request_id`. Refer to Chapter 2, “Displaying information,” for details on using `qstat`.

---

## Restarting checkpointed requests

CXbatch automatically restarts checkpointed requests when CXbatch is restarted. A request that fails to restart remains in the checkpointed state. If a request fails to restart, you can use the `qrestart` command to manually resubmit it for execution. A restarted request has the original privileges of the request’s owner.

The following conditions must be met before you can restart a request:

- The invoking user must be the owner of the request or must have CXbatch manager or operator privileges.
- The request must have been properly checkpointed.

The format for this command is

```
qrestart [-f] request_id [request_id ...]
```

where

`-f` forces the restart. Since you do not generally need to manually restart a checkpointed request unless it has failed to restart automatically, you will usually need to use this option.

`request_id` is the number assigned to the request in the queue. You can restart several requests with the same command by listing multiple `request_ids`.

By default, when a batch request is successfully submitted, CXbatch displays the identification number of the request on the terminal of the user submitting the request. You can also use the `qstat` command to find the `request_id`. Refer to Chapter 2, “Displaying information,” for details on using `qstat`.

Sometimes requests cannot be restarted. These types of requests fall into two categories:

- The request fails to meet requirements for checkpointability listed at the beginning of this section.
- One of the request’s processes is holding a nonrestartable request.

Refer to the `qrestart(1)` man page for more information.

---

# Transaction completion messages

# A

This appendix lists common codes that may be returned by any part of the CXbatch system, discusses the meaning of the codes, and describes actions you should take in response to the codes. These codes include error and success codes.

The format for an error message is

**TCM***machine\_code, message\_text*

where

*machine* indicates the type of machine on which the request was executing when the error message was issued. *machine* is either L for local or P for peer (remote) machine.

*code* is the mnemonic code for the error message.

*message\_text* is the text explaining the reason for the error.

The error messages in this appendix are listed alphabetically by the first letter in *code*.

The codes and message text in this appendix are followed by explanatory text detailing what caused the error, what happened to the request, and, in some cases, what action you should take.

Some explanatory text may indicate that "Quota information bits are present." This message indicates that the error involved a violation of a quota limit, and another error message is displayed with additional information about the limit violation. These additional messages are listed at the end of this appendix.

**TCML\_ACCESSDEN****Access denied at local host.**

The transaction cannot be performed because a queue denies access to a user invoking the `qsub` command.

A request was submitted directly to a queue that can only receive requests from other pipe queues.

A user submitted a request to a queue that does not have the user or the user's group in its access list and does not have unrestricted access.

The transaction failed and should not be retried.

**TCMP\_ACCESSDEN****Access denied at transaction peer.**

The transaction cannot be performed because a queue denies access to the owner of a request that is being queued remotely.

Tried to submit to `queue@host` where `nqsdemon@host` was started with the `-r` flag.

A request was submitted directly to `queue@host` (`qsub -q queue@host` remotely) that can only receive requests from other pipe queues.

A user submitted a request to a queue that does not have the user or the user's group in its access list and does not have unrestricted access.

The transaction failed and should not be retried.

**TCML\_ALREADACC****Already has access at local host.**

The transaction specified the addition of a user or group ID to a queue access set, and the user or group ID was already present in the queue access set.

The transaction failed.

**TCML\_ALREADEXI****Already exists at local host.**

The transaction specified the addition of a set element, and the element was already present in the set.

The transaction failed.

**TCMP\_BADCDTFIL**                      **Corrupt request control/data file(s) detected at transaction peer.  
Seek staff support.**

The transaction involved an operation on request control file(s) or data file(s), and the request control or data file(s) were found to be corrupt at the peer machine.

The transaction failed and should not be retried.

**TCMP\_CLIMIDUNKN**                      **Client machine id unknown at transaction peer.  
Seek staff support.**

The transaction cannot be performed because no machine ID can be determined on the remote machine for the client's network address.

A user submitted a request from *hostA* to *hostB* while *qmapmgr* on *hostB* did not have an *mid* entry for *hostA*.

**TCML\_COMPLETE**                      **Transaction complete at local host.**

The transaction concluded successfully.

This is a generic transaction completion code and can be used as the final result code from the local machine for any transaction when a more specific code does not exist.

The transaction succeeded.

**TCMP\_COMPLETE**                      **Transaction complete at transaction peer.**

The transaction concluded successfully.

This is a generic transaction completion code and can be used as the final result code from the local machine for any transaction when a more specific code does not exist.

The transaction succeeded.

**TCMP\_CONNBROKEN**                      **Network connection lost.  
Retry later.**

The transaction failed because the network connection established for the transaction was severed and the machine went down.

The transaction failed.

The transaction should be retried at a later time.

- TCMP\_CONNTIMOUT**                      **Network connection timeout.  
Retry later.**
- The transaction failed because of a timeout waiting for the transaction peer to respond, causing the machine to go down.
- The transaction failed.
- The transaction should be retried at a later time.
- 
- TCMP\_CONTINUE**                      **Subtransaction complete at transaction peer.**
- The transaction began successfully (i.e., access was granted, etc.), or the previous subtransaction of the transaction completed successfully. The server (in the client/server) is ready to perform the next subtransaction, as directed by the client.
- The transaction continues.
- 
- TCML\_EACCESS**                      **File access denied at local host.**
- The transaction involved a file operation that could not be performed because the protections set on the local machine denied access to the associated user.
- The transaction failed and should not be retried.
- 
- TCMP\_EACCESS**                      **File access denied at transaction peer.**
- The transaction involved a file operation that could not be performed because the protections set on the peer machine denied access to the associated user.
- The transaction failed and should not be retried.
- 
- TCML\_EFBIG**                      **File size limit exceeded at local host.**
- The transaction involved a file operation that would have increased the size of the file beyond the maximum supported at the local machine.
- The transaction failed and should not be retried.
- 
- TCMP\_EFBIG**                      **File size limit exceeded at transaction peer.**
- The transaction involved a file operation that would have increased the size of the file beyond the maximum supported at the peer machine.
- The transaction failed and should not be retried.

**TCML\_EISDIR**

**Operation on directory is prohibited at local host.**

The transaction involved an illegal operation on a directory based on the protections set on the local machine.

The transaction failed and should not be retried.

**TCMP\_EISDIR**

**Operation on directory is prohibited at transaction peer.**

The transaction involved an illegal operation on a directory based on the protections set on the peer machine.

The transaction failed and should not be retried.

**TCML\_ELOOP**

**Symbolic link translation limit exceeded at local host.**

The transaction required the translation of symbolic links on the local machine, and the symbolic link translation limit was exceeded on the local machine.

The transaction failed and should not be retried.

**TCMP\_ELOOP**

**Symbolic link translation limit exceeded at transaction peer.**

The transaction required the translation of symbolic links on the peer machine, and the symbolic link translation limit was exceeded on the peer machine.

The transaction failed and should not be retried.

**TCML\_ENFILE**

**Insufficient file descriptors at local host.  
Retry later.**

The transaction cannot be attempted or completed because of a file descriptor shortage on the local machine.

The transaction failed.

The transaction should be retried at a later time.

<b>TCMP_ENFILE</b>	<b>Insufficient file descriptors at transaction peer. Retry later.</b>
	The transaction cannot be attempted or completed because of a file descriptor shortage on the peer machine.
	The transaction failed. The transaction should be retried at a later time.
<b>TCML_ENOBUFS</b>	<b>Insufficient buffer space at local host. Retry later.</b>
	The transaction cannot be attempted because there are not enough buffers available on the local machine to establish the network connection required for the transaction at this time.
	The transaction failed. The transaction should be retried at a later time.
<b>TCMP_ENOBUFS</b>	<b>Insufficient buffer space at transaction peer. Retry later.</b>
	The transaction cannot be attempted because there are not enough buffers available on the peer machine to establish the network connection required for the transaction at this time.
	The transaction failed. The transaction should be retried at a later time.
<b>TCML_ENOENT</b>	<b>Component in file path does not exist at local host.</b>
	The transaction involved a file operation in which some element of the file path name did not exist at the local machine.
	The transaction failed and should not be retried.
<b>TCMP_ENOENT</b>	<b>Component in file path does not exist at transaction peer.</b>
	The transaction involved a file operation in which some element of the file path name did not exist at the peer machine.
	The transaction failed and should not be retried.

**TCML\_ENOMEM**

**Insufficient memory at local host.  
Retry later.**

The transaction cannot be performed because there is insufficient memory or swap space at the local machine.

The transaction failed. The transaction should be retried at a later time.

**TCMP\_ENOMEM**

**Insufficient memory at transaction peer.  
Retry later.**

The transaction cannot be performed because there is insufficient memory or swap space at the peer machine.

The transaction failed. The transaction should be retried at a later time.

**TCML\_ENOSPC**

**File system resource shortage at local host.  
Retry later.**

The transaction cannot be completed or performed because of a file system resource shortage on the local machine.

The transaction failed. The transaction should be retried at a later time.

**TCMP\_ENOSPC**

**File system resource shortage at transaction peer.  
Retry later.**

The transaction cannot be completed or performed because of a file system resource shortage on the peer machine.

The transaction failed. The transaction should be retried at a later time.

**TCML\_ENOTDIR**

**Invalid file path at local host.**

The transaction involved a file operation in which an element in the file path name was not a directory at the local machine.

The `qmgr Set CHECKpoint_directory directory-name` command specified a path that was not a directory.

The transaction failed and should not be retried.

**TCMP\_ENOTDIR****Invalid file path at transaction peer.**

The transaction involved a file operation in which an element in the file path name was not a directory at the peer machine.

The transaction failed and should not be retried.

**TCML\_EPERM****Permission denied at local host.**

The transaction involved a operation that could not be performed because the protections set on the local machine denied permission to the associated user.

The transaction failed and should not be retried.

**TCMP\_EPERM****Permission denied at transaction peer.**

The transaction involved a operation that could not be performed because the protections set on the peer machine denied permission to the associated user.

The transaction failed and should not be retried.

**TCML\_EPIPE****Fifo error at local host.****Retry later.**

A transaction has attempted to write on a pipe or socket which has no process attached to read the data. This usually happens because the read process exited prematurely or was killed.

The transaction failed. The transaction should be retried at a later time.

**TCMP\_EPIPE****Fifo error at transaction peer.****Retry later.**

A transaction has attempted to write on a pipe or socket which has no process attached to read the data. This usually happens because the read process exited prematurely or was killed.

The transaction failed. The transaction should be retried at a later time.

**TCML\_EROFS****Attempt to modify a read-only file system at local host.**

The transaction involved a file operation that would have altered a portion of a read-only file system at the local machine.

The transaction failed and should not be retried.

**TCMP\_EROF**

**Attempt to modify a read-only file system at transaction peer.**

The transaction involved a file operation that would have altered a portion of a read-only file system at the peer machine.

The transaction failed and should not be retried.

**TCML\_ERRORRETRY**

**Recoverable transaction failure at local host.  
Retry later.**

The transaction cannot be concluded successfully because of an error condition at the local machine that may not occur in the future.

This is a generic transaction completion code and can be used as the final result code from the local machine for any transaction when a more specific code does not exist. Explanatory text is displayed with this message.

The transaction failed. The transaction should be retried at a later time.

**TCMP\_ERRORRETRY**

**Recoverable transaction failure at transaction peer.  
Retry later.**

The transaction cannot be concluded successfully because of an error condition at the peer machine that might go away in the future.

This is a generic transaction completion code and can be used as the final result code from the local machine for any transaction when a more specific code does not exist. Explanatory text is displayed with this message.

The transaction failed. The transaction should be retried at a later time.

**TCML\_ETIMEDOUT**

**Connect(2) timeout at local host.  
Retry later.**

The transaction cannot be completed. A request for connection on the local machine failed because the connected party did not properly respond after a period of time.

The transaction failed. The transaction should be retried at a later time.

**TCML\_ETXTBSY**

**Text file operation denied at local host.  
Retry later.**

The transaction cannot be completed because the involved file operation dealt with a busy text file. The operation happened in circumstances under which such action is prohibited at the local machine.

The transaction failed. The transaction should be retried at a later time.

**TCMP\_ETXTBSY**

**Text file operation denied at transaction peer.  
Retry later.**

The transaction cannot be completed because the involved file operation dealt with a busy text file in circumstances under which such action is prohibited at the peer machine.

The transaction failed. The transaction should be retried at a later time.

**TCML\_FATALABORT**

**Non-recoverable transaction failure at local host.**

The transaction cannot be concluded successfully because of an error condition that will not go away in the foreseeable future at the local machine.

This is a generic transaction completion code and can be used as the final result code from the local machine for any transaction when a more specific code does not exist. Explanatory text is displayed with this message.

The transaction failed and should not be retried.

**TCMP\_FATALABORT**

**Non-recoverable transaction failure at transaction peer.**

The transaction cannot be concluded successfully because of an error condition that will not go away in the foreseeable future at the peer machine.

This is a generic transaction completion code and can be used as the final result code from the local machine for any transaction when a more specific code does not exist. Explanatory text is displayed with this message.

The transaction failed and should not be retried.

**TCML\_FIXBYNQS**

**Limit set to enforceable value at local host other than originally requested.**

The transaction specified the enforcement of a quota limit value for a batch queue that cannot be enforced by the local machine. The quota limit has instead been adjusted to fall within the limit enforceable at the local machine.

The transaction succeeds.

**TCML\_GRANFATHER**

**Transaction successful at local host but one or more requests were given a grandfather clause.**

The transaction modified a batch queue quota limit to be less than the limit set for one or more previously queued batch request limits.

The transaction succeeds.

**TCML\_INSHUTDOWN**

**CXbatch is shutting down at local host.**

The transaction was to test the internal state of CXbatch before continuing, and the local CXbatch daemon on the local machine is in the process of shutting down.

**TCML\_INSQUESPA**

**Insufficient space to queue request at local host.  
Retry later.**

The transaction failed because of insufficient queue space on the local machine. The transaction failed.

The local nqsdaemon either could not allocate memory for a new request or could not allocate a new transaction id.

The transaction should be retried at a later time.

**TCMP\_INSQUESPA**

**Insufficient space to queue request at transaction peer.  
Retry later.**

The transaction failed because of insufficient queue space on the peer machine.

The transaction failed. The transaction should be retried at a later time.

**TCML\_INSUFFMEM****Insufficient memory at local host.**

The transaction cannot complete because there is insufficient memory on the local machine.

The transaction failed.

**TCML\_INSUFFPRV****Insufficient privilege at local host.**

The transaction cannot be performed because CXbatch does not grant the associated user the appropriate privileges.

A `qmgr` command that required operator or manager privileges was attempted by a user without required privileges.

A user who was not a manager attempted to raise the priority of one of his/her requests.

The transaction failed.

**TCML\_INTERNERR****Internal CXbatch error at local host.****Seek staff support.**

The transaction could not be completed because of an internal error at the local machine.

The transaction failed and should not be retried.

**TCMP\_INTERNERR****Internal CXbatch error at transaction peer.****Seek staff support.**

The transaction could not be completed because of an internal error at the peer machine.

The transaction failed and should not be retried.

**TCML\_LOGFILERR****Cannot open or create new logfile at local host.**

The transaction designated the creation or opening of a new CXbatch log file, and the create or open operation failed.

The transaction failed.

**TCMP\_MAXNETCONN**      **Maximum network connection limit reached at transaction peer.  
Retry later.**

The transaction cannot begin because there are too many ongoing network transactions at the peer machine.

The transaction failed. The transaction should be retried at a later time.

**TCML\_MAXQDESTS**      **Maximum destination set cardinality reached at local host.**

The transaction specified the addition of a new queue destination for a pipe queue that would otherwise exceed the maximum number of queues allowed within a destination set for a single pipe queue. The maximum number of queue destinations in a destination set is 100.

The transaction failed.

**TCMP\_MIDCONFLCT**      **Machine id conflict between client and peer at transaction peer.  
Seek staff support.**

The transaction cannot begin because the respective machine ID values of the peer (server) machine and the local (client) machine do not match.

The transaction failed. The remote site should be marked as failed, and the transaction should not be retried.

**TCML\_NETDBERR**      **Local network database error at local host.  
Seek staff support.**

This completion code is returned when an error or an inconsistency is detected while a local client process accesses the network database on the client machine.

The transaction failed and should not be retried.

**TCMP\_NETDBERR**      **Local network database error at transaction peer.  
Seek staff support.**

This completion code is returned when an error or an inconsistency is detected while a server process accesses the network database on the server peer machine.

The transaction failed and should not be retried.

**TCML\_NETNOTSUPP**

**Networking not supported by implementation at local host.**

The transaction cannot begin because the local host implementation of CXbatch does not support the networking implementation required to perform the transaction.

The transaction failed and should not be retried.

**TCMP\_NETPASSWORD**

**Network password verification failure at transaction peer. Seek staff support.**

The transaction cannot begin because the client has failed to supply the proper CXbatch network server password to the remote server peer machine.

The transaction failed and should not be retried.

**TCML\_NOACCAUTH**

**No account authorization at local host.**

The transaction cannot be performed because there is no account database authorization at the local machine for the user associated with the transaction.

The transaction failed and should not be retried.

**TCMP\_NOACCAUTH**

**No account authorization at transaction peer.**

The transaction cannot be performed because there is no account database authorization at the peer machine for the user associated with the transaction.

The transaction failed and should not be retried.

**TCML\_NOACCNOW**

**Does not have access now at local host.**

The transaction specified the deletion of a user or group ID from a queue access set, and the specified user or group ID was not previously present in the set.

The transaction failed.

<b>TCML_NOESTABLSH</b>	<b>Unable to make connection with CXbatch daemon at local host. Retry later.</b>
	The transaction could not be performed because the connection between the client transaction process and the local CXbatch daemon at the local machine could not be established. The nqs daemons/net daemons are not running.
	The transaction failed. The transaction should be retried at a later time.
<b>TCMP_NOESTABLSH</b>	<b>Unable to make connection with CXbatch daemon at transaction peer. Retry later.</b>
	The transaction could not be performed because the connection between the transaction server and peer CXbatch daemon at the peer machine could not be established.
	The transaction failed. The transaction should be retried at a later time.
<b>TCML_NOLOCALDAE</b>	<b>CXbatch local daemon is not present at local host. Retry later.</b>
	The transaction cannot be performed because the CXbatch daemon at the local machine is not running.
	The transaction failed. The transaction should be retried at a later time.
<b>TCMP_NOLOCALDAE</b>	<b>CXbatch local daemon is not present at transaction peer. Retry later.</b>
	The transaction cannot be performed because the CXbatch daemon at the peer machine is not running.
	The transaction failed. The transaction should be retried at a later time.
<b>TCML_NOMOREPROC</b>	<b>Insufficient processes to perform transaction at local host. Retry later.</b>
	The transaction cannot begin because there are not enough processes available on the local host to perform the transaction.
	The transaction failed. The transaction should be retried at a later time.

<b>TCMP_NOMOREPROC</b>	<b>Insufficient processes to perform transaction at transaction peer. Retry later.</b>
	The transaction cannot begin because there are not enough processes available on the remote peer machine to perform the transaction.
	The transaction failed. The transaction should be retried at a later time.
<b>TCMP_NONETDAE</b>	<b>CXbatch net daemon is not present at transaction peer. Retry later.</b>
	The transaction cannot be performed because the CXbatch netdaemon at the peer machine is not running.
	The transaction failed. The transaction should be retried at a later time.
<b>TCMP_NONSECPORT</b>	<b>Non-secure network port verification error at transaction peer. Seek staff support.</b>
	The transaction cannot begin because the client process has connected to the remote CXbatch network daemon process using a nonsecure port.
	The transaction failed. The client server should be identified as failed, and the transaction should be retried after the client server program is repaired.
<b>TCML_NOPORTAVAI</b>	<b>No network port available at local host. Retry later.</b>
	The transaction cannot begin because no network port is available with which to perform the network operations required by the transaction at the local machine.
	The transaction failed. The transaction should be retried at a later time.

**TCML\_NOSUCHACC**            **No such account at local host.**

The transaction referred to a nonexistent account on the local machine, where the account must exist for the transaction to complete successfully.

An invalid account name was given to the `qmgr Add Manager` or `Add Users`, command. When using an account name, the account must currently exist in

`/etc/passwd` on the local machine. If the `[uid]` syntax is used, the account need not currently exist.

The transaction failed.

**TCML\_NOSUCHALI**            **No such queue alias at local host.**

The transaction involved an operation referring to a nonexistent queue alias at the local machine and therefore could not be performed.

The transaction failed and should not be retried.

**TCML\_NOSUCHDES**            **No such destination at local host.**

The transaction referred to a nonexistent queue destination for a pipe queue at the local machine, where the destination must exist for the transaction to complete successfully.

The transaction failed.

**TCML\_NOSUCHGRP**            **No such group at local host.**

The transaction referred to a nonexistent group at the local machine, where the group must exist for the transaction to complete successfully.

An invalid group name was given to the `qmgr Add Groups` command. When using a group name, the group must currently exist in the `/etc/group`. If the `[gid]` syntax is used, the `gid` need not currently exist.

The transaction failed.

**TCML\_NOSUCHMAC**            **No such machine.**

The transaction refers to a machine that is not known at the local host. The machine name must be known to the local host for the transaction to complete successfully.

The transaction failed.

- TCML\_NOSUCHMAN**      **No such manager/operator at local host.**
- The transaction refers to a nonexistent CXbatch manager or operator account. The account must have been previously defined as a CXbatch manager or operator for the request to complete successfully.
- An attempt was made to delete a CXbatch manager using the `qmgr DElete Manager` command while the account specified did not currently have manager privileges, but existed on the local machine.
- The transaction failed.
- TCML\_NOSUCHQUE**      **No such queue at local host.**
- The transaction involved an operation referring to a nonexistent queue at the local machine and could not be performed.
- The transaction failed and should not be retried.
- TCMP\_NOSUCHQUE**      **No such queue at transaction peer.**
- The transaction involved an operation referring to a nonexistent queue at the peer machine and could not be performed.
- The transaction failed and should not be retried.
- TCML\_NOSUCHQUO**      **No such quota supported at local host.**
- The transaction specified the setting of a batch queue quota limit that cannot be enforced at the local machine.
- The transaction failed.
- TCML\_NOSUCHREQ**      **No such request at local host.**
- The transaction could not be performed because it references a request that does not exist at the local machine.
- The transaction failed and should not be retried.
- TCMP\_NOSUCHREQ**      **No such request at transaction peer.**
- The transaction could not be performed because it references a request that does not exist at the peer machine.
- The transaction failed and should not be retried.

**TCML\_NOSUCHSIG****No such signal at local host.**

The transaction could not be performed because it referred to a signal that does not exist at the local machine.

The transaction failed and should not be retried.

**TCMP\_NOSUCHSIG****No such signal at transaction peer.**

The transaction could not be performed because it referred to a signal that does not exist at the peer machine.

The transaction failed and should not be retried.

**TCML\_NOTREQOWN****Not request owner at local host.**

The transaction could not be performed because it specified an operation on a request when the user associated with the transaction was not the request owner at the local machine.

A user without operator or manager privileges attempted to perform one of the following commands on a request not owned by this user:

```
qmgr MODify Request Priority
```

```
qmgr MOVE Request
```

```
qmgr CHkpnt Request
```

```
qmgr DElete Request
```

The transaction failed and should not be retried.

**TCMP\_NOTREQOWN****Not request owner at transaction peer.**

The transaction could not be performed because it specified an operation on a request when the user associated with the transaction was not the request owner at the peer machine.

The transaction failed and should not be retried.

**TCML\_PATHLEN****Resolved output filename for batch request at local host exceeds the maximum supported path length.**

The transaction could not be completed because a local transaction process reserved a queue slot on the local machine for a new batch request, and the resolved standard output or standard error path name of the request exceeds the maximum request path length supported by the CXbatch implementation at the local machine.

The transaction failed and should not be retried. The associated request must be deleted.

**TCMP\_PATHLEN**

**Resolved output filename for batch request at transaction peer exceeds the maximum supported path length.**

The transaction could not be completed because a remote transaction process reserved a queue slot on the remote machine for a new batch request, and the resolved standard output or standard error path name of the request exceeds the maximum request path length supported by the CXbatch implementation at the remote machine.

The transaction failed and should not be retried. The associated request must be deleted.

**TCML\_PEERARRIVE**

**Request arriving at local host.**

The transaction specified the deletion of a request that is in transit between two machines in the network, and the designated request is presently in the arriving state (to be delivered or is being delivered from a remote machine).

The transaction failed, and the coordinator process attempting the transaction must now retry the delete operation using the networked form of the transaction.

**TCML\_PEERDEPART**

**Request departing at local host.**

The transaction specified the deletion of a request that is in transit between two machines in the network, and the designated request is presently in the departing state (to be delivered or is being delivered to a remote machine).

The transaction failed, and the coordinator process attempting the transaction must now retry the delete operation using the networked form of the transaction.

**TCML\_PROTOFAIL**

**CXbatch protocol failure at local host.  
Seek staff support.**

The transaction could not be completed because of a CXbatch message protocol failure on the local machine.

The transaction failed and should not be retried.

**TCMP\_PROTOFAIL**                      **CXbatch protocol failure at transaction peer.  
Seek staff support.**

The transaction could not be completed because of a CXbatch message protocol failure between a local client transaction process and a remote server transaction process, or a protocol failure between the remote server process and the network daemon at the remote machine.

The transaction failed and should not be retried.

**TCML\_QUEDISABL**                      **Queue is disabled at local host.  
Retry later.**

The transaction involved an operation that could not be completed because the referenced queue was disabled at the local machine.

The transaction failed. The transaction should be retried at a later time.

**TCMP\_QUEDISABL**                      **Queue is disabled at transaction peer.  
Retry later.**

The transaction involved an operation that could not be completed because the referenced queue was disabled at the peer machine.

The transaction failed. The transaction should be retried at a later time.

**TCML\_QUEENABLE**                      **Queue is enabled at local host.**

The transaction specified the deletion of a local queue that is enabled. A queue must be disabled and empty before it can be deleted.

The transaction failed.

**TCML\_QUEHASREQ**                      **Queue has request(s) at local host.**

The transaction specified the deletion of a local queue that contains requests. A queue must be disabled and empty before it can be deleted.

The transaction failed.

**TCML\_QUOTALIMIT****Explicit request quota limits exceed maximums at local host.**

The transaction involved the queuing of a batch request or the modification of a batch request limit. The successful completion of the transaction would have caused the batch request to exceed one or more of the corresponding quota limits for that batch queue at the local machine.

Quota information bits are present.

The transaction failed and should not be retried.

**TCMP\_QUOTALIMIT****Explicit request quota limits exceed maximums at transaction peer.**

The transaction involved the queuing of a batch request or the modification of a batch request limit. The successful completion of the transaction would have caused the batch request to exceed one or more of the corresponding quota limits for that batch queue at the local machine.

Quota information bits are present.

The transaction failed and should not be retried.

**TCML\_REQCOLLIDE****Attempt to queue already existing request at local host.**

The transaction involved the queuing of a request that had the same request ID already existing at the local machine.

The transaction failed and should not be retried. The associated request must be deleted.

**TCMP\_REQCOLLIDE****Attempt to queue already existing request at transaction peer.**

The transaction involved the queuing of a request that had the same request ID already existing at the peer machine.

The transaction failed and should not be retried. The associated request must be deleted.

**TCML\_REQDELETE****Request deleted at local host.**

The transaction specified the deletion of a specific request at the local machine, and the request was deleted successfully.

The transaction succeeded.

- TCMP\_REQDELETE**            **Request deleted at transaction peer.**  
The transaction specified the deletion of a specific request at the peer machine, and the request was deleted successfully.  
The transaction succeeded.
- TCML\_REQMOVSTA**            **Request is not in an allowable state at local host.**  
The transaction involved moving a request from one queue to another. The state of the request was not one of the allowable states: queued, holding, or waiting.  
The transaction failed and should not be retried.
- TCML\_REQNOTHOLD**            **Request is not being held at local host.**  
The transaction involved releasing a held request, but the target request was not in the holding state.  
The transaction failed and should not be retried.
- TCML\_REQNOTQUE**            **Request is not in the queued state at local host.**  
The transaction involved putting a queued request on hold, but the target request was not in the queued state.  
The transaction failed and should not be retried.
- TCML\_REQOPHOLD**            **Request is being held by operator at local host.**  
The transaction involved releasing a held request, but the request was being held by an operator.  
The transaction failed and should not be retried.
- TCML\_REQRUNNING**            **Request is running at local host.**  
The transaction specified the deletion of a specific request at the local machine, but the request was already running, and the transaction did not specify a signal. The request is not deleted and continues execution.  
The transaction failed and should not be retried.

<b>TCMP_REQRUNNING</b>	<b>Request is running at transaction peer.</b>
	The transaction specified the deletion of a specific request at the peer machine, but the request was already running, and the transaction did not specify a signal. The request is not deleted and continues execution.
	The transaction failed and should not be retried.
<b>TCML_REQSIGNAL</b>	<b>Request was running and has been signalled at local host.</b>
	The transaction specified the deletion or signalling of a specific request at the local machine, and the request has been signalled successfully.
	The transaction succeeded.
<b>TCMP_REQSIGNAL</b>	<b>Request was running and has been signalled at transaction peer.</b>
	The transaction specified the deletion or signalling of a specific request at the peer machine, and the request has been signalled successfully.
	The transaction succeeded.
<b>TCML_ROOTINDEL</b>	<b>Root cannot be deleted at local host.</b>
	The transaction specified the deletion of root as a CXbatch manager at the local machine.
	The transaction failed.
<b>TCMP_RRFUNKNMID</b>	<b>Request refers to unknown machines at transaction peer.</b>
	The transaction involved the queuing of a request at a queue destination where the request referred to machine IDs not known to the transaction peer machine.
	The transaction failed (the request is not queued) and should not be retried.

**TCML\_SELMDUNKN**            **Local machine id unknown at local host.  
Seek staff support.**

The transaction could not be performed because the local client transaction process at the local machine was unable to determine the machine ID of the local machine.

The transaction failed and should not be retried. The containing queue should be stopped, and the request requeued, if the transaction is on behalf of a previously queued CXbatch request.

**TCMP\_SELMDUNKN**            **Local machine id unknown at transaction  
peer.  
Seek staff support.**

The transaction could not be performed because the local server transaction process at the peer machine was unable to determine the machine ID of the peer machine.

The transaction failed. The remote site should be marked as failed, and the transaction should not be retried.

**TCML\_SELREFDES**            **Attempt to create self-referential  
destination at local host.**

The transaction specified the addition of a pipe queue destination that would have caused a self-referential pipe queue destination set. This means that the destination set for a pipe queue contains the name of the pipe queue.

The transaction failed and should not be retried.

**TCML\_SHUTERROR**            **Shutdown error at local host.**

The transaction specified that the local CXbatch daemon at the local host be shut down, and an error occurred in the shutdown process.

The transaction failed.

**TCML\_SUBMITTED**            **Request successfully submitted at local  
host.**

The transaction involved queuing a request at the local machine, and the queuing transaction was successful.

Quota information bits are present.

The transaction succeeds.

**TCMP\_SUBMITTED****Request successfully submitted at transaction peer.**

The transaction involved queuing a request at the peer machine, and the queuing transaction was successful.

Quota information bits are present.

The transaction succeeds.

**TCML\_UNAFAILURE****Unanticipated transaction failure at local host.**

An unanticipated and totally unexpected error condition occurred when CXbatch tried to process the transaction at the local machine.

This is a generic transaction completion code and can be used as the code from the local machine for any transaction when a more specific code does not exist.

The transaction failed and should not be retried.

**TCMP\_UNAFAILURE****Unanticipated transaction failure at transaction peer.**

An unanticipated and totally unexpected error condition occurred when CXbatch tried to process the transaction at the peer machine.

This is a generic transaction completion code and can be used as the final result code from the local machine for any transaction when a more specific code does not exist.

The transaction failed and should not be retried.

**TCML\_UNRESTR****Access is currently unrestricted.**

This transaction specified that a queue access set be made unrestricted, and the queue access set for the named queue was already unrestricted.

The transaction failed.

**TCML\_WROQUETYP****Wrong queue type for transaction at local host.**

The transaction involved the queuing of a request, but the target queue on the local machine was of the wrong type for the request (e.g., submitting a batch request to a device queue and vice versa). This code is also seen when an operation that only relates to a pipe queue is attempted on a batch queue.

The transaction failed and should not be retried.

## TCMP\_WROQUETYP

### Wrong queue type for transaction at transaction peer.

The transaction involved the queuing of a request, but the target queue on the peer machine was of the wrong type for the request (e.g., submitting a batch request to a device queue by means of a pipe queue and vice versa). This code is also seen when an operation that only relates to a pipe queue is attempted on a batch queue.

The transaction failed and should not be retried.

The following are the quota limit violation information flags:

TCI_PP_CFLEXC	per-process core-file size limit
TCI_PP_CTLEXC	per-process CPU time limit
TCI_PP_DSLEXC	per-process data-segment limit
TCI_PP_MSLEXC	per-process memory size limit
TCI_PP_NELEXC	per-process nice execution priority limit
TCI_PP_PFLEXC	per-process permanent file size limit
TCI_PP_SSLEXC	per-process stack-segment limit
TCI_PP_TFLEXC	per-process temporary file size limit
TCI_PP_WSLEXC	per-process working set quota limit
TCI_PR_CTLEXC	per-request CPU time limit
TCI_PR_MSLEXC	per-request memory space limit
TCI_PR_NCPEXC	per-request CPU quota limit
TCI_PR_PFLEXC	per-request permanent file space limit
TCI_PR_TFLEXC	per-request temporary file space limit
TCI_NOIMPORT	import resource not available



---

# Request completion messages

# B

This appendix lists common codes that may be returned by a request and discusses the meaning and actions you should take in response to those codes. These codes include error and success codes.

The format for an error message is

*RCM\_code, message\_text*

where

*code* is the mnemonic code for the error message.

*message\_text* is the text explaining the reason for the error. Some error messages may also include action that CXbatch takes, such as "Request deleted."

The error messages in this appendix are listed alphabetically by the first letter in *code*.

The error codes and message text in this appendix are followed by explanatory text detailing what caused the error, what happened to the request, and, in some cases, what action you should take.

Some explanatory text may indicate that "Quota information bits are present." This message indicates that the error involved a violation of a quota limit, and another error message is displayed with additional information about the limit violation. These additional messages are listed at the end of this appendix.

**RCM\_2MANYENVARS**            **Too many environment variables to run batch request. Request not executed. Request deleted.**

Too many environment variables existed to execute the batch request.

No output files are returned.

The request is deleted.

**RCM\_2MANYSVARGS**            **Too many server arguments on server execve(). Request requeued.**

Too many server arguments in server execve().

No output files are returned.

The request is requeued, and the queue is stopped.

**RCM\_ABORTED**                **Request aborted via a signal. Request deleted.**

A shell process was terminated by a signal unrelated to CXbatch shutting down (unless something unpredictable happens, such as the request received SIGTERM signals and then killed itself with some signal besides SIGTERM or SIGKILL).

Output files (if any) are queued for return.

The request is deleted.

**RCM\_BADCDFIL**                **Corrupted request control and/or data files. Request not executed. Request files placed in CXbatch failed directory.**

An invalid request control file and/or data file was detected. No output files are returned. The request is placed in the failed directory on the local machine for later analysis.

**RCM\_BADSRVARG**                **Bad argument passed to request server. Request requeued.**

A bad argument or environment variable was given to the server.

No output files are returned.

The request is requeued, and the queue is stopped.

**RCM\_DELIVERED**

**Request has been successfully delivered to destination.**

The associated request, as previously routed by a pipe queue server, has been successfully delivered to its destination.

No output files are returned.

The request remains on the destination queue.

**RCM\_DELIVEREXP**

**Request delivery time expired.  
Request deleted.**

The associated request, as previously routed by a pipe queue server, has not been delivered to its destination. The original completion code from the server is RCM\_RETRYLATER, and the delivery expiration time on the request has been reached.

The original completion code is converted to this completion code when the expiration time has elapsed.

No output files are returned.

The request is deleted.

**RCM\_DELIVERFAI**

**Request could not be delivered.  
Request deleted.**

The associated request, as previously routed by a pipe queue server, could not be delivered to its destination. This completion code indicates a disastrous situation in which the request can never be delivered.

The information field of this return code must contain the TCM\_ code that caused the failure.

No output files are returned.

The request is deleted.

**RCM\_DELIVERRETX**

**Request was not delivered.  
Retry limit exceeded.  
Request deleted.**

The associated request, as previously routed by a pipe queue server, has not been delivered to its destination, and the retry limit for such a transaction has been reached.

The information field of this return code must contain the TCM\_ code that caused the failure.

No output files are returned.

The request is deleted.

<b>RCM_ENFILERUN</b>	<b>Unable to successfully start request because of a file descriptor shortage. Request requeued.</b>
	The request or transaction cannot be handled because of a file descriptor shortage on the local machine.
	No output files are returned.
	The request is requeued to run at a later time. All CXbatch queues on the local machine are stopped to conserve any remaining available file descriptors.
<b>RCM_ENOSPCRUN</b>	<b>Unable to successfully start request because of a file system resource shortage. Request requeued.</b>
	The request cannot be handled because of insufficient file system space on the local machine.
	No output files are returned.
	The request is requeued to run at a later time. All CXbatch queues on the local machine are stopped so as not to place additional file system space demands on the local machine.
<b>RCM_EXECUTING</b>	<b>Request executing.</b>
	This code is used in the internal message protocol between the server and shepherd processes when spawning a CXbatch batch request.
	This request completion code must never be returned.
<b>RCM_EXITED</b>	<b>Request exited normally.</b>
	The batch request server exited normally.
	Output files are queued for return.
	Local information concerning the request is deleted.
<b>RCM_IMPORTFAI</b>	<b>NFS import of directory failed. Request deleted.</b>
	<i>nqsimport</i> was unable to successfully perform the import.
	Output files are queued for return. The request is deleted.

**RCM\_INSUFFMEM**

**Insufficient memory to start request.  
Request queued.**

The server to handle the request could not be spawned because of insufficient memory or swap space.

No output files are returned.

The request is queued, and all queues are stopped.

**RCM\_INTERRUPTED**

**Server transaction processing aborted for  
CXbatch shutdown.  
Request queued.**

The network queue or pipe queue server stopped transaction processing because of a CXbatch shutdown.

No output files are returned.

The request is queued.

**RCM\_MIDUNKNOWN**

**Machine-id of request owner is no longer  
recognized by the execution machine.  
Request not executed.  
Request deleted.**

The local batch shell process, local network queue server, or local pipe server cannot identify the machine from which the request originated.

When the request was originally accepted by the local system, the owner machine ID (MID) was known to the local system. Since that time, however, the local host tables have changed, and the owner MID of the request is no longer recognized.

Output files are queued for return.

The request is deleted.

**RCM\_NETREQDEL**

**Request could not be successfully  
delivered.  
Previously routed request expired or was  
deleted at destination.  
Request deleted.**

The request, as previously routed by a pipe queue server, could not be delivered to its destination because the request expired or was deleted at the destination.

No output files are returned.

The request is deleted.

**RCM\_NOACCAUTH**                    **No account authorization on execution machine for mapped request owner user-id. Request not executed. Request deleted.**

The request cannot be handled because the local machine has no password entry for the request owner user ID.

Output files are queued for return.

This can occur if the local password database is changed, i.e., someone deletes an account from the password database after a request owned by the account has been queued.

The request is deleted.

**RCM\_NOMOREPROC**                **Insufficient processes available to spawn request. Request queued.**

The request cannot be executed because of a process shortage on the local machine.

No output files are returned.

The request is queued to run at a later time. All CXbatch queues on the local machine are stopped to conserve remaining available processes.

**RCM\_NONSECPORT**                **Server bind() error. Request queued.**

The pipe queue or network queue server connected to the transaction server on a nonsecure port. The queue server is flawed.

No output files are returned.

The request is queued, and the containing queue is stopped.

**RCM\_NORESTART**                **Interrupted request prohibits restart on CXbatch rebuild. Request deleted.**

The request was running at the time of a CXbatch shutdown or system crash, and the request was not defined as "restartable" by the request owner. This completion code is returned directly from the local CXbatch daemon and can only occur when CXbatch is being rebooted.

Output files are queued for return.

The request is deleted.

**RCM\_NOSVRETCODE**      **Server for request did not return a completion code. Request failed. Request files placed in CXbatch failed directory.**

The network queue server or pipe queue server failed to report a completion code when it exited.

No output files are returned.

The request is placed in the failed directory, and the appropriate device or queue is stopped.

**RCM\_PATHLEN**      **Resolved stdout or stderr pathname of batch request at destination exceeds the maximum supported length. Request deleted.**

The standard output or standard error path name of a batch request when resolved by the selected pipe queue destination exceeds the maximum request path length supported by the CXbatch implementation at the receiving machine.

No output files are returned.

The request is deleted.

**RCM\_PIPREQDEL**      **Request deleted.**

The routing of the request was interrupted by a delete request at the local machine.

No output files are returned.

The request is deleted.

**RCM\_REBUILDFAI**      **CXbatch rebuild failure. Request could not be requeued. Request deleted.**

A request that was not running at the time of an CXbatch shutdown or system crash could not be requeued because of an internal CXbatch error. This completion code is returned directly from the local CXbatch daemon and can only occur when CXbatch is being rebooted.

No output files are returned.

The request is placed in the CXbatch failed directory.

**RCM\_REQCOLLIDE**                      **Request collided with another previously existing request with the same request-id. The newer request has been deleted. Seek staff support.**

The request being routed collided with a previously existing request on a potential destination machine.

No output files are returned.

The request is deleted.

**RCM\_RETRYLATER**                      **Request transaction failed. Retry scheduled. Request queued.**

The transaction for this request or subrequest cannot be carried out successfully and will be retried later.

No output files are returned.

**RCM\_ROUTED**                              **Request successfully routed for delivery to destination.**

The request was successfully routed and queued at one of the remote destinations for the request, as selected by the pipe queue. Quota information bits are present.

No output files are returned.

**RCM\_ROUTEDLOC**                      **Request sent to local queue destination.**

The request was successfully routed and queued at one of the local destinations for the request, as selected by the pipe queue. Quota information bits are present.

No output files are returned.

**RCM\_ROUTEEXP**                      **Request routing time expired. Request deleted.**

The request has not been routed to a destination. The original completion code from the server is RCM\_RETRYLATER, and the routing expiration time on the request has been reached. The original completion code is converted to this completion code when the expiration time has elapsed.

No output files are returned.

The request is deleted.

**RCM\_ROUTEFAI**

**Request could not be routed.  
Request deleted.**

The request cannot be routed. No destination will accept it. Every destination selected for the request by the pipe queue server has been contacted, and each destination indicated that it will not accept the request. Information bits that define the set of failure conditions are present.

No output files are returned.

The request is deleted.

**RCM\_ROUTERETX**

**Request was not routed.  
Retry limit exceeded.  
Request deleted.**

The request was not routed to any destination, and the retry limit for this type of transaction has been reached. Failed pipe routing information bits that define the set of failure conditions are present.

No output files are returned.

The request is deleted.

**RCM\_SEREXEFAI**

**Request server execve() failed.  
Request requeued.**

Server execve() operation failed.

No output files are returned.

The request is requeued, and the queue is stopped.

**RCM\_SERVESIGERR**

**Server killed by unanticipated signal.  
Request requeued.**

Network or pipe server was killed by a signal that should not have killed it.

No output files are returned.

The request is requeued, and the device or queue is stopped as appropriate.

<b>RCM_SHEXEF2BIG</b>	<b>Too many environment variables to run request. Request deleted.</b>
	One or both of the argv[] or envp[] argument sets to the batch request shell exceeded the maximum size supported by the underlying UNIX implementation.
	Output files are queued for return.
	The request is deleted.
<b>RCM_SHUTDNABORT</b>	<b>Request aborted for CXbatch shutdown. The request was defined as unrestartable, and so the request has been deleted.</b>
	The server or shell process for the request exited because of a signal sent because of a CXbatch shutdown, and the request is not restartable.
	Output files are queued for return.
	The request is deleted.
<b>RCM_SHUTDNREQUE</b>	<b>Request aborted for CXbatch shutdown. The request has been requeued for later restart.</b>
	The server or shell process for the request exited because of a signal sent because of a CXbatch shutdown, and the request is restartable.
	When CXbatch is restarted, the request is requeued for continued execution.
<b>RCM_SSEXEF2BIG</b>	<b>Request shell execve() failed. Request requeued.</b>
	The attempt to execute the shell chosen by the system to interpret the batch request shell script failed.
	Output files are queued for return.
	The request is requeued, and the queue is stopped.
<b>RCM_STAGEOUT</b>	<b>Output file successfully returned to destination.</b>
	The output file was successfully returned to its intended destination.
	The output file is deleted.

**RCM\_STAGEOUTBAK**

**Output file could not be returned to primary destination.**

**Output file successfully returned to backup destination in user home directory on the execution machine.**

The output file could not be returned to its intended destination because of circumstances not allowing retry attempts. Instead, the file was successfully returned to its backup destination.

The information field of this return code must contain the TCM\_ code that caused the failure at the primary destination.

The output file is deleted.

The request completes.

**RCM\_STAGEOUTFAI**

**Output file could not be returned to primary or backup destination.**

The output file could not be returned to its intended or backup destination because of error conditions that prohibit retrying the operation.

The information field of this return code must contain the TCM\_ code that caused the failure at the primary destination.

The output file must be deleted.

**RCM\_UNABLETOEXE**

**Unable to execute request.  
Request deleted.**

The request could not be executed for reason(s) identified in the information bit vector of the completion code. Quota information bits are defined.

Output files are queued for return.

The request is deleted.

**RCM\_UNAFAILURE**

**Request failed.  
Request files placed in CXbatch failed  
directory.**

An unanticipated error condition occurred while CXbatch was trying execute the request or perform a transaction operation for the request.

If CXbatch was attempting to return an output file to the submitter, the associated output file is deleted, and the output file loss is recorded for the batch request.

All other operations result in the request being placed in the failed directory, and any associated output files are deleted.

**RCM\_USHBRKPNT**

**Unsupported shell breakpoint  
encountered.  
Request deleted.**

The shell chosen by the user to execute the batch request stopped at a breakpoint.

Output files are queued for return.

The request is deleted.

**RCM\_USHEXFAI**

**Request shell execve() failed.  
Request deleted.**

The attempt to execute the shell chosen by the system to interpret the batch request shell script failed.


Output files are queued for return.

The request is deleted.

The following are the request completion flags:

RCI_ACCESSDEN	Access denied
RCI_CLIMIDUNKN transaction peer	Client machine ID is unknown at transaction peer
RCI_EFBIG	File size limit exceeded
RCI_FATALABORT	Nonrecoverable transaction failure
RCI_MIDCONFLICT destination	Machine ID conflict between client and destination
RCI_NETNOTSUPP site.	Networking not supported at CXbatch site.
RCI_NETPASSWD	Network password verification error
RCI_NOSUCHQUE	No such queue
RCI_PEERINTERR peer	CXbatch internal error at transaction peer
RCI_PEERMIDUNKN transaction peer	Local machine ID is unknown at transaction peer
RCI_PEERNETDB peer	Network database error at transaction peer
RCI_PEERNOACATH peer	No account authorization at transaction peer
RCI_PROTOFAIL	CXbatch protocol failure
RCI_QUOTALIMIT maximums	Explicit request quota limits exceed maximums
RCI_RRFUNKNMID at transaction peer	Request refers to machine IDs unknown at transaction peer
RCI_WROQUETYP	Wrong queue type for request
RCI_UNAFAILURE	Unanticipated transaction failure





This appendix contains the CONVEX CXbatch man pages applicable to users presented in alphabetical order. The man pages included are:

- qchkpnt(1)
- qdel(1)
- qjlist(1)
- qlimit(1)
- qmgr(8)
- qps(1)
- qrestart(1)
- qstat(1)
- qsub(1)
- qwatch(8)

## NAME

*qchkpnt* – checkpoint CXbatch request(s).

## SYNOPSIS

*qchkpnt* [ **-e** *freq* ] [ **-f** ] *request-id* ...

## DESCRIPTION

*qchkpnt* checkpoints the requests whose *request-ids* are listed on the command line. The current state of all the processes comprising the request are saved in a set of checkpoint files. The checkpoint files are stored in the CXbatch *checkpoint directory*. A successfully checkpointed request will be restarted from its checkpointed state instead of being re-run.

Only running batch requests may be checkpointed. To checkpoint a request, the invoking user must be the owner of the request or have CXbatch operator privileges.

The options have the following meanings:

**-e** *freq* Normally the request is checkpointed immediately. When the **-e** option is present the request will be checkpointed periodically at intervals of *freq*.

The *freq* is specified as  $\langle [number] unit \rangle$ , where *number* is a positive integer and *unit* is [ *Hours* | *Days* | *Weeks* ]. A *number* of zero will cancel periodic checkpointing.

**-f** Force checkpoint even if one of the processes of the request hold non-checkpointable resources.

## DIAGNOSTICS

*qchkpnt* returns an exit status describing what it did. If there were no errors, the exit status is zero. If one or more of the requests weren't checkpointed, the exit status is the number of requests that weren't checkpointed. If a fatal error occurs and none of the requests are checkpointed (e.g., a syntax error), the exit status is one of the codes defined in  $\langle sysexits.h \rangle$ .

**EX\_USAGE** The command was used incorrectly, e.g., with the wrong number of arguments, a bad flag, a bad syntax in a parameter.

**EX\_OSFILE** Some batch system file does not exist, cannot be opened, or has an error.

**EX\_TEMPFAIL** Temporary failure; retry the command at a later time.

**EX\_NOPERM** You did not have sufficient permission to perform the operation.

**EX\_SOFTWARE** Too many request-ids were specified.

## SEE ALSO

*qsub*(1), *qrestart*(1), *chkpnt*(1), *restart*(1), *qmgr*(8)

## NOTES

CXbatch is an optional product; for more information, please contact your CONVEX sales representative.

## NAME

qdel - delete or signal CXbatch request(s).

## SYNOPSIS

qdel [ -k ] [ -signo ] [ -u username ] request-id[@host] ...

## DESCRIPTION

*qdel* deletes all queued CXbatch requests whose *request-ids* are listed on the command line. Additionally, if the flag **-k** is specified, the default signal of SIGKILL (9) is sent to any running request whose *request-id* is listed on the command line. This causes the receiving request to exit and be deleted. If the flag **-signo** is present, the specified signal is sent instead of the SIGKILL signal to any running request whose *request-id* is listed on the command line. *signo* can be either the signal number or the signal name. The signal names are as given in */usr/include/signal.h*, stripped of the common SIG prefix. In the absence of the **-k** and **-signo** flags, *qdel* will not delete a *running* CXbatch request.

To delete or signal a CXbatch request, the invoking user must be the owner (the submitter of the request) or the superuser. The only exception to this rule occurs when the invoking user has CXbatch operator privileges as defined in the CXbatch manager database. Under these conditions, the invoker may specify the **-u username** flag that allows the invoker to delete or signal requests owned by the user whose account name is *username*. When this form of the command is used, all *request-ids* listed on the command line are presumed to refer to requests owned by the specified user.

A CXbatch request is always uniquely identified by its *request-id*, no matter where it is in the network of the machines. A *request-id* is always of the form *seqno* or *seqno.hostname*, where *hostname* identifies the machine from whence the request was originally submitted, and *seqno* identifies the sequence number assigned to the request on the originating host. If the *hostname* portion of a *request-id* is omitted, the local host is always assumed. The local host is searched for each given *request-id*, unless a different host is specified with *@host*.

The *request-id* of any CXbatch request is displayed when the request is first submitted (unless the *silent* mode of operation for the given CXbatch command was specified). The user can also obtain the *request-id* of any request through the use of the *qstat(1)* command.

## DIAGNOSTICS

*qdel* returns an exit status describing what it did. If there were no errors, the exit status is zero. If one or more of the requests were not deleted, the exit status is the number of requests that weren't deleted. If a fatal error occurs and none of the requests are deleted (e.g., a syntax error), the exit status is one of the codes defined in *<sysexits.h>*.

EX_USAGE	The command was used incorrectly, e.g., with the wrong number of arguments, a bad flag, a bad syntax in a parameter.
EX_NOUSER	The user specified with <b>-u</b> did not exist.
EX_NOHOST	The host specified did not exist.
EX_OSFILE	Some batch system file does not exist, cannot be opened, or has an error.
EX_TEMPFAIL	Temporary failure; retry the request at a later time.
EX_NOPERM	You did not have sufficient permission to perform the operation.
EX_SOFTWARE	Too many request-ids were specified.

## CAVEATS

When a CXbatch request is signaled by the methods discussed above, the proper signal is sent to *all* processes comprising the CXbatch *request* that are in the same *process group*. Whenever a CXbatch request is spawned, a new *process group* is established for all processes in the request. However, should one or more processes of the request successfully execute a *setpgid()* system call,

such processes will *not* receive any signals sent by the *qdel(1)* command. This can lead to "rogue" request processes that must be killed by other means such as the *kill(1)* command.

**SEE ALSO**

*qjlist(1)*, *qlimit(1)*, *qstat(1)*, *qsub(1)*, *qmgr(8)*, *kill(2)*, *setpgrp(2)*, *signal(3c)*

**NOTES**

CXbatch is an optional product; for more information, contact your CONVEX sales representative.

**NAME**

qjlist - list the commands in a batch job.

**SYNOPSIS**

**qjlist** *request-id*[@*host*] ...

**DESCRIPTION**

*qjlist* lists the commands in each CXbatch request whose *request-ids* are listed on the command line. To list the commands in a CXbatch request, the invoking user must be the owner (the submitter of the request). The only exception to this rule occurs when the invoking user is the *superuser* or has CXbatch operator privileges as defined in the CXbatch manager database. Under these conditions, the invoker may specify any batch request.

A CXbatch request is always uniquely identified by its *request-id*. A *request-id* is always of the form *seqno* or *seqno.hostname*, where *hostname* identifies the machine from whence the request was originally submitted, and *seqno* identifies the sequence number assigned to the request on the originating host. If the *hostname* portion of a *request-id* is omitted, the local host is always assumed. The local host is searched for each given *request-id* unless a different host is specified with @*host*.

The *request-id* of any CXbatch request is displayed when the request is first submitted (unless the *silent* mode of operation for the given CXbatch command was specified). The user can also obtain the *request-id* of any request through the use of the *qstat(1)* command.

*qjlist* returns an exit status describing what it did. If there were no errors, the exit status is zero. If one or more of the requests were not listed, the exit status is the number of requests that weren't listed. If a fatal error occurs and none of the requests are listed (e.g., a syntax error), the exit status is one of the codes defined in *<syserrits.h>*.

EX_USAGE	The command was used incorrectly, e.g., with the wrong number of arguments, a bad flag, or bad syntax in a parameter.
EX_NOHOST	The host specified did not exist.
EX_OSFIL	Some batch system file does not exist, cannot be opened, or has an error.
EX_NOPERM	You did not have sufficient permission to perform the operation.

**SEE ALSO**

qdel(1), qlimit(1), qstat(1), qsub(1), qmgr(8)

**NOTES**

CXbatch is an optional product; for more information, contact your CONVEX sales representative.

## NAME

qlimit – show supported batch limits, and shell strategy for the named host(s).

## SYNOPSIS

qlimit [ *host-name ...* ]

## DESCRIPTION

*qlimit* displays the set of batch request resource limit types that can be directly enforced on the implied local host or named hosts and also the *batch request shell strategy* defined for the implied local host or named hosts.

If no *host-names* are given, the information displayed is relevant to only the local host. Otherwise, the supported batch request limits, and *batch request shell strategy* for each of the named hosts is displayed.

CXbatch supports many batch request resource limit types that can be applied to a batch request. However, not all UNIX implementations are capable of supporting the rather extensive set of limit types that CXbatch provides.

The set of limits applied to a batch request is always restricted to the set of limits that can be directly supported by the underlying UNIX implementation. If a batch request specifies a limit that cannot be enforced by the underlying UNIX implementation, the limit is ignored, and the batch request operates as though there were no limit (other than the obvious physical maximums) placed upon that resource type.

When an attempt is made to queue a batch request, each *limit-value* specified by the request (that can also be supported by the local UNIX implementation) is compared against the corresponding *limit-value* configured for the destination batch queue. If the corresponding batch queue *limit-value* for all batch request *limit-values* is defined as unlimited, or is greater than or equal to the corresponding batch request *limit-value*, the request can be successfully queued, provided that no other anomalous conditions occur. For request *infinity limit-values*, the corresponding queue *limit-value* must also be configured as infinity.

These resource limit checks are performed irrespective of the batch request arrival mechanism, either by a direct use of the *qsub(1)* command, or by the indirect placement of a batch request into a batch queue via a *pipe* queue. It is impossible for a batch request to be queued in a CXbatch batch queue if *any* of these resource limit checks fail.

Finally, if a request fails to specify a *limit-value* for a resource limit type supported on the execution machine, the corresponding *limit-value*, as configured for the destination queue, becomes the *limit-value* for the unspecified request limit.

Upon the successful queuing of a request in a batch queue, the set of limits under which the request will execute is frozen and will not be modified by subsequent *qmgr(8)* commands that alter the limits of the containing batch queue.

As mentioned above, this command also displays the *shell strategy* configured for the implied local host or named hosts. In the absence of a shell specification for a batch request, CXbatch must choose which shell should be used to execute that batch request. CXbatch supports three different algorithms, or *strategies*, to solve this problem that can be configured for each system by a system administrator, depending on the needs of the user's involved, and upon system performance criterion.

The three possible shell strategies are called:

- *fixed*
- *free*
- *login*

These shell strategies respectively cause the configured *fixed* shell to be exec'd to interpret all batch requests; cause the user's login shell as defined in the password file to be exec'd, which in turn chooses and spawns the appropriate shell for running the batch shell script; or cause only the user's login shell to be exec'd to interpret the script.

A shell strategy of *fixed* means that the same shell chosen by the system administrator is used to execute *all* batch requests (in the absence of a shell specification).

A shell strategy of *free* runs the batch request script *exactly* as would an interactive invocation of the script and is the default CXbatch shell strategy.

The strategies of *fixed* and *login* exist for host systems that are short on available free processes. In these two strategies, a single shell is exec'd, and that same shell is the shell that executes all of the commands in the batch request shell script.

When a shell strategy of *fixed* has been configured for a particular CXbatch system, the "fixed" shell that will be used to run *all* batch requests at that host is displayed.

#### DIAGNOSTICS

*qlimit* returns an exit status describing what it did. If there were no errors, the exit status is zero. If one or more of the hosts were invalid or couldn't be reached, the exit status is the number of hosts that were invalid or couldn't be reached. If a fatal error occurs, the exit status is one of the codes defined in `<sysexits.h>`.

**EX\_OSFILE**      Some batch system file does not exist, cannot be opened, or has an error.

#### SEE ALSO

qdel(1), qjlist(1), qstat(1), qsub(1), qmgr(8)

#### NOTES

CXbatch is an optional product; for more information, contact your CONVEX sales representative.

**NAME**

qmgr - CXbatch queue manager program

**SYNOPSIS**

**qmgr** [ *command* [ *options* ] ]

**DESCRIPTION**

*qmgr* is the primary utility used to configure, administer and operate the CXbatch system. *qmgr* recognizes three classes of users: managers, operators, and users. The superuser is a manager by default. All other managers and operators are assigned using *qmgr*. Managers can use any of the *qmgr* commands. Operators can use a subset of the commands. Users who are not granted manager or operator privileges can use only a few of the *qmgr* commands to display status information and manipulate their own requests.

The CXbatch system is made up of batch and pipe queues that transport and run batch requests. A batch queue holds requests for scheduled, perhaps delayed, processing. A pipe queue is a queue that can pass queued requests on to other pipe queues or batch queues. A batch request is a set of commands, or a shell script, that is to be run non-interactively with the results being returned to the user. The running of the request is handled by the CXbatch system rather than directly by the user.

*qmgr* can be invoked to run a single command given as the command line argument or, if no argument is given, as a command interpreter. As an interpreter, commands will be repeatedly prompted for and executed. The *qmgr* command set is extensive and will be discussed in functional groups below. Command keywords are case insensitive and can be abbreviated to their shortest unique substring. This is indicated in the command descriptions below by capitalizing the require substring. For example, when using the exit command, you must type at least 'ex' to distinguish it from the enable command. Thus the description for this command is:

**EXit**

Exit from the CXbatch manager subsystem.

**INFORMATIONAL COMMANDS**

The following commands return information about *qmgr* commands, CXbatch parameters and queue configuration. They are available to all users.

**HElp** [ *command* ]

Get help information. **HElp** without an argument displays information about what commands are available. **HElp** with an argument displays information about that command. The command may be partially specified as long as it is unique. A more complete help request yields more detailed information.

The **HElp** command provides information that is often more extensive than the command descriptions in this manual page! Use it.

**SHOw All**

Display the standard amount of information about "*limits supported*", *managers*, *parameters*, and *queues*. See below.

**SHOw LImits\_supported**

Display the list of CXbatch resource limit types that are meaningful on this machine. Note that users may request resource limits that are *not* meaningful on the machine where *qsub(1)* is invoked. If the request is to be executed on a remote machine where the limit is meaningful, then CXbatch honors it. Otherwise the unsupported limit is simply ignored.

**SHOw LOng Queues** [ *queue* [ *user* ] ]

Display in long format the status of all CXbatch queues on this host. If a *queue* is

specified, output is limited to that queue. If a *user* is specified, output downplays any requests not belonging to that user.

### SHOW Managers

Display the list of authorized CXbatch managers.

### SHOW Parameters

Display the general CXbatch parameters.

### SHOW Queues [ *queue* [ *user* ] ]

Display the status of all CXbatch queues on this host. If a *queue* is specified, output is limited to that queue. If a *user* is specified, output downplays any requests not belonging to that user.

## QUEUE MANAGEMENT

The following commands are used to define batch and pipe queues. They can be used only by CXbatch managers. For more information on the different types of queues, see the QUEUE TYPES section below.

**CR**eat**e** **B**atch\_queue *queue* **P**Riority = *n* [ **P**Ipeonly ]  
 [ **R**un\_limit = *n* ] [ **I**mpor**t**\_dir = {*Yes, No, Available*} ]  
 [ **S**hare\_policy **F**ixed = *user* ] [ **S**hare\_policy **U**ser ]

Define a batch queue named *queue* with inter-queue priority *n* (0..63). If **P**Ipeonly is specified, then requests may enter this *queue* only if their source is a pipe queue. The specification of a **R**un\_limit sets a ceiling on the maximum number of requests allowed to run in the batch queue at any given time. The default **R**un\_limit is one.

The specification of the **I**mpor**t**\_dir attribute determines the availability of the user's current working directory to a request at runtime. See the **S**E**T** **I**mpor**t**\_dir command for more information.

The specification of a **S**hare\_policy affects how the CPU usage of jobs running in this *queue* is charged. See the **S**E**T** **S**H**A**re\_policy **F**ixed and the **S**E**T** **S**H**A**re\_policy **U**ser commands for more information.

**CR**eat**e** **P**ipe\_queue *queue* **P**Riority = *n* **S**erver = ( *server* )  
 [ **D**estination = *destination* ]  
 [ **D**estination = ( *destination* [ , *destination* ... ] ) ]  
 [ **P**Ipeonly ] [ **R**un\_limit = *n* ]

Define a pipe queue named *queue* with inter-queue priority *n* (0..63) and associate it with a *server*. This is done by specifying an absolute path name to the program binary (*server*) and any arguments required by the program. After **D**estination appears a list of one or more *destination* queues that requests from this pipe *queue* may be sent to. If **P**Ipeonly is specified, then requests may enter this *queue* only if their source is a pipe queue. **R**un\_limit sets a ceiling on the maximum number of requests allowed to run in the pipe queue at any given time. The default **R**un\_limit is one.

### ADd Alias *alias queue*

Add the specified queue *alias* to "*queue*". A queue alias is an alternate name for a queue; any CXbatch command that requires a queue name will also work with an alias.

### DElete Alias *alias*

Delete the specified queue "*alias*". A queue alias is an alternate name for a queue; any CXbatch command that requires a queue name will also work with an alias.

### DElete Queue *queue*

The *queue* is deleted. To delete a *queue*, no requests may be present in the *queue*, and

the *queue* must be disabled. (See "DIisable Queue" below.)

**SEt DESCRIPTION** = (*description*) *queue*

Change the description of the named CXbatch queue.

**SEt PRiority** = *priority queue*

Specify the inter-queue *priority* of a *queue*.

#### QUEUE ACCESS

CXbatch supports queue access restrictions. For each queue access may be either *unrestricted* or *restricted*. If access is *unrestricted*, any request may enter the queue. If access is *restricted*, a request can only enter the queue if the requester or the requester's login group has been given access. Requests submitted by the superuser are an exception; they are always queued, even if the superuser has not explicitly been given access.

The following commands are used to grant access to queues. They can be used only by CXbatch managers.

In the **ADd** and **DElete** commands below a *group* or *user* can be specified in one of two ways: *name* or [*id*]. (The square brackets are literal characters used to indicate a gid or uid.)

**ADd Groups** = *group queue*

**ADd Groups** = ( *group* [ , *group* ... ] ) *queue*

The specified *group(s)* are added to the access list for *queue*.

**ADd Users** = *user queue*

**ADd Users** = ( *user* [ , *user* ... ] ) *queue*

The specified *user(s)* are added to the access list for *queue*.

**DElete Groups** = *group queue*

**DElete Groups** = ( *group* [ , *group* ... ] ) *queue*

The specified *group(s)* are deleted from the access list for *queue*.

**DElete Users** = *user queue*

**DElete Users** = ( *user* [ , *user* ... ] ) *queue*

The specified *user(s)* are deleted from the access list for *queue*.

**SEt NO\_Access** *queue*

Specify that no one can place requests in *queue*.

**SEt Unrestricted\_access** *queue*

Specify that no requests will be turned away from *queue* on the grounds of queue access restrictions.

#### BATCH QUEUE CONFIGURATION

The following commands are used to set the operating parameters for batch queues. They can only be used by CXbatch managers.

**SEt ACCOunting** = {*OFF*,*ON*} *queue*

Turn accounting on/off for a CXbatch batch queue. The queue named as a parameter of the command must already exist.

**SEt ACTivity\_id\_offset** = *offset queue*

Set the activity ID offset for a CXbatch batch queue. The queue named as a parameter of the command must already exist.

**SEt CHKpntable** = {*Yes*,*No*,*Available*} *queue*

Sets the status of the per-queue checkpoint resource. The queue must exist. The queue's

checkpoint resource value and the request's flags are used to determine if a request may be checkpointed.

If the checkpoint attribute is set to "Yes", then any request submitted to this queue will be checkpointed at CXbatch shutdown and may be checkpointed by the request owner or CXbatch operator, unless it explicitly requested not to be checkpointed. If this attribute is set to "Available", the request is checkpointable only if it specifically asked to be checkpointed. Requests in a queue with the checkpoint attribute set to "No" cannot be checkpointed by CXbatch.

**Set Import\_dir = {Yes,No,Available} queue**

Changes the **Import\_dir** attribute for a *queue*. The queue must already exist. If the **Import\_dir** attribute is set to *Yes*, any job submitted to this queue is run in the user's current working directory, unless the job specifically requests not to be imported. If the **Import\_dir** attribute is set to *Available*, only jobs that specifically request to be imported are imported. If it is set to *No*, jobs are run in the home directory. When importing, if a job is executed on a remote machine, CXbatch tries to access the local directory using NFS mounts. **NOTE:** CXbatch will make temporary NFS mounts into the /tmp filesystem. Care should be taken that any automatic clean-up operations on the /tmp filesystem do not traverse NFS mount points.

**Set MAXimum Request\_priority = priority queue**

Set a limit on request priorities on a per queue basis. Requests queued with a priority higher than the queue's maximum will have their priority lowered to the maximum.

#### BATCH QUEUE LIMITS

The following commands are used to set the limits batch queues impose on their requests. Every batch queue on the local host have each of the following limits associated with it at all times. If a request already in the queue has asked for more than a new limit, it is given a grandfather clause and allowed to retain its original limits. A request which specifies a particular limit may only enter a batch queue if the queue's corresponding limit is greater than or equal to the request's limit. See the section on **LIMITS** below for more information on the syntax on the various limit arguments. These commands can only be used by CXbatch managers.

**Set COREfile\_limit = ( limit ) queue**

Set a per-process maximum core file size *limit* for a batch queue against which the per-process maximum core file size limit for a request may be compared.

**Set DATA\_limit = ( limit ) queue**

Set a per-process maximum data segment size *limit* for a batch queue against which the per-process maximum data segment size limit for a request may be compared.

**Set NICE\_value\_limit = nice-value queue**

Set the UNIX *nice-value* limit for a batch queue, against which the *nice* value for a request may be compared. A request specifying a *nice-value* may only enter a batch queue if the queue's nice value is numerically less than (more willing to allow access to the CPU) or equal to the request's *nice* value. *nice-value* is an integer preceded by an optional negative sign.

**Set PER\_Process Cpu\_limit = ( limit ) queue**

Set a per-process maximum CPU time *limit* for a batch queue against which the per-process maximum CPU time limit for a request may be compared.

**Set PER\_Process Permfile\_limit = ( limit ) queue**

Set a per-process maximum permanent file size *limit* for a batch queue against which the

per-process maximum permanent file size limit for a request may be compared.

**SEt STack\_limit** = ( *limit* ) *queue*

Set a per-process maximum stack segment size *limit* for a batch queue against which the per-process maximum stack segment size limit for a request may be compared.

**SEt Working\_set\_limit** = ( *limit* ) *queue*

Set a per-process maximum working set size *limit* for a batch queue against which the per-process maximum working set size limit for a request may be compared.

## PIPE QUEUE CONFIGURATION

The following commands are used to set the operating parameters for pipe queues. They are only available to CXbatch managers.

**ADd DESTination** = *destination queue*

**ADd DESTination** = ( *destination* [ , *destination* ... ] ) *queue*

The specified *destination(s)* are added as valid destinations for a pipe queue named "*queue*".

**DElete DESTination** = *destination queue*

**DElete DESTination** = ( *destination* [ , *destination* ... ] ) *queue*

Delete the mappings from the pipe queue *queue* to the *destination* queues. All requests from the named *queue* being transferred to a deleted *destination* complete normally. If all *destinations* for a pipe *queue* are deleted in this manner, the pipe *queue* is effectively stopped.

**SEt DESTination** = *destination queue*

**SEt DESTination** = ( *destination* [ , *destination* ... ] ) *queue*

Associate one or more *destination* queues with a particular pipe *queue*.

**SEt Pipe\_client** = ( *client* ) *queue*

Associate a *pipe client* with a pipe *queue*. *client* should consist of the absolute path name to the program binary followed by any arguments required by the program.

## DESIGNATING MANAGERS AND OPERATORS

The following commands are used to specify CXbatch managers and operators. They can only be used by CXbatch managers.

A *manager* specification consists of an account name specification, followed by a colon, followed by either the letter *m* or the letter *o*. There are four ways to specify an account name:

```

local_account_name
[local_user_id]
[remote_user_id]@remote_machine_name
[remote_user_id]@[remote_machine_mid]

```

(The square brackets are literal characters used to indicate a uid or mid.) If the account name specification is followed by *:m*, the account is designated as a CXbatch *manager* account, capable of using all *qmgr* commands. If the account name specification is followed by *:o*, the account is designated as a CXbatch *operator* account, capable of only using those commands appropriate for a CXbatch *operator*. The *root* account always has full privileges.

**ADd Managers** *manager* ...

The specified *manager(s)* are added to the list of authorized CXbatch managers with privileges as specified.

**DElete Managers** *manager* ...

The specified *manager(s)* are deleted from the list of authorized CXbatch managers.

**SEt MANagers** *manager ...*

The list of authorized CXbatch managers is set to the specified *manager(s)*.

**GENERAL CXBATCH MANAGEMENT**

The following commands are used to set the general operating parameters of CXbatch. Only CXbatch managers can use these commands.

**SEt ACC\_logfile** *filename*

Change the file being used for CXbatch batch accounting.

**SEt Aid\_mask** = *mask*

Set the activity ID mask. Generally, this mask is the same as the spacing between aids in */etc/activities*. The following equation is used to determine the activity ID of the CXbatch job:

$$job\_aid = submitter\_aid - (submitter\_aid \% aid\_mask) + queue\_aid\_offset$$

where % is the modulus (remainder) function.

**Warning:** As shipped, the aid mask is one. When the mask is one, the above equation simplifies to  $job\_aid = submitter\_aid + queue\_aid\_offset$ . In this case, if a CXbatch job submits another CXbatch job, the second job has an activity ID of  $submitter\_aid + queue1\_aid\_offset + queue2\_aid\_offset$ . This is generally not desirable! Setting the activity ID mask to the spacing in */etc/activities* prevents this from happening.

**SEt CHEckpoint\_directory** *directory*

Specify the pathname of the checkpoint directory. All checkpoint files created by CXbatch after this command is issued will be placed in *directory*. Existing checkpoint files will not be moved. The *directory* must exist and be a valid directory.

**SEt DEBug level**

Set the debug *level*. The following values are valid:

- 0 No debug
- 1 Minimum debug
- 2 Maximum debug

**SEt DEFault Batch\_request Priority** *priority*

Set the default intra-batch-queue *priority*. This is *not* the UNIX execution time priority. This is the priority used if the user does not specify an intra-queue priority parameter on the *qsub(1)* command.

**SEt DEFault Batch\_request Queue** *queue*

Set the default batch *queue*. This is the queue used if the user does not specify a queue parameter on the *qsub(1)* command.

**SEt DEFault DESTination\_retry Time** *retry\_time*

Set the default number of hours that can elapse during which time a pipe queue destination can be unreachable, before being marked as completely failed.

**SEt DEFault DESTination\_retry Wait interval**

Set the default number of minutes to wait before retrying a pipe queue destination that was unreachable at the time of the last attempt.

**SEt Global Per\_user Run\_limit** = *run-limit*

Set the *global per-user run-limit* of the local system. The *global per-user run-limit* is the maximum number of requests that any given user can have running on the local system

at any given time. A *global per-user run-limit* of 0 turns off the enforcement of this limit.

**SEt MAIL *userid***

Specify the *userid* used to send CXbatch mail.

**SEt NO\_Default Batch\_request Queue**

Indicate that there is to be no default batch request queue.

**SEt SHELL\_strategy **FL**xed = ( *shell* )**

Specify that *shell* should be used to execute all batch requests. *shell* must be the absolute path name of a command interpreter. (See the SHELL STRATEGIES section below for more information.)

**SEt SHEll\_strategy **FR**ee**

Specify that the *free* shell strategy should be used to execute all batch requests. The *free* shell strategy duplicates the shell choice that would have been made if the batch request script had been executed interactively. Under this strategy, the user's *login shell* is allowed to determine the shell to be used to execute the batch request. The user's *login shell* is the shell named within the user's entry in the password file (see *passwd(4)*). (See the SHELL STRATEGIES section below for more information.)

**SEt SHEll\_strategy **Log**in**

Specify that the *login* shell strategy should be used to execute all batch requests. Under the *login* shell strategy, the user's login shell is used to execute the batch request. The login shell is the shell named in the password file (see *passwd(4)*). (See the SHELL STRATEGIES section below for more information.)

**GENERAL CXBATCH OPERATION**

The following commands are used for starting and shutting down CXbatch. They can be used by CXbatch managers or operators.

**SHUtdown [ **Force** ] [ *seconds* ]**

Shutdown CXbatch on the local host.

Each running checkpointable batch request is checkpointed, and, if the checkpoint succeeds, terminated. Successfully checkpointed requests will be restarted from their checkpointed state when CXbatch is rebooted.

After running requests are checkpointed, a SIGTERM signal is sent to each process of each request presently running. After the specified number of *seconds* of real time have elapsed, a SIGKILL signal is sent to all remaining processes for each request. If a *seconds* value is not specified, the delay is sixty seconds. Unlike **ABort Queue**, **SHUtdown** requeues all of the requests it kills, provided that the initial SIGTERM signal is caught or ignored by the running request.

When the optional **force** keyword is present CXbatch ignores any checkpoint errors incurred during the shutdown.

**STArt Cxbatch**

Start CXbatch on the local host. This command will fail if CXbatch is currently running on the local host.

**QUEUE OPERATION**

The following commands are used in operating queues in CXbatch. They can be used by CXbatch managers or operators.

**ABort Queue *queue* [ *seconds* ]**

All requests in the named queue that are currently running are aborted as follows. A SIGTERM signal is sent to each process of each request presently running in the named *queue*. After the specified number of *seconds* of real time have elapsed, a SIGKILL signal is sent to all remaining processes for each request running in the named *queue*. If a *seconds* value is not specified, the delay is sixty seconds. All requests aborted by this command are deleted, and all output files associated with the requests are returned to the appropriate destination.

#### **DI**sable Queue *queue*

Prevent any more requests from being placed in this queue.

#### **EN**able Queue *queue*

If the queue is already enabled, this command has no effect. Otherwise, the queue is enabled to accept new requests.

#### **MO**Ve Queue *queue1 queue2*

Move all requests currently in *queue1* to *queue2*. The request is moved regardless of any queue limit violations, access restrictions, or attribute violations.

#### **Pur**ge Queue *queue*

All queued requests are dropped from the queue and are irretrievably lost. Running requests in the *queue* are allowed to complete.

#### **ST**Art Queue *queue*

If the queue is already started, then nothing happens. Otherwise, the queue is started and requests in the queue are eligible for selection.

#### **ST**Op Queue *queue*

Any requests in the queue that are currently running are allowed to complete. All other requests are "frozen" in the queue. New requests can still be submitted to the queue, but are "frozen" like the other requests in the queue.

### **QUEUE OPERATING PARAMETERS**

The following commands are used to set the operating parameters of CXbatch queues. They can be used by CXbatch managers or operators.

#### **SEt PER\_User Run\_limit = run-limit queue**

Change the *per-user run-limit* of a CXbatch batch queue. The *per-user run-limit* determines the maximum number of requests that any given user can have running in the queue at any given time. A *per-user run-limit* of 0 turns off the enforcement of this limit.

#### **SEt Run\_limit = run-limit queue**

Change the *run-limit* of a CXbatch batch or pipe queue. The *run-limit* determines the maximum number of requests that are allowed to run in the queue at any given time.

#### **SEt SHARe\_policy Fixed = user queue**

Change the *share-policy* of a CXbatch batch queue. A queue with a fixed share policy will charge the CPU usage of jobs run in that queue to *account*. There are two ways to specify a user: *username* or [*id*]. (The square brackets are literal characters used to indicate a uid.)

#### **SEt SHARe\_policy User queue**

Change the *share-policy* of a CXbatch batch queue. A queue with a user share policy will charge the CPU usage of jobs run in that queue to the account from which the job was submitted.

**REQUEST OPERATIONS**

The following commands operate on CXbatch requests. They can only be executed by CXbatch managers or operators.

**MOVE Request *requestid* ... *queue***

Move the request(s) named by the *requestid(s)* to the named queue. The request is moved regardless of any queue limit violations, access restrictions, or attribute violations.

**RESume Request *requestid* ...**

Resume execution of suspended requests. Resumed requests start out in the *queued* state. Once the resumed request is about to enter the *running* state, it will be restarted from its checkpointed state instead of being re-run in its entirety.

**RUn Request *requestid* ...**

Force the request(s) named by the *requestid(s)* to begin executing immediately. If running the request would exceed the current run limit of the queue, then the queue's run limit will be increased by one until the request finishes executing.

**SUspend Request *requestid* ...**

Temporarily freeze execution of the named requests. The requests are checkpointed and terminated. Of course, a request that fails to checkpoint will continue to execute. Only checkpointable requests may be suspended in this manner.

**USER REQUEST OPERATIONS**

The following commands operate on CXbatch requests. These commands may be used by non-privileged user in dealing with their own requests. Otherwise, they are limited to use by CXbatch managers or operators.

**CHKpnt Request *requestid* ...**

Checkpoint the request(s) named by the *requestid(s)*. The state of the named batch request(s) is saved into a set of checkpoint files stored in the *checkpoint directory*.

**DElete Request *requestid* ...**

Delete the request(s) named by the *requestid(s)*. This command can delete both running and non-running requests. If a request is running, then all processes of the request are sent a SIGKILL signal.

**HOLd Request *requestid* ...**

Makes the specified request(s), named by *requestid(s)*, ineligible for running. The request(s) must be in the *queued* state prior to being held. The **RELease Request** command removes the effect of **hold**.

If a request is held by an operator, only an operator can release it.

**MODify Request *field* = *value requestid* ...**

Change the field of the request(s) specified by *requestid(s)* to be *value*.

The legal values for *field* are:

**Priority** - change the intra-queue request priority. A user can only decrease a request's priority, *operator* privileges are required to raise a request's priority.

**MOVE My\_request *requestid* ... *queue***

Move your request(s) named by the *requestid(s)* to the named queue. The request is not moved if any queue limits, access restrictions, or attributes would have prevented the request from being submitted to the new queue.

**RELease Request *requestid* ...**

Makes the specified request(s), named by *requestid(s)*, eligible for running. The request(s)

must be in the *holding* state prior to being released.

## QUEUE TYPES

CXbatch supports two different queue types that provide two very different functions. These two queue types are known as *batch* and *pipe*.

The queue type of *batch* can only be used to execute CXbatch batch requests. Only CXbatch batch requests created by the *qsub(1)* command can be placed in a *batch* queue.

Queues of type *pipe* are used to send CXbatch requests to other *pipe* queues or *batch* queues. In general, *pipe* queues act as the mechanism that CXbatch uses to transport *batch* requests to distant queues on other remote machines. It is also perfectly legal for a *pipe* queue to transport requests to queues on the same machine.

When a *pipe* queue is defined, it is given a destination set, that defines the set of possible destination queues for requests entered in that *pipe* queue. In this manner, it is possible for a *batch* request to pass through many pipe queues on its way to its ultimate destination, that must eventually be a queue of type *batch*.

Each *pipe* queue has an associated server. For each request handled by a *pipe* queue, the associated server is spawned which selects a queue destination for the request being handled, based upon the characteristics of the request and upon the characteristics of each queue in the destination set defined for the pipe queue.

Because a different server can be configured for each *pipe* queue, and *batch* queues can be endowed with the *pipeonly* attribute that will only admit requests queued via another *pipe* queue, it is possible for respective CXbatch installations to use *pipe* queues as a request class mechanism, placing requests that ask for different resource allocations in different queues, each of which can have different associated limits and priorities.

It is also completely possible for a pipe client (pipe queue server), when handling a request, to discover that no destination queue will accept the request, for various reasons that can include insufficient resource limits to execute the request or a lack of a corresponding account or privilege for queuing at a remote queue. In such circumstances, the request is deleted, and the user is notified by mail (see *mail(1)*).

## SHELL STRATEGIES

The execution of a batch request requires the creation of a shell process to interpret the shell script that defines the batch request. On many UNIX systems, there is more than one shell available (e.g., */bin/csh*, */bin/ksh*, */bin/sh*). To deal with this problem, CXbatch allows a shell path-name to be specified when a batch request is first submitted (*qsub* option *-s*).

If no particular shell is specified for the execution of the request, CXbatch must have some other means of deciding which shell to use when spawning the request. The solution to this dilemma has been to equip CXbatch with a batch request shell strategy that can be configured as necessary by the local system administrators.

The batch request shell strategy determines the shell to be used when executing a batch request on the local host that fails to identify any specific shell for its execution. Three such shell strategies can be configured for CXbatch, and they are known by the names of *fixed*, *free*, and *login*.

A shell strategy of *fixed* causes the request to be run by the *fixed shell*, the path name of which is configured by the system administrator. Thus, a particular CXbatch installation may be configured with a *fixed* shell strategy where the default shell used to execute all batch requests is defined as the Bourne shell.

A shell strategy of *free* causes the user's login shell (as defined in the password file), to be executed. This shell is, in turn, given a path name to the batch request shell script, and it is the user's login shell that actually decides which shell should be used to interpret the script. The *free* shell strategy therefore runs the batch request script exactly as would an interactive invocation of the script and is the default CXbatch shell strategy.

The third shell strategy of *login* causes the user's login shell (as defined in the password file) to be the default shell used to interpret the batch request shell script.

The strategies of *fixed* and *login* exist for host systems that are short on available free processes. In these two strategies, a single shell is executed, and that same shell is the shell that executes all of the commands in the batch request script (barring shell exec operations in any user startup files: *.profile*, *.login*, *.cshrc*).

The shell strategy as configured for any particular host can always be determined by the CXbatch *qlimit* command.

## LIMITS

CXbatch supports many batch request resource limit types that can be applied to a CXbatch batch queue. The configurability of these limits allows a CXbatch manager to set batch queue-specific resource limits that all batch requests in the queue must adhere to.

The syntax of a *limit* in commands of the form **SEt Some\_limit = ( limit ) queue** is quite flexible.

For *finite* CPU time limits, the acceptable syntax is as follows:

```
[[hours :] minutes :] seconds [.milliseconds]
```

Whitespace can appear anywhere between the principal tokens, with the exception that no whitespace can appear around the decimal point. **NOTE:** The *milliseconds* value may be ignored if the system does not support such granularity.

Example time "*limit-values*" are:

```
center, tab(;;) | l . 1234 : 58 : 21.29;- 1234 hrs 58 mins 21.290 secs 12345;- 12345 seconds 121.1;- 121.100 seconds 59:01;- 59 minutes and 1 second
```

For all other *finite* limits (with the exclusion of the *nice-value*), the acceptable syntax is:

```
.fraction [units]
```

or

```
integer [.fraction] [units]
```

where the "*integer*" and "*fraction*" tokens represent strings of up to eight decimal digits, denoting the obvious values. In both cases, the "*units*" of allocation may also be specified as one of the case insensitive strings:

```
center, tab(;;) | l . b;- bytes w;- words kb;- kilobytes (2^10 bytes) kw;- kilowords (2^10 words) mb;- megabytes (2^20 bytes) mw;- megawords (2^20 words) gb;- gigabytes (2^30 bytes) gw;- gigawords (2^30 words)
```

In the absence of any "*units*" specification, the units of bytes are assumed.

For all limit types with the exception of the *nice-value*, it is possible to state that no limit should be applied. This is done by specifying a "*limit*" of "unlimited", or any initial substring thereof.

The complications caused by batch request resource limits first show up when queuing a batch request in a batch queue. This operation is described in the following paragraphs.

If a batch request specifies a limit that cannot be enforced by the underlying UNIX implementation, the limit is ignored, and the batch request operates as though there were no limit (other than the obvious physical maximums) placed upon that resource type. (See the "*qlimit*"(1)

command to find out what limits are supported by a given machine.)

For each remaining finite limit that can be supported by the underlying UNIX implementation that is *not* a CPU "time-limit" or UNIX "nice-value", the "limit-value" is internally converted to the units of bytes or words, whichever is more appropriate for the underlying machine architecture.

As an example, a working set size limit value of 321 megabytes would be interpreted as  $321 \times 2^{20}$  bytes, provided that the underlying machine architecture was capable of directly addressing single bytes. Thus the original limit coefficient of 321 would become  $321 \times 2^{20}$ . On a machine that was only capable of addressing words, the appropriate conversion of  $321 \times 2^{20}$  bytes / # of bytes-per-word would be performed.

If the result of such a conversion would cause overflow when the coefficient was represented as a signed-long integer on the supporting hardware, the coefficient is replaced with the coefficient of:  $2^{N-1}$  where  $N$  is equal to the number of bits of precision in a signed long integer. For typical 32-bit machines, this default extreme limit would therefore be  $2^{31-1}$  bytes. For word addressable machines in the supercomputer class supporting 64-bit long integers, the default extreme limit would be  $2^{63-1}$  words.

Lastly, some implementations of UNIX reserve coefficients of the form:  $2^{N-1}$  as synonymous with infinity, meaning no limit is to be applied. For such UNIX implementations, CXbatch further decrements the default extreme limit so as to not imply infinity.

The identical internal conversion process as described in the preceding paragraphs is also performed for all finite limit-values specified with a particular batch request.

After each applicable request limit has been converted as described above, the resulting limit is then compared against the corresponding limit as configured for the destination batch queue. If the corresponding batch queue limit for all batch request limits is defined as unlimited, or is greater than or equal to the corresponding batch request limit, the request can be successfully queued, provided that no other anomalous conditions occur. For requests that ask for a limit of infinity, the corresponding queue limit must also be configured as infinity.

These resource limit checks are performed irrespective of the batch request arrival mechanism, either by a direct use of the "qsub"(1) command, or by the indirect placement of a batch request into a batch queue via a "pipe" queue. It is impossible for a batch request to be queued in a CXbatch batch queue if "any" of these resource limit checks fail.

Finally, if a request fails to specify a "limit" for a resource limit type that is supported on the execution machine, the corresponding "limit" as configured for the destination queue becomes the "limit" for the request.

Upon the successful queuing of a request in a batch queue, the set of limits under which the request will execute is frozen and will not be modified by subsequent qmgr(8) commands that alter the limits of the containing batch queue.

## DIAGNOSTICS

qmgr returns an exit status describing what the last qmgr command did. If there were no errors, the exit status is zero. If one or more of the operations failed, the exit status is the number of operations that failed. If a fatal error occurs and command completely fails, (ex, a syntax error), the exit status is one of the codes defined in <sysexit.h>.

- EX\_USAGE      The command was used incorrectly, e.g., with the wrong number of arguments, a bad flag, a bad syntax in a parameter.
- EX\_NOHOST    The host specified did not exist.
- EX\_OSFILE    Some batch system file does not exist, cannot be opened, or has an error.
- EX\_TEMPFAIL  Temporary failure; retry the request at a later time.

**SEE ALSO**

qdel(1), qlist(1), qlimit(1), qstat(1), qsub(1), qmapmgr(8)

**NOTES**

CXbatch is an optional product; for more information, contact your CONVEX sales representative.

**NAME**

*qps* - display process status of CXbatch related processes

**SYNOPSIS**

```
qps [ queue-name ... ]
qps -r request-id
qps -{p|q} process-id
```

**DESCRIPTION**

*qps* prints information about CONVEX CXbatch related processes.

If no queues are specified, information is printed about all CXbatch related processes, including the daemon processes. Otherwise, information is printed for the specified queues only. Only batch queues on the local system are significant; pipe queues and remote queues do not have processes running on the local system.

Information about processes pertaining to a particular request can be obtained by using the -*r request-id* option. Only a single request-id can be specified.

For each process, *qps* prints the queue name (QUEUE), the request ID (REQ), the process ID (PID), the state (STAT) of the process, CPU time (TIME) used by the process (including both user and system time) and which command is running (COMMAND). More information about these fields can be found on the *ps(1)* man page.

Using the -*p process-id* option, you can inquire whether a particular process is running from within the CXbatch system. If it is, a line is printed stating the queue and request to which the process is related and *qps* exits with a status of 0. Otherwise, you are informed that the process is not running from within the CXbatch system and *qps* exits with a status of 1. For the purposes of this inquiry, the top level daemons are not considered to be running under the CXbatch system, but the CXbatch shepherd processes are.

The -*q process-id* option is a silent version of the -*p* option. Nothing is printed, but the exit status of *qps* is set appropriately. Only a single process-id can be specified for either of these options.

**DIAGNOSTICS**

*qps* returns an exit status of 0 if no errors occur, unless the the -*p* or -*q* options are specified, in which case the exit status is 0 or 1 as described above. If a fatal error occurs, the exit status is one of the codes defined in *<syserrits.h>*.

- |          |   |
|----------|---|
| EX_USAGE | The command was used incorrectly, e.g., with the wrong number of arguments, a bad flag, a bad syntax in a parameter.                                  |
| EX_OSERR | An internal call to <i>ps(1)</i> failed. If this error occurs, first make sure a '/bin/lxwlg' returns valid output, then check the status of CXbatch. |

**SEE ALSO**

*ps(1)*, *qstat(1)*

**NOTES**

CXbatch is an optional product; for more information, contact your CONVEX sales representative.

**NAME**

*qrestart* - restart checkpointed CXbatch request(s).

**SYNOPSIS**

**qrestart** [-f] [ *request-id* ]

**DESCRIPTION**

*qrestart* restarts the requests whose *request-ids* are listed on the command line.

Only successfully checkpointed requests may be restarted. To restart a request, the invoking user must be the owner of the request or have CXbatch operator privileges.

*qrestart* is typically only used if the automatic restart of the request by CXbatch failed.

The options have the following meanings:

**-f** Force restart of request even if non-restartable conditions exist in any of the processes of the request.

**DIAGNOSTICS**

*qrestart* returns an exit status describing what it did. If there were no errors, the exit status is zero. If one or more of the requests were not restarted, the exit status is the number of requests that weren't restarted. If a fatal error occurs and none of the requests are restarted (e.g., a syntax error), the exit status is one of the codes defined in *<sysexits.h>*.

**EX\_USAGE** The command was used incorrectly, e.g., with the wrong number of arguments, a bad flag, a bad syntax in a parameter.

**EX\_OSFILE** Some batch system file does not exist, cannot be opened, or has an error.

**EX\_TEMPFAIL** Temporary failure; retry the command at a later time. **EX\_SOFTWARE** Too many request-ids were specified.

**REFERENCES**

*qsub*(1), *qchkpnt*(1), *chkpnt*(1), *restart*(1), *qmgr*(8)

**NOTES**

CXbatch is an optional product; for more information, please contact your CONVEX sales representative.

**NAME**

*qstat* - display status of CXbatch queue(s)

**SYNOPSIS**

*qstat* [ **-a** ] [ **-l** ] [ **-m** ] [ **-u user-name** ] [ **-x** ] [ *queue-name ...* ] [ *queue-name@host-name ...* ]

**DESCRIPTION**

*qstat* displays the status of CONVEX CXbatch queues.

If no queues are specified, the current state of each CXbatch queue on the local host is displayed. Otherwise, information is displayed for the specified queues only. Queues may be specified either as *queue-name* or *queue-name@host-name*. In the absence of a *host-name* specifier, the local host is assumed.

For each selected queue, *qstat* displays a *queue header* (information about the queue itself) followed by information about requests in the queue. Ordinarily, *qstat* shows only those requests belonging to the invoker. The following flags are available:

- a** Shows all requests.
- l** Requests are shown in a long format.
- m** Requests are shown in a medium-length format.
- u user-name** Shows only those requests belonging to *user-name*.
- x** The queue header is shown in an extended format.

The *queue header* always includes the queue-name, queue type, queue status (see below), an indication of whether or not the queue is *pipeonly* (accepts requests from pipe queues only), and the number of requests in the queue. An extended queue header also displays the priority and run limit of a queue, as well as the access restrictions, cumulative use statistics, server and destinations (if a pipe queue), and resource limits (if a batch queue).

By default, *qstat* displays the following information about a request: the *request-name*, the *request-id*, the owner, the relative request priority, and the current request state (see below). For running requests, the process group is also shown, as soon as it becomes available to the local CXbatch daemon.

*qstat -m* shows the following additional information: if the request was submitted with the constraint that it not run before a certain time and date, the constraining time and date are also displayed.

*qstat -l* shows the time at which the request was created, an indication of whether or not mail will be sent, where mail will be sent, and the user name on the originating machine. If a batch queue is being examined, resource limits, planned disposition of *stderr* and *stdout*, any advice concerning the command interpreter, and the umask value are shown.

The relative ordering of requests within a queue does not always determine the order in which the requests are run. The CXbatch request scheduler is allowed to make exceptions to the request ordering for the sake of efficient machine resource usage. However, requests appearing near the beginning of the queue have higher priority than requests appearing later, and are usually run before requests appearing later on in the queue.

**QUEUE STATE**

The general state of a queue is defined by two principal properties of the queue.

The first property determines whether or not requests can be submitted to the queue. If they can, the queue is said to be *enabled*. Otherwise the queue is said to be *disabled*. One of the words CLOSED, ENABLED, or DISABLED appears in the queue status field to indicate the respective queue states of: enabled (with no local CXbatch daemon), enabled (and local CXbatch daemon is present), and disabled. Requests can only be submitted to the queue if the queue is enabled and the local CXbatch daemon is present.

The second principal property of a queue determines if requests that are ready to run, but are not now presently running, will be allowed to run upon the completion of any currently running requests, and whether any requests are presently running in the queue.

If queued requests not already running are blocked from running, and no requests are presently executing in the queue, the queue is said to be *stopped*. If the same situation exists with the difference that at least one request is running, the queue is said to be *stopping*, where the requests presently executing will be allowed to complete execution, but no new requests will be spawned.

If queued requests ready to run are only prevented from doing so by the CXbatch request scheduler, and one or more requests are presently running in the queue, the queue is said to be *running*. If the same circumstances prevail with the exception that no requests are presently running in the queue, the queue is said to be *inactive*. Finally, if the CXbatch daemon for the local host upon which the queue resides is not running, but the queue would otherwise be in the state of *running* or *inactive*, the queue is said to be *shutdown*. The queue states describing the second principal property of a queue are therefore respectively displayed as STOPPED, STOPPING, RUNNING, INACTIVE, and SHUTDOWN.

### REQUEST STATE

The state of a request may be *arriving*, *holding*, *waiting*, *queued*, *routing*, *running*, *departing*, or *exiting*.

- *arriving*      The request is being enqueued from a remote host.
- *holding*        The request is presently prevented from entering any other state (including the *running* state), because a *hold* has been placed on the request.
- *waiting*        The request was submitted with the constraint that it not run before a certain date and time, and that date and time have not yet arrived.
- *queued*         The request is eligible to proceed (by *routing* or *running*).
- *routing*        The request has reached the head of a pipe queue and is receiving service.
- *departing*     A request is *departing* from the time the pipe queue turns to other work until the request has arrived intact at its destination.
- *running*        The request has reached its final destination queue and is actually executing.
- *exiting*        The batch request has completed execution and will exit from the system after the required output files have been returned (to possibly remote machines).

Imagine a batch request originating on a workstation, destined for the batch queue of a computation engine, to be run immediately. That request would first go through the states *queued*, *routing*, and *departing* in a local pipe queue. Then it would disappear from the pipe queue. From the point of view of a queue on the computation engine, the request would first be *arriving*, then *queued*, *running*, and finally *exiting*. Upon completion of the *exiting* phase of execution, the batch request would disappear from the batch queue.

### DIAGNOSTICS

*qstat* returns an exit status describing what it did. If there were no errors, the exit status is zero. If one or more of the queues were not listed, the exit status is the number of queues that weren't listed. If a fatal error occurs and none of the queues are listed (ex, a syntax error), the exit status is one of the codes defined in `<sysexits.h>`.

- |           |  |
|-----------|--|
| EX_USAGE  | The command was used incorrectly, e.g., with the wrong number of arguments, a bad flag, a bad syntax in a parameter. |
| EX_NOUSER | The user specified with <code>-u</code> did not exist.   |

EX\_OSFILE        Some batch system file does not exist, cannot be opened, or has an error.

**SEE ALSO**

qdel(1), qjlist(1), qlimit(1), qsub(1), qmgr(8)

**NOTES**

CXbatch is an optional product; for more information, contact your CONVEX sales representative.

## NAME

**qsub** – submit a CXbatch request.

## SYNOPSIS

**qsub** [ *flags* ] [ *script-file* ]

## DESCRIPTION

*qsub* submits a batch request to CONVEX CXbatch.

If no *script-file* is specified, the set of commands to be executed as a batch request is taken directly from the standard input file (*stdin*). In all cases however, the *script file* is spooled, so that later changes will *not* affect previously queued batch requests.

All of the flags that can be specified on the command line can also be specified within the first comment block inside the batch request *script file* as *embedded default flags*. Such flags appearing in the batch request *script file* set default characteristics for the batch request. If the same flag is specified on the command line, the command line flag (and any associated value) takes precedence over the *embedded* flag. See the section entitled LONG DESCRIPTION for more information on *embedded default flags*.

What follows is a terse definition of the flags implemented by the *qsub* command (see the section LONG DESCRIPTION for the complete definition and syntax used for each of these flags).

- a** – run request after stated time
- b** – set the billing activity for request
- c** – request is checkpointable
- cp** – specify periodic checkpoint frequency
- e** – direct *stderr* output to stated destination
- eo** – direct *stderr* output to the *stdout* destination
- h** – put request on hold after submitting
- i** – request requires current directory to be imported
- ke** – keep *stderr* output on the execution machine
- ko** – keep *stdout* output on the execution machine
- l** – run request under a login shell
- lr** – establish a resource limit
- mb** – send mail when the request begins execution
- me** – send mail when the request ends execution
- mu** – send mail for the request to the stated user
- ni** – don't import the current directory
- nr** – declare that batch request is not restartable
- nc** – request is not checkpointable
- o** – direct *stdout* output to the stated destination
- p** – specify intra-queue request priority
- q** – queue request in the stated queue
- r** – assign stated request name to the request
- s** – specify shell to interpret the batch request script
- t** – signal process when the job completes
- x** – export all environment variables with request
- y** – append accounting data to *stdout* output file
- z** – submit the request silently

If you request that a batch request be run under a login shell, the system and user startup files will be read (*/etc/login* and *~/.login* for C-shell or */etc/profile* and *\$HOME/.profile* for Bourne shell). A C-shell login shell will also read */etc/logout* and *~/.logout* while exiting. The *~/.cshrc* file is read by the C-shell regardless of whether or not it is executed as a login shell.

The environment string `ENVIRONMENT=BATCH` is added to the environment so that shell scripts (and the user's `.profile` (Bourne shell) or `.cshrc` and `.login` (C-shell) scripts), can test for batch request execution when appropriate, and not (for example) perform any setting of terminal characteristics, because a batch request is not connected to an input terminal. For example, if your login shell is C-shell, the following `.login` file prevents `stty`, `tset`, and `msgsg` from being run during batch jobs.

```
if (! $?ENVIRONMENT) then
    stty crt erase ^H kill ^U
    tset -Q
    msgsg -q
endif
```

If your login shell is Bourne shell, the following `.profile` file has the same effect.

```
if test "$ENVIRONMENT" != "BATCH"
then
    stty crt erase ^H kill ^U
    tset -Q
    msgsg -q
fi
```

When CXbatch is configured on more than one machine, it is possible for users to submit jobs to remote machines in the batch network. However, CXbatch must ensure that the submitter has permission to access the remote machine. This is accomplished with the *user equivalence* mechanism described in `hosts.equiv(5)` (the same method that is used by `rlogin` and `rsh`). If the machines in the batch network are not equivalenced in the `/etc/hosts.equiv` file, users must create a `.rhosts` file in their home directories. If this is not done, the submitter will not have access to the remote machine. Read the `hosts.equiv(5)` man page for more information.

## LONG DESCRIPTION

As described above, it is possible to specify *default flags* within the batch request *script file* that configure the default behavior of the batch request. The algorithm used to scan for such *embedded default flags* within a batch request script file is:

1. Read the first line of the *script file*.
2. If the current line contains only whitespace characters, or the first non-whitespace character of the line is ":", goto step 7.
3. If the first non-whitespace character(s) of the current line is not "#" or "\$!", goto step 8.
4. If the second non-whitespace character in the current line is *not* the "@" character, or the character immediately following the second non-whitespace character in the current line is *not* a "\$", goto step 7.
5. If no "-" is present as the character *immediately* following the "@\$" sequence, goto step 8.
6. Process the *embedded* flag, stopping the parsing process upon reaching the end of the line, or upon reaching the first unquoted "#" or "\$!" character(s).
7. Read the next *script file* line. Goto step 2.
8. End. No more *embedded* flags are recognized.

Here is an example of the use of *embedded* flags within the *script file*.

```

#
# Batch request script example:
#
# @$-a "11:30pm EDT" -lt "21:10, 20:00"
#     # Run request after 11:30 EDT by default,
#     # and set a maximum per-process CPU time
#     # limit of 21 minutes and ten seconds.
#     # Send a warning signal when any process
#     # of the running batch request consumes
#     # more than 20 minutes of CPU time.
# @$-lt 1:45:00
#     # Set a maximum per-process CPU time limit
#     # of one hour, and 45 minutes. (The
#     # implementation of CPU time limits is
#     # completely dependent upon the UNIX
#     # implementation at the execution
#     # machine.)
# @$-mb -me # Send mail at beginning and end of
#           # request execution.
# @$-q batch1 # Queue request to queue: batch1 by
#             # default.
# @$         # No more embedded flags.
#
make all

```

The following paragraphs give detailed descriptions of the *flags* supported by the *qsub* command.

**-a *date-time*** Do not run the batch request before the specified date and/or time. If a *date-time* specification is composed of two or more tokens separated by whitespace characters, the *date-time* specification must be placed within double quotes as in **-a "July, 4, 2026 12:31-EDT"** or otherwise escaped such that *qsub* and the shell will interpret the entire *date-time* specification as a single-character string. This restriction also applies when an embedded default **-a** flag with accompanying *date-time* specification appears within the batch request *script file*.

The syntax accepted for the *date-time* parameter is relatively flexible. Unspecified date and time values default to an appropriate value (e.g., if no date is specified, the current month, day, and year are assumed).

A date may be specified as a month and day (current year assumed), or the year can also be explicitly specified. It is also possible to specify the date as a weekday name (e.g. "*Tues*"), or as one of the strings: "*today*" or "*tomorrow*". Weekday names and month names can be abbreviated by any three-character (or longer) prefix of the actual name. An optional period can follow an abbreviated month or day name.

Time of day specifications can be given using a twenty-four hour clock, or "*am*" and "*pm*" specifications may be used. In the absence of a meridian specification, a twenty-four hour clock is assumed.

It should be noted that the time of day specification is interpreted using the precise meridian definitions whereby "*12am*" refers to the twenty-four hour clock time of 0:00:00, "*12m*" refers to noon, and "*12-pm*" refers to 24:00:00. Alternatively, the phrases "*midnight*" and "*noon*" are accepted as time of day specifications, where "*midnight*" refers to the time of 24:00:00.

A time zone may also appear at any point in the *date-time* specification. Thus, it is legal to say: "*April 1, 1987 13:01-PDT*". In the absence of a time zone specification, the local time zone is assumed, with daylight savings time being inferred when appropriate, based on the date specified.

All alphabetic comparisons are performed in a case insensitive fashion such that both "*WeD*" and "*wed*" refer to the day of Wednesday.

Some valid *date-time* examples are:

```
01-Jan-1986 12am, PDT
Tuesday, 23:00:00
11pm tues.
tomorrow 23:00-MST
```

**-b** [*group*].*activity*

Set the billing activity for request. The *group* and *activity* arguments refer to entries in the */etc/group* and */etc/activities* files respectively. The *group.activity* combination must correspond to an entry in the */etc/actwho* file. If *group* is omitted, the current primary group of the submitting user is assumed. A request always runs under the submitting user's default group. The *group* argument to this option is used only in verifying a user's access to the selected *activity*. See the *bill(1)* man page for more information.

**-c**

Specify that this request is checkpointable. A request queued with this flag is checkpointed automatically before a CXbatch shutdown, and may be explicitly checkpointed using CXbatch commands.

**-cp** *period*

Declare that this request should be checkpointed periodically by CXbatch at intervals of *period*. The *period* is specified as *<[number] unit>*, where *number* is a positive integer and *unit* is [ *Hours* | *Days* | *Weeks* ].

**-e** [*machine*.:][*/path/ stderr-filename*

Direct output generated by the batch request which is sent to the *stderr* file to the named [*machine*.:][*/path/ stderr-filename*

The brackets “[” and “]” enclose optional portions of the *stderr* destination *machine*, *path*, and *stderr-filename*.

If no explicit *machine* destination is specified, the destination machine defaults to the machine that originated the batch request or to the machine that will eventually run the request, depending on the respective absence or presence of the **-ke** flag.

If no *machine* destination is specified, and the path/filename does not begin with a “/”, the current working directory is prepended to create a fully qualified path name, provided that the **-ke** (keep *stderr*) flag is also absent. In all other cases, any partial path/filename is interpreted relative to the user's home directory on the *stderr* destination machine.

This flag cannot be specified when the **-eo** flag option is also present.

If the **-eo** and **-e** [*machine*.:][*/path/ stderr-filename* flag options are not present, all *stderr* output for the batch request is sent to the file whose name consists of the first seven characters of the *request-name* followed by the characters: “.e”, followed by the request sequence number portion of the *request-id* discussed below. In the absence of the **-ke** flag, this default *stderr* output file is placed on the machine that originated the batch request in the current working directory, as defined when the batch request was first submitted. Otherwise, the file is placed in the user's home directory on the execution machine.

**-eo**

Direct all output that would normally be sent to the *stderr* file to the *stdout* file for the batch request. This flag cannot be specified when the **-e** [*machine*.:][*/path/ stderr-filename* flag option is also present.

**-h**

Put request on hold after submitting. The request is put into a HOLD (user hold) rather than a QUEUED state at the time of submittal. The hold can be removed using the *qmgr(8)* ‘RELEASE Request’ command.

**-i**

Some jobs may require access to files located in the directory from which a job is

submitted. This option tells CXbatch that the current working directory should be imported before running this job. If the execution queue is on the same machine as the current working directory, CXbatch will change directories before starting the job. However, if the execution queue is on another machine, CXbatch will import the current directory with NFS. **NOTE:** CXbatch will make temporary NFS mounts into the /tmp filesystem. Care should be taken that any automatic clean-up operations on the /tmp filesystem do not traverse NFS mount points.

**-ke** In the absence of an explicit *machine* destination for the *stderr* file produced by a batch request, the *machine* destination chosen for the *stderr* output file is the machine that originated the batch request. In some cases, however, this behavior may be undesirable, so the **-ke** flag can be specified which instructs CXbatch to leave any *stderr* output file produced by the request on the machine that actually *executed* the batch request.

This flag is meaningless if the **-eo** flag is specified and cannot be specified if an explicit *machine* destination is given for the *stderr* parameter of the **-e** flag.

**-ko** In the absence of an explicit *machine* destination for the *stdout* file produced by a batch request, the *machine* destination chosen for the *stdout* output file is the machine that originated the batch request. In some cases, however, this behavior may be undesirable, and so the **-ko** flag can be specified which instructs CXbatch to leave any *stdout* output file produced by the request on the machine that actually *executed* the batch request.

This flag cannot be specified if an explicit *machine* destination is given for the *stdout* parameter of the **-o** flag.

**-l** The submitted request will be run under a login shell. This was the default behavior in the V1.0 release of CXbatch, but is now only done if explicitly requested. See the discussion above regarding shell start-up files and refer to the appropriate shell man page for details on the differences between login and non-login shells. **NOTE:** Running a job request under a login C-shell will cause the C-shell to issue a warning into the request's output file. This is a function of the C-shell that cannot be suppressed by CXbatch.

**-lx limit-argument**

Set a resource limit. Available limits are summarized in the following table.

Resource Limits		
Limit option	Limited resource	Enforcement
<b>-lc size-limit</b>	per-process corefile size	truncation
<b>-ld size-limit[,warn-limit]</b>	per-process data-segment	request denied
<b>-lf size-limit[,warn-limit]</b>	per-process perm-file	SIGXFSZ
<b>-lf size-limit[,warn-limit]</b>	per-request perm-space	none
<b>-lm size-limit[,warn-limit]</b>	per-process memory	none
<b>-lm size-limit[,warn-limit]</b>	per-request memory	none
<b>-ln nice-value</b>	per-process nice value	setpriority(2)
<b>-ls size-limit[,warn-limit]</b>	per-process stack-segment	request denied
<b>-lt time-limit[,warn-limit]</b>	per-process CPU time	SIGXCPU
<b>-lt time-limit[,warn-limit]</b>	per-request CPU time	none
<b>-lv size-limit[,warn-limit]</b>	per-process temp-file	none
<b>-lv size-limit[,warn-limit]</b>	per-request temp-space	none
<b>-lw size-limit</b>	per-process working set	paging

The *per-process corefile size limit* sets the maximum size of a corefile created by a process. The *per-process data-segment size limit* sets the maximum size to which a process's data-segment can grow. The *per-process perm-file size limit* sets the maximum size to which a permanent file written to by a process can grow. The *per-request perm-file space limit* sets the maximum file space which can be used by permanent files opened for writing by all of the processes in a batch request. The *per-process memory size limit* sets the maximum amount of memory a process can consume. The *per-request memory space limit* sets the maximum amount of memory which can be used by all of the processes in a batch request. The *per-process nice value* sets the nice value for all processes comprising the running batch request. The *per-process stack-segment size limit* sets the maximum size to which a process's stack-segment can grow. The *per-process CPU time limit* sets the maximum amount of CPU time a process can consume. The *per-request CPU time limit* sets the maximum amount of CPU time all of the processes that constitute a batch request can consume. The *per-process temp-file size limit* sets the maximum size to which a temporary file written to by a process can grow. The *per-request temp-file space limit* sets the maximum file space which can be used by temporary files opened for writing by all of the processes in a batch request. The *per-process working set size limit* sets the maximum amount of physical memory each process within a running batch request should use.

Enforcement of limits is done by the underlying UNIX implementation, normally through termination by a signal. Per-process limits are enforced only on the process exceeding the limit; per-request limits are enforced on all the processes which make up a request. The optional warning limits allow warning notification to be made processes where such warnings are supported by the underlying UNIX implementation. Not all UNIX implementations support all the resource limits listed above. The *qsub* command will pass all specified limits on to the destination machine. If a batch request specifies a limit, and the machine upon which the batch request is eventually run does not support the enforcement of that limit, the limit is simply ignored.

The 'Enforcement' column of the table above indicates the ConvexOS specific

enforcement policy of each resource limit. At this time, all *warn-limits* are treated identically to the hard limits. ConvexOS makes no distinction between permanent and temporary files; all files are treated as permanent.

See the section entitled **LIMITS** for more information on the implementation of batch request limits and for a description of the precise syntax of a the various limit arguments.

- mb** Send mail to the user on the originating machine when the request begins execution. If the **-mu** flag is also present, mail is sent to the user specified for the **-mu** flag instead of to the invoking user.
- me** Send mail to the user on the originating machine when the request has ended execution. If the **-mu** flag is also present, mail is sent to the user specified for the **-mu** flag instead of to the invoking user.

**-mu** *user-name*

Specify that any mail concerning the request should be delivered to the user *user-name*. *user-name* may be formatted either as *user* (containing no '@' characters), or as *user@machine*. In the absence of this flag, any mail concerning the request is sent to the invoker on the originating machine.

- ni** A queue can be configured so that it tries to import the originating directory of each job it runs. If this option is specified, CXbatch does not import the current directory.

- nc** Advise CXbatch that this request is not checkpointable. A request queued with this option will not be checkpointed at CXbatch shutdown, nor is it possible to checkpoint the request with any other CXbatch commands.

- nr** Declare that the request is non-restartable. If this flag is specified, the request is not restarted by CXbatch upon system boot if the request was running at the time of a CXbatch shutdown or system crash.

By default, CXbatch assumes that all requests are restartable, with the caveat that it is the responsibility of the user to ensure that the request will execute correctly if restarted, by use of appropriate programming techniques.

Requests that are not running are always preserved across host crashes and CXbatch shutdowns for later requeuing, with or without this flag.

When CXbatch is shutdown by an operator command to the *qmgr*(8) CXbatch control program, a **SIGTERM** signal is sent to all processes in all running CXbatch requests on the local host, and all queued CXbatch requests are barred from beginning execution. After a finite number of seconds have elapsed (typically sixty, but this value can be overridden by the operator), all remaining processes in all remaining running CXbatch requests are killed by the signal **SIGKILL**.

For a CXbatch request to be properly restarted after a CXbatch shutdown, the **-nr** flag must not be specified, and the spawned batch request shell must ignore **SIGTERM** signals (which is done by default). The spawned batch request shell also must not exit before the final **SIGKILL** arrives. Because the batch request shell is spawning commands and programs, waiting for their completion, this implies that the commands and programs being executed by the batch request shell must also be immune to **SIGTERM** signals, saving state as appropriate before being killed by the final **SIGKILL** signal.

See the **CAVEATS** section below for more discussion about the restartability of batch requests.

- o** [*machine*.:][[/]*path*/] *stdout-filename*

Direct output generated by the batch request which is sent to the *stdout* file to the named `[machine:][[/path/] stdout-filename`

The brackets “[” and “]” enclose optional portions of the *stdout* destination *machine*, *path*, and *stdout-filename*.

If no explicit *machine* destination is specified, the destination machine defaults to the machine that originated the batch request or to the machine that will eventually run the request, depending on the respective absence or presence of the `-ko` flag.

If no *machine* destination is specified, and the path/filename does not begin with a “/”, the current working directory is prepended to create a fully-qualified path name, provided that the `-ko` (keep *stdout*) flag is also absent. In all other cases, any partial path/filename is interpreted relative to the user’s home directory on the *stdout* destination machine.

If no `-o [machine:][[/path/] stdout-filename` flag is specified, all *stdout* output for the batch request is sent to the file whose name consists of the first seven characters of the *request-name* followed by the characters: “.o”, followed by the request sequence number portion of the *request-id* discussed below. In the absence of the `-ko` flag, this default *stdout* output file is placed on the machine that originated the batch request in the current working directory, as defined when the batch request was first submitted. Otherwise, the file is placed in the user’s home directory on the execution machine.

**-p *priority*** Assign an *intra-queue* priority to the request. The specified *priority* must be an integer, and must be in the range [0..63], inclusive. A value of 63 defines the highest *intra-queue* request priority, while a value of 0 defines the lowest. This priority does *not* determine the execution priority of the request. This priority is only used to determine the relative ordering of requests within a queue.

When a request is added to a queue, it is placed at a specific position within the queue such that it appears ahead of all existing requests whose priority is less than the priority of the new request. Similarly, all requests with a higher priority remain ahead of the new request when the queuing process is complete. When the priority of the new request is equal to the priority of an existing request, the existing request takes precedence over the new request.

If no *intra-queue* priority is chosen by the user, CXbatch assigns a default value.

The CXbatch manager can assign a maximum request priority on a per queue basis. The maximum request priority is a ceiling on priorities of requests submitted to that queue. A request that specifies a priority higher than the maximum for that queue has its priority lowered to the maximum.

**-q *queue-name*[@ *host*]**

Specify the queue to which the batch request is to be submitted. If no `-q queue-name[@ host]` specification is given, the user’s environment variable set is searched for the variable `QSUB_QUEUE`. If this environment variable is found, the character string value for `QSUB_QUEUE` is presumed to name the queue to which the request should be submitted. If the `QSUB_QUEUE` environment variable is not found, the request is submitted to the default batch request queue, *if* defined by the local system administrator. Otherwise, the request cannot be queued, and an appropriate error message is displayed. The *host* specifies the host where the queue resides. If no *host* is given, the local host is assumed. Not all hosts accept remote submissions.

**-r *request-name***

Assign the specified *request-name* to the request. In the absence of an explicit `-r`

*request-name* specification, the *request-name* defaults to the name of the *script file* (leading path name removed) given on the command line. If no *script file* was given, the default *request-name* assigned to the request is **STDIN**.

In all cases, if the *request-name* is found to begin with a digit, the character "R" is prepended to prevent a *request-name* from beginning with a digit. All *request-names* are truncated to a maximum length of 15 characters.

**-s** *shell-name*

Specify the absolute path name of the shell that is used to interpret the batch request script. This flag unconditionally overrides any *shell strategy* configured on the execution machine for selecting which shell to spawn in order to interpret the batch request script.

In the absence of this flag, the CXbatch system at the execution machine uses one of three distinct *shell choice strategies* for the execution of the batch request. Any one of the three strategies can be configured by a system administrator for each CXbatch machine.

The three shell strategies are called:

- *fixed*
- *free*
- *login*

These shell strategies respectively cause the configured *fixed* shell to be exec'd to interpret all batch requests; cause the user's login shell as defined in the password file to be exec'd, which in turn chooses and spawns the appropriate shell for interpreting the batch request script; or cause only the user's login shell to be exec'd to interpret the script.

A shell strategy of *fixed* means that the same shell (as configured by the system administrator) is used to execute all batch requests.

A shell strategy of *free* runs the batch request script *exactly* as would an interactive invocation of the script and is the default CXbatch shell strategy.

The strategies of *fixed* and *login* exist for host systems that are short on available free processes. In these two strategies, a single shell is exec'd, and that same shell is the shell that executes all of the commands in the batch request script.

The *shell strategy* configured for a particular CXbatch system can be determined by the *qlimit(1)* command.

**-t** *process-id* Signal process *process-id* when the job completes execution. The signal sent depends on how the job completed. If the job runs to normal completion, **SIGTERM** is sent. If the job is aborted while running, **SIGUSR1** is sent. If the job is deleted before it starts executing, **SIGUSR2** is sent.

**-x** Export all environment variables. When a batch request is submitted, the current values of the environment variables **HOME**, **SHELL**, **PATH**, **USER**, and **MAIL** are saved for later recreation when the batch request is spawned, as the respective environment variables **QSUB\_HOME**, **QSUB\_SHELL**, **QSUB\_PATH**, **QSUB\_USER**, and **QSUB\_MAIL**. Unless the **-x** flag is specified, no other environment variables are exported from the originating host for the batch request. If the **-x** flag option is specified, all remaining environment variables whose names do not conflict with the automatically exported variables are also exported with the request. These additional environment variables are recreated under the same name when the batch request is spawned.

**-y** Append accounting data to the stdout output file if accounting is enabled for the

queue in which the job runs. This data includes: queue, host, sequence number, remote host, submission time, start time, completion time, time spent executing in user mode, and time spent executing in the system.

- z Submit the batch request silently. If the request is submitted successfully, no messages are displayed indicating this fact. Error messages are, however, always displayed.

If the batch request is successfully submitted, and the -z flag has not been specified, the *request-id* of the request is displayed to the user. A *request-id* is always of the form *seqno.hostname*, where *seqno* refers to the sequence number assigned to the request by CXbatch, and *hostname* refers to the name of originating local machine. This identifier is used throughout CXbatch to uniquely identify the request, no matter where it is in the network.

The following events take place in the following order when a CXbatch *batch* request is spawned:

1. The process that will become the head of the *process group* for all processes comprising the batch request is created by CXbatch.
2. Resource limits are enforced.
3. The real and effective group-id of the process is set to the group-id as defined in the local password file for the request owner.
4. The real and effective user-id of the process is set to the real user-id of the batch request owner.
5. The user file creation mask is set to the value that the user had on the originating machine when the batch request was first submitted.
6. If the user explicitly specified a shell by use of the -s flag (discussed above), then that user-specified shell is chosen as the shell that will be used to execute the batch request script. Otherwise, a shell is chosen based upon the *shell strategy* as configured for the local CXbatch system. (See the earlier discussion of the -s flag for a description of the possible *shell strategies* that can be configured for a CXbatch system.)
7. The environment variables of HOME and SHELL, are set from the user's password file entry, as though the user had logged directly into the execution machine.
8. The environment string: ENVIRONMENT=BATCH is added to the environment so that shell scripts (and the user's *.profile* (Bourne shell) or *.cshrc* and *.login* (C-shell) scripts) can test for batch request execution when appropriate, and not (for example) perform any setting of terminal characteristics, because a batch request is not connected to an input terminal.
9. The environment variables of QSUB\_WORKDIR, QSUB\_HOST, QSUB\_REQNAME, and QSUB\_REQID are added to the environment. These environment variables equate to the obvious respective strings of the working directory at the time that the request was submitted, the name of the originating host, the name of the request, and the request *request-id*.
10. All of the remaining environment variables saved for recreation when the batch request is spawned are added at this point to the environment. When a batch request is initially submitted, the current values of the environment variables HOME, SHELL, PATH, USER, and MAIL are saved for later recreation when the batch request is spawned. When recreated however, these variables are added to the environment under the respective names QSUB\_HOME, QSUB\_SHELL, QSUB\_PATH, QSUB\_USER, and QSUB\_MAIL to avoid the obvious conflict with the local version of these environment variables. Additionally, all environment variables exported from the originating host by the -x option are added to the

environment at this time.

11. The current working directory is then set to the user's home directory on the execution machine, and the chosen shell is exec'd to execute the batch request script with the environment as constructed in the steps outlined above.

Unless the `-l` option is specified, the chosen shell is exec'd as a non-login shell and no start-up files (except the `~/cshrc` for a C-shell) are read. If the `-l` option is selected, the chosen shell is exec'd as though it were the *login* shell. If the Bourne shell is chosen to execute the script, `/etc/profile` and the user's `.profile` file are read. If the C-shell is chosen, `/etc/login` and the user's `.cshrc` and `.login` scripts are read and when the job exits, `/etc/logout` and the user's `.logout` script are read.

If the user did not specify a shell for the batch request, CXbatch chooses which shell is used to execute the shell script, based on the *shell strategy* as configured by the system administrator. (See the earlier discussion of the `-s` flag.)

In such a case, a *free* shell strategy instructs CXbatch to execute the login shell for the user (as configured in the password file). The login shell is in turn instructed to examine the shell script file and fork another shell of the *appropriate type* to interpret the shell script, behaving *exactly* as an interactive invocation of the script.

Otherwise no additional shell is spawned, and the chosen *fixed* or *login* shell sequentially executes the commands contained in the shell script file until completion of the batch request.

## QUEUE TYPES

CXbatch supports three different queue types that serve to provide three very different functions. These three queue types are known as *batch*, *pipe*, and *network*.

The queue type of *batch* can only be used to execute CXbatch *batch requests*. Only CXbatch *batch requests* created by the `qsub(1)` command can be placed in a *batch queue*.

Queues of type *pipe* are used to send CXbatch requests to other *pipe* queues or to request destination queues of type *batch*. In general, *pipe queues*, in combination with *network queues*, act as the mechanism that CXbatch uses to transport *batch* requests to distant queues on remote machines. It is also legal for a *pipe queue* to transport requests to queues on the *same* machine.

When a *pipe queue* is defined, it is given a *destination set* that defines the set of possible destination queues for requests entered in that *pipe queue*. In this manner, it is possible for a *batch* request to pass through many pipe queues on its way to its ultimate destination, which must eventually be a queue of type *batch*.

Each *pipe queue* has an associated *server*. For each request handled by a *pipe queue*, the associated server is spawned which must select a queue destination for the request being handled, based upon the characteristics of the request and upon the characteristics of each queue in the *destination set* defined for the pipe queue.

Because a different server can be configured for each *pipe* queue, and *batch* queues can be endowed with the *pipeonly* attribute that will only admit requests queued via another *pipe queue*, it is possible for respective CXbatch installations to use *pipe queues* as a *request class* mechanism, placing requests that ask for different resource allocations in different queues, each of which can have different associated limits and priorities.

It is also completely possible for a *pipe queue server*, when handling a request, to discover that no *destination queue* will accept the request, for various reasons that can include insufficient resource limits to execute the request or a lack of a corresponding account or privilege for queuing at a remote queue. In such circumstances, the request is deleted, and the user is notified by mail (see `mail(1)`).

The queue type of *network*, as alluded to earlier, is implicitly used by *pipe* queues to pass CXbatch requests between machines and is also used to handle queued file transfer operations.

## QUEUE ACCESS

CXbatch supports queue access restrictions. For each queue of queue type other than *network*, access may be either *unrestricted* or *restricted*. If access is *unrestricted*, any request may enter the queue. If access is *restricted*, a request can only enter the queue if the requester or the requester's login group has been given access to that queue (see *qmgr(8)*). Requests submitted by the superuser are an exception; they are always queued, even if the superuser has not explicitly been given access.

Use *qstat(1)* to determine who has access to a particular queue.

## LIMITS

CXbatch supports many batch request resource limit types that can be applied to a batch request. The existence of configurable resource limits allows a CXbatch user to set resource limits within which his or her request must execute.

The syntax used to specify a *limit-value* for one of the *limit-flags* (*-lx*), is quite flexible and describes values for three general limit categories. These three general categories respectively deal with *time limits*, *priority (nice value) limits* and *file/memory size limits*.

All the *limit-flags* expect a single *limit-argument*. If the *limit-argument* consists any whitespace which will cause the it to be passed to *qsub* as multiple tokens, it should be enclosed within double quotes or otherwise escaped such that is is passed a a single, character-string token. The optional *warn-limit* should also be included as part of the same token. This also applies to *limit-arguments* associated with *limit-flags* embedded within the batch request *script file*.

For *finite* CPU time limits, the *limit-value* is expressed in the reasonably obvious format:

```
[[hours :] minutes : ] seconds [.milliseconds]
```

Whitespace can appear anywhere between the principal tokens, with the exception that no whitespace can appear around the decimal point. **NOTE:** The *milliseconds* value may be ignored if the system on which the request is run does not support such granularity.

Example time *limit-values* are:

```
1234 : 58 : 21.29   - 1234 hrs 58 mins 21.290 secs
12345              - 12345 seconds
121.1             - 121.100 seconds
59:01            - 59 minutes and 1 second
```

Priority limits are expressed as a small, signed integer in the range acceptable for *nice values* on the execution machine. In general, increasingly negative *nice values* cause the relative execution priority of a process to increase, while increasingly positive *nice values* causes the relative priority to decrease. Because different UNIX implementations often support different finite ranges of *nice values*, CXbatch allows the specification of *nice values* that can eventually turn out to be outside the limits for the UNIX implementation running at the execution machine. In such cases, CXbatch simply binds the specified *nice values* limit to within the necessary range as appropriate.

For the *finite* size limits the acceptable syntax is:

```
.fraction [units]
```

or

```
integer [.fraction] [units]
```

where the *integer* and *fraction* tokens represent strings of up to eight decimal digits, denoting the obvious values. In both cases, the *units* of allocation may also be specified as one of the case

insensitive strings:

<i>b</i>	- bytes
<i>w</i>	- words
<i>kb</i>	- kilobytes ( $2^{10}$ bytes)
<i>kw</i>	- kilowords ( $2^{10}$ words)
<i>mb</i>	- megabytes ( $2^{20}$ bytes)
<i>mw</i>	- megawords ( $2^{20}$ words)
<i>gb</i>	- gigabytes ( $2^{30}$ bytes)
<i>gw</i>	- gigawords ( $2^{30}$ words)

In the absence of any *units* specification, the units of *bytes* are assumed.

For all limit types with the exception of the *nice* limit-value ( $-ln$ ), it is possible to state that no limit should be applied. This is done by specifying a *limit-value* of "unlimited" or any initial substring thereof. Whenever an *infinite limit-value* is specified for a particular resource type, the batch request operates as though no explicit limits have been placed upon the corresponding resource, other than by the limitations of the physical hardware involved.

The complications caused by *batch request* resource limits first show up when queuing a *batch request* in a *batch queue*. This operation is described in the following paragraphs.

If a batch request specifies a limit that cannot be enforced by the underlying UNIX implementation, the limit is simply ignored, and the batch request operates as though there were no limit (other than the obvious physical maximums) placed upon that resource type. (See the *qlimit(1)* command to find out what limits are supported by a given machine.)

For each remaining *finite* limit that can be supported by the underlying UNIX implementation that is *not* a CPU *time-limit* or UNIX *execution-time nice-value-limit*, the *limit-value* is internally converted to the units of *bytes* or *words*, whichever is more appropriate for the underlying machine architecture.

As an example, a *per-process memory size limit value* of 321 megabytes would be interpreted as  $321 \times 2^{20}$  bytes, provided that the underlying machine architecture was capable of directly addressing single bytes. Thus the original limit *coefficient* of 321 would become  $321 \times 2^{20}$ . On a machine that was only capable of addressing words, the appropriate conversion of  $321 \times 2^{20}$  bytes / #of-bytes-per-word would be performed.

If the result of such a conversion would cause overflow when the coefficient was represented as a *signed-long integer* on the supporting hardware, the coefficient is replaced with the coefficient of  $2^{N-1}$ , where  $N$  is equal to the number of bits of precision in a signed long integer. For typical 32-bit machines, this *default extreme limit* would therefore be  $2^{31-1}$  bytes. For word addressable machines in the supercomputer class supporting 64-bit long integers, the *default extreme limit* would be  $2^{63-1}$  words.

Lastly, some implementations of UNIX reserve coefficients of the form:  $2^{N-1}$  as synonymous with infinity, meaning no limit is to be applied. For such UNIX implementations, CXbatch further decrements the *default extreme limit* so as not to imply infinity.

The identical internal conversion process as described in the preceding paragraphs is also performed for each *finite limit-value* configured for a particular batch queue using the *qmgr(8)* program.

After all of the applicable *limit-values* have been converted as described above, each resulting *limit-value* is then compared against the corresponding *limit-value* as configured for the destination batch queue. If, for every type of limit, the batch queue *limit-value* is *greater than or equal to* the corresponding batch request *limit-value*, the request can be successfully queued, provided that no other anomalous conditions occur. For request *infinity limit-values*, the corresponding queue *limit-value* must also be configured as infinity.

These resource limit checks are performed irrespective of the batch request arrival mechanism, either by a direct use of the *qsub(1)* command or by the indirect placement of a batch request into a batch queue via a *pipe* queue. It is impossible for a batch request to be queued in a batch queue if *any* of these resource limit checks fail.

Finally, if a request fails to specify a *limit-value* for a resource limit type that is supported on the execution machine, the corresponding *limit-value* configured for the destination queue becomes the *limit-value* for the unspecified request limit.

Upon the successful queuing of a request in a batch queue, the set of limits under which the request will execute is frozen and will not be modified by subsequent *qmgr(8)* commands that alter the limits of the containing batch queue.

## CAVEATS

When a batch request is spawned, a new *process-group* is established such that all processes of the request exist in the same *process-group*. If the *qdel(1)* command is used to send a signal to a batch request, the signal is sent to all processes of the request in the created *process-group*. However, should one or more processes of the request choose to successfully execute a *setpgp(2)* system call, such processes will *not* receive any signals sent by the *qdel(1)* command. This can lead to “rogue” requests whose constituent processes must be killed by other means such as the *kill(1)* command.

It is extremely wise for all processes of a CXbatch request to catch any SIGTERM signals. By default, the receipt of a SIGTERM signal causes the receiving process to die. CXbatch sends a SIGTERM signal to all processes in the established *process-group* for a batch request as a notification that the request should be prepared to be killed, either because of an *abort queue* command issued by an operator using the *qmgr(8)* program, or because it is necessary to shutdown CXbatch and all running requests as part of a general shutdown procedure of the local host.

It must be understood that the spawned *shell* ignores SIGTERM signals. If the current immediate child of the shell does not ignore or catch SIGTERM signals, it will be killed by the receipt of such, and the shell will go on to execute the next command from the script (if there is one). In any case, the shell will not be killed by the SIGTERM signal, though the executing command will have been killed.

After receiving a SIGTERM signal delivered from CXbatch, a process of a batch request typically has sixty seconds to get its “house in order” before receiving a SIGKILL signal (though the sixty second duration can be changed by the operator).

All batch requests terminated because of an operator CXbatch shutdown request that did not specify the *-nr* flag are considered restartable by CXbatch and are requeued (provided that the batch request shell process is still present at the time of the SIGKILL signal broadcast as discussed above), so that when CXbatch is rebooted, such batch requests will be respawned to continue execution. It is, however, up to the user to make the request restartable by the appropriate programming techniques. CXbatch simply spawns the request again as though it were being spawned for the first time.

Upon completion of a batch request, a mail message can be sent to the submitter (see the discussion of the *-me* flag above). In many instances, the completion code of the spawned Bourne or C-Shell is displayed. This is the value returned by the shell through the *exit(2)* system call.

Lastly, there is no good way to echo commands executed by unmodified versions of the Bourne and C shells. While the Bourne and C shells can be spawned in such a fashion as to echo the commands they execute, it is often very difficult to tell an echoed command from genuine output produced by the batch request.

Thus, one of the better ways to write the shell script for a batch request is to place appropriate lines in the shell script of the form:

echo "*explanatory-message*"

where the echoed message should be a meaningful message chosen by the user.

#### DIAGNOSTICS

*qsub* returns an exit status describing what it did. If there were no errors, the exit status is zero. If a fatal error occurs and the request is not submitted (e.g., a syntax error), the exit status is one of the codes defined in *<sysexit.h>*, or one if none of the *<sysexit.h>* codes are appropriate.

- EX\_USAGE      The command was used incorrectly, e.g., with the wrong number of arguments, a bad flag, a bad syntax in a parameter.
- EX\_OSFILE     Some batch system file does not exist, cannot be opened, or has an error.
- EX\_TEMPFAIL   Temporary failure; retry the request at a later time.
- EX\_NOPERM     You did not have sufficient permission to perform the operation.
- EX\_NOINPUT    The script file does not exist or is not readable.

#### SEE ALSO

bill(1), mail(1), qdel(1), qlist(1), qlimit(1), qstat(1), qmgr(8) kill(2), setpgrp(2), signal(3c)

#### NOTES

CXbatch is an optional product; for more information, contact your CONVEX sales representative.

**NAME**

*qwatch* - watch status of CXbatch queue(s)

**SYNOPSIS**

**qwatch** [-i *interval*] [-c *count*]

**DESCRIPTION**

*qwatch* monitors the status of CONVEX CXbatch and periodically reports statistics about queues and/or requests.

Initially, *qwatch* displays queue headers for the first five queues. If there are more than five queues, additional pages can be displayed by typing the number of the desired page. Only five pages (25 queues) can be display by *qwatch*.

By typing the letter (a-y) corresponding to a queue, you can monitor the activity of that queue. The first page of the requests within the selected queue is displayed along with the queue header. If there are more that 16 requests in the queue, additional pages can be displayed by typing the number of the desired page. Only nine pages (144 requests) can be displayed by *qwatch*.

The command line switches, which may be given in any order, are:

- c *count*        If *count* is positive, *qwatch* automatically exits after *count* screen updates have been completed. If *count* is 0, *qwatch* keeps updating until stopped. The default value for *count* is 0.
- i *interval*     Specifies how many seconds to wait between each screen update. The default is five seconds.

The following keys are interpreted by *qwatch*:

- <SPACE>
- <RETURN>       Force *qwatch* to update the screen and restart the timer.
- <CTRL-L>       Replot the screen.
- [1-9]            Select a display page. (Only pages with information on them may be selected.)
- [a-y]            Select a queue. (Only available from queues display page.)
- Return to queues display page from requests display page.

The *SIGINT* signal, usually generated by <CTRL-C> or <DEL>, is used to exit *qwatch*. Additionally, *qwatch* can be interrupted and later resumed with the *SIGTSTP* signal, which is typically generated by <CTRL-Z> from *csh*.

**QUEUE STATE INFORMATION**

Refer to the *qstat(1)* manpage for an explanation of the information displayed by *qwatch*.

**DIAGNOSTICS**

*qwatch* returns an exit status describing what it did. If there were no errors, the exit status is zero. If a fatal error occurs, then the exit status is one of the codes defined in <*syszeits.h*>.

**EX\_UNAVAILABLE**

The terminal does not support the capabilities necessary for *qwatch* to work properly.

**EX\_USAGE**

The command was used incorrectly, e.g., with the wrong number of arguments, a bad flag, a bad syntax in a parameter.

**SEE ALSO**

*qstat(1)*, *syspic(8)*

**NOTES**

CXbatch is an optional product; for more information, contact your CONVEX sales representative.

---

# Index

---

## A

assistance xvi  
associated documents xv

---

## B

batch queue 1

---

## C

Checkpoint Restart  
and CXbatch 8

checkpointing

after submitting 66, 68  
overriding default 64, 65  
preventing restart 65  
restarting request 70  
resuming request 69  
submitting request 64, 65  
suspending request 69

commands

changing request priority 62  
checkpointing request 66, 68  
deleting queue request 58, 59  
general user qmgr 2  
manager qmgr 5  
moving request 61  
operator qmgr 4  
overriding checkpoint default 64, 65  
placing request on hold 60  
preventing restart 65  
removing hold on request 60  
restarting request 70  
resuming request 69  
submitting request 32, 33, 34  
submitting with checkpointing 64, 65  
suspending request 69  
viewing process status 28  
viewing queue limits 26

viewing queue status 10, 21  
viewing shell script contents 30  
COVUEbatch  
and CXbatch 7  
CXbatch  
and Checkpoint Restart 8  
and COVUEbatch 7  
and Share Scheduler 7  
incompatibilities with NQS 6

---

## E

embedded options 54

---

## H

help xvi

---

## I

information, supplemental xv

---

## M

man pages 113  
messages  
request completion 99  
transaction completion 71

---

## N

notational conventions xiv  
NQS  
incompatibilities with CXbatch 6

---

## O

ordering documents xv  
output  
  qjlist 30  
  qlimit 26  
  qps 29  
  qstat 10, 14, 16, 17  
  qwatch 22

---

## P

pipe queue 1  
Problems, reporting 71  
purpose of document xiii

---

## Q

qmgr utility  
  general user commands 2  
  manager commands 5  
  operator commands 4  
qsub command options  
  controlling resource limits 48  
  controlling run requests 34, 35, 36, 38, 40, 41, 42, 43  
  redirecting output files and error messages 44, 45, 46, 47  
  sending mail 50, 51, 52  
  signalling processes 53  
qsub commands  
  overriding checkpoint default 64, 65  
  preventing restart 65  
  submitting with checkpointing 64, 65  
queue  
  batch 1  
  pipe 1  
  viewing information 17  
  viewing status interactively 21  
  viewing status of 10  
queue request  
  changing priority 62  
  checkpointing 66, 68  
  controlling resource limits 48  
  controlling run options 35, 36, 38, 40, 41, 42, 43  
  controlling run requests 34  
  deleting 58, 59  
  moving 61  
  overriding checkpoint default 64, 65  
  placing on hold 60  
  preventing restart 65  
  redirecting output files and error messages 44, 45, 46, 47  
  removing hold 60  
  restarting 70  
  resuming 69

---

  sending mail 50, 51, 52  
  signalling processes 53  
  submitting 31, 32, 33, 34  
  submitting with checkpointing 64, 65  
  suspending 69  
  viewing information on 14

---

## R

Reporting problems 71  
request completion messages 99

---

## S

Share Scheduler  
  and CXbatch 7

---

## T

TAC xvi  
technical assistance xvi  
Technical Assistance Center xvi  
transaction completion messages 71  
typographic conventions xiv

---

## U

using this book xiii

---


## V

viewing  
  future run time 16  
  process status 29  
  queue information 17  
  queue limits 26  
  queue request information 14  
  queue status 10, 22  
  shell script contents 30

---



**Order Number**  
**DSW-183**



**Document Number**  
**710-002730-206**