

Release Notice
CONVEX CXbatch V1.1
Document No. 710-007830-001

April 1990

CONVEX Computer Corporation
Richardson, Texas USA

Release Notice, CONVEX CXbatch V1.1

Copyright 1990 CONVEX Computer Corporation

All rights reserved.

This document is copyrighted. This document, however, may be copied, duplicated, reproduced, translated, stored electronically, 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 (CONVEX) 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 OR ANY THIRD PARTY.

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

ConvexOS is a trademark of CONVEX Computer Corporation.

CONVEX CXbatch is a trademark of CONVEX Computer Corporation.

CONVEX Share Scheduler is a trademark of CONVEX Computer Corporation.

COVUE is a registered trademark of CONVEX Computer Corporation. COVUE consists of the following products: COVUEbatch, COVUEbinary, COVUEedt, COVUElib, COVUenet, and COVUEshell.

NFS is a trademark of Sun Microsystems, Inc.

UNIX is a registered trademark of AT&T Bell Laboratories.

Printed in the United States of America

Table of Contents

1 Release Notice	
1.1 Introduction	1-1
1.2 Contents of This Distribution	1-2
1.3 Notes and Cautions	1-3
1.4 Enhancements	1-5
1.5 Fixes	1-5
1.6 Known Software Problems	1-13
1.7 Documentation	1-13
1.8 Known Documentation Problems	1-14

List of Tables

1-1 Software for USA/Int'l Distribution	1-2
1-2 Documentation Release Package for Initial Installation	1-2

Release Notice

1.1 Introduction

This document describes the V1.1 Release of CONVEX CXbatch.

CXbatch allows users to submit non-interactive jobs for delayed execution on either the local machine or a remote machine in the CXbatch network. Queues can be configured such that jobs run only at night when the machine is lightly loaded, or they can be configured so that jobs are routed to the machine with the lightest load.

CXbatch is based on NASA's Network Queuing System (NQS). Many CONVEX enhancements have been added, including load balancing, batch accounting, automatic importing of files (via Network File System), and direct submission to remote machines.

The CXbatch package includes:

- *CONVEX CXbatch Release Notice*
- *CONVEX CXbatch Installation Procedure*
- *CONVEX CXbatch Concepts Manual*
- *CONVEX CXbatch User's Guide*
- *CONVEX CXbatch Programmer's Reference Manual*
- *CONVEX CXbatch System Manager's Guide*
- *CONVEX CXbatch System Manager Utilities Reference Manual*
- *CONVEX CXbatch Master Index*
- CONVEX CXbatch Distribution Tape

Always refer to this release notice before reporting questions or problems with CXbatch; your questions may be answered here. This release notice also lists fixes and workarounds that may save you time if you encounter a known problem.

An on-line copy of this release notice can be found in */usr/doc/v1.1cxbatch.relnotes*.

The remaining sections in this document describe the contents of this release.

- Section 1.2 describes the contents of this distribution
- Section 1.3 contains notes and cautions about use of the software
- Section 1.4 contains enhancements to previous functionality
- Section 1.5 describes fixes for previously reported problems
- Section 1.6 describes known software problems
- Section 1.7 describes the release documentation
- Section 1.8 contains corrections to the documentation

To install the software described in this document, consult the *Installation Procedures, CONVEX CXbatch V1.1* which was included with the CXbatch distribution.

1.2 Contents of This Distribution

The distribution package for this release of CONVEX CXbatch consists of:

- this release notice
- distribution media for the software
- installation instructions
- documentation

These items are described in the following tables.

Table 1-1: Software for USA/Int'l Distribution

Qty	Type	Part Number	Description
1	Mag Tape	710-003915-203	CONVEX CXbatch, installsw format

Table 1-2: Documentation Release Package for Initial Installation

Qty	Type	Part Number	Description
1	Document	710-007830-001	CONVEX CXbatch Release Notice
1	Document	710-007930-001	CONVEX CXbatch Installation Procedures
1	Document	710-006630-001	CONVEX CXbatch Concepts
1	Document	710-002730-203	CONVEX CXbatch User's Guide
1	Document	710-004430-002	CONVEX CXbatch Programmer's Reference Manual
1	Document	710-006730-001	CONVEX CXbatch System Manager's Guide
1	Document	710-006830-001	CONVEX CXbatch System Manager Utilities Reference Manual
1	Document	710-007630-001	CONVEX CXbatch Master Index

1.3 Notes and Cautions

This section contains useful information and words of caution about the product.

- The following are the prerequisites for this release of CXbatch:
 - Your system *must* be running V8.0 or later ConvexOS and Utilities.
 - Your system *must* be running V8.0 or later CONVEX Networking Utilities *if* you want to use the load balancing option, the automatic NFS import option, or the COVUE products.
 - Your system *must* be running V8.0 or later CONVEX NFS Utilities *if* you want to use the automatic NFS import option or the COVUE products.
- If the CONVEX Share Scheduler is installed on your system, you must create entries in */etc/passwd* for each batch queue to be created and assign shares for these entries. This is normally done using the *nu(8)* utility. This must be done for the default batch queues **BEFORE** installing the default configuration from the installsw tape. See the appropriate sections of the *Share Scheduler* and *CXbatch System Managers Guides*.
- The ConvexOS V8.0 release included special versions of *qmgr(8)* and *nqsdaemon(8)*. These have been superceded by this release of CXbatch. The */etc/rc* script included in the ConvexOS V8.0 contained about 20 lines of commands which kept these two modules up-to-date. These commands are no longer necessary and should be **REMOVED** or **COMMENTED OUT**. **FAILURE TO DO THIS COULD RESULT IN IMPROPER BEHAVIOR OF CXBATCH**. The lines in question begin with the comment '# attempt to recover 8.0 nqsdaemon and qmgr' and consist of two 'if' statements, each containing a nested 'if-then-else' statement. You should **carefully** edit your */etc/rc* file at this time.
- The NQS implementation supports a number of resource limits. CXbatch can only enforce the resource limits supported by the underlying ConvexOS operating system. The actual limits supported by CXbatch, or any other NQS batch implementation, is reported by the *qlimit* command or by the *qmgr* 'show limits_supported' command. All the standard NQS resource limits are accepted by the CXbatch *qsub* command and will be passed by CXbatch to the destination queue. If a requested resource limit is supported by the destination system, it will be enforced at run time. Otherwise, it will be silently ignored.

Some resource limits in CXbatch are supported only to the extent possible. For example, *cpu* time limits may be entered down to milliseconds with *qsub*, but because ConvexOS only supports granularity in seconds, milliseconds are ignored.

- When a request is selected to run, a shepherd process starts it by executing a top level shell which runs the submitted command script. When this top level shell process exits, the request is considered completed and the shepherd process cleans up by killing any remaining child processes started by the top level shell. This means that requests should not run any background commands without waiting for them to complete using the applicable shell commands.

One side effect of this behavior has been found when using the *mail(1)* or *binmail(1)* from within a CXbatch request. If the addressee of a mail message is on the local (execution) system, a sub-process is forked by mail to deliver the mail message and the mail command itself exits. If the CXbatch request completes before mail has been delivered, the delivery process may be killed by the shepherd process. This causes the mail to disappear without a trace. The work around for this situation is the use the *-v* option to *mail(1)* or sleep for a short period of time to allow the mail to be delivered.

- By default, CXbatch runs a job using a non-login shell. This can be overridden with

the '-l' option to qsub. If your job runs under a login shell, the */etc/login* (C shell) or */etc/profile* (Bourne shell and KornShell) file and your *.login* (C shell) or *.profile* (Bourne shell and KornShell) will be read. A job running using the C shell will also read */etc/logout* and *.logout*. The string `ENVIRONMENT=BATCH` is added to the environment so that shell scripts can test for batch request execution. With this variable, shell startup scripts (such as *.profile*, *.login*, and *.cshrc*) can test for batch request execution and not, for example, perform any setting of terminal characteristics.

For example, if your login shell is C shell, the following *.login* file prevents *stty*, *tset*, and *msgs* from running during batch jobs. **NOTE:** C shell will always read your *.cshrc* regardless of whether or not it is running as a login shell, so these precautions should be taken.

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

If your login shell is Bourne shell or KornShell, the following *.profile* file has the same effect.

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

- The CXbatch load balancing pipe client, *pipeldav*, works with CXbatch pipe queue destinations but not with non-CONVEX NQS destinations.
- When COVUE is used with CXbatch, certain COVUE commands cannot obtain queue and job information from pipe queue destinations that are running non-CONVEX NQS. There should not be any problems obtaining information with CXbatch destination queues.
- Although CXbatch is based on NQS, some NQS features are not supported:
 - device queues
 - queue complexes
 - *set global run limit* command
- The *list*, *submit*, *sbatch*, *ssp*, *move*, *remove*, *hold*, and *unhold* commands from the CONVEX Distributed Batch System are provided to ease the transition to CXbatch. These commands will not be included in future releases of CXbatch. They accept the same options that they did with the previous batch system, but some of the options are silently ignored. For more information consult the on-line man pages and Appendix A, "Conversion from the CONVEX Distributed Batch System," in the *CONVEX CXbatch System Manager's Guide*.
- When running jobs under the C shell, if a login shell is requested using the '-l' option to qsub, the message 'Warning: no access to tty...' will be included in the job output file. There is no way CXbatch can suppress this message. It would require a change to the C shell.
- When the 'Import_dir' attribute is set for a job request that is to run on a remote

machine, the current working directory of the request is made available to the remote system through the use of an NFS mount. These mounts are made onto temporary directories in the /tmp filesystem. Therefore, system administrators should pay particular attention to any automatic clean-up procedures for the /tmp filesystem and assure that such procedures do not traverse NFS mount points.

- Measures should be taken to insure that the file system containing /usr/spool/nqs has a reasonable amount of free disk space during regular queue processing. Running out of disk space can cause job status and output to be lost if output files cannot be copied back to a request's current working directory due to disk space problems. Mail notification can also be lost due to the inability to write into /usr/spool/mail due to disk space problems.
- While every effort is made to ensure that batch jobs are controlled and cannot disrupt the system, just as in any user environment a persistent user can circumvent some of the controlling mechanisms. For example, to get around a file size limit a user may catch SIGXFSZ and switch output files or a user might catch and ignore SIGXCPU to bypass a soft cpu limit. (Also see the Known Software Problems section of this document for more information on this problem.) System administrators have the responsibility to monitor and prevent system abuse through the use of online protection mechanisms (Disk Quotas, resource limits, etc.) and offline human management techniques.

1.4 Enhancements

This section summarizes enhancement included in CONVEX CXbatch V1.1.

- The addition of the 'qmgr modify request priority' command to allow requests within a queue to be re-sequenced.
- The addition of the 'qmgr set maximum request_priority' to allow queue managers to limit the priority which users can request with qsub. This allows 'priority space' to be reserved to run high priority requests using the 'qmgr modify request priority' command.
- The addition of the 'qsub -l' option to execute a request under a login shell. This was necessary because a request no longer runs under a login shell by default.
- The modification of the logdaemon to use the syslog facility.
- The removal of the 'qmgr set log_file' command. This was made obsolete by the inclusion of the syslog functionality.

1.5 Fixes

This section lists the bugs reported against CXbatch which are addressed by the CXbatch V1.1 release.

- *CXbatch* (PR-09334, PR-11440, PR-11962, PR-11988, PR-12070, PR-12129, PR-12365) — Need tool to change priority of jobs in queue instead of putting all jobs on hold, except one (or how every many) to run.

Resolution: The 'qmgr modify request priority' command has been added. This command allows a CXbatch operator to modify the intra-queue priority of a request.

- *qsub* (PR-09386) — In the *qsub*(1) manual entry, and elsewhere, it is written that the chosen shell is exec'd as though it were the login shell. Thus, *.profile* (or *.cshrc* and *.login*) are executed. However, it is not documented whether or not */etc/profile* (or */etc/login*) is executed like a login shell would execute. Also, it is not documented whether or not */etc/logout* and *.logout* are executed. If those files in */etc* are executed, then it should be documented that the system manager may have to change them to test for *\$ENVIRONMENT*.

Resolution: The documentation has been changed to include the requested information. The default functionality has changed somewhat in the V1.1 release, in that a request is no longer run under a login shell unless the new '-l' option to *qsub* is used.

- *qmgr* (PR-09762) — If more than one process runs in a batchjob, it should be possible to set a CPU-limit for the whole batchjob and not only for the processes running in this batchjob.

Resolution: The set of resource limits that can be established by the batch system is dependent on the resource limits supported by the underlying operating system. In particular, ConvexOS does not, at this time, support any limits on groups of related processes (commonly referred to as sessions). As new functionality is added to ConvexOS in this area, corresponding enhancements will be made to CXbatch. This report is being closed as a restriction.

- *qmgr* (PR-09780) — The *qmgr*(8) manual entry command description section states that the keyword characters shown in uppercase indicate the smallest possible abbreviation of the keyword. The help command is shown as "Help". However, in practice, the "H" alone is ambiguous and so the help command should be shown as "HElp". Also the show limits command is shown as "SHoW Limits_supported" but, in practice the "L" is ambiguous and "LIM" is required. Therefore the show limits command should be shown as "SHoW LIMits_supported".

Resolution: The *qmgr* man page and help file have been corrected.

- *nqsd daemon* (PR-09817, PR-11949) — If the */usr* partition is full, the stdout and stderr of a batch job will be empty, and no errors are reported/logged/mailed/etc. *nqsd daemon* should log an error and send mail (if possible) to the user.

Resolution: The *nqsd daemon* now checks the space in the file system where jobs are spooled. If there is not enough space it logs an error and sends mail to the user.

- *qdel* (PR-09965, PR-09991) — Superuser doesn't have to use the *-u* option to delete another person's job.

Resolution: The *qdel* manpage has been changed to state this fact.

- *move* (PR-09966) — With *hold* & *unhold*, the 'user' option can be used by a batch operator or manager. With *move* and *remove*, only root can use the 'user' option.

Resolution: No further development time is being spent on these commands as they will be removed from the next release of CXbatch. This report is being closed as obsolete.

- *qsub.man* (PR-09967) — Man page should state that milliseconds may be ignored in limits.

Resolution: The man page has been updated to indicate this fact.

- *qlimit* (PR-09968) — *qlimit* should display the hostname if multiple hosts are given
Resolution: *qlimit* displays the hostnames on a header line if multiple names are given.
- *qmgr* (PR-09969) — *qmgr* should check the validity of the pipe queue server.
Resolution: The *nqsdaemon* will not accept a pipe server that does not exist or is not executable. It also checks the validity of shells specified by 'set shell_strategy fixed'.
- *CXbatch* (PR-09973) — It would be nice if *cpus* were a batch resource that could be configured by the system administrator.
Resolution: The underlying ConvexOS operating system has no functionality to support this feature. This report is being closed as a restriction.
- *installation* (PR-09976) — There should be a way to de-install *CXbatch*.
Resolution: This request was made to allow regressing back to the old distributed batch system. Since distributed batch is no longer supported by Convex, this functionality is no longer necessary. This report is being closed as obsolete.
- *pipeldav* (PR-09977, PR-10568) — *pipeldav* should not use *rwhod* tables to get load average information. Another method should be used, perhaps *rpc*.
Resolution: *pipeldav* now determines the load average of remote hosts via the *rstat()* call.
- *CXbatch* (PR-09978) — The inodes in */usr/spool/nqs/private/root/...* should not be pre-allocated.
Resolution: The preallocation to the NQS inodes is a requirement due to the design of the NQS system. This report is being closed as a restriction.
- *CXbatch* (PR-09979) — We need to better document which *CXbatch* features won't work with other NQS batch systems.
Resolution: Every effort is being made to document Convex specific enhancements to the NQS batch standard.
- *nqsdaemon* (PR-09980) — If *CXbatch* used pseudo *ttys*, we wouldn't have the 'no tty' problems with *csh*, and users would have fewer problems with their *.login* files.
Resolution: *CXbatch* no longer starts batch jobs under a login shell. This means that the 'No access to tty ...' messages will disappear and that the user's *.login* file is not read.
- *qstat* (PR-09981) — *qstat* and *qmgr* always display units as plural. They should be singular if it is one unit.
Resolution: *qstat* and *qmgr* display units as singular if they are singular instead of always displaying units as plural.
- *remove* (PR-09982) — The 'user' option no longer works.
Resolution: No further development time is being spent on this command since it is being removed from the next release of *CXbatch*. This report is being closed as obsolete.

- *remove* (PR-09983) — The exit status of *remove* (and other skeleton programs) is always zero. They should exit with a proper status.

Resolution: No further development time is being spent on these commands since they are being removed from the next release of CXbatch. This report is being closed as obsolete.

- *CXbatch* (PR-09986) — CXbatch doesn't handle long login names well. You get a strange error message if the length exceeds the maximum (15 chars).

Resolution: *qsub* will not allow a batch request to be submitted if the user's name exceeds the maximum allowed. It displays a more useful error message; 'Username exceeds 15 characters'.

- *CXbatch* (PR-09989, PR-11316, PR-11686, PR-12199) — There is an overall run limit of 10 jobs. This should be raised for our system.

Resolution: The maximum run limit is now effectively limited only by the amount of memory available. The *qmgr* limits the number of running jobs per queue to 30,000.

- *qsub.man* (PR-09990) — The example:
 qsub -a "tomorrow 23-MST"
from the man page doesn't work.

Resolution: This appears to work as documented. This report is being closed as unreproducible.

- *qmgr.man* (PR-09992) — The man page for *qmgr* says that 'delete request' can only be used by a operator or manager. In fact, a user can use this command to delete their own requests.

Resolution: The *qmgr* manpage has been changed to state this fact.

- *qmgr* (PR-09993) — Contrary to documentation, the 'move my_request' command acts just like 'move request' when executed by root.

Resolution: This has been corrected and it now works as documented.

- *qmgr* (PR-09997) — The command 'purge queue' will, on rare occasions, claim to have succeeded, but fail to remove queued requests.

Resolution: The purge queue should take place before the abort q. Requests that were queued will start to execute after the abort q.

- *qmgr.hlp* (PR-09998) — The *qmgr* help file does not mention what privileges are required to use the 'add alias' and 'delete alias' commands.

Resolution: Added description of the privileges required for the add alias and delete alias commands to the *qmgr* help file.

- *CXbatch* (PR-10000, PR-12005) — I'd like a way to suppress the creation of empty stdout and stderr files. Also, I'd don't want *cs*h to print the 'Warning: no access to tty...' message.

Resolution: Empty stderr and stdout files are not returned. The 'no access to tty...' messages no longer appear.

- *qmgr* (PR-10217) — From within *qmgr*, if you ask for 'help set nice', it states that nice values are in the range of -20..20, except on system V, where 20 counts as 19. It makes no mention of the Convex scheme of -64..64 and should do so.

Resolution: The *qmgr* help screen has been corrected.

- *qmgr* (PR-10220) — The man page states:
`SEt ACCounting = {OFF,ON} <queue-name>`
 but this conflicts with
`SEt ACC_logfile <logfile-name>`

in both cases you must type the first 4, not 3, characters. thus it should say:
`SEt ACCounting = {OFF,ON} <queue-name>`

I'm not sure how to indicate the underscore is mandatory, as there is no upper case version of that character.

Resolution: Man page corrected.

- *qmgr* (PR-10276) — If you write a program that ignores the SIGXCPU signal and run it in a batchqueue, this program will continue running even when its cpu-limit is reached, thus making cpu limits on queues useless. This needs to be more strictly enforced, perhaps by a SIGKILL by the *qmgr*.

Resolution: The implementation of resource limits in the batch system is dependent on the functionality supported by the underlying operating system. As new functionality is added to ConvexOS in this area, corresponding enhancements will be made to CXbatch. There is a binary kernel patch available from the TAC which will cause SIGKILL to be sent to a process that surpasses its hard cpu-limit. This report is being closed as a restriction.

- *CXbatch SMG* (PR-10381) — Page 2-16 has incorrect file name--/usr/etc/exports should be /etc/exports (in Step 2). Step 2 (2-16) has incorrect format for both 6.2 and 7.0 unix for the exports file. For 6.2, it should be
`/usr hostA hostB hostC`
 for 7.0, it should be:
`/usr -access=hostA:hostB:hostC`

Step 3 has the incorrect name for executing; it should be /usr/etc/exportfs (for 7.0).

Resolution: The documentation has been corrected.

- *qsub* (PR-10457) — One user noticed that the batch system mounts users' file systems or working directories under /tmp. This should be specified VERY clearly in the documentation because /tmp is a file system that is often cleaned up by a find/rm in .crontab! This means that it is IMPERATIVE not to let the 'find' run down the batch-induced mount points (xdev or -type f) or else a user finds all of their older files vanishing!

Resolution: This information is being included in several places in the documentation.

- *CXbatch_users_guide* (PR-10562) — Page 4-1, Section 4.1 discusses the 'qstat' command. Table 4-1 shows options for the 'qdel' command instead of 'qstat'.

Resolution: This has been corrected.

- *qmgr* (PR-10665) — *qmgr* needs the ability to reorder jobs in a particular queue.

Resolution: This functionality has been achieved through the addition of the 'MODify

Request Priority' command in qmgr.

- *install* (PR-10688) — The installation script assumes that the host name of the YP master server should be the same as the YP domain name. This is very untrue. Another method of distinguishing this must be used.

Resolution: The installation script now gives you the opportunity to correct the YP master hostname if it is incorrect.

- *install* (PR-10688) — The installation script assumes that the source file for the 'services' YP map is in /etc/services. Ours is in /etc/yp/src/services. One could prompt the installer (possibly with reasonable defaults) for the above information or else instruct them to make the changes to the services file manually (much less wear and tear on the install script writer).

Resolution: The installation procedure now gives you the opportunity to enter an alternate name for the services file.

- *submit* (PR-10694) — The submit man page for CXbatch V1.0 says that users are notified of job completion in a manner similar to write(1). This is not supported in CXbatch and the paragraph should be removed from the submit man page.

Resolution: Removed from man page.

- *qsub* (PR-10870) — The message 'Warning: no access to tty; thus no job control in this shell...' is considered a bug.

Resolution: This has been fixed.

- *qsub* (PR-10871) — You get an empty error file, rather than no error file, if there are no errors. This should be fixed or documented.

Resolution: Empty stderr or stdout files are no longer returned.

- *nqsdemon* (PR-10898) — The nqsdemon executes jobs with a default path that doesn't include /usr/convex.

Resolution: The nqsdemon now includes /usr/convex in the default path.

- *qsub(1)* (PR-10899) — The man page for qsub reads (in part):

7. The environment variables of HOME, SHELL, PATH, LOGNAME (not all systems), USER (not all systems), and MAIL are set from the user's password file entry, as though the user had logged directly into the execution machine.

This is incorrect, PATH and MAIL are NOT set from the password file. In particular, PATH is first set from a default value in nqsdemon.

Resolution: The man page has been corrected.

- *CPU-limit* (PR-10979, PR-11500, PR-12341) — With the new batch system there is no way of specifying a CPU-limit on a per-request basis. Our customers urgently need this feature. Batch jobs can run for ever in a CPU-limited Queue if you catch SIGXCPU, which makes CPU-limits on queues useless. Will there be another solution to this problem than a kernel patch?

Resolution: The the set of resource limits that can be established by the batch system

is dependent on the resource limits supported by the underlying operating system. In particular, ConvexOS does not, at this time, support any limits on groups of related processes (commonly referred to as sessions). As new functionality is added to ConvexOS in this area, corresponding enhancements will be made to CXbatch. There is a binary kernel patch available from the TAC which will cause SIGKILL to be sent to a process that surpasses its hard cpu-limit. This report is being closed as a restriction.

- *doc* (PR-10979) — I could not find a clear statement in the CXbatch documentation that per-request-CPU-limits don't work. There also should be clear statements for things that don't work.

Resolution: The documentation now states that the qlimit command and the qmgr 'show limits_supported' command are the definitive sources for which limits are supported on a given machine.

- *install* (PR-11254) — The installation script for CXbatch (when enhancements are shipped) needs a message displayed along with the question DO YOU WANT THE DEFAULT CONFIGURATION? telling that the current CXbatch database will be destroyed if the default configuration is selected. Make it clear that entering 'NO' will retain the current database. Install Doc should emphasize this, too.

Resolution: The installation script now has an expanded message about this operation.

- *qmgr* (PR-11359) — qmgr should support an option to change a job priority once it is in the queue, enabling that job to be the next job to start.

Resolution: This functionality has been achieved through the addition of the "MODify Request Priority" command to qmgr.

- *qmgr* (PR-11422) — Need to be able to re-set intra-queue values of queued jobs. The manager and operator need to be able to set a maximum intraqueue priority that users can specify on the qsub cmd and need to be able to re-set this value on queued jobs in order to move some jobs up or down in the queue.

Resolution: The 'qmgr modify request priority' command has been added. This command allows the CXbatch operator to change the intra-queue priority of a request.

- *installsw* (PR-11536) — The install for CXbatch assumes that your machine has networking that networking is running. Many installs are done in single user. The install script needs to start inetd before it tries to add the mid. Otherwise, the install fails, and you have to add the mid's by hand. This becomes very painful if you are also running yp.

Resolution: The installation script has been enhanced to take these factors into account.

- *nice in CXbatch* (PR-11548, PR-11796) — The per process nice value is limited to the range -20 to 20. When using the old distributed batch system, we had one of our queues set to nice value of 48. With CXbatch, we are unable to do this.

Resolution: The nqsdaemon and qmgr now support the CONVEX maximum and minimum nice values of 64..-64.

- *cpu per request limit* (PR-11552) — The CXbatch -IT option does not work. Option -IT specifies the per-request-time limit. (Not supported by convex but supported by cray).

Resolution: The set of resource limits that can be established by the batch system is dependent on the resource limits supported by the underlying operating system. In

particular, ConvexOS does not, at this time, support any limits on groups of related processes (commonly referred to as sessions). As new functionality is added to ConvexOS in this area, corresponding enhancements will be made to CXbatch. This report is being closed as a restriction.

- *CXbatch_doc* (PR-11571) — The CXbatch User's Guide, on page 2-20, section 2.5.2, contains a statement that not all UNIX applications support per-request limits, machines that do not support the limits ignore the limit, and that CONVEX machines do not support the limits and ignore them. The man page for CXbatch, distributed with the release notes, under its discussion of the -IT flag on page 10, needs the statement added that CONVEX machines do not support the limits and ignore them.

Resolution: The documentation now states that the qlimit command and the qmgr 'show limits_supported' command are the definitive sources for which limits are supported on a given machine.

- *shutdown* (PR-11597) — It is too easy to confuse the system shutdown command with the qmgr shutdown command and shut the system down to single user instead of shutting down CXbatch as desired--easy to forget "qmgr"

Resolution: The qmgr shutdown command is part of the standard NQS qmgr, and thus Convex hesitated to change it. A change in shutdown(8) in ConvexOS 8.0, however, has mostly solved this problem, in that the user has to confirm a system shutdown. This report is being closed as a restriction.

- *help screen CXbatch* (PR-11920) — The help screen is inconsistent. All lines should begin with Help.

Resolution: The help screen has been fixed.

- *cpu-time limit* (PR-11951) — Implement the per_request cpu time limit to help balance use of system.

Resolution: The set of resource limits that can be established by the batch system is dependent on the resource limits supported by the underlying operating system. In particular, ConvexOS does not, at this time, support any limits on groups of related processes (commonly referred to as sessions). As new functionality is added to ConvexOS in this area, corresponding enhancements will be made to CXbatch. This report is being closed as a restriction.

- *CXbatch* (PR-12131) — The default value of PATH is in nqsdaemon. The default path in CXbatch is ":/usr/ucb:/bin:/usr/bin". We believe the path should also include /usr/convex.

Resolution: The nqsdaemon now includes /usr/convex in the default path.

- *qmgr* (PR-12275) — If the share utility is running on the system and you attempt to create a CXbatch queue that is not in the password file, you get a core dump instead of the error message that there is no entry in the password file for the queue you are creating with qmgr. This happened on mikey and saki--mikey is at 8.0.0.2 and saki is at 8.0.0.4.

Resolution: This problem was corrected in the nqsdaemon and qmgr included in the Share installation. There were problems that prevented the new qmgr from being correctly installed. This should no longer be a problem.

1.6 Known Software Problems

This section lists the known problems with CXbatch V1.1 as of March 5, 1990. Any problems reported after this date may not be reflected in the document. Please refer to this list before reporting a problem in order to ensure that it has not been reported previously. Serious problems include workarounds if they are known.

- If your batch shell is */bin/sh* and your execution directory is not \$HOME, *sh* does not read your *.profile*.
- If a batch job tries to increase a resource limit above the job's hard limit with the *cs*h commands *limit* or *unlimit*, *cs*h prints the cryptic error message "Not owner." Also, because of a bug in *cs*h, any commands in the *.login* file after the failing *limit* command are not interpreted.
- The skeleton programs provided with CXbatch (*list*, *submit*, *sbatch*, *ssp*, *move*, *remove*, *hold*, and *unhold*) are not robust. They ignore many options that were supported in CONVEX Distributed Batch. Consult the on-line man pages for complete information.
- There is a known problem in the ConvexOS V8.0 kernel which causes a SIGXCPU rather than a SIGKILL to be sent to a process which surpasses its hard cpu-limit. If a user process catches or ignores SIGXCPU, the process will continue to run. A binary kernel patch is available from the TAC to correct this problem.
- Beginning with the ConvexOS V8.0 release, it is no longer possible to set most resource limits to 0 (zero). CXbatch V1.1 does not yet account for this change and will return an error via mail whenever a request attempts to set a resource limit to 0.
- Under certain configurations involving Yellow Pages, standard error output from CXbatch requests may be lost. If you experience this problem, please contact your TAC representative.
- Under heavy stress, *netdaemon* may shut down due to an *accept()* error. If this happens, all batch managers receive mail notifying them of the problem, and an error is logged in the batch log file. The problem can be worked around by shutting down and restarting CXbatch.
- There is no way to unset a description on a queue after adding one with the *qmgr* 'set description' command.

1.7 Documentation

The documentation provided to you as part of the CXbatch package includes man pages for the utilities and the following manuals:

- *CONVEX CXbatch Concepts*
- *CONVEX CXbatch User's Guide*
- *CONVEX CXbatch Programmer's Reference Manual*
- *CONVEX CXbatch System Manager's Guide*
- *CONVEX CXbatch System Manager Utilities Reference Manual*
- *CONVEX CXbatch Master Index*

Man pages for each of the CXbatch utilities may be found on line in section 1; man pages for the daemons and system administration programs may be found on line in section 8. Printed copies of all man pages are included in the distribution package.

1.8 Known Documentation Problems

There are no known documentation problems to describe.