

UNICOS® Configuration Administrator's Guide

S-2303-10011

© 1995–2002 Cray Inc. All Rights Reserved. This manual or parts thereof may not be reproduced in any form unless permitted by contract or by written permission of Cray Inc.

U.S. GOVERNMENT RESTRICTED RIGHTS NOTICE

The Computer Software is delivered as "Commercial Computer Software" as defined in DFARS 48 CFR 252.227-7014.

All Computer Software and Computer Software Documentation acquired by or for the U.S. Government is provided with Restricted Rights. Use, duplication or disclosure by the U.S. Government is subject to the restrictions described in FAR 48 CFR 52.227-14 or DFARS 48 CFR 252.227-7014, as applicable.

Technical Data acquired by or for the U.S. Government, if any, is provided with Limited Rights. Use, duplication or disclosure by the U.S. Government is subject to the restrictions described in FAR 48 CFR 52.227-14 or DFARS 48 CFR 252.227-7013, as applicable.

Autotasking, CF77, Cray, Cray Ada, Cray Channels, Cray Chips, CraySoft, Cray Y-MP, Cray-1, CRInform, CRI/*TurboKiva*, HSX, LibSci, MPP Apprentice, SSD, SuperCluster, UNICOS, UNICOS/mk, and X-MP EA are federally registered trademarks and Because no workstation is an island, CCI, CCMT, CF90, CFT, CFT2, CFT77, ConCurrent Maintenance Tools, COS, Cray Animation Theater, Cray APP, Cray C90, Cray C90D, Cray CF90, Cray C++ Compiling System, CrayDoc, Cray EL, Cray Fortran Compiler, Cray J90, Cray J90se, Cray J916, Cray J932, CrayLink, Cray MTA, Cray MTA-2, Cray MTX, Cray NQS, Cray/REELlibrarian, Cray S-MP, Cray SSD-T90, Cray SV1, Cray SV1ex, Cray SV2, Cray SX-5, Cray SX-6, Cray T90, Cray T94, Cray T916, Cray T932, Cray T3D, Cray T3D MC, Cray T3D MCA, Cray T3D SC, Cray T3E, CrayTutor, Cray X-MP, Cray XMS, Cray-2, CSIM, CVT, Delivering the power . . . , DGauss, Docview, EMDS, GigaRing, HEXAR, IOS, ND Series Network Disk Array, Network Queuing Environment, Network Queuing Tools, OLNET, RQS, SEGLDR, SMARTE, SUPERLINK, System Maintenance and Remote Testing Environment, Trusted UNICOS, and UNICOS MAX are trademarks of Cray Inc.

DECnet, VAX, and VAXBI are trademarks of Digital Equipment Corporation. IBM is a trademark and VM is a product of International Business Machines Corporation. ESCON is a trademark of International Business Machines Corporation. Kerberos is a trademark of the Massachusetts Institute of Technology. NASTRAN is a trademark of National Aeronautics and Space Administration. NFS is a trademark of Sun Microsystems, Inc. NSC is a trademark of Network Systems Corporation. REELlibrarian is a trademark of Sceptre Corporation. UNIX, the "X device," X Window System, and X/Open are trademarks of The Open Group in the United States and other countries. All other trademarks are the property of their respective owners.

Cray UNICOS Version 10.0 is an X/Open Base 95 branded product.

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

New Features

UNICOS® Configuration Administrator's Guide

S-2303-10011

This document was updated for the UNICOS 10.0.1.1 operating system release, specifically:

- The `NSIDEBUF` kernel parameter of the `config.h` file was changed so that it is now valid and user changable on Cray SV1ex systems.
- The `BTE_STATE` keyword was removed from the CSL parameter file because it was only valid during development.
- On GigaRing based systems, the common disk parameters `PDDMAX` and `PDDSLMAX` can now each have a minimum of 0 specified.

Record of Revision

<i>Version</i>	<i>Description</i>
9.0	May 1995 Original Printing. Documentation to support the UNICOS 9.0 release running on all Cray computer systems. This manual contains the contents of and supersedes the information formerly provided in the “Configuring UNICOS” section in the <i>UNICOS System Administration</i> , publication SG-2113 8.0
9.1	November 1995 Revision to support the UNICOS 9.1 release.
9.2	December 1996 Revision to support the UNICOS 9.2 release.
9.3	May 1997 Revision to support the UNICOS 9.3 release.
10.0	October 1997 Revision to support the UNICOS 10.0 release.
10.0.0.3	September 1998 Revision to support the UNICOS 10.0.0.3 release.
10.0.0.4	January 1999 Revision to support the UNICOS 10.0.0.4 release.
003	July 1999 Revision to support the UNICOS 10.0.0.6 release.
004	February 2000 Revision to support the UNICOS 10.0.0.7 release.
10008	November 2000 Revision to support the UNICOS 10.0.0.8 release.
10010	October 2001 Revision to support the UNICOS 10.0.1.0 release.
10011	May 2002 Revision to support the UNICOS 10.0.1.1 release.

Contents

	<i>Page</i>
Preface	ix
UNICOS System Administration Publications	ix
Related Publications	x
Ordering Documentation	xii
Conventions	xii
Reader Comments	xiv
Introduction [1]	1
Special GigaRing Configuration Rules for the Cray SV1ex System	2
Minimum GigaRing Configuration	2
SSD Memory Access	3
Back Door I/O Rules	4
SuperRing Configuration Rules	4
<code>sn.h</code> File [2]	7
Required Steps	8
Optional Steps	9
<code>config.h</code> File [3]	11
<code>config.mh</code> File [4]	31
CSL Parameter File [5]	37
CSL Syntax	37
Identifiers	37
Constants	38
Operators, Separators, and Comments	38
CSL Usage	39
S-2303-10011	iii

	<i>Page</i>
The Boot Process	39
Configuration Verification	39
Error Message Syntax	39
CSL Parameter File Sections	40
dumpinfo Section	41
gigaring Section	42
gr_route Subsection	43
gr_union Subsection	43
ios_e Section	44
IOP Declarations	45
HISP Declarations	46
IOP Boot Declarations	46
Example of ios_e Section	46
mainframe Section	47
Common Parameters	48
Parameters for Cray SV1 series GigaRing Based System	50
Parameters for Cray SV1 series VME Systems Only	52
Parameters for GigaRing Based Systems Only	52
Parameters for Model E Based Systems Only	54
Example of the mainframe Section for a GigaRing Based System	54
Example of the mainframe Section for a Model E Based System	55
unicos Section	55
Example for a GigaRing Based System	59
Additional Parameters for a Cray SSD-T90 System	60
Additional Parameters for a Cray SV1ex System with SSD-I	60
Example for a Model E Based System	60
filesystem Section	61
Physical Device Definition	62
Device Node Definition	73
Cray SSD-T90 Description	74

	<i>Page</i>
Cray SV1ex Auxiliary Memory (SSD-I) Description	75
Root, Swap, and Secondary Data Segment (SDS) Devices	76
Example of the <code>filesystem</code> Section Containing a RAM File System	76
Example of the <code>filesystem</code> Section for a GigaRing Based System	77
Example of the <code>filesystem</code> Section for a GigaRing Based System with Disk Devices Configured for Third-party I/O	77
Example of the <code>filesystem</code> Section for a Cray SSD-T90 Device	78
Example of the <code>filesystem</code> Section for a SSD-I Device	79
Example of the <code>filesystem</code> Section for Model E Based Systems	80
Example of the <code>filesystem</code> Section Containing an SSD for Model E Based Systems	81
<code>network</code> Section	81
Network Parameters	82
Customized Network Device Prototypes for Model E Based Systems	87
<code>device type</code> for Model E and Model V Based Systems	88
Device Types	89
Device Formats	90
<code>logical path</code> for Model E Based Systems	94
<code>direction</code> Argument for Model E Based Systems	95
<code>device type</code> Statement for Model E Based Systems	95
<code>np_spec</code> Statement for Model E Based Systems	95
<code>network</code> Section Example Common to All Systems	96
<code>network</code> Section Example for GigaRing Based Systems	96
<code>network</code> Section Example for Model E Based Systems	97
<code>network</code> Section Example for Model V Based Systems	98
<code>revision</code> Section	99
Runtime Configuration Scripts and Files [6]	101
<code>at.deny</code> and <code>at.allow</code> Files	102
<code>bcheckrc</code> Script	102
<code>brc</code> and <code>coredd</code> Scripts	103
<code>cron.deny</code> and <code>cron.allow</code> Files	103

	<i>Page</i>
cshrc File	103
daemons File	104
fstab File	105
gettydefs File	106
ghippi#.arp File	107
gr#.arp File	107
group File	107
hycf.local_network Files	108
inittab File	108
interfaces File	109
issue File	109
ldchlist File	109
motd File	110
netstart Script	110
netvar.conf File	111
passwd File	111
profile File	112
rc Script and rcoptions File	112
shutdown Script	113
umountem Script	113
Appendix A Reserved Keywords in CSL	115
Index	119
Tables	
Table 1. Kernel Parameters in config.h (Common)	12
Table 2. Kernel Parameters in config.h (Common) Not in ICMS	15
Table 3. Dynamic Kernel Memory Parameters	16
Table 4. Shared Memory Kernel Parameters in config.h	18
Table 5. Kernel Parameters in config.h for Process Error Thresholds	19
Table 6. IPC Kernel Parameters in config.h	19

	<i>Page</i>
Table 7. MLS Kernel Parameters in <code>config.h</code>	22
Table 8. TCP/IP Kernel Parameter in <code>config.h</code>	28
Table 9. NFS Kernel Parameters in <code>config.h</code>	28
Table 10. CSL Boot-time Equivalents for NFS Kernel Parameters	29
Table 11. General System Parameters (Common)	32
Table 12. General System Parameters (GigaRing Based Systems Only)	32
Table 13. General System Parameters (Model E Based Systems Only)	33
Table 14. Kernel Generation Parameters	33
Table 15. Kernel Subsystem Parameters	35
Table 16. Programming Environment Parameters	36
Table 17. CPU/MSP Correlation	51
Table 18. I/O Module 0 Channels	53
Table 19. Example of Complete Channel Assignments	54
Table 20. Online Tape Parameters	56
Table 21. Table Size Parameters	56
Table 22. Maximum Limits Parameters	56
Table 23. Disk Parameters (Common)	57
Table 24. Disk Parameters (GigaRing Based Systems Only)	58
Table 25. Disk and Block Transfer Engine Parameters (Cray SV1ex Systems with SSD-I Only)	58
Table 26. Disk Parameters (Model E Based Systems Only)	59
Table 27. ION Unit Bit and Range Numbers (GigaRing Based Systems Only)	59
Table 28. Start and Length Units for Physical Storage Device Types	63
Table 29. Disk Information (GigaRing Based Systems Only)	64
Table 30. Target Memory Type Values	75
Table 31. Network Parameter Values (Common)	83
Table 32. Network Parameter Values (GigaRing Based Systems)	84
Table 33. Network Parameter Values (Common to Model E and Model V Based Systems)	85
Table 34. Network Parameter Values (Model V Based Systems)	86
Table 35. Network Parameter Values (Model E Based Systems)	86
Table 36. Standard Interface Configuration Templates	89

	<i>Page</i>
Table 37. Network Device Types (GigaRing Based Systems)	90
Table 38. Network Device Type (Model E Based Systems)	90
Table 39. Network Device Type (Model V Based Systems)	90

This publication provides information about the UNICOS configuration files created when the UNICOS version 10.0 operating system is installed and configured.

This manual documents UNICOS version 10.0 running on Cray systems. It contains information needed in the administration of various UNICOS features available to all UNICOS systems.



Warning: Starting with the UNICOS release 10.0, the term *Cray ML-Safe* replaces the term *Trusted UNICOS*, which referred to the system configuration used to achieve the UNICOS 8.0.2 release evaluation. Because of changes to available software, hardware, and system configurations since the UNICOS 8.0.2 system release, the term *Cray ML-Safe* does not imply an evaluated product, but refers to the currently available system configuration that closely resembles that of the evaluated Trusted UNICOS 8.0.2 system.

For the UNICOS 10.0 release, the functionality of the Trusted UNICOS system has been retained, but the `CONFIG_TRUSTED` option, which enforces conformance to the strict B1 configuration, is no longer available.

UNICOS System Administration Publications

Information on the structure and operation of a Cray computer system running the UNICOS operating system, as well as information on administering various products that run under the UNICOS operating system, is contained in the following documents:

- *General UNICOS System Administration* contains information on performing basic administration tasks as well as information about system and security administration using the UNICOS multilevel security (MLS) feature. This publication contains chapters documenting file system planning, UNICOS startup and shutdown procedures, file system maintenance, basic administration tools, crash and dump analysis, the UNICOS MLS feature, and administration of online features.
- *UNICOS Resource Administration*, contains information on the administration of various UNICOS features available to all UNICOS systems. This publication contains chapters documenting accounting, automatic incident reporting (AIR), the fair-share scheduler, file system quotas, file system monitoring, system activity and performance monitoring, and the Unified Resource Manager (URM).

- *UNICOS Configuration Administrator's Guide*, provides information about the UNICOS kernel configuration files and the runtime configuration files and scripts.
- *UNICOS Networking Facilities Administrator's Guide*, contains information on administration of networking facilities supported by the UNICOS operating system. This publication contains chapters documenting TCP/IP for the UNICOS operating system, the UNICOS network file system (NFS) feature, and the network information system (NIS) feature.
- *NQE Administration*, describes how to configure, monitor, and control the Cray Network Queuing Environment (NQE) running on a UNIX system.
- *Kerberos Administrator's Guide*, contains information on administration of the Kerberos feature, a set of programs and libraries that provide distributed authentication over an open network. This publication contains chapters documenting Kerberos implementation, configuration, and troubleshooting.
- *Tape Subsystem Administration*, contains information on administration of UNICOS and UNICOS/mk tape subsystems. This publication contains chapters documenting tape subsystem administration commands, tape configuration, administration issues, and tape troubleshooting.

Related Publications

The following man page manuals contain additional information that may be helpful.

- *UNICOS User Commands Reference Manual*
- *UNICOS System Calls Reference Manual*
- *UNICOS File Formats and Special Files Reference Manual*
- *UNICOS Administrator Commands Reference Manual*
- *UNICOS System Libraries Reference Manual*

The following publication is useful for establishing connectivity between the High Performance Parallel Interface (HIPPI) network of a Cray mainframe and any host that has a physical path to any of the network interfaces of the Cray L7R.

- *Cray L7R Release Overview and Software Installation Guide*, contains information on the Cray L7R release and details regarding software installation and configuration for the Cray L7R. This publication contains chapters

documenting an overview of the release, purpose and function of the Cray L7R, system and network configuration requirements, software installation and configuration instructions, and troubleshooting.

Design specifications for the UNICOS multilevel security (MLS) feature are based on the trusted computer system evaluation criteria developed by the U. S. Department of Defense (DoD). If you require more information about multilevel security on UNICOS, you may find the following sources helpful:

- DoD Computer Security Center. *A Guide to Understanding Trusted Facility Management* (DoD NCSC-TG-015). Fort George G. Meade, Maryland: 1989.
- DoD Computer Security Center. *Department of Defense Trusted Computer System Evaluation Criteria* (DoD 5200.28-STD). Fort George G. Meade, Maryland: 1985. (Also known as the *Orange book*.)
- DoD Computer Security Center. *Trusted Network Interpretation of the Trusted Computer System Evaluation Criteria* (DoD NCSC-TG-005-STD). Fort George G. Meade, Maryland: 1987. (Also known as the *Red book*.)
- DoD Computer Security Center. *Summary of Changes, Memorandum for the Record* (DoD 5200.28-STD). Fort George G. Meade, Maryland: 1986.
- DoD Computer Security Center. *Password Management Guidelines* (CSC-STD-002-85). Fort George G. Meade, Maryland: 1985.
- Wood, Patrick H. and Stephen G. Kochan. *UNIX System Security*. Hasbrouck Heights, N.J.: Hayden Book Company, 1985.

Ordering Documentation

To order software documentation, contact the Cray Software Distribution Center in any of the following ways:

E-mail:

orderdsk@cray.com

Web:

<http://www.cray.com/craydoc/>

Click on the [Cray Publications Order Form link](#).

Telephone (inside U.S., Canada):

1-800-284-2729 (BUG CRAY), then 605-9100

Telephone (outside U.S., Canada):

Contact your Cray representative, or call +1-651-605-9100

Fax:

+1-651-605-9001

Mail:

Software Distribution Center
Cray Inc.
1340 Mendota Heights Road
Mendota Heights, MN 55120-1128
USA

Conventions

The following conventions are used throughout this document:

<u>Convention</u>	<u>Meaning</u>
command	This fixed-space font denotes literal items, such as file names, pathnames, man page names, command names, and programming language elements.
<i>variable</i>	Italic typeface indicates an element that you will replace with a specific value. For instance, you may replace <i>filename</i> with the name <i>datafile</i> in your program. It also denotes a word or concept being defined.

user input	This bold, fixed-space font denotes literal items that the user enters in interactive sessions. Output is shown in nonbold, fixed-space font.
[]	Brackets enclose optional portions of a syntax representation for a command, library routine, system call, and so on.
. . .	Ellipses indicate that a preceding element can be repeated.

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

<u>Term</u>	<u>Definition</u>
Cray PVP systems	All configurations of Cray parallel vector processing (PVP) systems, including Cray SV1 series systems.
Cray MPP systems	All configurations of the Cray T3E series. The UNICOS operating system is not supported on Cray T3E systems. Cray T3E systems run the UNICOS/mk operating system.

Reader Comments

Contact us with any comments that will help us to improve the accuracy and usability of this document. Be sure to include the title and number of the document with your comments. We value your comments and will respond to them promptly. Contact us in any of the following ways:

E-mail:

swpubs@cray.com

Telephone (inside U.S., Canada):

1-800-950-2729 (Cray Customer Support Center)

Telephone (outside U.S., Canada):

Contact your Cray representative, or call +1-715-726-4993 (Cray Customer Support Center)

Mail:

Software Publications
Cray Inc.
1340 Mendota Heights Road
Mendota Heights, MN 55120-1128
USA

Introduction [1]

Properly configuring a system is very important and should be done carefully to ensure the efficient performance of your UNICOS computer system. *UNICOS System Configuration Using ICMS*, and the online help files associated with the menu system provide the information you need to configure the UNICOS operating system. This manual provides additional details about the UNICOS kernel configuration files and the run-time configuration files and scripts.

Note: It is strongly recommended that you use the UNICOS installation and configuration menu system (ICMS) to maintain your system configuration.

UNICOS kernel configuration information appears in the following files:

- `sn.h`, which contains parameters that define machine-specific characteristics of your mainframe.
- `config.h`, which contains parameters that define the UNICOS kernel.
- `config.mh`, which contains parameters used to include and exclude kernel subsystems.
- Configuration specification language (CSL) parameter file, which contains statements that specify hardware and software characteristics for your system.

Note: As of the UNICOS 9.2 release, ICMS is not used for installation but is used to maintain configuration.

The default values presented within this document are not necessarily the values a site must use. It is up to the system support staff to determine optimal values for the site. If help is needed, call your local Cray service representative.

The *Cray ML-Safe configuration* refers to a configuration of the UNICOS system that supports processing at multiple security labels and system administration using only non-super user administrative roles. The Cray ML-Safe configuration consists of the subset of UNICOS software that offers these capabilities.



Warning: Starting with the UNICOS 10.0 release, the term *Trusted UNICOS* has been replaced by the term *Cray ML-Safe*. Because of changes to available software, hardware, and system configurations since the UNICOS 8.0.2 system release, the term *Cray ML-Safe* does not imply an evaluated product, but refers to the currently available system configuration that closely resembles that of the evaluated *Trusted UNICOS 8.0.2* system.

1.1 Special GigaRing Configuration Rules for the Cray SV1ex System

The following sections describe rules that apply to all Cray SV1ex systems.

1.1.1 Minimum GigaRing Configuration

All Cray SV1ex systems must follow the following minimum configuration rules:

1. Each node must have at least one Cray-supplied disk drive for storing the UNICOS operating system. All new Cray SV1ex systems are shipped with fiber disks. Existing Cray J90 systems that are being upgraded with new processor modules can use existing fiber, IPI, or SCSI disks. Existing Cray SV1 systems that are being upgraded with only new processor modules or are being upgraded with both new processor and new memory modules can use existing fiber, IPI, or SCSI disks.
2. The Processor 0 GigaRing channel is always present and is used to boot the system. Therefore, this ring must include the disk drive(s) that contain UNICOS.
3. Each node requires an Ethernet connection for logging in. This requires one MPN with an ETN-11.
4. The first PC-10 cabinet (joined to the right side of the processor cabinet, as viewed from the front) contains the GigaRing channels from processors 0-3. The GigaRing bulkhead in this PC-10 cabinet should be filled starting with Processor 0. Do not skip processors.
5. Whenever possible fill the first PC-10 cabinet (processors 0-3) in the system first.
6. If present on the system, the second PC-10 cabinet (joined to the left side of the Cray SV1 series cabinet, as viewed from the front) contains the GigaRing channels from processors 4-7.
7. Ring channels must be configured with at least one ION (MPN or SPN).
8. MPNs are not allowed on shared rings that are used to connect multiple mainframe nodes (that is, a ring connecting a Cray T90 system to a Cray SV1ex system). GigaRing channels with multiple mainframe nodes must include at least one SPN (FCN, HPN, IPN, or ESN).

1.1.2 SSD Memory Access

Cray SV1ex systems have a new faster memory subsystem. Depending on system configuration, the Cray SV1ex can have extended memory, a portion of which includes an auxiliary memory known as an internal static storage device (SSD-I) and a high speed block transfer engine (BTE). SSD-I is internal to the Cray SV1ex main memory modules. Main memory occupies the lower 32 Gbytes of the memory subsystem, while SSD memory occupies the upper address space. Depending on the system configuration, each mainframe cabinet can have an SSD that is 0-, 32-, or 96-Gbytes. All memory words are 64 bits wide.

The main purpose of SSD is to temporarily store data sets of a job to speed up access to data sets. It is used essentially as the SSD-T (Cray SSD-T90) and SSD-E (Model E SSD) are used on other Cray PVP systems. Supported uses for SSD-I include fast swap space, file system space, logical device (disk) cache, and secondary data segment (SDS). The BTE provides a CPU-controlled data path for direct memory to memory transfers, for example, between main memory and extended memory or even within main memory.

SSD memory is accessible in two ways: front door and back door. All systems that have an SSD have front door access capability, which allows the SSD to be used as a swap device and/or SSD file system. Only system configured with back door access have full SSD functionality, such as direct disk transfer to and from SSD space and `ldcache`.

Front door access is defined as the movement of data between main memory and the internal SSD (SSD-I). This is accomplished by using `ssread` and `sswrite` system calls for access to SDS space. Control logic transfers words directly between Main memory and SSD. The transfer rate depends on buffer alignment, that is, BTE versus data movement via CPU. SSD-resident file systems and swap I/O also use front door access.

Back door access is defined as the movement of data directly between SSD-I and disk devices. Back door transfers can occur over either a conventional point-to-point GigaRing or a SuperRing consisting of two to four mainframe GigaRing adapters. Back door transfer rates are a function of the number of mainframe nodes, disk controllers, and disk devices, connected to the GigaRing or SuperRing channel. Whereas, the number of rings that can be configured is dependent on the number of processor modules in the system.

1.1.3 Back Door I/O Rules

Back door I/O can be provided on a Cray SV1ex system as follows:

1. The minimum configuration is a point-to-point ring consisting of one mainframe GigaRing connection and one ION, where the ION is an FCN, IPN, HPN, or disk or network MPN (not tape).
2. Alternatively, a SuperRing configuration, consisting of two to four mainframe GigaRing connections and one or more IONs on a single ring, can in some cases provide improved I/O bandwidth. See Section 1.1.4, page 4 for more information.
3. Although a ring that is set up to allow back door I/O can also be used for I/O transfers between disk and main memory, no networking or tape IONs can exist on the same ring. Any attempt to open a tape device on a ring configured with back door capability will fail and generate an error message on the system console.
4. Back door I/O access to ldcache in the SSD is limited to file systems that exist entirely on disk partitions that reside on properly configured back door-capable rings, including both primary and alternate I/O paths.
5. Back door I/O is required to use the SSD for secondary data segment (SDS) space (to support SDS space swapping).

1.1.4 SuperRing Configuration Rules

The following rules apply to all systems that have a SuperRing configuration:

1. A SuperRing is defined to be a single GigaRing consisting of two to four mainframe GigaRing channel adapters (on either processor modules or I/O modules) and one or more IONs on a single ring. Although a SuperRing can be configured to support front door (regular) I/O, a SuperRing is primarily intended to support back door (direct disk to SSD) I/O transfers.
2. There is no maximum supported configuration on a single SuperRing. However, if the total number of mainframe GigaRing adapters and IONs on a ring is greater than 6, performance might be degraded significantly.
3. A SuperRing configuration of three mainframe GigaRing connections with two FCNs provides the best balanced overall capability for back door I/O. The FCNs can move data at ~300 Mbyte/s and the mainframe nodes operate at 200 Mbyte/s.

4. Multiple SuperRings can be configured on a single mainframe; however, a SuperRing can not be a shared ring between multiple mainframe nodes.
5. On systems that use an 8x8 backplane, all the mainframe GigaRing adapters that make up a single SuperRing must be on the same side of the mainframe. (CPU slots 0-3 are on one side, 4-7 are on the other.)
6. On systems that use a 4x4 backplane, only one SuperRing consisting of either two or three mainframe GigaRing adapters is possible unless the system boot ring is configured as a SuperRing.

sn.h File [2]

Note: Typically, this file is stored in `/usr/src/uts/cf.xxxx/sn.h`. It is strongly recommended that you use the UNICOS installation and configuration menu system (ICMS) to maintain your system configuration, rather than manually editing this file. (For more information on the ICMS, see the online help files and the *UNICOS System Configuration Using ICMS*.) If special circumstances require that you set the parameters in `sn.h` manually, use the procedures in this chapter.

The `sn.h` file contains parameters that define machine-specific characteristics of your mainframe. You must change some of these parameters to reflect your system's characteristics.

You will set the parameters by using the following menu:

```
Configure System
  ->Mainframe hardware configuration
```

If you manually edit `sn.h`, you must define the following parameters:

- Mainframe serial number
- Main memory size
- Number of banks of main memory
- Number of bits per chip in main memory

If your system is not fully configured with CPUs, you should also define the following parameters:

- Number of CPUs
- Number of mainframe clusters

Appropriate defaults for the remaining parameters in the `sn.h` file are set automatically, based upon your machine's serial number. You will not need to set these parameters.

2.1 Required Steps

The following steps are required if you edit `sn.h` manually:

1. Set your mainframe serial number by defining the `SN` parameter to be that value.

Example:

```
#define SN          4025
```

2. Set the physical memory size parameter (`MEMORY`) in decimal words to the highest addressable word.

Example:

```
#define MEMORY      512*MEGAWD-1
```

3. Set the `NBANKS` parameter to the number of memory banks in your mainframe:

Example:

```
#define NBANKS      1024
```

4. Set the number of bits per memory chip used in main memory by defining the `CHIPSZ` parameter with one of the following values:

<u>Bits/chip</u>	<u>CHIPSZ value</u>
65536	M64KCH
262144	M256KCH
1048576	M1MCH
2097152	M2MCH
4194304	M4MCH
8388608	M8MCH
16777216	M16MCH
33554432	M32MCH – (not valid for UNICOS kernel)
67108864	M64MCH

Example:

```
#define CHIPSZ      M1MCH
```

2.2 Optional Steps

The following steps are optional. Appropriate default values are assigned according to the mainframe serial number.

1. Set the mainframe type parameter (MFTYPE) to one of the following values.

<u>Option</u>	<u>Description</u>
CRAY_TS	Cray T90
CRAYMP	Cray SV1 series Cray J90se

Example:

```
#define MFTYPE      CRAYT90
```

2. Set the mainframe subtype (MFSUBTYP) parameter to the appropriate option.

<u>Option</u>	<u>Description</u>
YMPJ90	Cray Y-MP architecture, including Cray SV1 series, and Cray J90se series
T4XXX	Cray T94
T16XXX	Cray T916
T32XXX	Cray T932

Example:

```
#define MFSUBTYP   T900XX
```

3. If your system is not fully configured, define the number of CPUs (NCPU) and cluster registers (MAXCLUS).

Example:

```
#define NCPU        8
#define MAXCLUS     9
```

4. Set the clock period to the frequency in hertz (cycles per second) by defining HZ; the default values are based on MFSUBTYP and mainframe serial number.

Example:

```
#define HZ          HZ_416
```


config.h File [3]

Note: It is strongly recommended that you use the UNICOS installation and configuration menu system (ICMS) to maintain your system configuration, rather than manually editing this file. (For more information on the ICMS, see the online help files associated with the tool and *UNICOS System Configuration Using ICMS*.) Typically, this file is stored in `/usr/src/uts/cf.xxxx/config.h`.

This chapter summarizes parameters found in the `/usr/src/uts/c1/cf/config.h` file.

Note: The default values shown reflect the settings used for initial installations; appropriate values for upgrades will be changed automatically as part of the upgrade conversion process. Sites performing an upgrade should only need to change values manually if they are enabling features or changing hardware.

These parameters pertain to general configuration, the UNICOS multilevel security (MLS) feature, and TCP/IP. To set the general configuration and TCP/IP parameters, use the following menu:

```
Configure System
->Kernel Configuration
```

To set the MLS parameters, use the following menu:

```
Configure System
->Multilevel Security (MLS) Configuration
```

In general, a value of 0 disables a parameter and a value of 1 enables it.

Table 1 through Table 2 summarize resource and configuration parameters for the UNICOS kernel. Some parameters are not available in the ICMS and may therefore be edited only by those sites with a source license.

Table 1. Kernel Parameters in `config.h` (Common)


Parameter	Default value	Description
CDLIMIT	010000000000	Maximum file size in blocks (see <code>ulimit(2)</code>).
DMODE	1	Default offline file retrieval mode for sites running data migration. 1 indicates automatic retrieval; 0 indicates manual-only retrieval.
QUEUE_IO_WITHOUT_NICE_PRIORITY_VALUES	0	Enable/disable the use of nice priorities on the I/O queues that access mirrored disks, striped disks, and non-GigaRing high performance parallel interface (HIPPI) disks. A value of 0 means that a low-priority process may not get access to a disk that is running with the priority-ordered I/O queue. A non-zero value means that even a low-priority process will have access to a disk because the priority order is not being used.
EXTDCORE	0	Extended core file naming.
FLUSHONPANIC	1	The capability to flush buffered data to target I/O devices before the system is halted.
LDCHCORE	0	Memory clicks reserved for <code>ldcache</code> blocks. These are used for <code>ldcache</code> blocks specified as type <code>MEM</code> . This value can be changed at boot time in the <code>CSL</code> parameter file.
LINK_MAX	1000	Maximum number of directory entries and maximum number of links to one file.
MAXASYN	192	Maximum number of asynchronous I/O structures allowed per process. <code>MAXASYN</code> has considerable performance impact on applications utilizing asynchronous I/O (for example, <code>listio</code> and <code>aqio</code> system calls).
MAXPIPE	20	Maximum number of blocks allocated per pipe on the pipe device.
MAXRAH	8	Maximum number of read-ahead blocks per read operation.
NASYN	800	Number of asynchronous I/O headers.

Parameter	Default value	Description
NBLK_FCTR	20	Factor for maximum request size to/from system buffer cache. The default of 20 means that 1/20th (5%) of the system buffer cache can be allocated to a single request.
NBUF_FCTR	20	Number of cache blocks based on memory size. The default of 20 means that 1/20th (5%) of the memory is assigned to buffer cache. This can be changed at boot time to an absolute number of cache blocks with the NBUF parameter in the CSL parameter file. There is now a default maximum of 10000 system buffers. This limit can be overridden by explicitly defining NBUF in the sn.h file, config.h file, or startup parameter file, or by altering NBLK_FCTR in config.h.
NCALL	NPROC	Number of callout table entries.
NCLIST	1000	Maximum number of <code>clists</code> . These are terminal I/O buffers.
NCRBUF	4	Maximum number of checkpoint and restart buffers.
NC1INODE	NINODE	Number of in-core inodes for NC1FS (file system migration parameter). Must be equal to NINODE. This parameter is not in the ICMS.
NC_NAMLEN	15	Maximum size of path-name component that will be cached in the directory name look-up cache (DNLC).
NC_SIZE	1024	Number of entries in directory name look-up cache. This should be less than NINODE.
NEXECS	4	Maximum number of parallel <code>exec</code> operations with maximum arguments. This parameter is not in ICMS.
NFILE	2100	Maximum number of files that can be open at one time system-wide.
NFLOCKS	400	Number of file-lock regions.

Parameter	Default value	Description
NHBUF_FCTR	4	Number of hash entries in the hash table for system cache, based on number of blocks assigned to system cache. The default of 4 specifies 1 hash list for every 4 system buffers.
NINODE	1500	Number of unique files that can be open at one time. Also the size of the in-core inode table.
NLDCH	0 or 2000	Number of SSD logical device cache headers, 0 for systems without SSD and 2000 for systems with SSD. This can be changed at boot time in the CSL parameter file.
NLDMAP	250	Number of file systems open concurrently.
NMNT	250	Number of slice structures.
NMOUNT	75	Number of file systems (local, NFS, DFS, and so on) that can be mounted.
NPBUF	500	Number of physical I/O headers. This can be changed at boot time in the CSL parameter file.
NPLCHCTL	0	Number of partition cache headers. This can be changed at boot time in the CSL parameter file. This parameter is necessary only if you specify the <code>-M CMEM</code> option of the <code>pcache(8)</code> command.
NPROC	650	Maximum number of processes that may be active simultaneously.
NPTY	128	Number of pseudo terminals (tty/pty pairs). This is also the maximum number of active terminal sessions.
NQUOTA	1400 or 0	If your site is running file system quotas, 1400 is the default number of in-core quota entries; otherwise, it is set to 0.
NSESS	300	Maximum number of sessions open simultaneously.
NSIDEBUF	48	Number of blocks for the SSD side door buffer. All sites with an SSD should have a side door buffer assigned. This also applies to Cray SV1ex systems.
NSU_ASYN	4	Number of asynchronous I/O headers reserved for the system.

Parameter	Default value	Description
NTEXT	100	Maximum number of shared-text programs that can be running simultaneously.
NUSERS	200	Maximum number of users defined for the fair share scheduler.
PLCHCORE	0	Memory clicks reserved for partition cache blocks. These are used for partition cache blocks specified as type MEM. This can be changed at boot time in the CSL parameter file.
TAPE_MAX_PER_DEV	65536	Maximum number of bytes per device for buffered I/O.
TAPE_MAX_CONF_UP	4	Maximum number of tapes that may be configured up.
XTRASEC	3	Number of CPU seconds that a process or session may use following receipt of a SIGCPULIM signal before the SIGKILL signal is sent to terminate the process or session.

Table 2. Kernel Parameters in config.h (Common) Not in ICMS

Parameter	Default value	Description
ACNICE	HZ	The lowest process nice value is saved as pc_acnice after the specified number of seconds of CPU time.
FSLG_BUFSIZE	400	Number of file system log records.
K_OPEN_MAX	16384	Maximum number of open files per process.
KSTACKSIZE	1500 or 2000	Kernel stack size, 2000 if you have the DCE Distributed File Service (DFS) installed, 1500 if you do not.
		 Caution: Do not change this value. This parameter does not appear in ICMS because it should not be changed.
MAX_UNLINKED_BYTES	26214400	Maximum bytes allowed in unlinked files for checkpoint and restart.
NC1MINRAW	20	Multiplier for automatic mixed buffer.

Parameter	Default value	Description
NC1INODE	NINODE	Number of in-core inodes for NC1FS (file system migration parameter). Must be equal to NINODE.
NEXECS	4	Maximum number of parallel exec operations with maximum arguments.
U_MAXPACK	16	Maximum number of outstanding user packets allowed. For Cray SV1 series and Cray J90se systems only.

Table 3 lists the dynamic kernel memory parameters.

Table 3. Dynamic Kernel Memory Parameters

Parameter	Default value	Description
KM_CHM_PADSZ	$((KSTACKSIZE + KM_WPU - 1) \gg KM_WPU_SHIFT) * (NPROC / 2)$	Size of the memory padding. The total size (in KM_UNITS) must be large enough for NPROC/2 kernel stack entries. Compilation of lowmem.c will fail if this constraint is not met.
KM_EXPAND_UNITS	128	The number of units each expansion will add. Because expansion space is acquired from coremap, it must be a multiple of coremap units. That is: $(KM_EXPAND_UNITS * KM_WPU) == \text{multiple of } (MEMKLIK * NWPC)$
KM_NO_THRASH	$(4 * KM_EXPAND_UNITS)$	The number of units that must be reached before a contraction will take place. This prevents contractions from reclaiming expansion space that may be requested within a short time span. Expansions can be expensive if they require shuffling memory.


Parameter	Default value	Description
		<p>An expansion may require moving the bitmap, which is stored in dynamic kernel memory, in order to increase its capacity. The frequency of such moves is dependent upon the values of <code>KM_WPU</code> and <code>KM_EXPAND_UNITS</code>.</p> <p>For example, if <code>KM_EXPAND_UNITS</code> is 512, each expansion requires 8 additional words of bitmap memory to manage the space. If <code>KM_WPU</code> is 16, this will cause the bitmap to be reallocated and moved every two expansions. If <code>KM_WPU</code> is 128, the bitmap will need to be reallocated every 8 expansions.</p> <p>In a given configuration of 640 initial units, 128 words per unit, and 128 units acquired with each expansion, the bitmap will not move until the 120th expansion; that is, 640 units divided by 64 bits per word results in 10 words required being initially required for the bitmap. However, since the bitmap is also allocated in an area of memory equal to the value of <code>KM_WPU</code>, 128 words are reserved for the bitmap. This leaves 118 extra words in the bitmap unit. Since each expand of 128 units requires an additional 2 words, 59 expands can be done before the bitmap capacity is exceeded and the bitmap will be grown (and moved if necessary).</p>
<code>KM_UNITS</code>	1408	<p>Initial dynamic kernel memory units from <code>coremap</code>.</p> <p> Caution: The value of <code>(KM_UNITS * KM_WPU)</code> should be a multiple of <code>(MEMKLIK * NWPC)</code> to avoid wasting space acquired from <code>coremap</code>. That is, at boot time, a chunk of memory is acquired from <code>coremap</code> to be managed at a finer granularity.</p>
<code>KM_WPU</code>	128	Words per unit. This value must be a power of 2.
<code>KM_WPU_SHIFTC</code>	7	Number of shifts to multiply or divide by the <code>KM_WPU</code> value.

Table 4 lists the parameters in `config.h` for the Cray T90 series shared memory feature configuration. These parameters are significant only for the Cray T90 series.

Table 4. Shared Memory Kernel Parameters in `config.h`

Parameter	Default value	Description
SHMMAX	128	Maximum size in clicks of a shared segment. Because memory is allocated in <code>MEMKLIK</code> units, this should be a multiple of <code>MEMKLIK</code> . This value must be greater than or equal to 0.
SHMMIN	1	Minimum size in clicks of a shared segment. This value must be greater than or equal to 0 and less than the <code>SHMMAX</code> value.
SHMMNI	20	Maximum number of shared memory identifiers available concurrently in the system (that is, the size of the <code>shmem</code> table). This value must be greater than 0.
SHMSEG	2	Maximum number of shared segments that a process can have attached simultaneously. This value must be 0, 1, or 2. (The upper bound is a hardware restriction.)

Table 5 lists the parameters in `config.h` that are used for setting process error thresholds. These values represent the number of allowable errors of each type (program range, operand range, and error exit) per connect to a CPU. The counters are reset on each pass through `resume()`. The default value is 500 for all systems.

If the values are set too high, the counts may never be reached (especially on a busy system), and the failing process will not be aborted. The definition of unreachable will vary according to application and system load. Setting the values to 0 disables this feature.

Table 5. Kernel Parameters in config.h for Process Error Thresholds

Parameter	Default value	Description
MAXUSRERR	500	Maximum number of error exits that may be encountered before the process is aborted. Setting the value to disables this feature.
MAXUSRORE	500	Maximum number of operand range errors that may be encountered before the process is aborted. Setting the value to disables this feature.
MAXUSRPRE	500	Maximum number of program range errors that may be encountered before the process is aborted. Setting the value to 0 disables this feature.

Table 6 lists the parameters in config.h that are used for interprocess (IPC) semaphore defines and message defines.

Table 6. IPC Kernel Parameters in config.h

Parameter	Default value	Description
MSGMAX	2048	Maximum size of a message in bytes. This value must be greater than 0 and less than the MSGMNB value.
MSGMNB	4096	Maximum number of bytes that may be in a message queue. This value must be greater than 0 and less than 999999.
MSGMNI	40	Maximum number of message queues that may be in the system at one time (that is, the size of the msgque table). This value must be greater than 0.
MSGSEG	1024	Number of message segments. The size of the message area is MSGSEG * MSGSSZ bytes. (This will be rounded up to MEMKLIK.) This value must be greater than 0 or equal to the MSGTQL value. The upper limit is 9999.

Parameter	Default value	Description
MSGSSZ	32	Message segment size. The message area is split up into segments of size MSGSSZ bytes. Each message is allocated the number of segments required to hold the message. This value must be greater than 0 and less than the MSGMAX value, and it must be a multiple of 8 bytes.
MSGTQL	100	Number of message headers in the system. There is one message header per active message. This value must be greater than or equal to the MSGMNI value. The upper limit is 9999.
SEMAEM	16384	Maximum adjust-on-exit value that can be set on a <code>semop(2)</code> system call by specifying the <code>SEM_UNDO</code> flag. This value must be greater than 0 and less than the <code>SEMVMX</code> value.
SEMMNI	40	Maximum number of semaphore sets that may be in the system at one time (that is, the size of the <code>sema</code> table). This value must be greater than 0 and less than 1000.
SEMMNS	100	Number of semaphores in the system (that is, the size of the <code>sem</code> table). A semaphore set may consist of one or more semaphores. This value must be greater than or equal to the <code>SEMMNI</code> value.
SEMMNU	40	Number of undo structures in the system (that is, the size of the <code>semu</code> table). One entry is used by each process that registers semaphore operations to be undone at process termination. This value must be greater than 0 and less than the <code>NPROC</code> value. A value close to the <code>SEMMNI</code> value will conserve table space. (The undo structure size in words is $3 * SEMUME + 2$.)
SEMMSL	25	Maximum number of semaphores per set. This value must be greater than 0 and less than the <code>SEMMNI</code> value.

Parameter	Default value	Description
SEMOPM	25	Maximum number of operations per <code>semop(2)</code> system call. Operations can be performed on the semaphores within one set. This value must be greater than 0 and less than or equal to the <code>SEMMSL</code> value.
SEMUME	20	Maximum number of undo entries per process. Each undo entry in the undo structure allows for an undo operation on one semaphore in a semaphore set. This value must be greater than 0 and less than or equal to the <code>SEMMNS</code> value. To conserve table space, set <code>SEMUME</code> to a value less than the <code>SEMMNI</code> value. (The undo structure size in words is $3 * SEMUME + 2$.)
SEMVMX	32767	Maximum value of an individual semaphore. This value must be greater than 0 and less than $2^{**}32$.

Table 7 lists the parameters in `config.h` for the UNICOS multilevel security (MLS) feature. These parameters are significant only if your site enables the UNICOS MLS feature. See *General UNICOS System Administration*, for more information on setting these parameters and on the UNICOS MLS feature.

Note: In general, a value of 0 disables a parameter and a value of 1 enables it.

Table 7. MLS Kernel Parameters in `config.h`

Parameter	Default value	Description
COMPART_ACTIVE_DEFAULT	0	Generic MLS defaults.
COMPART_VALID_DEFAULT	0	Generic MLS defaults.
CONSOLE_MSG	0	Capability to send a message to <code>/dev/console</code> when any user exceeds <code>MAXLOGS</code> login attempts.
DECLASSIFY_DISK	0	Deactivates declassification of the disk; activates when the value is not 0.
DECLASSIFY_PATTERN	0	Declassifies disk write pattern.
DELAY_MULT	0	Multiplier for failed logins. When set to 1 (enabled), the next login attempt is delayed by (<i># of failed logins</i>) * (<code>LOGDELAY</code>) seconds. Otherwise, the delay period is not multiplied by any factor, and the next login is delayed <code>LOGDELAY</code> seconds.
DEV_ENFORCE_ON	0	Capability to enforce strict device labeling.
DISABLE_ACCT	0	Disables the login account when <code>MAXLOGS</code> is exceeded. Used with <code>DISABLE_TIME</code> .
DISABLE_TIME	-1	Duration in seconds for which the user is disabled when the <code>MAXLOGS</code> is exceeded. If -1 (or any negative value), the user is disabled indefinitely.
FORCED_SOCKET	NORMAL	When set to <code>NORMAL</code> , <code>syslogd(8)</code> can use sockets and pipes.
FSETID_RESTRICT	1	If nonzero, only a privileged process can create <code>setuid / setgid</code> files.
LOGDELAY	0	Maximum number of seconds to delay between login attempts.
MAXLOGS	0	Maximum number of failed login attempts allowed before a user is disabled if the <code>DISABLE_ACCT</code> is set to 1.
MAXSLEVEL	0	Maximum security level allowed on the system.
MINSLEVEL	0	Minimum security level allowed on the system.
MLS_INTEGRITY	NORMAL	Integrity code in <code>secure.c</code> . This setting is not in the ICMS.

Parameter	Default value	Description
MLS_OBJ_RANGES	NORMAL	Allow object ranges outside the system label range in <code>mount</code> and <code>setdevs</code> .
NFS_REMOTE_RW_OK	1	Enables (nonzero) or disables (zero) NFS-mounting of a remote file system in read-write mode.
NFS_SECURE_EXPORT_OK	1	Enables (nonzero) or disables (zero) exporting a secure file system with NFS.
NFS_SECURE_PORTMON	1	Enables (nonzero) or disables (zero) whether NFS clients are required to use privileged ports (<code>ports < IPPORT_RESERVED</code>) in order to get NFS services. This, along with <code>NFS_PORTMON</code> (see Table 9, page 28), must be disabled for non-privileged access on MLS systems.
OVERWRITE_COUNT	3	Declassifies disk overwrite count.
PASS_MAXSIZE	8	Maximum size for machine-generated passwords.
PASS_MINSIZE	8	Minimum size for machine-generated passwords.
PRIV_SU	1	When set to 1, root (UID 0) has privilege.
RANDOM_PASS_ON	0	When set to yes (1), machine-generated passwords are enabled for all users. Setting to no (0) disables this.
SANITIZE_PATTERN	0	The pattern used to scrub disks if <code>SECURE_SCRUB</code> is enabled. The default pattern is 0.
SECURE_NET_OPTIONS	(<code>NETW_SOCKET_COMPAT</code> <code>NETW_RCMD_COMPAT</code>)	

Parameter	Default value	Description
		<p>Options for running TCP/IP on a secure system. The value may be any combination of the following:</p> <ul style="list-style-type: none"> • <code>NETW_STRICT_B1</code>, in which strict B1 evaluation rules are applied by the system; therefore, TCP sessions are restricted to a single label, regardless of the type of remote host. • <code>NETW SOCK_COMPAT</code>, in which sockets are automatically made multilevel if the creating process has <code>PRIV_SOCKET</code>. • <code>NETW_RCMD_COMPAT</code>, in which traditional <code>hosts.equiv</code> and <code>.rhosts</code> behavior is allowed. <p>If <code>NETW_RCMD_COMPAT</code> is not used, remote logins with <code>rlogin(1)</code> to <code>root</code> are disallowed and all other logins must be to the same user name and require the same user ID on both systems.</p> <p>The default value of represents the combination of <code>NETW SOCK_COMPAT</code> and <code>NETW_RCMD_COMPAT</code>. To add the strict B1 evaluation rules, you should add the <code>NETW_STRICT_B1</code> value; that is, <code>(NETW SOCK_COMPAT NETW_RCMD_COMPAT NETW_STRICT_B1)</code>.</p>
<code>SECURE_MAC</code>	0	When set to 1, enforces system high/system low MAC.
<code>SECURE_OPERATOR_CONSOLE</code>	<code>"/dev/console"</code>	Secure console for use by the MLS operator. This parameter is unused.
<code>SECURE_PIPE</code>	NORMAL	Enables/disables "read down" on pipes.
<code>SECURE_SCRUB</code>	NORMAL	Enables (SECURE) or disables (NORMAL) scrubbing of data blocks on file deletion.
<code>SECURE_SYSTEM_CONSOLE</code>	" "	Secure console for MLS administration. This parameter is unused.

Parameter	Default value	Description
SLG_ACT_NFS	SLGOFF	Enables (SLGON) or disables (SLGOFF) logging of NFS activity.
SLG_ACT_NQS	SLGOFF	Enables (SLGON) or disables (SLGOFF) logging of NQS activity.
SLG_ALL_NAMI	SLGOFF	Enables (SLGON) or disables (SLGOFF) logging of all mkdir, rmdir, link, and rm calls.
SLG_ALL_RM	SLGOFF	Enables (SLGON) or disables (SLGOFF) logging of all remove requests.
SLG_ALL_VALID	SLGOFF	Enables (SLGON) or disables (SLGOFF) logging of all access requests.
SLG_BUFSIZE	163840	Size of the security log buffer. The default value holds approximately 1000 security log records. Must be a multiple of 4.
SLG_CF_NET	SLGOFF	Enables (SLGON) or disables (SLGOFF) logging of network configuration changes.
SLG_CF_NQS	SLGOFF	Enables (SLGON) or disables (SLGOFF) logging of NQS configuration changes.
SLG_CF_UNICOS	SLGON	Enables (SLGON) or disables (SLGOFF) logging of UNICOS configuration changes.
SLG_DIR	/usr/adm/sl	Full path name of directory where security logs are kept.
SLG_DISCV	SLGON	Enables (SLGON) or disables (SLGOFF) logging of discretionary access violations.
SLG_FILE	slogfile	File name of active security log.
SLG_FILEXFR	SLGON	Enables (SLGON) or disables (SLGOFF) logging of all file transfers.
SLG_FPREFIX	s.	Prefix of retired security logs.
SLG_JEND	SLGON	Enables (SLGON) or disables (SLGOFF) logging of job ends.
SLG_JSTART	SLGON	Enables (SLGON) or disables (SLGOFF) logging of job starts.
SLG_LINKV	SLGON	Enables (SLGON) or disables (SLGOFF) logging of all link violations.

Parameter	Default value	Description
SLG_LOG_AUDIT	SLGON	Enables (SLGON) or disables (SLGOFF) logging of audit criteria changes.
SLG_LOG_CHDIR	SLGON	Enables (SLGON) or disables (SLGOFF) logging of all <code>chdir(1)</code> requests.
SLG_LOG_CRL	SLGOFF	Enables (SLGON) or disables (SLGOFF) logging of Cray/REELlibrarian activity.
SLG_LOG_DAC	SLGON	Enables (SLGON) or disables (SLGOFF) logging of discretionary access control (DAC) activities.
SLG_LOG_IPNET	SLGON	Enables (SLGON) or disables (SLGOFF) logging of Internet Protocol (IP) layer activity.
SLG_LOG_OPER	SLGON	Enables (SLGON) or disables (SLGOFF) logging of operator actions.
SLG_LOG_PRIV	SLGOFF	Enables (SLGON) or disables (SLGOFF) logging of all privileges used in system calls.
SLG_LOG_SECSYS	SLGON	Enables (SLGON) or disables (SLGOFF) logging of non-inode security system calls.
SLG_LOG_SHUTDOWN	SLGON	Enables (SLGON) or disables (SLGOFF) logging of system shutdown.
SLG_LOG_STARTUP	SLGON	Enables (SLGON) or disables (SLGOFF) logging of system startup.
SLG_LOG_TAPE	SLGOFF	Enables (SLGON) or disables (SLGOFF) logging of tape activity.
SLG_LOG_TCHG	SLGON	Enables (SLGON) or disables (SLGOFF) logging of system time changes.
SLG_MANDV	SLGON	Enables (SLGON) or disables (SLGOFF) logging of mandatory access requests.
SLG_MAXSIZE	8192000	Size of security log to retire. When <code>SLG_FILE</code> is larger than this, it is moved to a file prefixed with <code>SLG_FPREFIX</code> .
SLG_MKDIRV	SLGON	Enables (SLGON) or disables (SLGOFF) logging of all <code>mkdir</code> violations.
SLG_NETWV	SLGON	Enables (SLGON) or disables (SLGOFF) logging of network access violations.

Parameter	Default value	Description
SLG_PATH_TRACK	SLGON	Enables (SLGON) or disables (SLGOFF) tracking of all path names on accesses.
SLG_PHYSIO_ERR	SLGOFF	Enables (SLGON) or disables (SLGOFF) logging of all physical I/O errors.
SLG_REMOVEV	SLGON	Enables (SLGON) or disables (SLGOFF) logging of all <code>rm</code> violations.
SLG_RMDIRV	SLGON	Enables (SLGON) or disables (SLGOFF) logging of all <code>rmdir</code> violations.
SLG_STATE	SLGOFF	Enables (SLGON) or disables (SLGOFF) security log.
SLG_SUID_RQ	SLGON	Enables (SLGON) or disables (SLGOFF) logging of <code>setuid</code> system calls.
SLG_SULOG	SLGON	Enables (SLGON) or disables (SLGOFF) logging of all <code>su(1)</code> command attempts.
SLG_T_PROC	SLGON	Enables (SLGON) or disables (SLGOFF) logging of trusted process activity.
SLG_USER	SLGOFF	Enables (SLGON) or disables (SLGOFF) logging of user names and passwords for failed login attempts.
SYSTEM_ADMIN_CONSOLE	<code>"/dev/console"</code>	Default console for MLS administration. The value of this parameter must be <code>/dev/console</code> .
SYSVCOMPS	0	Bit map that defines valid system compartments. By default, system compartments are turned off.

Table 8 contains the parameter in `config.h` that relates to TCP/IP. See the *UNICOS Networking Facilities Administrator's Guide*, for more information on setting this parameter.

Table 8. TCP/IP Kernel Parameter in `config.h`

Parameter	Default value	Description
TCP_NMBSPACE	3800	Controls the number of MBUFs (TCP/IP managed memory buffers). For information about appropriate values, see <i>UNICOS Networking Facilities Administrator's Guide</i> . This pool of buffers is allocated at boot time in module <code>/usr/src/uts/tcp/kern/uipc_mbuf</code> and cannot be increased while the system is running. The ICMS places this value in both the <code>config.h</code> file and the CSL parameter file. The value in the CSL parameter file overrides the value in <code>config.h</code> .

Table 9 shows the NFS kernel parameters in `config.h`.

Table 9. NFS Kernel Parameters in `config.h`

Parameter	Default value	Description
NFS3_ASYNC_MAX	64	The amount of data (in 4096-byte blocks) that will be written per file asynchronously. After this value is exceeded, all writes will be synchronous. (2000 * 4096 = 8,192,000 bytes.)
NFS3_ASYNC_TIME	120	The amount of time (in seconds) that data will be held in the NFS asynchronous write cache on the client.
NFS_DUPTIMEOUT	3	Time interval in seconds during which duplicates will not be replayed.
NFS_MAXDATA	32768	Maximum user data read or written by way of NFS.
NFS_MAXDUPREQS	1200	Size of the duplicate request cache. This value must be large enough so that the entry will still be there when the first retry comes in.

Parameter	Default value	Description
NFS_NUM_RNODES	256	Number of NFS nodes that may be read from (known as <i>rnodes</i>) at any one time in the kernel. Each active NFS file or directory requires an <i>rnode</i> . Files may have multiple requests outstanding at any one time. If the number of active NFS files exceeds the <code>NFS_NUM_RNODES</code> value, the <i>rnodes</i> are shared among active files.
NFS_PORTMON	0	If set non-zero, then NFS clients are required to use privileged ports (<code>ports < IPPORT_RESERVED</code>) in order to get NFS services.
NFS_PRINTINTER	0	Time interval between two messages on the console window when the server is not responding.
CNFS_STATIC_CLIENTS	8	Number of CNFS static client handles.
NFS_STATIC_CLIENTS	8	Number of NFS static client handles.
CNFS_TEMP_CLIENTS	8	Number of CNFS temporary client handles.
NFS_TEMP_CLIENTS	8	Number of NFS temporary client handles.

Some of the NFS kernel parameters have CSL boot-time equivalents, as shown in Table 10. These boot-time equivalents override the values that were set in the `config.h` file when the kernel was built. The CSL tags and their values will override the built-in kernel values, and change the performance of the UNICOS operating system.

Table 10. CSL Boot-time Equivalents for NFS Kernel Parameters

CSL tag	config.h
<code>nfs3_async_max</code>	<code>NFS3_ASYNC_MAX</code>
<code>nfs3_async_time</code>	<code>NFS3_ASYNC_TIME</code>
<code>nfs_duptimeout</code>	<code>NFS_DUPTIMEOUT</code>
<code>nfs_maxdata</code>	<code>NFS_MAXDATA</code>
<code>nfs_maxdupreqs</code>	<code>NFS_MAXDUPREQS</code>
<code>nfs_num_rnodes</code>	<code>NFS_NUM_RNODES</code>
<code>nfs_portmon</code>	<code>NFS_PORTMON</code>

CSL tag	config.h
nfs_printinter	NFS_PRINTINTER
cnfs_static_clients	CNFS_STATIC_CLIENTS
nfs_static_clients	NFS_STATIC_CLIENTS
cnfs_temp_clients	CNFS_TEMP_CLIENTS
nfs_temp_clients	NFS_TEMP_CLIENTS

config.mh File [4]

Note: It is strongly recommended that you use the install tool to maintain your system configuration, rather than manually editing this file. For more information on the UNICOS installation and configuration menu system (ICMS), refer to the online help files and to the *UNICOS System Configuration Using ICMS*.

The config.mh file contains values that are used to set the following:

- General system parameters
- Kernel generation parameters
- Kernel subsystem parameters,
- Programming environment parameters

Default values are either literals or 1 (which enables a subsystem) or 0 (which disables a subsystem). The following tables describe the parameters and their default values, grouped by type and listed in alphabetical order.

Note: In general, a value of 0 disables a parameter, and a value of 1 enables it.

These parameters are set with the following ICMS menus:

```
Configure System
  ->Major Software Configuration
```

```
Configure System
  ->Major Hardware Configuration
```

```
Build/Install System
  ->Build options
```

```
Configure System
  ->IOS Configuration (Model E based systems only)
```

Table 11. General System Parameters (Common)

Parameter	Default value	Description
CONFIG_DIAGDIR	<i>/ce</i>	Directory where the online diagnostics are kept. This must be a full path name.
CONFIG_ID	UNICOS	Identification of the operating system.
CONFIG_NODE	<i>node_name</i>	Node name of the system, by which it is known to a communications network. See <code>uname(1)</code> .
CONFIG_SN	<i>snnumber</i>	Serial number of the system to generate.
CONFIG_SYS	<i>system_name</i>	Name of the system. This information is used by the <code>uname(1)</code> command. Typically, <i>system_name</i> is the same as the serial number.
CONFIG_TARGET	(null)	Target of the system to generate. If null, the existing default is used. See <code>target(1)</code> .
CONFIG_TMPDIR	<i>/tmp</i>	Name of the desired temporary-file directory. If the <code>TMPDIR</code> environment variable is already defined, this value is not used. See <code>tmpnam(3)</code> .
CONFIG_VERSION	<i>version</i>	Version name of the system. If unspecified, it defaults to the value of the environment variable <code>LOGNAME</code> . See <code>uname(1)</code> .

Table 12. General System Parameters (GigaRing Based Systems Only)

Parameter	Default value	Description
CONFIG_ID	UNICOS	Identification of the operating system.
CONFIG_IOS_F	1	GigaRing support.
CONFIG_MK	0	UNICOS/mk operating system (used for Cray T3E systems only). The default is UNICOS.

Table 13. General System Parameters (Model E Based Systems Only)

Parameter	Default value	Description
CONFIG_ID	UNICOS	Identification of the operating system.
CONFIG_IOS_F	0	GigaRing support.
CONFIG_IOSA_SN	0	IOS-E serial numbers.
CONFIG_IOSB_SN	0	IOS-E serial numbers.
CONFIG_MK	0	UNICOS/mk operating system (used for Cray T3E systems only). The default is UNICOS.
CONFIG_NIOS	1	Number of IOS-E machines (0, 1, or 2).

Table 14. Kernel Generation Parameters

Parameter	Default value	Description
CONFIG_CAM_CPP_LOC	CONFIG_GENBIN/./reqs/cc/mppcpp	The environment variable CAM_CPP_LOCATION is set to the value of CONFIG_CAM_CPP_LOC. The cray-t3e assembler uses this variable.
CONFIG_CPP	CONFIG_GENBIN/./reqs/cc/cpp	Sets the nmake(1) CPP variable to the generation cpp.
CONFIG_CPSAVE	0	The cpset(8) -o option capability. Most released software does not use this option, in which case this parameter has no effect.
CONFIG_GCC	CONFIG_GENBIN/cc	Sets the C default to the generation cc compiler.
CONFIG_GEN_SEGDIR	CONFIG_GENBIN/./lib/segdirs	The environment variable GEN_SEGDIR is set to the value of CONFIG_GEN_SEGDIR. The segldr(1) command uses this variable.
CONFIG_GENBIN	/usr/gen/bin	Full path of the directory that contains the generation software.

Parameter	Default value	Description
CONFIG_GENCMDS	CONFIG_GCC CONFIG_CPP	A list of products that must exist as executables. The existence of the listed products is verified when any part of UNICOS is regenerated using <code>nmake(1)</code> .
CONFIG_GENPROD_RULES	0	Repeatable relocatables capability.
CONFIG_MPP_CPP	CONFIG_GENBIN/./reqs/cc/mppcpp	Sets the <code>nmake(1)</code> CPP variable to the generation <code>mppcpp</code> for the <code>cray-t3e</code> target.
CONFIG_PACKAGE	0	Certain <code>nmake(1)</code> targets are disabled when this is enabled. Used for packaging purposes only. 1 indicates that this is a packaging build.
CONFIG_PATH	CONFIG_GENBIN:/bin:/usr/bin:/usr/ucb	Generation software directory paths. The environment variable <code>PATH</code> is set to the value of <code>CONFIG_PATH</code> .
CONFIG_RLS_MAJOR	<i>integer</i>	UNICOS major release number.
CONFIG_RLS_MINOR	<i>integer</i>	UNICOS minor release number.
CONFIG_RLS_REVISION	<i>integer</i>	UNICOS revision release number.
CONFIG_SUPPORT_DIR	(null)	Full path of the directory that contains the support software.
CONFIG_TRGBIN	/usr/gen/trg	Full path of the directory that contains the targeting software.
CONFIG_UMASK	022	The <code>umask(1)</code> setting to be used during the build.
CONFIG_XLIBS	0	Build and install cross-targeted libraries. This is not available on Cray T90 IEEE mainframes.
CONFIG_XLIBTARGET	<i>target</i>	Desired set of cross-targeted libraries, such as <code>cray-c90</code> for use on a Cray T90 mainframe. This is not available on Cray T90 IEEE mainframes.

Parameter	Default value	Description
CONFIG_MIXED	0	Build and install cross-targeted libraries for a mixed mode (Cray floating-point and IEEE) Cray T90 CPU system.
CONFIG_MIXEDTARGET	<i>target</i>	Desired set of cross-targeted libraries for a mixed mode system, such as <code>cray-ts,ieee</code> for use on a Cray floating-point Cray T90 mainframe.

Table 15. Kernel Subsystem Parameters

Parameter	Default value	Description
CONFIG_BBG	0	Bus Based Gateway (BBG).
CONFIG_BMM	0	Bit matrix multiply functional unit (<code>uts</code> kernel).
CONFIG_CRL	0	Cray/REELlibrarian.
CONFIG_CVT	1	Cray Visualization Toolkit (CVT). If this parameter is set to 1, <code>CONFIG_X11</code> must also be set to 1.
CONFIG_DFS	1	Distributed Computing Environment (DCE) distributed file system (DFS).
CONFIG_DM	0	Cray Data Migration Facility (DMF).
CONFIG_ELS	0	Cray SV1 series and Cray J90se support.
CONFIG_FQUOTAS	1	File quotas.
CONFIG_HPI3	0	IPI-3/HIPPI packet driver capability in the kernel.
CONFIG_HSX	0	HSX device driver (<code>uts</code> kernel).
CONFIG_IPI3	1	IPI-3/IPI capability in the kernel.
CONFIG_KERBEROS	0	Kerberos support.
CONFIG_MPP	1	CRAY MPP system is attached.
CONFIG_NETMON	1	Network monitor.
CONFIG_NETTOLS	0	Network testing tools.
CONFIG_NFS	1	Network File System (NFS).

Parameter	Default value	Description
CONFIG_NFS3	1	NFS version 3 Protocol (NFS3). CONFIG_NFS must be configured if this is set to 1.
CONFIG_NFSKRB	0	NFS with Kerberos authentication. If this parameter is set to 1, CONFIG_NFS and CONFIG_KERBEROS must also be set to 1.
CONFIG_OWS	1	Operator workstation (OWS). This parameter is not used and is provided for compatibility purposes.
CONFIG_RPC	1	Remote process control system (RPC).
CONFIG_TAPE	1	Online tape capability in the kernel.
CONFIG_TCP	1	TCP/IP network system. This parameter must always be set to 1.
CONFIG_TRUSTED	0	Trusted UNICOS.
CONFIG_X11	1	X11 Window Management System.
CONFIG_YP	0	Yellow pages.

Table 16. Programming Environment Parameters

Parameter	Default value	Description
CONFIG_CRAYLIBS	CONFIG_GENBIN/./lib	The environment variable GEN_CRAYLIBS is set to the value of CONFIG_CRAYLIBS. The compilers and assembler use this variable.
CONFIG_LD_STD_DIR	CONFIG_GENBIN/./lib/cld	The environment variable GEN_LD_STD_DIR is set to the value of CONFIG_LD_STD_DIR. The cld script uses this variable.

CSL Parameter File [5]

This chapter describes the configuration specification language (CSL) parameter file, which contains statements that specify hardware and software characteristics for your system.

Note: It is strongly recommended that you use the installation tool to maintain your system configuration, rather than manually editing this file. For more information on the UNICOS installation and configuration menu system (ICMS), see the online help files and the *UNICOS System Configuration Using ICMS*

When the configuration is activated, a copy of the CSL parameter file is stored in `/etc/config/param`. For all Cray systems, the administrator must manually transfer the parameter file to the workstation.

5.1 CSL Syntax

There are three classes of tokens that make up CSL:

- Identifiers
- Constants
- Operators, separators, and comments

White space (horizontal tabs, new lines, carriage returns, and spaces) separates individual tokens.

5.1.1 Identifiers

An *identifier* is a sequence of digits and letters that specify either special keywords (such as `CONFIG`) or specific objects (such as a physical device). The digits and characters can be enclosed by double quotes. The underscore (`_`) and dash (`-`) are interpreted as letters. Uppercase and lowercase letters are interpreted differently; that is, CSL identifiers are case-sensitive. There are no restrictions on the first character of an identifier.

For example, each of the following is a valid identifier:

```
tmp
STI
abc_d
```

```
29-A1-31
Dump
0x1246
"507"
_xyz
```

There are two classes of identifiers:

- *Keyword identifiers* have special meaning in CSL and cannot be used to name other things. For a list of reserved keywords, see:
- *Object identifiers* name specific objects. Objects are divided into three classes:
 - Physical devices
 - Logical devices
 - Slices

Each class of object has its own name space. That is, each object in a given class must have a unique name, but objects in different classes can share the same name.

5.1.2 Constants

All constants are positive integers. An integer consists of 1 digit or a sequence of digits. If the first digit of a constant is (zero), the constant is interpreted as octal; otherwise, the constant is assumed to be decimal. The use of digits 8 or 9 in an octal constant causes an error.

The following are all examples of valid constants:

```
012345      12      44673      0003455
```

5.1.3 Operators, Separators, and Comments

The allowable operators and separators in CSL are as follows:

```
{          }          ;          '          ,
```

The meanings of these operators and separators depend on the context in which they are used.

To intersperse comments between objects, begin and end the comment text as follows:


```
/* This is the comment text. */
```

5.2 CSL Usage

This section provides general information on CSL usage.

CSL statements are located in the CSL parameter file. CSL statements are used to extend or override the default configuration.

All statements in the CSL parameter file must be terminated by a semicolon.

5.2.1 The Boot Process

UNICOS processes CSL statements in order of appearance in the CSL parameter file at boot time.

When processing is finished, a copy of the CSL file is placed in the `/CONFIGURATION` file.

5.2.2 Configuration Verification

You can verify configurations by using the following menu selection:

```
Configure System
->Disk Configuration
   ->Verify the disk configuration
```

You can also use the user-level program `econfig(8)`, which shows error messages for any errors in the configuration.

The `econfig` program accepts only valid CSL statements as input. It is recommended that you verify the CSL statements by using the `econfig` command before booting a new configuration to prevent receiving errors during CSL processing.

5.2.3 Error Message Syntax

If, while processing CSL statements, UNICOS detects a condition warranting your attention, a message describing the condition is written to the system console. Conditions warranting a message are divided into three levels of severity:

- Informational messages, which do not affect the boot process.

- Error messages, which inform you of a problem in statement syntax or configuration consistency that prevents the system from booting. Directive processing continues, but the system panics when statement processing completes.
- Panic messages, which inform you of a problem that prevents the completion of statement processing. The system panics immediately.

The syntax of a condition message is as follows:

severity: message_text, location

<i>severity</i>	Values are INFO, ERROR, and PANIC.
<i>message_text</i>	The message proper.
<i>location</i>	Problem location. The <i>location</i> is expressed as a line number in the CSL parameter file.

5.3 CSL Parameter File Sections

The statements in the CSL parameter file are organized in the file by functionality:

<u>Section</u>	<u>Description</u>
dumpinfo	Cray IOS Model E based system dump parameters
gigaring	GigaRing configuration parameters
ios_e	Cray IOS Model E based system and Cray SV1 Model V configuration parameters
mainframe	Physical mainframe characteristics parameters
unicos	UNICOS kernel parameters
filesystem	Physical storage devices and file system layout parameters
network	Network parameters and device parameters
mpp	Massively parallel processing (MPP) parameters (Model E based systems only)
revision	Revision identifier parameters
t90_config	Cray T90 configuration parameters



Caution: There are specific naming requirements for naming dump devices and dump locations. For information about these requirements, see *General UNICOS System Administration*.

Note: The default parameter file contains sections that are I/O-specific and therefore will not be used by all systems. You should remove the unused sections from your mainframe's parameter file to avoid potential problems:

- For GigaRing based systems, remove the `dumpinfo` and `ios_e` sections.
- For Cray Model E based systems, remove the `gigaring` section.

5.3.1 dumpinfo Section



Caution: There are specific naming requirements for naming dump devices and dump locations. For information about these requirements, see *General UNICOS System Administration*.

Note: This section does **not** apply to GigaRing based systems or to Model V based systems.

The `dumpinfo` section of the CSL parameter file defines the types and ranges of memory to dump when invoking the operator workstation (OWS) `dumpsys(8)` command.

The `dumpinfo` section is specified in the CSL parameter file using the following syntax:

```
dumpinfo {
    range_declaration_1;
    range_declaration_2;
    .
    .
    .
}
```

A range declaration is specified using the following syntax:

```
type range is start to stop units ;
```

type Either `memory` or `SSD`, indicating central memory or SSD memory to be dumped.

<i>start</i>	Integer indicating the beginning location of a memory section to be dumped.
<i>stop</i>	Integer indicating the ending location of a memory section to be dumped.
<i>units</i>	Units in which the start and stop values are to be computed: words, Mwords, or Gwords.

Up to four declarations of each *type* are allowed in the `dumpinfo` section.

The following example shows a typical `dumpinfo` section of a CSL parameter file:

```
dumpinfo {  
    memory range is 0 to 16 Mwords;  
}
```

Note: If the first memory range specification is not large enough to contain the entire kernel space, the mainframe dump routine will override the site specification and try to dump the full kernel space when a system dump is performed. This can lead to a truncated dump if the dump device is not large enough to contain the expanded dump. A truncated dump will not include executing exchange packages, CPU registers, or user areas. This reduces the usefulness of the dump. To override this behavior, and to enforce the site-specified values, insert a memory range specification of 0 to 0 as the first range. Sites that have LDCHCORE, PLCHCORE, and/or a random access memory (RAM) disk configured commonly experience this, because the LDCHCORE, PLCHCORE, and RAM disk space is included within the kernel space.

For example:

```
dumpinfo {  
    memory range is 0 to 0 Mwords;  
    memory range is 0 to 16 Mwords;  
}
```

5.3.2 gigaring Section

Note: This section does **not** apply to Cray Model E based systems or Cray Model V based systems.

All Cray SV1ex systems must follow some minimum configuration rules; see Section 1.1, page 2.

GigaRing configuration in UNICOS is done in two places:

- Internal routing and path selection is done in the `gigaring` section of the parameter file.
- Physical channel designation is done in the `mainframe` section of the parameter file.

The `gigaring` section consists of two subsections:

- `gr_route`
- `gr_union`

5.3.2.1 `gr_route` Subsection

The `gr_route` subsection provides a means of internal routing and path selection known as *source routing*. Source routing chooses the channel used to route traffic to a given ring. Where more than one mainframe GigaRing channel exists on a ring, messages are routed in a round-robin fashion.

Routing information is declared in the `gr_route` subsection. For example:

```
gigaring {
    gr_route {
        ring 06 {
            channel 024;
        }
        ring 07 {
            channel 034;
            channel 044;
        }
    }
}
```

There can be up to 8 routes per ring. By default, the first route declared is designated as the primary route. Valid ring numbers are decimal integers in the range 0 through 127. Valid node numbers are decimal integers in the range 1 through 63.

5.3.2.2 `gr_union` Subsection

A *GigaRing union* is a logical representation of a device that has more than one ring and node designation. For example, a Cray SSD-T90 with four GigaRing connections (and therefore four ring and node addresses) can be represented by one `gr_union` ring and node address. This applies only to devices for which the mainframe acts as the master for the direct memory access (DMA) operation.

This applies to the Cray SSD-T90 and also to the Cray SSD-I on Cray SV1ex systems.

By convention, ring 0200 is reserved for GigaRing union devices. There is a maximum of 16 nodes (node numbers through 017) reserved for these devices. Devices that are configured as GigaRing union devices allow their device drivers to query the low-level master DMA driver for physical ring and node addresses. The device driver can then route the DMA requests by targeting one physical ring/node address. This is known as *destination routing*. DMA requests can be scheduled by targeting the least busy channel. The retrying of requests in error can be targeted to another destination.

The GigaRing union device allows for ease of configuration and backward compatibility with UNICOS device node methodology.

An example of a `gr_union` declaration is as follows:

```
gigaring {
    gr_union {
        ring 0200, node 01 {
            ring 06, node 04;
            ring 06, node 05;
            ring 07, node 04;
            ring 07, node 05;
        }
    }
}
```

In this example, a GigaRing union logical device designated as ring 0200, node 01, will be created and will consist of four physical destinations.

5.3.3 `ios_e` Section

Note: This section does not apply to GigaRing based systems.

The `ios_e` section defines the topology of the IOS-E hardware. Specify the characteristics using the following menu selection:

```
Configure System
->IOS configuration
```

This section includes the following information:

- Number of clusters that constitute the IOS-E

- Number and type of I/O processors (IOPs) in each cluster
- High-speed (HISP) channel information (channel, target, and operating mode)
- Full OWS path names for each IOP binary

The `ios_e` section is specified in the CSL parameter file by the following statement:

```
ios_e { list of cluster declarations }
```

A cluster declaration is specified by the following statement:

```
cluster ordinal { list of cluster statements }
```

The following types of IOS cluster statements exist:

- IOP declarations
- HISP declarations
- IOP boot declarations

5.3.3.1 IOP Declarations

IOP declarations are specified by the following statement:

```
iop;
```

The following are valid IOPs:

- `muxiop` or `miop`
- `eiop 0`
- `eiop 1`
- `eiop 2`
- `eiop 3`
- `eiop 4`

Note: `eiop 4`, `muxiop`, and `miop` are equivalent declarations except for Model V based systems. For Model V based systems, `eiop` indicates the VME controller type and may range from 0 through 50. See the *UNICOS Basic Administration Guide for Cray J90se and Cray SV1 Series Systems*.

5.3.3.2 HISP Declarations

Note: This section does **not** apply to Cray Model V based systems.

HISP declarations must be correct for the mainframe type on which they are declared, or system problems may occur. HISP declarations are specified by the following statement:

```
channel ordinal is hisp ordinal to chan_target , mode chan_type ;
```

HISP channel targets are as follows:

- mainframe
- SSD

The only channel type is C90, which means 200 Mbyte control / 200 Mbyte data.

5.3.3.3 IOP Boot Declarations

Note: This section does **not** apply to Cray Model V based systems.

IOP boot declarations are specified by the following statement:

```
boot iop with "pathname" ;
```

If the *pathname* is fully resolved (to the root directory), it is used. If a partial path name is specified, it is appended to the OWS value for `ROOTDIR`, as specified in the OWS configuration file.

5.3.3.4 Example of `ios_e` Section

The following example shows the `ios_e` section of a CSL parameter file:

```
ios_e {  
    cluster 0 {  
        miop; eiop 0; eiop 1; eiop 2; eiop 3;    }  
}
```



```

channel 010 is hisp 0 to mainframe, mode C90;
channel 014 is hisp 1 to SSD          , mode C90;

boot miop with "/home/ows7001/cri/os/ios/iopmux";
boot eiop 0 with "/home/ows7001/cri/os/ios/eiop.comm";
boot eiop 1 with "/home/ows7001/cri/os/ios/eiop.bmx";
boot eiop 2 with "/home/ows7001/cri/os/ios/eiop.dca2";
boot eiop 3 with "/home/ows7001/cri/os/ios/eiop.dca1";
}
cluster 1 {
miop; eiop 0; eiop 1; eiop 2; eiop 3;
channel 010 is hisp 0 to mainframe, mode C90;
channel 014 is hisp 1 to SSD          , mode C90;

boot miop with "/home/ows7001/cri/os/ios/iopmux";
boot eiop 0 with "/home/ows7001/cri/os/ios/eiop.dca1";
boot eiop 1 with "/home/ows7001/cri/os/ios/eiop.dca1";
boot eiop 2 with "/home/ows7001/cri/os/ios/eiop.dca2";
boot eiop 3 with "/home/ows7001/cri/os/ios/eiop.hippi";
}
}

```

5.3.4 mainframe Section

The mainframe section defines the following hardware parameters:

- I/O modules
- Number of CPUs (less those used as I/O modules)
- Number of multistreaming processors (MSPs)
- Number of mainframe cluster registers
- Size of the mainframe memory
- Size of the SSD-I memory
- Channel information
 - Physical channel for a GigaRing environment
 - Low-speed channel information (channel and target) for Model E based systems

- Very high speed (VHISP) channel information for Model E based systems

Set these parameters by using the following menu selection for Model E based systems:

```
Configure System
  ->Mainframe hardware configuration
```

The mainframe section is specified in the CSL parameter file by the following statement:

```
mainframe { list of hardware definitions }
```

5.3.4.1 Common Parameters

The following parameters are common to GigaRing based systems and Model E based systems:

- Number of CPUs (less those used as I/O modules)
- Number of mainframe cluster registers
- Size of memory
- Down CPUs
- CPU Cache

5.3.4.1.1 Number of CPUs

The number of CPUs is specified by the following statement:

```
value cpus ;
```

value is the physical number of CPUs in the machine. Specifying a smaller value will cause UNICOS to use only that many CPUs, starting at CPU 0.

5.3.4.1.2 Number of Mainframe Cluster Registers

The number of mainframe cluster registers is specified by the following statement:

```
value clusters ;
```

value is the number of clusters. If this is not specified, it defaults to 1 greater than the `cpus` value.

5.3.4.1.3 Size of Memory

The mainframe memory size is specified by the following statement:

```
value units memory;
```

units may be either `words` or `Gwords`. Typically, *value* is set to the physical amount of memory in the machine, but it can be set to a smaller value. This may be useful when experiencing memory problems.

Your customer engineer can tell you the memory configuration for your machine.

5.3.4.1.4 Size of Extended Memory

The Cray SV1ex mainframe extended memory size is specified by the following statement:

```
value units xmemory;
```

units can be `words`, `Mwords`, or `Gwords`. Typically, *value* is set to the physical amount of memory in the machine, but it can be set to a smaller value. This may be useful when experiencing memory problems.

Your customer engineer can tell you the extended memory configuration for your machine.

5.3.4.1.5 Down CPUs

The CPUs to be placed in a down condition at deadstart time are specified by the following statements:

```
down cpu cpuname;  
down cpus cpuname1 cpuname2 ...;
```

cpunameX is the logical number of the CPU.

5.3.4.1.6 CPU Cache

The ability to disable CPU cache, either system-wide or for the specified CPUs, is specified by the following statements:

```
cacheoff user;  
cacheoff kernel;  
cacheoff user cpu cpuname;  
cacheoff user cpus cpuname1 cpuname2 ...;  
cacheoff kernel cpu cpuname;  
cacheoff kernel cpus cpuname1 cpuname2 ...;
```

The `cacheoff user;` and `cacheoff kernel;` directives disable user cache or kernel cache on a system-wide basis. The remaining directives disable user cache or kernel cache for a specific CPU or set of CPUs. *cpunameX* is the logical number of the CPU on which cache will be disabled.

For example, the following statement will disable kernel cache in all CPUs and will disable user cache in CPUs 1, 3 and 21:

```
cacheoff kernel;  
cacheoff user cpus 1 3 21;
```

5.3.4.2 Parameters for Cray SV1 series GigaRing Based System

- Multistreaming Processors (MSPs)
- I/O Modules

5.3.4.2.1 Multistreaming Processors (MSPs)

The maximum number of multistreaming processors (MSPs) is specified by the following statement:

```
msp_number msp;
```

The variable *msp_number* specifies the maximum number of MSPs in the machine. *msp_number* is an integer in the range from 0 through 6.

The maximum number of MSPs configurable is constrained by the number of physical CPUs, as shown in Table 17.

Table 17. CPU/MSP Correlation

Number of CPUs	Maximum Number of MSPs
4	0
8	1
12	2
16	3
20	4
24	5
28	6
32	6

5.3.4.2.2 I/O modules

An I/O module is a CPU module that will be used primarily to provide increased I/O bandwidth and/or I/O capacity on a SV1 series system. An I/O module will have a GigaRing channel connection. I/O modules perform operating system services, such as I/O interrupt handling, clock interrupt handling, system process handling (such as network and utility). I/O modules are not available for user job assignment.

SV1 I/O modules are specified by the following statement:

```
iomodules module1 module2 ...;
```

modulenameX is the logical number of the CPU.

The variables *modulename1* and *modulename2* specify the CPU modules to use as I/O modules. For example, if the value in *modulename1* is 5 and *modulename2* is 6, the statement

```
iomodules 5 6;
```

configures CPU modules 5 and 6 (using 0-based module numbering) as I/O modules to perform operating system services. In this configuration, CPUs 16 - 23 would be used to perform operating system services.

5.3.4.3 Parameters for Cray SV1 series VME Systems Only

The channels to be used for memory high performance parallel interface (HIPPI) are specified by the following statement:

```
channel channel_number is memhippi to mainframe;
```

channel_number specifies memory HIPPI input channel.

For example:

```
channel 040 is memhippi to mainframe;
```

This parameter is required for all Cray SV1 series and VME systems.

5.3.4.4 Parameters for GigaRing Based Systems Only

The physical channel configuration declares a physical channel to be a GigaRing channel. It creates a GigaRing port by assigning a ring number and node number to a given mainframe channel number.

```
channel ordinal is gigaring to ring ring_number, node node_number ;
```

ordinal is the channel number. *ring_number* must be an integer in the range through 127. *node_number* must be an integer in the range 1 through 63. For every channel entry in the `gr_route` and `gr_union` sections, there must be a corresponding channel entry in the `mainframe` section.

At UNICOS boot time, a `cnode` structure is declared to represent each GigaRing port. The ring and node numbers will be read and verified from the GigaRing node, or, in the case of a direct connect, be set according to the ring and node numbers specified.

The following channel numbers are valid for Cray J90se GigaRing based systems and Cray SV1 series GigaRing based systems:

024	064
034	074
044	0104
054	0114

The following channel numbers are valid for Cray T90 GigaRing based systems (all even octal integers from 0100 through 0176):

0100	0120	0140	0160
0102	0122	0142	0162
0104	0124	0144	0164
0106	0126	0146	0166
0110	0130	0150	0170
0112	0132	0152	0172
0114	0134	0154	0174
0116	0136	0156	0176

Table 18 shows the channels for I/O module 0.

Table 18. I/O Module 0 Channels

Erred packet queue 0	Erred packet queue 1	Incoming packet queue 0	Incoming packet queue 1	Outgoing packet queue	DMA TIB/TCB buffer	Ring number
0100	0101	0200	0201	0300	0301	GR 0
0102	0103	0202	0203	0302	0303	GR 1
0104	0105	0204	0205	0304	0305	GR 2
0106	0107	0206	0207	0306	0307	GR 3
0110	0111	0210	0211	0310	0311	GR 4
0112	0113	0212	0213	0312	0313	GR 5
0114	0115	0214	0215	0314	0315	GR 6
0116	0117	0216	0217	0316	0317	GR 7

Table 19 shows the channel number ranges for a Cray T932 configuration of four IO2 modules.

Table 19. Example of Complete Channel Assignments

Erred packet queue 0	Erred packet queue 1	Incoming packet queue 0	Incoming packet queue 1	Outgoing packet queue	DMA TIB/TCB buffer	IO2 module number
0100-0116	0101-0117	0200-0216	0201-0217	0300-0316	0301-0317	IO2 - 0
0120-0136	0121-0137	0220-0236	0221-0237	0320-0336	0321-0337	IO2 - 1
0140-0136	0141-0157	0240-0256	0241-0257	0340-0356	0341-0357	IO2 - 2
0160-0176	0161-0177	0260-0276	0261-0277	0360-0376	0361-0377	IO2 - 3

5.3.4.5 Parameters for Model E Based Systems Only

The Channel declarations are for Cray Model E based systems only. Channel declarations are either for a low-speed channel pair or a very high speed (VHISP) channel. A channel declaration is specified by the following statement:

```
channel ordinal is channel_type to channel_target ;
```

ordinal is the channel number, for which the preferred format is octal. *channel_type* is either `lowspeed` or `VHISP`. A low-speed *channel_target* is the cluster ordinal or `pseudo TCP`. The VHISP *channel_target* is `SSD`. You no longer must specify `pseudo TCP` for TCP/IP to be operational. If `pseudo TCP` is included in a parameter file, the setting will be ignored by the UNICOS kernel during the booting of the mainframe.

5.3.4.6 Example of the mainframe Section for a GigaRing Based System

The following shows an example of the `mainframe` section of the CSL parameter file for a GigaRing based system:

```
mainframe {
    16 cpus;
    2 msp;
    1024 Mwords memory;
    channel 024 is gigaring to ring 6, node 1;
    channel 034 is gigaring to ring 7, node 5;
    channel 044 is gigaring to ring 7, node 6;
}
```


5.3.4.7 Example of the mainframe Section for a Model E Based System

The following shows an example of the mainframe section of the CSL parameter file in a Model E based system:

```
mainframe {  
    2 cpus;  
    32 Mwords memory;  
    channel 16 is lowspeed to cluster 0;  
    channel 18 is lowspeed to cluster 1;  
    channel 1 is vhispc 0 to SSD;  
}
```

5.3.5 unicos Section

The unicos section sets certain tunable parameters. Set these parameters by using the following menu selection:

```
Configure System  
->UNICOS Kernel Configuration
```

The unicos section is specified in the CSL parameter file by the following statement:

```
UNICOS { list of tunable parameters } ;
```

A UNICOS tunable parameter is specified by the following statement:

```
value parameter ;
```

All systems have the same tunable parameters for online tapes, table sizes, and maximum limits. Table 20 through Table 22 show these parameters.

Table 20. Online Tape Parameters

Parameter	Description
TAPE_MAX_CONF_UP	Maximum number of tape devices that can be configured up at the same time.
TAPE_MAX_DEV	Maximum number of tape devices.
TAPE_MAX_PER_DEV	Maximum number of bytes allocated per tape device.

Table 21. Table Size Parameters

Parameter	Description
LDCHCORE	Memory clicks reserved for ldcache blocks assigned as type MEM.
NLDCH	Number of ldcache headers.
NPBUF	Number of physical I/O buffers.

Table 22. Maximum Limits Parameters

Parameter	Description								
FAULT_RESPONSE	<p>Control the fault tolerance feature. The <i>value</i> is a bit mask that controls the various options in the feature. Each bit enables or disables an option, as follows:</p> <table border="0"> <tr> <td>Bit</td> <td>Option Controlled (1=on, 0=off)</td> </tr> <tr> <td>01</td> <td>Fault tolerance feature.</td> </tr> <tr> <td>02</td> <td>Panic when detecting a lock held by a down CPU.</td> </tr> <tr> <td>04</td> <td>Down CPU cache on a user cache parity error.</td> </tr> </table> <p>The following example disables the feature: 0 FAULT_RESPONSE</p> <p>The default for <i>value</i> is 07 (all options selected).</p> <p>The following example enables the basic feature and downing of CPU cache on a user cache parity error, but does not panic when detecting a lock held by a down CPU: 05 FAULT_RESPONSE</p>	Bit	Option Controlled (1=on, 0=off)	01	Fault tolerance feature.	02	Panic when detecting a lock held by a down CPU.	04	Down CPU cache on a user cache parity error.
Bit	Option Controlled (1=on, 0=off)								
01	Fault tolerance feature.								
02	Panic when detecting a lock held by a down CPU.								
04	Down CPU cache on a user cache parity error.								

Parameter	Description
GUESTMAX	Maximum number of guest systems.
NBUF	Number of buffer headers. There is now a default maximum of 10000 system buffers. This limit can be overridden by explicitly defining NBUF in the sn.h file, config.h file, or startup parameter file, or by altering NBLK_FCTR in config.h.
NGRT	Maximum number of guest resource table entries.

GigaRing based systems and Model E based systems have common and unique disk parameters. Table 23 through Table 26 show these parameters.

Table 23. Disk Parameters (Common)

Parameter	Description
LDDMAX	Maximum number of logical disk devices (LDDs). This value also limits the minor number for this type. The maximum minor number is LDDMAX -1.
MDDSLMAX	Maximum number of mirrored disk device (MDD) slices. This value also limits the minor number for this type. The maximum minor number is MDDSLMAX -1.
PDDMAX	Maximum number of physical disk devices (PDDs); minimum is 0.
PDDSLMAX	Maximum number of PDD slices; minimum is 0. This value also limits the minor number for this type. The maximum minor number is PDDSLMAX -1.
RDDSLMAX	Maximum number of RAM disk device (RDD) slices. This value also limits the minor number for this type. The maximum minor number is RDDSLMAX -1.
SDDSLMAX	Maximum number of striped disk device (SDD) slices. This value also limits the minor number for this type. The maximum minor number is SDDSLMAX -1.

Table 24. Disk Parameters (GigaRing Based Systems Only)

Parameter	Description
SSDTMAX	Maximum number of SSD-T90 devices. The value must be an integer in the range 0 through 8. 0 is the default.
SSDTSLMAX	Maximum number of slices that can be allocated to the SSD-T90 device. This value must be an integer in the range 0 through $2^{15} - 1$, depending upon the value set for SSDTMAX. The default is 32. This value also limits the minor number for this type. The maximum minor number is SSDTSLMAX - 1.
XDDMAX	Maximum number of physical devices. The default is 32.
XDDSLMAX	Maximum number of physical slices. The default is 256.

Table 25. Disk and Block Transfer Engine Parameters (Cray SV1ex Systems with SSD-I Only)

Parameter	Description
SSDIMAX	Maximum number of SSD-I devices. The value must be either 0 or 1. 0 is the default.
SSDISLMAX	Maximum number of slices that can be allocated to the SSD-I device. This value must be an integer in the range 0 through $2^{15} - 1$, depending upon the value set for SSDIMAX. The default is 32. This value also limits the minor number for this type. The maximum minor number is SSDISLMAX - 1.
XDDMAX	Maximum number of physical devices. The default is 32.
XDDSLMAX	Maximum number of physical slices. The default is 256.

Table 26. Disk Parameters (Model E Based Systems Only)

Parameter	Description
HDDMAX	Maximum number of HIPPI disk devices (HDDs).
HDDSLMAX	Maximum number of HDD slices. This value also limits the minor number for this type. The maximum minor number is HDDSLMAX - 1.
SSDSSLMAX	Maximum number of SSD driver (SSDD) slices. This value also limits the minor number for this type. The maximum minor number is SSDSSLMAX - 1.

Table 27 shows the valid unit bit and range numbers for IONs.

Table 27. ION Unit Bit and Range Numbers (GigaRing Based Systems Only)

ION type	Bits	Range
FCN	Loop ID 0-7	0-127
HPN	Facility bits 0-8	0-127
IPN	Unit bits on a daisy chain 0-2	0-7
MPN	Device ID 0-7	0-14
	Logical unit number 0-7	
RAID	RAID partition bits 9-15	0-127

5.3.5.1 Example for a GigaRing Based System

The following is an example unicos section for a GigaRing based system:

```
unicos {
    2480    NBUF;                /* System buffers */
    100    NLDCH;              /* Ldcache headers */
    2000    LDCHCORE;          /* Ldcache memory */
    150    LDDMAX;             /* Max. number of LDD devices */
    150    PDDMAX;             /* Max. number of PDD devices */
    256    PDDSLMAX;           /* Max. number of DISK slices */
    32     XDDMAX               /* Max. number of xdisk devices */
    256    XDDSLMAX            /* Max. number of xdisk slices */
    8      SDDSLMAX;           /* Max. number of SDD slices */
}
```

```

8      MDDSLMAX;          /* Max. number of MDD slices */
4      RDDSLMAX;          /* Max. number of RAM slices */
30     TAPE_MAX_CONF_UP;  /* Max. number of tapes configured */
65536  TAPE_MAX_PER_DEV; /* Max. tape buffer size */
}

```

5.3.5.2 Additional Parameters for a Cray SSD-T90 System

The following is an example of additional parameters required for a Cray SSD-T90 system:

```

unicos {
    2      SSDTMAX;
    16     SSDTSLMAX;
}

```

5.3.5.3 Additional Parameters for a Cray SV1ex System with SSD-I

The following is an example of additional parameters required for a Cray SV1ex system with SSD-I:

```

unicos {
    1      SSDIMAX;
    16     SSDISLMAX;
}

```

5.3.5.4 Example for a Model E Based System

The following is an example unicos section for a Model E based system:

```

unicos {
2480   NBUF;              /* System buffers */
100    NLDCH;             /* Ldcache headers */
2000   LDCHCORE;         /* Ldcache memory */
150    LDDMAX;           /* Max. number of LDD devices */
150    PDDMAX;           /* Max. number of PDD devices */
256    PDDSLMAX;         /* Max. number of DISK slices */
8      SDDSLMAX;         /* Max. number of SDD slices */
8      MDDSLMAX;         /* Max. number of MDD slices */
4      RDDSLMAX;         /* Max. number of RAM slices */
4      SSDDSLMAX;        /* Max. number of SSD slices */
30     TAPE_MAX_CONF_UP; /* Max. number of tapes configured */
65536  TAPE_MAX_PER_DEV; /* Max. tape buffer size */
}

```

5.3.6 filesystem Section

The `filesystem` section includes the following:

- Description of the physical devices in the system:
 - `xdisks`, `disks`, `RAM`, `SSD-T` (Cray `SSD-T90`) for GigaRing based systems and `SSD-I` (Cray `SV1ex` auxiliary memory) for Cray `SV1ex` systems `filesystem` section
 - `disks`, `RAM`, and `SSD` for Model E based systems (`SSD-E`)
- Description of the device nodes in the system:
 - `XDD`, `QDD`, `PDD`, `RDD`, `MDD`, `HDD`, and `SDD` for GigaRing based systems
 - `PDD`, `LDD`, `RDD`, `MDD`, `HDD`, and `SDD` for Model E based systems
- Identification of the root, swap, `SDS`, and dump devices
- Description of the Cray `SSD-T90` (GigaRing based systems only)
- Description of the `SSD-I` (Cray `SV1ex` systems only)

Set these parameters by using the following menu selection:

```
Configure System
->Disk configuration
```

For information on striping file systems, see *General UNICOS System Administration*.

The beginning of the `filesystem` section is indicated by the following line in the CSL parameter file:

```
filesystem {
```

The following sections describe each portion of the `filesystem` section of the parameter file. For more detailed information on physical device specification, see *General UNICOS System Administration*.

Set these parameters by using the following menu selection:

```
Configure System
->Disk Configuration
  ->Physical Devices
```

5.3.6.1 Physical Device Definition

The following sections describe the definition of physical storage devices for GigaRing based systems and Model E based systems. Supported disk types include: IPI-2, SCSI, Fibre Channel, HIPPI, Ethernet (10Base-T and 100Base-T), FDDI, ATM DC3, Network disk, and Special (RAM, SSD, SSD-T, SSD-I).

5.3.6.1.1 Slice Definitions for Physical Storage Devices



Warning: Due to the fact that the UNICOS kernel and `econfig` do not know the size of `xdd` devices, checks for a slice spilling over the end of the drive and for gaps after the last slice do not work for `xdd` (nor for `hdd` devices).

Each physical storage device can be segmented into one or more pieces; each piece is a *slice*. For each slice, a *minor device number*, a starting device address, and the slice length must be specified. The starting device address and the slice length can be specified in units of sectors or blocks. A *block* is defined to be 4096 bytes (512 64-bit words). A *sector* is an integer multiple of some number of blocks that varies with different physical storage devices. The ratio of the device's sector size to the size of a block is defined as one `iounit`. For example, a physical disk device may be formatted with a sector size of 16,384 bytes, giving it an `iounit` value of 4.

Note: `iounit` is not valid for SSD-I.

Generally, the start and length of a disk slice is specified in units of sectors. If the unit value is `blocks`, the start and length values must be integer multiples of the `iounit` value. The `iounit` value for a physical disk storage device is determined by the device type designator in the physical disk declaration.

Generally, the start and length of an `xdisk` slice are specified in units of sectors. If `blocks`, the start and length values must be integer multiples of the `iounit` value. The `iounit` value for a physical disk storage device is determined by the `iounit` designator in the physical `xdisk` declaration. In the case of an `xdisk` where no `iounit` is declared, the `iounit` value defaults to 1.

The start and length of a slice of a RAM disk are specified in `blocks`.

The start and length of an `ssd` slice can be specified in units of either `blocks` or `sectors`. If `sectors`, the `iounit` value of the SSD is determined by the

optional `type` designator in the physical SSD declaration. If a `type` is not specified, the `iounit` value defaults to 1.

The start and length of an `ssdt` or `ssdi` slice can be either `blocks` or `sectors`. If `sectors` are used, the `iounit` of the Cray SSD-T90 device is determined by the optional `iounit` designator in the physical Cray SSD-T90 device declaration. If an `iounit` designator is not specified, the `iounit` value defaults to 1.

The following example shows the two forms for slice configuration:

```
slice_type slice_name {
    minor minor_number;
    sector starting_sector_number;
    length length_in_sectors sectors;
}
```

```
slice_type slice_name {
    minor minor_number;
    block starting_block_number;
    length length_in_blocks blocks;
}
```

Device minor numbers must be unique among all slices of a given storage device type. They must be greater than 0 and less than the maximum number of slices specified in the `unicos` section of the parameter file.

Table 28 shows the preferred units and optional start and length units for the various physical storage device types, as well as the parameter in the `unicos` section that determines the maximum value allowable for the minor device number.

Table 28. Start and Length Units for Physical Storage Device Types

Physical storage device type	Preferred units	Optional units	Maximum minor number -1
disk	sectors	blocks	PDDSLMAX
xdisk	sectors	blocks	XDDSLMAX
hippi_disk	sectors	blocks	HDDSLMAX
RAM	blocks	(none)	RDDSLMAX

Physical storage device type	Preferred units	Optional units	Maximum minor number -1
SSD	blocks	sectors	SSDDSLMAX
SSD-T90	blocks	sectors	SSDTSLMAX
SSD-I	blocks	sectors	SSDISLMAX

5.3.6.1.2 Physical Device Definition for GigaRing Based Systems



Warning: Due to the fact that the UNICOS kernel and `econfig` do not know the size of xdd devices, checks for a slice spilling over the end of the drive and for gaps after the last slice do not work for xdd (nor for hdd devices).

The following types of physical storage devices are available for GigaRing based systems:

- Random access memory (RAM)
- Physical storage devices (disk and xdisk)
- Solid-state storage device for a Cray T90 system (Cray SSD-T90)
- Auxiliary memory (SSD-I) for a Cray SV1ex system

Table 29 summarizes disk information for GigaRing based systems.

Table 29. Disk Information (GigaRing Based Systems Only)

Disk type	ION	PCA type	Driver	Node residence	Major number	CSL type	mkspice value
IPI-2	IPN	SPN	qdd	/dev/pdd	dev_qdd	disk	YES
SCSI	MPN	MPN	xdd	/dev/xdd	dev_xdd	xdisk	NO
Fibre	FCN	SPN	xdd	/dev/xdd	dev_xdd	xdisk	NO
HIPPI	HPN	MPN	xdd	/dev/xdd	dev_xdd	xdisk	NO

The **RAM device definition** has the following format (which is identical to that in a Model E based system):

```

RAM ram_name
  { length length_number units ;
    pdd pdd_slice_name
      { minor minor_number
        block starting_block_number
        length length_in_blocks blocks
      }
  }

```

<i>ram_name</i>	Name of the RAM, which must be unique among all devices.
<i>length_number</i>	Size of the RAM, specified in <i>units</i> .
<i>units</i>	One of the following: blocks, words, or Mwords.
<i>pdd_slice_name</i>	Name of the slice.
<i>minor_number</i>	Minor number of the slice, which must be unique across the device type.
<i>starting_block_number</i>	Block number where the slice starts.
<i>length_in_blocks</i>	Length of the slice in blocks.

The **physical storage device definition** has the following format in a GigaRing based system:

```

disk device_name {
  type type;
  iopath {
    ring ring_number;
    node node_number;
    channel channel_number;
  }
  unit disk_unit_number;
  pdd pdd_slice_name {
    minor minor_number;
    sector starting_sector_number;
    length length_in_sectors sectors;
  }
}

```

<i>device_name</i>	Name of the physical storage device, which must be unique among all devices.
<i>type</i>	Type of the physical storage device.

<i>ring_number</i>	Number of the I/O path ring.
<i>node_number</i>	Number of the I/O path node.
<i>channel_number</i>	Number of the I/O path channel. The channel specified is the channel in the peripheral channel adapter (PCA). You must include a leading 0 to specify the channel number in octal form.
<i>disk_unit_number</i>	Number of the disk. For disk devices that can be daisy chained, the unit number specifies the physical unit number of the device. It is recommended that start and length for disk devices be expressed in sectors.
<i>pdd_slice_name</i>	Name of the slice, which must be unique among all slices for all devices.
<i>minor_number</i>	Minor number of the slice, which must be unique across the device type.
<i>starting_sector_number</i>	Starting sector number of the slice.
<i>length_in_sectors</i>	Length of the slice in sectors.

The **SSD-T (ssdt)** definition applies only to GigaRing based systems. It has the following format:

```

ssdt ssdt_name {
    iounit iounit_value;
    tmttype tmttype_number;
    [lunit lunit_number;]
    iopath {
        ring ring_number;
        node node_number;
    }

    [piopaths piopaths_bitmask];
    xdd xdd_name {
        minor minor_number;
        block starting_block_number;
        length length_in_blocks blocks;
    }
    [piopaths piopaths_bitmask];
    xdd xdd_name {
        minor minor_number;
        block starting_block_number;
        length length_in_blocks blocks;
    }
}

```

<i>ssdt_name</i>	Name of the Cray SSD-T90, which must be unique among all devices.
<i>iounit_value</i>	Block size in 512-byte units.
<i>tmttype_number</i>	Target memory type, which determines the memory address and configuration characteristics for a given Cray SSD-T90. This value must be an integer in the range through 255. The default is 1.
<i>lunit_number</i>	(Optional) The logical unit number of the Cray SSD-T90; this is only used if there is more than one Cray SSD-T90 device connected to the system, or if you want to split a single Cray SSD-T90 into more than one logical device. Normally, a system has only one Cray SSD-T90 and <i>lunit</i> defaults to 0.
<i>ring_number</i>	GigaRing ring number for the Cray SSD-T90 device. (This can be a logical ring/node address in the form of a <i>gr_union</i> device; see Section 5.3.2.2, page 43.)

<i>node_number</i>	GigaRing node number for the Cray SSD-T90 device. (This can be a logical ring/node address in the form of a <code>gr_union</code> device; see Section 5.3.2.2, page 43.)
<i>piopaths_bitmask</i>	<p>(Optional) A bit mask representing the number of physical paths (GigaRing channels) used to split a single request across multiple Cray SSD-T90 connections to a Cray T90 mainframe. Each bit in the bit mask is a possible path:</p> <ul style="list-style-type: none">• 03 represents 2 paths, which is standard for 512-Mword Cray SSD-T90 devices• 017 represents 4 paths, which is standard for 1024-Mword Cray SSD-T90 devices <p>Using this parameter can increase the bandwidth of individual large requests but will result in higher system overhead and may decrease overall Cray SSD-T90 throughput. The <code>piopaths</code> parameter is typically used for specific applications (such as NASTRAN) that need the additional bandwidth and cannot use asynchronous requests.</p> <p>By default, individual user requests are scheduled in a round-robin fashion across the connection and no parallel I/O is done. There are no <code>piopath</code> values, which gives the best system throughput for multiple user requests.</p>
<i>xdd_name</i>	Name of the slice, which must be unique among all slices for all devices.
<i>minor_number</i>	Minor number of the slice, which must be unique across the device type.
<i>starting_block_number</i>	Starting block number of the slice.
<i>length_in_blocks</i>	Length of the slice in blocks.

The **SSD-I (`ssdi`) definition** applies only to Cray SV1ex, which are all GigaRing based systems. It has the following format:

```

ssdi ssdi_name {
    tmttype tmttype_number;
    iopath {
        ring ring_number;
        node node_number;
    }
    xdd xdd_name {
        minor minor_number;
        block starting_block_number;
        length length_in_blocks blocks;
    }
    xdd xdd_name {
        minor minor_number;
        block starting_block_number;
        length length_in_blocks blocks;
    }
}

```

<i>ssdi_name</i>	Name of the SSD-I, which must be unique among all devices.
<i>tmttype_number</i>	Target memory type, which determines the memory address and configuration characteristics for SSD-I. This value must be 3 for SSD-I. There is no default.
<i>xdd_name</i>	Name of the slice, which must be unique among all slices for all devices.
<i>minor_number</i>	Minor number of the slice, which must be unique across the device type.
<i>starting_block_number</i>	Starting block number of the slice.
<i>length_in_blocks</i>	Length of the slice in blocks.

The **xdisk definition** applies only to GigaRing based systems. It has the following format:

```
xdisk xdisk_name {
    iounit iounit_value;
    iopath {
        ring ring_number;
        node node_number;
        channel channel_number;
    }
    unit disk_unit_number
    xdd xdd_slice_name {
        minor minor_number;
        sector starting_sector_number;
        length length_in_sectors;
    }
}
```

<i>xdisk_name</i>	Name of the physical storage device, which must be unique among all devices.
<i>iounit_value</i>	Sector size in 4096-byte I/O units.
<i>ring_number</i>	GigaRing ring number of the peripheral channel adapter (PCA).
<i>node_number</i>	GigaRing ring node number of the PCA.
<i>channel_number</i>	Number of the I/O path channel. The channel specified is the channel in the PCA. You must include a leading 0 to specify the channel number in octal form.
<i>disk_unit_number</i>	Device unit number.
<i>xdd_slice_name</i>	Name of the slice, which must be unique among all slices for all devices.
<i>minor_number</i>	Minor number of the slice, which must be unique across the device type.
<i>starting_sector_number</i>	Starting sector number of the slice.
<i>length_in_sectors</i>	Length of the slice in sectors.

5.3.6.1.3 Physical Device Definition for Model E Based Systems

The following types of physical devices are available for Model E based systems:

- Random access memory (RAM)
- Physical storage devices

- Solid-state storage device (SSD)

The **RAM device definition** has the following format (which is identical to that in a GigaRing based system):

```
RAM ram_name
  { length length_number units ;
    pdd pdd_slice_name
      { minor minor_number
        block starting_block_number
          length length_in_blocks blocks
      }
  }
```

<i>ram_name</i>	Name of the RAM, which must be unique among all devices.
<i>length_number</i>	Size of the RAM, specified in <i>units</i> .
<i>units</i>	One of the following: blocks, words, or Mwords.
<i>pdd_slice_name</i>	Name of the slice.
<i>minor_number</i>	Minor number of the slice, which must be unique across the device type.
<i>starting_block_number</i>	Block number where the slice starts.
<i>length_in_blocks</i>	Length of the slice in blocks.

The **physical storage device definition** has the following format for a Model E based system:

```
disk device_name {
  type type;
  iopath {
    cluster cluster_number;
    eiop eiop_number;
    channel channel_number;
  }
  unit disk_unit_number;
  pdd pdd_slice_name {
    minor minor_number;
    sector starting_sector_number;
    length length_in_sectors sector;
  }
}
```

<i>device_name</i>	Name of the physical storage device, which must be unique among all devices.
<i>type</i>	Type of the physical storage device.
<i>cluster_number</i>	Number of the I/O cluster.
	Note: For Cray Model V based systems, <i>cluster</i> specifies the VME IOS, and <i>eiop</i> specifies the controller within the VME IOS. For more information, see the <i>Cray Scalable I/O Messages</i> .
<i>eiop_number</i>	Number of the EIOP I/O processor.
<i>channel_number</i>	Number of the I/O path channel. You must include a leading 0 to specify the channel number in octal form.
<i>disk_unit_number</i>	Number of the disk. For disk devices that can be daisy chained, the unit number specifies the physical unit number of the device. It is recommended that start and length for disk devices be expressed in sectors.
<i>pdd_slice_name</i>	Name of the slice.
<i>minor_number</i>	Minor number of the slice, which must be unique across the device type.
<i>starting_sector_number</i>	Starting sector number of the slice.
<i>length_in_sectors</i>	Length of the slice in sectors.

The **solid-state storage device (SSD) definition** has the following format:

```

SSD ssd_name
  { length length_number units ;
    pdd pdd_slice_name {
      minor minor_number;
      block starting_block_number;
      length length_in_blocks blocks;
    }
  }

```

<i>ssd_name</i>	Name of the SSD, which must be unique among all devices.
<i>length_number</i>	Size of the SSD.

<i>units</i>	One of the following: blocks, Mwords, or sectors.
<i>pdd_slice_name</i>	Name of the slice, which must be unique among all slices for all devices.
<i>minor_number</i>	Minor number of the slice, which must be unique across the device type.
<i>starting_block_number</i>	Starting block number of the slice.
<i>length_in_blocks</i>	Length of the slice in blocks.

5.3.6.2 Device Node Definition

The following device nodes can be defined in the `filesystem` section of the CSL parameter file:

<u>Device type</u>	<u>Description</u>
pdd	Physical device for use with the CSL type disk.
ldd	Logical device.
sdd	Striped device.
mdd	Mirrored device.
qdd	Physical device for GigaRing based systems; used to divide <code>xdisk</code> entries in the <code>filesystem</code> section.
xdd	Physical disk device for GigaRing based systems for use with the CSL type <code>xdisk</code> .

The only limitation is that any slice used in a node definition must have been defined in a physical storage device definition. For more information on mirrored and striped devices, see *General UNICOS System Administration*.

Set the QDD and PDD parameters by using the following menu selection:

```
Configure System
  ->Disk Configuration
    ->Physical Device Slices (/dev/pdd entries)
```

Set the LDD parameters by using the following menu selection:

```
Configure System
  ->Disk Configuration
    ->Logical Devices (/dev/dsk entries)
```

Set the SDD parameters by using the following menu selection:

```
Configure System
->Disk Configuration
   ->Striped Devices (/dev/sdd entries)
```

Set the MDD parameters by using the following menu selection:

```
Configure System
->Disk Configuration
   ->Mirrored Devices (/dev/mdd entries)
```

Set the XDD parameters by using the following menu selection:

```
Configure System
->Disk Configuration
   ->Physical Device Slices on GigaRing Systems
      (/dev/xdd entries)
```

Each node definition has the following syntax:

```
node_type name {
    minor number;
    device slice;
```

The node type and device are one of the device types defined in the previous list. The name is site-configurable. The minor number is required and must be unique across the device type. The slice is a name of a slice (or slices or other device node definition) previously defined in your CSL parameter file.

5.3.6.3 Cray SSD-T90 Description

The Cray SSD-T90 configuration reflects the explicit nature of the GigaRing, as opposed to the implicit VHISP form used for Model E based systems. A Cray SSD-T90 device can have either a physical GigaRing `iopath` or a GigaRing union `iopath`.

Note: The Cray SSD-T90 has an `iopath` designator (without a channel number) in its declaration.

The `lunit` (logical unit) parameter is the device ordinal that will be assigned to this physical Cray SSD-T90. With one Cray SSD-T90, the `lunit` parameter is optional and defaults to 0. If more than one Cray SSD-T90 is designated, `lunit`

is required, must be unique among SSDs, and must be smaller than the maximum number of Cray SSD-T90 devices. The maximum number of Cray SSD-T90 devices is designated by the `SSDTMAX` parameter in the `unicos` section of the parameter file. See Section 5.3.6.1, page 62, for the format.

The `tmttype` parameter is the target memory type associated with the Cray SSD-T90. The target memory type determines memory address and configuration characteristics for a given Cray SSD-T90. Table 30 shows the `tmttype` values.

Table 30. Target Memory Type Values

<code>tmttype</code> value	Description
8	An 8-processor system
9	A 16-processor system

5.3.6.4 Cray SV1ex Auxiliary Memory (SSD-I) Description

The Cray SV1ex system has a faster, extended memory, a portion of which includes an auxiliary memory called SSD-I, and a high speed block transfer engine (BTE). SSD-I is internal to the Cray SV1ex memory modules. It is used essentially as the SSD-T (Cray SSD-T90) and SSD-E (Model E SSD) are used on other Cray PVP systems. Supported uses for SSD-I include fast swap space, file system space, logical device (disk) cache, and secondary data segment (SDS). The BTE provides a CPU-controlled data path for direct memory to memory transfers, for example, between main memory and extended memory or even within main memory.

The maximum number of SSD-I devices is designated by the `SSDIMAX` parameter in the `unicos` section of the parameter file. See Section 5.3.6.1, page 62, for the format.

The `tmttype` parameter is the target memory type associated with SSD-I memory. The target memory type determines memory address and configuration characteristics for the SSD-I. It must be set to a value of 3 for SSD-I.

5.3.6.5 Root, Swap, and Secondary Data Segment (SDS) Devices

The statements for the root, swap, and SDS devices have the following syntax in the `filesystem` section of the CSL parameter file for GigaRing based systems:

```
rootdev is ldd name;  
swapdev is ldd name;  
sdsdev is xdd name;  
dmpdev is xdd name;
```

`sdsdev` must be an SSD, Cray SSD-T90 slice, or SSD-I slice; `dmpdev` must be an XDD definition; the others must be an LDD definition.

These statements have the following syntax for Model E based systems:

```
rootdev is ldd name;  
swapdev is ldd name;  
sdsdev is pdd name;  
dmpdev is ldd name;
```

Set these parameters by using the following menu selection:

```
Configure System  
->Disk Configuration  
->Special System Device Definitions
```

5.3.6.6 Example of the `filesystem` Section Containing a RAM File System

The following example shows the `filesystem` section of a working CSL parameter file containing a RAM file system:

Note: Additional CSL tags are required in the `unicos` section. See Section 5.3.5, page 55.

```
filesystem {  
  RAM ramdev {length 10240 blocks;  
    pdd ram {minor 3; block 0; length 10240 blocks;}  
  }  
}
```

5.3.6.7 Example of the filesystem Section for a GigaRing Based System

The following example shows the filesystem section of a working CSL parameter file for a GigaRing based system:

Note: Additional CSL tags are required in the unicos section. See Section 5.3.5, page 55.

```
filesystem {
  /* Physical device configuration */
  xdisk d04026.3 { iounit 1;
    iopath { ring 04; node 02; channel 06; } unit 03;
    pdd 04026.3_usr_i { minor 231; sector 0; length 444864 sectors; }
    pdd 04026.3_ccn { minor 232; sector 444864; length 222432 sectors; }
  }
  xdisk d03020.0 { iounit 1;
    iopath { ring 03; node 02; channel 0; } unit 0;
    xdd mpn.s400 { minor 17; sector 0; length 102000 sectors; }
    xdd mpn.roote { minor 18; sector 102000; length 250000 sectors; }
    xdd mpn.usre { minor 19; sector 352000; length 250000 sectors; }
  }
  /* HPN Device */
  xdisk d257.200.78.223 { iounit 22;
    iopath { ring 02; node 05; channel 07; }
    unit 40136; ifield 0337;
    xdd hpn.1 { minor 23; block 0; length 250000 blocks; }
  }
  /* Logical device configuration */
  ldd usr_i { minor 86; xdd 04026.3_usr_i; }
  ldd ccn { minor 40; xdd 04026.3_ccn ; }
  ldd root_e { minor 17; xdd mpn.roote ; }
  ldd usr_e { minor 30; xdd mpn.usre ; }
  ldd swap { minor 2; xdd mpn.s400 ; }
  ldd hpn { minor 10; xdd hpn.1 ; }
  rootdev is ldd root_e;
  swapdev is ldd swap;
}
```

5.3.6.8 Example of the filesystem Section for a GigaRing Based System with Disk Devices Configured for Third-party I/O

The following example shows the filesystem section of a working CSL parameter file for a GigaRing based system with disk devices configured for third-party I/O.

Note: Additional CSL tags are required in the `unicos` section. See Section 5.3.5, page 55.

```
filesystem {
  /* Physical device configuration */
  xdisk d04026.3 { iounit 1;
    iopath { ring 04; node 02; channel 06; } unit 03;
    pdd 04026.3_usr_i { minor 231; sector 0; length 444864 sectors; }
    pdd 04026.3_ccn { minor 232; sector 444864; length 222432 sectors; }
  }
  xdisk d03020.0 { iounit 1;
    iopath { ring 03; node 02; channel 0; } unit 0;
    xdd mpn.s400 { minor 17; sector 0; length 102000 sectors; }
    xdd mpn.roote { minor 18; sector 102000; length 250000 sectors; }
    xdd mpn.usre { minor 19; sector 352000; length 250000 sectors; }
    xdd mpn.dump { minor 136; sector 807360; length 100000 sectors; }
  }
  /* HPN Device */
  xdisk d257.200.78.223 { iounit 22;
    iopath { ring 02; node 05; channel 07; }
    unit 40136; ifield 0337;
    xdd hpn.1 { minor 23; block 0; length 250000 blocks; }
  }
  /* Logical device configuration */
  ldd usr_i { minor 86; xdd 04026.3_usr_i ; }
  ldd ccn { minor 40; xdd 04026.3_ccn ; }
  ldd root_e { minor 17; xdd mpn.roote ; }
  ldd usr_e { minor 30; xdd mpn.usre ; }
  ldd swap { minor 2; xdd mpn.s400 ; }
  ldd hpn { minor 10; xdd hpn.1 ; }
  rootdev is ldd root_e;
  swapdev is ldd swap;
  dmpdev is xdd mpn.dump;
}
```

5.3.6.9 Example of the `filesystem` Section for a Cray SSD-T90 Device

The following example shows the `filesystem` section of a working CSL parameter file for a Cray SSD-T90 device.

Note: Additional CSL tags are required in the `unicos` section. See Section 5.3.5, page 55.

```
filesystem {
```



```

ssdt SSDT00 { iounit 1; tmttype 1; lunit 0;
    iopath { ring 0200; node 01; }
    length 512 Mwords;
    xdd ssd_blk0 { minor 2; block 0; length 131072 blocks; }
    xdd ssd_blk1 { minor 3; block 131072; length 131072 blocks; }
}
sdsdev is xdd ssd_blk1;
}

```

5.3.6.10 Example of the filesystem Section for a SSD-I Device

The following example shows the filesystem section of a working CSL parameter file for an SSD-I device.

Note: The length can be 4 Gwords or 8 Gwords. Additional CSL tags are required in the unicost section. See Section 5.3.5, page 55.

```

filesystem {
    /*
    * SSD-I configuration.
    */
    ssdi ssd-svlex {
        tmttype 3;
        length 4 Gwords;
        xdd ssdi_test {
            minor 1;
            block 0;
            length 97152 blocks;
        }
        xdd ssdi_swap {
            minor 2;
            block 97152;
            length 1000000 blocks;
        }
        xdd ssdi_sds {
            minor 3;
            block 1097152;
            length 1000000 blocks;
        }
    }
}

```

5.3.6.11 Example of the filesystem Section for Model E Based Systems

The following example shows the filesystem section of a working CSL parameter file for Model E based systems:

```
filesystem {
  disk d0230.0 {type DD62; iopath{cluster 0; eiop 2; channel 030;} unit 0;
    pdd root.0      {minor 10; sector 0; length 166824 sectors;}
    pdd usr.1       {minor 11; sector 166824; length 166824 sectors;}
    pdd core        {minor 12; sector 333648; length 333648 sectors;}
  }
  disk d0232.0 {type DD62; iopath{cluster 0; eiop 2; channel 032;} unit 0;
    pdd root.1      {minor 15; sector 0; length 166824 sectors;}
    pdd usr.0       {minor 16; sector 166824; length 166824 sectors;}
    pdd scratch     {minor 17; sector 333648; length 333648 sectors;}
  }
  disk d0234.0 {type DD62; iopath{cluster 0; eiop 2; channel 034;} unit 0;
    pdd src         {minor 20; sector 0; length 667296 sectors;}
  }
  disk dl230.0 {type DD62; iopath{cluster 1; eiop 2; channel 030;} unit 0;
    pdd mfs0        {minor 60; sector 0; length 166824 sectors;}
    pdd scr1        {minor 61; sector 166824; length 500472 sectors;}
  }
  disk dl232.0 {type DD62; iopath{cluster 1; eiop 2; channel 032;} unit 0;
    pdd mfs1        {minor 62; sector 0; length 166824 sectors;}
    pdd scr2        {minor 63; sector 166824; length 333648 sectors;}
    pdd dump        {minor 64; sector 500472; length 166824 sectors;}
  }
  disk dl234.0 {type DD62; iopath{cluster 1; eiop 2; channel 034;} unit 0;
    pdd stripe0     {minor 70; sector 0; length 667296 sectors;}
  }
  disk dl236.0 {type DD62; iopath{cluster 1; eiop 2; channel 036;} unit 0;
    pdd stripe1     {minor 71; sector 0; length 667296 sectors;}
  }
  sdd strfs {minor 1; pdd stripe0;
             pdd stripe1;}
  mdd mfs   {minor 1; pdd mfs0;
             pdd mfs1;}
  ldd root0 { minor 10; pdd root.0 ;}
  ldd usr   { minor 11; pdd usr.0 ;}
  ldd src   { minor 12; pdd src ;}
  ldd core  { minor 13; pdd core ;}
  ldd dump  { minor 14; pdd dump ;}
  ldd bkroot { minor 15; pdd root.1 ;}
```

```

ldd bkusr      { minor 16; pdd usr.1      ;}
ldd tmp        { minor 21; pdd tmp0      ;
                pdd tmp1              ;}
ldd usrtmp     { minor 27; pdd usrtmp    ;}
ldd scratch    { minor 35; pdd scratch   ;
                pdd scr1              ;
                pdd scr2              ;}
ldd mir.fs     { minor 30; mdd mfs       ;}
ldd str.fs     { minor 40; sdd strfs     ;}
rootdev is ldd root0;
swapdev is ldd swap;
dmpdev is ldd dump;
}

```

5.3.6.12 Example of the `filesystem` Section Containing an SSD for Model E Based Systems

The following example shows the `filesystem` section of a working CSL parameter file containing an SSD:

Note: Additional CSL tags are required in the `unicos` section. See Section 5.3.5, page 55.

```

filesystem {
  SSD ssddev {length 512 Mwords;
    pdd ssd_0a {minor 2; block 0; length 524288 blocks;}
    pdd ssd_0b {minor 3; block 524288; length 524288 blocks;}
  }
  ldd swap { minor 23; pdd ssd_0a ;}
  swapdev is ldd swap;
  sdsdev is pdd ssd_0b;
}

```

5.3.7 network Section

The `network` section defines network devices and network parameters. You can configure them by using the following menu:

```

Configure System
->UNICOS Kernel Configuration
  ->Communication Channel Configuration

```

The `network` section includes the following information:

- Descriptions of network parameters

- Customized network device prototypes (Model E based systems only)
- Descriptions of each specific network device using standard templates or customized prototypes

The `network` section is specified in the CSL parameter file in the following manner:

```
network {  
    number network_parameter_statement;  
    custom_network_device_specification_statement;  
    physical_network_device_statement;  
}
```

The three statements are repeated as necessary to describe the network device configuration.

The following sections describe the three statement types that compose the `network` section.

5.3.7.1 Network Parameters

You can set the network parameters by using the following menu selections:

```
Configure System  
->UNICOS Kernel Configuration  
->Network Parameters
```

and

```
Configure System  
->Network Configuration  
->TCP/IP Configuration  
->TCP Kernel Parameters Configuration
```

The network parameter statement has the following format:

```
number network_parameter_statement;
```

The *number* is a valid CSL number for the network statement. *network_parameter_statement* argument can have the values described in Table 31 through Table 35, page 86.

Table 31. Network Parameter Values (Common)

Parameters	Description
atmarp_entries	Size of the asynchronous transfer mode (ATM) address resolution protocol (ARP) table.
atmarp_recv	Amount of socket space used for receive for ATM ARP traffic. Must be a power of 2.
atmarp_send	Amount of socket space used for send for ATM ARP traffic. Must be a power of 2.
cnfs_static_clients	Maximum number of active Cray NFS static clients.
cnfs_temp_clients	Maximum number of active Cray NFS temporary clients.
harp_entries	Size of the High Performance Parallel Interface (HIPPI) address resolution protocol (ARP) table.
harp_recv	Amount of socket space used for receive for HIPPI ARP traffic. Must be a power of 2.
harp_send	Amount of socket space used for send for HIPPI ARP traffic. Must be a power of 2.
hidirmode	Sets permissions or file mode for the following directories: /dev/ghippi# (GigaRing based system) and /dev/hippi (Model E based system).
hifilemode	Sets permissions or file mode for the following files: /dev/ghippi#/* (GigaRing based system) and /dev/hippi/* (Model E based system).
nfs_duptimeout	Time interval in seconds during which duplicate requests received by the NFS server will be dropped.
nfs_maxdata	Maximum amount of data that can be transferred in an NFS request (the NFS data buffer size).
nfs_maxdupreqs	Size of the NFS server's duplicate request cache.
nfs_num_rnodes	Size of the NFS client's NFS file-system-dependent node table (rnode table for NFS).

Parameters	Description
nfs_portmon	Enables (nonzero) or disables (zero) whether NFS clients are required to use privileged ports (<code>ports < IPPORT_RESERVED</code>) in order to get NFS services.
nfs_printinter	Time in seconds between server not responding message to a down server.
nfs_static_clients	Number of static client handles reserved for NFS client activity.
nfs_temp_clients	Number of dynamically allocated client handles that can be used for NFS client activity when all of the static client handles are in use.
nfs_wcredmax	Maximum number of credential structures.
nfs3_async_max	Maximum amount of data (in 4096-byte blocks) that will be written per file asynchronously. After this value is exceeded, all writes will be synchronous. The default is 2000 (2000 * 4096 = 8,192,000 bytes).
nfs3_async_time	The amount of time (in seconds) that data will be held in the NFS async write cache on the client.
tcp_numbspace	Number of clicks of memory set aside for TCP/IP managed memory buffers (Mbufs).

Table 32. Network Parameter Values (GigaRing Based Systems)

Parameters	Description
ithreshold	(Optional) Specifies the multiplier (an integer in the range 1 through 4) used to configure the input response table size. The default is 1. This parameter is ignored for <code>gr</code> interfaces.
othreshold	(Optional) Specifies the multiplier (an integer in the range 1 through 4) used to configure the output response table size. The default is 1. This parameter is ignored for <code>gr</code> interfaces.

Parameters	Description
maxinputs	Maximum number of asynchronous read I/O requests that can be issued to the I/O node. This must be an integer value in the range 1 through 256. The default is 16.
maxoutputs	Maximum number of write I/O requests that can be issued to the ION. This must be an integer value in the range 1 through 256. The default is 16.
maxusers	Maximum number of applications that can share the HIPPI device. This parameter applies only to HIPPI devices; it must be set to 1 for other interfaces. For HIPPI devices, the value must be an integer in the range 1 through 8. The default is 2.

Table 33. Network Parameter Values (Common to Model E and Model V Based Systems)

Parameters	Description
fdmaxdevs	Maximum number of fiber distributed data interface (FDDI) FCA-1 channel adapters.
himaxdevs	Maximum number of channels that you can configure for HIPPI.
himaxpaths	Maximum numbers of logical paths per channel that you can configure for HIPPI.
npmaxdevs	Maximum number of channels that can be configured for low-speed channel communication devices.

Table 34. Network Parameter Values (Model V Based Systems)

Parameters	Description
atmmxdevs	Maximum number of ATM devices.
enmxdevs	Maximum number of Ethernet devices.

Table 35. Network Parameter Values (Model E Based Systems)

Parameters	Description
fddirmode	Sets permissions of file mode for /dev/fddi* directories.
fdfilemode	Sets permission of file mode for /dev/fddi*/* files.
fdmaxpaths	Maximum number of logical paths per physical FDDI channel.
npdirmode	Sets the permissions for subdirectories of /dev/comm.
npfilemode	Sets the permissions for /dev/comm/CHAN/lpXXfiles.
npito	Input time-out.
npmaxpaths	Total number of logical paths for low-speed channel communication devices.
npmaxppd	Maximum number of paths per device for low-speed channel communications devices.
npoto	Output time-out.
nprthresh	Reads threshold for low-speed channels (number of reads that can be queued to the IOS at one time).
nperto	Default read time-out (per path).
npwthresh	Writes threshold for low-speed channels (number of writes that can be queued to the IOS at one time).

5.3.7.2 Customized Network Device Prototypes for Model E Based Systems

The customized network device prototype lets you define characteristics for low-speed devices unique to your site. The network device prototype statement has the following syntax:

```

np_spec interface_type {
    device type device_type;
    channel mode channel_mode;
    driver type driver_type;
    driver mode driver_mode_number;
    device function function_number;
    direction timeout timeout_number;
}
    
```

<i>interface_type</i>	Type of the interface as defined in Table 36, page 89.
<i>device_type</i>	Type of the device as defined in Section 5.3.7.3.
<i>channel_mode</i>	One of the following values: 12MB 12MB_LP 6MB
<i>driver_type</i>	One of the following values: FY LCP MP NSC_MP PB RAW
<i>driver_mode_number</i>	Number of the driver mode.
<i>function_number</i>	Number of the function.
<i>direction</i>	Either input or output (only used with HIPPI devices).
<i>timeout_number</i>	Time-out value.

5.3.7.3 `device type` for Model E and Model V Based Systems

The `device type` statement, which defines the network device type in IOS Model E and Model V systems, has the following format:

```
device type device;
```

The following are valid device types:

<u>Device type</u>	<u>Description</u>
A130	NSC A130 devices
CNT	CNT devices
DIAG	Diagnostic devices
DX4130	FDDI connection
DX8130	FDDI connection
EN643	IP router
FEI_3	FEI-3 devices
FEI_3FY	FEI-3FY devices
FEI_4	FEI-4 devices
FEI_CN	FEI-CN devices
FEI_DS	FEI-1 data-streaming IBM connection
FEI_UC	FEI-1 user driver channel
FEI_VA	VAX FEI-1, port A
FEI_VB	VAX FEI-1, port B
FEI_VM	Cray devices
N130	NSC N130 devices
VME	VME-bus devices

Cray provides you with configuration templates for a number of standard interfaces, as shown in Table 36.

Table 36. Standard Interface Configuration Templates

Interface type	Device type	Channel mode	Driver type	Driver mode	Device func	Input time-out	Output time-out	Read time-out
S_DIAG6	DIAG	6MB	RAW	0	0	100	100	100
S_DIAG12	DIAG	12MB	RAW	0	0	100	100	100
S_DIAG12L	DIAG	12MB_LP	RAW	0	0	100	100	100
S_FEICN	FEI_CN	6MB	LCP	0	0	100	100	600
S_FEIDS	FEI_DS	6MB	LCP	0	0	100	100	600
S_FEIUC	FEI_UC	6MB	LCP	0	0	100	100	600
S_FEIVA	FEI_VA	6MB	LCP	1	0	100	100	600
S_FEIVB	FEI_VB	6MB	LCP	2	0	100	100	600
S_FEIVM	FEI_VM	6MB	LCP	0	0	100	100	600
S_FEI3	FEI_3	6MB	MP	0	0	100	100	600
S_FEI312	FEI_3	12MB	MP	0	0	100	100	600
S_VAXBI	VAX_BI	12MB	MP	1	0	100	100	600
S_N130X	N130	12MB	PB	0	0	100	100	600
S_EN643X	EN643	12MB	PB	0	0	100	100	600
S_DX4130X	DX4130	12MB	PB	0	0	100	100	600
S_DX8130X	DX8130	12MB	PB	0	0	100	100	600
S_FEI3FY	FEI_3FY	6MB	FY	010	0	100	100	600
S_FEI4	FEI_4	6MB	FY	010	0	100	100	600
S_A130X	A130	12MB	NSC_MP	0	0	100	100	600

If your site has a network device that does not match these definitions, you must customize a device definition by using the customized network device prototypes, as described in this section. The menu that you would use to do so is as follows:

```
Communication Channel Configuration
->Custom network device specification
```

5.3.7.4 Device Types

The following tables describe the types of devices.

Note: For more information about ATM, see the *Asynchronous Transfer Mode (ATM) Administrator's Guide*.

Table 37. Network Device Types (GigaRing Based Systems)

Device type	Description
gfddi	GigaRing FDDI device.
gatm	GigaRing asynchronous transfer mode (ATM) device.
gether	GigaRing Ethernet device.
ghippi	GigaRing HIPPI device.
gr	GigaRing TCP/IP device (provides direct mainframe-to-mainframe communication).

Table 38. Network Device Type (Model E Based Systems)

Device type	Description
fddev	IOS-E FDDI device.
hi	High-speed HIPPI device.
npdev	Low-speed channel; if this device is specified, then the <i>number</i> argument is the ordinal of the network device.

Table 39. Network Device Type (Model V Based Systems)

Device type	Description
atmdev	ATM device.
endev	Ethernet device.
fddev	FDDI device.

5.3.7.5 Device Formats

The following sections describe device formats.

5.3.7.5.1 Device Formats for GigaRing Based Systems

The following formats apply to GigaRing based systems:

```

gatm number {
    iopath {
        iopath_information
    }
    maxusers number;
    maxinputs number;
    maxoutputs number;
    [ithreshold number;]
    [othreshold number;]
}
    
```

```

gether number {
    iopath {
        iopath_information
    }
    maxusers number;
    maxinputs number;
    maxoutputs number;
    [ithreshold number;]
    [othreshold number;]
}
    
```

```

gfddi number {
    iopath {
        iopath_information
    }
    maxusers number;
    maxinputs number;
    maxoutputs number;
    [ithreshold number;]
    [othreshold number;]
}
    
```

```
ghippi number {  
    iopath {  
        iopath_information  
    }  
    maxusers number;  
    maxinputs number;  
    maxoutputs number;  
    [ithreshold number];  
    [othreshold number];  
}
```

```
gr number {  
    iopath {  
        ring ring_number;  
        node node_number;  
    }  
}
```

In the `gr` device format, *ring_number* identifies the GigaRing that TCP/IP will use to communicate between mainframes, and *node_number* identifies this mainframe's assigned node ID.

5.3.7.5.2 Device Format for Model E Based Systems

The following formats apply to Model E based systems only:

```
fddev number {  
    padcnt number;  
    treq number;  
    maxwrt number;  
    maxrd number;  
    iopath {  
        iopath_information  
    }  
}
```

Note: There are two formats for `hidev`: one for input and one for output.

```

hidev number {
    iopath {
        iopath_information
    }
    logical path number {
        flags number;
        I_field number;
        ULP_id number;
    }
    flags number;
    input;
    device type type;
}

```

```

hidev number {
    iopath {
        iopath_information
    }
    logical path number {
        flags number;
        I_field number;
        ULP_id number;
    }
    flags number;
    output;
    device type type;
}

```

```

npdev number {
    iopath {
        iopath_information
    }
    np_spec name;
}

```

5.3.7.5.3 Device Formats for Model V Based Systems

The following formats apply to Model V based systems only:

```

atmdev number {
    iopath {
        iopath_information
    }
}

```

```

endev number {
    iopath {
        iopath_information
    }
}

```

```

fddev number {
    iopath {
        iopath_information
    }
}

```

5.3.7.6 logical path for Model E Based Systems

The `logical path` statement has the following format:

```

logical path number {
    identifier number
}

```

If you have a high-speed HIPPI network device, you could use the *identifier* argument, which has the following values:

<u>Identifier</u>	<u>Description</u>
UL_id	Upper-level protocol identifier; it is a value between 0 and 255 that identifies the protocol or user of a HIPPI packet. The value is the first byte of each HIPPI packet and is used with <code>input</code> to direct incoming packets to the appropriate process (TCP/IP uses a value of 004).
I_field	A data value that is passed over the HIPPI channel when a connection is made. In <code>HXCF_DISC</code> mode, the <code>I-field</code> value precedes each packet. The

flags	<p>I-field value selects the output port of an NSC P_8 or PS_32 switch.</p> <p>Flags related to the specific HIPPI logical device. The only flag available at this time is 02000 (HXCF_IND). When this flag is set, the driver interprets the first word of each user output buffer as the I-field value. The driver puts the incoming I-field value in the first word of the user input buffer.</p>
-------	--

5.3.7.7 *direction* Argument for Model E Based Systems

The *direction* argument has the following syntax:

```
direction ;
```

The *direction* argument (used only with HIPPI devices) can be either input or output.

5.3.7.8 *device type* Statement for Model E Based Systems

The *device type* statement has the following syntax:

```
device type device ;
```

device defines the device connected to the HIPPI channel.

5.3.7.9 *np_spec* Statement for Model E Based Systems

The *np_spec* statement has the following syntax:

```
np_spec name
```

name specifies the customized network device specification describing the network device connected to the low-speed channel.

5.3.7.10 network Section Example Common to All Systems

The following example configures basic network parameters that are needed for all systems:

```
network {
    8 nfs_static_clients;
    8 nfs_temp_clients;
    8 cnfs_static_clients;
    8 cnfs_temp_clients;
32768 nfs_maxdata;
    256 nfs_num_rnodes;
1200 nfs_maxdupreqs;
    3 nfs_duptimeout;
    0 nfs_printinter;
8000 tcp_nmbospace;
0700 hidirmode;
0600 hifilemode;
}
```

5.3.7.11 network Section Example for GigaRing Based Systems

The following is an example of a network section for a GigaRing based system:

```
network {
    gether 0 {
        iopath { ring 01; node 02; channel 05; }
        maxusers 1;
        maxinputs 64;
        maxoutputs 64;
    }
    gfddi 0 {
        iopath { ring 01; node 02; channel 04; }
        maxusers 1;
        maxinputs 64;
        maxoutputs 64;
    }
    gatm 0 {
        iopath { ring 01; node 03; channel 01; }
        maxusers 1;
        maxinputs 64;
        maxoutputs 64;
    }
    ghippi 0 {
```

```

        iopath { ring 01; node 04; channel 02; }
        maxusers 4;
        maxinputs 64;
        maxoutputs 64;
    }

    gr 0 {
        iopath { ring 01; node 05; }
    }
}

```

5.3.7.12 network Section Example for Model E Based Systems

The following example configures FDDI, HIPPI, and NP devices and only those custom network device specifications (np_spec) that are needed:

```

network {
    4 nmaxdevs;
    32 nmaxpaths;
    16 nmaxppd;
    100 npito;
    100 npoto;
    600 nprto;
    10 nprthresh;
    10 npwthresh;
    4 himaxdevs;
    8 himaxpaths;
    1 fdmaxdevs;
    16 fdmaxpaths;
    0700 fddirmode;
    0600 fdfilemode;
    0700 npdirmode;
    0600 npfilemode;
    np_spec FEI3FY {
        device type FEI_3FY;
        channel mode 6MB;
        driver type FY;
        driver mode 8;
        device function 0;
        input timeout 100;
        output timeout 100;
        read timeout 600;
    }
}

```

```
npdev 0 {
    iopath { cluster 0; eiop 0; channel 030; }
    np_spec FEI3FY;
}
hidev 0 {
    iopath { cluster 0; eiop 3; channel 030; }
    logical path 0 { flags 00; I_field 00; ULP_id 00; }
    flags 00;
    input;
    device type PS_32;
}
hidev 1 {
    iopath { cluster 0; eiop 3; channel 032; }
    logical path 0 { flags 00; I_field 00; ULP_id 00; }
    flags 00;
    output;
    device type PS_32;
}
fddev 0 {
    treq 8;
    padcnt 3;
    maxwrt 10;
    maxrd 10;
    iopath { cluster 0; eiop 0; channel 034; }
}
}
```

5.3.7.13 network Section Example for Model V Based Systems

The following example configures Ethernet, FDDI, and ATM devices for Model V based systems.

```
network {
    2 himaxdevs;
    4 himaxpaths;
    1 fdmaxdevs;
    1 enmaxdevs;
    0 npmaxdevs;
    2 atmmaxdevs;
    131072 atmarp_recv;
    65536 atmarp_send;
    1024 atmarp_entries;
```

```
    endev 0 {  
        iopath { cluster 1; eiop 0; channel 020;  
    }  
  
    atmdev 0 {  
        iopath { cluster 0; eiop 0; channel 020;  
    }  
  
    atmdev 1 {  
        iopath { cluster 3; eiop 0; channel 020;  
    }  
  
    fddev 0 {  
        iopath { cluster 2; eiop 0; channel 020;  
    }  
}
```

5.3.8 revision Section

The `revision` section marks the CSL parameter file with a site-defined name for identification purposes, particularly for programs and other Cray products. The `revision` string is set automatically when you use the ICMS.

The `revision` section is specified in the CSL parameter file by the following statement:

```
revision text_string
```

The *text_string* should be a string that is significant for your site and allows you to identify the file.

Runtime Configuration Scripts and Files [6]

This chapter describes scripts and files that are important in the configuration of various aspects of a UNICOS system. These scripts and files are executed or read during system initialization or when changing run levels. For a discussion of system initialization and run-level configuration, see *General UNICOS System Administration*.

Each of the following sections discusses one script or file. The scripts and files appear alphabetically by file name (not path name).

The following scripts and files are discussed:

- `/usr/lib/cron/at.deny` and `/usr/lib/cron/at.allow`
- `/etc/bcheckrc`
- `/etc/brc` and `/etc/coredd`
- `/usr/lib/cron/cron.deny` and `/usr/lib/cron/cron.allow`
- `/etc/cshrc`
- `/etc/config/daemons`
- `/etc/fstab`
- `/etc/gettydefs`
- `/etc/ghippi#.arp` (GigaRing based systems only)
- `/etc/gr#.arp` (GigaRing based systems only)
- `/etc/group`
- `/etc/hycf.local_network` (Model E based systems only)
- `/etc/inittab`
- `/etc/config/interfaces`
- `/etc/issue`
- `/etc/config/ldchlist`
- `/etc/motd`
- `/etc/netstart`

- `/etc/config/netvar.conf`
- `/etc/passwd`
- `/etc/profile`
- `/etc/rc`
- `/etc/config/rcoptions`
- `/etc/shutdown`
- `/etc/umountem`

6.1 `at.deny` and `at.allow` Files

You can alter the `/usr/lib/cron/at.deny` and `/usr/lib/cron/at.allow` files. These files define users who are permitted or excluded from using the `at(1)` command.

Users are permitted to use `at` if their login appears in the file `/usr/lib/cron/at.allow`. If that file does not exist, the file `/usr/lib/cron/at.deny` is checked to determine if the user should be denied access to `at`. If neither file exists, only `root` is allowed to submit a job. An empty `at.allow` file means no user is allowed to use `at`; an empty `at.deny` file means no user is denied the use of `at`.

UNICOS is released with an empty `at.deny` file. The `allow/deny` files consist of one user name per line. Change one file or the other to allow or deny user access to `at`.

6.2 `bcheckrc` Script

The `/etc/bcheckrc(8)` script performs the commands necessary (such as setting the system time or checking file system consistency) before the `rc` script (see `brc(8)`) is to be executed and the file systems are to be mounted. See the `bcheckrc(8)` man page for more information.

Note: This script is not intended to be modified directly. To modify the execution of this script, change the options in `/etc/config/rcoptions`.

6.3 `brc` and `coredd` Scripts

The `brc(8)` script is used for processing system dumps. It copies system dumps to a separate file system by executing the `coredd(8)` shell script.

Note: These scripts are not intended to be modified directly. To modify the execution of this script, change the options in `/etc/config/rcoptions`.

See the `brc(8)` and `coredd(8)` man pages for more information.

6.4 `cron.deny` and `cron.allow` Files

You can alter the `/usr/lib/cron/cron.deny` and `/usr/lib/cron/cron.allow` files during configuration. These files define users who are permitted or excluded from using the `crontab(1)` command.

Users are permitted to use `cron` if their login appears in the file `/usr/lib/cron/cron.allow`. If that file does not exist, the file `/usr/lib/cron/cron.deny` is checked to determine if the user should be denied access to `cron`. If neither file exists, only `root` is allowed to submit a job. An empty `cron.allow` file means no user is allowed to use `cron`; an empty `cron.deny` file means no user is denied the use of `cron`.

UNICOS is released with an empty `cron.deny` file. The files consist of one user name per line. Change one file or the other to allow or deny user access to `cron`.

6.5 `cshrc` File

The `/etc/cshrc` file is the equivalent of `/etc/profile` (see Section 6.21, page 112) for `csh` (the C shell). The release version of `/etc/cshrc` sets the user's file creation mode, prints the `/etc/motd` file on the user's screen, and tells the user if mail and news exists.

If the `/etc/cshrc` file exists and your login shell is `/bin/csh`, the file is executed by the shell upon login before your session begins. Then, if your login directory contains a file named `.cshrc`, it is executed by the shell before the session begins. See the `cshrc(5)` and `csh(1)` man pages for more information.

6.6 daemons File

The `/etc/config/daemons` file is used by the `sdaemon(8)` command for controlling the starting and stopping of system daemons. (The `sdaemon(8)` command is used by the system startup procedures supplied with UNICOS to start daemon processes necessary for system operation.) To access this file through the UNICOS installation and configuration menu system (ICMS), use the following menu selection:

```
Configure system
->System daemons configuration
  ->System daemons table
```

An entry in the `/etc/config/daemons` file has the following format:

<i>daemongroup tag groupstart kill pathname arguments</i>

<i>daemongroup</i>	The group to which this daemon belongs.
<i>tag</i>	A convenient tag to use when referring to the daemon. You can use the name of the daemon's executable file or any other name you choose.
<i>groupstart</i>	Indicates whether this daemon should be started when its group is started. A value of <code>YES</code> indicates this daemon will be started as part of its group. A value of <code>ASK</code> indicates that the operator will be asked at startup time if the daemon should be started. Any other value (<code>NO</code> is suggested) indicates that the daemon will not be started as part of its group. A daemon with a <i>groupstart</i> of <code>NO</code> may be started individually.
<i>kill</i>	A value indicating how the daemon process may be terminated.
<i>pathname</i>	The path name of the daemon executable.
<i>arguments</i>	Command-line arguments to be used when starting the daemon.

A full explanation of the format of the `/etc/config/daemons` file and its use by the `sdaemon(8)` command may be found on the `sdaemon(8)` man page.

The following *daemongroup* values are used by the system startup procedure supplied with UNICOS to start groups of daemons in the following order:

SYS1	Non-networking daemon processes started before the networks are present
TCP	Daemon processes related to TCP/IP
NFS	Daemon processes related to NFS
SYS2	Non-networking daemon processes started after the networks are present

You may choose to define additional groups to suit your local configuration.

6.7 `fstab` File

The `/etc/fstab` file (see `fstab(5)`) contains a list of file systems that the `bcheckrc(8)` script checks by invoking the `mfscck(8)` command. UNICOS programs read `/etc/fstab`, but do not write information to it; you must create and maintain the information in this file.

Note: The `mfscck(8)` command was formerly called `gencat(8)`.

Use the following menu selection to change this file:

```
Configure System
  ->File System (fstab) Configuration
```

However, if you want to change the file manually, this section contains the information you need.

The `/etc/fstab` file is an ASCII file. The fields are separated by white space (tabs or spaces); each group is separated from the next by a newline character. Each entry in `fstab` has the following format:

<i>filesystem directory type options frequency passnumber</i>

The first and last fields are used by `mfscck` (see the `mfscck(8)` man page for more information). See also the routines `getfsent`, `getfsspec`, `getfstype`, and `getfsfile` described on the `getfsent(3)` man page for information on reading records from `/etc/fstab`.

The `/etc/fstab` file contains entries to make mounting file systems easier. When the `mount(8)` command is invoked with only a directory argument or an incomplete argument list, it searches `fstab` for the missing arguments.

The *type* field must have one of the following values. Each value specifies a file system type.

<u>Type</u>	<u>File system</u>
NC1FS	UNICOS file system
NFS	Network file system
PROC	/proc file system

See the `fstab(5)` man page for further information on `fstab`. See also the `mount(8)` man page.

6.8 `gettydefs` File

The `/etc/gettydefs` file contains information used by the `getty(8)` command to set up the speed and terminal settings for a line. `getty` is used for the IOS-E terminals.

The information in `gettydefs` also specifies the appearance of the login prompt.

The `/etc/gettydefs` file is an ASCII file. The fields are separated by pound signs (#); each group is separated from the next by a newline character. Each entry in `gettydefs` has the following format:

label#initial_flags#final_flags#login_prompt#next_label

The fields contain the following information:

<i>label</i>	String against which <code>getty</code> tries to match its second argument.
<i>initial_flags</i>	Initial <code>ioctl(2)</code> settings to which the terminal is to be set if a terminal type is not specified with <code>getty</code> .
<i>final_flags</i>	Flags that are set just before <code>getty</code> executes <code>login(1)</code> .
<i>login_prompt</i>	The login prompt. All characters and white space in the field are included in the prompt.
<i>next_label</i>	The desired speed, specified by a break character. If this field does not contain the speed, <code>getty</code> searches for the entry with <i>next_label</i> as its <i>label</i> field and sets up the terminal for those settings.

See the `gettydefs(5)` and `getty(8)` man pages for detailed information on the `/etc/gettydefs` file. Also see the `login(1)`, `vi(1)`, and `ioctl(2)` man pages for more information on `gettydefs`.

6.9 `ghippi#.arp` File

The `/etc/ghippi#.arp` file lets you configure the hardware addresses of known hosts on local networks that are directly connected to a GigaRing based system via HIPPI. *X* is the ordinal of the HIPPI interface.

For more information, see *UNICOS Networking Facilities Administrator's Guide*.

6.10 `gr#.arp` File

The `/etc/gr#.arp` file lets you configure the hardware addresses of known hosts on local networks that are connected using host-to-host GigaRing interfaces. *X* is the ordinal of the GigaRing interface.

For more information, see *UNICOS Networking Facilities Administrator's Guide*.

6.11 `group` File

The `/etc/group` file (see `group(5)`) is provided to translate group names to group IDs and group IDs to names. It is not used for user validation.

The `/etc/group` file is an ASCII file that contains the following information for each group:

- Group name
- Encrypted password
- Numerical group ID (GID)
- Comma-separated list of user names allowed in the group

The `password` field is not used in UNICOS and is set to an asterisk (*). See the `group(5)` man page for more information on the `/etc/group` file and `udb(5)` and *General UNICOS System Administration*, for more information on the user database (UDB) and the `group` file.

6.12 hycf .local_network Files

The `/etc/hycf.local_network` files let you configure the hardware addresses of known hosts on local networks that are directly connected to an Model E based system. There is one file for each local network connected to the mainframe.

6.13 inittab File

The `/etc/inittab` file controls the actions of the `init(8)` process. It contains information on processes and scripts that are to be executed at the various run levels. See *General UNICOS System Administration*, for a description of UNICOS run levels and run-level configuration.

The `/etc/inittab` file is composed of position-dependent entries that have the following format:

id: *rstate*: *action*: *process*

Each entry is delimited by a newline character; however, a backslash (\) preceding a newline character indicates a continuation of the entry. Up to 512 characters for each entry are permitted. Comments may be inserted in the *process* field, using a # sign at the beginning of the field.

The entry fields are as follows:

<u>Field</u>	<u>Description</u>												
<i>id</i>	1 to 4 characters that uniquely identify an entry.												
<i>rstate</i>	Run level in which this entry is to be processed.												
<i>action</i>	Keywords that specify how to treat the process in the <i>process</i> field. The following actions are recognized by <code>init</code> :												
	<table border="0"> <tr> <td><code>boot</code></td> <td><code>bootwait</code></td> <td><code>generic</code></td> </tr> <tr> <td><code>initdefault</code></td> <td><code>off</code></td> <td><code>once</code></td> </tr> <tr> <td><code>ondemand</code></td> <td><code>respawn</code></td> <td><code>sysinit</code></td> </tr> <tr> <td><code>timezone</code></td> <td><code>wait</code></td> <td></td> </tr> </table>	<code>boot</code>	<code>bootwait</code>	<code>generic</code>	<code>initdefault</code>	<code>off</code>	<code>once</code>	<code>ondemand</code>	<code>respawn</code>	<code>sysinit</code>	<code>timezone</code>	<code>wait</code>	
<code>boot</code>	<code>bootwait</code>	<code>generic</code>											
<code>initdefault</code>	<code>off</code>	<code>once</code>											
<code>ondemand</code>	<code>respawn</code>	<code>sysinit</code>											
<code>timezone</code>	<code>wait</code>												
<i>process</i>	An <code>sh(1)</code> command to be executed.												

Be sure to change the time zone entry in the `/etc/inittab` file to reflect your site's time zone. The ICMS will copy your time zone entry if you are performing an upgrade installation, but initial installations must make the change to

`/etc/inittab` manually. For instructions on changing the time zone, see *General UNICOS System Administration*.

For more information about the structure of the `/etc/inittab` file, see `inittab(5)` in the *UNICOS File Formats and Special Files Reference Manual*.

6.14 interfaces File

The `/etc/config/interfaces` file contains the list of TCP/IP interfaces that may be configured using the `initif(8)` command. The `initif` command is used by the system startup procedures supplied with UNICOS. A full explanation of the format of the `/etc/config/interfaces` file may be found on the `initif(8)` man page and in the *UNICOS Networking Facilities Administrator's Guide*. To access this file through the ICMS, use the following menu selection:

```
Configure system
  ->Network configuration
    ->General network configuration
      ->Network interface configuration
```

6.15 issue File

The `/etc/issue` file (see `issue(5)`) is displayed before the login prompt. It is used to echo any important messages users might need to know prior to logging in, such as the following:

```
The system is dedicated.
Please do not log on.
```

See *General UNICOS System Administration*, for more information on the `/etc/issue` file.

6.16 ldchlist File

The `/etc/config/ldchlist` file is used by the `rc` system start-up script supplied with UNICOS to initialize the logical device cache to be used during normal system operation. To access this file through the ICMS, use the following menu selection:

```

Configure system
  ->Disk configuration
    ->logical device cache
    
```

For more information, see the description of `rc` on the `brc(8)` man page and in *General UNICOS System Administration*.

6.17 motd File

The `/etc/motd` file (see `motd(5)`) is the UNICOS message-of-the-day file. It generally contains the UNICOS release level and any important messages that must be conveyed to each user as they log in. It is released containing a warning to unauthorized users. The following is the released version of the `/etc/motd` file:

```
cat /etc/motd
```

```

This is a private computer facility.  Access for any reason must be
specifically authorized by the owner.  Unless you are so authorized,
your continued access and any other use may expose you to criminal
and/or civil proceedings.
    
```

6.18 netstart Script

The `netstart` script is executed by the `rc` script to initialize UNICOS networking software. It starts each type of UNICOS networking software (for example, TCP/IP or OSI) by calling the appropriate startup script or command for that software type.

The list of scripts executed by `/etc/netstart` is as follows:

<u>Command</u>	<u>Description</u>
<code>/etc/netstart.pre</code>	The script to be run before <code>netstart</code> to accommodate local requirements.
<code>/etc/nwmstart</code>	Command used to initialize the underlying network media
<code>/etc/tcpstart</code>	Command to initialize TCP/IP software.
<code>/etc/nfsstart</code>	Command to initialize the network file system (NFS) software.

<code>/etc/ypstart</code>	Command to initialize network information service (NIS) software. (NIS was formerly known as yellow pages.)
<code>/etc/netstart.pst</code>	The script to be run after <code>netstart</code> to accommodate local requirements.

See the `netstart(8)` man page for additional information about the `netstart` script.

6.19 netvar.conf File

The `/etc/config/netvar.conf` file contains a list of command-line arguments for the `netvar(8)` command. These arguments are used by the `/etc/tcpstart` script supplied with UNICOS to initialize TCP/IP kernel variables at system startup. For a complete list of arguments to the `netvar(8)` command and further information about its operation, see the `netvar` man page and the *UNICOS Networking Facilities Administrator's Guide*.

To access this file through the ICMS, use the following menu selection:

```
Configure system
  ->Network configuration
    ->TCP/IP network configuration
      ->kernel parameters
```

6.20 passwd File

The `/etc/passwd` file contains one entry per user. Each entry consists of the login name, encrypted password (this field is set to an asterisk (*) or to `* ,pw->pw_age` if password aging is active), user ID (UID), group ID (GID), a comment field, initial working directory, and the program the user uses as a shell.

The `/etc/passwd` file exists in UNICOS for compatibility with predecessor systems and plays no role in the validation and management of system users. That function is performed by the user database (UDB). For information on the relationship of `/etc/passwd` to the UDB, see *General UNICOS System Administration*, and the `udb(5)`, `udbgen(8)`, and `udbsee(1)` man pages. For the format of `/etc/passwd` and more information on its content, see the `passwd(5)` man page.

6.21 profile File

If the `/etc/profile` file exists, it is executed by the standard shell upon login before your session begins. Then, if your login directory contains a file named `.profile`, it is executed by the shell before the session begins. The file `.profile` is useful for setting exported environment variables.

6.22 rc Script and rcoptions File

The `rc` (see `brc(8)`) script is executed by the `init(8)` command using an entry in the `/etc/inittab` file (see `inittab(5)`) when changing the run level from single-user to multiuser mode.

To access this file through the ICMS, use the following menu selection:

```
Configure system
  ->Startup (/etc/rc) configuration
```

Note: The `rc` script is not intended to be edited directly. To modify the execution of this script, change the options in `/etc/config/rcoptions`. To perform tasks not controlled by `rcoptions`, create the following files:

- `/etc/rc.pre`
- `/etc/rc.mid`
- `/etc/rc.post`

These files are executed before, during, and after the `rc` script.

The UNICOS `rc` script is intended to provide a high degree of system startup configuration without modification of the script. This includes mounting file systems; activating accounting logging, error logging, and system activity logging; and starting system daemons. To do this, `rc` references a number of subsidiary configuration files and calls various configurable system utilities.

For a complete discussion of the capabilities and configuration of the `rc` script supplied with UNICOS, see the `brc(8)` man page and *General UNICOS System Administration*.

The `/etc/config/rcoptions` file contains a list of environment variable definitions that control the startup configuration of the system. These definitions are read in by the `rc` (see `brc(8)`) script supplied with UNICOS. For a list of environment variables and explanation of their effect on system startup, see the description of the `/etc/rc` script in *General UNICOS System Administration*.

6.23 shutdown Script

The `shutdown(8)` script is a part of the UNICOS operation procedure. It primarily terminates all currently running processes in an orderly and cautious manner. The procedure is as follows:

1. All users logged on the system are notified by a broadcast message to log off the system. You may display your own message at this time. Otherwise, the standard file save message is displayed.
2. All user processes are killed and daemons are shut down.
3. `shutdown` unmounts all file systems.

After the script completes, the dynamic blocks of all file systems are updated before the system is to be stopped (see `ldsync(8)` and `sync(1)`). This must be done before rebooting the system to ensure file system integrity.

After `shutdown` completes, the system is in single-user mode and is essentially equivalent to the state of the system after a reboot procedure. The release version of `shutdown` also shuts down the NQS daemon so that jobs can be recovered when the system reboots.

See *General UNICOS System Administration* for more information on the `shutdown` script.

6.24 umountem Script

The `umountem(8)` script unmounts all mounted file systems except the current root file system. It is called by the `shutdown(8)` shell procedure to ensure that all file systems are truly unmounted before going to single-user mode. This script may also be called manually.

Reserved Keywords in CSL [A]

The following keywords (listed from left to right, top to bottom) have special meaning in the configuration specification language (CSL):

12MB	12MB_LP	6MB
A130	A400	ALL
CNT	Cray	D90
DA301	DA302	DA60
DA62	DD10	DD11
DD19	DD29	DD3
DD301	DD302	DD304
DD308	DD309	DD314
DD318	DD39	DD4
DD40	DD41	DD42
DD49	DD50	DD501
DD5I	DD5S	DD60
DD61	DD62	DD6S
DDAS2	DDES DI	DDIMEM
DDL DAS	DD_U	DIAG
DX4130	DX8130	EN643
FAULT_RESPONSE	FEI_3	FEI_3FY
FEI_4	FEI_CN	FEI_DS
FEI_UC	FEI_VA	FEI_VB
FEI_VM	FY	GATEWAYS
GUESTMAX	Gword	Gwords
HD16	HD32	HD64
HDDMAX	HDDSLMAX	I/O
IEEE	IMP_3	IMP_4

IMP_5	IMP_6	IMP_7
I_field	LCP	LDCHCORE
LDDMAX	LLC	MDDSLMAX
MP	Mword	Mwords
N130	N400	NBUF
NGRT	NLDCH	NPARTITION
NPBUF	NPLCHCTL	NPOOL
NSC_MP	NTRANSACT	NTTYDEV
NYPEDEV	OLNET	PB
PBUFSIZE	PDDMAX	PDDSLMAX
PLCHCORE	PS_32	P_8
RAM	RAW	RD-1
RDDSLMAX	REGT	SBUFSIZE
SDDSLMAX	SMT	SSD
SSDDSLMAX	SSDTMAX	SSDTSLMAX
SSDIMAX	SSDISLMAX	SSD_E
SSD_F	T3D	T90_config
TAPE_MAX_CONF_UP	TAPE_MAX_DEV	TAPE_MAX_PER_DEV
TCP	ULP_id	VAXBI
VME	XDDMAX	XDDSLMAX
YMP	access	alternate
atmarp_entries	atmarp_recv	atmarp_send
atmdev	atmmaxdevs	bank
bias	block	blocks
boot	bootable	both
c100d100	c100d200	c200d200
channel	cluster	clusters
cnfs_static_clients	cnfs_temp_clients	configure
cpu	cpus	cylinder

cylinders	device	disk
dmpdev	down	driver
dumpinfo	eiop	endev
enmaxdevs	enmaxpaths	fddev
fddirmode	fdfilemode	fdmaxdevs
fdmaxpaths	fiber_vhisp	filesystem
flags	floating	function
gateway	gatm	gether
gfddi	ghippi	gigaring
gr	gr_route	gr_union
group	half	halfs
halves	harp_entries	harp_recv
harp_send	hdd	hidev
hidirmode	hifilemode	himaxdevs
himaxpaths	hippi_disk	hisp
ifield	input	io_connect
io_cpus	io_cpus_mask	iomodule
iomodules	iopath	ios_e
iounit	is	ithreshold
kernel	ldd	length
logical	logical_machine	lower
lowspeed	lunit	mainframe
maintenance	maxinputs	maxoutputs
maxrd	maxusers	maxwrt
mdd	memory	minor
miop	mixed	mode
mpp	mtu	msps
muxiop	network	nfs3_async_max
nfs3_async_time	nfs_duptimeout	nfs_maxdata

nfs_maxdupregs	nfs_num_rnodes	nfs_printinter
nfs_static_clients	nfs_temp_clients	nfs_wcredmax
no	nocache	no-direction
node	np_spec	npdev
npdirmode	npfilemode	npito
npmaxdevs	npmaxpaths	npmaxppd
npoto	nprthresh	nprto
npwthresh	othreshold	output
owner	ows	padcnt
path	pdd	phase_III
physical	piopaths	point
primary	profile	pseudo
qdd	range	read
register	reserve	revision
rft	ring	rootdev
route	sdd	sdsdev
section	sector	sectors
select	serial	spare
spares	ssdi	ssdt
start	stream	subsection
swapdev	system	tcp_nmbSPACE
timeout	tntype	to
track	tracks	treq
type	unicos	unit
upper	user	user_channel
vhisP	with	word
words	write	xdd
xdisk	xmemory	

A

Access request logging, 25
ACNICE, 15
Address resolution protocol (ARP), 83
Adjust-on-exit value, 20
ARP
 See Address resolution protocol
Async write cache, 28
Asynchronous I/O headers, 12
Asynchronous I/O headers reserved for the system, 14
Asynchronous I/O structures allowed per process, 12
Asynchronous transfer mode, 86
Asynchronous transfer mode (ATM) address resolution protocol (ARP), 83
asynchronously written data, 28
at.allow file, 102
at.deny file, 102
ATM
 See Asynchronous transfer mode
ATM device
 Model V based, 90
atmarp_entries parameter, 83
atmarp_recv parameter, 83
atmarp_recvspace parameter, 83
atmarp_send parameter, 83
atmmaxdevs parameter, 86
Audit criteria logging, 26
Automatic mixed buffer, 15
Auxiliary memory (SSD-I)
 description, 75

B

Back door I/O access, 3
Back door I/O configuration rules, 4
Bit matrix multiply functional unit, 35
Bits per memory chip, 8

Boot process in CSL, 39
boot statement, 46
brc shell script, 103
Buffer cache factor, 13
Buffer headers, 57
Buffer size for NFS data, 83
Buffers, 84
Bus Based Gateway, 35

C

C default, 33
Cache, 83
Cache blocks based on memory size, 13
Cache size of duplicate requests (NFS_MAXDUPREQS), 28
Callout table entries, 13
CDLIMIT, 12
Channel assignments, 54
Channel declarations, 54
 GigaRing based systems, 52
 IOS-E based systems, 54
Channel information, 45
channel statement
 GigaRing based systems, 52
 Model E based systems, 54
chdir(1) request logging, 26
Checkpoint and restart buffers, 13
Chip size, 8
CHIPSZ, 8
Client handles in CNFS, 29
Client handles in NFS, 29
clist, 13
Clock period, 9
cluster, 55
cluster statement, 45
Clusters, 44
cnfs_static_clients, 30
CNFS_STATIC_CLIENTS, 29

- cnfs_static_clients parameter, 83
- cnfs_temp_clients, 30
- CNFS_TEMP_CLIENTS, 29
- cnfs_temp_clients parameter, 83
- Comments in CSL, 38**
- COMPART_ACTIVE_DEFAULT, 22
- COMPART_VALID_DEFAULT, 22
- CONFIG_BEG, 35
- CONFIG_BMM, 35
- CONFIG_CPP, 33
- CONFIG_CPSAVE, 33
- CONFIG_CRAYLIBS, 36
- CONFIG_CRL, 35
- CONFIG_CVT, 35
- CONFIG_DFS, 35
- CONFIG_DIAGDIR, 32
- CONFIG_DM parameter, 35
- CONFIG_ELS, 35
- CONFIG_FQUOTAS, 35
- CONFIG_GCC, 33
- CONFIG_GEN_SEGDIR, 33
- CONFIG_GENBIN, 33
- CONFIG_GENCMDS, 34
- CONFIG_GENPROD_RULES, 34
- config.h configuration file, 1, 11
 - CNFS_STATIC_CLIENTS, 29
 - CNFS_TEMP_CLIENTS, 29
 - MSGMNI, 19
 - MSGSSZ, 20
 - NFS kernel parameters, 28**
 - NFS_PRINTINTER, 29
 - NFS_STATIC_CLIENTS, 29
 - NFS3_ASYNC_TIME, 28
 - NPBUF, 14
 - NPLCHCTL, 14
 - security log (SLG_STATE), 27**
 - SLG_STATE, 27
 - static client handles in CNFS**
 - (CNFS_STATIC_CLIENTS), 29
 - static client handles in NFS**
 - (NFS_STATIC_CLIENTS), 29
 - temporary client handles**
 - (CNFS_TEMP_CLIENTS), 29
 - temporary client handles**
 - (NFS_TEMP_CLIENTS), 29
 - U_MAXPACK, 16
 - CONFIG_HPI3, 35
 - CONFIG_HSX, 35
 - CONFIG_ID, 32
 - GigaRing based systems, 32**
 - Model E based systems, 33**
 - CONFIG_IOS-F
 - GigaRing based systems, 32**
 - Model E based systems, 33**
 - CONFIG_IOSA_SN, 33
 - CONFIG_IOSB_SN, 33
 - CONFIG_IPI3, 35
 - CONFIG_KERBEROS, 35
 - CONFIG_LD_STD_DIR parameter, 36
 - config.mh configuration file for UNICOS kernel
 - subsystems, 1, 31**
 - CONFIG_CRAYLIBS, 36
 - CONFIG_MIXED, 35
 - CONFIG_MIXEDTARGET, 35
 - CONFIG_MK
 - GigaRing based systems, 32**
 - Model E based systems, 33**
 - CONFIG_MPP, 35
 - CONFIG_MPP_CPP, 34
 - CONFIG_NETMON parameter, 35
 - CONFIG_NETTOLS, 35
 - CONFIG_NFS parameter, 35
 - CONFIG_NFS3, 36
 - CONFIG_NFSKRB, 36
 - CONFIG_NIOS, 33
 - CONFIG_NODE, 32
 - CONFIG_OWS, 36
 - CONFIG_PACKAGE, 34
 - CONFIG_PATH, 34
 - CONFIG_RLS_MAJOR, 34
 - CONFIG_RLS_MINOR, 34
 - CONFIG_RLS_REVISION, 34
 - CONFIG_RPC, 36

-
- CONFIG_SN, 32
 - CONFIG_SUPPORT_DIR, 34
 - CONFIG_SYS, 32
 - CONFIG_TAPE, 36
 - CONFIG_TARGET, 32
 - CONFIG_TCP, 36
 - CONFIG_TMPDIR, 32
 - CONFIG_TRUSTED, 36
 - CONFIG_VERSION, 32
 - CONFIG_X11, 36
 - CONFIG_XLIBS, 34
 - CONFIG_XLIBTARGET, 34
 - Configuration**
 - shell scripts, 110, 112
 - SSD-I, 79
 - /CONFIGURATION, 39
 - Configuration logging**, 25
 - Configuration specification language**
 - boot-time equivalents for NFS kernel parameters, 29
 - Configuration specification language (CSL)**
 - parameter file, 1, 37, 57, 83–84
 - asynchronous transfer mode (ATM) address resolution protocol (ARP), 83
 - ATM devices, 86
 - atmarp_entries parameter, 83
 - atmarp_recv parameter, 83
 - atmarp_send parameter, 83
 - atmmaxdevs parameter, 86
 - boot process and, 39
 - buffer headers, 57
 - buffer size, 83
 - cnfs_static_clients parameter, 83
 - cnfs_temp_clients parameter, 83
 - Cray NFS clients, 83
 - customized network device prototypes for Model E based systems, 87
 - device formats
 - GigaRing based systems, 91
 - IOS-E based systems, 92
 - disk, 63
 - disk parameters
 - common, 57
 - GigaRing based systems, 58
 - Model E based systems, 59
 - SSD-I systems, 58
 - dumpinfo section, 41
 - duplicate request cache, 83
 - endev device, 90
 - error message syntax, 39
 - Ethernet devices, 86
 - fault tolerance, 56
 - FAULT_RESPONSE parameter, 56
 - fddev device, 90
 - fdmaxdevs parameter, 85
 - fiber distributed data interface (FDDI) (FCA-1)
 - channel adapter, 85
 - gatm device, 90
 - gether device, 90
 - gfddi device, 90
 - ghippi device, 90
 - GigaRing device types, 90
 - gigaring section, 42
 - gr_route subsection, 43
 - gr_union subsection, 44
 - gr device, 90
 - guest resource table entries, 57
 - guest systems, 57
 - GUESTMAX parameter, 57
 - harp_entries parameter, 83
 - harp_recv parameter, 83
 - harp_send parameter, 83
 - HDDMAX parameter, 59
 - HDDSLMAX parameter, 59
 - hi device, 90
 - hidirmode parameter, 83
 - hifilemode parameter, 83
 - High Performance Parallel Interface (HIPPI), 83, 85
 - himaxdevs parameter, 85
 - himaxpaths parameter, 85
 - HIPPI directory mode, 83
 - HIPPI disk device (HDD) slices, 59
 - HIPPI disk devices (HDDs), 59

- hippi_disk device type, 63
- interface configuration templates, 89
- ios_section
 - clusters, 45
 - IOP declarations, 45
- ios_e section, 44
 - HISP declarations, 46
 - IOP boot declarations, 46
- IPI-2 disk type, 64
- ithreshold parameter, 84
- ldcache blocks, 56
- ldcache headers, 56
- LDCHCORE parameter, 56
- LDDMAX parameter, 57
- logical disk devices (LDDs), 57
- low-speed channel communication devices, 85
- mainframe configuration
 - menu, 48
- mainframe section, 47
 - channel declarations (GigaRing based systems), 52
 - channel declarations (Model E based systems), 54
 - common parameters, 48
 - CPUs, 48
 - Cray SV1 series GigaRing based System parameters, 50
 - mainframe cluster registers, 49
 - memory size, 49
- maximum limits parameters, 57
- maxinputs parameter, 85
- maxoutputs parameter, 85
- maxusers parameter, 85
- MDDSLMAX parameter, 57
- mirrored disk device (MDD) slices, 57
- NBUF parameter, 57
- network device types
 - GigaRing based, 90
 - Model E based, 90
 - Model V based, 90
- network parameters
 - common, 83
 - GigaRing based, 85
 - Model V based, 86
- network section, 81
 - menu, 81
- NFS requests, 83
- nfs_duptimeout parameter, 83
- nfs_maxdata parameter, 83
- nfs_maxdupreqs parameter, 83
- nfs_num_rnodes parameter, 83
- nfs_printinter parameter, 84
- nfs_static_clients parameter, 84
- nfs_temp_clients parameter, 84
- nfs_wcredmax parameter, 84
- nfs3_async_max parameter, 84
- nfs3_async_time parameter, 84
- NGRT parameter, 57
- NLDCH parameter, 56
- NPBUF parameter, 56
- npdev device, 90
- npfilemode parameter, 86
- npmaxdevs parameter, 85
- npmaxpaths parameter, 86
- npmaxppd parameter, 86
- npoto parameter, 86
- nprthresh parameter, 86
- nprto parameter, 86
- npwthresh parameter, 86
- othreshold parameter, 84
- output time-out, 86
- parameter file sections, 40
 - common, 47, 55, 61, 81
 - GigaRing based systems, 42
 - Model E and Model V based systems, 44
 - Model E based systems, 41
- PDDMAX parameter, 57
- PDDSLMAX parameter, 57
- physical devices, 58
- physical disk devices (PDDs), 57
- physical I/O buffers, 56
- physical slices, 58
- RAM device type, 63
- RAM disk device (RDD) slices, 57

- RDDSLMAX parameter, 57
- revision section, 99
- rnode table for NFS, 83
- SDDSLMAX parameter, 57
- semicolon in, 39
- Socket space, 83
- SSD device type, 64
- SSD-I device type, 64
- SSD-I devices, 58
- SSD-I slice allocation, 58
- SSD-T90 device type, 64
- SSD-T90 devices, 58
- SSD-T90 slice allocation, 58
- SSDD slices, 59
- SSDDSLMAX parameter, 59
- SSDIMAX parameter, 58
- SSDISLMAX parameter, 58
- SSDTMAX parameter, 58
- SSDTSLMAX parameter, 58
- statement terminator, 39
- striped disk device (SDD) slices, 57
- syntax
 - comments, 38
 - constants, 38
 - identifier, 37
 - operators, 38
 - separators, 38
- table size parameters, 56
- TAPE_MAX_CONF_UP parameter, 56
- TAPE_MAX_DEV parameter, 56
- TAPE_MAX_PER_DEV parameter, 56
- TCP/IP managed memory buffers, 84
- tcp_numbspace parameter, 84
- termination of statements, 39
- token classes, 37
- unicos section, 55
 - menu, 55
- usage, 39
- verification of configurations, 39
- XDDMAX parameter, 58
- XDDSLMAX parameter, 58
- xdisk, 63
- Console for MLS administration
 - (SYSTEM_ADMIN_CONSOLE), 27
- CONSOLE_MSG, 22
- Constants in CSL, 38
- Contraction value, 16
- coredd shell script, 103
- coremap initial dynamic kernel memory units, 17
- CPP, 34
- CPP variable, 33
- CPUs, 9, 48
- Cray Data Migration Facility (DMF), 35
- Cray floating-point and IEEE), 35
- Cray J90se support, 35
- Cray MPP system, 35
- Cray NFS clients, 83
- Cray SV1 series support, 35
- Cray Visualization Toolkit, 35
- Cray/REELlibrarian, 35
- Cray/REELlibrarian activity logging, 26
- cray-t3e assembler, 33
- CRL
 - See Cray/REELlibrarian
- cron.allow file, 103
- cron.deny file, 103
- Cross-targeted libraries, 34
- cshrc file, 103
- CSL
 - See Configuration specification language (CSL) parameter file
- CSL reserved keywords, 115
- Customized network device prototypes for Model E based systems, 87
- CVT
 - See Cray Visualization Toolkit
- Cycles per second, 9
- D**
 - daemons file, 104
 - Data Migration Facility
 - See Cray Data Migration Facility
 - data read or written by way of NFS, 28
 - Data written asynchronously, 28

DCE

See Distributed Computing Environment

Declarations in CSL

HISP, 46

IOP, 45

DECLASSIFY_DISK, 22

DECLASSIFY_PATTERN, 22

DELAY_MULT, 22

Destination routing, 44

/dev/comm, 86

DEV_ENFORCE_ON, 22

/dev/fddi*, 86

Device labeling, 22

device node definition, 73

device type definition, 73

device type statement, 88

DFS

See Distributed Computing Environment

Directory name look-up cache (DNLC)

number of entries, 13

path name, 13

Disable a parameter, 11

DISABLE_ACCT, 22

DISABLE_TIME, 22

Discretionary access control (DAC) logging, 26

Discretionary access violations, 25

disk, 62–63

Distributed Computing Environment (DCE)

distributed file system (DFS) (CONFIG_DFS), 35

Distribution Center,

DMA driver, 44

DMF

See Cray Data Migration Facility

DMODE, 12

Dump requirements, 41

dumpinfo section, 41

Duplicate request cache, 83

Duplicate request cache size, 28

Duplicate time-out, 28

Dynamic kernel memory parameters, 16

E

econfig(8) command and configuration verification, 39

eiop statement, 45

Enable a parameter, 11

enmaxdevs, 86

Environment variables

definitions in rcoptions file, 112

Error exit maximum, 19

Error message syntax for CSL, 39

Error messages and CSL, 40

/etc/config/param, 37

Ethernet device, 90

Ethernet devices, 86

Examples

filesystem section, 76

Cray SSD-T90, 78

GigaRing based, 77, 80

SSD for Model E based, 81

SSD-I, 79

third-party I/O, 77

ios_e section, 46

mainframe section

GigaRing based systems, 54

Model E based systems, 55

unicos

GigaRing based systems, 59

Model E based systems, 60

SSD-I (additional parameters), 60

SSD-T90 (additional parameters), 60

exec operations, 13

exec operations with maximum arguments, 16

executables (CONFIG_GENCMDS), 34

Expansion units, 16

EXTDCORE, 12

Extended core file naming, 12

Extended memory

auxiliary memory (SSD-I), 75

size, 49

xmemory, 49

F

Failed login attempt logging, 27
 Failed logins, 22
 Fair share scheduler, number of users, 15
 FAULT_RESPONSE parameter, 56
 FCN bit and range numbers, 59
 fddev device, 90
 fddev device format, 92
 FDDI channel, 86
 FDDI device, 90
 fddirmode, 86
 fdfilemode, 86
 fdmaxdevs parameter, 85
 fdmaxpaths, 86
 Fiber distributed data interface (FDDI) (FCA-1)
 channel adapter, 85
 File mode permission, 86
 File quotas, 35
 File retrieval mode, 12
 File size in blocks, 12
 File system log records, 15
 File systems open concurrently, 14
 File systems that can be mounted, 14
 file transfer logging, 25
 File-lock regions (NFLOCKS), 13
 Files that can be open at one time system-wide
 (NFILE), 13
 filesystem section, 61
 Flush buffered data to target I/O devices, 12
 FLUSHONPANIC, 12
 FORCED_SOCKET, 22
 Front door I/O access, 3
 FSETID_RESTRICT parameter, 22
 FSLG_BUFSIZE, 15
 fstab file, 105

G

gatm device, 90
 gatm device format, 91
 General configuration, 11
 Generation parameters, 31
 Generation software (CONFIG_GENBIN), 33

Generation software directory paths, 34
 gether device, 90
 gether device format, 91
 getty command, 106
 gettydefs file, 106
 gfddi device format, 91
 ghippi device, 90
 ghippi device format, 91
 ghippi#.arp file, 107
 GigaRing
 configuration rules for Cray SV1ex, 2
 GigaRing configuration parameters (gigaring
 section), 42
 GigaRing device types, 90
 gigaring section, 42
 GigaRing support
 GigaRing based systems, 32
 Model E based systems, 33
 gr device, 90
 gr device format, 92
 gr_route subsection, 43
 gr_union subsection
 SSD-I, 44
 SSD-T90, 44
 group file, 107
 grX.arp file, 107
 Guest resource table entries, 57
 Guest systems, 57
 GUESTMAX parameter, 57

H

Handles, 84
 harp_entries parameter, 83
 harp_send parameter, 83
 Hash entries (NHBUF_FCTR), 14
 HDD
 See High performance parallel interface disk
 devices
 HDDMAX parameter, 59
 HDDSLMAX parameter, 59
 Hertz, 9
 hi device, 90

hidev device format, 92
 hidirmode parameter, 83
 hifilemode parameter, 83
 High Performance Parallel Interface (HIPPI), 83, 85
 High performance parallel interface (HIPPI) disk devices, 59
 High-speed (HISP), 45
 High-speed channel (HISP) declarations in CSL, 46
 High-speed HIPPI device, 90
 himaxdevs parameter, 85
 himaxpaths parameter, 85
 HIPPI directory mode, 83
 hippy_disk device type, 63
 HISP
 See High-speed (HISP) channel information
 HPN bit and range numbers, 59
 HSX device driver, 35
 hycf.* file, 108
 HZ, 9

I

I/O clusters, 44
 I/O error logging, 27
 I/O module 0 channels, 53
 I/O modules, 51
 I/O node unit bit and range numbers, 59
 I/O processor
 boot declarations, 46
 I/O processor (IOP)
 declarations in CSL, 45
 I/O processors, 45
 ICMS
 See Installation and configuration menu system (ICMS)
 Identifiers in CSL, 37
 IEEE and Cray floating-point, 35
 In-core inode table, 14
 In-core inodes, 13
 In-core inodes for NC1FS (file system migration parameter), 16

In-core quota entries, 14
 Informational messages and CSL, 39
 initif command
 use of interfaces file, 109
 inittab file, 108
 Install tool
 See Installation and configuration menu system (ICMS)
 Installation, 1
 Installation and configuration menu system (ICMS), 1
 Integrity code in *secure.c*, 22
 Interface configuration templates, 89
 interfaces file, 109
 Internal routing, 43
 internal SSD (SSD-I)
 See SSD-I
 Internet Protocol (IP) layer activity logging, 26
 Introduction, 1
 IOC
 See I/O clusters
 ION
 See I/O node
 IOP
 See I/O processors
 iopath, 66, 69
 IOS-E hardware topology definition, 44
 IOS-E machines, 33
 ios_e section, 44
 IOS-E serial numbers, 33
 iounit, 66, 69
 IPC kernel parameters, 19
 IPI-2 disk type, 64
 IPI-3/HIPPI packet driver capability, 35
 IPI-3/IPI capability, 35
 IPN bit and range numbers, 59
 issue file, 109
 ithreshold parameter, 84

J

job end logging, 25
 Job start logging, 25

K

K_OPEN_MAX, 15
 Kerberos support, 35
 Kernel generation parameters, 31
 Kernel memory parameters, 16
 Kernel parameters, 1
 common, 12
 Kernel subsystem parameters, 31
 Kernel subsystems
 See config.mh configuration file
 Keyword identifiers in CSL, 38
 KM_CHM_PADSZ parameter, 16
 KM_EXPAND_UNITS, 16
 KM_NO_THRASH parameter, 16
 KM_UNITS parameter, 17
 KM_WPU parameter, 17
 KM_WPU_SHIFTC parameter, 17

L

ldcache blocks, 12
 Ldcache headers, 56
 LDCHCORE, 12, 42
 LDCHCORE parameter, 56
 ldchlist file, 110
 LDD
 Logical disk devices, 57
 ldd logical device, 73
 LDDMAX parameter, 57
 Libraries, cross-targeted, 34
 Limits parameters, 57
 link call logging, 25
 Link violation logging, 25
 LINK_MAX, 12
 Links to a file, 12
 Log records for the file system, 15
 LOGDELAY, 22
 Logging parameters for MLS, 25
 Logical device cache (ldcache)
 back door I/O access, 4
 Logical device cache headers, 14
 Logical devices in CSL, 38
 Logical disk devices (LDDs), 57

logical path statement, 94
 Login account disabling, 22
 login attempts, 22
 Login attempts, 22
 Login failure logging, 27
 Low-speed channel, 90
 Low-speed channel communication devices, 85

M

Machine-specific characteristics, 1
 Mainframe cluster registers, 48
 Mainframe memory size, 49
 extended memory, 49
 mainframe section, 47
 Mainframe serial number, 8
 Mainframe subtype, 9
 Mainframe type, 9
 Major release number, 34
 Managed memory buffers, 28, 84
 Mandatory access request logging, 26
 MAX_UNLINKED_BYTES parameter, 15
 MAXASYN, 12
 maxinputs parameter, 85
 MAXLOGS, 22
 maxoutputs parameter, 85
 MAXPIPE, 12
 MAXRAH, 12
 MAXSLEVEL, 22
 maxusers parameter, 85
 MAXUSRERR, 19
 MAXUSRORE parameter, 19
 MAXUSRPRE parameter, 19
 MDD
 See Mirrored disk device
 mdd mirrored device, 73
 MDDSLMAX parameter, 57
 MEMORY, 8
 Memory banks, 8
 Memory buffers, 84
 Memory buffers for TCP/IP, 28
 Memory padding, 16
 memory range, 42

- Memory size, 49
- Menu system
 - See Installation and configuration menu system (ICMS)
- Message header number, 20
- Message interval, 29
- Message queue maximum bytes, 19
- Message queue maximum number, 19
- Message segment number, 19
- Message segment size, 20
- Message size, 19
- Messages and CSL, 39
- MFSUBTYP, 9
- MFTYPE, 9
- Minor release number, 34
- MINSLEVEL, 22
- Mirrored disk device (MDD) slices, 57
- Mixed buffer, 15
- Mixed mode (Cray floating-point and IEEE), 35
- mkdir call logging, 25
- mkdir violation logging, 26
- MLS, 11
- MLS administration console, 27
- MLS kernel parameters, 22
- MLS_INTEGRITY, 22
- MLS_OBJ_RANGES, 23
- motd file, 110
- MPN bit and range numbers, 59
- MSGMAX, 19
- MSGMNB, 19
- MSGMNI, 19
- msgque table size, 19
- MSGSEG, 19
- MSGTQL, 20
- Multilevel security (MLS), 11
- Multistreaming Processors (MSPs), 50
- muxiop statement, 45

- N**
- Name of the system, 32
- NASYN, 12
- NBANKS, 8
- NBLK_FCTR, 13
- NBUF parameter, 57
- NBUF_FCTR, 13
- NC_NAMLEN, 13
- NC_SIZE, 13
- NC1FS in-core inodes, 16
- NC1INODE, 13, 16
- NC1MINRAW, 15
- NCALL, 13
- NCLIST, 13
- NCPU, 9
- NCRBUF, 13
- netvar.conf file, 111
- Network access violation logging, 26
- Network device prototypes for Model E based systems, 87
- Network device types
 - Model E based, 90
- Network File System, 35
- Network File System with Kerberos authentication, 36
- Network monitor, 35
- Network parameters
 - common, 83
 - IOS-E, 86
 - menu, 82
 - network section, 81
- Network testing tools, 35
- NEXECS, 13, 16
- NFILE, 13
- NFLOCKS, 13
- NFS
 - See Network File System
 - NFS activity logging, 25
 - NFS kernel parameters
 - See config.h configuration file
 - NFS kernel parameters and their CSL boot-time equivalents, 29
 - NFS requests, 83
 - NFS services
 - NFS clients, 23, 29, 84
 - NFS version 3 Protocol, 36

- nfs_duptimeout, 29
- nfs_duptimeout parameter, 83
- NFS_DUPTIMEOUT parameter, 28
- nfs_maxdata, 29
- NFS_MAXDATA, 28
- nfs_maxdata parameter, 83
- nfs_maxdupreqs, 29
- NFS_MAXDUPREQS, 28
- nfs_maxdupreqs parameter, 83
- NFS-mounting of a remote file system in read-write mode, 23
- nfs_num_rnodes parameter, 83
- nfs_num_rnodes&, 29
- nfs_portmon, 84
- NFS_PORTMON, 29
- nfs_printinter, 30
- NFS_PRINTINTER, 29
- nfs_printinter parameter, 84
- NFS_REMOTE_RW_OK parameter, 23
- NFS_SECURE_EXPORT_OK, 23
- NFS_SECURE_PORTMON, 23
- nfs_static_clients, 30
- NFS_STATIC_CLIENTS, 29
- nfs_static_clients parameter, 84
- nfs_temp_clients, 30
- NFS_TEMP_CLIENTS, 29
- nfs_temp_clients parameter, 84
- nfs_wcredmax parameter, 84
- nfs3_async_max, 29
- NFS3_ASYNC_MAX, 28–29
- nfs3_async_max parameter, 84
- nfs3_async_time, 29
- NFS3_ASYNC_TIME, 28
- nfs3_async_time parameter, 84
- NGRT parameter, 57
- NHBUF_FCTR, 14
- Nice priority values, 12
- Nice value, 15
- NINODE, 14
- NLDCH, 14
- NLDCH parameter, 56
- NLDMAP, 14
- NMNT, 14
- NMOUNT parameter, 14
- Node name of the system, 32
- Node number, 29
- non-inode security system call logging, 26
- np_spec statement, 87
- NPBUF, 14
- NPBUF parameter, 56
- npdev device, 90
- npdev device format, 93
- npdirmode, 86
- npfilemode parameter, 86
- npito, 86
- NPLCHCTL, 14
- npmaxdevs, 85
- npmaxpaths parameter, 86
- npmaxppd parameter, 86
- npoto parameter, 86
- NPROC, 14
- nprthresh parameter, 86
- nprtto parameter, 86
- NPTY, 14
- npwthresh parameter, 86
- NQS activity logging, 25
- NQUOTA, 14
- NSESS, 14
- NSIDEBUF, 14
- NSU_ASYN, 14
- NTEXT, 15
- Number of CPUs, 9
- NUSERS, 15
- O**
- Object identifiers in CSL, 38
- Object ranges, 23
- Offline file retrieval mode, 12
- ONFIG_CAM_CPP_LOC, 33
- Online diagnostics location, 32
- Online tape capability, 36
- Online tape parameters, 56
- Open files per process, 15
- Operand range error maximum, 19

Operating mode information, 45
 Operating system identification, 32
 GigaRing based systems, 32
 Model E based systems, 33
 Operator action logging, 26
 Operator workstation, 36
 Operators in CSL, 38
 othreshold parameter, 84
 Output time-out, 86
 OVERWRITE_COUNT, 23
 OWS
 See Operator workstation
 OWS path names, 45

P

Packaging, 34
 Panic messages and CSL, 40
 Parameter file, 37
 Parameter file sections
 See Configuration specification language (CSL)
 parameter file
 Partition cache blocks, 15
 Partition cache headers, 14
 PASS_MAXSIZE, 23
 PASS_MINSIZE, 23
 passwd file, 111
 Password size maximum, 23
 Password size minimum, 23
 Passwords enabled, 23
 Path name access tracking, 27
 Path selection, 43
 PDD
 See Physical disk devices
 pdd physical device, 73
 PDDMAX parameter, 57
 PDDSLMAX parameter, 57
 Permissions, 86
 Physical devices, 58
 Physical devices in CSL, 38
 Physical disk devices (PDDs), 57
 Physical I/O buffers, 56
 Physical I/O error logging, 27

Physical I/O headers (NPBUF), 14
 Physical memory size, 8
 Physical slices, 58
 physical storage device definition, 65, 71
 Pipe device block allocation, 12
 pipes, 22
 PLCHCORE, 15, 42
 PRIV_SU, 23
 Privilege logging, 26
 Process error thresholds, 19
 Processes that may be active simultaneously
 (NPROC), 14
 profile file, 112
 Program range error maximum, 19
 Programming environment parameters, 31, 36
 Prototype network devices for Model E based
 systems, 87
 Pseudo terminals (tty/pty pairs), 14

Q

qdd physical device, 73
 QUEUE_IO_WITHOUT_NICE_PRIORITY_VALUES, 12
 Quota entries, 14
 Quotas for files, 35

R

RAID bit and range numbers, 59
 RAM
 See Random access memory
 RAM device definition, 64, 71
 RAM device type, 63
 Random access memory (RAM) disk, 42
 Random access memory (RAM) disk device (RDD)
 slices, 57
 RANDOM_PASS_ON, 23
 Range numbers
 unit bit and, 59
 rcoptions file, 112
 RDD
 See Random access memory (RAM) disk device
 RDDSLMAX parameter, 57
 Read down on pipes, 24

- Read-ahead blocks per read operation, 12
 - Release number, 34
 - release number (CONFIG_RLS_REVISION), 34
 - Remote process control system, 36
 - Remove request logging, 25
 - Repeatable relocatables capability, 34
 - Request cache, 83
 - Request cache size, 28
 - Reserved keywords in CSL, 115
 - Restart and checkpoint buffers, 13
 - Retired security logs, 25
 - Retrieval mode for offline files, 12
 - Revision identifiers in CSL, 37
 - revision section, 99
 - rm call logging, 25
 - rm violation logging, 27
 - rmdir call logging, 25
 - rmdir violation logging, 27
 - Rnode number, 29
 - Rnode table for NFS, 83
 - Root privilege, 23
 - Routing, 43
 - RPC
 - See Remote process control system
- S**
- SANITIZE_PATTERN, 23
 - scrub pattern for disks, 23
 - Scrubbing of data, 24
 - sdaemon command
 - use of daemons file, 104
 - SDD
 - See Striped disk device
 - sdd striped device, 73
 - SDDSLMAX parameter, 57
 - Secure file system export, 23
 - SECURE_MAC parameter, 24
 - SECURE_OPERATOR_CONSOLE, 24
 - SECURE_PIPE, 24
 - SECURE_SCRUB, 24
 - SECURE_SYSTEM_CONSOLE parameter, 24
 - security, 27
 - Security level minimum, 22
 - Security levelmaximum, 22
 - Security log buffer size, 25
 - Security log name, 25
 - Security log size, 26
 - security logs location, 25
 - Security logs, retired, 25
 - SEMAEM, 20
 - Semaphore individual value, 21
 - Semaphore sets, 20
 - Semaphores in the system, 20
 - Semaphores per set, 20
 - SEMMNI, 20
 - SEMMNS, 20
 - SEMMNU, 20
 - SEMMSL, 20
 - semop(2) system call adjust-on-exit value, 20
 - semop(2) system call operations, 21
 - SEMOPM, 21
 - SEMUME, 21
 - SEMVMX, 21
 - Separators in CSL, 38
 - Serial number, 32
 - Sessions open simultaneously, 14
 - setuid / setgid files, 22
 - setuid system call logging, 27
 - Shared memory identifiers, 18
 - Shared memory kernel parameters, 18
 - Shared memory parameters, 18
 - Shared segment size maximum, 18
 - shared segment size minimum, 18
 - Shared segments, 18
 - Shared-text programs that can be running
 - simultaneously, 15
 - shmem table size, 18
 - SHMMAX, 18
 - SHMMIN, 18
 - SHMMNI, 18
 - SHMSEG, 18
 - shutdown shell script, 113
 - Side door buffer, 14

SIGCPULIM signal, number of CPU seconds that a process or session may use following , 15

SLG_ACT_NFS, 25

SLG_ACT_NQS, 25

SLG_ALL_NAMI, 25

SLG_ALL_RM, 25

SLG_ALL_VALID, 25

SLG_BUFSIZE, 25

SLG_CF_UNICOS, 25

SLG_DIR, 25

SLG_DISCVconfig.h configuration file
discretionary access violations (SLG_DISCV), 25

SLG_FILE, 25

SLG_FILEXFR, 25

SLG_FPREFIX, 25

SLG_JEND, 25

SLG_JSTART, 25

SLG_LINKV, 25

SLG_LOG_AUDIT, 26

SLG_LOG_CHDIR, 26

SLG_LOG_CRL, 26

SLG_LOG_DAC parameter, 26

SLG_LOG_IPNET parameter, 26

SLG_LOG_OPER, 26

SLG_LOG_PRIV, 26

SLG_LOG_SECSYS parameter, 26

SLG_LOG_SHUTDOWN, 26

SLG_MANDV, 26

SLG_MAXSIZE parameter, 26

SLG_MAXSIZEconfig.h configuration file, 26

SLG_MKDIRV, 26

SLG_NETWV, 26

SLG_PATH_TRACK, 27

SLG_PHYSIO_ERR, 27

SLG_REMOVEV, 27

SLG_RMDIRV, 27

SLG_STATE, 27

SLG_SUID_RQ, 27

SLG_SULOG, 27

SLG_T_PROC, 27

SLG_USER, 27

SLGOFF, 26

slice
ssdi, 63
ssdt, 63

Slice structures, 14

Slices in CSL, 38

SN, 8

sn.h configuration file, 1, 7

Socket space, 83

sockets, 22

solid-state storage device (SSD) definition, 72

SSD, 41, 46, 54-55

SSD device type, 64

SSD logical device cache headers, 14

SSD memory access, 3

SSD side door buffer, 14

SSD-I, 3
disk parameters, 58
filesystem section, 61, 79
gr_union subsection, 44
mainframe section, 47
ssdi definition, 68
ssdi_name, 69
SSDIMAX parameter, 58, 60
SSDISLMAX parameter, 60
tmttype_number, 69

SSD-I description, 75

SSD-I device type, 64

SSD-I devices, 58

SSD-I slice allocation, 58

SSD-T90
filesystem section, 61
gr_union subsection, 44
SSDTMAX parameter, 60
SSDTSLMAX parameter, 60

SSD-T90 description, 74

SSD-T90 device type, 64

SSD-T90 devices, 58

SSD-T90 slice allocation, 58

SSDDSLMAX parameter, 59

ssdi (SSD-I) definition, 68

ssdi_name
SSD-I device name, 69

-
- SSDIMAX parameter, 58
 - SSDISLMAX parameter, 58
 - ssdt (SSD-T) definition, 66
 - SSDTMAX parameter, 58
 - SSDTSLMAX parameter, 58
 - Static client handles in CNFS
 - (CNFS_STATIC_CLIENTS), 29
 - Static client handles in NFS
 - (NFS_STATIC_CLIENTS), 29
 - Striped disk device (SDD) slices, 57
 - Structures in the system, 20
 - su(1) command attempt logging, 27
 - Subdirectory permissions, 86
 - Subsystem parameters, 31
 - Subsystems
 - See config.mh configuration file
 - SuperRing
 - configuration rules, 4
 - definition, 4
 - Support software directory, 34
 - System compartments, 27
 - System dump parameters, 41
 - system high/system low MAC, 24
 - System parameters, 31
 - System shutdown logging, 26
 - System workstation
 - See Operator workstation
 - SYSTEM_ADMIN_CONSOLE, 27
 - SYSVCOMPS, 27
- T**
- Table size parameters, 56
 - Tape capability, 36
 - Tape devices, 56
 - Tape parameters, 56
 - TAPE_MAX_CONF_UP parameter, 56
 - TAPE_MAX_DEV parameter, 56
 - TAPE_MAX_PER_DEV parameter, 56
 - Target information, 45
 - Target of the system to generate, 32
 - TCP/IP, 11
 - TCP/IP kernel parameters, 28
 - TCP/IP managed memory buffers, 28, 84
 - TCP/IP network system, 36
 - TCP_NMBSPACE, 28
 - tcp_numbspace parameter, 84
 - Temporary client handles in CNFS, 29
 - temporary client handles in CNFS
 - (CNFS_TEMP_CLIENTS), 29
 - Temporary client handles in NFS, 29
 - temporary client handles in NFS
 - (NFS_TEMP_CLIENTS), 29
 - Temporary-file directory, 32
 - Terminal I/O buffers, 13
 - Time out, 86
 - tmttype, 66, 68
 - tmttype_number
 - SSD-I target memory type, 69
 - Trusted process activity logging, 27
 - Trusted UNICOS, 36
 - tty/pty pairs, 14
 - Tunable parameters, 55
- U**
- U_MAXPACK, 16
 - umountem shell script, 113
 - Undo entries per process, 21
 - UNICOS installation, 1
 - unicos section, 55
 - UNICOS statement, 55
 - UNICOS/mk operating system
 - GigaRing based systems, 32
 - IOS-E based systems, 33
 - Unique files that can be open at one time, 14
 - Unlinked files for checkpoint and restart, 15
 - User number defined for the fair share
 - scheduler, 15
 - user packets, 16
 - uts kernel, 35
- V**
- Version name, 32

W

Words per unit of dynamic kernel memory, 17

Write cache, 28

X

xdd, 66, 68

XDDMAX parameter, 58

xddphysical device, 73

XDDSLMAX parameter, 58

xdisk, 62–63, 69

xdisk definition, 69

xmemory, 49

XTRASEC, 15