

# HP Instant Capacity user's guide for Version 10.x

HP-UX 11i v3, HP-UX 11i v2, and HP-UX 11i v1



© Copyright 2000-2010 Hewlett-Packard Development Company, L.P

**Legal Notices**

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

UNIX is a registered trademark of The Open Group.

---

# Table of Contents

About This Guide.....	11
Intended Audience.....	11
New and Changed Information in This Edition.....	11
Publishing History.....	12
Document Organization.....	13
Typographic Conventions.....	14
Location of Instant Capacity Information.....	15
Instant Capacity user's guide for Version 10.x.....	15
Locating the Instant Capacity release notes for Version 10.x.....	15
Manpages.....	16
HP Encourages Your Comments.....	16
1 Introduction to Instant Capacity.....	17
Instant Capacity Summary.....	18
Terminology.....	19
Overview.....	20
Software Product Overview.....	20
System Overview.....	20
Instant Capacity System Hardware.....	21
Instant Capacity Software.....	21
Utility Pricing Solutions Portal.....	22
Instant Capacity Administration System.....	22
Instant Capacity Database.....	22
Other System Management Commands.....	22
Most Recent Instant Capacity Product Versions and Supported Platforms.....	23
Past Versions and Supported Operating Systems.....	25
2 Getting Started.....	27
Instant Capacity Requirements.....	28
Program Requirements.....	28
Software Requirements.....	28
HP-UX 11i v1 Requirements and Dependent Products.....	28
Required Patches for HP-UX 11i v1.....	29
HP-UX 11i v2 Requirements and Dependent Products.....	29
Required Patches for HP-UX 11i v2.....	29
HP-UX 11i v3 Requirements and Dependent Products.....	30
OpenVMS Version 8.4 Requirements.....	30
Network Requirements.....	30
Email Requirements.....	30
System Contact Requirement.....	31
Usage Rights Requirement.....	31
Instant Capacity Integration with Virtual Partitions (HP-UX).....	31
Instant Capacity Components.....	33
Overview.....	33
Processors and Cores.....	33
Cell Boards and Memory.....	33
Global Instant Capacity.....	35
Instant Capacity Codewords.....	36
Temporary Instant Capacity.....	37

Instant Capacity Compliance and Enforcement.....	38
Configuration Change Notification.....	39
Core Activation.....	40
Increasing Processing Capacity by Purchasing RTUs.....	40
Instant Capacity Cell Board.....	41
Instant Capacity Software Validation.....	42
On HP-UX Systems.....	42
On OpenVMS Systems.....	42
Status Reporting on Instant Capacity Systems.....	43
Time Zone Considerations.....	44
<b>3 Installing and Removing Instant Capacity Software.....</b>	<b>45</b>
Installing Instant Capacity Software.....	46
Installing from the HP-UX Media (HP-UX 11i v1, 11i v2, and 11i v3):.....	46
Installing from the HP Software Depot (For HP-UX 11i v1, 11i v2, and 11i v3):.....	47
For All HP-UX Installations.....	47
Installing Instant Capacity on OpenVMS Systems.....	47
Reinstalling Instant Capacity Software.....	48
Preserving current Instant Capacity information.....	48
Removing Instant Capacity Software.....	49
<b>4 Using Instant Capacity to Manage Processing Capacity.....</b>	<b>51</b>
Checking the Status of your Instant Capacity System.....	52
Setting System Contact Information.....	54
Applying a Right To Use (RTU) Codeword.....	55
Activating Cores.....	57
Example Core Activation Session.....	58
Deactivating Cores.....	59
Example Deactivation Session for Hardware-Partitionable Systems.....	59
Overriding Deferred Activation and Deactivation.....	61
Load-Balancing Active Cores.....	62
Understanding and Managing Intended Active Values.....	63
Activations and Deactivations in a Virtual Partition Environment.....	64
When to Use the vparmodify or icapmodify Commands.....	64
What is Unused Capacity.....	64
Static Virtual Partitions.....	64
Activating Cores in a Virtual Partition Environment.....	65
Deactivating Cores in a Virtual Partition Environment.....	65
Boot Time Compliance.....	66
Assigning a Cell to a Partition.....	67
Unassigning a Cell from a Partition.....	69
Software Application Considerations.....	71
Test Activation of Cores Using Temporary Capacity.....	72
Replacement of Failed Cores.....	73
HP-UX LPMC and HPMC.....	73
LPMC Deactivations in Virtual Partitions.....	73
Failed Monarch Processors (HP-UX).....	73
Replacement of Failed Cores on OpenVMS.....	73
Failed OpenVMS Primary Processors.....	73
<b>5 Temporary Instant Capacity.....</b>	<b>75</b>
Temporary Instant Capacity Overview.....	76
Ordering Temporary Instant Capacity.....	78

HP-UX Licensing and Support with Temporary Instant Capacity.....	78
OpenVMS Licensing and Support with Temporary Instant Capacity.....	78
Using Temporary Instant Capacity.....	79
Acquiring and Configuring Temporary Instant Capacity.....	79
Using Temporary Instant Capacity.....	79
Temporary Capacity and Virtual Partitions.....	81
Tracking Usage of Temporary Instant Capacity.....	83
Temporary Instant Capacity Warning Period.....	85
Temporary Instant Capacity Expiration and Compliance Enforcement.....	86
Temporary Instant Capacity Exceptions.....	87
Error for Activation with Insufficient Temporary Capacity.....	87
Temporary Capacity Balance Needing Action.....	87
Temporary Capacity Negative Balance.....	87
Temporary Capacity Enforcement.....	87
<b>6 Instant Capacity Cell Board.....</b>	<b>91</b>
Instant Capacity Cell Board.....	92
Overview.....	92
Ordering Instant Capacity Cell Board.....	93
HP-UX and OpenVMS License and Support.....	94
Acquiring Usage Rights for Instant Capacity Cell Board.....	95
Instant Capacity Cell Board and Considerations of Core Usage Rights.....	96
Activation of an Instant Capacity Cell Board.....	97
Accidental Activation of an Instant Capacity Cell Board.....	98
Instant Capacity Cell Board Activation Exception Error.....	99
Instant Capacity Cell Board and Temporary Instant Capacity.....	100
<b>7 Global Instant Capacity.....</b>	<b>101</b>
Global Instant Capacity Overview.....	102
Global Instant Capacity Requirements.....	104
Global Instant Capacity Group Managers.....	105
Standby Group Managers.....	105
Global Instant Capacity Grouping Rules.....	107
Global Instant Capacity Sharing Rights.....	108
Creating Global Instant Capacity Groups.....	109
Global Instant Capacity Resource Sharing.....	112
Effect of Temporary Capacity.....	112
Status Reporting.....	113
Global Instant Capacity and Temporary Capacity.....	115
Temporary Capacity and Freed Usage Rights.....	115
Temporary Capacity and Status Reporting.....	115
Temporary Capacity Prefetch.....	116
Removing a Global Instant Capacity Group Member.....	117
Reinstalling a Group Member.....	118
Group Manager Availability (No Standby Manager).....	119
Group Manager Failover Considerations.....	120
Upgrades and Global Instant Capacity.....	121
Adding New Partitions.....	121
Rights Seizure.....	122
When to Migrate Usage Rights and When to Seize Usage Rights.....	122
Effects of Rights Seizure.....	122
Down Partitions with Powered-On Cells.....	123
Temporary Capacity and Rights Seizure.....	123

Other Considerations.....	124
Additional HA Solutions.....	124
Summary of Rights Seizure.....	124
Considerations for Multiple Groups.....	127
Additional Considerations.....	128
<b>8 Using Instant Capacity on HP Integrity Superdome 2.....</b>	<b>129</b>
Overview.....	129
iCAP: From compliance check paradigm to self enforcement.....	129
Important Considerations.....	130
Sizing Partitions.....	130
Resizing Partitions.....	130
Memory iCAP.....	130
Installing iCAP on HP Integrity Superdome 2.....	130
iCAP commands.....	131
Viewing number of active cores.....	131
Generating snapshot.....	131
Setting usage rights for a partition.....	132
Deactivating Cores.....	132
Activating Cores.....	132
Using iCAP Memory.....	132
iCAP Use Cases.....	133
Core migration between partitions in a complex.....	133
Core migration across complexes (GiCAP).....	134
<b>9 Troubleshooting.....</b>	<b>137</b>
Handling Compliance Exceptions.....	137
Troubleshooting the Instant Capacity Software.....	140
Diagnosing Email Configuration.....	142
<b>10 Frequently Asked Questions.....</b>	<b>143</b>
Instant Capacity Software.....	144
Instant Capacity Hardware.....	147
Global Instant Capacity.....	148
<b>1 Instant Capacity HP-UX Manpages.....</b>	<b>149</b>
icap(5).....	150
icapmanage(1M).....	159
icapmodify(1M).....	174
icapnotify(1M).....	180
icapstatus(1M).....	182
icapid(1M).....	189
<b>A Special Considerations.....</b>	<b>191</b>
Assumed Values in icapstatus Command.....	192
Assumed Processor Values.....	192
Assumed Memory Values.....	192
Upgrading to Instant Capacity version B.06.x or later (HP-UX).....	193
Dual-Core Support in Instant Capacity Systems.....	195
New Partition Creation and Instant Capacity.....	196
Implications of Removing a Cell from an Instant Capacity System .....	197

Shutting Down a Partition with Instant Capacity Cores.....	198
Instant Capacity and Reinitializing the nPartition (Genesis Partitions).....	199
par Commands from PC System Management Station.....	200
Instant Capacity Compatibility with Processor Sets (HP-UX).....	201
Overview.....	201
Scope of the Instant Capacity Software Interacting with psets.....	201
psets on nPars.....	201
psets on vPars.....	201
Configuring Email on Instant Capacity Systems.....	202
Email Requirements.....	202
Email Configuration.....	202
Steps to Confirm or Diagnose Email Configuration.....	204
Configuring Instant Capacity's From Email Address.....	204
Configuring Your Server to Send but Not Receive Email.....	205
Testing Email Transmission of the Asset Report.....	205
Measurement Software and Instant Capacity Systems.....	206
HP OpenView Measurement Products.....	206
Other Measurement Software.....	206
Dynamic Processor Resilience (HP-UX).....	207
Security Issues.....	208
Customer protections which iCAP assumes to be in place.....	208
Disabling the iCAP daemon (HP-UX).....	208
Customer Security Requirements.....	208
Security Tuning Options.....	208
<b>B Considerations for OpenVMS Systems.....</b>	<b>209</b>
CLI Support on OpenVMS.....	210
HP-UX Style Commands.....	210
OpenVMS Command Mapping.....	210
OpenVMS iCAP Files.....	211
DCL ICAP Command.....	212
DCL Commands.....	212
ICAP ACTIVATE.....	212
ICAP APPLY.....	213
ICAP DEACTIVATE.....	214
ICAP RECONCILE.....	215
ICAP SET.....	216
ICAP SHOW.....	218
ICAP_SERVER.....	219
Special OpenVMS-Specific Features and Considerations.....	220
Core Activation and Deactivation.....	220
Email Considerations.....	220
Restrictions.....	221
 Glossary.....	 223
 Index.....	 229

---

# List of Figures

1-1	Instant Capacity System Elements.....	21
4-1	Permanent Activation of Instant Capacity Components.....	56
4-2	Partition premodification state: One cell assigned with 3 active and 1 inactive cores, and usage rights for 2 additional cores.....	67
4-3	Premodification state: Unassigned cell with 4 inactive cores and no usage rights.....	67
4-4	Partition postmodification state: Cell 2 assigned to partition.....	67
4-5	Partition premodification state: Three cells with 2 active and 2 inactive cores in each, and 6 expected inactive cores.....	69
4-6	Partition postmodification state: Cell 3 is unassigned (total of 6 active cores remaining).....	69
4-7	Partition postmodification state: Unassigned Cell 3 with 4 inactive cores.....	69
4-8	Partition premodification state: Three cells with 3 active and 1 inactive cores in each, and 3 expected inactive cores.....	69
4-9	Partition postmodification state: Unassigned Cell 3 (total of 8 active cores are set).....	70
4-10	Postmodification state: Unassigned Cell 3 with 4 inactive cores, with usage rights available for 1 additional core.....	70
5-1	Using Temporary Instant Capacity: Temporary Activation of Cores Without Usage Rights.....	77
6-1	nPartition premodification state: One cell assigned with 1 active core and 3 inactive cores; the complex has no additional core usage rights.....	99
6-2	nPartition premodification state: Instant Capacity cell (Cell 2) with 4 inactive cores.....	99
6-3	nPartition requested state: Instant Capacity Cell (Cell 2) cannot be activated in nPartition.....	99
7-1	Using Global Instant Capacity.....	103
8-1	iCAP Memory configuration option.....	133

---

# List of Tables

1-1	Most Recent Instant Capacity Versions and Supported Platforms.....	23
6-1	Cell Board Activation Not Requiring Additional Core Usage Rights.....	96
6-2	Cell Board Activation Requiring Additional Core Usage Rights.....	96
8-1	Use Case 1 – Initial Configuration.....	133
8-2	Use Case 1 – Configuration of Complex After Core Migration.....	134
8-3	Use Case 2 – Initial Configuration.....	134
8-4	Use Case 2 – Initial Configuration.....	134
8-5	Configuration of Complex1 post core migration.....	134
8-6	Configuration of Complex2 post core migration.....	135
10-1	Email sent by the Instant Capacity software.....	145
10-2	Asset reporting email sent by the Instant Capacity software.....	146
A-1	Removing a Cell and Decreasing Inactive Cores.....	197
B-1	HP-UX and OpenVMS Command Equivalents.....	210
B-2	OpenVMS iCAP Files.....	211

---

## List of Examples

2-1	Configuration Change Notification email for Instant Capacity System (not vPar).....	39
4-1	Sample Session of icapstatus (on HP-UX).....	53
4-2	Applying an RTU Codeword (HP-UX).....	55
4-3	Activating an Additional Core (HP-UX).....	58
4-4	Deactivating an Active Core (HP-UX).....	60
4-5	Correcting an Incorrect Number of Deferred Active Cores (HP-UX).....	61
4-6	Undoing an Accidental Deferred Activation (HP-UX).....	61
4-7	vPar Boot-Time Compliance Message.....	66
5-1	Applying a Temporary Capacity Codeword (HP-UX).....	79
5-2	Activating an Instant Capacity Core with Temporary Capacity (HP-UX).....	80
5-3	Temporary Capacity Information from icapstatus Command (HP-UX).....	83
5-4	Temporary Capacity Expiration Reminder.....	84
5-5	Error Message for Activation with Insufficient Temporary Capacity (HP-UX).....	87
5-6	Error Message for Temporary Capacity Partial Enforcement.....	88
5-7	Error Message for Temporary Capacity Complete Enforcement.....	89
7-1	Applying a Sharing Rights Codeword.....	108
7-2	Creating a Group.....	109
7-3	Adding a Member to a Group.....	110
7-4	Example Output of icapstatus for a Group Member System.....	111
9-1	Exception Report for More Cores Active than Expected.....	138
A-1	Example Edit to Sendmail Configuration (/etc/mail/sendmail.cf).....	203

---

# About This Guide

The *HP Instant Capacity user's guide for Version 10.x* provides you with the most recent information for using the Instant Capacity Version 10 software. This document describes Instant Capacity Version B.11.31.10.00.00 on HP-UX 11i v3 systems, B.11.23.10.00.00 on HP-UX 11i v2 systems, B.11.11.09.02.00 on HP-UX 11i v1 systems, and Instant Capacity Version 9.x on OpenVMS 8.4 Integrity servers.

The latest version of this document can be found online at: [www.hp.com/go/hp-icap-docs](http://www.hp.com/go/hp-icap-docs).

This chapter covers the following topics:

- “Intended Audience” (page 11)
- “New and Changed Information in This Edition” (page 11)
- “Publishing History” (page 12)
- “Document Organization” (page 13)
- “Typographic Conventions” (page 14)
- “Location of Instant Capacity Information” (page 15)
- “HP Encourages Your Comments” (page 16)

## Intended Audience

All personnel with system administrator access (that is, with `root` login privileges on HP-UX systems, or Administrators on OpenVMS systems) to an Instant Capacity system should read and understand the contents of this document and the implications of managing an Instant Capacity system.

Administrators are expected to have knowledge of HP-UX, or OpenVMS Integrity operating system concepts, commands, and configuration.

This document is not a tutorial.

## New and Changed Information in This Edition

This edition of *HP Instant Capacity user's guide for Version 10.x* includes a new chapter that provides information about iCAP support on HP Integrity Superdome 2. For detailed information about iCAP support on HP Integrity Superdome 2, see [Chapter 8 \(page 129\)](#)

## Publishing History

The document printing date and part number indicate the document's current edition. The printing date will change when a new edition is printed. Minor changes may be made at reprint without changing the printing date. The document part number will change when extensive changes are made. Document updates may be issued between editions to correct errors or document product changes. To ensure that you receive the updated or new editions, you should subscribe to the appropriate product support service. See your HP sales representative for details. You can find the latest version of this document on line at:

[www.hp.com/go/hp-icap-docs](http://www.hp.com/go/hp-icap-docs).

<b>Manufacturing Part Number</b>	<b>Supported Operating Systems</b>	<b>Supported Versions</b>	<b>Edition Number</b>	<b>Publication Date</b>
T1428-90074	HP-UX 11i v2	B.11.23.10.00.00	3	<i>September 2010</i>
	HP-UX 11i v3	B.11.31.10.00.00		
B9073-90201	HP-UX 11i v1	B.11.11.09.00.00	2	<i>November 2009</i>
	HP-UX 11i v2	B.11.23.09.00.00		
	HP-UX 11i v3	B.11.31.09.00.00		
	HP OpenVMS Version 8.4	9.x		

# Document Organization

This user's guide is not designed to be read from front to back in its entirety. To get an understanding of Instant Capacity version 9.x, you should read this chapter, *Chapter 1 - Instant Capacity Overview*, and *Chapter 2 - Getting Started*. After reading these chapters, you can utilize the table of contents and index (in back) for specific topics of interest.

- *About This Guide* provides an introduction to this guide and locating Instant Capacity information.
- *Chapter 1, Introduction to Instant Capacity* provides an overview of the Instant Capacity system.
- *Chapter 2, Getting Started* describes Instant Capacity requirements, concepts and methods, and related software topics.
- *Chapter 3, Installing and Removing Instant Capacity Software* contains procedures on how to install and reinstall Instant Capacity software.
- *Chapter 4, Using Instant Capacity to Manage Processing Capacity* explains how to view system status, apply codewords, activate and deactivate cores, assign and unassign cells, and HP's test activation policy for Instant Capacity.
- *Chapter 5, Temporary Instant Capacity* gives you details on what temporary capacity is, and how to order and use it.
- *Chapter 6, Instant Capacity Cell Board* provides details on what Instant Capacity Cell Board is, and how to order and use it.
- *Chapter 7, Global Instant Capacity* gives you details on Global Instant Capacity, describes the concept of groups and shared resources, and how to order and use it.
- *Chapter 8, Using Instant Capacity on HP Integrity Superdome 2* provides information about iCAP support on HP Integrity Superdome 2.
- *Chapter 9, Troubleshooting* gives you step by step procedures to resolve problems with the Instant Capacity software and other related configurations.
- *Chapter 10, Frequently Asked Questions* contains questions and answers to common Instant Capacity software topics.
- *Instant Capacity HP-UX Manpages* contains the actual HP-UX manpages for `icap`, `icapmanage`, `icapmodify`, `icapnotify`, `icapstatus`, and `icapid`.
- *Appendix A, Special Considerations* describes assumed values in `icapstatus` output, upgrading to Instant Capacity B.06.x or later software, dual core support, creating and shutting down partitions, implications of removing a cell board from an Instant Capacity system, `par` commands with PC SMS, integration with Psets, configuring email, testing email transmission of an asset report, measurement software, dynamic processor resilience, and security related issues.
- *Appendix B, Considerations for OpenVMS Systems* contains information for running Instant Capacity on OpenVMS systems.
- *Glossary* explains Instant Capacity systems and software terms.

# Typographic Conventions

This document uses the following conventions:

<code>%</code> , <code>\$</code> , or <code>#</code>	A percent sign represents the C shell system prompt. A dollar sign represents the system prompt for the Bourne, Korn, and POSIX shells. A number sign represents the superuser prompt.
<code>audit(5)</code>	A manpage. The manpage name is <i>audit</i> , and it is located in Section 5.
Command	A command name or qualified command phrase.
Computer output	Text displayed by the computer.
<b>Ctrl+x</b>	A key sequence. A sequence such as <b>Ctrl+x</b> indicates that you must hold down the key labeled <b>Ctrl</b> while you press another key or mouse button.
ENVIRONMENT VARIABLE	The name of an environment variable, for example, <code>PATH</code> .
[ERROR NAME]	The name of an error, usually returned in the <code>errno</code> variable.
<b>Key</b>	The name of a keyboard key. <b>Return</b> and <b>Enter</b> both refer to the same key.
<i>Glossary term</i>	The defined use of an important word or phrase.
<b>User input</b>	Commands and other text that you type.
<i>Variable</i>	The name of a placeholder in a command, function, or other syntax display that you replace with an actual value.
[]	The contents are optional in syntax. If the contents are a list separated by <code> </code> , you must choose one of the items.
{ }	The contents are required in syntax. If the contents are a list separated by <code> </code> , you must choose one of the items.
...	The preceding element can be repeated an arbitrary number of times.
Ⓢ	Indicates the continuation of a code example.
	Separates items in a list of choices.
WARNING	A warning calls attention to important information that if not understood or followed will result in personal injury or nonrecoverable system problems.
CAUTION	A caution calls attention to important information that if not understood or followed will result in data loss, data corruption, or damage to hardware or software.
IMPORTANT	This alert provides essential information to explain a concept or to complete a task
NOTE	A note contains additional information to emphasize or supplement important points of the main text.

## Location of Instant Capacity Information

### Instant Capacity user's guide for Version 10.x

You can find the *HP Instant Capacity user's guide for Version 10.x* in the following locations:

- On the HP Documentation website:  
[www.hp.com/go/hp-icap-docs](http://www.hp.com/go/hp-icap-docs)  
These are the most current and the most current localized versions of the user guide.
- September 2010 HP-UX 11i v1 Instant Information media (English and Japanese only)
- September 2010 HP-UX 11i v2 Instant Information media (English and Japanese only)
- September 2010 HP-UX 11i v3 Instant Information media (English and Japanese only)
- In the Instant Capacity 10.x HP-UX software product:  
`/usr/share/doc/icapRelNotes.pdf`
- On the OpenVMS Version 8.4 Documentation CD (English only)

### Locating the Instant Capacity release notes for Version 10.x

You can find the *Instant Capacity release notes for Version 10.x* in the following locations:

- On the HP Documentation website:  
[www.hp.com/go/hp-icap-docs](http://www.hp.com/go/hp-icap-docs)  
These are the most current release notes and the most current localized versions.
- September 2010 HP-UX 11i v1 Instant Information media (English and Japanese only)
- September 2010 HP-UX 11i v2 Instant Information media (English and Japanese only)
- September 2010 HP-UX 11i v3 Instant Information media (English and Japanese only)
- In the Instant Capacity 10.x HP-UX software product:  
`/usr/share/doc/icapRelNotes.pdf`  
This document is not localized and is available in English only.
- On the OpenVMS Version 8.4 Documentation CD (English only)

## Manpages

---



**NOTE:** The information contained in this section applies only to HP-UX systems. It does not apply to Integrity servers running OpenVMS.

---

See “Instant Capacity HP-UX Manpages” (page 149) for more information about the following manpages:

- `icap(5)`: overview of the Instant Capacity commands and their usage
- `icapmanage(1M)`: how to manage Global Instant Capacity (GiCAP) groups
- `icapmodify(1M)`: how to manage processing capacity in your Instant Capacity system, change system contact information, and apply codewords
- `icapnotify(1M)`: how to manage asset reporting and configuration notification
- `icapstatus(1M)`: how to display processing capacity status, usage information, and system information
- `icapd(1M)`: overview of the Instant Capacity daemon, which provides the software a complexwide view of processing capacity

## HP Encourages Your Comments

HP encourages your comments concerning this document. We are committed to providing documentation that meets your needs. Send any errors found, suggestions for improvement, or compliments to:

[docsfeedback@hp.com](mailto:docsfeedback@hp.com)

Include the document title, manufacturing part number, and any comment, error found, or suggestion for improvement you have concerning this document.

---

# 1 Introduction to Instant Capacity

This chapter covers the following topics:

- “Instant Capacity Summary” (page 18)
- “Overview” (page 20)

For more in-depth information, see `icap(5)`.

# Instant Capacity Summary

HP Instant Capacity software provides the ability to instantly increase or decrease computing capacity on specified HP enterprise servers.



**NOTE:** HP Instant Capacity for HP 9000 and HP Integrity Servers, also known as Instant Capacity or iCAP, was known in earlier versions as Instant Capacity on Demand, or iCOD. Although the commands, warning messages and error messages refer to the software as iCAP, some internal files might still refer to iCOD.

**NOTE:** For simplicity and commonality, this book uses the HP-UX commands in all examples. For information about OpenVMS command equivalents, see [Appendix B \(page 209\)](#).

With Instant Capacity, you initially purchase an HP enterprise server with a specified amount of active processing capacity, and a specified amount of inactive processing capacity. This amount can vary based on your sales contract with HP.

Processing capacity consists of the system components:

- *Processors* containing *cores*
- *Cell boards*
- Memory

For each type of component, the number of components that can be active is equal to the number of usage rights applied to the complex for that type of component. Components that are purchased with a part number identifying them as “Instant Capacity” and without the label “Right to Use” come without usage rights. Components that are not labeled “Instant Capacity” implicitly include usage rights that can be applied to any component of that type that is installed on the complex.

Prior to activation of an inactive component, you must obtain additional *usage rights*. The fundamental method is to purchase usage rights by purchasing the appropriate Instant Capacity products that include the label “Right to Use (RTU)”. HP then supplies an *RTU codeword*. When the codeword is applied to the HP enterprise server, the inactive component can be activated.

Additional methods for activating components for which usage rights have not been purchased include:

- If an HP-UX server is a member of a *Global Instant Capacity (GiCAP)* group, and if extra capacity is available from other members of the group, capacity can be “borrowed” from another member of the group. For information about GiCAP, see [Chapter 7](#).
- You can purchase additional temporary capacity and apply the *Temporary Instant Capacity (TiCAP)* codeword in order to activate one or more cores temporarily. For more information on TiCAP, see [Chapter 5](#). If a server is a member of a GiCAP group, temporary capacity can be shared among members of the group.
- You can temporarily activate one or more inactive cores using the *Instant Access Capacity (IAC)* provided with the initial purchase of the Instant Capacity component. Instant Access Capacity is the same as TiCAP except it is automatically provided with an Instant Capacity component and cannot be purchased separately. It provides an immediate buffer of temporary capacity in case extra capacity is needed before there is time to purchase an RTU or a TiCAP codeword, or to set up a GiCAP group on an HP-UX system.



**IMPORTANT:** It is always a good idea to keep some quantity of temporary capacity in reserve. Purchase of codewords may take one or more days, so having a buffer of temporary capacity allows you to avoid delays in activation of additional cores. Instant Access Capacity provides this buffer initially, but as that capacity is depleted, ongoing purchases of additional temporary capacity are recommended to replenish it.

The Instant Capacity software product is a part of the HP Utility Pricing Solutions (formerly On Demand Solutions) program.

This document provides you with the most recent information about using the Instant Capacity version 9.x software to manage processing capacity in your HP enterprise server.

---



**NOTE:** All personnel with system administrator access (that is, root login privileges) to an Instant Capacity system should read and understand the contents of this document and the implications of increasing or decreasing processing capacity.

---

## Terminology

For commonly used terms associated with the HP Utility Pricing Solutions program, see the Glossary.

# Overview

## Software Product Overview

The following Instant Capacity software products are associated with HP's Utility Pricing Solutions program:

- **iCOD**: HP product number B9073BA (HP-UX)
- **iCAP**: HP OpenVMS product number BA484AA

This user's guide contains information for Instant Capacity version 10.x on HP-UX systems and version 9.x on OpenVMS systems.

Instant Capacity must be run on a partitionable system. In an *Integrity Virtual Machines (VM)* environment, Instant Capacity software provides meaningful functionality only on the *VM Host* system; it does not run on a *virtual machine* (also known as a "guest").

The Instant Capacity product has been available on HP-UX since March 2000 (version B.01.00).

The HP-UX version 9.x software can be obtained from the following HP Software Depot website (search for "Instant Capacity"):

<http://www.hp.com/go/softwaredepot>

## System Overview

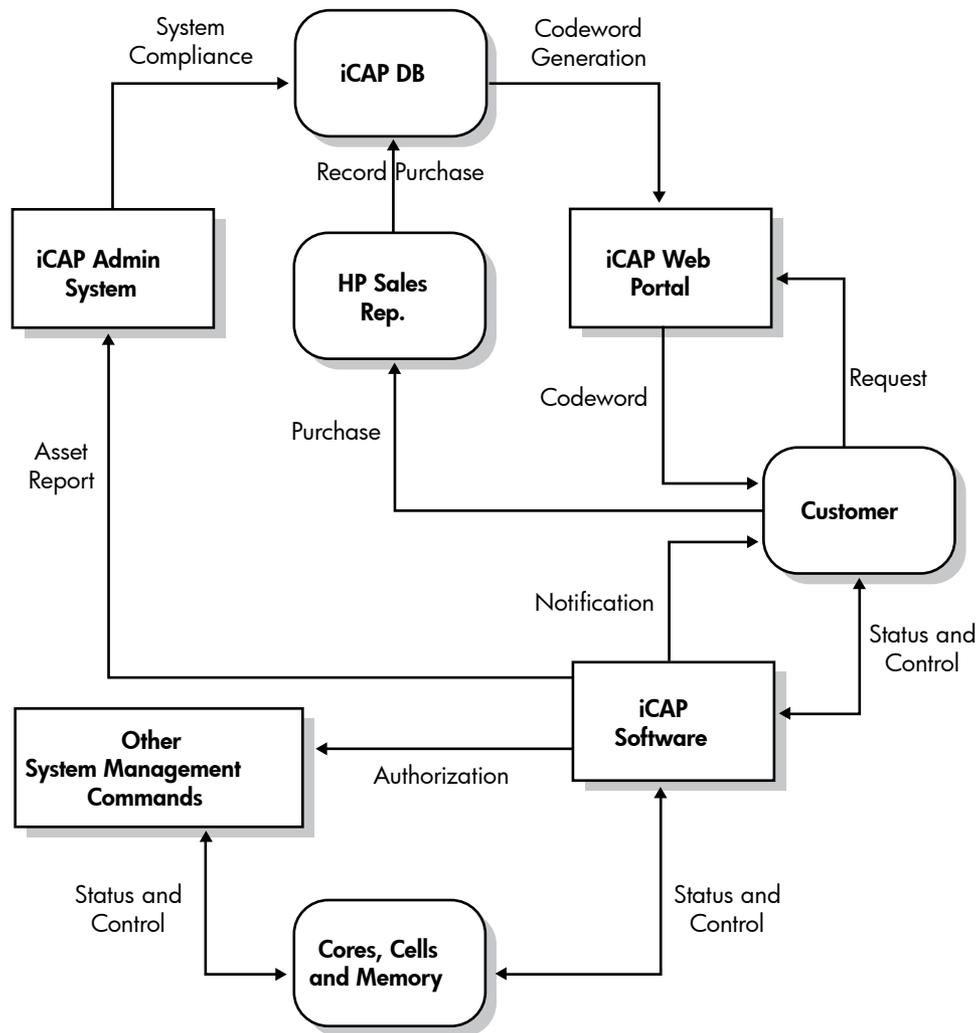
Instant Capacity version 9.x consists of the following elements:

- Instant Capacity system hardware (including cells, cores, and memory)
- Instant Capacity software
- **Utility Pricing Solutions Portal**
- Instant Capacity Administration System
- Instant Capacity database
- Other system management commands

For more information about Instant Capacity concepts and methods, see Chapter 2: "Getting Started" (page 27).

Figure 1-1 (page 21) shows how the Instant Capacity system elements and processes are linked.

**Figure 1-1 Instant Capacity System Elements**



## Instant Capacity System Hardware

The Instant Capacity system hardware is made up of the following components:

- Cell boards
- Processors, which contain cores
- Memory

Every Instant Capacity system contains a combination of these components that are purchased either with usage rights (and available for activation) or without usage rights (must be inactive). Note that although you purchase processors, the Instant Capacity software monitors and manages cores.

## Instant Capacity Software

The Instant Capacity software provides the means to:

- Increase or decrease (load balance) system processing capacity (`icapmodify` command).
- View status and configuration of the system components (`icapstatus` command).
- Administer system identification and notification information (`icapmodify` command).
- If configured, send system asset reports through encrypted email to HP (`icapd` daemon on HP-UX, `ICAP_SERVER` process on OpenVMS).
- Send configuration change notification, through encrypted email, to the specified system contact.

- Monitor and report system compliance (`icapd` daemon on HP-UX, `ICAP_SERVER` process on OpenVMS).
- Manage Global Instant Capacity groups (`icapmanage` command).

For details about these commands on HP-UX, see “Instant Capacity HP-UX Manpages” (page 149). For the OpenVMS equivalents of these commands, see “DCL Commands” (page 212).

## Utility Pricing Solutions Portal

The Utility Pricing Solutions (or Instant Capacity) portal is located at the HP web site:

<http://www.hp.com/go/icap/portal>

After you purchase a component without usage rights, HP sends you a letter containing instructions on how to obtain an RTU codeword from the Utility Pricing Solutions portal.

## Instant Capacity Administration System

If asset reporting is configured, the `icapd` daemon sends asset reports, in the form of encrypted email messages, to the Instant Capacity Administration System, which saves information in the Instant Capacity database.

## Instant Capacity Database

The Instant Capacity database is a repository on an HP (internal) corporate server that tracks system resources and provides the information for codeword generation. Note that this database should not be confused with a Global Instant Capacity database, which is created on a customer Group Manager system. See Chapter 7 (page 101) for details about Global Instant Capacity.

## Other System Management Commands

Other system management commands (such as `vparmodify`, `parCLI` and `parMgr` on HP-UX) provide an interface to modify system configuration that affects Instant Capacity contractual compliance.

## Most Recent Instant Capacity Product Versions and Supported Platforms

Table 1-1 lists the current versions of Instant Capacity and the platforms supported for each version.

**Table 1-1 Most Recent Instant Capacity Versions and Supported Platforms**

Software and Version	Operating System Version	Supported Hardware Platforms	Notes
iCAP B.11.31.10.00.00 (B9073BA)	HP-UX 11i v3	HP Integrity servers: <ul style="list-style-type: none"> <li>• HP Integrity Superdome 2</li> <li>• Superdome</li> <li>• rx8640</li> <li>• rx8620</li> <li>• rx7640</li> <li>• rx7620</li> </ul> HP 9000 servers: <ul style="list-style-type: none"> <li>• Superdome</li> <li>• rp8440</li> <li>• rp8420</li> <li>• rp8400</li> <li>• rp7440</li> <li>• rp7420</li> <li>• rp7410</li> </ul>	Available on: <ul style="list-style-type: none"> <li>• <a href="http://www.hp.com/go/softwaredepot">http://www.hp.com/go/softwaredepot</a></li> <li>• September 2009 HP-UX 11i v3 Operating Environment media</li> </ul>
iCAP B.11.23.10.00.00 (B9073BA)	HP-UX 11i v2	HP Integrity servers: <ul style="list-style-type: none"> <li>• Superdome</li> <li>• rx8640</li> <li>• rx8620</li> <li>• rx7640</li> <li>• rx7620</li> </ul> HP 9000 servers: <ul style="list-style-type: none"> <li>• Superdome</li> <li>• rp8420</li> <li>• rp8400</li> <li>• rp7420</li> <li>• rp7410</li> </ul>	Available on: <ul style="list-style-type: none"> <li>• <a href="http://www.hp.com/go/softwaredepot">http://www.hp.com/go/softwaredepot</a></li> <li>• September 2009 HP-UX 11i v2 Applications Software media</li> </ul>
iCAP B.11.11.09.00.00 (B9073BA)	HP-UX 11i v1	HP 9000 servers: <ul style="list-style-type: none"> <li>• Superdome</li> <li>• rp8440</li> <li>• rp8420</li> <li>• rp8400</li> <li>• rp7440</li> <li>• rp7420</li> <li>• rp7410</li> </ul>	Available on: <ul style="list-style-type: none"> <li>• <a href="http://www.hp.com/go/softwaredepot">http://www.hp.com/go/softwaredepot</a></li> <li>• September 2009 HP-UX 11i v1 Applications Software media</li> </ul>
iCAP 9.x (BA484AA)	HP OpenVMS Version 8.4	HP Integrity servers: <ul style="list-style-type: none"> <li>• Superdome</li> <li>• rx8640</li> <li>• rx8620</li> <li>• rx7640</li> <li>• rx7620</li> </ul>	Available on: <ul style="list-style-type: none"> <li>• OpenVMS Version 8.4 Operating System media</li> </ul>

For commonality across all platforms supported by Instant Capacity, generic references to version numbers are in the form “9.x” in this document. Because the operating system version is also

incorporated in the depot version number on newer releases of HP-UX, specific references to version numbers may be more precise; for example, B.11.23.09.00.00 for the HP-UX 11i v2 depot.

## Past Versions and Supported Operating Systems

Previous versions of the Instant Capacity software are as follows:

- B.01.00 (on HP-UX 11.00)
- B.02.x (on HP-UX 11.00 and 11i v1)
- B.03.x (on HP-UX 11i v1)
- B.04.x (on HP-UX 11.00 and 11i v1)
- B.05.00 (on HP-UX 11.00 and 11i v1)
- B.06.x (on HP-UX 11i v1 and 11i v2)
- B.07.x (on HP-UX 11i v1 and 11i v2)
- B.08.00 (on HP-UX 11i v1 and 11i v2)
- B.08.00.01 (on HP-UX 11i v1 and 11i v2)
- B.08.01.01 (on HP-UX 11i v1, 11i v2 and 11i v3)
- B.08.02 (on HP-UX 11i v1, 11i v2 and 11i v3)
- B.08.02.01 (on HP-UX 11i v1, 11i v2 and 11i v3)
- B.08.03.00 (on HP-UX 11i v1, 11i v2, and 11i v3)
- B.08.03.01 (on HP-UX 11i v1, 11i v2, and 11i v3)
- 8.0 (on OpenVMS Version 8.3)
- 8.1 (on OpenVMS Version 8.3-1H1)



---

## 2 Getting Started

This chapter covers the following topics:

- “Instant Capacity Requirements” (page 28)
- “Instant Capacity Components” (page 33)
- “Global Instant Capacity” (page 35)
- “Instant Capacity Codewords” (page 36)
- “Temporary Instant Capacity” (page 37)
- “Instant Capacity Compliance and Enforcement” (page 38)
- “Configuration Change Notification” (page 39)
- “Core Activation” (page 40)
- “Instant Capacity Cell Board” (page 41)
- “Instant Capacity Software Validation” (page 42)
- “Status Reporting on Instant Capacity Systems” (page 43)
- “Time Zone Considerations” (page 44)

For more in-depth information, see [icap\(5\)](#).

# Instant Capacity Requirements

## Program Requirements

To participate in the Instant Capacity version 9.x program, you must comply with the following conditions of the HP Utility Pricing Solutions program:

- Maintain the HP Instant Capacity software on each HP-UX or OpenVMS *partition* in the system. The Instant Capacity software is a nonintrusive, low-overhead software module that resides on the partition.
- Migrate to later Instant Capacity software versions as they become available.

For specifics about your individual program requirements, see the Utility Pricing Solutions contract from HP or your authorized channel partner.



---

**IMPORTANT:** Participants in the Utility Pricing Solutions program who do not meet these requirements may be in breach of contract. This can result in unnecessary expense for both the program participant and HP.

---

## Software Requirements

Every HP-UX and OpenVMS partition in an Instant Capacity system is required to have the Instant Capacity software installed and the `icapd` daemon running (on HP-UX systems), or the `ICAP_SERVER` process running (on OpenVMS systems). You must maintain the Instant Capacity software until all program requirements are fulfilled and until the system is no longer considered an Instant Capacity system. If you are using GiCAP, all systems must be running Instant Capacity version 9.x.



---

**NOTE:** For necessary HP-UX updates, install the latest Operating Environments Update Release (OEUR) or Application Release (AR) update, if possible. This ensures that you install all required products and versions. Instant Capacity does not work correctly with HP-UX releases prior to December 2004.

---

Your Instant Capacity system is ordered and shipped with all of the following required software and dependent products, with the exception of the *vPars* software. If your system's operating system is reinstalled or installed with Ignite-UX, ensure that the following software requirements are satisfied.

## HP-UX 11i v1 Requirements and Dependent Products

The following software is required for Instant Capacity version 9.x on HP-UX 11i v1:

- HP-UX 11i v1
- iCOD software bundle B9073BA (version 9.x) — On the 11i v1 Applications Software media, or download the most recent version from the HP Software Depot at <http://www.hp.com/go/softwaredepot>.
- The kernel configuration must include the `diag2` module.
- WBEM B8465BA bundle (version A.02.05 or higher)
- nParProvider bundle (version B.12.01.06.02 or higher, available from the OE media)
- If you have a virtual partitioned environment, the Virtual Partitions (vPars) software (bundle T1335AC) must be version A.03.05 or higher. vPars is available separately and is not included with the OE.
- (GiCAP only) The CIM Server configuration property `sslClientVerificationMode` must be set to a value of "optional" on all GiCAP Group Managers and on all OS instances of all member systems. (The CIM Server may need to be restarted if the property was not previously set to this value.) For details, see `cimconfig(1M)`. Note that communication between

the managers and members of groups is established using SSL certificates that are supplied by the GiCAP software.

## Required Patches for HP-UX 11i v1

The following patches are required for Instant Capacity version 9.x on HP-UX 11i v1:

- PHKL\_22987: S700\_800 11.11 pstat() patch (Use only if your system runs MeasureWare software)
- PHKL\_23154: S700\_800 11.11 dflush() patch
- PHKL\_25218: S700\_800 11.11 PDC Call retry, PDC SCSI\_PARMS, iCOD hang fix
- PHKL\_26232: S700\_800 11.11 Psets Enablement patch, FSS iCOD patch
- PHKL\_30197: S700\_800 11.11 Psets & vPar, Reboot hangs, serial number
- PHCO\_24396: S700\_800 11.11 /etc/default/tz patch
- PHCO\_24477: S700\_800 11.11 sar (1M) patch
- PHCO\_29832: S700\_800 11.11 reboot(1M) patch
- PHCO\_29833: S700\_800 11.11 killall(1M) patch



**IMPORTANT:** For the most recent required patches, see the Instant Capacity Installation page on the HP Software Depot (search for “B9073BA”):

<http://www.hp.com/go/softwaredepot>

---

## HP-UX 11i v2 Requirements and Dependent Products

The following software is required for Instant Capacity version 9.x on HP-UX 11i v2:

- HP-UX 11i v2
- iCOD software bundle B9073BA (version 9.x) — On the 11i v2 Applications Software media, or download the most recent version from the HP Software Depot at <http://www.hp.com/go/softwaredepot>.
- The kernel configuration must include the `diag2` module.
- WBEM B8465BA bundle (version A.02.05 or higher)
- nParProvider bundle (version B.23.01.06.02 or higher, available from the OE media)
- If you have a virtual partitioned environment, the Virtual Partitions (vPars) software (bundle T1335BC) must be version A.04.01 or higher. vPars is available separately and is not included with the OE.
- (GiCAP only) The CIM Server configuration property `sslClientVerificationMode` must be set to a value of “optional” on all GiCAP Group Managers and on all OS instances of all member systems. (The CIM Server may need to be restarted if the property was not previously set to this value.) For details, see *cimconfig*(1M). Note that communication between the managers and members of groups is established using SSL certificates that are supplied by the GiCAP software.



**NOTE:** On HP-UX 11i v2 systems, updated firmware might be required by the nPar or Virtual Partition software, as documented for these products.

---

## Required Patches for HP-UX 11i v2

The following patches are required for Instant Capacity version 9.x on HP-UX 11i v2:

- PHKL\_33752: S700\_800 11.23 Itanium 2 Processor speed reporting patch required only for IA systems
- PHCO\_34721: killall(1M) patch
- PHKL\_35174: `_CS_PARTITION_IDENT` patch required for GiCAP on rx8640 and rx7640 systems



---

**IMPORTANT:** For the most recent required patches, see the Instant Capacity Installation page on the HP Software Depot (search for “B9073BA”):

<http://www.hp.com/go/softwaredepot>

---

## HP-UX 11i v3 Requirements and Dependent Products

The following software is required for Instant Capacity version 9.x on HP-UX 11i v3:

- HP-UX 11i v3
- iCOD software bundle B9073BA (version 9.x) — Installed automatically when the HP-UX 11i v3 Operating Environment (OE) is installed, or download the most recent version from the HP Software Depot at <http://www.hp.com/go/softwaredepot>.
- The kernel configuration must include the `diag2` module.
- `WBEMSVCS` bundle (version A.02.05 or higher)
- `nParProvider` bundle (version B.31.01.07.01 or higher, available from the OE media)
- If you have a virtual partitioned environment, the Virtual Partitions (vPars) software (bundle T1335BC) must be version A.05.01 or higher. vPars is available separately and is not included with the OE.
- (GiCAP only) The CIM Server configuration property `sslClientVerificationMode` must be set to a value of “optional” on all GiCAP Group Managers and on all OS instances of all member systems. (The CIM Server may need to be restarted if the property was not previously set to this value.) For details, see `cimconfig(1M)`. Note that communication between the managers and members of groups is established using SSL certificates that are supplied by the GiCAP software.

## OpenVMS Version 8.4 Requirements

The following software is required for Instant Capacity version 9.x on OpenVMS Version 8.4:

- HP OpenVMS Version 8.4 for Integrity Servers or later
- iCAP software bundle BA484AA (version 9.x) — Included with OpenVMS Version 8.4 and automatically installed on relevant systems when Version 8.4 is initially installed
- `WBEMCIM` bundle (version X2.92 or higher) — Optionally installed with OpenVMS Version 8.4
- `nParProvider` bundle — Installed with OpenVMS Version 8.4
- `SSL` bundle (version SSL 1.4 or higher) — Optionally installed with OpenVMS Version 8.4

## Network Requirements

GiCAP and iCAP do not use any network ports directly. However, version 9.0 GiCAP uses `WBEM/SSL` for the communication protocol, requiring use of the CIM HTTPS port 5989. GiCAP also uses `ping` (which uses port 7 for the echo protocol) and host lookup services (`getaddrinfo` and `getnameinfo`) which may use the Domain Name System (DNS), Network Information Service (NIS), or a local host file to provide the lookup function.

## Email Requirements

For some configurations, previous versions of the Instant Capacity software (prior to B.07.x) required email connectivity to HP in order to send asset reports as encrypted email messages. Instant Capacity version B.07.x and later no longer require email connectivity or asset reporting, however, you can choose to configure it because it can be useful for viewing complexwide asset information at the Utility Pricing Solutions portal (<http://www.hp.com/go/icap/portal>).



---

**NOTE:** Starting with iCAP version 8.02, asset reporting is turned off by default for new installations. However, for reinstallation and upgrades the asset reporting remains at its previous setting; if your system was manufactured prior to iCAP version 8.02, asset reporting might still be turned on. Unless you turn off asset reporting or configure the email connectivity, error messages are logged when the software attempts to send asset reports. The `icapstatus` command displays the current configuration for asset reporting. You turn asset reporting on or off with the `icapnotify -a` command. For details about how to configure email connectivity, see “Configuring Email on Instant Capacity Systems” (page 202).

---

## System Contact Requirement

Your organization can ensure the successful management of Instant Capacity systems by designating a system contact. The system contact receives the following types of email notifications from the Instant Capacity software:

- Core activation or deactivation
- Compliance exception
- Temporary capacity expiration
- Instant Capacity enforcement
- Virtual partition boot time compliance
- Hardware incompatibility in a GiCAP group (see “Upgrades and Global Instant Capacity” (page 121))

If a system contact is not specified, notification of core changes are not sent from the Instant Capacity partition. However, the `root` account on HP-UX and the `SYSTEM` account on OpenVMS continue to receive other notification messages.

## Usage Rights Requirement

A system managed under the Instant Capacity program can include one or more components (core, cell, or memory) that are without usage rights. Before you can use these components, you must obtain additional usage rights. Usage rights can either be purchased from HP, or, if the system is a member of a GiCAP group, borrowed temporarily from another member of the group, as described in the section “Global Instant Capacity Sharing Rights” (page 108).

Purchase of usage rights from HP is managed through the use of RTU codewords. Contact your HP sales representative in order to purchase component-specific usage rights. After such a purchase, HP sends you a letter informing you how to retrieve the RTU codeword from the HP Utility Pricing Solutions website:

<http://www.hp.com/go/icap/portal>

After the RTU codeword is retrieved from the Utility Pricing Solutions portal, the RTU codeword is applied to your server by the use of the `icapmodify -C` command. When the codeword is applied, component-specific usage rights on the system are increased, allowing you to activate one or more additional components.

## Instant Capacity Integration with Virtual Partitions (HP-UX)

The minimum required versions of vPars software for HP-UX systems are as follows:

- HP-UX 11i v1: vPars version A.03.05
- HP-UX 11i v2: vPars version A.04.01
- HP-UX 11i v3: vPars version A.05.01

Each of these versions provides a *virtual partition* environment which is tightly integrated with Instant Capacity, making it less likely for a complex to be misconfigured or to violate contractual compliance.

The Instant Capacity software must be installed on all virtual partitions in an Instant Capacity system.

For details about virtual partitions, see the *Installing and Managing HP-UX Virtual Partitions* manual on the HP Documentation website:

<http://docs.hp.com>

# Instant Capacity Components

## Overview

The Instant Capacity software monitors and enforces compliance with contractual agreements. It authorizes or denies activation of system components (cores, cells, memory) based on a complexwide database of usage rights. For details about acquiring additional usage rights, see “Usage Rights Requirement” (page 31).

Activation of components is restricted according to complexwide compliance for each component type. A complex is in a compliant state when the number of active components of a given type does not exceed the number of that component’s available usage rights on the complex.

## Processors and Cores

Although you purchase Instant Capacity processors for your system, the Instant Capacity software monitors and manages the total number of cores. For example, if you have a dual-core Instant Capacity processor, two cores must remain inactive on the complex.

The Instant Capacity software enforces compliance for cores by comparing the actual number of *inactive cores* with the expected number of inactive cores (the number of cores without usage rights) for the entire complex, according to the contract with HP. Available core usage rights can be used to activate any core in an active cell board. Note also that temporary capacity can be used to activate cores beyond the number of available core usage rights for the complex, but only for a limited period of time.



---

**NOTE:** Unless a system participates in a GiCAP group (see [Chapter 7](#)), usage rights are complexwide (single node for OpenVMS) only. If components are moved from one complex to another, the counts of allowable active and inactive components do not change for either complex. In particular, the number of “expected inactive” components of each type does not change if components are removed. This means that the removal of *inactive* components from a complex can cause that complex to be out of compliance with the Instant Capacity contract because there are fewer visible inactive components than the complex-wide count of components without usage rights. The complex might even become unusable if, for example, enough other cores must be made inactive to meet compliance resulting in insufficient active cores to have at least one active core per configured cell.

---

## Cell Boards and Memory

Instant Capacity offers you a way to have additional (inactive) cell-board capacity in your system for growing business needs. When the need arises, these cell boards, which contain memory and cores, are available for instant activation and use after reboot or *cell online activation* when additional cell-board usage rights are purchased from HP and an RTU codeword is applied. As with cores, the Instant Capacity software enforces compliance by comparing the number of actual *inactive cells* with the expected number of inactive cells (the number of cells without usage rights for the entire complex).

The cell-board, memory, and core usage rights are tracked separately. To activate an Instant Capacity cell, you must acquire sufficient cell usage rights, as well as sufficient memory usage rights to enable all the memory attached to the cell. You cannot activate a cell board without activating all attached memory, so when you purchase an RTU for a cell you need to purchase an RTU for the cell’s memory. These are normally bundled together in a single purchase.

Depending on the need, you might want to activate one or more cores at the same time the cell and memory are activated, so you might also need to acquire additional core usage rights. After a cell board is activated, all of the cores on the cell board are available for activation if the complex has enough available core usage rights or temporary capacity. Since usage rights for all types of

components can be conveyed with a single RTU codeword, it is particularly useful to anticipate the core and memory needs when purchasing cell-board usage rights.

As with other components, the Instant Capacity software enforces compliance for memory by comparing the actual amount of inactive memory with the expected inactive memory.



**IMPORTANT:** You must have one active core for each active cell board.

---

# Global Instant Capacity

Global Instant Capacity, or GiCAP, provides HP customers with the flexibility to move usage rights for Instant Capacity components within a group of servers. It also provides “pooled” temporary capacity across the group. This provides more cost-effective *high availability*, including disaster recovery capabilities, more adaptable *load balancing*, and more efficient and easier use of temporary capacity.

Global Instant Capacity is built on the concept of a server group, or GiCAP group. A GiCAP group is a list of servers that are allowed to share Instant Capacity usage rights. In addition, at least one HP-UX system running Instant Capacity must be designated as the GiCAP Group Manager. The Group Manager hosts the GiCAP software that maintains a database of information about the group and about group resources (usage rights and temporary capacity). A GiCAP group is managed using the command `icapmanage` on the Group Manager system.

While GiCAP is part of Instant Capacity and is installed at the same time as Instant Capacity, it is not enabled during installation. To create groups and share resources across groups, you must purchase GiCAP *sharing rights*, acquire the GiCAP codeword from the HP Utility Pricing Solutions portal (<http://www.hp.com/go/icap/portal>), and apply the associated codeword to the Group Manager system. Application of the sharing rights codeword to the Group Manager system enables the addition of members with Instant Capacity components to groups. In addition, you must acquire grouping rules from the portal and apply those rules to the Group Manager system. GiCAP sharing rights and grouping rules are described in [Chapter 7 \(page 101\)](#). All GiCAP Group Manager systems and group members must run Instant Capacity version 9.0 or later.

Instant Capacity allows deactivation of cores on non-Instant Capacity systems (those without any Instant Capacity components), allowing such systems to participate in a GiCAP group and loan usage rights to Instant Capacity systems. A GiCAP group consists of servers that are allowed to be grouped together according to a set of grouping rules defined by HP. These grouping rules must be acquired from the Utility Pricing Solutions portal and applied to the Group Manager system before groups can be created. For more information, see “[Global Instant Capacity Grouping Rules](#)” (page 107).

## Instant Capacity Codewords

Instant Capacity uses codewords for several purposes: to adjust available usage rights for system components (RTU codewords), to apply temporary capacity to the system, and to apply sharing rights to a GiCAP Group Manager system to enable the creation of one or more groups.

All types of codewords must be purchased as specific product numbers from HP. After purchase, the actual codeword (an encrypted string) can be retrieved from the Utility Pricing Solutions web portal. When you retrieve codewords from the portal, you must provide the sales order number for the codeword purchase and the serial number of the system. For more information about purchasing codewords, see the Utility Pricing Solutions web portal information in “System Overview” (page 20).

Once obtained from the portal, the rules and uses of GiCAP codewords are very different from other types of Instant Capacity codewords. The GiCAP codewords are described in “Global Instant Capacity Sharing Rights” (page 108) and are referred to as “GiCAP codewords”. Other types of codewords are referred to as “Instant Capacity codewords”.

The following Instant Capacity codewords are applied to a server using the `icapmodify -C` command:

- Core RTU
- Cell board RTU
- Memory RTU
- Temporary capacity

For details about using temporary capacity codewords, see “Acquiring and Configuring Temporary Instant Capacity” (page 79).

Application of an RTU codeword adjusts the number of component-specific usage rights on the system. As a result, more components can be active simultaneously.



**IMPORTANT:** Instant Capacity codewords are based on both the serial number of a system and a unique sequencing value for that server. These codewords must be applied in the sequence in which they are obtained for a particular server. They can be applied to any partition on the server.

---

## Temporary Instant Capacity

You can purchase an amount of temporary processing capacity for your Instant Capacity system. Temporary Instant Capacity, or TiCAP, is purchased in units of processing days. TiCAP allows one or more cores beyond the count allowed by the available usage rights to be activated for up to the specified period of prepaid minutes without requiring the purchase of additional usage rights.

You can activate and deactivate cores according to your needs until the activation time equals your prepaid temporary capacity duration. For example, with a prepaid duration of 30 days of temporary capacity, you can activate one core for 30 days, or four cores for one hour a day for 180 days (or any combination that totals 43,200 minutes).



**NOTE:** Temporary Instant Capacity cannot be used to activate inactive Instant Capacity cell boards or memory. To activate cores on an inactive cell board, you must first activate the cell board and its memory either by using newly purchased cell and memory usage rights or by deactivating a cell board and memory elsewhere in the server.

---

Your temporary capacity balance is decreased only when you are using more cores than normally allowed by your available core usage rights. The charge against temporary capacity is not associated with specific cores or partitions. That is, if you use temporary capacity to activate one core in partition A, and then you deactivate any core in partition B, the complex stops using temporary capacity.

The Instant Capacity software uses the debiting of temporary capacity to track the noncompliance of a system, as described in “Instant Capacity Compliance and Enforcement” (page 38).

TiCAP can be added to an Instant Capacity system by purchasing and applying a temporary capacity codeword (available from the Utility Pricing Solutions portal) using the `icapmodify -C` command.

The `icapstatus` command provides information about the amount of temporary capacity time remaining on the complex.

For more information about temporary capacity, see Chapter 5: “Temporary Instant Capacity” (page 75). For a discussion of temporary capacity in a GiCAP group, see Chapter 7.

# Instant Capacity Compliance and Enforcement

The Instant Capacity software primarily maintains complexwide information about the usage rights and activation of system components. The software monitors the number of usage rights for the entire complex for each type of component. (If you are using Global Instant Capacity, the software also maintains groupwide information about usage rights. For more information about GiCAP, see Chapter 7.)

The Instant Capacity software uses the debiting of temporary capacity as a compliance enforcement mechanism on the following systems:

- A system where TiCAP has been applied (the TiCAP balance is nonzero)
- A GiCAP group member system (or a system which had previously been a GiCAP group member and has a nonzero TiCAP balance)
- A system which has had a TiCAP debit within the past 24 hours

If temporary capacity is negative, it means the system has gone out of compliance. Also, if the temporary capacity balance is negative and the number of active cores exceeds the number of core usage rights for the complex, automatic deactivation of cores might occur at boot time.

The Instant Capacity software authorizes activation of cores, cells, and memory if allowed by the system usage rights. You cannot activate additional components if that action takes the system out of compliance.

For example, if your Instant Capacity contract specifies that your server must contain 12 cores with usage rights and 4 cores without usage rights, you can have up to 12 cores activated at any one time, and 4 cores must be inactive at all times.

The Instant Capacity software can determine the following compliance aspects:

- Whether a system is in compliance or out of compliance with the Instant Capacity contract
- The number of additional cores that can be activated
- The number of additional cells and the quantity of memory that can be activated

The enforcement methods used by the software include:

- Not allowing activations that cause the system to be out of compliance
- Deactivating cores at boot time:
  - Automatically deactivate cores at boot time if temporary capacity is exceeded and the number of active cores continues to exceed the number of core usage rights for the complex (see “Temporary Instant Capacity Expiration and Compliance Enforcement” (page 86)).
  - Prevents a virtual partition from booting if the number of assigned cores is greater than the number of *intended active* cores for the nPartition (see “Boot Time Compliance” (page 66)).
- For GiCAP members in a running (booted) virtual partition environment, if the number of assigned cores is greater than the number of intended active cores, the `icapmodify` command may be disallowed. In this case, to bring the vPar database into compliance, deactivate cores using the `vparmodify` command.
- On OpenVMS systems, the `ICAP_SERVER` process dynamically deactivates active cores that exceed the number of core usage rights for the complex.

# Configuration Change Notification

Specifying an increase or decrease in the number of active cores (using the `icapmodify` command) causes a core configuration change. An email notification is sent to the system contact when a change occurs that affects the configuration of cores.

If you do not want an email notification to be sent when configuration changes are made, disable this feature by using the following command on HP-UX systems:

```
/usr/sbin/icapnotify -n off
```

Example 2-1 shows a configuration change email notification the Instant Capacity software sends to the system contact. Note that if the operation is a deferred configuration change, “previous” and “current” show equal values; only the value for “Number of cores to be active after reboot” reflects the requested change.

## **Example 2-1 Configuration Change Notification email for Instant Capacity System (not vPar)**

---

Subject: Instant Capacity Configuration Change Notification

A configuration change has been made to the following system:  
super.corp.com

One or more cores were activated.

Details of the change include:

Time of change:	09/08/08 11:00:00
Deferred change:	No
Previous number of active cores:	3
Current number of active cores:	4
Number of cores to be active after reboot:	4

Description of change: New fiscal year increase  
Person making change: Mary Jones  
System contact email: mjones@corp.com

If you are the system contact and do not want to receive this type of notification in the future, it can be disabled by executing the following command on the system in question:

```
/usr/sbin/icapnotify -n off
```

To turn notification on, execute:

```
/usr/sbin/icapnotify -n on
```

---

## Core Activation

As previously described, an Instant Capacity system contains a specified quantity of activated processing capacity (cells, cores, and memory) and a specified amount of deactivated processing capacity. Systems can have fewer active components than they have rights to activate. Such systems can instantly activate additional components without the need to purchase an RTU, up to the number of component usage rights on the system.

### Increasing Processing Capacity by Purchasing RTUs

When processing demands change significantly, you can enable use of additional system components using the following procedure:

1. Purchase additional usage rights for a component type, by sending a purchase order to HP for an Instant Capacity RTU product. Soon after HP receives your purchase order, you receive a letter from HP that contains information about retrieving the RTU codeword from the Utility Pricing Solutions web portal.
2. Acquire the RTU codeword from the Utility Pricing Solutions web portal:  
<http://www.hp.com/go/icap/portal>
3. Apply the RTU codeword by entering the `icapmodify -C` command (note the `-C` option is uppercase) on any partition in the complex.
4. Activate a component. Depending on the type of component, do one of the following:
  - Activate a core in a hard partition (*nPartition*) by entering the `icapmodify -a` command. Note: For details about activating a core in a virtual partition, see “Instant Capacity Integration with Virtual Partitions (HP-UX)” (page 31).
  - Activate a cell board by using the `parmodify` or `parmgr` command. For details about activation of cell boards (and memory), see “Activation of an Instant Capacity Cell Board” (page 97).



**IMPORTANT:** To avoid a delay in activation of additional cores, purchase and keep in reserve some quantity of temporary capacity for the system. Temporary capacity can be used to instantly and temporarily activate cores while waiting for an RTU codeword. For details about temporary capacity, see Chapter 5. You can also temporarily activate one or more cores using the Instant Access Capacity (IAC) provided with the purchase of Instant Capacity processors.

---

# Instant Capacity Cell Board

Instant Capacity Cell Board allows you to have additional (inactive) cell-board capacity in your system for growing business needs. When the need arises, you can purchase additional cell and memory usage rights; then the inactive cell boards, which contain memory and cores, are available for instant activation and use.

The Instant Capacity software monitors and enforces the count of inactive cell boards (those without usage rights) throughout the complex. Inactive cell boards assigned to a partition have the `use-on-next-boot` flag set to `n` (no).

After applying a cell board and memory codeword to convey additional usage rights, you may activate any cell board with memory not exceeding the amount allowed by the available memory usage rights, and configure the cell into an nPartition. This is controlled by setting the `use-on-next-boot` flag to `y` (yes) with the `parmodify` command after the cell is configured and assigned to the partition, and then rebooting the nPartition.



**IMPORTANT:** An *active cell* board must have a minimum of one *active core*; therefore, you must also have sufficient core usage rights on the complex to activate at least the one core per cell board. The Instant Capacity software redistributes active cores across all cell boards of the partition. Temporary Instant Capacity can be used to activate this core but the system needs to have at least one core usage right for each active cell. Also, if you are creating a new partition for the new cell board, additional constraints apply and you might need to acquire additional core usage rights. For more information see “New Partition Creation and Instant Capacity” (page 196).

---

For more information about Instant Capacity Cell Board, see Chapter 6: “Instant Capacity Cell Board” (page 91).

# Instant Capacity Software Validation

## On HP-UX Systems

The Instant Capacity software (HP-UX product B9073BA) is installed by HP manufacturing on instantly ignited HP-UX systems. The Instant Capacity software can also be installed by an HP service representative on existing (supported) HP-UX enterprise servers as an *add-on*.



---

**NOTE:** The Instant Capacity software is automatically installed when the HP-UX 11i v3, 11i v2, or 11i v1 Operating Environment (OE) is installed.

---

To verify the Instant Capacity software is installed on your system, execute the following HP-UX command:

```
/usr/sbin/swlist -l product iCOD
```

Output similar to the following is displayed:

```
iCOD B.11.23.09.00 HP-UX iCOD Instant Capacity
```

To verify that the Instant Capacity software installation has not been corrupted, execute the following HP-UX command:

```
/usr/sbin/swverify iCOD
```

The command will display the message `Verification succeeded`.

## On OpenVMS Systems

On OpenVMS systems, the Instant Capacity software is automatically installed on partitionable systems when the OpenVMS Version 8.3 or later operating system is installed. You should not need to install the Instant Capacity software separately on OpenVMS systems. Onsite Instant Capacity installation by an HP service representative after the initial installation of OpenVMS Version 8.3 is not an option for OpenVMS systems.

To verify that the Instant Capacity software is installed and configured, run the following commands:

```
$ @sys$manager:ICAP$CLI_UTILS.COM CONFIG_CHECK
$ show log ICAP$CONFIGURED
"ICAP$CONFIGURED" = "TRUE" (LNM$JOB_nnnnnnnn)
```

## Status Reporting on Instant Capacity Systems

You can use the `icapstatus` command to view the status of your Instant Capacity system. With no options, the `icapstatus` command provides the following information:

- Version number of the Instant Capacity software
- System identification information (system ID, serial number, current product number, unique ID)
- Email addresses for both the system contact (mail to the local system) and a “From” address (mail from the local system)
- Asset reporting status (on or off)
- Temporary capacity warning period (in days)
- Exception status (indicates whether complex is in an exception state)
- If part of a GiCAP group, group membership information, including borrow/loan status
- Local partition information
- Instant Capacity resource information for the entire complex
- Allocation of Instant Capacity resources among the hard partitions

For information about what is reported by the `icapstatus` command, see “Checking the Status of your Instant Capacity System” (page 52). For details about the `icapstatus` command, see `icapstatus(1M)`.

## Time Zone Considerations

On HP-UX systems, the `icapd` daemon performs routine Instant Capacity software tasks on a daily basis. A partition's local time zone setting affects what time zone the `icapd` daemon uses for the timing of these tasks. Be sure that the time zone is set properly to ensure synchronization among the partitions.

Because the HP-UX `icapd` daemon is started by `init`, the `/etc/default/tz` file must contain the desired time zone specification. By default, the time zone is set to `EST5EDT`. You can specify the time zone the `icapd` daemon uses to interpret noon and midnight by modifying the entry in the `/etc/default/tz` file.

On OpenVMS systems, the `ICAP_SERVER` uses the time zone settings defined by the `SYS$STARTUP:TDF$UTC_STARTUP.COM` file. To view the time zone settings, enter the command `@sys$manager:utc$time_setup "show"`. Enter the command `@sys$manager:utc$time_setup` and follow the menu instructions to modify the time zone setting for the iCAP partition.

---

## 3 Installing and Removing Instant Capacity Software

This chapter covers the following topics:

- “Installing Instant Capacity Software” (page 46)
- “Reinstalling Instant Capacity Software” (page 48)
- “Removing Instant Capacity Software” (page 49)

# Installing Instant Capacity Software

## Factory Integrated Systems

The Instant Capacity software is installed by HP on all HP enterprise servers, even those without Instant Capacity components. To verify that the software is installed and configured, use the following HP-UX command :

```
/usr/sbin/swverify iCOD
```

The command output displays the message `Verification succeeded.`



**IMPORTANT:** The Instant Capacity software is automatically installed when the HP-UX 11i v1, 11i v2, or 11i v3 Operating Environment (OE) is installed. If any partition in the system has version B.06.x or later installed, then all partitions in the system with Instant Capacity software must be running version B.06.x or later. For details about upgrading the software from a version previous to B.06.x, see “Upgrading to Instant Capacity version B.06.x or later (HP-UX)” (page 193).

For HP-UX you generally do not need to install the Instant Capacity software. However, if you do, it is available from the following media/location:

- HP software depot at <http://www.hp.com/go/softwaredepot>
- September 2008 HP-UX 11i v3 Operating Environments (OE) media
- September 2008 HP-UX 11i v2 Applications Software media
- September 2008 HP-UX 11i v1 Applications Software media

If the system is a Group Manager or a member of a GiCAP group, prior to installing the iCAP version 9.0 software, ensure that **WBEM** version A.02.05 or later is installed, set the CIM Server configuration property `sslClientVerificationMode` to “optional”, and restart the `cimserver` on the Group Manager, the standby Group Manager (if any), and on all HP-UX OS instances of all member systems. Because this configuration attribute is not dynamic, you must specify the `-p` option, and you must restart the `cimserver` to set the value. For details, see `cimconfig(1M)`.

Use the following commands to set `sslClientVerificationMode` to “optional” and restart the `cimserver`:

```
# cimconfig -s sslClientVerificationMode=optional -p
# cimserver -s; cimserver
```

Use the following command to verify the value of `sslClientVerificationMode` after a restart:

```
# cimconfig -g sslClientVerificationMode -c -p
Current value: optional
Planned value: optional
```

## Installing from the HP-UX Media (HP-UX 11i v1, 11i v2, and 11i v3):

Follow this procedure to install the Instant Capacity software on your HP-UX system from either the appropriate HP-UX Applications Software or OE media:

1. Log in as `root`.
2. Determine the DVD drive device file by entering the following command:  

```
ioscan -fnC disk
```
3. Insert the appropriate HP-UX Applications Software or OE DVD into the DVD drive.
4. Mount the DVD drive to the desired directory. The following example uses the `/dev/dsk/c1t2d0` device file (from step 2) and the `/dvd` directory. To mount the DVD drive, enter a similar command as:  

```
mount -r /dev/dsk/c1t2d0 /dvd
```
5. Install the B.09.x bundle B9073BA from the DVD:

```
swinstall -s /dvd B9073BA
```

6. Continue with “For All HP-UX Installations” (page 47).

## Installing from the HP Software Depot (For HP-UX 11i v1, 11i v2, and 11i v3):

1. Search for B9073BA on the HP Software Depot website: <http://www.hp.com/go/softwaredepot>
2. Select the link that appeared as a result of your search, and follow the instructions on the Installation page.
3. Continue with “For All HP-UX Installations” (page 47).

## For All HP-UX Installations

After you have successfully installed the Instant Capacity software using the `swinstall` command, perform the following procedure to validate your installation:

1. Execute the `/usr/sbin/icapstatus` command.
2. Verify that the `icapstatus` command output indicates the correct number of components without usage rights for cells, cores, and memory.

If any number is incorrect, contact your local HP Response Center and request iCAP assistance.

3. Log in as `root`
4. Set the system contact information by entering the following command:  

```
/usr/sbin/icapmodify -c <contact_email_address>
```
5. If you want to configure asset reporting, then ensure that outgoing mail can be sent to HP mail servers from your system, even if the system is behind a firewall. See “Diagnosing Email Configuration” (page 142).
  - Test the transmission of your asset report, via email to HP, by entering the following command:  

```
/usr/sbin/icapnotify <reply_address>
```

The `icapnotify` command sends an asset report to HP, `root`, and the supplied reply address.

HP responds with an email to the reply address after the asset report is received.
  - Use an email client to verify the return email from HP to the reply email address.

## Installing Instant Capacity on OpenVMS Systems

On OpenVMS systems, the Instant Capacity software is automatically installed on partitionable systems when the OpenVMS Version 8.3 or later operating system is installed. You should not need to install the Instant Capacity software separately on OpenVMS systems. Instant Capacity hardware components have already been configured at the factory before delivery. The Instant Capacity software is included on the OpenVMS Version 8.3 Operating Systems media.

To verify that the Instant Capacity software is installed and configured, run the following OpenVMS commands:

```
$ @sys$manager:ICAP$CLI_UTILS.COM CONFIG_CHECK
$ show log ICAP$CONFIGURED
"ICAP$CONFIGURED" = "TRUE" (LNM$JOB_nnnnnnn)
```

# Reinstalling Instant Capacity Software

## Preserving current Instant Capacity information

If you reinstall HP-UX on a partition with Instant Capacity (for example, installing HP-UX by either cold-installing or installing from a golden image), you must perform the following steps. Otherwise, all information in the Instant Capacity configuration file (change history, system contact information) is lost.



---

**IMPORTANT:** If the system is a member of a GiCAP group, it should be removed from the group before reinstalling HP-UX, and added back into the group after the reinstallation. “Reinstalling a Group Member” (page 118) describes a workaround if HP-UX is reinstalled while a system is a member of a GiCAP group. If the system is a GiCAP Group Manager, additional steps described in Chapter 7 are necessary to configure the Group Manager after a reinstallation of HP-UX.

---

1. Before the reinstall action, manually save your Instant Capacity data and processor allocation history by backing up the following files:
  - a. HP-UX: `/etc/.iCOD_data`  
OpenVMS: `SYS$SYSTEM:SYS$IICAP.DAT`
  - b. HP-UX: `/var/adm/icap.log`  
OpenVMS: `SYS$MANAGER:ICAP.LOG`
  - c. HP-UX: `/var/adm/icap.log.old`
  - d. OpenVMS: `SYS$MANAGER:ICAP$CORES.DAT`These files are restored in step 3.
2. Install the appropriate HP-UX or OpenVMS OE from its media onto the partition. The Instant Capacity software bundle B9073BA is installed automatically when the HP-UX OE is installed, and the Instant Capacity software bundle BA484AA is installed automatically on OpenVMS systems.
3. Restore your Instant Capacity data and processor allocation history files:
  - a. HP-UX: `/etc/.iCOD_data`  
OpenVMS: `SYS$SYSTEM:SYS$IICAP.DAT`
  - b. HP-UX: `/var/adm/icap.log`  
OpenVMS: `SYS$MANAGER:ICAP.LOG`
  - c. HP-UX: `/var/adm/icap.log.old`
  - d. OpenVMS: `SYS$MANAGER:ICAP$CORES.DAT`

## Removing Instant Capacity Software



---

**IMPORTANT:** Do not attempt to remove the Instant Capacity software.

---



---

## 4 Using Instant Capacity to Manage Processing Capacity

This chapter covers the following topics:

- “Checking the Status of your Instant Capacity System” (page 52)
- “Setting System Contact Information” (page 54)
- “Applying a Right To Use (RTU) Codeword” (page 55)
- “Activating Cores” (page 57)
- “Deactivating Cores” (page 59)
- “Overriding Deferred Activation and Deactivation” (page 61)
- “Load-Balancing Active Cores” (page 62)
- “Understanding and Managing Intended Active Values” (page 63)
- “Activations and Deactivations in a Virtual Partition Environment” (page 64)
- “Assigning a Cell to a Partition” (page 67)
- “Unassigning a Cell from a Partition” (page 69)
- “Software Application Considerations” (page 71)
- “Test Activation of Cores Using Temporary Capacity” (page 72)
- “Replacement of Failed Cores” (page 73)

# Checking the Status of your Instant Capacity System

You can use the `icapstatus` command to view the status of your Instant Capacity system. The `icapstatus` command, issued without options, provides the following information:

- Version number of the Instant Capacity software
- System identification information (system ID, serial number, product number, unique ID)
- System contact email address
- Instant Capacity From: email address
- Asset reporting status (on or off)
- Temporary capacity warning period (in days)
- Exception status (indicates if complex is in an exception state)
- If a member of a GiCAP group, membership information and borrow/loan status of usage rights
- Local virtual partition status (if applicable):
  - Total number of assigned cores
  - Number of active assigned cores
  - Number of inactive assigned cores
  - Additional cores that can be assigned with current usage rights
  - Number of cores that could be assigned with additional usage rights
  - Number of cores that can be assigned with temporary capacity
  - Number of cores currently unavailable for assignment
- Local nPartition status (if not a virtual partition):
  - Date and time the command was issued
  - Total number of configured cores
  - Number of intended active cores
  - Number of active cores
  - Number of inactive cores
  - Additional cores that can be activated with current usage rights
  - Number of cores that could be activated with additional usage rights
  - Number of cores that can be activated with temporary capacity
  - Number of cores that are deconfigured or attached to inactive cells
- Instant Capacity resource summary:
  - Number of cells without usage rights
  - Number of inactive cells
  - Amount of memory without usage rights
  - Amount of inactive memory
  - Number of cores without usage rights
  - Number of inactive cores
  - Number of cores using temporary capacity
  - Number of cores that must be deactivated (insufficient usage rights)
  - Temporary capacity available
  - Projected temporary capacity expiration
- Allocation of Instant Capacity resources among hard partitions:
  - nPar ID
  - Total cores
  - Intended active cores
  - Actual active cores
  - Inactive cores

- Inactive memory
- Inactive cells
- Runs iCAP (indicates whether the hard partition contains compatible Instant Capacity software)
- nPar name

For details about the `icapstatus` command and its output, see `icapstatus(1M)`.

### Example 4-1 Sample Session of `icapstatus` (on HP-UX)

```
> /usr/sbin/icapstatus

Software version:          B.11.31.09.00.00.71
System ID:                 supericod
Serial number:             1234567890
Product number:           A6912A
Unique ID:                 fffff-fff-ffffff-ffff
System contact e-mail:    mjones@corp.com
From e-mail:               Set to the default ('adm')
Asset reporting:          on
Temporary capacity warning period: 15 days
Exception status:         No exception

Local nPartition Status (09/17/08 12:34:56)
-----
Total number of configured cores:          8
Number of Intended Active cores:          7
Number of active cores:                   6
Number of inactive cores:                 2
Additional cores that can be activated with current usage rights: 1
Number of cores that could be activated with additional usage rights: 1
Number of cores that can be activated with temporary capacity: 0
Number of cores that are deconfigured or attached to inactive cells: 0

Instant Capacity Resource Summary
-----
Number of cells without usage rights:      0
Number of inactive cells:                 0
Amount of memory without usage rights:    0.0 GB
Amount of inactive memory:                0.0 GB
Number of cores without usage rights:      4
Number of inactive cores:                 6
Number of cores that must be deactivated (insufficient usage rights): 0
Temporary capacity available:              0 days, 0 hours, 0 minutes

Allocation of Instant Capacity Resources among the nPartitions
-----
              Intended Actual
nPar Total Active  Active =====Inactive===== Runs
ID  Cores Cores   Cores Cores  Memory   Cells iCAP nPar Name
==== =====
0    8     5     4    4    0.0 GB    0  Yes  Partition 0
1    8     7     6    2    0.0 GB    0  Yes  Partition 1 (local)
N/A  0    N/A    N/A  N/A    0.0 GB    0  N/A  Unassigned Cells
```

## Setting System Contact Information

HP recommends that you specify a system contact's email address on each partition in your system. On OpenVMS systems, the email address may be a logical pointing to a distribution list.

If specified, the system contact receives the following types of Instant Capacity email:

- Configuration change notification when cores are activated or deactivated
- Compliance exception notification
- Temporary capacity expiration notification
- Instant Capacity enforcement notification
- Virtual partition boot time compliance notification
- If participating in a GiCAP group and a hardware incompatibility is detected (see "Upgrades and Global Instant Capacity" (page 121))



**NOTE:** Instant Capacity email messages are sent to the system contact email address, if specified, and the root account on the partition. Most notifications are also written to the system log.

---

To specify the Instant Capacity system contact's email address, use the `icapmodify -c` command. Note that you must specify a valid Internet email address. For example:

```
> /usr/sbin/icapmodify -c mjones@corp.com  
The contact email address has been set to mjones@corp.com.
```



**NOTE:** The email address specified for the system contact can be an email alias if you want multiple recipients to receive Instant Capacity email messages.

---

## Applying a Right To Use (RTU) Codeword

Unless you have a balance of Instant Access Capacity (IAC) or temporary capacity (or usage rights available from a GiCAP group), you must purchase additional usage rights before activation of an inactive core. To purchase additional usage rights:

1. Contact your HP sales representative to purchase the appropriate Instant Capacity RTU products.
2. Acquire an RTU codeword from the Utility Pricing Solutions web portal at <http://www.hp.com/go/icap/portal>.
3. Apply the RTU codeword using the `icapmodify -C` command. Note that the `-C` option is in uppercase.

Example Example 4-2 shows how to apply a core RTU codeword using the `icapmodify -C` command.

### Example 4-2 Applying an RTU Codeword (HP-UX)

---

```
> /usr/sbin/icapmodify -C mDwvj0M.fbHhKC9.byrgc8k.Pc7PjMt-1cp63H9.29xrDLU.g2CJhQM.RzuYyd6
```

The following valid codeword has been applied to the complex:

```
Right to Use Codeword
1 Core(s)
```

Use `icapstatus(1M)` to see the results of the application of this codeword.

NOTE: Application of Right to Use codewords does not result in the activation of components. Use `icapstatus(1M)` to see the results of the application of this codeword.

For cores: Use `icapmodify(1M)` to activate the cores.

For cell boards and memory: Use `parmodify(1M)` to activate cells by setting the `use_on_next_boot` flag to 'y' or use `parmgr(1M)`.

---

Application of the codeword increments the count of Additional cores that can be activated with current usage rights (limited by the Number of inactive cores) in the output of the `icapstatus` command.

---

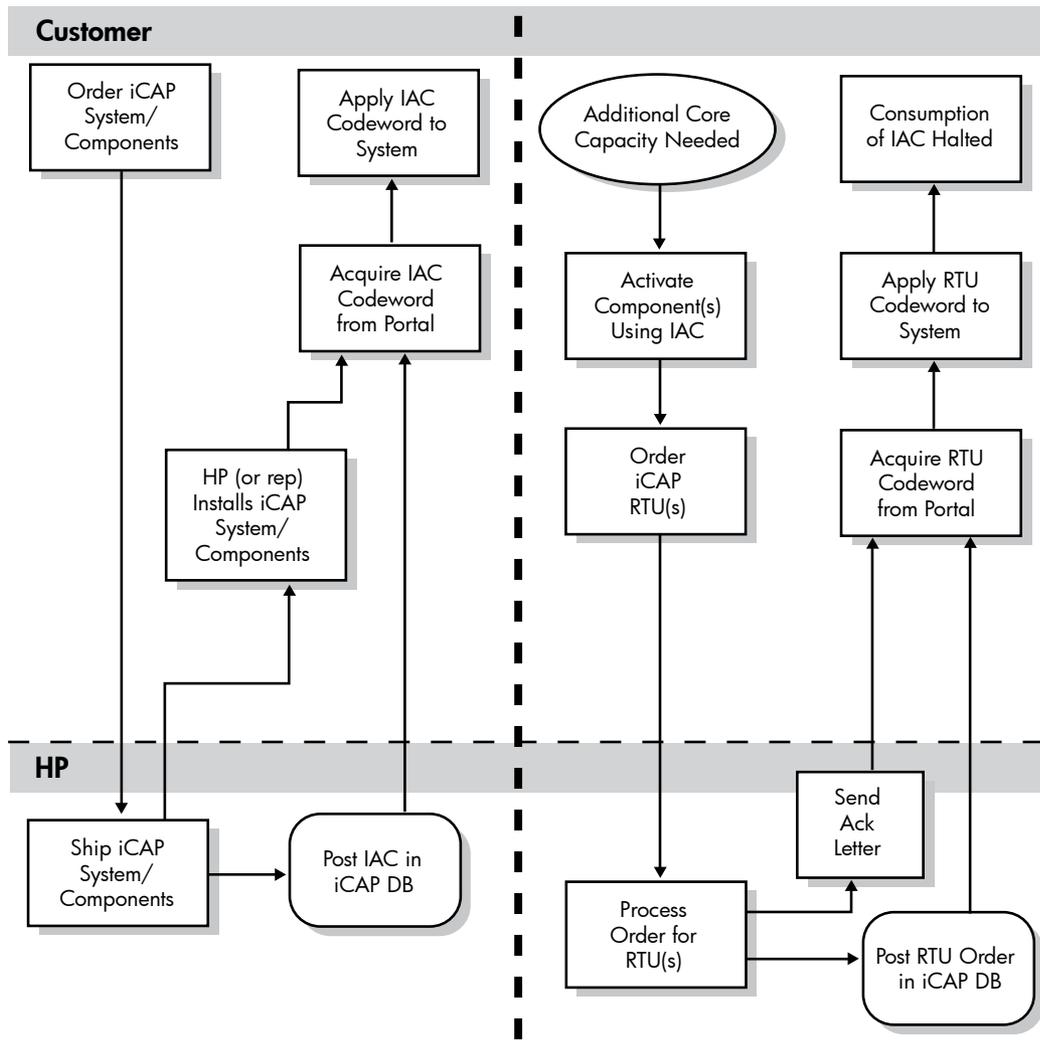


**IMPORTANT:** RTU codewords are based on both the serial number of a system and a unique sequencing value for that server. These codewords must be applied in the sequence in which they are obtained for a particular server. They can be applied to any partition on the server.

---

Figure 4-1 (page 56) illustrates the process of ordering Instant Capacity components and ordering and applying usage rights.

**Figure 4-1 Permanent Activation of Instant Capacity Components**



## Activating Cores

The `icapmodify` command provides the ability to increase processing capacity instantly by activating cores with available usage rights in Instant Capacity systems. At any time, any number of inactive cores with usage rights can be activated, as long as sufficient usage rights are available.

Whether you are activating or deactivating cores, the `icapmodify` command adjusts only the number of dynamic cores, and it does not explicitly identify specific cores.

### Activating Cores in nPartitions

The Instant Capacity software provides two types of activation:

- *Instant* (`icapmodify` command's default behavior) — activation occurs immediately.
- *Deferred* (`icapmodify -D`) — activation occurs after the next reboot of the partition.

Instant activation of cores occurs when the `icapmodify` command is used with either the `-a` or `-s` option, and the `-D` option is not specified. If the `-s` option is used, the number of cores specified must be greater than the number of currently active cores.

Deferred activation of cores occurs when the `icapmodify` command is used with both the `-D` option and either the `-a` or `-s` option. With the deferred option (`-D`), core activation occurs after a reboot of the partition. The reboot (and the core activation) can take place at a planned time. For example, if you activate cores in deferred activation mode and schedule a partition reboot to occur on the first day of the next month, the cores are activated at that time.



**IMPORTANT:** If you shut down a partition for 12 hours or more, it should be powered off or deactivated to avoid additional charges. To power off the partition, execute the `PE` command from the system MP.

On HP-UX systems, always use the `shutdown -R -H` command when shutting down or rebooting an Instant Capacity partition. If the partition is already shut down, use the `rr` command from the system MP to reset cells for reconfiguring. For information about the `shutdown` command, see `shutdown(1M)`.

On OpenVMS systems, always use the `sys$system:shutdown.com` command procedure when shutting down or rebooting an Instant Capacity partition.

---

Although a deferred activation does not immediately change the number of active cores, the intended active value is changed. This change does affect compliance checking, even before the partition reboot occurs. In particular, compliance checking is calculated as if the activation had not been deferred.



**IMPORTANT:** On OpenVMS Instant Capacity systems, HP strongly recommends that you activate cores using the `icapmodify` or the `ICAP SET` command. The use of the `START CPU` command on an Instant Capacity system can result in unintended consequences, such as a reduction of available temporary capacity. Another unintended side effect might be the adjustment in core usage across the complex, depending on the intended core settings on the partition where the `START CPU` command was issued.

---

To activate one or more inactive cores, use the `icapmodify` command as `root`. For more information about this command, see `icapmodify((1M))`.

### Constraints

The Instant Capacity software does not activate cores that are marked for deconfiguration. Also, you cannot use Instant Capacity to activate more cores than are configured in the current nPartition. If you want more cores, you must modify the nPartition with the `parmodify` command. You can use Instant Capacity to activate more cores than are configured into the current virtual partition, but only if the associated nPartition contains enough unassigned cores

and the virtual partition allows enough additional cores to fulfill the request. Otherwise, use the `parmodify` command to reconfigure the `nPartitions`, or use the `vparmodify` command to remove cores from other virtual partitions within the same `nPartition` (essentially adding to the unassigned pool).

## Example Core Activation Session

Example 4-3 shows how to activate an additional core in an `nPartition` environment. At the beginning of this activation session, there are a total of 4 cores in the partition; 2 cores are activated and 2 are inactive, but usage rights have been acquired to activate at least one inactive core. In this example, 1 additional core is activated, leaving the partition with 3 active cores and 1 inactive core.

### Example 4-3 Activating an Additional Core (HP-UX)

---

```
> /usr/sbin/icapmodify -a 1 "Add CPU for new FY: Bill P."
```

```
3 cores are intended to be active and are currently active.
```

---

In this example, note the following:

- The core activation is instant (that is, a reboot is not required).
- The double-quoted text serves as an audit trail of why the activation was done and who performed it. This information is optional and is written to the Instant Capacity log file (`var/adm/icap.log`) if provided.



**NOTE:** To defer the activation until the next reboot, add the `-D` option to the command. For more information, see `icapmodify(1M)`.

---

The `icapmodify` command allows you to activate additional cores with the `-a` option, or set the total number of active cores with the `-s` option. For example, the `icapmodify -a 2` command activates two additional cores in a partition. The `icapmodify -s 2` command sets the total number of active cores in a partition to 2.

For details about software application implications when activating additional cores, see “Software Application Considerations” (page 71).

# Deactivating Cores

You can decrease processing capacity instantly on HP enterprise servers with the Instant Capacity software (even on servers with sufficient usage rights for all cores to be active simultaneously). Any number of active cores can be deactivated at any time, within the following partition constraints. Core deactivation can be useful for load balancing cores in nPartitions (hard partitions) of Instant Capacity systems.

## Deactivating Cores in nPartitions

The software provides two types of core deactivation:

- *Instant* (`icapmodify` command's default behavior) — deactivation occurs immediately.
- *Deferred* (`icapmodify -D`) — deactivation occurs after the next reboot of the partition.

Instant deactivation of cores occurs when the `icapmodify` command is used with the `-d` option but without the `-D` option.



---

**IMPORTANT:** On OpenVMS Instant Capacity systems, HP strongly recommends that you deactivate cores using the `icapmodify` or the `ICAP SET` command. The use of the `STOP CPU` command on an Instant Capacity system can result in unintended consequences, such as a reactivation of the core when an Instant Capacity reconciliation transaction is requested.

---

Deferred deactivation of cores occurs when the `icapmodify` command is used with both the `-D` and `-d` options. With the deferred option (`-D`), core deactivation occurs after a reboot of the partition. The scheduled timing of the reboot (and the core deactivation) can take place at a planned time. For example, if you deactivate cores in deferred activation mode and schedule a partition reboot to occur on the first day of the next month, the cores are deactivated at that time. Since deferred deactivation does not immediately decrease the number of active cores, compliance checking is not affected by deferred deactivation.

To deactivate one or more active cores, use the `icapmodify` command as `root`. For details, see `icapmodify(1M)`.

## Partition Constraints

An nPartition must have a minimum of one active core for each active cell. Deactivation of cores is limited by this rule. If the deactivation applies to a virtual partition, additional constraints might apply, such as the minimum number of cores specified for the virtual partition.

## Example Deactivation Session for Hardware-Partitionable Systems

Example 4-4 (page 60) shows how to deactivate an active core. At the beginning of this deactivation session, there are a total of 4 cores in the partition; 3 cores are active and 1 is inactive. In this example, 1 active core is deactivated, leaving the partition with 2 active cores and 2 inactive cores. As with activation, you do not specify a particular core to be deactivated. You specify only the number of cores to be deactivated.

#### Example 4-4 Deactivating an Active Core (HP-UX)

---

```
> /usr/sbin/icapmodify -d 1
```

```
2 cores are intended to be active and are currently active.
```

---



**NOTE:** In Example 4-4, the core deactivation is instant (that is, does not require a reboot). To defer deactivation until the next reboot, use the `-D` option with the command. For details, see `icapmodify(1M)`.

---

The `icapmodify` command allows you to either deactivate cores with the `-d` option or to set the total number of active cores with the `-s` option. For example, the `icapmodify -d 1` command deactivates 1 additional core in a partition. The `icapmodify -s 2` command sets the total number of active cores in a partition to 2.

# Overriding Deferred Activation and Deactivation

---



**NOTE:** Although this section discusses only activation of cores, the discussion applies also to deactivation of cores.

---

If you have performed a deferred core activation (using the `icapmodify -D` command), and the intended number of active cores specified is no longer desirable, you can override the (pending) deferred activation by performing another deferred or instant `icapmodify` operation. This second operation overrides the first activation.

You might experience one of the following scenarios with deferred activation:

- The deferred number of active cores was incorrect and you want it to be correct when the system reboots.
- The entire deferred operation was accidental and you want to undo it.

Example 4-5 and Example 4-6 describe how to override these situations.

## Example 4-5 Correcting an Incorrect Number of Deferred Active Cores (HP-UX)

---

1. On your system or partition you currently have 2 active cores and 2 inactive cores. You need 4 active cores, so you perform a deferred activation for 2 additional active cores by entering the following command:

```
/usr/sbin/icapmodify -D -a 2
```

2. Later, and prior to a system reboot, you realize that you need only 3 active cores (not 4). You can override the action in step 1 by entering the following command:

```
/usr/sbin/icapmodify -D -s 3
```

The `-s` option in step 2 sets the number of active cores. The activation takes place after the next system reboot due to the `-D` option. You could also perform step 2 without the `-D` option so that the `icapmodify` operation is instant.

---

## Example 4-6 Undoing an Accidental Deferred Activation (HP-UX)

---

1. On your system or partition, you currently have 2 active cores and 2 inactive cores. You accidentally perform a deferred activation for 1 additional active core by entering the following command:

```
/usr/sbin/icapmodify -D -a 1
```

2. Later, and prior to a system reboot, you realize that you didn't want to activate the additional core (which would give you 3 active cores) and you want the number of active cores to be 2. You can override the action in step 1 by entering the following command:

```
/usr/sbin/icapmodify -a 0
```

The `-a 0` option in step 2 overrides the previous (deferred) `icapmodify` command, which was executed in step 1. The `-a` option is relative to the number of active cores, not to the intended number of active cores.

You could accomplish the same result as step 2 with the following command:

```
/usr/sbin/icapmodify -s 2
```

---

## Load-Balancing Active Cores

Active cores can be redistributed across any or all partitions of a hardware-partitionable system if those partitions contain inactive cores.

For example, consider a system with two partitions:

- Partition 1 has 5 active cores and 3 inactive cores.
- Partition 2 has 8 active cores and 0 inactive cores.

You need to add processing power to Partition 1 because of application demand and you notice that the active cores in Partition 2 are underutilized.

Deactivating an active core in Partition 2 decreases the number of active cores in that partition, and activating one of the cores in Partition 1 increases the number of active cores in that partition. The total number of active cores in the complex is the same at the end of this operation.



**IMPORTANT:** To remain in compliance, you must perform the deactivation operation first.

---

The two partitions are left with the following:

- Partition 1 now has 6 active cores and 2 inactive cores.
- Partition 2 now has 7 active cores and 1 inactive core.

Redistribution of active cores does not affect compliance. This is because you do not change the overall number of active cores in the complex. If the complex was in compliance prior to the redistribution, it remains in compliance. Make sure that you have proper licensing for all proprietary and nonproprietary software when performing load balancing.

# Understanding and Managing Intended Active Values

The Instant Capacity software maintains a value for each nPartition of a complex called *intended active*. Fundamentally, the intended active value is the number of cores intended to be active after a reboot of the nPartition.

The concept and proper manipulation of intended active is critical because:

- It determines the number of cores which will be active upon the boot of an nPartition.
- The value is used to establish the allocation of core usage rights for an nPartition within the total number of core usage rights purchased for the complex. That is, it represents the intended distribution of core resources for each nPartition.

When the intended active value changes, the distribution of usage rights across the hard partitions is changed, and the core resources originally intended for one nPartition might be diverted to another nPartition. In fact, this is the mechanism that allows for load balancing of resources to occur. Within a GiCAP group, changes to the intended active value allow resources to be load balanced across the entire group. Naturally, it is important to be sure that this is an intentional redistribution of resources.

Note that even inactive nPartitions have an intended active value that represents a “reservation” of usage rights for that partition. This is necessary to ensure every configured nPartition has core resources in order to boot.



---

**IMPORTANT:** An nPartition that was created but never booted to HP-UX will implicitly reserve usage rights for all configured cores, regardless of any intended active value set for the nPartition.

---

Within a virtual partition environment, the intended active value is especially critical because a value that does not conform to the virtual partition assignments can result in a virtual partition environment that cannot be booted. For more information see “Boot Time Compliance” (page 66).

In a virtual partition environment, the behavior of the `icapmodify -a`, `-d` and `-s` options is different when there is unused capacity represented by an intended active value that exceeds the sum of the cores assigned to the virtual partitions in an nPartition, as described in “Activations and Deactivations in a Virtual Partition Environment” (page 64).

# Activations and Deactivations in a Virtual Partition Environment

Instant Capacity can be present on HP-UX systems or partitions where virtual partition technology is employed. In a virtual partition environment, cores that are not assigned to any virtual partition are considered inactive (in addition to other classes of inactive cores). Unassigned cores can be assigned (activated) or deassigned (deactivated) using either the `icapmodify` command or the `vparmodify` command, depending on the type of adjustment needed and the level of logging or reporting desired.

One important consideration is that `vparmodify` can be used to activate or deactivate cores in other virtual partitions within the `nPartition`; `icapmodify` only activates or deactivates cores within the current virtual partition (the partition where the command is invoked). The `vparmodify` command does not change the value for the number of intended active cores for the `nPartition`.

Another consideration is that core assignment via the `vparmodify` command does not result in Instant Capacity logging of the activation, email configuration change notification, or transmission of an asset report to HP.



**NOTE:** Deferred activations and deactivations are not supported in any virtual partition environment.

Instant Capacity always consumes unused capacity before it consumes additional usage rights when activating cores. Instant Capacity always releases usage rights from unused capacity before it releases usage rights by deactivating cores.

---

## When to Use the `vparmodify` or `icapmodify` Commands

When usage rights freed by deactivating cores are to be used in a different `nPartition`, use the `icapmodify` command. When usage rights freed by deactivating cores are to be used by another virtual partition within the same `nPartition`, use the `vparmodify` command. The following sections provide details on activating and deactivating cores within `nPartition` and virtual partition environments.

## What is Unused Capacity

In a virtual partition environment, the Instant Capacity software may allocate or deallocate cores from what is called “unused capacity”. *Unused capacity* is the difference between *actual active* and *intended active* cores, and it is a side effect of using `vparmodify` to deactivate cores while not immediately activating the cores in another virtual partition within the same `nPartition`. This is usually a transient state, since a user typically migrates cores from one virtual partition to another with a deactivate command followed by an activate command. However, unused capacity can persist if, for example, a utility such as `gWLM` ignores an error status from an activation and leaves the previously deactivated cores in an unassigned state. The Instant Capacity software always takes unused capacity into account when a request is made to activate or deactivate cores such that it attempts to eliminate the discrepancy between intended active and actual active.

## Static Virtual Partitions

If a virtual partition is static (that is, if its resources cannot be migrated, added, deleted, or modified) and you attempt to activate or deactivate cores, the Instant Capacity software displays a message indicating that the configuration cannot be modified.



---

**NOTE:** The `icapstatus` command output indicates that the number of cores that can be assigned (to the local virtual partition) is zero if the static resource attribute for the local virtual partition is set.

---

## Activating Cores in a Virtual Partition Environment

In a virtual partition environment, the `icapmodify` command must be used to modify processing capacity when you are making any adjustment to an nPartition or to multiple nPartitions. When you execute the `icapmodify` command to activate a core, the Instant Capacity software verifies that the request can be satisfied. If so, the *local nPartition's* intended active number is increased, and the appropriate number of cores are added to the local virtual partition.

When cores in a virtual partition are activated using the `icapmodify` command, if unused capacity is available, the cores are activated from the unused capacity balance before increasing the active cores. If enough cores are available to meet the request, the proper number of cores are added to the local virtual partition. If there is any unused capacity available prior to the activation, the result is that the actual active count changes by more than the intended active count, and the intended active count may not change at all. However, the discrepancy between intended active and actual active is reduced. If there is no unused capacity available, the desired number of cores will be activated in the virtual partition.

If you are adjusting core assignments across virtual partitions in a single nPartition, use the `vparmodify` command for the best coordination between the Instant Capacity software and the vPars software, and for optimized performance. The `vparmodify` command is the fastest and most efficient way to adjust capacity within virtual partitions of a single nPartition, but it does not affect the intended active count for the nPartition. Therefore, it cannot be used to migrate unused capacity either to or from other nPartitions. When you execute the `vparmodify` command to activate a core, the command verifies with the Instant Capacity software how many cores are available for activation. This number is calculated as the difference between the local nPartition's intended active number and the total number of cores assigned to the vPars database.

## Deactivating Cores in a Virtual Partition Environment

In a virtual partition environment, the `icapmodify` command must be used to modify processing capacity when you are making any adjustment to an nPartition or to multiple nPartitions. When you execute the `icapmodify` command to deactivate a core, the Instant Capacity software verifies that the request can be satisfied. If so, the local nPartition's intended active number is decreased, and the appropriate number of cores are removed from the local virtual partition.

When the `icapmodify` command is used in a virtual partition, it checks with the Instant Capacity software to determine how many cores are available for deactivation. This number is calculated as the difference between the local nPartition's intended active number and the total number of cores assigned to the vPars database. If there is any unused capacity available prior to the deactivation, the result is that the actual active count changes by more than the intended active count, and the intended active count may not change at all. However, the discrepancy between intended active and actual active is reduced.

If you are adjusting core assignments across virtual partitions in a single nPartition, use the `vparmodify` command for the best coordination between the Instant Capacity software and the vPars software, and for optimized performance. The `vparmodify` command is the fastest and most efficient way to adjust capacity within virtual partitions of a single nPartition, but it does not affect the intended active count for the nPartition and it therefore cannot be used to migrate unused capacity either to or from other nPartitions. When you execute the `vparmodify` command to deactivate a core, authorization is not required from the Instant Capacity software.

Whether you are activating or deactivating cores, the `icapmodify` command adjusts only the number of dynamic cores, and it does not explicitly identify specific cores.

## Boot Time Compliance

A compliance check is performed whenever a virtual partition is booted. If the total number of cores assigned to all virtual partitions in the current vPar database exceeds the nPartition's intended active core count, the Instant Capacity software notifies the vPar monitor. The monitor prevents any virtual partition from booting until the user first performs a hard partition boot and then modifies either the vPar configuration or the Instant Capacity intended active count for the nPartition. [Example 4-7](#) shows a sample boot-time compliance message sent when a virtual partition is prevented from booting.

### **Example 4-7 vPar Boot-Time Compliance Message**

---

To: root@parl.yourorg.com  
Subject: vPar Boot Time Compliance

This message is being sent to inform you that a vpar is not being allowed to boot because doing so would take this complex out of compliance from an Instant Capacity perspective. The number of cores assigned to this vPar database (/stand/vpdb) exceeds the number of intended active cores by 1. To correct this problem, boot this partition back into an nPartition and modify the vPars assigned to this database or modify the number of intended active cores for this nPartition.

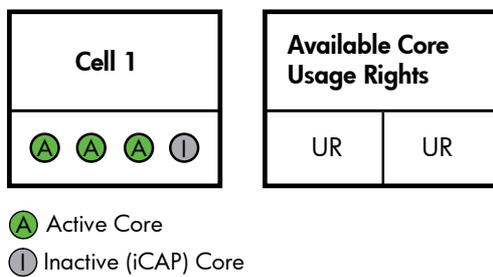
---

## Assigning a Cell to a Partition

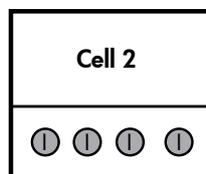
A cell can be assigned to a partition only if sufficient cell usage rights are available across the complex, as well as sufficient memory usage rights to enable activation of all the memory on the cell, and sufficient usage rights for at least one core of the cell to be active.

In an Instant Capacity system, when a cell is assigned to a partition, depending on the number of core usage rights available in the system when the cell is assigned, the number of intended active cores for the partition automatically changes. Figure 4-2, Figure 4-3, and Figure 4-4 illustrate a single partition with one assigned and one unassigned cell.

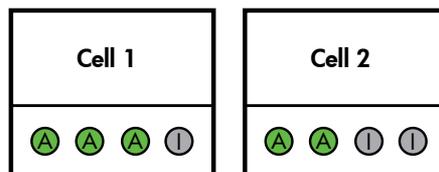
**Figure 4-2 Partition premodification state: One cell assigned with 3 active and 1 inactive cores, and usage rights for 2 additional cores**



**Figure 4-3 Premodification state: Unassigned cell with 4 inactive cores and no usage rights**



**Figure 4-4 Partition postmodification state: Cell 2 assigned to partition**



When Cell 2 is assigned to the partition, the number of intended active cores for the partition is automatically changed to 5. When the partition is rebooted, 5 cores in the partition are activated.

When an unassigned cell is assigned to a partition via the `parmodify` command, the Instant Capacity software determines the number of available core usage rights in the complex, and uses this number to activate as many cores as possible in the new cell. (This number typically corresponds to the `icapstatus` value for `Additional cores` that can be activated with current usage rights or the actual number of cores, whichever is smaller. In order to assign an inactive cell to a partition, this value must be nonzero.)

However, when a new partition is created via the `parcreate` command, core usage rights must be available for *all* cores configured on the new partition's cells, and core usage rights are automatically assigned to these new cells. For more information see “New Partition Creation and Instant Capacity” (page 196).

Note that when determining available usage rights, the Instant Capacity software calculates the usage rights consumed by each partition to be the greater of either the actual active or intended active cores for the partition.



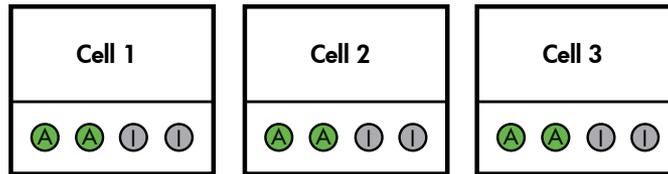
**NOTE:** Cell boards are assigned to specific partitions and cannot be shared between partitions. All cores on a cell board are accessible by only the partition to which the cell board is assigned. Cores on one cell board cannot be shared across multiple partitions.

---

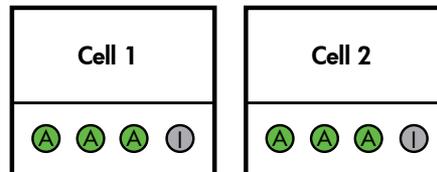
## Unassigning a Cell from a Partition

When a cell is unassigned from a partition in an Instant Capacity system, the number of intended active cores in the partition decreases only if the number of cores being removed with the cell is greater than the number of expected inactive cores in the partition. In [Figure 4-5](#), [Figure 4-6](#), and [Figure 4-7](#) showing a single partition system with 3 cells, the number of intended active cores remains the same because the number of cores with the removed cell (4) does not exceed the total number of expected inactive cores in the partition (6).

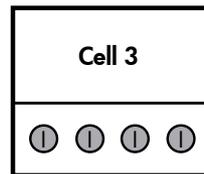
**Figure 4-5 Partition premodification state: Three cells with 2 active and 2 inactive cores in each, and 6 expected inactive cores**



**Figure 4-6 Partition postmodification state: Cell 3 is unassigned (total of 6 active cores remaining)**



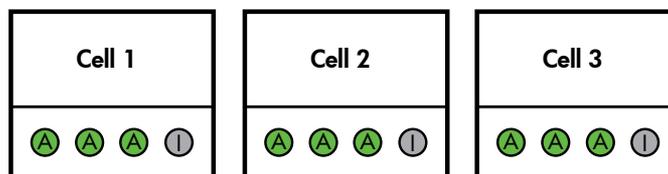
**Figure 4-7 Partition postmodification state: Unassigned Cell 3 with 4 inactive cores**



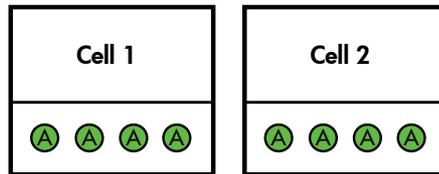
When Cell 3 is unassigned from the partition, the number of intended active cores for the partition remains at 6. When the partition is rebooted, a total of 6 cores are activated. Cell 3 becomes an unassigned cell with 4 inactive cores, essentially freeing up usage rights that are distributed among the remaining cells.

In [Figure 4-8](#), [Figure 4-9](#), and [Figure 4-10](#), the number of cores removed (4) is greater than the number of expected inactive cores in the partition (3). When this happens, the number of intended active cores is automatically set to the total number of remaining cores in the partition (8).

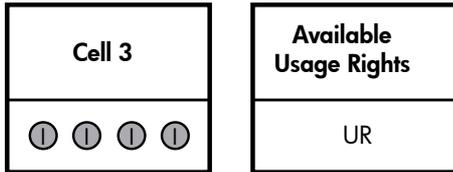
**Figure 4-8 Partition premodification state: Three cells with 3 active and 1 inactive cores in each, and 3 expected inactive cores**



**Figure 4-9 Partition postmodification state: Unassigned Cell 3 (total of 8 active cores are set)**



**Figure 4-10 Postmodification state: Unassigned Cell 3 with 4 inactive cores, with usage rights available for 1 additional core**



When Cell 3 is unassigned from the partition, the number of intended active cores changes from 9 to 8 (because 8 is the total number of cores remaining in the partition). When the partition is rebooted, a total of 8 cores are activated. Cell 3 becomes an unassigned cell with 4 inactive cores, and (unused) usage rights are available for 1 additional core for the complex.



**NOTE:** If you intend to migrate a cell from one partition to another, you can control the number of core usage rights available to the cell (in the new partition) by deactivating cores in the partition you removed the cell from. By deactivating cores, you make core usage rights available to the entire complex.

## Software Application Considerations

Some software applications size themselves based on the number of available cores when the application is started. If an application is running when an additional core is activated, the application might not recognize the newly activated core as available for processing. Therefore, you might need to do one of the following for optimal performance with this type of application:

- Restart the application in order for it to recognize the presence of newly activated cores.
- Reconfigure the application prior to restarting it, for maximum performance benefits of the newly activated core.
- Use the deferred activation option when activating cores so that processors are activated only in conjunction with system reboots. For details, see `icapmodify(1M)` for details



**IMPORTANT:** When you activate a core, the number of active cores in the system increases. Consequently, a license upgrade might be required for some of the software from HP or other vendors on your system.

---

## Test Activation of Cores Using Temporary Capacity

You might want to test your software application for proper operation and improved performance by activating an additional core. The use of temporary capacity or Instant Access Capacity (IAC) is required for activation of a core without usage rights for testing purposes. For details, see Chapter 5: “Temporary Instant Capacity” (page 75).

The following testing guidelines are meant to be an aid to your test plan. You might need to get consulting help to develop a more detailed plan.

1. Test your applications for proper functionality and performance first by testing with the number of inactive cores equal to the number of cores without usage rights. (The system should already be configured this way.) Be sure to check measurement tools that monitor core usage.
2. Acquire temporary capacity for the necessary amount of core test activation.
3. Use temporary capacity to activate one or more inactive cores to be used while your applications are running.
4. Confirm that measurement tools, which monitor processing usage, account for the newly activated cores.
5. Verify that applications are benefiting from the performance of the extra cores (as per your expectations for your applications). Some applications might need to be restarted or reconfigured to take advantage of the newly activated cores.
6. When you are finished with your testing, deactivate cores until the number of inactive cores again matches the number of cores without usage rights, thereby stopping the usage of temporary capacity.
7. Use `icapstatus` to verify that no cores are consuming temporary capacity.

# Replacement of Failed Cores

## HP-UX LPMC and HPMC

If an active core fails with a Low Priority Machine Check (LPMC) in a partition with Instant Capacity, its processing capacity is replaced instantly by an inactive core, if any are available in the partition. The failed core is marked for deconfiguration during the next system reboot.

For additional considerations in a virtual partition environment, see “LPMC Deactivations in Virtual Partitions” (page 73).

If an active core fails with a High Priority Machine Check (HPMC), then upon reboot the failed core is deconfigured and its processing capacity is instantly replaced by an inactive core, if any are available in the partition.



---

**NOTE:** In both of the preceding scenarios, replace the failed core in a timely manner using your normal hardware support process.

---

## LPMC Deactivations in Virtual Partitions

In a vPar environment, if the LPMC monitor deactivates a core, it automatically replaces the failing core with an Instant Capacity core from the free pool, if such a core is available. The failing core remains in the virtual partition until either the virtual partition or the virtual partition monitor is rebooted.

For more information about LPMC in vPars, see the various white papers at the HP Documentation web site (search for “LPMC”):

<http://docs.hp.com>

## Failed Monarch Processors (HP-UX)

*Monarch processors* that are failing with an LPMC are not instantly replaced. When a monarch processor experiences an LPMC, the LPMC monitor marks the processor for deconfiguration; however, the LPMC monitor cannot deactivate the processor until the system is rebooted.

Deactivation of a monarch processor is not possible because it is the controlling processor of the operating system (CPU 0). Therefore, the system cannot replace a (failing) monarch processor.

If your system has only one active processor, that processor is considered a monarch processor and it cannot be replaced on line. A reboot of the system is required to replace the failing monarch processor.

If there are multiple active processors in your system, one of them is designated as the monarch processor, and the other (nonmonarch) processors can be replaced on line. If the monarch processor fails, it cannot be replaced without a reboot.

## Replacement of Failed Cores on OpenVMS

If a core is experiencing correctable errors, shut it down and start up another Instant Capacity core, thereby keeping the active-core count constant.

If a core experiences a fatal problem leading to a system crash, upon reboot you can start another Instant Capacity core, thereby replacing the failed core and keeping the active-core count constant.

## Failed OpenVMS Primary Processors

An OpenVMS primary processor that is failing cannot be instantly replaced.

If your system has only one active processor, it is considered a primary processor and it cannot be replaced on line. A reboot of the system is required to replace the failing primary processor.

If there are multiple active processors in your system, one of them is designated as the primary processor and the other (nonprimary) processors can be replaced on line. If the primary processor fails, it cannot be replaced without a reboot.

---

## 5 Temporary Instant Capacity

This chapter covers the following topics:

- “Temporary Instant Capacity Overview” (page 76)
- “Ordering Temporary Instant Capacity” (page 78)
- “Using Temporary Instant Capacity” (page 79)
- “Temporary Capacity and Virtual Partitions” (page 81)
- “Tracking Usage of Temporary Instant Capacity” (page 83)
- “Temporary Instant Capacity Warning Period” (page 85)
- “Temporary Instant Capacity Expiration and Compliance Enforcement” (page 86)
- “Temporary Instant Capacity Exceptions” (page 87)

# Temporary Instant Capacity Overview

You can purchase an amount of temporary capacity (TiCAP) time for inactive cores without usage rights in your Instant Capacity system. Temporary capacity can be purchased in units of processing days. Temporary capacity allows one or more inactive cores to be activated for up to the specified period of prepaid processing time, without requiring permanent usage rights for the cores.

You can activate and deactivate inactive cores as necessary until the elapsed activation time equals the duration of your prepaid temporary capacity. For example, with a prepaid duration of 30 days of temporary capacity, you can activate one core for 30 days or four cores for 1 hour a day for 180 days (or any combination that totals 43,200 minutes).

Temporary capacity activations are persistent. This means that activations using temporary capacity survive in a partition that is rebooted. You must deactivate cores to stop consumption of temporary capacity. The cores you deactivate need not be on the same partition as those you activated to start consuming temporary capacity.



**NOTE:** Temporary capacity credits can be used on any partition in the complex for which they were purchased. Temporary capacity credits are not transferable from one system to another unless the systems are in the same Global Instant Capacity group. For details of temporary capacity in a GiCAP group, see [Chapter 7](#).

---

If temporary capacity is depleted and you continue to have more active cores than core usage rights across the complex, on the next reboot of any partition in the complex the software automatically deactivates one or more cores in order to bring the system into a state closer to compliance. The Instant Capacity software deactivates as many cores as necessary to either stop consumption of temporary capacity or to bring the partition to the minimum number of required active cores (one per active cell board).

---

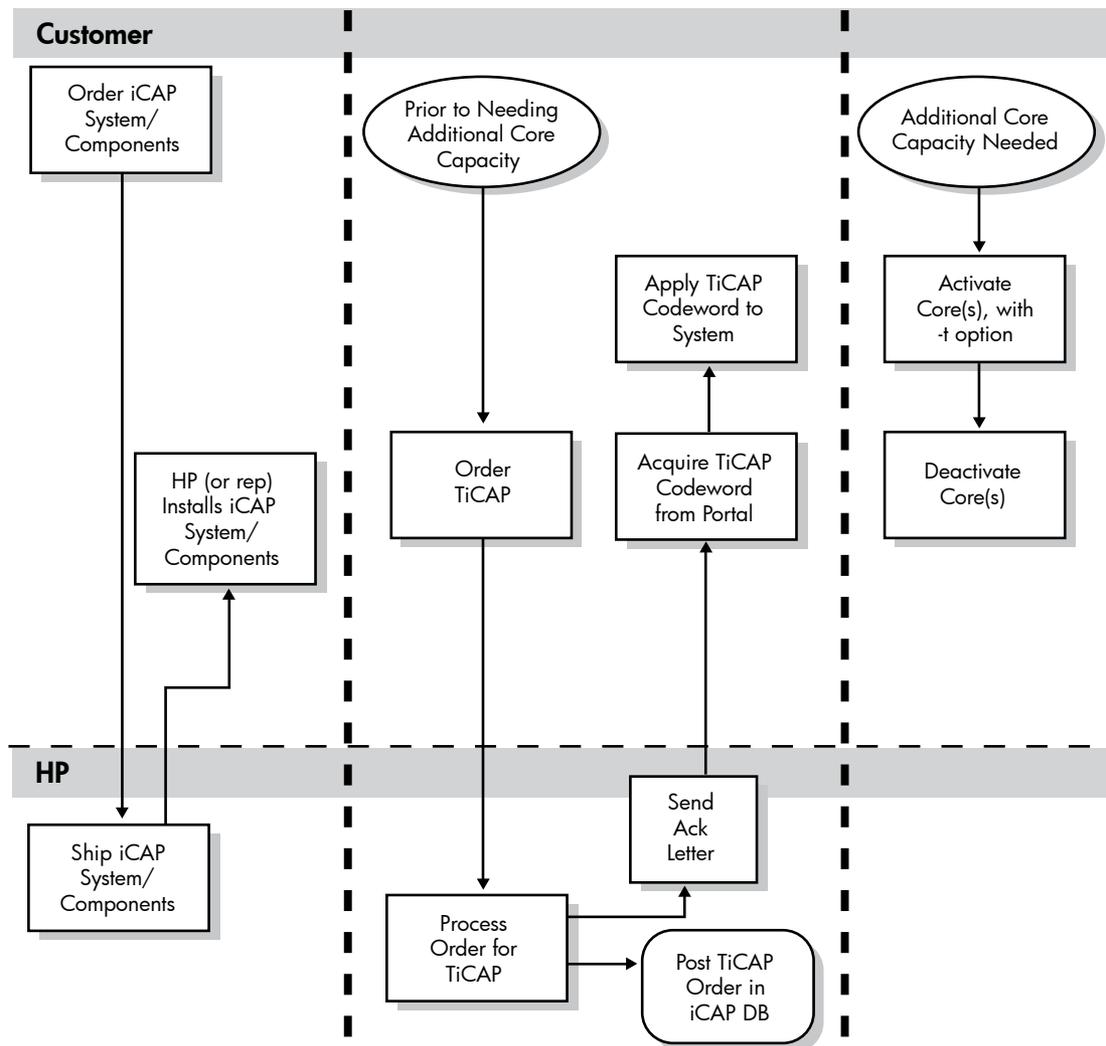


**IMPORTANT:** Temporary capacity can be used to activate Instant Capacity cores on a temporary basis only. It cannot be used to activate Instant Capacity cell boards or Instant Capacity memory.

---

Figure 5-1 illustrates the process of purchasing and applying temporary capacity.

**Figure 5-1 Using Temporary Instant Capacity: Temporary Activation of Cores Without Usage Rights**



## Ordering Temporary Instant Capacity

To add temporary capacity credits to a system, order the desired quantity of the Temporary Instant Capacity product for your type of server. The system serial number is required for orders of Temporary Instant Capacity.

Instant Capacity cores that are added to an existing system can include some additional temporary capacity called Instant Access Capacity (IAC). Over time, the IAC can be consumed and you might want to also order additional temporary capacity in order to continue activating cores on a temporary basis. Temporary capacity is purchased in units of processing days, each of which can be used to activate a single core for 24 hours (either continuously or spread over several days), or multiple cores for portions that add up to a single day. Then you must follow the configuration procedures for each partition (see “Acquiring and Configuring Temporary Instant Capacity” (page 79)).

### HP-UX Licensing and Support with Temporary Instant Capacity

When you purchase temporary capacity, the temporary HP-UX license-to-use is included when Instant Capacity cores are activated using temporary capacity. You might also need software licenses for nonproprietary software. Check with your application software vendor for licensing requirements. Since licensing requirements can change without notice, check your contract to understand the details of software licensing with temporary capacity.

### OpenVMS Licensing and Support with Temporary Instant Capacity

When you purchase temporary capacity, the temporary OpenVMS license-to-use is included when Instant Capacity cores are activated using temporary capacity. The OpenVMS License Management Facility recognizes when Temporary Instant Capacity cores are activated, and they treat the usage as compliant. For nonproprietary software that uses per-core licensing, check with the vendor for licensing requirements.

# Using Temporary Instant Capacity

## Acquiring and Configuring Temporary Instant Capacity

To add temporary capacity to a system that contains Instant Capacity cores (cores without usage rights), follow this procedure:

1. Order the desired amount of Temporary Instant Capacity for your type of server by submitting a purchase order to HP. Be sure to specify the system serial number.

After you purchase an amount of temporary capacity, HP sends you a letter that tells you how to acquire a temporary capacity codeword and apply it to the system.

2. Acquire the temporary capacity codeword from the Utility Pricing Solutions portal (<http://www.hp.com/go/icap/portal>)
3. Apply the temporary capacity codeword using the `icapmodify -C` command, on any partition of the server, as shown in Example 5-1.

### Example 5-1 Applying a Temporary Capacity Codeword (HP-UX)

---

```
> /usr/sbin/icapmodify -C vnyqD.qjieC7e.LaLdQGH.4aCNYBp-BQk3w9n.jfDhpvz.LiEB58C.7Q3dca2
```

The following valid codeword has been applied to the complex:

```
Temporary Capacity Codeword  
30 days 0 hours 0 minutes
```

Use `icapstatus(1M)` to see the results of the application of this codeword.

---



**NOTE:** Instant Capacity codewords are based on both the serial number of a system and a unique sequencing value for that server. These codewords must be applied in the sequence in which they are obtained for a particular server. They can be applied to any partition on the server.

---

4. Optional: To view temporary capacity balances on the portal, configure your partition for email connectivity to HP. For details, see “Configuring Email on Instant Capacity Systems” (page 202).

## Using Temporary Instant Capacity

Auditing of temporary capacity is done at the complex level on Instant Capacity systems that support partitioning and that are not part of a GiCAP group. For information about GiCAP, see “Global Instant Capacity and Temporary Capacity” (page 115).

In a system that contains Instant Capacity cores with the temporary capacity RTU codeword applied, you can activate one or more cores in an nPartition and allow them to use temporary capacity by entering the `icapmodify -t -a<number_of_cores>` command.

In Example 5-2, two cores are currently active in the partition. You want to activate a third core, but there are no available usage rights for activation. Because temporary capacity is available on the system, you can activate a third core with it.

## Example 5-2 Activating an Instant Capacity Core with Temporary Capacity (HP-UX)

---

```
> /usr/sbin/icapmodify -t -a 1
```

```
3 cores are intended to be active and are currently active.
```

```
Number of cores using temporary capacity: 1  
Projected temporary capacity expiration: 12/22/08 08:00:00
```

---



**NOTE:** Temporary capacity cannot be used to activate Instant Capacity cores in inactive Instant Capacity cell boards. You must purchase additional usage rights for the cell board and perhaps also for the memory of the cell board.

---

To decrease the use of temporary capacity or to stop using it entirely, deactivate the appropriate number of cores. (Do not use the `icapmodify -t` command when deactivating cores.) Temporary capacity is no longer used when the number of active cores is equal to or less than the number of cores with usage rights, across the complex.

To deactivate one or more cores in an nPartition, use the `icapmodify -d <number>` command.

# Temporary Capacity and Virtual Partitions

If temporary capacity is being consumed in any virtual partition environment (previously authorized with `icapmodify -a n -t`), deactivating a core with the `vparmodify -d` command temporarily reduces the consumption of temporary capacity, although there may be a delay of up to 30 minutes for the consumption of temporary capacity to cease. A subsequent core activation using the `vparmodify -a` command increases consumption of temporary capacity, if the activation results in more active cores than available core usage rights.

Use the `icapmodify -d` command to immediately decrease or cease the use of temporary capacity. It is not necessary to use the `-t` option with the `-d` option. Use the `icapmodify -a -t` command to resume temporary capacity consumption.

The following is an example of the output from an `icapstatus` command on a partitionable system containing vPars.

```
> /usr/sbin/icapstatus

Software version:          B.11.31.09.00.00.71
System ID:                 zoo6
Serial number:             USR4020003
Product number:           A6093A
Unique ID:                 Z3e0ec8e078cd3c7b
System contact e-mail:    mjones@corp.com
From e-mail:               Set to the default ('adm')
Asset reporting:          on
Temporary capacity warning period: 15 days
Exception status:         No exception

Local Virtual Partition Status
-----
Total number of assigned cores:          4
Number of active assigned cores:         4
Number of inactive assigned cores:       0
Additional cores that can be assigned with current usage rights: 2
Number of cores that could be assigned with additional usage rights: 1
Number of cores that can be assigned with temporary capacity: 0
Number of cores currently unavailable for assignment: 0

Local nPartition Status (09/17/08 12:34:56)
-----
Total number of configured cores:        8
Number of Intended Active cores:         3
Number of active cores:                   5
Number of inactive cores:                 3

Instant Capacity Resource Summary
-----
Number of cells without usage rights:     0
Number of inactive cells:                 0
Amount of memory without usage rights:    0.0 GB
Amount of inactive memory:                0.0 GB
Number of cores without usage rights:     4
Number of inactive cores:                 6
Number of cores that must be deactivated (insufficient usage rights): 0
Temporary capacity available:              0 days, 0 hours, 0 minutes

Allocation of Instant Capacity Resources among the nPartitions
-----
          Intended Actual
nPar  Total Active  Active =====Inactive===== Runs
ID    Cores Cores   Cores Cores  Memory  Cells iCAP nPar Name
=====
  0     4     4     4     0    0.0 GB    0 Yes zoo0
  1     4     4     4     0    0.0 GB    0 Yes zoo1
  2     8     4     4     4    8.0 GB    1 Yes zoo2
  3     4     0     4     0    0.0 GB    0 Yes zoo3a
  4     0     4     4     0    0.0 GB    0 Yes zoo5
  5     2     2     4     2    0.0 GB    0 Yes zoo7
  6     3     3     4     1    0.0 GB    0 Yes zoo6 (local)
```

8	0	1	4	3	0.0 GB	0 Yes	zoo8
9	0	4	4	0	0.0 GB	0 Yes	zoo9
10	0	4	4	0	0.0 GB	0 Yes	zoo10
11	0	4	4	0	0.0 GB	0 Yes	zoo11
12	0	4	4	0	0.0 GB	0 Yes	zoo12
13	0	4	4	0	0.0 GB	0 Yes	zoo13
N/A	8	N/A	N/A	8	4.0 GB	2 N/A	Unassigned Cells

# Tracking Usage of Temporary Instant Capacity

The `icapstatus` command provides the following information about the use of temporary capacity on the system:

- Amount of temporary capacity remaining (in days, hours, and minutes)
- Number of cores using temporary capacity (number of active cores without usage rights)
- Projected temporary capacity expiration date and time (based on the current temporary capacity consumption rate)

You can find this information in the `Instant Capacity Resource Summary` section of the `icapstatus` command output.

**Example 5-3** shows the display of temporary capacity information from the `Instant Capacity Resource Summary` section of the `icapstatus` command output:

## Example 5-3 Temporary Capacity Information from `icapstatus` Command (HP-UX)

---

```
> /usr/sbin/icapstatus
```

```
Instant Capacity Resource Summary
```

```
-----  
Number of cells without usage rights:          0  
Number of inactive cells:                     0  
Amount of memory without usage rights:        0.0 GB  
Amount of inactive memory:                    0.0 GB  
Number of cores without usage rights:         4  
Number of inactive cores:                     3  
Number of cores using temporary capacity:     1  
Temporary capacity available:                 10 days, 1 hours, 40 minutes  
Projected temporary capacity expiration:      12/31/08 16:00:00  
-----
```

If the system is part of a GiCAP group, additional group temporary capacity might also be available, and is visible only via the `icapmanage` command on the Group Manager system.

## Temporary Capacity Expiration Reminder

The Instant Capacity software calculates when the temporary capacity balance will expire, based on the current consumption rate. After the temporary capacity balance reaches a certain residual number of days (see “Temporary Instant Capacity Warning Period” (page 85)), a reminder email message is automatically sent to the system contact, if one is specified, and `root`. These messages are sent on a daily basis until temporary capacity expires. **Example 5-4** shows a temporary capacity expiration reminder email message. Note that the output from `icapstatus` during the warning period includes a warning about the expiration of temporary capacity.

## Example 5-4 Temporary Capacity Expiration Reminder

---

To: root@parl.yourorg.com  
Subject: Temporary Capacity Expiration Reminder

```
*****  
**** Failure to perform the following steps will result in the complex ****  
**** attempting to deactivate cores on any booting partitions until ****  
**** the complex is in compliance with the Instant Capacity contract. ****  
*****
```

This message is being sent to remind you that your Instant Capacity complex (containing the partition krmt10b) has 3 cores currently consuming temporary capacity (TiCAP) and that the temporary capacity balance at the current consumption rate is projected to expire on or around:

12/31/08 16:00:00

You can view the current temporary capacity balance and consumption rate, by using the `icapstatus` command.

To adjust the number of calendar days to receive the temporary capacity warning before temporary capacity actually expires use: `icapmodify -w`

Before the temporary capacity balance runs out, you must perform one of the following steps:

1. Purchase additional temporary capacity and apply the temporary capacity codeword to the complex.
2. Deactivate cores until the number of inactive cores on the complex matches the number of cores without usage rights, reported by `icapstatus`.
3. Purchase additional core usage rights to match the number of cores currently consuming temporary, and apply the Right to Use codewords to the complex so that they can be permanently activated.

See the Instant Capacity user's guide at `/usr/share/doc/icapUserGuide.pdf` for more information.

---

## Temporary Instant Capacity Warning Period

By default, the Instant Capacity software sends the expiration reminder when the temporary capacity balance is projected to expire within 15 days. You can adjust that warning period by specifying a different value with the `icapmodify` command, using the `-w` option. For example, the following command specifies a longer warning period, for more advance notice:

```
> icapmodify -w 20
```

The Temporary Capacity Warning Period has been successfully set to 20 days.

# Temporary Instant Capacity Expiration and Compliance Enforcement

---



**IMPORTANT:** If you leave cores without usage rights activated beyond the purchased temporary capacity duration, the software automatically deactivates one or more cores on the next reboot of any partition in the complex.

---

After the temporary capacity is depleted and you continue to have more active cores than usage rights across the complex, a notice similar to the following appears at the bottom of the `icapstatus` output:

```
WARNING: Temporary capacity has expired and this complex is out
of compliance with the Instant Capacity contract because there
are 2 more active cores than there are core usage rights.
Deactivation of cores may occur during partition reboot to
bring the complex into compliance. In order to avoid the
deactivation of cores upon reboot, you need to take corrective
action immediately. Either deactivate 2 core(s), apply
additional temporary capacity codewords, or purchase and apply
Right to Use codewords for 2 core(s).
```

As stated in the warning, if cores without usage rights continue to be used, then on the next reboot of any partition in the complex, the software automatically deactivates one or more cores in order to bring the system into a closer state of compliance. The Instant Capacity software deactivates as many cores as is necessary to either stop consumption of temporary capacity or to bring the partition to the minimum number of required active cores. You must purchase either additional temporary capacity or the appropriate number of usage rights (RTU codewords) to be in full compliance.

For examples of the error messages that are sent as a result of compliance enforcement, see “Temporary Instant Capacity Exceptions” (page 87).

Auditing of temporary capacity is done at the complex (OpenVMS node) level on Instant Capacity systems that support partitioning (and are not part of a GiCAP group). Although temporary capacity might have been purchased for use by a specific partition, it is available to all partitions in the complex or OpenVMS node (or to the entire GiCAP group, if applicable).

## Temporary Instant Capacity Exceptions

### Error for Activation with Insufficient Temporary Capacity

You cannot activate an Instant Capacity core with temporary capacity unless there is a sufficient balance of temporary capacity available. For details on how to increase the temporary capacity balance, see “Acquiring and Configuring Temporary Instant Capacity” (page 79).

Example 5-5 shows an error message for attempting to activate an inactive core without usage rights and without a sufficient balance of temporary capacity:

#### **Example 5-5 Error Message for Activation with Insufficient Temporary Capacity (HP-UX)**

---

```
> /usr/sbin/icapmodify -t -a 1
```

```
ERROR:  Operation not approved because there is not enough temporary capacity  
        to satisfy the request. Activations require at least 30 minutes  
        worth of temporary capacity per core consuming temporary capacity.
```

---

### Temporary Capacity Balance Needing Action

If the temporary capacity balance reaches 30 minutes or less, the `icapstatus` command output displays “less than 30 minutes” in the `Exception status` field (at the beginning of the `icapstatus` output). When this state occurs, you need to take corrective action immediately by doing one of the following:

- Deactivate Instant Capacity cores that are using temporary capacity.
- Apply additional temporary capacity codewords.
- Acquire additional core usage rights and apply the RTU codeword.

### Temporary Capacity Negative Balance

A complex is out of compliance with the Instant Capacity contract if a negative balance of temporary capacity occurs.

The Instant Capacity software sends an exception report (via email) if there is a negative balance of temporary capacity. Exception information is also written to the `syslog` file. For details of the exception report for a negative temporary capacity balance, see “Handling Compliance Exceptions” (page 137).

If you continue to have more active cores than core usage rights across the complex, a negative capacity balance results in a compliance enforcement action, as described in “Temporary Instant Capacity Expiration and Compliance Enforcement” (page 86). If there is a negative temporary capacity balance but the number of cores with usage rights is greater than or equal to the number of active cores, then the complex remains in an exception state, but without (additional) enforcement action.

### Temporary Capacity Enforcement

When the temporary capacity balance is depleted and you continue to have more active cores than core usage rights across the complex, an enforcement action occurs at partition reboot to bring the system into a state closer to compliance (by deactivating one or more cores). Example 5-6 shows the message that is sent when enforcement results in a partially compliant state but temporary capacity continues to be depleted. Example 5-7 shows the message that is sent when the enforcement is able to deactivate enough cores so that temporary capacity is no longer used.

## Example 5-6 Error Message for Temporary Capacity Partial Enforcement

---

To: root@par1.yourorg.com  
Subject: Instant Capacity enforcement notice

This message is being sent to inform you that, due to expiration of temporary capacity, 1 additional core(s) were deactivated on your Instant Capacity system (containing the partition par1) to bring the complex into compliance.

Prior to deactivation, the number of active cores exceeded the number of available core usage rights by 3. 3 core(s) without usage rights were found to be active in the complex. This state was likely the result of having activated Instant Capacity core(s) using temporary capacity (TiCAP), and then allowing the TiCAP balance to expire prior to deactivation of the core(s).

As a result, the intended active value was reduced by 1 and 1 core(s) were deactivated.

There are currently 3 active core(s) and 1 core usage rights. This complex is not in compliance with the Instant Capacity contract. Other partitions may also experience core deactivation upon reboot until compliance is restored. To bring the system back into compliance now, perform one or more of the following steps:

1. Purchase additional temporary capacity and apply the temporary capacity codeword(s) to the complex.
2. Deactivate cores until no cores are consuming temporary capacity.
3. Purchase additional usage rights to match the number of cores consuming temporary capacity and apply the Right to Use codewords to the complex so that they can be permanently activated.

To activate these 1 core(s) again, you can perform one of the following actions:

1. Purchase additional temporary capacity and apply the TiCAP codeword(s) to the complex, and use temporary capacity to activate the core(s).
2. Deactivate cores in other partitions after the complex is in compliance. This frees up core usage rights which can be used to activate cores on this partition.

You can view the current temporary capacity compliance of your system by using the `icapstatus` command.

See the Instant Capacity user's guide at `/usr/share/doc/icapUserGuide.pdf` for more information.

---

## Example 5-7 Error Message for Temporary Capacity Complete Enforcement

---

To: root@parl.yourorg.com  
Subject: Instant Capacity enforcement notice

This message is being sent to inform you that, due to expiration of temporary capacity, 1 core(s) were deactivated on your Instant Capacity complex (containing the partition parl) to bring the complex into compliance with the Instant Capacity contract.

Prior to deactivation, the number of active cores exceeded the number of available usage rights by 1. 1 core(s) without usage rights were found to be active in the complex. This state was likely the result of having activated Instant Capacity core(s) using temporary capacity (TiCAP) and then allowing the temporary capacity balance to expire prior to deactivation of the core(s).

As a result, the intended active value was reduced by 1 and 1 core(s) were deactivated. To activate these 1 core(s) again, you can perform one of the following actions:

1. Purchase additional temporary capacity, apply the TiCAP codeword(s) to the complex, and use temporary capacity to activate the core(s).
2. Deactivate cores in other partitions. This frees up core usage rights which can be used to activate cores on this partition.

There are currently 3 active core(s) and 3 core usage rights. This complex is now compliant with the Instant Capacity contract.

You can view the current temporary capacity compliance of your system by using the `icapstatus` command.

See the Instant Capacity user's guide at `/usr/share/doc/icapUserGuide.pdf` for more information.

---



---

## 6 Instant Capacity Cell Board

This chapter covers the following topics:

- “Instant Capacity Cell Board” (page 92)
- “Ordering Instant Capacity Cell Board” (page 93)
- “HP-UX and OpenVMS License and Support” (page 94)
- “Acquiring Usage Rights for Instant Capacity Cell Board” (page 95)
- “Instant Capacity Cell Board and Considerations of Core Usage Rights” (page 96)
- “Activation of an Instant Capacity Cell Board” (page 97)
- “Accidental Activation of an Instant Capacity Cell Board” (page 98)
- “Instant Capacity Cell Board Activation Exception Error” (page 99)
- “Instant Capacity Cell Board and Temporary Instant Capacity” (page 100)

# Instant Capacity Cell Board

## Overview

Instant Capacity Cell Board offers a way to have additional (inactive) cell board capacity in your system for growing business needs. When the need arises, you acquire the necessary usage rights in order to activate and use the cell boards, which contain memory and processors or cores.

An Instant Capacity cell board is configured at HP manufacturing already assigned to an nPartition (hard partition) with its `use-on-next-boot` flag set to `n` (no), so it does not participate in the boot of the nPartition.

When you are ready to activate a cell board, you can increase the cell usage rights by either purchasing the appropriate Right to Use (RTU) products, or by borrowing usage rights if you are using Global Instant Capacity to share usage rights within a group of servers. To purchase usage rights, submit a purchase order to HP for the appropriate RTU products to increase the cell usage rights available on the complex, as well as sufficient usage rights for all of the memory on the cell board and, depending on the available usage rights and existing complex configuration, usage rights for one or more additional cores. Then, the cell board is available for activation and participation in the boot of the nPartition. This is controlled by setting the `use-on-next-boot` flag to `y` (yes) with the `parmodify` command and rebooting the nPartition.



---

**NOTE:** If usage rights for the cell board, its memory, and at least one core are insufficient, the Instant Capacity software prevents the cell board from being configured to participate (become active) in the boot of an nPartition.

---

Any cell board, whether or not usage rights are available for activation, can be assigned to an nPartition with the `use-on-next-boot` flag set to `n` (no).

Because an active cell board must have a minimum of 1 active core, prior to activation of a cell board one of the following must be true:

- Usage rights for at least one additional core must be available in the complex. There must be at least one active core per cell board. The Instant Capacity software redistributes active cores across all cell boards in the partition.
- Usage rights for at least one additional core must be purchased and the RTU codeword must be applied to the complex.
- If the complex is a member of a Global Instant Capacity (GiCAP) group, usage rights for at least one additional core must be available from the group.

After a cell board is activated, all of the cores on the cell board can be activated, depending on the availability of core usage rights. You might need to acquire additional core usage rights in order to activate additional cores from the newly activated cell board.

For information about assigning and unassigning a cell board to an nPartition, see “Assigning a Cell to a Partition” (page 67) and “Unassigning a Cell from a Partition” (page 69).

Ask your HP sales representative about availability of the Instant Capacity Cell Board product.

# Ordering Instant Capacity Cell Board

To order the Instant Capacity Cell Board product, do the following:

- Order the appropriate HP product number for the cell board for your specific class of HP server.
- Order the appropriate HP product number for the entire amount of Instant Capacity memory on the cell board.
- Order the appropriate HP product numbers and quantities of Instant Capacity processors for the number of additional cores to activate on the cell board (for details about core usage rights see “Instant Capacity Cell Board and Considerations of Core Usage Rights” (page 96)).



**NOTE:** It is highly recommended that you have the same number of processors or cores and amount of memory on all cell boards in a given hard partition (nPartition). For optimal performance, each nPartition should have cell boards with identical numbers of processors or cores and amounts of memory (otherwise, system performance can be unpredictable).

---

Rules for ordering memory help ensure that the Instant Capacity cell board matches the amount of memory in the other cell boards in a given nPartition.

## HP-UX and OpenVMS License and Support

You do not initially pay for HP-UX and OpenVMS license and support fees on an Instant Capacity cell board.

When you acquire the usage rights for a cell board by purchasing the cell board RTU product, there is an additional cost for the incremental HP-UX or OpenVMS license and support for each core that is activated. That is, the HP-UX or OpenVMS license and support costs are based on a “per active core” basis and are not included as part of the cell board RTU.

If activation of an Instant Capacity cell board does not increase the number of active cores, then you do not have to pay incremental HP-UX or OpenVMS license and support fees.

Your system must be properly licensed for the HP-UX or OpenVMS Operating Environment (OE) when activating the Instant Capacity cell board. You might also need software licenses for nonproprietary software. Consult your application software vendor about licensing requirements.

## Acquiring Usage Rights for Instant Capacity Cell Board

Before activation of an (inactive) Instant Capacity cell board, you must acquire (purchase, or borrow from a GiCAP group) additional usage rights from HP. To purchase additional usage rights:

- Order the appropriate HP RTU product for the cell board for your specific class of HP server.
- Order the appropriate HP RTU product for the entire amount of Instant Capacity memory on the cell board.
- Order the appropriate HP RTU products for the number of additional cores you want to activate on the cell board. This number depends on several factors (and in some cases might not be required); examine this need before ordering any cell board usage rights. For more details about how to determine this number, see “Instant Capacity Cell Board and Considerations of Core Usage Rights” (page 96).

HP sends you a letter that tells you on how to acquire RTU codewords for the purchased components. The letter also describes how to apply the codewords to the system to increase the usage rights on the complex. The steps are as follows:

1. Acquire the appropriate RTU codeword (cell board, memory, core) from the Utility Pricing Solutions portal (<http://www.hp.com/go/icap/portal>).
2. Apply the appropriate RTU codeword (cell board, memory, core) using the `icapmodify -C` command.



**NOTE:** If multiple RTU products are purchased at one time, a single codeword can be generated that incorporates multiple usage rights for the different components.

---



**IMPORTANT:** RTU codewords are based on both the serial number of a system and a unique sequencing value for that server. These codewords must be applied in the sequence in which they are obtained for a particular server. They can be applied to any partition on the server.

---

For an alternative to purchasing usage rights, see [Chapter 7](#) for a discussion of GiCAP and how you can borrow usage rights from other members of a GiCAP group.

# Instant Capacity Cell Board and Considerations of Core Usage Rights

At least one core usage right must be available for an Instant Capacity cell board you want to activate. Each active cell board must have at least 1 active core. However, you do not necessarily need to acquire additional core usage rights. No additional core usage rights are required unless the requirement for a minimum of 1 core per active cell board cannot otherwise be met. If the number of active cores in an nPartition equals (or exceeds) the number of cell usage rights, then you do not need to purchase additional core usage rights.



**NOTE:** The following examples assume that the number of intended active cores for the nPartition remains constant before and after cell board activation.

For example, consider an nPartition with 1 active cell board with all 4 cores active, and 1 inactive cell board with 4 inactive cores. The number of intended active cores for the nPartition is 4. No additional core usage rights are available. Activation of the inactive cell board on the complex results in 2 active cores per cell board after reboot. That is, the Instant Capacity software distributes the available core usage rights across the 2 active cell boards in the partition. The requirement that at least 1 core is active on each cell board is satisfied, so purchase of additional core usage rights is not necessary. You can choose to activate additional cores in the partition. If you do so, you should purchase core usage rights at the same time as the cell usage rights. In this instance, it is not required.

**Table 6-1 Cell Board Activation Not Requiring Additional Core Usage Rights**

State	Active Cell Board Cores	Inactive Cell Board Cores	Notes
Before Cell Board Activation	4 active	4 inactive	No additional core usage rights are available on the complex.
After Cell Board Activation	2 active, 2 inactive	2 active, 2 inactive	No additional core usage rights are required because the number of core usage rights is greater than the number of active cell boards.

In a different scenario, activation of an additional cell usage right can cause the number of core usage rights to be below the minimum (1 active core per active cell board) and necessitate the acquisition of additional core usage rights. For example, consider an nPartition with 1 active cell board containing 1 active core and 3 inactive cores, and 1 inactive cell board with 4 inactive cores. The number of intended active cores is 1 and there are no available core usage rights on the complex. In this case, purchase of an additional cell usage right for the inactive cell board requires purchase of an additional core usage right in order to meet the minimum requirement of 1 core per active cell board.

**Table 6-2 Cell Board Activation Requiring Additional Core Usage Rights**

State	Active Cell Board Cores	Inactive Cell Board Cores	Notes
Before Cell Board Activation	1 active, 3 inactive	4 inactive	No core usage rights are available on the complex.
After Cell Board Activation	1 active, 3 inactive	1 active, 3 inactive	One additional core usage right is acquired because the number of core usage rights is less than the number of active cell boards.

## Activation of an Instant Capacity Cell Board

An Instant Capacity cell board is usually assigned to an nPartition; however, the cell board does not participate in the boot of the nPartition. Activating an Instant Capacity cell board is a two-step process:

1. Set the cell board's `use-on-next-boot` flag to `y` (yes) using the `parmodify` command.
2. Reboot the nPartition (using the `shutdown -r` command on HP-UX), or (on HP-UX 11.31 OEUR 0709 or later) a *Cell Online Activation*.

For example, to change the `use-on-next-boot` flag to `y` on the Instant Capacity cell board in cabinet 0, slot 5, in nPartition 3, execute the following command:

```
/usr/sbin/parmodify -p 3 -m 0/5::y:
```

If core usage rights are available in the complex, the number of intended active cores is increased to as high a number as possible, limited by the number of cores in the newly activated cell board. The available core usage rights are automatically used in the cell activation. If there are no available core usage rights, the number of active cores remains the same.

After you have set the cell board's `use-on-next-boot` flag to `y`, and have rebooted the nPartition, you can use the `icapmodify` command to activate cores that are listed as `Additional cores that can be assigned with current usage rights` (as reported by the `icapstatus` command).

Activating an Instant Capacity cell board causes at least one core to become active on that cell board after reboot.

For details about adding and configuring cells in nPartitions, see the *nPartition Administrator's Guide*.

## Accidental Activation of an Instant Capacity Cell Board

If you accidentally activate an Instant Capacity cell board, you can deactivate it by following this two-step procedure:

1. Set the cell board's `use-on-next-boot` flag to `n` (no) by using the `parmodify` command.
2. Reboot the `nPartition`, or (on HP-UX 11.31 OEUR 0709 or later) a Cell Online Deactivation. This step is not necessary if there was no reboot after the activation.

For example, to change the `use-on-next-boot` flag to `n` on the Instant Capacity cell board in cabinet 0, slot 5, in `nPartition` 3, execute the following command:

```
/usr/sbin/parmodify -p 3 -m 0/5::n:
```

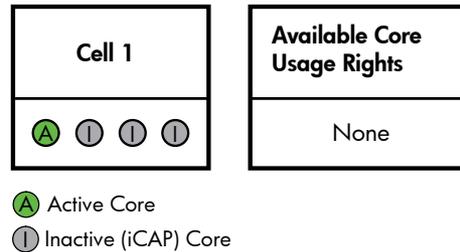
In this command, the `n` flag sets the cell board's `use-on-next-boot` flag to "no" and causes the cell board to not participate in the `nPartition` when it is booted.

For details about adding and configuring cells in `nPartitions`, see the *nPartition Administrator's Guide*.

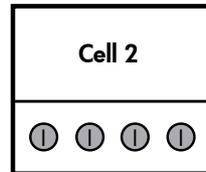
# Instant Capacity Cell Board Activation Exception Error

When you attempt to activate an Instant Capacity cell board in an nPartition, depending on the number of core usage rights that are currently available in the complex, there is a chance the number of intended active cores for the nPartition is out of compliance and the activation fails. Figure 6-1, Figure 6-2, and Figure 6-3 illustrate this.

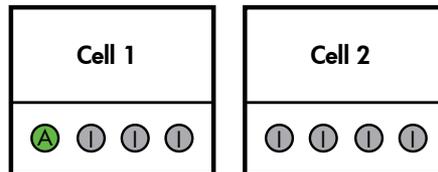
**Figure 6-1 nPartition premodification state: One cell assigned with 1 active core and 3 inactive cores; the complex has no additional core usage rights**



**Figure 6-2 nPartition premodification state: Instant Capacity cell (Cell 2) with 4 inactive cores**



**Figure 6-3 nPartition requested state: Instant Capacity Cell (Cell 2) cannot be activated in nPartition**



In this case, the `parmodify` command fails. This is because the nPartition would have 2 active cell boards, and therefore must have at least 2 active cores. With only one core usage right, the nPartition is out of compliance.

To activate the Instant Capacity cell board and be in compliance, you must first purchase an RTU for an additional core usage right, or deactivate a core in another partition, if possible.

## Instant Capacity Cell Board and Temporary Instant Capacity

You can activate cores only on activated cell boards for which cell board usage rights have been acquired. This is true for both permanent activation of a core and temporary activation of a core using temporary capacity.

To acquire usage rights for an Instant Capacity cell board, you must acquire usage rights for the cell board and the entire amount of memory it contains. For details, see [“Acquiring Usage Rights for Instant Capacity Cell Board”](#) (page 95).

---

# 7 Global Instant Capacity

This chapter covers the following topics:

- “Global Instant Capacity Overview” (page 102)
- “Global Instant Capacity Requirements” (page 104)
- “Global Instant Capacity Group Managers” (page 105)
- “Global Instant Capacity Grouping Rules” (page 107)
- “Global Instant Capacity Sharing Rights” (page 108)
- “Creating Global Instant Capacity Groups” (page 109)
- “Global Instant Capacity Resource Sharing” (page 112)
- “Global Instant Capacity and Temporary Capacity” (page 115)
- “Removing a Global Instant Capacity Group Member” (page 117)
- “Reinstalling a Group Member” (page 118)
- “Group Manager Availability (No Standby Manager)” (page 119)
- “Group Manager Failover Considerations” (page 120)
- “Upgrades and Global Instant Capacity” (page 121)
- “Rights Seizure” (page 122)
- “Considerations for Multiple Groups” (page 127)
- “Additional Considerations” (page 128)

# Global Instant Capacity Overview

Global Instant Capacity, or GiCAP, provides you with the flexibility to move usage rights for Instant Capacity components within a group of servers. It also provides “pooled” temporary capacity across the group. This has several potential benefits: cost-effective high availability, more adaptable load balancing, and more efficient and easier use of temporary capacity. A GiCAP Group is managed using the `icapmanage` command.

Global Instant Capacity provides several benefits:

- **Cost-effective high availability.** In case of planned or unplanned down time, you can transfer usage rights from a failed partition or a failed member complex to one or more other servers in the group that are providing backup availability. Without GiCAP, the only way to provide this *failover* scenario is to provision each server with an adequate amount of temporary capacity.
- **Load balancing.** To provide adaptability and to accommodate changing demands, usage rights can be transferred between servers in a group. For example, a server with extra, unused capacity can release usage rights to activate additional components on an overloaded server that needs extra capacity.
- **Pooled temporary capacity.** Temporary capacity usage rights can be shared across servers for better efficiency and ease of use. By pooling temporary capacity, there is less need to provision temporary capacity for each server.

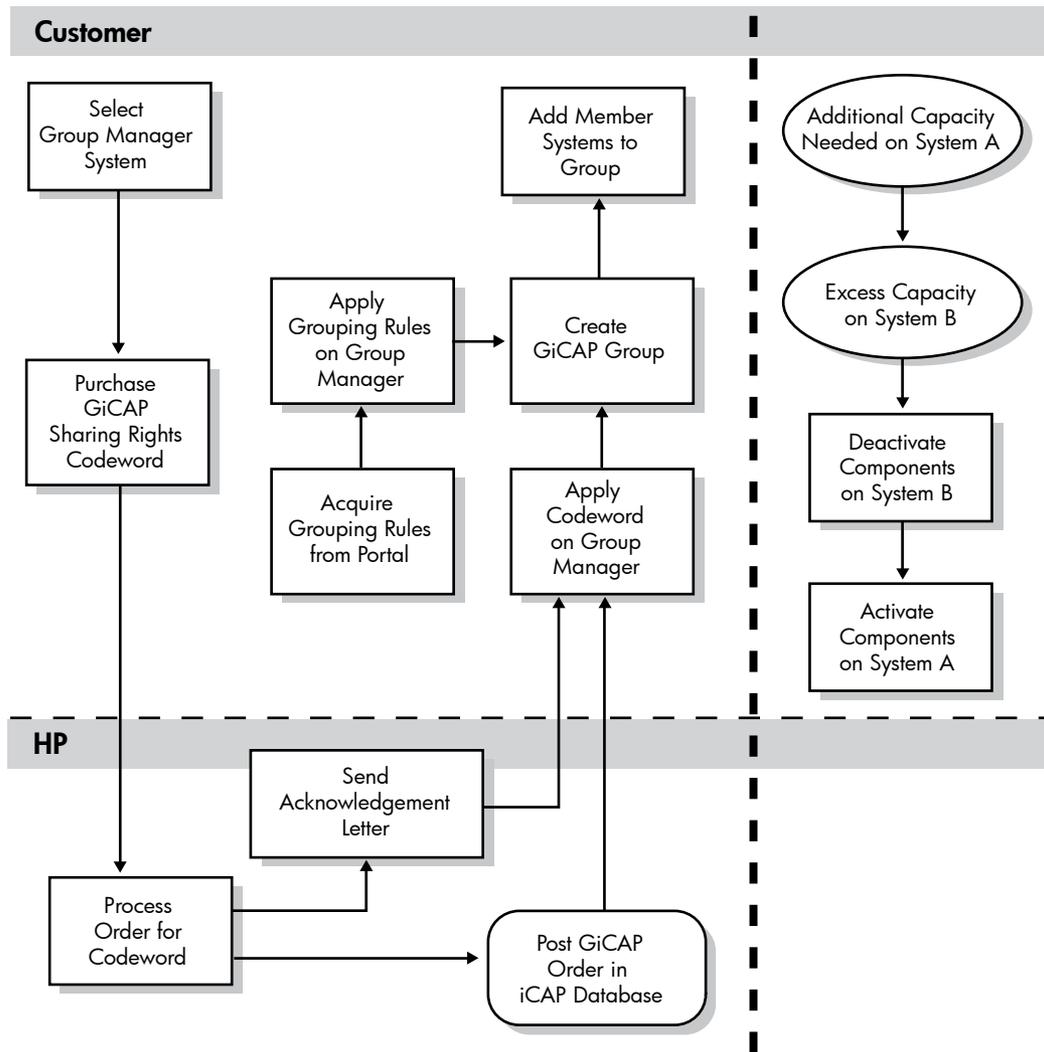
Global Instant Capacity is part of Instant Capacity version 8.x and later on HP-UX systems, and is enabled by purchasing a special GiCAP sharing rights codeword. After purchase, the codeword can be retrieved from the HP Utility Pricing Solutions web portal:

<http://www.hp.com/go/icap/portal>

When you retrieve the codeword, you must provide the sales order number for the codeword purchase, the serial number of the Group Manager system, and partition information for the Group Manager. Applying this codeword on a system running Instant Capacity enables the creation of a GiCAP group with members.

Figure 7-1 illustrates the process of configuring and using Global Instant Capacity.

**Figure 7-1 Using Global Instant Capacity**



## Global Instant Capacity Requirements

To use Global Instant Capacity, the Group Manager and all partitions on all servers in the group must be running Instant Capacity version 9.0 or later. Every GiCAP group member must be hardware-compatible with other GiCAP group members as determined by the GiCAP grouping rules.

WBEM version A.02.05 or higher must be installed on the Group Manager and on all member systems in order to use GiCAP. The CIM Server configuration property `sslClientVerificationMode` must be set to a value of “optional” on all GiCAP Group Managers and on all OS instances of all member systems. (The CIM Server may need to be restarted if the property was not previously set to this value.) For details, see *cimconfig(1M)*. Note that communication between the managers and members of groups is established using SSL certificates that are supplied by the GiCAP software.

To create a GiCAP group with members, you must purchase GiCAP sharing rights, acquire the GiCAP codeword from the HP Utility Pricing Solutions portal, and apply the associated codeword to the Group Manager system. GiCAP sharing rights are described in “Global Instant Capacity Sharing Rights” (page 108).

For a complete list of all software requirements, see “Instant Capacity Requirements” (page 28).

# Global Instant Capacity Group Managers

For each group, an HP-UX system must be designated as an active GiCAP Group Manager. This system maintains information about the group, group resources, and grouping rules. Use the `icapmanage` commands on a Group Manager system only.

The active Group Manager must be an HP-UX system running Instant Capacity software version 9.0 or later. The system running the Group Manager does not require any Instant Capacity components, and it does not need to be a partitionable system. A Group Manager cannot be run on a virtual machine (also known as a guest). The system must have a machine-readable serial number, as displayed by the shell command `getconf CS_MACHINE_SERIAL`.

The Group Manager can be run on either a partitionable or a nonpartitionable system. If run on a partitionable system, changing the configuration of the partitions can result in the Group Manager becoming inoperative. For optimal availability, run the Group Manager on a server or partition that is not a member of any GiCAP group, although it does not need to be dedicated to only the Group Manager software. A Group Manager can manage multiple groups, but it is recommended that it manages a single group for recoverability and ease of use.

## Standby Group Managers

The active GiCAP Group Manager can designate a standby Group Manager to be available for use as an active Group Manager if the current Group Manager is unavailable. The requirements and recommendations defined for the Group Manager also apply to the standby Group Manager; it is a Group Manager with “standby” status. This standby Group Manager is set up using the `icapmanage -a -S` command on the active Group Manager system.

When the `icapmanage -a -S` command is issued, it sets up a `cron` job on the active Group Manager that periodically transfers a copy of the GiCAP database to the standby Group Manager. Prior to activating the standby Group Manager as the active Group Manager, the GiCAP database on the standby Group Manager should match the database on the active Group Manager. Although this is done automatically on a periodic basis, under some circumstances (such as intermittent problems with network communication between the active Group Manager and the standby Group Manager) the database on the standby Group Manager may be out of date. In this case, you might need to manually transfer a copy of the database from the active Group Manager to the standby Group Manager. This is done using the `icapmanage -t` command on the active Group Manager.

The standby Group Manager can take control of GiCAP group management from the active Group Manager using the command `icapmanage -Q` on the standby Group Manager. This allows GiCAP group operations to continue if the GiCAP Group Manager is unable to function. The new active Group Manager attempts to contact all group members, informing them that this system is the active Group Manager. It also attempts to contact the previous Group Manager to make it the standby Group Manager. The current system becomes the active Group Manager regardless of whether or not these attempted contacts succeed or fail.

When the standby Group Manager becomes the active Group Manager, the `cron` job set up on the former active (and now standby) Group Manager is removed, and a new `cron` job is set up on the currently active Group Manager to continue the periodic transfer of a copy of the GiCAP database to the standby Group Manager.

The standby Group Manager is removed using the `icapmanage -r -S` command on the active Group Manager system. The active Group Manager updates all member hosts about the removal, discontinues all GiCAP database transfer operations, and directs the standby Group Manager to remove its copy of the GiCAP database.

The use of the `icapmanage` command is limited on a standby Group Manager. Only the `-Q` and `-s` options are allowed, all other options are restricted to use only on the active Group Manager.

In some circumstances, the cron job set up by the `icapmanage -a -S` command may stop running, perhaps due to another operation on the Group Manager system. The cron entry can be restored by any of the following methods:

- Manually edit the `crontab` file and re-enter the `icapmanage -t` instruction. This is the preferred method.
- Remove and re-add the standby Group Manager.
- Run `icapmanage -Q` on the active Group Manager. This is not recommended because it might have some unexpected side effects.

# Global Instant Capacity Grouping Rules

A GiCAP group consists of a list of server complexes that are allowed to share Instant Capacity usage rights (for cores, cell boards, and memory) and temporary capacity. Other than performance considerations, there are no particular constraints on the number of servers allowed in a group, but there are grouping rules defined by HP to specify the types of servers allowed to group together.



**NOTE:** The Instant Capacity installation does not include a default grouping rules file. Grouping rules must be downloaded from the HP Utility Pricing Solutions portal:

<http://www.hp.com/go/icap/portal>

Grouping rules are defined according to server class. The price structure of usage rights is also based on server class. Since GiCAP pools usage rights, these usage rights can be used on any type of server, regardless of the server class for which they were originally purchased. Therefore, grouping rules were created to define the classes of servers allowed to share usage rights.

Use the `icapmanage -R` command to view the hardware grouping information. When used in combination with a list of host names, this command reports all hardware types with which the hosts might be grouped. If the hosts are not compatible with each other, no hardware is reported. Without a list of host names, this command reports all supported hardware and grouping rules.

Under some circumstances you might need to acquire newer grouping rules from the portal (for example, when adding new hardware not previously covered by the grouping rules currently in use). Install the encrypted rules file on the Group Manager system using the `icapmanage -i` command. The `icapmanage -R` command shows the version number of the grouping rules installed, and when you download a new grouping rules file use `icapmanage -R -U` to compare the version and supported systems before deciding to install the new grouping rules.

The following command installs a grouping rules file, retrieved from the portal, on the Group Manager system:

```
> icapmanage -i -U FSTL012234_gicap.encrypt
```

# Global Instant Capacity Sharing Rights

Although GiCAP is part of Instant Capacity and is installed at the same time as Instant Capacity, it is not enabled during installation. To share resources across groups, you must purchase GiCAP *sharing rights*, acquire the GiCAP codeword from the HP Utility Pricing Solutions portal (<http://www.hp.com/go/icap/portal>), and apply the associated codeword to the Group Manager system. Application of the sharing rights codeword to the Group Manager system enables the addition of members with Instant Capacity components to groups.

You purchase at least as many GiCAP sharing rights as the total number of cores without usage rights across all the potential group members. Members can be added to a GiCAP group as long as there are sufficient sharing rights available, and as long as the grouping rules indicate hardware compatibility.

Unlike other Instant Capacity codewords, GiCAP codewords must be generated for and applied to a specific partition if the Group Manager is on a partitionable system. This means that to retrieve the codeword, you must specify the purchase order number, the system serial number and partition information, if any. Use the `icapmanage -s` command on the Group Manager system to retrieve the serial number and nPar ID, or vPar code that is applicable.

GiCAP codewords have a sequence value and must be applied in the order in which they were generated for the Group Manager system. However, GiCAP codewords are sequenced independently from any other types of Instant Capacity codewords that might be generated for the same system, and can therefore be applied independently from Instant Capacity codewords.

Example 7-1 shows how to apply a sharing rights codeword.

## Example 7-1 Applying a Sharing Rights Codeword

---

```
> icapmanage -C R8J2DBW.5UTxyWQ.2MekJ43.G5cdTVP.1-m9kvweQ.AYqEXym.wj3dyLj.Fbtg7s1
```

The following valid codeword has been applied to the complex:

```
Global Instant Capacity Sharing Rights Codeword  
32 Sharing Rights
```

Use `icapmanage(1M)` to see the results of the application of this codeword.

---

# Creating Global Instant Capacity Groups

After the sharing rights codeword and the grouping rules have been applied to the Group Manager, a GiCAP group can be created by issuing the `icapmanage` command using the `-a`, `-g`, and `-m` options. Use the `-a` option to add members, the `-g` option to select the group name, and the `-m` option to specify a name for the new member along with a list of hosts running on the system. The list of hosts must include at least one host per nPartition or virtual partition on the system.

Note that a single partition of a complex cannot join a GiCAP group; rather you must specify all partitions of a complex when adding a group member. All partitions on a group member must be running HP-UX. An Instant Capacity server can join a group if the Group Manager has at least as many GiCAP sharing rights as the total number of cores without usage rights on that server. Members can be added to a GiCAP group as long as there are sufficient GiCAP sharing rights available and it is permitted by the grouping rules. Each member that joins the group decreases the available GiCAP sharing rights by the number of cores without usage rights contributed by that member complex.

Although the size of GiCAP groups is not specifically restricted, performance of group-related functions is affected by the number of group members and the number of partitions for each member server, as well as by the types of hardware involved. A larger number of group members can cause an increase in startup time for the Group Manager and can also affect the performance of `icapmodify` commands when a transfer of usage rights occurs. If temporary capacity is being used, then the size of the group may also increase the amount of communication time needed for tracking of temporary capacity.

When adding groups to a Group Manager system, you can use the `icapmanage -T` command to test hardware compatibility for one or more host systems in order to determine which groups the systems can join. When used in combination with the `-g` option to specify a group name, this command tests whether the specified host systems have hardware that is compatible with the group. Without the `-g` option, this command reports which groups of all the groups managed by this Group Manager have hardware which is compatible with the host systems. The host names do not have to be from the same complex, but in order to best predict the possibility of being able to join a group, the list of hosts should include all the nPartitions for a particular complex. If the hosts are not compatible with each other, no groups are reported as having compatible hardware.

You can create multiple GiCAP groups that can be managed by the same Group Manager or by different Group Manager systems. Systems with no any Instant Capacity components can be part of a GiCAP group. Deactivating resources on these systems allows them to loan usage rights to other members in the group.

Example 7-2 shows how to create a group and show group status.

## Example 7-2 Creating a Group

---

```
> icapmanage -a -g one
```

```
Group one added.
```

```
> icapmanage -s
```

```
Software version:          B.09.00
32 GiCAP Sharing Rights: 0 in use, 32 available
Group ID: one
Group Members:
    No members found
```

---

Example 7-3 updates the grouping rules for all groups managed by the Group Manager, tests whether a server complex has hardware compatible with group “one”, and adds a member called

“IT” to that group. When you first add new members to a group, you are prompted for the root password for each specified host. The password is used only to establish secure communication and is not saved or stored.

### **Example 7-3 Adding a Member to a Group**

---

```
> icapmanage -i -U /tmp/GiCAP.rules
> icapmanage -T node.corp.com -g one
root@mypar.node.corp.com's password:
The server(s) are compatible with GiCAP group one

> icapmanage -a -m IT:node.corp.com -g one
Member IT added to group one.
```

---

Example 7-4 shows the output of the `icapstatus` command for a group member system.

## Example 7-4 Example Output of icapstatus for a Group Member System

> /usr/sbin/icapstatus

Software version: B.11.31.09.00.00.71  
System ID: node  
Serial number: USR4020003  
Product number: A6093A  
Unique ID: Z3e0ec8e078cd3c7b  
System contact e-mail: mjones@corp.com  
From e-mail: Set to the default ('adm')  
Asset reporting: on  
Temporary capacity warning period: 15 days  
Exception status: No exception

Member zoo6 in GiCAP group MyGroup

-----  
Active Group Manager: node1.corp.com  
Standby Group Manager: node2.corp.com  
Borrowed core usage rights: 0  
Borrowed cell usage rights: 0  
Borrowed memory usage rights: 0.0 GB

Local nPartition Status (09/17/08 12:34:56)

-----  
Total number of configured cores: 8  
Number of Intended Active cores: 2  
Number of active cores: 2  
Number of inactive cores: 6  
Additional cores that can be activated with current usage rights: 0  
Number of cores that could be activated with additional usage rights: 6  
Number of cores that can be activated with temporary capacity: 6  
Number of cores currently unavailable for assignment: 0

Instant Capacity Resource Summary

-----  
Number of cells without usage rights: 0  
Number of inactive cells: 0  
Amount of memory without usage rights: 0.0 GB  
Amount of inactive memory: 0.0 GB  
Number of cores without usage rights: 6  
Number of inactive cores: 6  
Number of cores using temporary capacity: 0  
Temporary capacity available: 60 days, 0 hours, 0 minutes  
Projected temporary capacity expiration: N/A

Allocation of Instant Capacity Resources among the nPartitions

-----  
                  Intended Actual  
nPar  Total  Active  Active  =====Inactive=====  Runs  
ID    Cores  Cores  Cores  Cores  Memory  Cells  iCAP  nPar  Name  
====  =====  =====  =====  =====  =====  =====  =====  =====  =====

0	8	# 2	2	6	0.0 GB	0	Yes	zoo0	(local)
1	8	8	8	0	0.0 GB	0	Yes	zoo1	
N/A	0	N/A	N/A	N/A	0.0 GB	0	N/A	Unassigned	Cells

# Value for Intended Active not specifically set and defaults to the number of configured cores in active cells.

---

# Global Instant Capacity Resource Sharing

Once a group is established, Instant Capacity resources (core, cell board, memory usage rights, and temporary capacity) can be shared among all the members of the group. Resource sharing can occur in several ways:

- During creation of the group, some members might have unused usage rights, so that by simply joining the group, additional usage rights are available for use by any member of the group.
- Even if there are no unused usage rights across the group, a member of the group can deactivate resources (cores, cells, or memory) to make additional usage rights available for activation by any other member in the group.
- Temporary capacity from all members of the group is available for use by any member of the group.
- *Rights seizure* allows the sharing of usage rights from failed partitions or failed servers.

Usage rights are shared by deactivating resources on one group member and then activating the resources on another member of the group. The system on which the resources are deactivated is lending usage rights to the activating (or borrowing) system. Activations and deactivations are performed on the individual member systems using the usual `icapmodify` commands (for cores) or `parmodify` commands (for cells) to effect this “loan” operation (also sometimes referred to as a transfer of usage rights).

Any temporary capacity available to individual members of the group is combined into a larger pool of temporary capacity that is available for consumption by any and all members of the group, as needed. Initiating usage of shared temporary capacity is the same as with individually purchased TiCAP: group members use the `icapmodify -a -t` command to activate shared temporary capacity. This differs from the sharing of usage rights in that temporary capacity is never a “loan” to be returned; it is always depleted through usage over time.

## Example: Core Rights Sharing

In this scenario, core usage rights are not immediately available to any member of the group `mygroup`. Group member `member1` has an immediate need for more processing power. However, group member `member2` can loan a core usage right by deactivating one core.

First, `member2`, currently with 8 active cores, deactivates one core:

```
member2> icapmodify -d 1
7 cores are intended to be active and are currently active.
```

The core usage right from `member2` is now available to any member of the group and can be used to activate an additional core on `member1`:

```
member1> icapmodify -a 1
8 cores are intended to be active and are currently active.
```

The output of the `icapstatus` command on loaning system `member2` shows that the Number of Intended Active Cores and Number of active cores have decreased by 1, and the Number of inactive cores and Number of cores without usage rights have increased by 1. On borrowing system `member1`, the Number of Intended Active Cores and Number of active cores have increased by 1, and the Number of inactive cores and Number of cores without usage rights have decreased by 1.

The output of `icapmanage -s` on the Group Manager system shows that the total number of cores without usage rights for the group has not changed.

## Effect of Temporary Capacity

On systems where usage rights and temporary capacity are available, Instant Capacity uses usage rights before it uses temporary capacity. If temporary capacity is being used on at least one member system, a component on another member is deactivated, and a component on a third

member system needs to be activated, the usage rights made available by the deactivated component can be taken by the system using temporary capacity. In this case you might need to use the `-t` option to the `icapmodify` command to activate the component on the third member system by using temporary capacity.

Whenever you have more active cores than the number of core usage rights, the temporary capacity balance is depleted as a mechanism for tracking noncompliance of the group, even if TiCAP has not been purchased for or applied to any member of the group. This differs from the behavior of TiCAP on a complex which is not a member of a group, where TiCAP is decremented only if TiCAP had been specifically purchased for the complex. Within a GiCAP group, temporary capacity is used as an additional compliance mechanism to support the high availability features of a group.

Because group members are automatically considered to be users of temporary capacity, to avoid unexpected TiCAP depletion in a group, it is important to avoid the situations that cause the Instant Capacity software to make assumptions that all cores might be active on a remote nPartition.

If a member is removed from the group, the TiCAP balance on that complex will continue to be used as a compliance mechanism (decremented whenever the number of active cores exceeds the number of core usage rights), unless the TiCAP balance on the system is exactly 0.

## Status Reporting

Usage rights and temporary capacity can sometimes be temporarily assigned to the Group Manager, which can result in difficulty interpreting some of the data from the `icapstatus` command. The total temporary capacity reported for the group by `icapmanage -s` might not equal the sum of temporary capacity reported by each member system. This is because the Group Manager prefetches an amount of temporary capacity in anticipation of needing it for a future operation, so the temporary capacity might not be immediately assigned to a member system. Also, individual counts of cores, cells, and memory without usage rights for each member of the group might not add up to the total counts of cores, cells, and memory without usage rights for the group. In all cases, totals reported (by `icapmanage -s`) for group temporary capacity and usage rights are the important values that represent available resources for the group. Use the verbose (`-v`) option to the `icapmanage -s` command to see resources held by the Group Manager.

The `icapstatus` command reports usage rights borrowed or loaned by the system from or to other members of the group. Borrowed rights are those that are currently resident on the member and that originated elsewhere. Loaned rights originate on the system but are currently resident somewhere else in the group. They might be either in use or unassigned on another group member, or they might be unassigned on the Group Manager itself. The `icapmanage -s` command displays the status of the entire group, while the `icapstatus` command displays an isolated view of a single member.

Unassigned usage rights are free to be moved to other member systems, even if a particular activation request fails. For example, consider a group where there are no free usage rights. If member `m1` releases 2 core usage rights by deactivating 2 cores, and member `m2` tries to activate 3 cores, the activation request fails (not enough usage rights), after the Group Manager has already migrated 2 core usage rights from `m1` to `m2`. The borrow and loan values for `m1` and `m2` show that the loan has occurred, even though the activation on `m2` failed to activate any cores. A subsequent activation of 2 cores on `m2` will be successful and will occur more quickly because the 2 core usage rights are already assigned to `m2`, or because the rights are moved elsewhere in the group if requested.

## Example: Cell/Memory Sharing

In this scenario, member1 of the group mygroup has an inactive cell it wants to activate, but no usage rights are available on the system. However, member2 of the group has available usage rights.

First, the output from the icapstatus command on member1 shows that no cell or memory usage rights are available:

```
Instant Capacity Resource Summary
-----
Number of cells without usage rights:          1
Number of inactive cells:                     1
Amount of memory without usage rights:        16.0 GB
Amount of inactive memory:                    16.0 GB
Number of cores without usage rights:         8
Number of inactive cores:                     8
```

The output of the icapstatus command on member2 shows that memory and cell usage rights are available, along with 1 core usage right:

```
Instant Capacity Resource Summary
-----
Number of cells without usage rights:          1
Number of inactive cells:                     2
Amount of memory without usage rights:        16.0 GB
Amount of inactive memory:                    32.0 GB
Number of cores without usage rights:         7
Number of inactive cores:                     8
```

In this situation, it is not necessary to deactivate components on member2 because the system is underutilized and usage rights are available. The usage rights from member2 are available to member1, and the cell can be activated:

```
member1> parmodify -p 0 -m 2::y:
```

After a reboot or (on HP-UX 11.31 OEUR 0709 or later) a Cell Online Activation, all of the cell and memory components on member1 are active.

## Global Instant Capacity and Temporary Capacity

Temporary capacity can be shared across servers for better efficiency and ease of use. Temporary capacity within a GiCAP group is always available to all members of a group without the need to purchase temporary capacity for each server. You can exercise some control over how “willing” each GiCAP member system is to share temporary capacity by setting its “temporary capacity warning period”. No member's temporary capacity balance is decreased below the member's warning period until all members of the group are within their warning period.

### Example: Activation Using Pooled Temporary Capacity

In this scenario, `member1` of `mygroup` has 2 active cores and needs to activate 6 more cores, but only 5 usage rights are available in the group. There is no temporary capacity available on `member1`. Other members of `mygroup` have a sufficient amount of temporary capacity, so you can activate the cores using temporary capacity:

```
member1> icapmodify -a 6 -t
```

```
8 cores are intended to be active and are currently active.
```

```
Number of cores using temporary capacity:          1
Projected temporary capacity expiration: Less than 30 minutes
```

Notice that 5 additional cores are permanently activated with the available usage rights, and only the last core is activated with TiCAP. Initially, only 30 minutes of temporary capacity are transferred to `member1`, since 30 minutes of temporary capacity are transferred per core activated with temporary capacity. Every 30 minutes the daemon determines whether temporary capacity is depleted and acquires more from the group as needed.

## Temporary Capacity and Freed Usage Rights

When a complex consumes temporary capacity, the Instant Capacity daemon periodically decrements a complex's temporary capacity balance. Before doing so it contacts the Group Manager to determine whether there are core usage rights available on other group members. If not, temporary capacity continues to be consumed. If usage rights are available anywhere in the group, they are transferred to the complex using temporary capacity in order to stop temporary capacity consumption on that complex.

Between the time the core usage rights are made available and the Instant Capacity daemon monitors temporary capacity consumption, the `icapstatus` command reports that temporary capacity is being consumed when it is not. When the transfer of usage rights is completed, the `icapstatus` output is updated on both systems to reflect the transfer. Because of the delay, the changes might appear to be unrelated to a user-initiated operation, but they are due to the previously initiated deactivation that freed up core usage rights.

## Temporary Capacity and Status Reporting

The temporary capacity balance reported by `icapstatus` on a group member reflects only the temporary capacity that has been applied or transferred to that system via the Group Manager. You might still receive temporary capacity expiration warning messages even though more temporary capacity is available in the group.

Temporary capacity is transferred to group members in 30-minute blocks. When a block of temporary capacity has been consumed, the Group Manager continues to transfer group temporary capacity to the system every 30 minutes as long as it is available. However, the `icapstatus` command on the local system might report temporary capacity as expired even though it is still being used to activate cores, as shown in the `icapstatus` listing of “Number of cores using temporary capacity”.

## Temporary Capacity Prefetch

Because temporary capacity is pooled for the group, adjustments to the temporary capacity balance can be made even when it is not being consumed. For performance reasons, the Group Manager anticipates potential future use of temporary capacity and might prefetch an amount of temporary capacity from one or more member systems. Although temporary capacity will not be used on a member system unless the `-t` option was specified with the `icapmodify` command, an `icapmodify` command without the `-t` option might still result in an adjustment of the temporary capacity balance for members of the group. When this happens, the overall temporary capacity balance for the group does not change, but there might be differences in the allocation to individual member systems.

## Removing a Global Instant Capacity Group Member

Before removing a member from a GiCAP group, all the borrowed usage rights must be returned and all outstanding loans reclaimed. Do this by deactivating resources on the member about to be removed. Loaned usage rights are reclaimed by deactivating enough resources elsewhere in the group to cover the loan. The reclamation of loaned usage rights on the member about to be removed does not require the activation of resources on that member. As long as the usage rights are not in use elsewhere in the group, removing the member results in the member reclaiming its loans.

There is no constraint on temporary capacity because a removed member will take with it whatever temporary capacity is assigned to it.

When a member is removed from a group, some number of sharing rights are released and made available for future use. The number freed is equal to the number of cores without usage rights on the removed member.

The following commands remove member IT from its group, and then remove group one:

```
> icapmanage -r -m IT  
Member IT removed.
```

```
> icapmanage -r -g one  
Group one removed.
```

A group cannot be removed until all members in that group have been removed.

## Reinstalling a Group Member

If a group member is reignited, use the `icapmanage` command as follows to reestablish communication between the newly ignited group member and the Group Manager:

```
icapmanage -a -m <member>:<host> -g <group_name>
```

See the *HP Instant Capacity Release Notes* for more details about reestablishing SSL certificates.

## Group Manager Availability (No Standby Manager)

If the active Group Manager becomes unavailable and a standby Group Manager is not defined or not used, management of the GiCAP group is unavailable until the Group Manager is restored or replaced. The GiCAP group members continue to operate as isolated Instant Capacity systems, using whatever usage rights and temporary capacity they had available when the Group Manager became unavailable. A GiCAP group member using borrowed usage rights can continue using those usage rights. A GiCAP group member that has loaned usage rights to other members in the GiCAP group cannot recover those usage rights until the Group Manager is restored.

If a standby Group Manager is not used, loaned and borrowed usage rights remain in place in GiCAP group members when the active Group Manager is unavailable. All usage rights held by the Group Manager itself are left unavailable until the Group Manager is restored.

Normally, usage rights are not held by the Group Manager, but this can occur after a member is removed from a GiCAP group, or if the Group Manager becomes unavailable during the transfer of usage rights from one GiCAP group member to another.

## Group Manager Failover Considerations

If the active Group Manager system becomes unavailable and a standby Group Manager has previously been defined, the standby Group Manager can be used to take over GiCAP group operations from the Group Manager. If both the active Group Manager and the standby Group Manager are unavailable, or if the active Group Manager fails and the administrator chooses not to have the standby Group Manager take over group operations, then the usage rights and temporary capacity remain as per allocated to each group member, as described in the previous section “Group Manager Availability (No Standby Manager)” (page 119).

However, if a standby Group Manager has been defined, an administrator can have a standby Group Manager take control at any time using the `icapmanage -Q` command. While this can be done routinely, for example to allow shutting down a functioning active Group Manager for maintenance, normally this command is issued either when the active Group Manager has failed, or when a network outage has made it unable to communicate with critical group members. When a standby manager is told to take control, it attempts to update all members and the current active Group Manager so that group operations can proceed smoothly.

However, in the case of a failure, it is possible that the `icapmanage -Q` command is unable to contact the active Group Manager and some members of the groups that it now manages. When this happens, the previously active Group Manager remains active, unaware of the change of control. This is referred to as a “bifurcated” (or “split”) GiCAP group. Members that were reachable by the standby Group Manager when it took control cannot accept commands from the old active Group Manager; but unreachable members continue to consider it active. Control operations can be carried out on both active Group Managers, each communicating with the members that it (and only it) can reach. Groups and members can be added or removed on both (subject to the set of members each can command), and sharing rights can be added on both. In some cases this can be valuable; for example, when two data centers each remain functional but some intervening network link has been broken. Each isolated set of systems can proceed with independent disaster recovery operations within their group subset.

At some point, communication is restored and the split groups must be rejoined. This is accomplished through issuing new `icapmanage -Q` command. It can be executed on either active Group Manager to confirm that Group Manager as the active Group Manager and demote the other to standby status. Be aware that doing this loses all database changes made on the demoted Group Manager during the time that the group was split. There is no method to merge the two databases, and in particular any new sharing rights applied to the Group Manager designated now as standby are lost.

## Upgrades and Global Instant Capacity

Be careful before upgrading or changing hardware or operating systems for any member of a GiCAP group. If a member of a GiCAP group changes hardware in such a way that the hardware is no longer compatible with the group, then the group is considered to be out of compliance and group functions are restricted.

Also, the number of available sharing rights is adjusted whenever an Instant Capacity codeword is applied to a GiCAP member system that modifies the number of cores without usage rights on that member. (RTU and AddOn codewords for cores cause such adjustments.)

If more sharing rights are in use than were purchased for the Group Manager, all groups managed by that Group Manager are out of compliance and all group functions are restricted until the problem is resolved. The problem can be resolved by purchasing and applying additional sharing rights to the Group Manager, purchasing and applying core usage rights to one or more group members, or removing one or more group members from their group.

When such an incompatibility is detected, the GiCAP Group Manager sends email to the local root account and to the registered contact email address for each member of the group.

## Adding New Partitions

When reconfiguring a member system by adding or deleting an nPartition, use the `icapmanage -u` command to add or remove host names to or from the member. For example, enter this command to add a new nPartition with hostname `hostC` to member `memberA`:

```
icapmanage -u -m memberA -h hostC
```

To remove a host name, prefix the host name with “!”, as in `!hostC`. You can also add or remove a list of hosts; separate the additions from the removals with a “!”, as in `hostA,hostB!hostC,hostD` to add `hostA` and `hostB` while removing `hostC` and `hostD`.

If the Group Manager runs on a partitionable system, changing the configuration of the partitions may render the Group Manager inoperative. This is because the Group Manager database is tied to the serial number of the complex, the nPartition ID number, and a virtual partition ID number. Although adding or modifying partitions does not change these IDs, they are changed if a partition is deleted and then re-created on the Group Manager.

## Rights Seizure

GiCAP disaster recovery was introduced in Instant Capacity version 8.02.01 and provides the ability to “seize” core usage rights from a GiCAP member that is off line because of some disaster, and transfer them to the Group Manager. Then, using normal activation commands, these usage rights can be used to activate additional processor cores on other group members to increase capacity. Rights seizure can be used to provide disaster recovery when all partitions of a GiCAP member fail. It also enables the restoration of usage rights to the member once contact has been restored.

When a failure occurs on a partition in an active group member, use the `icapmanage -x` command to acquire core usage rights from the specified host to make them available to other group members. This is known as rights seizure. The specified host must be known to the GiCAP Group Manager (it appears in the output of the `icapmanage -s` command) and not currently running. That and any other (vPar) hosts associated with the same hard partition are verified to be unreachable by the GiCAP Group Manager. The `icapmanage -x` command verifies with each known host on the hard partition. The hard partition containing the specified host is left with one core usage right per active cell. Any core usage rights in excess of this are made available for use elsewhere in the GiCAP group.

The seizure of core usage rights from a fully unavailable member is temporary. After 10 days, the usage rights are automatically restored to the member from which they were seized. Seizure of core usage rights from a member with at least one accessible partition (at the time of rights seizure) do not expire. The `icapmanage -z` command allows you to restore previously seized core usage rights to a specified host.

The following command seizes core usage rights from a partition that is unavailable, so that the rights are available for other group member activations:

```
icapmanage -x mypar1.node.corp.com
```

## When to Migrate Usage Rights and When to Seize Usage Rights

Here are some basic guidelines for determining when to seize usage rights or when to simply *migrate* those rights:

- Planned downtime and load balancing

Whenever possible, migrate usage rights by deactivating cores in one partition and then activating cores in another partition. The involved partitions do not need to be part of the same member server. For example, if maintenance is planned such that a partition is unavailable for a period of time, deactivate cores in that partition before it becomes unavailable. The usage rights from this partition are available to the entire GiCAP group during this maintenance period.

- Unplanned downtime

When an nPartition or server goes down unexpectedly, seize the usage rights from that nPartition to make them available to the remaining GiCAP group members.

## Effects of Rights Seizure

Rights seizure takes almost all the available usage rights from the hard partition containing the specified host, leaving only enough usage rights to be able to boot (one core usage right for each cell configured for use-on-next-boot (UONB)). The partition has the value `intended active` set to the required minimum (one core per configured cell) and the value `actual active` set to zero. The number of core usage rights seized is equal to the difference between the new value for `intended active` and the greater of the old values for `intended active` and `actual active`. These seized core usage rights are available for use elsewhere in the GiCAP group.

Although rights can be seized from any hard partition that is unavailable, the Instant Capacity software makes some additional restrictions when all partitions of a complex are unavailable.

As a result, there are different behaviors and constraints depending on whether any partitions can be contacted on the specified member complex.

If, at the time of rights seizure, at least one member partition can be contacted, then the software is able to make an immediate adjustment to the available core usage rights, just as if a normal migration operation using `icapmodify -d` had been performed before the specified hard partition stopped running. This makes core usage rights available for potential loans to other member systems. In this situation, the seized core usage rights do not have an expiration date. However, because other member partitions are running Instant Capacity software, you can assume that the unreachable partition is using all cores on cells configured for that partition. Unless an explicit restore operation is performed, when the failed partition is rebooted it will have only the minimum number of core usage rights left after the rights seizure. Because of this, cells in partitions from which usage rights have been acquired should be rebooted or made inactive within 12 hours. If this is not done, the partition may begin to consume temporary capacity. If temporary capacity is not available, the complex may no longer be in compliance with the Instant Capacity contract. Cells may be made inactive by removing them from the partition, shutting down the partition from within the OS by using `shutdown -R -H`, or with the `MP RR` command.

If, at the time of rights seizure, all member partitions are unreachable, the rights seizure is deferred and must be viewed as a limited and immediate loan of usage rights from the specified partition to the group. This loan of seized usage rights expires in 10 days. Upon expiration, usage rights are automatically restored to the member partitions from which they were seized. The expiration date for a rights seizure operation effectively terminates the period during which the core usage rights are available to other group members for purposes of disaster recovery. If none of the member partitions are reachable by the expiration date for a particular member, the usage rights are automatically restored (reassigned) to the member partition (or complex, in the case of unassigned seized rights) from which they were seized. However, note that if the seized usage rights have been redeployed to other members and are not released at the expiration time, the group may go out of compliance or temporary capacity may be used to maintain compliance.

If any partition of the inaccessible member from which rights seizures were deferred reconnects to the group before the expiration date, the seized core usage rights (for all partitions) are finalized as a loan from the member to the group, the expiration date is no longer relevant, and the usage rights can thereafter be manipulated with normal `icapmodify` operations.

The `icapmanage -x` operation can be performed once for each hard partition on the member.

## Down Partitions with Powered-On Cells

Partitions that are not running an Instant Capacity daemon are assumed to be running another OS and using all cores on cells configured in the partition. The Instant Capacity software can avoid this assumption when all cells configured in the partition are powered down. Unless an explicit restore operation is performed, when the failed partition is rebooted it will have only the minimum number of core usage rights left after the rights seizure. Because of this, cells in partitions from which usage rights have been seized should be rebooted or made inactive within 12 hours. If this is not done, the partition may begin to consume temporary capacity. If temporary capacity is not available, the complex may no longer be in compliance with the Instant Capacity contract. Cells can be made inactive by removing them from the partition, shutting down the partition from within the OS by using `shutdown -R -H`, or with the `MP RR` command.

## Temporary Capacity and Rights Seizure

Instant Capacity is designed to stop using temporary capacity and instead take advantage of usage rights if any become available. If temporary capacity is in use during a failover sequence, the activation on the failover node might need to specify the use of temporary capacity. After the rights have been seized and before the activation on the failover node, there is a window of time when the iCAP daemons (on other partitions in the group) can wake up and start using the

seized usage rights in order to stop using temporary capacity. If this happens, the seized usage rights are no longer available for the failover activation.

By specifying the use of temporary capacity on failover activation, you guarantee that the core activations needed for failover will occur. The total temporary capacity consumption across the group remains the same, even though the temporary capacity might be consumed on the failover server instead of on the original server.

If, for some reason, it is important to keep the temporary capacity usage to a particular server, you can manually deactivate the temporary capacity-consuming cores before doing the failover activations, and then reactivate the temporary capacity usage after the failover activations. For similar reasons, it might also be important to make sure other activations are not occurring simultaneously during a failover sequence.

## Other Considerations

Rights seizure can be used as part of an automatic failover system, but be sure that resources are seized appropriately and in a manner that does not cause problems when the problem is corrected. The Instant Capacity software determines that the partition is down based on whether the `ping` command is unsuccessful for the partition. In some circumstances, `ping` might be unsuccessful but the system remains functional (for example, if a network connection is interrupted). In this case, rights seizure may be inappropriate and leave workloads without the necessary resources.

There are no restrictions regarding the type of environment from which rights were seized. Rights that are seized from either an nPartition environment or a virtual partition environment can be deployed on any member of the GiCAP group, regardless of the target environment type.

## Additional HA Solutions

Although rights seizure is the main method used in a GiCAP *high availability* (HA) scenario, there are additional methods and configurations allowing GiCAP to provide you with HA solutions:

- You can use GiCAP to move capacity from one or more nonproduction servers, such as test servers, during a failover situation. You have a set of standby servers that are part of a group and can pool their resources to provide failover capability.
- You can combine GiCAP with temporary capacity. If you have a set of nonproduction servers, and some of those servers contain temporary capacity, temporary capacity from the entire group is made available during a failover situation. Temporary capacity is not seizable, so any HA solution involving temporary capacity must ensure that it is available on the standby server.
- Systems with full usage rights can also be a part of a GiCAP group and can be used as donor systems, contributing usage rights to the group and allowing additional activations on member systems with Instant Capacity components.

## Summary of Rights Seizure

The following summarizes usage rights seizure on GiCAP systems:

- iCAP can seize usage rights from an unavailable server.  
Usage rights can be seized from a group member even if the entire server is unavailable. In this case, usage rights expire after 10 days and revert to the original partition.
- Only core usage rights can be seized.

While GiCAP allows migration of all types usage rights between member servers (cores, cells, and memory), seizure from a failed partition is only for core usage rights. If cell or memory usage rights are needed, partitioning tools such as the `parmodify` command can be used either remotely or on an accessible partition to remove one or more cells from the unavailable partition, making usage rights available. Changing the use-on-next-boot (UONB)

flag to false for an inactive cell also releases the cell usage right and associated memory usage rights.

- Maximum core usage rights are seized.

There is no way to specify the number of core usage rights that are seized. Instant Capacity seizes the maximum number possible, while ensuring that the nPartition can still be booted. This seizure results in a partition with one core usage right for each active cell upon reboot of the server.

- Expiration of seized usage rights (complete server failure)

If at the time of an attempted rights seizure all member partitions are unreachable, the right seizure is instead treated as a loan of usage rights from the specified member to the group. The loan expires 10 days from the first use of the `icapmanage -x` command. If none of the member partitions are reachable by the expiration date for a particular member, the usage rights are automatically restored (or reassigned) to the member from which they were seized. If, however, one of those member partitions reconnects to the group before the expiration, the loan can be committed, with the usage rights actually transferred from the member.

If a system cannot reconnect to the Group Manager within the 10-day period, you can extend the expiration of the seized usage rights. This extension is applied to the Group Manager through the use of a codeword obtained from HP support.

- Core usage rights are always seized from the nPartition, even when the specified host is a virtual partition.
- Usage rights can be seized from an nPartition only when:
  - The nPartition is down.
  - Assuming the nPartition is hosting virtual partitions, all the virtual partitions are down.

In either of these cases, rights seizure is the only technique available for migrating core usage rights from the nPartition.

- Because rights are seized at the nPartition level, when an entire server is unavailable, to seize all rights from the server, the `icapmanage -x` command must be run on each nPartition in the server.
- An nPartition must be unavailable for usage rights to be seized.

Usage rights can be seized only if the partition is unavailable as determined by the `ping` command.

- Partial server failure

If, at the time of rights seizure at least one nPartition of the server is still running, Instant Capacity can make an immediate adjustment to the available core usage rights, and the seized core usage rights do not have an expiration date. However, because there are other nPartitions running Instant Capacity software, the unreachable partition might be assumed to be using all cores on cells configured for that partition, resulting in possible noncompliance or temporary capacity usage. To avoid this, cells in partitions from which usage rights have been seized should be made inactive within 12 hours.

- When seizing usage rights from an nPartition that contains virtual partitions:
  - All the virtual partitions must be down before rights seizure is allowed.
  - Usage rights cannot be restored to the virtual partitions before they boot. Otherwise, you may not be able to reboot certain virtual partitions, depending on the vPars database definition and the allocation of usage rights among the virtual partitions. Run the restoration command `icapmanage -z` specifying any virtual partition of the nPartition from which rights were seized. Failure to restore usage rights to the virtual partition before connecting to the Group Manager can leave the virtual partition in a nonbootable state, requiring additional steps to fix the vPars database before the virtual partitions can be booted.

- If there are multiple VM guests for the hard partition, they will be affected by the reduction of usage rights. In an HP Integrity VM environment, specify rights seizure only for the VM host, not the guests.
- The Group Manager must have network connectivity.  
If the Group Manager is also unavailable, and a standby Group Manager is not defined, it is not possible to transfer usage rights to a standby server.

For additional information about rights seizure, see the *Cost-effective high-availability solutions with HP Instant Capacity on HP-UX* white paper on the HP Documentation website:

<http://docs.hp.com>

## Considerations for Multiple Groups

You can create multiple GiCAP groups, and they can be managed by the same Group Manager or by different Group Manager systems. Note that if a Group Manager has an associated standby Group Manager, the standby Group Manager functions as a standby for all the groups managed by that Group Manager.

A server complex can be a member of only one GiCAP group at a time. To participate in a different group, it must be removed from one group before being added to the other group.

Sharing rights can never be transferred between two active Group Manager systems. As you create new groups or add new members to existing groups, you might need to purchase and apply additional sharing rights to the relevant active Group Manager systems. This is necessary even if the member has been moved from another Group Manager that now has excess sharing rights.

Sharing rights can never be applied to a Group Manager with standby status. If a standby Group Manager is requested to take control, the sharing rights of the former active Group Manager move to it. If additional sharing rights need to be applied before failing back to the original Group Manager, they must be purchased specifically for the new active Group Manager system (formerly the standby Group Manager) and acquired from the portal. Once applied, the new total of sharing rights moves to the originally active Group Manager if it is requested to retake control, although only if the new active Group Manager is able to propagate its GiCAP database to the originally active Group Manager.

## Additional Considerations

Systems that do not have any Instant Capacity components can be part of a GiCAP group. Deactivating resources on these systems allows them to loan usage rights to other members in the group.

Members of a GiCAP group do not have to be located near each other. IP connectivity between the members, the Group Manager, and the standby Group Manager (if any), sufficient GiCAP sharing rights, and adherence to the GiCAP grouping rules are the only constraints.

For systems with multiple network cards, you can specify the additional network paths as additional hosts when defining member systems in a group, for enhanced connectivity. However, there might be a significant performance penalty because the Instant Capacity software occasionally must access all the multiple hosts in that case. You cannot specify the Group Manager and the standby Group Manager such that they are different network paths to the same system. An OS instance can host one and only one GiCAP database, regardless of the number of hostnames by which that OS instance might be reachable.

WBEM version A.02.05 or higher must be installed on the Group Manager and on all member systems in order to use Global Instant Capacity. The CIM Server configuration property `sslClientVerificationMode` must be set to a value of “optional” on all GiCAP Group Managers and on all OS instances of all member systems. (The CIM Server may need to be restarted if the property was not previously set to this value.) For details, see *cimconfig(1M)*. Note that communication between the managers and members of groups is established using SSL certificates that are supplied by the GiCAP software.

If the GiCAP Group Manager system becomes unavailable, the standby GiCAP Group Manager can take over GiCAP group operations from the Group Manager. If both the Group Manager and the standby Group Manager are unavailable, usage rights and temporary capacity remain as allocated to each group member. Within a server complex, the usage rights can be deployed to other partitions, but movement of usage rights between complexes cannot occur.

On a system with full usage rights, the `icapd` daemon does not constantly verify the system configuration. It can take up to 12 hours for each partition of a system converted from non-iCAP to iCAP to discover that it is now an iCAP system. During this time, `icapmanage -s` shows asterisks for the new member, and the `icapstatus` command issued on that system shows “Runs iCAP” as “N”. To force a faster conversion, kill the `icapd` daemon and wait for it to restart.

---

# 8 Using Instant Capacity on HP Integrity Superdome 2

This chapter covers the following topics:

- “Overview” (page 129)
- “Important Considerations” (page 130)
- “Installing iCAP on HP Integrity Superdome 2” (page 130)
- “iCAP commands” (page 131)
- “iCAP Use Cases” (page 133)

## Overview

### iCAP: From compliance check paradigm to self enforcement

The Instant Capacity (iCAP) software on HP Integrity Superdome 2 does not perform a compliance check on the complex to check if the complex meets the iCAP contract. This functionality will be added in a later version. In order to comply with the iCAP terms and conditions, you must self enforce the compliance check with the help of tools and options provided by the Instant Capacity software. The following use cases describe how to perform self enforcement in various scenarios.

#### Case 1 — The customer has purchased some number of iCAP cell blades

1. Boot the partitions on the complex.
2. Set the usage rights to be used for every partition individually using the iCAP tools and options.

The total of all assigned usage rights assigned on the complex must not exceed the number of cores on active cell blades, that is, eight times the number of active cell blades.

#### Case 2 — The customer has purchased some number of iCAP cell blades and some number of Rights to Use (RTU) for the processors on the cell blades

1. Go to the HP Utility Pricing Portal at [www.hp.com/go/icap/portal](http://www.hp.com/go/icap/portal).

This portal handles all of iCAP transactions.

2. Sign in to the portal using your Passport ID.

If you are a first time user, you must register on this portal to obtain a Passport ID.

3. Click **Create codewords**.



**NOTE:** For self enforcement, codewords are not provided but the process used for inventory tracking mirrors what is used for codewords.

---

4. Follow the prompts as if you were going to generate an RTU code word. An audit snapshot of the complex is required to complete the process.
5. Subscribe to receive a mail or keep a record of the transaction.
6. Adjust the core usage rights on the partitions to take advantage of the additional computation capacity.

#### Case 3 — GiCAP

To use GiCAP in this model of self-enforcement, a group manager is not required. However, it may be needed in a later version.

1. Go to the HP Utility Pricing Portal at [www.hp.com/go/icap/portal](http://www.hp.com/go/icap/portal).

This portal handles all of the iCAP transactions.

2. Sign in to the portal using your Passport ID.  
If you are a first time user, you must register on this portal to obtain a Passport ID.
3. Click **Create codewords**.



---

**NOTE:** For self enforcement, codewords are not provided but the process used for inventory tracking mirrors what is used for codewords.

---

4. Follow the prompts as if you were going to generate a GiCAP code word.
5. Subscribe to receive a mail or keep a record of the transaction.

You can move the usage rights across complexes by disabling cores on one complex and activating them on another as long as the total does not exceed the rights purchased.

## Important Considerations

### Sizing Partitions

Size the partition for maximum cores it might require (this can be more than the available RTUs). For activating the cores dynamically using iCAP, the cores must be assigned to the partition. Hence, the partitions must be over provisioned with the maximum requirement. This consideration might be simple when creating nPars due to the fact that it is assigned at the granularity of cells and not cores.

For example, if you know that a vPar may need a maximum of 10 cores at any point in time, create a vPar with 10 cores. The number of usage rights can be assigned to this vPar after booting it to HP-UX with iCAP Version 10.00. If the partition is to be assigned with only five usage rights, it can be done using the iCAP commands, so that only five cores are active on the vPar at any point in time. iCAP ensures that the number of active cores are equal to assigned usage rights (Intended Active). You can dynamically activate or deactivate cores limited by the partition size set initially.

### Resizing Partitions

If a partition has to be resized to change the number of cores that it was originally assigned with, the cores have to be brought down and re-adjusted. After you re-adjust the assigned cores, HP recommends you to set the usage rights assigned to the partition to ensure compliance.

### Memory iCAP

The iCAP memory is available on HP Integrity Superdome 2 servers. The iCAP offering on memory is very similar to the current iCAP offering. You can purchase an HP Integrity Superdome 2 cell blade with all iCAP memory or all active memory. All inactive memory on an iCAP cell blade has to be activated at the same time. Until the memory is activated, the iCAP cell blade is considered to be logically powered off. In order to comply with the iCAP Terms & Conditions, HP recommends you to power off the iCAP cell blade that contains iCAP memory. After purchasing enough memory RTUs to activate all of the iCAP memory on the cell blade, follow the same process in the core usage section to update the iCAP inventory on the HP Utility Pricing Portal.

## Installing iCAP on HP Integrity Superdome 2

The installation and upgradation of iCAP on HP Integrity Superdome 2 remains the same as on other iCAP supported systems. However, for iCAP Version 10.00 and later, on HP-UX 11i v3 HP Integrity systems, patch PHCO\_39831 must be installed since it is a co-requisite.

## iCAP commands

iCAP supports a limited set of command options on HP Integrity Superdome 2. The user is presented with command options that are sufficient to self enforce the compliance.

### Viewing number of active cores

iCAP ensures that the number of active cores are equal to the number of usage rights on a partition. You can view the number of active cores on the partition using the following command:  
machinfo

The number of logical processors displayed under the CPU info section is the number of active cores on the partition.



---

**NOTE:** If Hyper-threading is on, the machinfo command reports twice the number of actual active cores, as logical processors. For example, if the same output as shown below was captured after turning on Hyper-threading, there would be 14 logical processors reported.

---

Sample output:

```
# machinfo
CPU info:
  2 Intel(R) Itanium(R) Processor Family processors (1.47 GHz, 20 MB)
    4.8 GT/s QPI, CPU version C0
    7 logical processors

Memory: 8043 MB (7.85 GB)

Firmware info:
  Firmware revision: 001.040.000
  FP SWA driver revision: 1.18
  IPMI is supported on this system.
  BMC firmware revision: 0.05

Platform info:
  Model: "ia64 hp Superdome2 SD16"
  Machine ID number: f6cf-f0e7-4ed6-9718-1165c2f
  Machine serial number: testing

OS info:
  Nodename: beer07b
  Release: HP-UX B.11.31
  Version: U (unlimited-user license)
  Machine: ia64
  ID Number: 411663
  vmunix_release_version:
  @(#) $Revision: vmunix: B.11.31_LR FLAVOR=perf
```

### Generating snapshot

An audit snapshot is to be presented at the Utility Pricing Portal while purchasing usage rights for iCAP resources. The audit snapshot is an encrypted representation of the iCAP resources on the complex. This can be generated using the icapstatus command as shown below:

```
icapstatus -s
```

Sample output:

```
# icapstatus -s
This audit snapshot string may be presented to the HP Portal or to WTEC in
order to establish the current iCAP state of the complex. The string is an
encrypted and signed representation of the global icapstatus information.
Product number: *****
Serial number: ******
Audit snapshot:
```

```
*****.4DyUMVs.dj4XXTe.UwR1J5X.8pBUyEE.X9VMA1x.ggskF9A.p0Ae4tx.yeK7iza.pc22CEf.
NRbfVf5.GeD11Tc.r7qKKNz.7WDY9RN.YjKpTyK.atfR7Li.Hhna33J.vvfLae8.FvBgYmM.*****
```

## Setting usage rights for a partition

Self enforcement requires you to decide how many usage rights are to be assigned for a partition and set it. You can use the `icapmodify` command for this purpose. The number of active cores is set to the number specified as a command argument. You must run the command after you boot to a partition for the first time and subsequently after resizing the partition:

```
icapmodify -s <number of core usage rights>
```

Sample output:

```
# icapmodify -s 4
4 cores are intended to be active and are currently active.
```

## Deactivating Cores

You must deactivate the cores in order to maintain compliance when usage rights are moved. You can deactivate the cores until a minimum number required by the partition is reached:

```
icapmodify -d <number of cores to deactivate>
```



---

**IMPORTANT:** In HP Integrity Superdome 2, an active partition must have a minimum of one *active core*. This is unlike the previous generation of servers where a partition with x cell boards would required a minimum of x core usage rights (one per cell board).

---

Sample output:

```
# icapmodify -d 1
3 cores are intended to be active and are currently active.
```

## Activating Cores

Cores can be dynamically activated on a partition to address any additional demand. As mentioned before, iCAP does not perform any compliance check to see if the total active cores exceed the total usage rights purchased. It is required that the cores should be available to the partition for activating them. that is, core activation request will fail if the request exceeds the assigned number of cores for the partition. Execute the following command to activate the cores:

```
icapmodify -a <number of cores to activate>
```

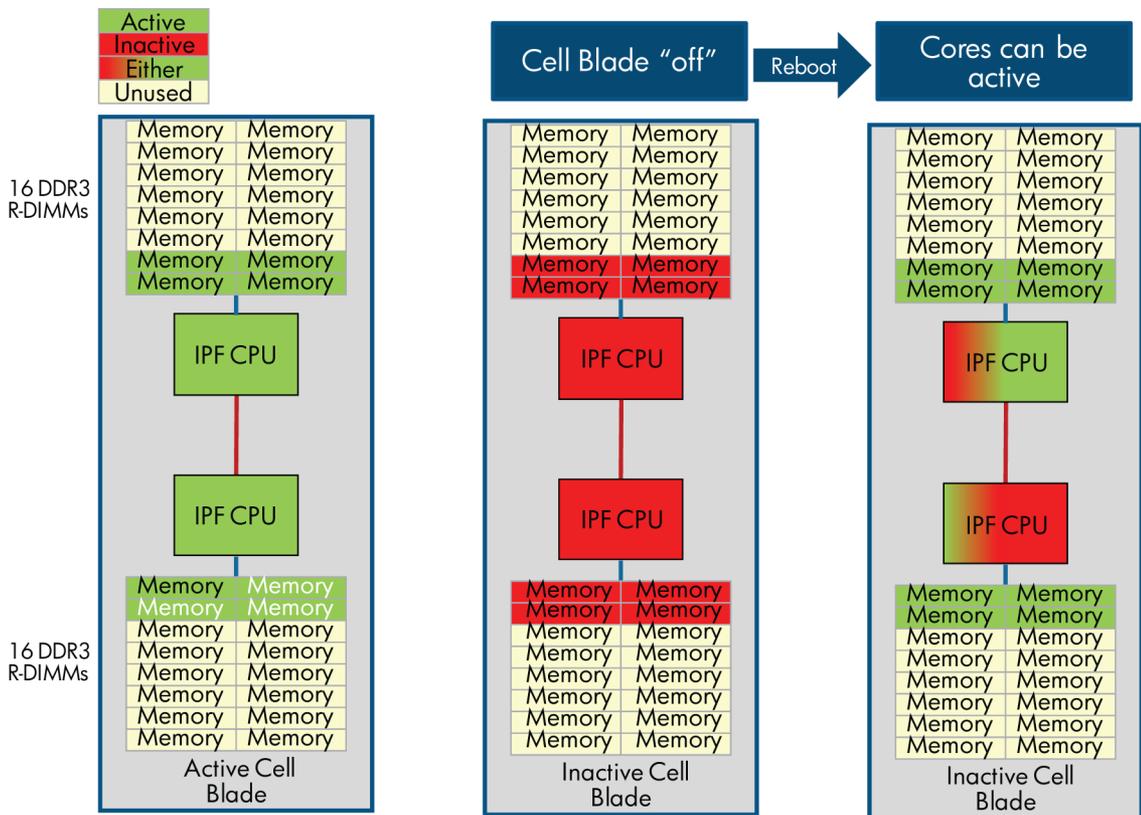
Sample output:

```
# icapmodify -a 1
4 cores are intended to be active and are currently active.
```

## Using iCAP Memory

The iCAP memory is offered in units of 16 GB, and the RTU is available in units of 4 GB. During the time of activation, you must purchase enough iCAP RTUs to activate all the memory on the cell blade. After the purchase, the iCAP cell blade can be turned on. You must reboot the partition on which the iCAP memory is activated for the iCAP changes to take effect. The cores in the cell blade with inactive memory need not be activated for the memory to be activated. After adding a cell blade to a partition to activate the memory, you must adjust the number of active cores to ensure compliance. Figure 8-1 describes the configuration options for iCAP memory.

**Figure 8-1 iCAP Memory configuration option**



## iCAP Use Cases

### Core migration between partitions in a complex

Active cores can be redistributed across any or all partitions of the complex if those partitions contain inactive cores.

Consider a complex with 2 nPartitions, nPar1 and nPar2. Table 8-1 displays the Initial Configuration of the complex:

**Table 8-1 Use Case 1 – Initial Configuration**

Partition	Total Cores	Total Active Cores
nPar1	16	12
nPar2	8	2

Core migration from nPar1 to nPar2 can be done as follows:

In nPar1:

```
# icapmodify -d 2
10 cores are intended to be active and are currently active.
```

In nPar2:

```
# icapmodify -a 2
4 cores are intended to be active and are currently active.
```



**IMPORTANT:** To remain in compliance, you must perform the deactivation operation first.

Core migration between the partitions is now complete.

Configuration of Complex after core migration:

**Table 8-2 Use Case 1 — Configuration of Complex After Core Migration**

Partition	Total Cores	Total Active Cores
nPar3	16	10
nPar4	8	4



**IMPORTANT:** The total number of active cores in the complex must not change at the end of this operation.

## Core migration across complexes (GiCAP)

Consider two complexes, complex1 and complex2 each with 2 partitions. The partitions nPar1 and nPar2 are contained in complex1; nPar3 and nPar4 are contained in complex2.

Table 8-3 contains initial configuration of Complex1.

**Table 8-3 Use Case 2 — Initial Configuration**

Partition	Total Cores	Total Active Cores
nPar1	16	12
nPar2	16	10

Table 8-4 contains Initial configuration of Complex2.

**Table 8-4 Use Case 2 — Initial Configuration**

Partition	Total Cores	Total Active Cores
nPar3	16	10
nPar4	8	8

Core migration from complex2 to complex1 may be done as follows:

Choose a partition(s) in complex2 where you would like to reduce the active core count. In this use case, we have chosen nPar4.

In nPar4:

```
# icapmodify -d 3
5 cores are intended to be active and are currently active.
```

The usage rights obtained from complex2 can be distributed among the partition(s) of complex1. In this use case, we distribute it among nPar1 and nPar2.

In nPar1:

```
# icapmodify -a 2
14 cores are intended to be active and are currently active.
```

In nPar2:

```
# icapmodify -a 1
11 cores are intended to be active and are currently active.
```

Core migration between the complexes is complete.



**IMPORTANT:** To remain in compliance, you must perform the deactivation operation first.

**Table 8-5 Configuration of Complex1 post core migration**

Partition	Total Cores	Total Active Cores
nPar1	16	14
nPar2	16	11

**Table 8-6 Configuration of Complex2 post core migration**

Partition	Total Cores	Total Active Cores
nPar3	16	10
nPar4	8	5



**IMPORTANT:** The sum of total number of active cores in the complexes involved in core migration must not change at the end of this operation. In the above example, the sum of the total active cores in the complexes is 40 (22 in complex1 and 18 in complex2). It must remain same after the core migration.



---

## 9 Troubleshooting

This chapter covers the following topics:

- “Handling Compliance Exceptions” (page 137)
- “Troubleshooting the Instant Capacity Software” (page 140)
- “Diagnosing Email Configuration” (page 142)

### Handling Compliance Exceptions

A complex can get out of compliance with the Instant Capacity contract if any of the following occurs:

- More cells are active than expected (not enough inactive cells).
- More memory is active than expected (not enough inactive memory).
- More cores are active than expected (not enough inactive cores).
- There is a negative temporary capacity balance.
- (GiCAP) Sharing rights are insufficient.
- (GiCAP) Hardware is added that is incompatible with the group.



**NOTE:** Your system might be out of compliance because it has different Instant Capacity software products installed. For example, if a partition has the old product B9073AA installed (Instant Capacity versions B.03.x through B.05.x) and another partition in the same system has the new product B9073BA installed (Instant Capacity version B.06.00 or greater), the B9073BA software determines that all components in partitions that have B9073AA installed are active. For details of correcting this noncompliant state, see “Upgrading to Instant Capacity version B.06.x or later (HP-UX)” (page 193).

---

The Instant Capacity software sends an exception report (via email) if one of these exception conditions occurs. Exception information is also written to the system log file. In some cases, compliance is enforced by deactivating cores at boot time. For more details about enforcement, see “Temporary Instant Capacity Expiration and Compliance Enforcement” (page 86) and virtual partition “Boot Time Compliance” (page 66).

Example 9-1 shows an email exception report for having more cores active than expected.

## Example 9-1 Exception Report for More Cores Active than Expected

---

To: root@parl.yourorg.com  
Subject: Instant Capacity Exception Report

This message is being sent to inform you that your Instant Capacity complex (containing the partition parl) is in an exception state based on the following detected exceptions:

More cores active than expected

This complex is out of compliance with the Instant Capacity contract. The listed exceptions must be corrected as soon as possible.

'More cores active than expected' means that the number of active cores across the complex exceeds the number of core usage rights. For details of core usage, use the `icapstatus` command. This exception state may be corrected by: deactivating cores until the number of inactive cores matches the global number of cores without usage rights, as reported by `icapstatus`. Alternately, additional core usage rights can be purchased for permanent activation, or temporary capacity (TiCAP) can be purchased and applied to the complex.

NOTE: When a system is in an exception state, many system management operations are likely to fail. These include, but are not limited to: the ability to activate cores, the ability to manage hard partitions (`nPars`), the ability to manage virtual partitions (`vPars`).

NOTE: One or more of the exceptions listed in this mail may be due to assumptions made because of an inability to get complete information (see `icapstatus` output for details). In some cases, exception states arise when partitions are not shut down properly, or have been loaded without Instant Capacity software. To eliminate these possibilities, do the following:

- 1) always use the "shutdown" command when shutting down a partition.
- 2) boot any partitions that may have been shutdown improperly.
- 3) ensure that all cells in the system are powered on.
- 4) ensure that Instant Capacity software is properly loaded and configured on all partitions.

NOTE: An exception related to cells, memory, or cores may occur if a cell containing inactive components is removed from the complex (e.g. for repairs or upgrades). Because Instant Capacity compliance requires that the number of inactive components on the complex must match the number of components without usage rights, you may need to adjust the number of inactive components on the complex if a cell containing inactive components is removed.

See the Instant Capacity user's guide at `/usr/share/doc/icapUserGuide.pdf` for more information.

---

As mentioned earlier, you can also receive an exception report for other exception conditions. Here are the other conditions and examples of the appropriate exception report content:

- More cells active than expected

'More cells active than expected' means that the number of active cells across the complex exceeds the number of cell usage rights. To find out how many inactive cells are expected on the complex, run `icapstatus` and look at the global number of cells without usage rights. This exception may be corrected by using `parmodify` to set the `use_on_next_boot` flag for an assigned cell to "n", followed by a partition reboot. Alternately, cells may be

turned off after a partition reboot, unassigned from partitions, or additional cell usage rights may be purchased for permanent activation.

- More memory active than expected

'More memory active than expected' means that the amount of active memory across the complex exceeds the available memory usage rights. To find out how much inactive memory is expected on the complex, run `icapstatus` and look at the global amount of memory without usage rights. Typically, this exception occurs when a newly added cell without usage rights is activated, but it may also occur when a cell with a small amount of memory is deactivated and replaced with a cell with a greater amount of memory. To correct this exception, one or more cells will have to be deactivated in order to deactivate the appropriate amount of memory. This can be done by using `parmodify` to set the `use_on_next_boot` flag for an assigned cell to "n", followed by a partition reboot. Alternately, cells may be turned off after a partition reboot, unassigned from partitions, or the appropriate amount of memory usage rights may be purchased for permanent activation.

- Negative temporary capacity balance

'Negative temporary capacity balance' means that the authorized temporary capacity (TiCAP) balance on the system has been depleted and continued use of core(s) without usage rights has caused additional (unauthorized) temporary capacity to be consumed. To correct this exception, first correct any 'More cores active than expected' exception so that the temporary capacity balance does not continue to grow more negative. Then, purchase additional temporary capacity and apply the temporary capacity codeword to the complex. Alternately, purchase additional usage rights to match the number of core(s) consuming temporary capacity and apply the Right to Use (RTU) codewords to the complex.

# Troubleshooting the Instant Capacity Software

If the Instant Capacity software is not functioning, perform the following steps:

1. Verify that the Instant Capacity software is installed and not corrupted. On HP-UX systems, enter the following command:

```
/usr/sbin/swverify iCOD
```

You should see the message `Verification succeeded.` in the output of the `swverify` command.

On OpenVMS systems, enter the following commands to verify that the Instant Capacity and WBEM software are installed and configured:

```
$ @sys$manager:ICAP$CLI_UTILS.COM CONFIG_CHECK  
$ show log ICAP$CONFIGURED  
"ICAP$CONFIGURED" = "TRUE" (LNM$JOB_nnnnnnnn)
```

```
$ pipe product show hist | search sys$pipe WBEMCIM  
HP I64VMS WBEMCIM X2.92-A090624 Full LP Install Val 05-OCT-2009
```

2. (HP-UX) If the `swverify` command in step 1 displays an error, reinstall the Instant Capacity software. For details, see “Installing Instant Capacity Software” (page 46).

(OpenVMS) If the Instant Capacity software is not configured, configure it using the `SYS$MANAGER:ICAP$CONFIGURE.COM` procedure. If the WBEMCIM software is not installed, install it using the `PRODUCT INSTALL` utility.

3. Verify that the status of your Instant Capacity system or partition is correct by entering the following command:

```
HP-UX: /usr/sbin/icapstatus  
OpenVMS: $ICAP SHOW STATUS
```

The Instant Capacity software version is displayed as something similar to “B.11.23.09.00”. If you do not have a version 9.x installed, install the Instant Capacity 9.x software.

If there is a discrepancy between the number of reported components with or without usage rights and your Instant Capacity contract, contact your local HP Response Center and request iCAP assistance.

4. Ensure that the required processes for Instant Capacity are running.

On HP-UX systems, verify that the `icapd` daemon is running on the system or partition by entering the following command:

```
/usr/bin/ps -e | grep icapd
```

The command indicates that the `icapd` daemon is running on the partition. If it is not running, view the system log file (`syslog`) for `icapd` error messages and take the appropriate action.

On OpenVMS systems, verify that the `ICAP_SERVER` and `WBEM CIMSERVER` processes are running:

```
$ pipe show sys | search sys  
$pipe ICAP_SERVER%SEARCH-I-NOMATCHES, no strings matched
```

```
$ pipe show sys | search sys$pipe CIMSERVER  
202046D CIMSERVER HIB 10 8335702 0 00:22:46.98 75250 77655 M
```

A line of output lists the process information. In this example, the `CIMSERVER` process is running and the `ICAP_SERVER` process is not running.

5. (HP-UX) Make sure that the kernel driver `diag2` is built into the kernel.
6. (HP-UX) Make sure that the `nParProvider` bundle is installed.
7. Make sure that the required WBEM provider modules are installed and running.

On HP-UX systems, the WBEM B8465BA bundle (version A.02.05 or higher) must be installed.

On OpenVMS systems, verify the WBEM installation by using the following commands:

```
$ cimprovider := $WBEM_OPT:[wbem.bin]cimprovider
$ pipe cimprovider -l | search sys$pipe -
"HP_NParProviderModule", "HP_iCAPProviderModule",
"HP_iCODProviderModule"
HP_NParProviderModule
HP_iCAPProviderModule
HP_iCODProviderModule
```

All three of the provider modules must be loaded.

8. (HP-UX) Make sure partition commands, such as `parstatus`, are working. For failures in virtual partitions, check the `vPar` commands such as `vparstatus`.
9. View the Instant Capacity log file and syslog file for any error messages. On HP-UX systems, these files are `/var/adm/icap.log` and `/var/adm/syslog/syslog.log`. On OpenVMS systems, these files are `sys$manager:icap.log` and `sys$manager:operator.log`.

## Additional Troubleshooting Steps for Email Connectivity

If you are using asset reporting, perform the following additional troubleshooting steps to make sure that HP is able to receive an email message from the Instant Capacity software:

1. Execute this command:  

```
/usr/sbin/icapnotify <reply_address>
```

Where `reply_address` is the email address where you want HP to send a confirmation message.
2. Verify that HP replies with a confirmation email message, to the reply address specified in step 1.

If you do not receive the confirmation email message in step 2 from HP, then your partition is unable to send email over the internet to the `hp.com` domain. For details on troubleshooting your email configuration, see “Diagnosing Email Configuration” (page 142).

# Diagnosing Email Configuration

Follow these steps to confirm the email configuration or to aid in debugging the configuration:

1. Send an email message from your system to an email address in the same domain (intranet) and confirm receipt of the email message.
2. Send an email message from your system to an email address outside of your domain (to the internet, for example, to a yahoo or hotmail email address) and confirm receipt of the email message.
3. Send an email message from your system to someone at HP (for example, a HP representative in a local account team) and confirm the person at HP received the email message.
4. As root, execute the command:

```
/usr/sbin/icapnotify <reply_address>
```

This command sends an email message to the HP audit application. HP sends a confirmation email message to the reply address that is specified. Receipt of the confirmation email message confirms successful email configuration.

5. If the previous steps are all successful, but asset reports are still not visible at the HP portal, examine your email configuration to determine whether outgoing messages are automatically being modified or appended, such as to include something like a privacy notice. Additions or modifications to encrypted asset reports can cause the portal to reject them.

If any of these steps do not produce the correct result, see “[Configuring Email on Instant Capacity Systems](#)” (page 202) for details on how to correctly configure email connectivity for Instant Capacity.

---

# 10 Frequently Asked Questions

This chapter covers frequently asked questions on the following topics:

- “Instant Capacity Software” (page 144)
- “Instant Capacity Hardware” (page 147)
- “Global Instant Capacity” (page 148)

## Instant Capacity Software

**What software product is required for Instant Capacity on Itanium-based servers running HP-UX?** The HP software bundle for the Instant Capacity version 9.x software, on Itanium-based servers running HP-UX 11i v1, 11i v2, or 11i v3, is HP product number B9073BA.

**Can one HP enterprise server be under both a Pay per use (PPU) and Instant Capacity contract at the same time?** No, the PPU and Instant Capacity software bundles are mutually exclusive. Both can be installed on the same HP enterprise server, but because the server can be purchased using either PPU or Instant Capacity (but not both), the server can be configured only for the purchased pricing solution.

**Where can I find the Instant Capacity 9.x software bundle?** The Instant Capacity 9.x software bundle (B9073BA for HP-UX systems, BA484AA for OpenVMS systems) is installed at the factory for new HP-UX systems and is automatically installed on OpenVMS Version 8.4 systems when the operating system is installed. If you need to install the software, it is available from the following:

- (HP-UX) HP Software Depot web site: <http://www.hp.com/go/softwaredepot> (search for “Instant Capacity”)
- September 2008 HP-UX 11i v3 Operating Environment (OE) media
- September 2008 HP-UX 11i v2 Applications Software media
- September 2008 HP-UX 11i v1 Applications Software media
- November 2009 OpenVMS Version 8.4 Operating System Media

For details about installing the Instant Capacity 9.x software bundle, see “Installing Instant Capacity Software” (page 46).

**One of my HP-UX or OpenVMS applications has compatibility issues with the Instant Capacity software. How do I correct the problem?** The application might have a problem when cores are activated or deactivated. Some applications size themselves at system startup based on the number of active cores and they do not adjust for core increases or decreases. For more information, see “Software Application Considerations” (page 71).

**We want to utilize temporary capacity on our Itanium-based server. What system configuration is necessary and how do we acquire temporary capacity?** First, purchase the Temporary Instant Capacity (TiCAP) product from your HP sales representative, acquire and apply the TiCAP codeword, and then activate additional cores with temporary capacity. If you want asset reporting in order to view temporary capacity balances on the Utility Pricing Solutions portal, make sure that email is properly configured for the system you plan on using temporary capacity. For details about email configuration, see “Diagnosing Email Configuration” (page 142). For details about temporary capacity, see Chapter 5: “Temporary Instant Capacity” (page 75).

**How much history is retained in the Instant Capacity log files?** The Instant Capacity log files retain up to 2 MB of Instant Capacity events. An Instant Capacity event occurs and is written to the log files when one of the following occurs:

- The Instant Capacity software sends an asset report to HP (daily at noon, if configured).
- A partition with Instant Capacity is shut down.
- A partition with Instant Capacity is started.
- A partition with Instant Capacity has a configuration change (that is, a core is activated or deactivated).
- A codeword is applied.
- Usage rights are seized from a GiCAP system.

You can view all events in the Instant Capacity log files in the `/var/adm/icap.log` or `/var/adm/icap.log.old` file on HP-UX systems, and in the `sys$manager:icap.log` file on OpenVMS systems. GiCAP events can be viewed on the Group Manager system in `/var/adm/GiCAP.log` (see “Global Instant Capacity” (page 148)).

**How can I obtain codewords for newly purchased usage rights if the Utility Pricing Solutions portal is down?** If the Utility Pricing Solutions portal is down, contact the HP Response Center. The Response Center can create an emergency codeword via the Instant Capacity codeword backup tool.

**What licensing is required for the Instant Capacity software?** For Instant Capacity version 9.x, to activate additional components (cores, cell boards, or memory), you must acquire additional usage rights individually. For details, see “Usage Rights Requirement” (page 31)s. To create a Global Instant Capacity group, you must purchase sharing rights. For more information about sharing rights, see “Global Instant Capacity Sharing Rights” (page 108).

**The resulting configuration of my Instant Capacity system does not agree with what I ordered from HP. How did this configuration change occur?** The Instant Capacity software can control the granularity of processor activation or deactivation to the single-core level. The Instant Capacity ordering and manufacturing rules often do not allow such fine granularity.

The Instant Capacity ordering rules dictate the quantity of cores with and without usage rights in the cell boards. Because the Instant Capacity software distributes the core usage rights (for a given partition) in a manner that optimizes loads across all cells, the resultant configuration might be different than the original order. However, the number of cores with and without usage rights matches what was ordered.

For example, suppose you order an rx8620 server with 2 cell boards, in which the first cell board contains 4 active cores with usage rights, and the second cell board contains 2 active cores with usage rights and 2 inactive cores without usage rights, for a total of 6 active and 2 inactive cores. At run time, the Instant Capacity software balances the distribution of active cores across the cell boards so that each cell has 3 active cores with usage rights and 1 inactive core without usage rights.

**How does Instant Capacity interact and coexist with partitions running software other than HP-UX?** Instant Capacity is supported only on HP-UX and OpenVMS for Integrity systems. If other partitions of an Instant Capacity system are running another operating system, all the system components in the non-HP-UX and OpenVMS partitions appear to the Instant Capacity software as active components (with usage rights). When verifying the correct number of inactive components without usage rights, only the HP-UX and OpenVMS partitions are examined.

**What email is sent by the Instant Capacity software?** The following table lists the email messages sent to the system from the Instant Capacity software. On OpenVMS systems, the iCAP software agent is ICAP\_SERVER rather than icapd.

**Table 10-1 Email sent by the Instant Capacity software**

Triggered By	Email Message
icapmodify (if a configuration change occurs)	Information about the configuration change is sent to the system contact, if specified, and if change notification is set to “on”.
icapd (daily, when the projected TiCAP balance expiration is less than the warning period: by default, when less than 15 days)	A temporary capacity expiration notification is sent to the system contact, if specified, and root.
icapd (daily, if more than expected cores, memory, cells, are active; also if TiCAP has a negative balance)	An exception report (for noncompliance) is sent to the system contact, if specified, and root.
icapd (if one or more cores are deactivated at boot time to enforce compliance)	An Instant Capacity enforcement message is sent to the system contact, if specified, and root.
vPars startup (when the virtual partition has more cores assigned to it than the number of intended active cores for the nPartition)	Information about why the virtual partition is not being allowed to boot is sent to the system contact, if specified, and root.

**Table 10-1 Email sent by the Instant Capacity software (continued)**

Triggered By	Email Message
icapmodify (when a codeword is applied to a GiCAP member which modifies the number of cores without usage rights on that member, and available sharing rights are insufficient, triggering a lockout on all groups; or when incompatible hardware is added)	A warning message is sent to the system contact, if specified, and root, stating that the Group Manager must adjust the number of sharing rights.
icapmanage (when usage rights are seized from an unavailable server)	An informational message is sent to indicate that core usage rights have been acquired from an inactive group member.
icapmanage (when seized usage rights expire)	An informational message is sent to indicate that usage rights seized from an inactive group member have expired and therefore the seized usage rights are reverted to the original group member.
icapmodify (when a group is noncompliant due to a lack of sharing rights or for hardware noncompliance)	An exception report (for noncompliance) is sent to the system contact, if specified, and root.

If asset reporting is configured, and the system has email connectivity to HP, these messages are sent to HP:

**Table 10-2 Asset reporting email sent by the Instant Capacity software**

Triggered By	Email Message
icapnotify (on demand)	An asset report is sent to the reply address, root, and HP (the asset report sent to HP is encrypted).
System startup and system shutdown	An encrypted asset report is sent to HP.
icapd (daily at noon)	An encrypted asset report is sent to HP.

**Does the upgrade of one partition on a server to Instant Capacity version 9.x mean that every partition must be upgraded?** The HP Utility Pricing Solutions contract requires all participants to run the latest version of the Instant Capacity software, and to maintain the Instant Capacity software on each partition in the system. Participants who do not meet these requirements may be in breach of contract. However, there may be some circumstances where you cannot upgrade to the latest version of Instant Capacity. In these cases, it is possible to run mixed versions, but HP does not test or guarantee interoperability except within the documented upgrade procedures and restrictions. Earlier versions of Instant Capacity prior to version 6 need to be upgraded, as described in “Upgrading to Instant Capacity version B.06.x or later (HP-UX)” (page 193). However, to use the GiCAP features, the Group Manager, the standby Group Manager (if any), and all OS instances on all member systems must be running version 9.0 or later of the Instant Capacity software.

## Instant Capacity Hardware

**Can a faulty cell board be replaced with an inactive Instant Capacity cell board?** Yes. First, deactivate the failed cell board by using the `parmodify` command and rebooting. Then activate the inactive iCAP cell board and reboot. In this situation, you do not need to obtain an RTU to activate the cell board.

# Global Instant Capacity

**Does HP know the configuration of the GiCAP groups?** No. GiCAP group data is stored on the GiCAP Group Manager, which runs in the customer's data center. The group configuration is limited by HP grouping rules, but information about groups or group members is not communicated to HP.

**Is GiCAP migration supported on a completely unavailable server?** Yes. GiCAP Disaster Recovery is supported as long as the Group Manager is running Instant Capacity version 9.0 or later, and all known hosts on member servers in the group are running Instant Capacity version 9.0 or later.

**How much history is retained in the GiCAP log files?** The GiCAP log files retain up to 2 MB of GiCAP events. Some example GiCAP events that might be logged include:

- A GiCAP group is created or removed.
- A member is added to or removed from a GiCAP group.
- A member is updated through addition or removal of a host.
- A standby Group Manager is added or removed.
- A standby Group Manager takes control from the active Group Manager.
- A codeword is applied that modifies the available Sharing Rights.
- The grouping rules are replaced.
- Usage rights or temporary capacity is transferred between a member and the GiCAP Group Manager.

You can view all events in the GiCAP log files in the `/var/adm/gicap.log` file on the Group Manager system.

# Instant Capacity HP-UX Manpages

# iCAP(5)

## NAME

iCAP - Instant Capacity software for HP-UX

## DESCRIPTION

The HP Instant Capacity program provides services for instantly increasing or decreasing processing capacity on supported HP servers to meet varying system demands. An Instant Capacity server is an HP cellular (partitionable) server that is governed by an Instant Capacity contract constraining the number of cores, cell boards, and memory that must remain inactive at all times. The total number of each type of component that can be active at any time is referred to as the number of usage rights for that component type. The Instant Capacity software provides the ability to dynamically assign usage rights where they are needed among the various partitions of a complex through the use of commands like `icapmodify(1M)` and `parmodify(1M)`, and it provides the means to load balance system processing capacity through these management commands. When the processing demand significantly changes, you can purchase additional component usage rights to increase the overall number of components that can be activated at any time.

Instant Capacity is a part of the HP Utility Pricing Solutions (UPS) program, and was formerly known as iCOD. For detailed information about Instant Capacity, see the *HP Instant Capacity user's guide* located at `/usr/share/doc/icapUserGuide.pdf`.

## Initializing an Instant Capacity Server

Instant Capacity software is installed by HP manufacturing on instantly ignited systems, and is a required component of all HP-UX systems. You can upgrade to later versions by downloading it from <http://www.hp.com/go/softwaredepot> (search for product ID B9073BA), or by installing it from an OE or AE update.

An Instant Capacity server needs no additional initialization, but some of the following configuration steps might be useful:

1. Execute the `icapmodify -c` command to configure contact information (email address) for configuration change notification and exception reports.
2. If using temporary capacity, execute the `icapmodify -w` command to configure the temporary capacity warning period.
3. If the server is a GiCAP Group Manager or a member of a GiCAP group, use the `cimconfig` command (see `cimconfig(1M)`) to set the CIM Server configuration property `sslClientVerificationMode` to "optional". You might need to restart the CIM Server (see `cimserver(1M)`) if the property was not previously set to this value.

## Codewords

Instant Capacity uses codewords for several purposes: to adjust available usage rights for system components, to apply an amount of "temporary capacity" to the system, and to apply "sharing rights" to a Global Instant Capacity (GiCAP) Group Manager to enable creation of one or more groups of member servers that can share usage rights.

All types of codewords must be purchased as specific product numbers from HP. After purchase, the codeword (an encrypted string) must be retrieved from the Utility Pricing Solutions web portal and applied to the appropriate system. Codewords are generated specifically for the system for which they were purchased. You must always specify at least the system serial number and the purchase order number in order to retrieve a codeword from the portal. Application and use of GiCAP codewords are different from other iCAP codewords and are described in the section "GiCAP Sharing Rights".

Prior to activating an Instant Capacity component, a Right to Use (RTU) codeword must be applied to the complex to increase the component-specific usage rights on the complex. After a

fee has been paid to HP for the type and number of components that are to be activated, RTU codewords are made available through the HP Utility Pricing Solutions portal (<http://www.hp.com/go/icap/portal>).

Instant Capacity codewords (such as RTU codewords) are applied to a complex using the `icapmodify` command on any partition of the complex. iCAP codewords are generated with a sequence number, and all iCAP codewords for a particular complex must be applied in the order in which they were generated.

After the appropriate codewords have been applied to a complex, additional components in the complex may be activated, up to the number of component usage rights granted by the applied codewords. Depending on their type, components are activated using the `icapmodify` command (if activating cores), or other commands including `parmodify` (see `parmodify(1M)`) and `parmgr` (see `parmgr(1M)`).

In addition to RTU codewords, cores can be activated with temporary capacity. Temporary capacity codewords allow the activation of more cores than allowed by the usage rights on the complex, but only for a limited time.

## Software Removal

Instant Capacity software cannot be removed. Other software products depend on it to approve configuration changes to the system.

## Status of Instant Capacity Components

Information about the Instant Capacity components on a complex and the available usage rights for each type of component can be obtained by invoking the `icapstatus` command. This command also provides information about the amount of temporary capacity presently in use, the projected expiration of the temporary capacity, and counts of active and inactive components.

One status field that is important to understand is the `intended active` (IA) value for each `nPartition`. Fundamentally, the `intended active` value is the number of cores intended to be active after a reboot. As processing needs change, the IA value is modified by commands like `icapmodify`, `parcreate`, and `parmodify` to represent the new distribution of active cores across the `nPartitions` and thus, the new numbers to activate on reboot of the `nPartitions`.

Note that the status field `intended active` can differ from the status field `actual active` (AA) in some cases, including the following:

- In an `nPartition`, activations or deactivations can be deferred, meaning that the change is pending until the next reboot.
- If there are deconfigured (failed) cores, it may be impossible for the `nPartition` to reach the configured IA value. The IA value is not affected by deconfigured cores, but the AA value may be lower in this case.
- In a virtual partition, core assignments can be increased or decreased with either the `icapmodify` command or the `vparmodify` command, but only `icapmodify` changes the `intended active` for the `nPartition`. `vparmodify` activation commands are constrained by the `intended active` value, but deactivation commands are allowed to deactivate below IA, so that the number of cores actually assigned to all the virtual partitions of the `nPartition` (represented by the AA status field) may be less than the IA value for the `nPartition`. This results in something called “unused capacity”; but typically happens only during the window of time that resources are being shifted across the virtual partitions of an `nPartition`. Overall, this mechanism allows for quicker shifting of resources within virtual partitions.

If the complex is a member of a GiCAP group, the `icapstatus` command provides information about group membership, including any borrow or loan status of usage rights. Detailed information about GiCAP groups can be obtained by invoking the `icapmanage -s` command on a Group Manager system.

## Virtual Partitions

Instant Capacity has a minimum version dependency on vPars A.03.05. For versions of vPars before A.03.05, the `icapmodify` command for activating or deactivating cores in a virtual partition fails with an error message indicating the vPars version dependency.

Instant Capacity can be present on systems or partitions where virtual partition technology is employed. In a virtual partition environment, cores that are not assigned to any virtual partition are considered inactive (in addition to other classes of inactive cores). Unassigned cores can be assigned (activated) or deassigned (deactivated) using either the `icapmodify` command or the `vparmodify` command, depending on the type of adjustment needed, the version of vPars being used, and the level of logging or reporting desired.

One important consideration is that `vparmodify` can be used to activate or deactivate cores in other virtual partitions within the nPartition; `icapmodify` only activates or deactivates cores within the current virtual partition (the partition where the command is invoked). Another consideration is that core assignment via the `vparmodify` command does not result in logging of the activation, email configuration change notification, or transmission of an asset report to HP.

However, the most important consideration is that the `icapmodify` command must be used in a virtual partition environment when you are making any adjustment to an nPartition. If you are adjusting core assignments across virtual partitions in a single nPartition, use the `vparmodify` command for the best coordination between the Instant Capacity software and the vPars software, and for optimized performance. The `vparmodify` command is the fastest and most efficient way to adjust capacity within virtual partitions of a single hard partition, but it does not affect the `intended active` count for the nPartition. Therefore, it cannot be used to migrate unused capacity either to or from other nPartitions.

Note that with vPars A.03.05 or greater, a compliance check is performed whenever a virtual partition is booted. If the total number of cores assigned to all virtual partitions in the current vPar database exceeds the nPartition's `intended active` core count, the Instant Capacity software notifies the vPar monitor, and the monitor prevents any virtual partition from booting until the user performs a hard partition boot and modifies either the vPar configuration or the Instant Capacity `intended activecount` for the nPartition.

For more information about virtual partitions, see *vparmodify(1M)*.

## HP Integrity Virtual Machines (Integrity VM)

In an Integrity VM environment, Instant Capacity software provides meaningful functionality only on the VM Host; it does not run on a virtual machine (also known as a “guest”). In particular, Instant Capacity commands will report an error if attempted from a guest. A GiCAP Group Manager cannot be run on a guest, nor can a guest be specified in the host list for a GiCAP group member.

## Processor Sets

In an environment where processor sets are being used, the `icapmodify` command activates Instant Capacity cores into the default processor set and deactivates cores from only the default processor set. Activation or deactivation of cores in nondefault processor sets is a two-step operation. The first step involves the user migrating the cores into or out of the default processor set; the second step is the activation or deactivation of those cores using the `icapmodify` command.

For more information about processor sets, see *psrset(1M)*.

## Temporary Capacity (TiCAP) Program

Customers can purchase an amount of temporary capacity time. This temporary capacity can be used to activate one or more cores beyond the number for which usage rights have been purchased. These extra cores can remain active until they consume the available temporary

capacity time. This allows temporary activation of cores without requiring the purchase and activation of an RTU codeword for permanent activation.

Whenever an Instant Capacity component without usage rights is purchased, an amount of Instant Access Capacity (IAC) might also be included. Instant Access Capacity is exactly the same as temporary capacity, except it is automatically provided with an Instant Capacity component and is not separately purchased. It provides an immediate buffer of temporary capacity in case extra capacity is needed before there is time to purchase either an RTU codeword, a temporary capacity codeword, or to setup a Global Instant Capacity (GiCAP) group.

Temporary capacity can be added to the complex by applying a temporary capacity codeword (available from the HP Utility Pricing Solutions portal) using the `icapmodify` command. Information about the amount of temporary capacity time remaining on a complex can be obtained by executing the `icapstatus` command. A warning is also sent via email when the temporary capacity balance is expected to be depleted within a certain period of time.

The `icapmodify` command allows activating a core using temporary capacity only if at least 30 minutes of temporary capacity is available for each core being activated.

Whenever temporary capacity has been applied to a system (or if the complex is part of a GiCAP group), extra care must be taken to avoid situations that cause the Instant Capacity software on one nPartition of a group member to make assumptions that all cores might be active on another nPartition of the member, for example, when the nPartition is making boot-related configuration changes (at EFI or BCH). If left in this state for more than 12 hours, unexpected temporary capacity may be consumed on the complex.

## Instant Capacity Cell Board

Instant Capacity Cell Board offers a way to have additional (inactive) cell board capacity for your system. These Instant Capacity cell boards, which contain memory and cores, can be activated after a cell RTU codeword is obtained from the HP Utility Pricing Solutions portal and is applied to the complex using the `icapmodify` command. To activate a cell, you must have usage rights for all memory attached to the cell and at least one core.

## Global Instant Capacity

Global Instant Capacity (GiCAP) provides HP customers with the flexibility to move usage rights for Instant Capacity components within a group of servers. It also provides “pooled” temporary capacity across the group. This has several potential benefits: cost-effective high availability, more adaptable load balancing, and more efficient and easier use of temporary capacity.

For example, in case of planned or unplanned down time, a customer can transfer usage rights from a failed partition on one server to one or more other servers in the group that are providing backup availability, thus allowing additional activations of iCAP components on the backup servers. Without GiCAP, the only way to provide this failover scenario is to provision each server with an adequate amount of temporary capacity.

A similar scenario exists for load balancing. Rather than using temporary capacity whenever a server is overloaded (peak profiles for all workloads on a server), usage rights can be transferred from other servers in the GiCAP group that have extra capacity. These borrowed usage rights enable new component activations on the overloaded system.

Pooled temporary capacity for a group of servers is more efficient because all temporary capacity is available to all servers in the GiCAP group. It is also easier to manage if temporary capacity needs to be applied to only one member of the group and monitored across the group instead of monitoring TiCAP for each member complex.

## GiCAP Groups

Global Instant Capacity is built on the concept of a server group, or GiCAP group. The group consists of a list of server complexes that are allowed to share Instant Capacity usage rights (for cores, cell boards, and memory) and temporary capacity. There are no constraints on the number

of servers allowed in a group, but grouping rules defined by HP specify the types of servers allowed to group together.

## GiCAP Group Manager

For each group, an HP-UX system must be designated as an active Global Instant Capacity Group Manager. It is this system that maintains information about the group, group resources, and grouping rules. The `icapmanage` commands are intended to be used only on a Group Manager system to manage one or more GiCAP groups.

The active Group Manager must be an HP-UX system running the Instant Capacity software version 9.0 or later. The system running the Group Manager does not need to have any Instant Capacity components, nor does it need to be a partitionable system. The system must have a machine-readable serial number, as displayed by the shell command `getconf CS_MACHINE_SERIAL`. It is recommended that the Group Manager not be on a partition that is a member of any GiCAP group, and that it manages a single group. If run on a partitionable system, changing the configuration of the partitions may result in the GiCAP Manager becoming inoperative.

## Standby GiCAP Group Manager

The active GiCAP Group Manager can designate a standby Group Manager. This standby Group Manager can take control of GiCAP group management from the active Group Manager using the command `icapmanage -Q`. This allows GiCAP group operations to continue if the GiCAP Group Manager is unable to function.

Note that the requirements and recommendations defined for the Group Manager also apply to the standby Group Manager; it is a Group Manager with “standby” status. The Group Manager in charge of the group has “active” status. If the standby Group Manager takes control, the intention is that its status becomes “active” and the former Group Manager’s status becomes “standby”. For more details, see the “Group Manager Failover Considerations” section.

## GiCAP Group Members

Every operating system on a GiCAP group member must be running Instant Capacity version 9.0 or later. Every GiCAP group member must be hardware-compatible with other GiCAP group members as determined by the GiCAP grouping rules.

## GiCAP Grouping Rules

Once you have determined which system will host the active Group Manager, you must acquire grouping rules from the portal and install the encrypted file on the active Group Manager system using the `icapmanage -i` command. Under some circumstances (for example, when adding new hardware not covered by the grouping rules currently in use), you might need to acquire newer grouping rules from the portal.

## GiCAP Sharing Rights

To create a GiCAP group with members, you must purchase GiCAP sharing rights, acquire the GiCAP codeword from the HP Utility Pricing Solutions portal (<http://www.hp.com/go/icap/portal>), and apply the associated codeword to the active Group Manager system. You purchase at least as many GiCAP sharing rights as the total number of cores without usage rights across all the potential group members. Members can be added to a GiCAP group as long as sufficient sharing rights are available, and as long as the grouping rules indicate hardware compatibility.

Unlike other iCAP codewords, GiCAP codewords must be generated for, and applied to, a specific partition if the active Group Manager is on a partitionable system. This means that in order to retrieve the codeword, you must specify the purchase order number, the system serial number and partition information, if any. Use the `icapmanage -s` command on the active Group Manager system to get the applicable serial number and nPar ID, or vPar code.

GiCAP codewords also have a sequence value and must be applied in the order in which they were generated for the Group Manager system. However, GiCAP codewords are sequenced independently from any other types of iCAP codewords that might be generated for the same system, and can therefore be applied independently from iCAP codewords.

## GiCAP Group Creation

After the sharing rights codeword and the grouping rules have been applied to the active Group Manager, a GiCAP group can be created by issuing the `icapmanage` command using the `-a`, `-g` and `-m` options. Members are added by issuing the `icapmanage` command using the `-a` option, the `-g` option to select the group name, and the `-m` option to specify a name for the new member along with a list of hosts running on the system. The list of hosts must include at least one host per nPartition or virtual partition on the system.

Note that a single partition of a complex cannot join a GiCAP group; all partitions of a complex must be specified when creating a group member. All partitions on a group member must be running HP-UX or OpenVMS. Each member that joins the group decreases the available GiCAP sharing rights by the number of cores without usage rights contributed by that member complex.

## GiCAP Resource Sharing

Once a group is established, Instant Capacity resources (core, cell board, memory usage rights, and temporary capacity) can be shared among all the members of the group.

Usage rights are shared by deactivating resources on one group member, and then activating resources on another member of the group. In effect, the system on which the resources were deactivated is loaning usage rights to the activating (or borrowing) system. The activation and deactivation commands are done using the usual `icapmodify` and `parmodify` commands on the individual member systems to effect this “loan” operation (also sometimes referred to as a transfer of usage rights).

Any temporary capacity available to individual members of the group is combined into a larger pool of temporary capacity that is available for consumption by any and all members of the group, as needed. Usage of shared temporary capacity is the same as with individually purchased TiCAP: group members use the `icapmodify -a -t` command to activate shared temporary capacity. This differs from the sharing of usage rights in that temporary capacity is never a loan to be returned; it is always depleted through its usage over time.

## GiCAP Member Removal

Before removing a member from a GiCAP group, all the borrowed usage rights must be returned, and all outstanding loans must be reclaimed. Borrowed usage rights are returned by deactivating resources on the member about to be removed. Loaned usage rights are reclaimed by deactivating enough resources elsewhere in the group to cover the loan. The reclamation of loaned usage rights on the member about to be removed does not require the activation of resources on that member. As long as the usage rights are not in use elsewhere in the group, removing the member results in the member reclaiming its loans.

There is no such constraint with respect to temporary capacity. A removed member will take with it whatever temporary capacity is currently assigned to it.

When a member is removed from a group, some number of sharing rights are released and become available for future use. The number freed is equal to the number of cores without usage rights on the removed member.

## Upgrades and GiCAP

Care must be exercised before upgrading or changing hardware or operating systems for any member of a GiCAP group. If a member of a GiCAP group changes hardware in such a way that the hardware is no longer compatible with the group, then the group is considered to be out of compliance and group functions are restricted as described in the section “Group Compliance”.

Also, note that the number of available sharing rights is adjusted whenever an iCAP codeword is applied to a GiCAP member system which modifies the number of cores without usage rights on that member. (RTU and AddOn codewords for cores cause such adjustments.)

If available sharing rights go negative (more in use than were purchased for the Group Manager), then all groups managed by that Group Manager are out of compliance and all group functions are restricted until the problem is resolved. The problem can be resolved by purchasing and applying additional sharing rights to the Group Manager, by purchasing and applying core usage rights to one or more group members, or by removing one or more group members from their group.

## **Groups and the TiCAP Balance**

Whenever you have more active cores than the number of core usage rights, the temporary capacity balance is depleted as a mechanism for tracking noncompliance of the group, even if TiCAP has not been purchased for or applied to any member of the group. This differs from the behavior of TiCAP on a complex which is not a member of a group, where TiCAP is decremented only if TiCAP had been specifically purchased for the complex. Within a GiCAP group, temporary capacity is used as an additional compliance mechanism to support the high availability features of a group.

Because group members are automatically considered to be users of temporary capacity, to avoid unexpected TiCAP depletion in a group, it is important to avoid the situations that cause the Instant Capacity software to make assumptions that all cores might be active on a remote nPartition, as described previously in the Temporary Capacity section.

If a member is removed from the group, the TiCAP balance on that complex will continue to be used as a compliance mechanism (decremented whenever the number of active cores exceeds the number of core usage rights), unless the TiCAP balance on the system is exactly 0.

## **Group Compliance**

When a group is out of compliance, the group is said to be “locked”. Sharing of usage rights and temporary capacity between members of the group is not allowed, although members can return borrowed usage rights or reclaim loaned rights. Members cannot be added to the group, but members can be removed from the group as long as the members deactivate any cores using borrowed usage rights and as long as GiCAP is able to reclaim any loaned usage rights.

When such an incompatibility is detected, the GiCAP Group Manager sends email to the local root account and to the registered contact email address for each member of the group.

## **Multiple Group Considerations**

You can create multiple GiCAP groups, and they can be managed by the same Group Manager or by different Group Manager systems. Note that if a Group Manager has an associated standby Group Manager, the standby Group Manager functions as a standby for all the groups managed by that Group Manager.

A server complex can only be a member of a single GiCAP group at a time. In order to participate in a different group, it must be removed from its group before being added to the other group.

Sharing rights can never be transferred between two active Group Manager systems. As you create new groups or add new members to existing groups, you might need to purchase and apply additional sharing rights to the relevant active Group Manager systems. This is necessary even if the member has been moved from another Group Manager that now has excess sharing rights.

Sharing rights can never be applied to a Group Manager with standby status. If a standby Group Manager is requested to take control, the sharing rights of the former active Group Manager move to it. If additional sharing rights need to be applied before failing back to the original Group Manager, they must be purchased specifically for the new active Group Manager system (formerly the standby Group Manager) and acquired from the portal. Once applied, the new total of sharing

rights moves to the originally active Group Manager if it is requested to retake control, although only if the new active Group Manager is able to propagate its GiCAP database to the originally active Group Manager. For more details, see the “Group Manager Failover Considerations” section.

## Group Manager Failover Considerations

If the active Group Manager system becomes unavailable, the standby Group Manager can take over GiCAP group operations from the Group Manager. If both the active Group Manager and the standby Group Manager are unavailable, or if the active Group Manager fails and the `icapmanage -Q` command was not issued to make the standby Group Manager the active Group Manager, usage rights and temporary capacity remain as per allocated to each group member. Within a server complex, the usage rights can be deployed to other partitions, but movement of usage rights and temporary capacity between complexes cannot occur.

An administrator can have a standby Group Manager take control at any time using the `icapmanage -Q` command. While this can be done routinely, for example to allow shutting down a functioning active Group Manager for maintenance, normally this command is issued either when the active Group Manager has failed, or when a network outage has made it unable to communicate with critical group members. When a standby manager is told to take control, it attempts to update all members and the current active Group Manager so that group operations can proceed smoothly.

However, in the case of a failure, it is possible that the `icapmanage -Q` command is unable to contact the active Group Manager and some members of the groups that it now manages. When this happens, the previously active Group Manager remains active, unaware of the change of control. This is referred to as a “bifurcated” (or “split”) GiCAP group. Members that were reachable by the standby Group Manager when it took control cannot accept commands from the old active Group Manager; but unreachable members continue to consider it active. Control operations can be carried out on both active Group Managers, each communicating with the members that it (and only it) can reach. Groups and members can be added or removed on both (subject to the set of members each can command), and sharing rights can be added on both. In some cases this can be valuable; for example, when two data centers each remain functional but some intervening network link has been broken. Each isolated set of systems can proceed with independent disaster recovery operations within their group subset.

At some point, communication is restored and the split groups must be rejoined. This is accomplished through issuing a new `icapmanage -Q` command. It can be executed on either active Group Manager to confirm that Group Manager as the active Group Manager and demote the other to standby status. Be aware that doing this loses all database changes made on the demoted Group Manager during the time that the group was split. There is no method to merge the two databases, and in particular any new sharing rights applied to the Group Manager designated now as standby are lost.

## Additional GiCAP Considerations

Systems that have no Instant Capacity components can be part of a GiCAP group. Deactivating resources on these systems allows them to loan usage rights to other members in the group.

Members of a GiCAP group do not have to be located near each other. The only constraints are IP connectivity between the members, the Group Manager, and the standby Group Manager (if any), sufficient GiCAP sharing rights, and adherence to the GiCAP grouping rules.

For systems with multiple network cards, you can specify the additional network paths as additional hosts when defining member systems in a group, for enhanced connectivity. However, there might be a significant performance penalty because the Instant Capacity software occasionally must access all the multiple hosts in that case. You cannot specify the Group Manager and the standby Group Manager such that they are different network paths to the same system. An OS instance can host one and only one GiCAP database, regardless of the number of hostnames by which that OS instance might be reachable.

WBEM version A.02.05 or higher must be installed on the Group Manager and on all member systems in order to use Global Instant Capacity. The CIM Server configuration property `sslClientVerificationMode` must be set to a value of “optional” on all GiCAP Group Managers and on all OS instances of all member systems. (The CIM Server may need to be restarted if the property was not previously set to this value.) For details, see *cimconfig(1M)*. Note that communication between the managers and members of groups is established using SSL certificates that are supplied by the GiCAP software.

## Compliance

In general, a complex is in a compliant state when the number of active components of a given type does not exceed the number of usage rights associated with the type of component. The one exception is that the number of active cores is allowed to exceed the number of core usage rights as long as there is a sufficient positive balance of temporary capacity. A negative balance always indicates a system which is out of compliance.

A GiCAP group is in a compliant state as long as all the members are in a compliant state, all the members of the group continue to have compatible hardware as determined by the hardware grouping rules, and the number of sharing rights installed on the GiCAP manager is equal or greater than the total number of cores without usage rights on complexes managed by the Group Manager.

## SEE ALSO

*icapmodify(1M)*, *icapnotify(1M)*, *icapstatus(1M)*, *icapmanage(1M)*, *icapd(1M)*, *parmgr(1M)*, *parmodify(1M)*, *parolrad(1M)*, *vparmodify(1M)*.

# icapmanage(1M)

## NAME

icapmanage - Global Instant Capacity (GiCAP) management commands for GiCAP groups.

## SYNOPSIS

```
icapmanage -i -U <rule_file>
icapmanage -C <codeword>
icapmanage -a -g <group_name>
icapmanage -r -g <group_name>
icapmanage -T <hostlist> [-g <group_name>]
icapmanage -a -m <member_name>:<hostlist> -g <group_name>
icapmanage -r -m <member_name>
icapmanage -s -g <group_name> [-b] [-v]
icapmanage -Q [-n]
icapmanage -R [<hostlist>] [-U <rule_file>]
icapmanage -a -S <host>
icapmanage -r -S <host>
icapmanage -t
icapmanage -u -m <member_name> -h [<hostlist>] [!<hostlist>]
icapmanage -x <host>
icapmanage -z <host>
```

## DESCRIPTION

A Global Instant Capacity (GiCAP) group consists of a Group Manager system and one or more partitionable complexes which are the members. The primary purpose of a GiCAP group is to enable group members to share Instant Capacity usage rights (for cores, cell boards, and memory) and temporary capacity. A Group Manager has either active or standby status. An active Group Manager coordinates and monitors resource sharing, has had hardware grouping rules applied, and has had GiCAP sharing rights purchased for and applied to it. Optionally, an active Group Manager can designate as a standby Group Manager any HP-UX system that runs Instant Capacity software but has not itself had GiCAP sharing rights applied. The standby Group Manager has the ability to take control as an active Group Manager if the original active Group Manager should become unusable for any reason.

Because the members of a group are partitionable complexes and GiCAP must be able to contact all of the partitions, GiCAP members must be defined by specifying a `<hostlist>` to describe the host OS instances of all the partitions.

The symbol `<hostlist>` is a string of one or more host names or IP addresses with the form `<host> [, <host>] . . .`

A host can be a vPar or nPar OS instance, but it cannot be a virtual machine (“guest”). When using HPVM on a prospective GiCAP group member, you must specify the name of the HPVM host OS instance, not the name of an HPVM guest OS instance.

The `icapmanage` command, which is used to create, manage, and remove groups, can be invoked only on a Group Manager system. On a standby Group Manager, only the `-Q` and `-s` options are allowed, all other options are restricted to use only on the active Group Manager. The `icapmanage` command should not be invoked on a group member system.

The `icapmanage` command can be used to install a grouping rules file, apply a GiCAP sharing rights codeword, create and remove GiCAP groups, test if a server can be added to a GiCAP group, update a GiCAP group by adding or removing members, show grouping rules and supported hardware, seize core usage rights from member partitions of a GiCAP group to be used by another member of the group, restore seized core usage rights to the original member partition, update an existing GiCAP member by adding or removing hosts, set up and activate a standby Group Manager, transfer the GiCAP database to the standby Group Manager, and remove a standby Group Manager.

WBEM version A.02.05 or higher must be installed on the Group Manager, the standby Group Manager (if any), and on all OS instances of all member systems in order to use Global Instant Capacity. Similarly, the CIM Server configuration property `sslClientVerificationMode` must be set to a value of "optional" on all Group Managers and member OS instances. (The CIM Server may need to be restarted if the property was not previously set to this value.) For details, see `cimconfig(1M)`. Note that communication between the managers and members of groups is established using SSL certificates that are supplied by the GiCAP software.

For a complete overview of Global Instant Capacity, see `icap(5)`, and for more detailed information see the *HP Instant Capacity user's guide for Version 10.x* located at `/usr/share/doc/icapUserGuide.pdf`.

## Options

`icapmanage` recognizes these options and arguments:

- a Add a GiCAP group, add a member to a group, or designate a system as a standby Group Manager.  
To create a new group, use the `-a` option, along with the `-g` option to name the new group.  
To add a member to a GiCAP group, use the `-a` option, along with the `-m` option to specify a member name and list of hostnames, and the `-g` option to specify the group name. The list of hostnames must contain a representative host for each nPartition or virtual partition of the complex, and each host must be up and the Group Manager must be able to contact it. (To avoid potential problems, the Group Manager should also be able to contact any standby manager that is defined.) To add multiple members to a GiCAP group, you must run `icapmanage -a -m -g` multiple times; you cannot specify multiple members in a single command.  
To add a standby Group Manager, use the `-a` option, along with the `-S` option to identify the host system that is to serve as the standby Group Manager.
- b Provide brief status information only. Show summary group information without information specific to the Group Manager and the members.
- g *<group\_name>* Specify a GiCAP group name for a GiCAP operation.
- h Used with the `-u` option to identify hosts to be added to or removed from an existing GiCAP group member's list. A hostlist identifies hosts to be added. A hostlist preceded by an exclamation point (!) identifies hosts to be removed. There must be no intervening blanks within a hostlist, and if both add and remove lists are specified, the entire string must be contiguous.

- i Install a grouping rules file on a Group Manager system.
- m Add a member (a partitionable complex) to a group, with name *member\_name*. Specify an OS instance (*host*) for each nPartition or virtual partition of the complex (do not specify virtual machine or guest OS instances).  

<member\_name>:<host>[, <host>] . . .

A member of a group must encompass all nPar and vPar OS instances of a complex, and each OS instance specified as a *host* must be accessible (ping-able) in order for the command to succeed. When you first add a member to a group, you are prompted for the root password for each specified host. The password is used for initial communication only and is not saved or stored.

If you are using the HP Virtual Server Environment (VSE), it may be helpful to define the *member\_name* to be the same as the VSE complex name (which includes the serial number).

A member complex can be added to a group if it is not in an exception state, if there are enough GiCAP sharing rights available to match the number of cores without usage rights on that complex, and if the complex's hardware is compatible with the grouping rules and with the other members of the group.
- m <member\_name> Specify the member name when removing a member from a GiCAP group.
- n Prevents `icapmanage -Q` from attempting to exchange SSL keys with member hosts which the new active Group Manager cannot contact. This option allows `icapmanage -Q` to be issued from a script without the script having to deal with the possibility of root password prompts. If `-n` is specified and there are member hosts which the new active group manager cannot contact, the active manager will not be able to use those hosts for GiCAP operations. If a GiCAP group member has no hosts which the active manager can contact, that member will not be able to participate in the GiCAP group.
- r Remove a member from a group, remove a GiCAP group, or remove a system from use as a standby Group Manager.  

When used with the `-m` option to specify the member name, removes that member from a GiCAP group. In order to do the remove, at least one host defined for the member must be up and the Group Manager must be able to contact it. (To avoid potential problems, the Group Manager should also be able to contact any standby manager that is defined.) Note that a member cannot be removed from a group until any borrowed usage rights are returned to the group and any loaned usage rights are returned to the member. Removal of a member from a group releases sharing rights and makes them available for future use.

When used in combination with the `-g` option, removes the specified GiCAP group. All members must be removed before the group can be removed.

When used in combination with the `-S` option, removes the identified host system from use as a standby Group Manager.

- `-s` Request status about one or more GiCAP groups. Specification without any additional options displays group and member information for all GiCAP groups managed by this Group Manager. Use the `-g <group_name>` option to limit the information to the named group only. Use `-b` to display group-level information only. Use `-v` to include additional manager-specific information and more detailed group and member information, especially partition-specific information for members. For more information about fields that are displayed see “Status Information”.
- `-t` Transfer a copy of the GiCAP database from the active Group Manager to a standby Group Manager. The active Group Manager transfers the GiCAP database to the standby Group Manager whenever the active Group Manager makes changes to the database. In addition, the Instant Capacity software uses this command as part of an automatically generated `cron` job (see `cron(1M)`) to ensure that this copy takes place if something should prevent normal replication of the GiCAP database to the standby Group Manager. While you should not need to issue this command, you may wish to do so occasionally, especially if there are intermittent problems with network communications between the active Group Manager and the standby Group Manager. You may also wish to issue this command before having the standby Group Manager take over (via the `icapmanage -Q` command on the active Group Manager) for planned downtime on the active Group Manager. The transfer operation works only on the active Group Manager system.
- `-u` Update the list of hosts (OS instances) associated with an existing GiCAP group member. This command is typically used when the vPar or nPar configuration of a member complex changes. The `-m` option identifies the member to be updated. The `-h` option identifies the hosts to be added or removed.

**Note:** You are prompted for the root password for each host to be added. The password is used for initial communication only and is not saved or stored.

For a given command execution, host removals are processed before host additions. As a result, the same host can be removed and readded with one update operation. With removes or adds applied, the resulting list must contain a representative host for each nPartition or virtual partition of the associated server.

When adding hosts, each host must be up and the Group Manager must be able to contact it. (To avoid potential problems, the Group Manager should also be able to contact any standby manager that is defined.)

-v

Provide verbose status information. Include all levels of information (group, manager and member). For Group Managers, include resources being held by the Group Manager including temporary capacity. For members, include borrow and loan information as well as partition-specific information such as the allocation of resources among the hard partitions, and partition-specific information about seized or seizable usage rights (see `icapmanage -x`). This option is ignored if the `-b` option is also specified.

This verbose option also attempts to contact every host of each group member to check that two-way communication can be established between the issuing group manager (either active or standby) and each such host. The `icapmanage -s` command without `-v` attempts to contact multiple hosts on a group member only if that is necessary to find a host that responds. In either case, `icapmanage -s` reports which hosts it failed to contact. Use of the `-v` option can be useful as a general health check of group communication.

-x <host>

Acquire core usage rights from the nPartition containing the specified host to make them available to other group members, also known as rights seizure. Rights seizure takes almost all the available usage rights from the hard partition containing the specified host. The specified host must be known to the GiCAP manager (it appears in the output of `icapmanage -s`) and not currently running. If other hosts are associated with the same hard partition, they cannot be currently running. The `icapmanage -x` operation can be performed once for each hard partition on the member. For confirmation, the `icapmanage -x` command verifies each known host on the hard partition. The hard partition containing the specified host is left with one core usage right per active cell. Any core usage rights in excess of this amount becomes available for use elsewhere in the GiCAP group. For example, if a partition of 2 cells is currently consuming 6 core usage rights, `icapmanage -x <host>` makes 4 core usage rights (6 - 2) available for reassignment.

While rights can be seized from any hard partition that is unavailable, the Instant Capacity software makes some additional restrictions when all partitions of a complex are unavailable. As a result, there are different behaviors and constraints depending on whether or not a partition can be contacted on the specified member complex.

If, at the time of rights seizure, at least one member partition can be contacted, then the software is able to make an immediate adjustment to the available core usage rights, just as if an `icapmodify -d` operation had been performed before the specified hard partition stopped running. This makes core usage rights available for potential loans to other member systems. In this situation, the seized core usage rights do not have an expiration date. However, because there are other member partitions

running iCAP software, the unreachable partition may be assumed to be using all cores on cells configured for that partition. Because of this, cells in partitions from which usage rights have been acquired should be rebooted or made inactive within 12 hours. If this is not done, the partition can begin to consume temporary capacity. If temporary capacity is not available, the complex might no longer be in compliance with the iCAP contract. Cells can be made inactive by removing them from the partition, shutting down the partition from within the OS by using `shutdown -R -H`, or with the `MP RR` command.

If, at the time of rights seizure, all member partitions are unreachable, the rights seizure is deferred and must be viewed as a limited and immediate loan of usage rights from the specified partition to the group. This loan of seized usage rights expires in 10 days. Upon expiration, usage rights are automatically restored to the member partitions from which they were seized. The expiration date for a rights seizure operation effectively terminates the period during which the core usage rights are available to other group members for purposes of disaster recovery. If none of the member partitions are reachable by the expiration date for a particular member, the usage rights are automatically restored (reassigned) to the member partition (or complex, in the case of unassigned seized rights) from which they were seized. However, if the seized usage rights have been redeployed to other members and are not released at expiration time, the group might go out of compliance, or temporary capacity might be used to maintain compliance.

If any partition of the inaccessible member from which rights seizures were deferred reconnects to the group before the expiration date, then the seized core usage rights (for all partitions) are finalized as a loan from the member to the group, the expiration date is no longer relevant, and the usage rights can thereafter be manipulated with normal `icapmodify` operations.

While rights seizure operations can be performed in a virtual partition environment, the rights seizure always operates on, and affects, the entire nPartition. This means:

- Rights seizure can be performed only if all the virtual partitions for an nPartition are down.
- For a given nPartition, you can specify any of the virtual partition host names as the target of a rights seizure operation or as the target of a usage rights restore operation.
- Because rights seizure leaves only a minimum of core usage rights with the nPartition, it is likely that the remaining number of core usage rights is not sufficient to satisfy the number of cores assigned to each virtual partition in the nPartition. This means that the virtual partitions likely cannot be booted (due to noncompliance) once the original failure is corrected.

To avoid this problem, usage rights must be restored (using the `-z` option) before failback.

`-z <host>`

Restore previously seized core usage rights to the nPartition containing the specified host. Core usage rights must be available in the GiCAP group or the command fails.

This option can be useful particularly when core usage rights have been seized from systems running vPars, because the restoration of core usage rights may be necessary in order to be able to reboot the vPar, depending on the vPars database definition and the allocation of usage rights among the vPars. If this is a potential problem, the restoration command must be performed before rebooting any vPar within an nPar from which rights were seized.

`-C <codeword>`

GiCAP codeword application. This option allows the user to apply a GiCAP codeword to a Group Manager system. (This option cannot be used to apply an iCAP codeword such as an RTU or TiCAP codeword.) For GiCAP sharing rights codewords, first purchase the GiCAP codeword from HP. The number of rights purchased must equal or exceed the number of cores without usage rights for all planned members for all groups managed by the Group Manager. Next, retrieve the codeword from the HP Utility Pricing Solutions portal and apply it to the Group Manager system. Unlike iCAP codewords, GiCAP codewords are generated for a specified partition on a Group Manager system, and can only be applied to that partition. Like iCAP codewords, GiCAP codewords are also generated in a sequence and must be applied in the order they are generated for the Group Manager partition. However, GiCAP codewords are sequenced independently from any iCAP codewords for the same complex, and can be applied independently from any such iCAP codewords. Application of the GiCAP codeword allows members to be added to one or more GiCAP groups.

`-Q`

Make a standby GiCAP Group Manager the active Group Manager for all group members. The new active Group Manager attempts to contact all group members, informing them that this system is the active Group Manager. The new Group Manager also attempts to contact the previous Group Manager to make it the standby Group Manager. If the new Group Manager cannot contact a group member, that group member will continue to be managed by its current manager. This may result in the member not being able to loan or borrow usage rights or it may result in a split group. The current system becomes the active Group Manager regardless of whether these attempted contacts succeed or fail.

When attempting to contact all group members, the new active Group Manager may find that it must establish SSL communication with a member host. If so, it will prompt for that host's root password so that it can exchange SSL keys with the host (unless the `-n` option is also specified). HP recommends that SSL communication between the

standby Group Manager and all member hosts has been previously established before the use of the `icapmanage -Q` command. Normally this occurs automatically when adding a new member to a group, updating the host list for an existing member, or adding a standby Group Manager.

This command can be issued on an active Group Manager. In this case, the active Group Manager continues as the active Group Manager. It attempts to contact all group members and the previously active Group Manager and informs them that this system is the active Group Manager. This is useful if a prior `icapmanage -Q` command failed to contact some group members or the previously active Group Manager.

Note that if the active Group Manager manages multiple groups, all members of all such groups must be contacted during this operation.

`-R [<host>[, <host>] ...]`

Report hardware grouping information. When used with a list of host names, report all hardware types with which the hosts might be grouped. The hosts can be specified using either the IP address or the name of the host. If the hosts are not compatible with each other, no hardware is reported. Without a list of host names, report all supported hardware and grouping rules. Specification of the `-U` option reports hardware associated with the specified rule file instead of the installed rule file, and can be used to evaluate a new grouping rules file before installing it with the `-i` option.

`-S`

When used in combination with the `-a` option, identifies a host system to be set up as a standby Group Manager. This command will prompt for the root password of the host so that the active Group Manager can exchange SSL keys with the designated standby host, establishing the ability to communicate with it. The active Group Manager proceeds to update all member hosts of all managed groups with information about the standby Group Manager, establish secure communications between the member hosts and the standby Group Manager, and set up a `cron` job responsible for the periodic transfer of the GiCAP database to the standby Group Manager. Note that a transfer also occurs whenever the database is updated on the active Group Manager until such time as the standby Group Manager takes control or is removed from use as a standby Group Manager.

When used in combination with the `-r` option, identifies a host system which is to be removed from use as a standby Group Manager. The active Group Manager updates all member hosts of all managed groups about the removal, discontinues the GiCAP database transfer operations (removes the `cron` job created when the standby Group Manager was set up), and directs the identified host system to remove its copy of the GiCAP database.

A standby Group Manager can be added anytime after the active Group Manager has grouping rules installed on it (before or after the application of sharing rights to the active Group Manager, before or after groups are created). A Group Manager in standby status can be removed at any time.

-T <host>[, <host>] . . .

Test hardware compatibility for one or more host systems in order to determine which groups the systems can join. When used with the -g option to specify a group name, tests whether the specified host systems have hardware compatible with the group. Without the -g option, report which of all the groups managed by this Group Manager have hardware compatible with the host systems. A host can be specified using either the IP address or the name of the host. The host names do not have to be from the same complex, but to best predict the possibility of joining a group, the list of hosts should include all the nPartitions for a particular complex. If the hosts are not compatible, no groups are reported as having compatible hardware.

-U <rule\_file>

Specify the filename of a rule file.

### Status Information

This section describes the fields that might be displayed when `icapmanage -s` is used to show status. The exact choice of options used in combination with the -s option determines how much of the information is displayed.

First, the display shows the software version number of the Global Instant Capacity software and the current date and time. This is followed by the identification information for the Group Manager: serial number, nPar ID (if any), and vPar code (if any). This identification information is necessary when requesting a sharing rights codeword from the portal. Next, the Group Manager status (active or standby) is displayed. On an active Group Manager, this is followed by the name of the dedicated standby Group Manager (if any). On a standby Group Manager, this is followed by the name of the associated active Group Manager. Next, the display shows the number of sharing rights that have been purchased and applied to this Group Manager, and how many sharing rights are currently in use versus the number still available to accommodate addition of new members or new groups with new members.

### Symbols used in the display

The output of the `icapmanage -s` command uses symbols next to the value of some fields in certain circumstances. These symbols are as follows:

- \* Indicates an estimated number of active cores, memory or cells, because full information is not available (for example, when a host is not reachable, or on some platforms when cells are powered off).
- # Indicates that all cores in the partition, rather than any specific number, have been requested to be active. The number indicates the current number of cores in the partition.

### Information displayed for each GiCAP group

These values can be displayed for each group managed by the Group Manager.

Group <group\_name>: This field displays the name of the GiCAP group.

Group Members: This summarizes the name of each member in the group, and also shows the host names comprising each member complex.

Instant Capacity Resource Summary for the group	This section shows values which are summed across all group members.
Number of cells without usage rights:	This field displays the total number of cells across all group members which must remain inactive because usage rights have not been purchased.
Number of inactive cells:	This field displays the actual number of inactive cells across the group. It does not include counts for inactive cells associated with inaccessible members (no partitions could be contacted).
Amount of memory without usage rights:	This field displays the total amount of memory which must remain inactive across all group members because usage rights have not been purchased.
Amount of inactive memory:	This field displays the actual amount of inactive memory across the group. It does not include counts for inactive memory associated with inaccessible members (no partitions could be contacted).
Number of cores without usage rights:	This field displays the total number of cores which must remain inactive across all group members because usage rights have not been purchased.
Number of inactive cores:	This field displays the actual number of inactive cores across the group. It does not include counts for inactive cores associated with inaccessible members (no partitions could be contacted).
Number of cores using temporary capacity:	This field displays the number of cores using temporary capacity within the group. It does not include values for unreachable members (no partitions could be contacted).
Temporary capacity available:	This field displays the amount of pooled group temporary capacity that is available to any member of the group.
Projected temporary capacity expiration:	This field displays the date and time that temporary capacity is projected to expire for the group at the present consumption rate.
Temporary capacity balance:	This field displays the total amount of temporary capacity assigned to the group. This value is different from "Temporary capacity available", and is displayed only when members with temporary capacity cannot be contacted by the group manager (their temporary capacity is not available to the group but is part of the group total).
Unassigned cell usage rights:	This field displays the number of cells that can be activated immediately in the group because of usage rights that are not in use. This field is displayed only when the -v option is specified.
Unassigned memory usage rights:	This field displays the amount of memory that can be activated immediately in the group

because of usage rights that are not in use. This field is displayed only when the `-v` option is specified.

Unassigned core usage rights: This field displays the number of cores that can be activated immediately in the group because of usage rights that are not in use. This field is displayed only when the `-v` option is specified.

Seized core usage rights: This summary field displays the number of expiring core usage rights that were seized from inaccessible members (see `icapmanage -x`). This field is displayed for all status options (`-s`, `-b`, `-v`), but only when usage rights were seized from member systems where none of the partitions could be contacted.

Expiration of seized core usage rights: This summary field displays the earliest expiration date for core usage rights that were seized from members where none of the partitions could be contacted (see `icapmanage -x`). The expiration date for a rights seizure operation effectively terminates the period during which the core usage rights are available to other group members for disaster recovery. This field is displayed for all status options (`-s`, `-b`, `-v`), but only when usage rights were seized from member systems where none of the partitions could be contacted.

Information Displayed for the Group Manager This section describes resources that are temporarily held by the Group Manager. The resources are available to all members of the group and include cell, memory, and core usage rights, and temporary capacity. Manager-held resources usually represent a transient state, for example when usage rights have been released (or seized using `icapmanage -x`) from one member of a group but have not yet been activated on another target member system. These values are displayed only when the `-v` option (verbose) is specified.

Information displayed for each member of the group This section is repeated for each member of the group if the `-b` option is not specified. With a few exceptions, the information in this section consists of a subset of the information provided by the `icapstatus` command on the member systems. In all cases, a Resource Summary minimally includes the values for cells, memory, and cores without usage rights for the member system, and the balance of temporary capacity assigned to the member complex. If any partition of the complex can be contacted, the resource summary also contains counts of inactive cells, memory and cores, the count of cores using temporary capacity, and the projected date of temporary capacity expiration. If the `-v` option is specified and any partition of the member can be contacted, the member resource summary includes counts of borrowed cell, memory, and core usage rights, and partition-specific information is included in the "Allocation of Instant Capacity Resources among the partitions" section.

If core usage rights have been seized from the member, the count of cores without usage rights is adjusted by the number of seized usage

rights, to reflect the overall balance of cores without usage rights that must be maintained in the group.

If none of the partitions of the member can be contacted, then the Resource Summary contains limited information, using the last-known values for the member. Only the fields listed for unreachable members have values that are included in the group totals. For example, the group total for inactive cores does not include counts for unreachable members, but the group total for number of cores without usage rights includes last-known values for unreachable members.

Additionally, information about rights seizure can be displayed for members where none of the partitions can be contacted:

**Seized core usage rights:** This summary field displays the total number of expiring core usage rights that were seized from all partitions of this member (using `icapmanage -x` operations when none of the partitions could be contacted). This member summary field is not displayed when the `-v` option is specified (because the `-v` option includes a detailed list instead).

**Expiration of seized core usage rights:** This summary field displays the earliest expiration date for core usage rights that were seized from all partitions of this member (using `icapmanage -x` operations). The expiration date for a rights seizure operation effectively terminates the period during which the core usage rights are available to other group members for disaster recovery.

This member summary field is not displayed when the `-v` option is specified. However, if multiple rights seizure operations occurred for this member, there might be multiple expiration dates. To see a more detailed list of partition-specific rights seizure counts and expiration dates, use the `-v` option.

**Unassigned core usage rights seized:** This field displays the number of unused core usage rights (not previously assigned to any partition) that were seized from this member. This field is displayed only when the `-v` option is specified.

You can use `icapmanage -z` to restore core usage rights that were previously seized from partitions. It does not restore previously seized core usage rights if those rights were not assigned to a partition at the time of seizure. In other respects however, these unassigned seized usage rights are similar to usage rights seized from partitions: if the member reconnects to the group manager before the expiration date, the seized usage rights are completed as a loan to the group. If the expiration date occurs before reconnection, the seized usage rights revert to

	the member system from which they were seized.
Core usage rights seized from npar <n>:	This field displays the number of expiring core usage rights that were seized from a specific partition of the member. This field is displayed only when the -v option is specified, and it is repeated for each partition where usage rights were seized for disaster recovery, if none of the member partitions could be contacted.
Expiration:	This field displays the expiration date for core usage rights that were seized with icapmanage -x operations from the preceding named partition, or from the complex in the case of unassigned usage rights. The expiration date for a rights seizure operation effectively terminates the period during which the core usage rights are available to other group members for purposes of disaster recovery.  This field is displayed only when the -v option is specified and only for systems where usage rights were seized for disaster recovery, if none of the member partitions could be contacted.

### Guidelines for Interpretation of Usage Rights Counts

In general, the Instant Capacity software is designed to enforce the concept that a certain number of resources must remain inactive in order to stay in compliance with the iCAP contract. For that reason, many of the display fields for `icapstatus` and also for `icapmanage` are negative counts of usage rights, such as Cores without usage rights and Cells without usage rights.

At the same time, the GiCAP software is intended to allow the transfer of usage rights among members of the group. So, seizure of usage rights, borrows and loans of usage rights are described as positive counts of usage rights, such as Number of core usage rights seized and Borrowed core rights. The Group Manager also sometimes holds usage rights during transient states, so all the Group Manager resources are described as positive counts of core usage rights.

To properly interpret the `icapmanage` display of usage rights, it is sometimes necessary to total the usage rights using both positive and negative (subtractive) values. For example, if the group contains two members, each having 3 cores without usage rights, then the total for the group is 6 cores without usage rights. If a rights are seized from one member, for example taking 2 core usage rights away, then those 2 core usage rights are held by the Group Manager before the 4 are used by another member of the group. During that transitional period, the total for the group remains 6 cores without usage rights, but the member total for the unreachable member is adjusted to be 5 cores without usage rights (2 were taken away from the unreachable member) and 3 cores without usage rights (for the other member). At the same time, the Group Manager holds 2 usage rights. The total for cores without usage rights is calculated as  $6=5+3-2$ .

In this example, “counts of cores without usage rights” are also adjusted as a result of rights seizure operations in order to maintain the contractually required balance of cores without usage rights. Similar adjustments are made every time a resource (cores, cells, memory) is borrowed from or loaned to other members of the group. These adjustments are automatically reflected in the `icapstatus` output on the member systems; in the case of unreachable members and rights seizure, you might see the adjustment first in the display from the `icapmanage` output.

## EXTERNAL INFLUENCES

### Environment Variables

- LANG determines the locale to use for the locale categories when both LC\_ALL and the corresponding environment variable (beginning with LC\_) do not specify a locale. If LANG is not set or is set to the empty string, a default of "C" is used (see *lang(5)*).
- LC\_CTYPE determines the interpretation of single- and multiple-byte characters.
- LC\_MESSAGES determines the language in which messages are displayed.

If any internationalization variable contains an invalid setting, `icapmanage` behaves as if all internationalization variables are set to C (see *environ(5)*).

### International Code Set Support

Single- and multiple-byte character code sets are supported.

### RETURN VALUE

`icapmanage` exits with one of these values:

- 0 Command succeeded.
- >0 Command failed; error message sent to STDERR.

### FILES

<code>/var/adm/GiCAP.log</code>	Log file for GiCAP operations and messages.
<code>/etc/opt/iCAP/GiCAP.rules</code>	Encrypted file containing grouping rules used by the Group Manager.
<code>/etc/opt/iCAP/GiCAP_db</code>	Encrypted file containing information about sharing rights and information about each group managed by the Group Manager.
<code>/etc/opt/iCAP/GiCAP.configFile</code>	Present on every host of a GiCAP group member. It contains information needed by the iCAP software to contact the Group Manager for the host.
<code>/etc/opt/iCAP/GiCAP_keygen</code>	Script used to establish (or reestablish) SSL certificates.
<code>/etc/opt/iCAP/GiCAPcert.pem, /etc/opt/iCAP/* .pem, /etc/opt/iCAP/GiCAPpkey.pem</code>	SSL certificate files.
<code>/etc/opt/iCAP/GiCAP.checkScavengedMember.xml</code>	Script used by the Group Manager to periodically check for members reconnecting to the group after a seizure operation.

### EXAMPLES

Install a new grouping rules file:

```
icapmanage -i -U /tmp/GiCAP.rules
```

Purchase a sharing rights codeword from HP with rights equal to or greater than the number of cores without usage rights for all planned members of the group. Retrieve the codeword from the portal, and apply the sharing rights codeword to the Group Manager system.

```
icapmanage -C R8J2DBW.5UTxyWQ.2MekJ43.G5cdTVP.1-m9kvweQ.AYqEXym.wj3dyLj.Fbtg7s1
```

Create a GiCAP group named ADMIN1:

```
icapmanage -a -g ADMIN1
```

Test whether a server complex has hardware that is compatible with the group:

```
icapmanage -T mypar1.node.hp.com,mypar2.node.hp.com -g ADMIN1
```

Add a member called IT to the ADMIN1 group. Supply the root password for each of these partitions in response to the prompts:

```
icapmanage -a -m IT:mypar1.node.hp.com,mypar2.node.hp.com -g ADMIN1
root@mypar1.node.hp.com's password:
root@mypar2.node.hp.com's password:
```

Designate host mystandby for use as a standby Group Manager:

```
icapmanage -a -S mystandby.node.hp.com
```

Remove host mystandby from use as a standby Group Manager:

```
icapmanage -r -S mystandby.node.hp.com
```

Take control as an active Group Manager:

```
icapmanage -Q
```

Show the full status of the ADMIN1 group:

```
icapmanage -s -g ADMIN1 -v
```

Seize core usage rights from a partition that is unavailable to make them available to other group member activations:

```
icapmanage -x mypar1.node.hp.com
```

Restore previously seized usage rights to the partition from the previous example:

```
icapmanage -z mypar1.node.hp.com
```

Report supported hardware and grouping rules for a specific grouping rules file:

```
icapmanage -R -U /tmp/GiCAP.rules
```

Add host mypar3 to existing member IT:

```
icapmanage -u -m IT -h mypar3.node.hp.com
```

Add host mypar4 to existing member IT, and remove host mypar3:

```
icapmanage -u -m IT -h mypar4.node.hp.com!mypar3.node.hp.com
```

Remove group member IT from its group:

```
icapmanage -r -m IT
```

Remove the ADMIN1 group:

```
icapmanage -r -g ADMIN1
```

## AUTHOR

icapmanage was developed by HP.

## SEE ALSO

icapmodify(1M), icapnotify(1M), icapstatus(1M), icapd(1M), icap(5), *cimconfig*(1M), *cron*(1M).

# icapmodify(1M)

## NAME

`icapmodify` - Activate and deactivate cores. Specify system contact email address. Change Instant Capacity configuration information. Specify Instant Capacity from email address. Specify system identifier. Specify temporary capacity warning period. Apply codewords.

## SYNOPSIS

```
icapmodify -c <contact_email_address>
icapmodify -C <codeword>
icapmodify -f <from_email_address>
icapmodify -i <system_id>
icapmodify -r
icapmodify -w <warning_days>
icapmodify -a <n> [-D] [-t] [desc[:user_name]]
icapmodify -d <n> [-D] [desc[:user_name]]
icapmodify -s <n> [-D] [-t] [desc[:user_name]]
```

Obsolescent:

```
icod_modify -c <contact_email_address>
icod_modify -C <codeword>
icod_modify -f <from_email_address>
icod_modify -i <system_id>
icod_modify -r
icod_modify -w <warning_days>
icod_modify -a <n> [-D] [-t] [desc[:user_name]]
icod_modify -d <n> [-D] [desc[:user_name]]
icod_modify -s <n> [-D] [-t] [desc[:user_name]]
```

## DESCRIPTION

Use `icapmodify` to activate or deactivate cores, specify system contact or Instant Capacity “from” email address, apply iCAP codewords, change the system identifier, specify a warning notification period before temporary capacity expires, and change Instant Capacity configuration information.

Note that the deprecated `icod_modify` command performs identical functions to the `icapmodify` command and is maintained for backward compatibility.

For detailed information on the use of this command, activation and deactivation of Instant Capacity components, compliance, and temporary capacity, see the *HP Instant Capacity user's guide for Version 10.x* located at `/usr/share/doc/icapUserGuide.pdf`.

## Compliance

The `icapmodify` command does not allow activation of cores beyond the number of available core usage rights. Additional usage rights are granted through the application of either an RTU codeword or a temporary capacity codeword. In general, a complex is in a compliant state when the number of active components of a given type does not exceed the number of usage rights associated with the type of component. The one exception is that the number of active cores is allowed to exceed the number of core usage rights as long as there is a sufficient positive balance of temporary capacity.

## Intended Active

Changes to the number of `intended active` cores through the use of this command are persistent (survive system reboot). The `intended active` number is the number of cores that the Instant Capacity software attempts to activate at system boot time. It is adjusted by use of the `-a`, `-d`, and `-s` options. The number of `intended active` cores for each partition is displayed using the `icapstatus` command (see `icapstatus(1M)`).

## Virtual Partitions

When activating or deactivating cores within virtual partitions, special considerations apply. You can use either the `icapmodify` command or the `vparmodify` command, depending on the type of adjustment needed and the level of logging or reporting desired. For example, core assignment via the `vparmodify` command does not result in logging of the activation, email configuration change notification, or transmission of an asset report to HP.

Instant Capacity has a minimum version dependency on vPars A.03.05. For versions of vPars before A.03.05, the `icapmodify` command for activating or deactivating cores in a virtual partition fails with an error message citing the vPar version dependency.

For vPars versions A.03.05 or greater, the `icapmodify` command must be used in a virtual partition environment when you are making any adjustment to an nPartition. To adjust core assignments across virtual partitions in a single nPartition, use the `vparmodify` command (`-a` and `-d` options) for the best coordination and for optimized performance. The `vparmodify` command does not affect the `intended active` number for the nPartition, and it therefore cannot be used to migrate unused capacity either to or from other nPartitions. If `vparmodify` is used to deactivate cores in a virtual partition, and they are not activated in another virtual partition, the deactivation will not free usage rights and they are held as unused capacity by the nPartition.

## Options and Arguments

The `icapmodify` command recognizes the following options and arguments:

- `-a <n>` Immediately activates *n* additional cores for this nPartition, as long as the end result does not take the complex out of compliance. This option also increases the number of `intended active` cores by *n* for the nPartition. If specified within a virtual partition, it also results in the assignment of additional cores to the local vPar.
- `-c <contact_email_address>` Sets the system contact email address. This is the email address that receives configuration change notification and exception reports. If you want multiple recipients to receive these reports, that this can be an email alias.
- `-C <codeword>` iCAP codeword application. This option allows the user to apply an iCAP codeword received from the HP Utility Pricing Solutions portal. Application of codewords only provides usage rights for Instant Capacity components; it does not activate any components. This option cannot be used to apply GiCAP codewords. For details about GiCAP codewords, see `icapmanage(1M)`.
- `-d <n>` Immediately deactivates *n* cores, if possible. Instant Capacity software must leave at least one core active for each configured cell in a partition — this is a firmware and OS requirement. That is, in a partition of 4 cells, attempts to reduce the active core count below 4 will fail. This option also reduces the number of `intended active` cores by *n* for the nPartition. If specified within a virtual partition,

it deassigns the specified number of cores from the local vPar.

**-D** Defers a core activation or deactivation until the next reboot. This option modifies the default behavior of the `-d`, `-s`, and `-a` options, which is to activate or deactivate cores instantly.



---

**NOTE:** This option is not supported within a virtual partition.

Deferred operations are not cumulative. If there is a pending deferred operation, a subsequent activation or deactivation request (`-s`, `-a`, or `-d`), deferred or not, cancels the pending deferred request and resets the values for `intended active` and `actual active` based on the request and the current value for `actual active`.

---

**-f** *<from\_email\_address>* Set Instant Capacity “from” email address. Causes all Instant Capacity email correspondence from this system to appear to be sent from *from\_email\_address*. Specifying an empty string (“”) returns to default behavior, which is to send from the `adm` user on the local system. The address specified must be DNS resolvable by HP.

**-i** *<system\_id>* Set system identifier used during asset reporting. The default setting for the system identifier is the hostname of the Instant Capacity system. This value can be returned to the default setting by specifying an empty string (“”). The system identifier is a string that users specify to help track and distinguish their systems.

**-r** Reconcile. Activate or deactivate cores (subject to compliance limits) to bring the system to a state where the `intended active` number of cores are active. The `-r` option can also be used to undo a deferred (`-D`) operation, causing an immediate activation or deactivation.

**-w** *<warning days>* Set temporary capacity warning period to desired number of days. If not specified, the default warning period is 15 days. The Instant Capacity software calculates when the temporary capacity expires based on the current consumption rate. When the temporary capacity balance is projected to be depleted within the warning period, a warning message is sent by email to the system contact if specified, and root. If temporary capacity is depleted and you continue to have more active cores than core usage rights across the complex, on the next reboot of any partition in the complex the software automatically deactivates one or more cores in order to bring the complex into a more compliant state. Instant Capacity software deactivates as many cores as is necessary to either stop consumption of temporary capacity or to bring the partition to the minimum number of required active cores.

**-s** *<n>* Sets the number of active cores and the number of intended active cores to *n*, as long as the end result does not take the complex out of compliance. In an `nPartition` environment, depending on the value of *n*, this option

works exactly as the `-a` option (if  $n$  is greater than the current number of active cores), or exactly as the `-d` option (if  $n$  is less than the current number of active cores). Specifying a value of  $n$  less than the number of cells in a partition will fail. In a virtual partition environment, if unused capacity is not available, this option will use the value of  $n$  as the desired number of cores to be active in the local virtual partition. If unused capacity is available, `icapmodify` will first activate cores from unused capacity before increasing the active cores. For further details of using this option in an virtual partition environment, see the Virtual Partition and Unused Capacity sections of Chapter 4 of the *HP Instant Capacity user's guide for Version 10.x* located at `/usr/share/doc/icapUserGuide.pdf`.

The `icapmodify` command fails if it is unable to set the value exactly as requested. However, a failed request can still make a partial change to the number of intended active cores.

`-t`

Authorize use of temporary capacity. This option, combined with either the `-a` or the `-s` option, specifies that a core activation is allowed to consume temporary capacity. Temporary capacity is consumed when the number of active cores exceeds the number of core usage rights. It is no longer used when the number of active cores is decreased to no more than the number of core usage rights available to the complex. Use `icapmodify -d` or `-s` to reduce or stop the use of temporary capacity. It is not necessary to use the `-t` option when using the `-d` option. If a previous activation via `icapmodify` resulted in temporary capacity being consumed in a virtual partition environment, deactivating a core with a `vparmodify` command temporarily reduces the consumption of temporary capacity. A subsequent core activation using `vparmodify` increases consumption of temporary capacity if the activation results in more active cores than core usage rights.

`desc`

Optional description to help customers identify this configuration change. This description becomes part of the Instant Capacity logfile (`var/adm/icap.log`) entry documenting the activation or deactivation. This description is also contained in the configuration change notification email.

`user_name`

Optional string identifying the person performing the core activation or deactivation. This can be any ASCII string, and becomes part of the Instant Capacity logfile (`var/adm/icap.log`) entry documenting the activation or deactivation. The string specified here is also contained in the configuration change notification email.

## UPGRADES

The `icapmodify` command fails if the system is in a state where a software upgrade is incomplete (the software on the system has been upgraded from a version earlier than B.06.00, but an upgrade codeword issued by the HP Utility Pricing Solutions portal was not applied to the complex). The

only option that can be used when the complex is in this state is the `-C` option, which accepts the upgrade codeword.

## EXTERNAL INFLUENCES

### Environment Variables

- `LANG` determines the locale to use for the locale categories when both `LC_ALL` and the corresponding environment variable (beginning with `LC_`) do not specify a locale. If `LANG` is not set or is set to the empty string, a default of `"C"` is used (see *lang(5)*).
- `LC_CTYPE` determines the interpretation of single- and multiple-byte characters.
- `LC_TIME` determines the date and time strings output.
- `LC_MESSAGES` determines the language in which messages are displayed.

If any internationalization variable contains an invalid setting, `icapmodify` behaves as if all internationalization variables are set to `C` (see *environ(5)*).

### International Code Set Support

Single- and multiple-byte character code sets are supported. However, input to the command must be entered using ASCII characters only.

## RETURN VALUE

The `icapmodify` command exits with one of these values:

- 0 Command succeeded.
- >0 Command failed; error message sent to `STDERR`.

## FILES

`/var/adm/icap.log`

## EXAMPLES

Instantly activate one core with "Add horsepower now" as the description and "Super User" as the user name:

```
icapmodify -a 1 "Add horsepower now:Super User"
```

Activate two cores (deferred until the next reboot) with "Add horsepower after reboot" as the description and "Super User" as the user name:

```
icapmodify -D -a 2 "Add horsepower after reboot:Super User"
```

Instantly activate one core, using temporary capacity if necessary, with "Temp use of one core" as the description and "Super User" as the user name:

```
icapmodify -t -a 1 "Temp use of one core:Super User"
```

Instantly activate or deactivate cores to specify 8 active cores (and 8 intended active cores) with "Set active cores to 8" as the description and "Super User" as the user name:

```
icapmodify -s 8 "Set active cores to 8:Super User"
```

Deactivate one core at the next reboot with "Less horsepower after reboot" as the description and "Super User" as the user name:

```
icapmodify -D -d 1 "Less horsepower after reboot:Super User"
```

Apply an iCAP codeword:

```
icapmodify -C 7y5ejVS.P5CuwXu.XaTyDVP.7Tx0Mvc-J783H9b.yWT5Weu.69JPu$u.vVV685a5
```

Set the Instant Capacity from\_email\_address to admin@research.corp.com:

```
icapmodify -f admin@research.corp.com
```

Set the system\_id to Asset\_Num\_234:

```
icapmodify -i Asset_Num_234
```

Set the system contact email address to super\_user@corp.com:

```
icapmodify -c super_user@corp.com
```

## **AUTHOR**

icapmodify was developed by HP.

## **SEE ALSO**

icapnotify(1M), icapstatus(1M), icapmanage(1M), icapd(1M), icap(5), *vparmodify*(1M).

# icapnotify(1M)

## NAME

`icapnotify` - Test email connectivity to HP for Instant Capacity (iCAP) systems. Request a confirmation response email from HP. Turn configuration change notification and asset reporting on or off.

## SYNOPSIS

```
icapnotify <reply_address>
```

```
icapnotify -a on/off
```

```
icapnotify -n on/off
```

Obsolescent:

```
icod_notify <reply_address>
```

```
icod_notify -a on/off
```

```
icod_notify -n on/off
```

## DESCRIPTION

When a *reply\_address* is specified, `icapnotify` sends an asset report via email to HP, root, and the specified email address. Confirmation email is sent from HP to the specified reply email address indicating that HP received the asset report email.

Note that asset reporting is optional but can be useful for viewing complexwide asset information at the HP Utility Pricing Solutions portal (<http://www.hp.com/go/icap/portal>).

The deprecated command `icod_notify` provides identical functionality to the `icapnotify` command and is maintained for backward compatibility.

For detailed information on email configuration and requirements, see the *HP Instant Capacity user's guide for Version 10.x* located at `/usr/share/doc/icapUserGuide.pdf`.

## Options

The `icapnotify` command recognizes these options and arguments:

`-a on/off` Turns email asset reporting on or off. This option is used to specify if the Instant Capacity software should send asset reports to HP via email. The `icapstatus` command displays the present setting for this option.

`-n on/off` Turns configuration change notification on or off. If specified *on*, executing the `icapmodify` command results in a configuration change notification email sent to the system contact email address summarizing any configuration change. Configuration change notification email can be turned off by specifying *off*. Configuration change notifications are not sent if the system contact email address is not set.

## EXTERNAL INFLUENCES

### Environment Variables

- `LANG` determines the locale to use for the locale categories when both `LC_ALL` and the corresponding environment variable (beginning with `LC_`) do not specify a locale. If `LANG` is not set or is set to the empty string, a default of `C` is used (see *lang(5)*).
- `LC_CTYPE` determines the interpretation of single- and multiple-byte characters.
- `LC_MESSAGES` determines the language in which messages are displayed.

If any internationalization variable contains an invalid setting, `icapnotify` behaves as if all internationalization variables are set to `C` (see *environ(5)*).

## International Code Set Support

Single- and multiple-byte character code sets are supported.

## RETURN VALUE

The `icapnotify` command exits with one of these values:

- 0 Command succeeded.
- >0 Command failed; error message sent to `STDERR`.

## EXAMPLES

Test email connectivity with HP by sending an asset report to HP, `root`, and `super_user@corp.com`, and request a confirmation email from HP to be sent to `super_user@corp.com`:

```
icapnotify super_user@corp.com
```

Turn email asset reporting on:

```
icapnotify -a on
```

Turn email asset reporting off:

```
icapnotify -a off
```

Turn configuration change notification on:

```
icapnotify -n on
```

Turn configuration change notification off:

```
icapnotify -n off
```

## AUTHOR

`icapnotify` was developed by HP.

## SEE ALSO

`icapmodify(1M)`, `icapstatus(1M)`, `icapmanage(1M)`, `icapd(1M)`, `icap(5)`.

# icapstatus(1M)

## NAME

`icapstatus` - Display Instant Capacity (iCAP) status and system information.

## SYNOPSIS

`icapstatus`

`icapstatus -s`

Deprecated:

`icod_stat`

`icod_stat -s`

## DESCRIPTION

The `icapstatus` command displays Instant Capacity status and configuration information, counts, status, and allocation of Instant Capacity components (cores, memory, and cells) for an Instant Capacity system. If the system is a member of a Global Instant Capacity (GiCAP) group, membership information and status on any borrowed or loaned usage rights is displayed. Optionally, system snapshot information containing encrypted audit data is displayed. The deprecated `icod_stat` command performs the identical functions and is maintained for backward compatibility.

For further information see the *HP Instant Capacity User's Guide for Version 9.x* located at `/usr/share/doc/icapUserGuide.pdf`.

## Symbols used in the display

`icapstatus` uses symbols next to the value of some fields in certain circumstances. These symbols are:

- \* Indicates an estimated number of active cores or active memory. For more information on what these estimates might be, see the *HP Instant Capacity user's guide for Version 10.x* located at `/usr/share/doc/icapUserGuide.pdf`.
- # Indicates that all cores in the partition have been requested to be active. The number indicates the current number of cores in the partition.

## icapstatus display fields

If no options are specified, `icapstatus` displays:

Software Version:	This field displays the version of the Instant Capacity client software on the local system.
System ID:	This field displays the user-specified system identifier that the Instant Capacity system uses when reporting the system state to HP, and when HP refers to this system in correspondence to the system contact. The default value for this field is the hostname of the Instant Capacity system. To change this value, use the <code>icapmodify -i</code> command (see <code>icapmodify(1M)</code> ).
Serial number:	This field displays the hardware serial number of the Instant Capacity complex.
Product number:	This field displays the hardware product number of the Instant Capacity complex.
Unique ID:	This field displays a unique identifier for the Instant Capacity complex (generated by HP).

System contact email:	This field displays the email address for the person who should receive configuration change notification and exception reports for the local system. This field is set via the <code>icapmodify -c</code> command.
From email:	This field displays the email address that will be specified as the sender of all Instant Capacity initiated email correspondence for the local system. This field is set via the <code>icapmodify -f</code> command. If not set, email will be sent from the <code>adm</code> user on the local system.
Asset reporting:	This field indicates if the Instant Capacity software on the local system is presently configured to send email asset reports to HP. This is configured using the <code>icapnotify -a</code> command.
Temporary capacity warning period:	This field displays the number of days constituting the temporary capacity warning period for the complex. When the temporary capacity balance is projected to be depleted within this number of days, a warning message is sent by email to the system contact and to root. The value of the warning period can be set using the <code>icapmodify -w</code> command.
Exception Status:	This field indicates if the complex is presently in an exception state. A complex is usually in an exception state when the number of active components of a given type (cores, cells, memory) exceeds the number of available usage rights for that component type. The one exception is that the number of active cores is allowed to exceed the number of core usage rights as long as there is a sufficient positive balance of temporary capacity. A negative temporary capacity balance is always an exception state.

### Information displayed for a GiCAP group

The following status information is displayed when the system is a member of a GiCAP group.

Member <mname> in GiCAP group <gname>	This section heading identifies the member name (<mname>) and the name of the GiCAP group (<gname>) that includes this system as a member.
Active Group Manager:	This field identifies the name of the active Group Manager system where you can get more detailed information about the group, by invoking the <code>icapmanage</code> command (see <code>icapmanage(1M)</code> ).
Standby Group Manager:	This field identifies the name of the standby Group Manager system, if a standby has been defined for the group. The standby Group Manager has the ability to take control of the group if the active Group Manager should become unavailable for any reason. The field is displayed as "N/A" if no standby Group Manager has been defined for the group.
Borrowed/Loaned core usage rights:	This field identifies the count of core usage rights that were either borrowed from or loaned to the GiCAP group. This value must be 0 in order to remove the member from the GiCAP group.
Borrowed/Loaned cell usage rights:	This field identifies the count of cell board usage rights that were either borrowed from or loaned to the GiCAP

group. This value must be 0 in order to remove the member from the GiCAP group.

Borrowed/Loaned memory usage rights: This field identifies the count of memory usage rights that were either borrowed from or loaned to the GiCAP group. This value must be 0 in order to remove the member from the GiCAP group.

The output of `icapstatus` reflects the results of any GiCAP group operations: the borrowing or loaning of component usage rights or the transfer of temporary capacity. The output does not take into account any component usage rights or temporary capacity that is available on other GiCAP group members. Thus, `icapstatus` on group member A may report that there are no additional cores that may be activated with current usage rights even though there is an inactive core on group member A and an available core usage right on group member B. Use the `icapmanage` command on the Group Manager system to get more complete information about available group resources.

When a group member is using temporary capacity and core usage rights are made available on another group member through the use of `icapmodify -d <n>`, there may be a delay between the time the core usage rights are made available and the time the core usage rights move to the group member using temporary capacity. This also means that there may be a delay before the final result is reflected in the output of `icapstatus` for each group member involved in the transfer of the core usage rights.

### Information displayed for the local virtual partition

The following status is displayed when `icapstatus` is run on a virtual partition. Note that some of the displayed information pertains specifically and only to the local virtual partition (such as the “number of active assigned cores” or “number of inactive assigned cores”). However, because usage rights and temporary capacity are always calculated globally across the entire complex, other local values involving these items (such as the “can be assigned” or “could be assigned” values) are the result of calculations using count values across all the partitions.

Total number of assigned cores:	This field displays the total number of cores assigned to the local virtual partition.
Number of active assigned cores:	This field displays the number of assigned cores in the local virtual partition that are currently active.
Number of inactive assigned cores:	This field displays the number of assigned cores in the local virtual partition that are currently inactive.
Additional cores that can be assigned with current usage rights:	This field displays the number of unassigned cores in the hard partition that are not currently assigned to any virtual partition and can be instantly assigned, because enough usage rights are available.
Number of cores that could be assigned with additional usage rights:	This field displays the number of cores that are available for assignment to the virtual partition if additional usage rights are purchased, or additional usage rights were borrowed from a GiCAP group.
Number of cores that can be assigned with temporary capacity:	This field displays the number of additional cores (beyond the number allowed by purchased usage rights or currently borrowed GiCAP usage rights) that can be assigned to the virtual partition and activated using temporary capacity currently available on the complex. When assigning and activating a core using temporary capacity, <code>icapmodify</code> assumes the core is active for at least 30 minutes. Thus, if a complex has a small temporary capacity balance, it may not be possible to activate all the inactive cores in a partition using temporary capacity. Also, temporary

capacity is not used as long as there are available core usage rights on the complex (or in the group, if a member of a GiCAP group), even if the `-t` option is used with `icapmodify` for an activation.

Number of cores currently unavailable for assignment:

This field displays the number of unassigned cores in the hard partition that are not assigned to the local virtual partition and cannot be instantly assigned. This number includes cores in inactive cells and deconfigured cores. When using versions of vPars before A.03.05, bound processors at the time the local virtual partition booted, if unbound later, cannot be instantly assigned to the local virtual partition without an intervening reboot (and assuming usage rights are available).

### Information displayed for the local nPartition

In addition to displaying the current date and time as part of the heading, the following status is displayed when `icapstatus` is run on a hard partition. Much of this information is also displayed when `icapstatus` is run on a virtual partition, except as otherwise specified. Note that some of the displayed information pertains specifically and only to the local hard partition (such as the “number of active cores” or “number of inactive cores”). However, because usage rights and temporary capacity are always calculated globally across the entire complex, other local values involving these items (such as “can be activated” or “could be activated” values) are the result of calculations using count values across all the partitions. The heading for this section includes a datestamp showing the date and time the `icapstatus` command was issued.

Total number of configured cores: This field displays the number of cores physically present in the hard partition.

Number of intended active cores: This field displays the number of cores requested to be active for this hard partition; this is the number of cores that are activated during a boot operation. Typically, this is the number that results from the execution of an `icapmodify` command. Other commands such as `parmodify` and `parcreate` can also affect this value. `vparmodify` does not affect this value.

Number of active cores: This field displays the current number of cores active in the hard partition.

Number of inactive cores: This field displays the current number of inactive cores in the hard partition.

Additional cores that can be activated with current usage rights: This field displays the number of cores that are immediately available for activation using existing core usage rights on the complex. This information is not displayed on a virtual partition.

Number of cores that could be activated with additional usage rights: This field displays the number of cores that are available for activation if additional core usage rights were purchased or additional usage rights were borrowed from a GiCAP group. These cores are also available for temporary activation if temporary capacity is available. This information is not displayed on a virtual partition.

Number of cores that can be activated with temporary capacity: This field displays the number of additional cores (beyond the number allowed by the purchased usage rights or currently borrowed GiCAP usage rights) that can be activated using temporary capacity currently available on the complex. When activating a core using temporary

capacity, `icapmodify` assumes the core is active for at least 30 minutes. Thus, if a complex has a small temporary capacity balance, it may not be possible to activate all the inactive cores in a partition using temporary capacity. Also, temporary capacity is not used as long as there are available core usage rights on the complex, even if the `-t` option is used with `icapmodify` for an activation. This information is not displayed on a virtual partition.

Number of cores that are deconfigured or attached to inactive cells:

This field displays the number of cores that cannot be activated by Instant Capacity software. This includes cores in inactive cells, deconfigured cores, and failed cores deactivated because of LPMCs (Low Priority Machine Checks). This information is not displayed on a virtual partition.

### Instant Capacity Resource Summary

The following status is displayed for the entire complex:

Number of cells without usage rights:

This field displays the number of configured cells in excess of the number of cell usage rights applied to the complex (purchased rights or borrowed from a GiCAP group). Therefore, this number represents the count of cells which are expected to be inactive.

Number of inactive cells:

This field displays the current number of inactive cells in the complex.

Amount of memory without usage rights:

This field displays the total amount of configured memory in excess of the amount of memory usage rights applied to the complex (purchased rights or borrowed from a GiCAP group). This amount of memory is expected to be inactive.

Amount of inactive memory:

This field displays the current amount of inactive memory in the complex.

Number of cores without usage rights:

This field displays the number of configured cores in excess of the number of core usage rights applied to the complex (purchased rights or borrowed from a GiCAP group). This is the number of cores that are expected to be inactive. (If fewer cores are inactive than are expected to be inactive, the complex is either consuming temporary capacity or is out of compliance.)

Number of inactive cores:

This field displays the current number of inactive cores in the complex.

Number of cores using temporary capacity:

This field displays the number of cores presently consuming temporary capacity in the complex.

Number of cores that must be deactivated (insufficient usage rights):

This field displays the number of active cores in excess of the number of available usage rights on the complex. For the system to be in compliance, this number of cores must be deactivated.

Temporary capacity available:

This field displays the amount of temporary capacity available. This balance is displayed in days, hours, and minutes. This value does not reflect any possible use of pooled temporary capacity from a GiCAP group, if the system is a member of a group.

Projected temporary capacity expiration:

This field displays the date and time that temporary capacity is projected to expire at the present consumption rate. This value does not reflect any possible use of pooled temporary capacity from a GiCAP group, if the system is a member of a group.

## Allocation of Instant Capacity Resources Among the nPartitions

The following table displays how Instant Capacity components are distributed among the partitions in the complex:

nPar ID:	This field displays the partition number for the row of data.
Total cores:	This field displays the total number of cores physically present for the hard partition.
Intended active cores:	This field displays the number of cores requested to be active for this hard partition; this is the number of cores that will be activated during a boot operation. Typically, this is the number that results from the execution of an <code>icapmodify</code> command. Other commands such as <code>parmodify</code> and <code>parcreate</code> can also affect this value. <code>vparmodify</code> does not affect this value.
Actual active cores:	This field displays the current number of active cores for the hard partition. In a virtual partition, the count represents the total number of cores assigned to all the virtual partitions.
Inactive cores:	This field displays the current number of inactive cores in the hard partition. In a virtual partition, the count represents the total number of cores not assigned to any virtual partition.
Inactive memory:	This field displays the current amount of inactive memory in the hard partition.
Inactive cells:	This field displays the current number of inactive cells in the hard partition.
Runs iCAP:	This field indicates whether the hard partition contains compatible Instant Capacity software.
nPar name:	This field displays the partition name corresponding to the row of information that is displayed. If the row is the local partition, the word "local" follows the partition name. For cells that are not assigned to partitions, the "Unassigned Cells" label is displayed.

## Options

`icapstatus` recognizes these options:

- s Displays system snapshot information. This option displays the product and serial number for this system, as well as a snapshot string that can be entered into the HP Utility Pricing Solutions portal for periodic auditing purposes.

## EXTERNAL INFLUENCES

### Environment Variables

- `LANG` determines the locale to use for the locale categories when both `LC_ALL` and the corresponding environment variable (beginning with `LC_`) do not specify a locale. If `LANG` is not set or is set to the empty string, a default of `C` is used (see `lang(5)`).
- `LC_TIME` determines the date and time strings output.
- `LC_MESSAGES` determines the language in which messages (other than the date and time strings) are displayed.

If any internationalization variable contains an invalid setting, `icapstatus` behaves as if all internationalization variables are set to C (see *environ(5)*).

## **RETURN VALUE**

The `icapstatus` command exits with one of these values:

- 0           Command succeeded.
- 2           Command succeeded; system is not an Instant Capacity system.
- >0, !=2    Command failed; error message sent to `STDERR`.

## **AUTHOR**

`icapstatus` was developed by HP.

## **SEE ALSO**

`icapmodify(1M)`, `icapnotify(1M)`, `icapmanage(1M)`, `icapd(1M)`, `icap(5)`.

# icapd(1M)

## NAME

icapd - Instant Capacity (iCAP) daemon.

## SYNOPSIS

icapd

## DESCRIPTION

The `icapd` (formerly `icodd`) daemon is installed and started as part of the Instant Capacity software on all potential iCAP systems, and respawns itself if killed. If this daemon is not running, other Instant Capacity commands fail. The operations this daemon performs are vital in keeping the complexwide view of the Instant Capacity state current. The following entry is added to `/etc/inittab` in order to have `icapd` start and respawn itself:

```
icap:23456:respawn:/usr/sbin/icapd # Instant Capacity daemon
```

This daemon is not started on hardware that is not supported under the Instant Capacity program. If `icapd` is installed and running on a system with Instant Capacity components present (cores, cells, or memory), it sends daily asset report email to HP (if so configured), tracks temporary capacity, sends exception notifications, and maintains a healthy Instant Capacity state.

For more information about the functions that `icapd` performs for Instant Capacity systems, see the *HP Instant Capacity user's guide* located at `/usr/share/doc/icapUserGuide.pdf`.

The `icapd` daemon reports errors via `syslog` (see `syslog(3C)`). Exception notification email is sent to `root` and to the system contact email address (configured via the `icapmodify` command (see `icapmodify(1M)`).

The `icapd` daemon performs periodic operations based on the time of day. The `icapd` daemon is spawned by `init` and gets its time zone specification from the `/etc/default/tz` file. By default, the time zone specified in `/etc/default/tz` is `EST5EDT`. You can specify which time zone the `icapd` daemon uses to interpret its current time by modifying the `/etc/default/tz` file. For details about the time zone format, see `environ(5)`. A restart of the `icapd` daemon is required before the new time zone value takes effect (that is, kill the `/usr/sbin/icapd` process).

## AUTHOR

`icapd` was developed by HP.

## SEE ALSO

`icapmodify(1M)`, `icapstatus(1M)`, `icapnotify(1M)`, `icapmanage(1M)`, `icap(5)`.



---

# A Special Considerations

This appendix covers the following topics:

- “Assumed Values in icapstatus Command” (page 192)
- “Upgrading to Instant Capacity version B.06.x or later (HP-UX)” (page 193)
- “Dual-Core Support in Instant Capacity Systems” (page 195)
- “New Partition Creation and Instant Capacity” (page 196)
- “Implications of Removing a Cell from an Instant Capacity System ” (page 197)
- “Shutting Down a Partition with Instant Capacity Cores” (page 198)
- “Instant Capacity and Reinitializing the nPartition (Genesis Partitions)” (page 199)
- “par Commands from PC System Management Station” (page 200)
- “Instant Capacity Compatibility with Processor Sets (HP-UX)” (page 201)
- “Configuring Email on Instant Capacity Systems” (page 202)
- “Measurement Software and Instant Capacity Systems” (page 206)
- “Dynamic Processor Resilience (HP-UX)” (page 207)
- “Security Issues” (page 208)

## Assumed Values in `icapstatus` Command

The `icapstatus` command might make assumptions on the number of active cores and amount of active memory, depending on certain system conditions. If values are assumed, the `icapstatus` command output contains an asterisk next to the appropriate field.

### Assumed Processor Values

Occasionally, the output of the `icapstatus` command contains an asterisk (\*) next to the value in the `Actual Active Cores` field (under the section `Allocation of Instant Capacity Resources among the nPartitions`). The asterisk appears in the output whenever a nonlocal partition appears to be active, but the `icapd` daemon is not reporting system information. Some examples that can cause this situation include:

- The absence of a compatible version of the Instant Capacity software on a nonlocal partition
- No space left in `/var` on the nonlocal partition prevents `icapd` from communicating system information
- A nonlocal partition shut down for an extended period of time with shutdown or reboot, without the `-R -H` options (which brings the cells to an inactive state)
- A nonlocal partition running an operating system that is not HP-UX or OpenVMS (for example, Windows or Linux)

In these cases, the Instant Capacity software on other partitions identifies all cores in the nonlocal partition as active. In addition to the asterisks in the output of `icapstatus`, this can affect the following:

- Temporary capacity consumption
- The ability to change the complex with `parmodify`, `parmgr`, or `parcreate` commands
- Core activation with the `icapmodify` command
- System noncompliance



---

**NOTE:** The number of active cores is always known for a local partition.

---

If a nonlocal partition appears to be inactive, the number of active cores reported by the `icapstatus` command is zero. For example, if the hardware for a nonlocal partition is inactive, `icapstatus` considers the partition as inactive and reports the number of active cores as zero.

### Assumed Memory Values

When a cell board is powered off, regardless of whether the cell is in the local partition or in a nonlocal partition, the amount of inactive memory is assumed to be 2 GB and is reported as such by the `icapstatus` command. If the amount of memory without usage rights configured for the cell exceeds 2 GB and the cell is powered down and remains assigned to the partition, the system goes out of compliance.

An asterisk (\*) might appear next to the value of the `Inactive Memory` field in the `icapstatus` output, in the section `Allocation of Instant Capacity Resources among the nPartitions`.

## Upgrading to Instant Capacity version B.06.x or later (HP-UX)

The first time a version of the codeword-based B9073BA Instant Capacity software (B.06.00 or later) is loaded onto a system where the old B9073AA software (version B.03.x through B.05.x) has been in use, the new software requires the system to go through an upgrade process.

This process involves transferring Instant Capacity inventory information from HP to the system through the application of an upgrade codeword. This allows the system (using the B9073BA software) to keep track of the number of components without usage rights.

Prior to December 2003, the Instant Capacity software product (B9073AA) was loaded from Support Plus media and was not present on the 11i v1 and 11i v2 Operating Environment (OE) media. In December 2003, an updated version of the Instant Capacity software product (B9073BA) was placed on the 11i v1 and 11i v2 OE media.

The Instant Capacity version B.09.x software (B9073BA) is automatically installed when the HP-UX 11i v3 OE is installed. When a newer version of the HP-UX OE is installed in a partition, the Instant Capacity software in the OE is automatically updated in that partition.

After the Instant Capacity software is updated in one partition from an older B9073AA version (B.03.x through B.05.x) to the newer B9073BA version (B.06.00 or later), the Instant Capacity software in all of the system's other partitions needs to be updated before the new software is completely operational. Until that time, the new software does not have visibility of the status of partitions running the old software. As a result, it is likely that exception email will be generated daily from the new software until all updates are completed.

After the installation of new Instant Capacity B9073BA software on the first partition, perform the following steps.



---

**NOTE:** The following HP-UX procedures use the obsolete commands, such as `icod_stat` and `icod_modify`, because these commands are common to all versions of Instant Capacity (B.06.x, or later). On version B.09.x, these commands could be `icapstatus` or `icapmodify`, which have identical results.

---

1. Execute the following command and record the output:

```
/usr/sbin/icod_stat -s
```

You can copy and paste the output of the `icod_stat -s` command and save it in a text file. You need this information later in step 4.

2. Go to the HP Utility Pricing Solutions web portal at <http://www.hp.com/go/icap/portal>. Log in to the portal and click the link to get codewords.
3. Click the **Upgrade Codeword** link to get an upgrade codeword.
4. Supply the requested information from the `icod_stat -s` command output gathered in step 1, and get the upgrade codeword from the portal. You need the codeword in the next step.
5. Apply the upgrade codeword, generated by the Utility Pricing Solutions portal, to the system by executing this command:

```
/usr/sbin/icod_modify -C codeword
```

where *codeword* is the upgrade codeword that was supplied by the portal in step 4. This is easily accomplished by copying and pasting the codeword (that is generated by the portal) to the system where you are executing the `icod_modify -C codeword` command. The upgrade codeword needs to be applied only once on the entire system.

6. Install or upgrade all other partitions with the Instant Capacity B9073BA software (version B.06.00 or later). This software is on the 11i v1, 11i v2 and 11i v3 OE media, starting with the December 2003 media, and is automatically installed when the 11i v1, 11i v2, or 11i v3 OE is installed. The software is also available from the HP web site at <http://www.hp.com/go/softwaredepot> (search for B9073BA). If you do not complete this step, in the absence of other information, the Instant Capacity software assumes that all partitions that were not upgraded are fully active. If all partitions in a system are not upgraded, the Instant Capacity software might determine the system to be in an exception state.

7. Execute the following command:

```
/usr/sbin/icod_stat
```

Inspect the `icod_stat` command output for the line that indicates the `Exception status` (near the top of the output). If it displays “No exception”, your system is in compliance. Inspect the remainder of the output to see the distribution of your active and inactive cores in the system and, if you want to make changes, modify it using the `/usr/sbin/icod_modify` command.

8. If the `Exception status` in step 7 indicates that there are more cores, cells, or memory active than expected, you need to deactivate the appropriate number of each component in order to bring the system into compliance with the Instant Capacity contract.
9. From the output in Step 7, verify that the correct system contact information is specified. If necessary, update the contact information by executing:

```
/usr/sbin/icod_modify -c contact_email_address
```

10. Ensure you have installed all of the required HP-UX patches on each partition. See “Required Patches for HP-UX 11i v1” (page 29) or “Required Patches for HP-UX 11i v2” (page 29) for details. If necessary, obtain missing patches from the web site: <http://us-support2.external.hp.com>. Detailed instructions on installing each patch are available on this web site.
11. *If email connectivity is required (when upgrading to B.06.x), or if asset reporting is being used with B.07.x or later versions:* Execute the following command:

```
/usr/sbin/icod_notify reply_address
```

This command verifies that your `sendmail` configuration is capable of sending email messages to the `hp.com` domain. The `icod_notify` command needs to be executed with an email address as an argument. The email address is the address HP uses when responding to the email sent from the Instant Capacity system and is the email address to which the HP acknowledgement message is sent.

Using any email reader, verify receipt of the acknowledgement email from HP, which typically is sent within 1 hour.

## Dual-Core Support in Instant Capacity Systems

Each HP cellular complex has 4 sockets. With dual core processing, each socket accepts a CPU module that contains 2 processor cores. You can upgrade an Instant Capacity Superdome system to a dual-core system by replacing the cell boards and processors. Contact your HP service representative for details on upgrading to dual core processors.

The Instant Capacity software supports dual core processors. The software treats each core individually and allows the following actions at the core level:

- Applying Right to Use (RTU) codewords
- Activating and deactivating
- Load balancing across partitions
- Configuring in virtual partitions

## New Partition Creation and Instant Capacity

You can assign a cell to an existing partition even if the cell contains cores without usage rights (Instant Capacity processors), as long as there are enough available core, cell, and memory usage rights to cover activation of the cell, its memory, and at least one of the cores on the cell. You can verify that the partition has valid Instant Capacity software installed, and that the partition is running an operating system capable of activating and deactivating cores.

However, when a new partition is created, these verifications are not possible; therefore, the Instant Capacity software does not allow the partition to be created unless it contains only components with usage rights. That is, creation of the new partition fails if any of the cell boards contain Instant Capacity components without usage rights.

# Implications of Removing a Cell from an Instant Capacity System

The Instant Capacity software tracks the expected number of inactive components (cores, cells, and memory) in a complex and knows the actual number of active and inactive components. The complex is in compliance if the actual number of inactive components meets or exceeds the expected number of inactive components.

The complex is out of compliance if the actual number of inactive components is less than the expected number of inactive components and no temporary capacity is available.

However, a complex can also get out of compliance if a cell is removed from the complex. For example, if a cell contains inactive cores that are contributing to compliance, and the cell is removed, there will be fewer inactive cores in the complex. This can result in the complex being out of compliance, and temporary capacity might begin to be debited.

For example, a complex contains 2 cells, with 2 partitions having 2 inactive and 2 active cores each. The Instant Capacity software expects the complex to have 4 inactive cores. If one of the cells (0) experiences a hardware problem, and you remove the cell, the complex is left with only 1 cell that contains 2 active and 2 inactive cores. The complex is now out of compliance because 4 inactive cores are expected to be in the complex, yet there are only 2 inactive cores.

**Table A-1 Removing a Cell and Decreasing Inactive Cores**

State	Partition (Cell) 0	Partition (Cell) 1	Notes
Before cell 0 is removed	2 active, 2 inactive	2 active, 2 inactive	4 inactive cores expected (in compliance)
After cell 0 is removed	0 active, 0 inactive	2 active, 2 inactive	4 inactive cores expected (out of compliance)

In this example, all cores in the removed cell are identified as active. This causes the complex to be out of compliance because the complex has 2 more active cores than it has core usage rights. This results in the complex consuming 2 hours of temporary capacity for each hour that the complex remains in this state. Deactivating another core from cell 1 decreases the amount of temporary capacity being consumed, but because at least 1 core must be active per active cell, this complex cannot remain in compliance unless of temporary capacity is used.

Note that removal of a cell, followed by a reboot of the affected partition, does not affect either the intended active number for the partition or the required number of inactive cores, which is determined by the overall availability of core usage rights across the complex. While the cell is absent, temporary capacity can be consumed if the number of inactive cores is less than the expected number of inactive cores. Having additional temporary capacity allows this system to remain in compliance even in the presence of a cell hardware failure.

## Shutting Down a Partition with Instant Capacity Cores

The Instant Capacity software saves information about the number of active cores for each partition, and this information expires over time. If the partition is not active (but the hardware is powered up), Instant Capacity software on other partitions assumes that all cores in the inactive partition are active unless it can detect otherwise. For details about these assumed processor values, see “Assumed Values in `icapstatus` Command” (page 192).

These are the general rules the Instant Capacity software uses:

- If the partition is to be shut down for less than 12 hours, no action is necessary.
- If the partition is to be shut down for more than 12 hours, consider powering the cells off, or shutting the partition down using the `shutdown -R -H` command.
- If the partition crashes or shuts down abnormally, reboot the partition within 12 hours or power it off.



**IMPORTANT:** If you shut down a partition for 12 hours or more, also power it off to avoid additional charges. To power off the partition, execute the `PE` command from the system MP.

On HP-UX systems, always use the `shutdown` command when shutting down or rebooting an Instant Capacity partition. For information about the `shutdown` command, see `shutdown(1M)`.

On OpenVMS systems, always use the `sys$system:shutdown.com` procedure when shutting down or rebooting an Instant Capacity partition.

---

## Instant Capacity and Reinitializing the nPartition (Genesis Partitions)

Any use of the `CC` command at the service processor level has the potential to overwrite the Instant Capacity configuration, and is therefore not recommended on Instant Capacity systems. In particular, creating a Genesis Partition on an Instant Capacity system is not recommended because it causes the system to be out of compliance.

If you clear the configuration of a system that has Instant Capacity components (components without usage rights), you must acquire and apply a codeword that restores a lost iCAP configuration to remain in compliance with your contract. Because the codeword can be generated only by using a recent audit snapshot of the system, generate an audit snapshot with `icapstatus -s` before doing the reset or, if asset reporting is configured, make sure that a recent asset report was sent to the portal.

## par Commands from PC System Management Station

Use of `par` commands (such as `parmodify` or `parcreate`) can cause changes to a complex that affect the Instant Capacity state of the complex. Therefore, if a `par` command is executed on an Instant Capacity complex from a PC System Management Station (SMS), the command must be directed towards a HP-UX partition in order to succeed so that the Instant Capacity software can authorize the change. The `par` commands support this functionality through the `-h` option.

For more information, see *parmodify(1M)*, *parcreate(1M)*, and *parremove(1M)*.

# Instant Capacity Compatibility with Processor Sets (HP-UX)

## Overview

The Instant Capacity software successfully coexists with processor sets (psets).

To coexist with psets, the Instant Capacity software activates and deactivates cores in only the default processor set. Cores in nondefault processor sets are not activated or deactivated.



**NOTE:** There must be at least one core in the default processor set. The last remaining core in the default processor set is unavailable for deactivation.

---

## Scope of the Instant Capacity Software Interacting with psets

The Instant Capacity software does not provide any additional functionality to specifically support adding or removing cores from a specific pset.

### psets on nPars

In an nPar environment where psets are present, the Instant Capacity software activates and deactivates cores in only the default pset.

Cores can be manually migrated to the default pset for purposes of deactivation, or from the default pset to other psets after activation.

### psets on vPars

In a vPar environment, the Instant Capacity software passes the request for a core activation or deactivation to the `vparmodify` command.

With vPars version A.04.01 or later, vPars is pset-aware to some extent. When adding cores to a running vPar, a core is always added to the default pset; thereafter, the core can be moved to another pset. When removing cores, `vparmodify` chooses cores from the default pset first. If not enough exist in the default pset to satisfy the request, `vparmodify` chooses the remaining cores arbitrarily without regard to pset membership.

With vPars versions earlier than A.04, no special consideration is given to psets from the `vparmodify` command's perspective. Therefore, when using vPars, cores in nondefault psets must be bound cores. Otherwise, a core designated for deactivation by `vparmodify` might be selected from an unexpected pset.

# Configuring Email on Instant Capacity Systems

## Email Requirements

Previous versions of the Instant Capacity software required email connectivity to HP in order to send asset reports as encrypted email messages. Starting with version B.07.x, Instant Capacity software does not require email connectivity or asset reporting. However, you can choose to configure it because it can be useful for viewing complex-wide asset information at the HP Utility Pricing Solutions portal.



---

**NOTE:** Email asset reporting is set to “off” by default when the Instant Capacity software is installed. You turn asset reporting on or off with the `icapnotify -a` command. You can view the current setting of email asset reporting in the `Asset reporting` field, near the beginning of the `icapstatus` command output.

---

The following are requirements for email connectivity:

- The Instant Capacity system or partition must have the `sendmail` utility installed and configured so that it can send email to the `hp.com` domain.
  - The domain name in the Instant Capacity `FROM` email address, for the email sent from the Instant Capacity system to HP, must be DNS resolvable by HP. For details, see “Configuring Instant Capacity’s From Email Address” (page 204).
- 



**IMPORTANT:** On OpenVMS systems, SMTP mail must be configured for email connectivity. For information about configuring SMTP mail, see the documentation for your TCP/IP provider.

---

Although the `sendmail` configuration and routing can vary, the system must have the ability to send email to the `hp.com` domain.

The ability to receive email from HP is optional but is useful for testing the ability to send email to HP. For more information, see “Configuring Your Server to Send but Not Receive Email” (page 205). For more information about `sendmail`, see `sendmail(1M)`.

The `sendmail` utility is part of the HP-UX core and is installed with the HP-UX operating system. However, you must follow the `sendmail` configuration process to complete its installation. For information, see the appropriate chapter in the following documents which are located on the HP Documentation website (<http://docs.hp.com>):

- For HP-UX 11i v1: *Installing and Administering Internet Services* (B2355-90685)
- For HP-UX 11i v2 and 11i v3: *HP-UX Internet Services Administrator’s Guide* (B2355-91060)

To access either of the documents, select **I/O Cards and Networking Software -> Internet Services**.

### On Partitionable Systems

To enable asset reporting, configure email connectivity on each partition. This makes it easier to redistribute cores across partitions later (that is, load balance). For details, see “Load-Balancing Active Cores” (page 62).

---



**IMPORTANT:** The email is rejected by the mail servers at HP if, for email sent from the Instant Capacity system to HP, the domain name in the `From` field is not DNS resolvable by HP. Also, because asset reports are encrypted and must be decrypted at the HP portal, the decryption process might not work correctly if outgoing email sent from your system is automatically modified in any way, for example, to include a privacy notice.

---

## Email Configuration

### Before you start

If you decide to enable email connectivity, your Instant Capacity system must be network accessible to HP mail servers that are outside your company's firewalls. If your Instant Capacity

system is on an isolated network, email from the system does not reach HP. This causes your system to be out of compliance with your Instant Capacity contract if you are using temporary capacity (TiCAP).

### The sendmail application

The `sendmail` application is used by the Instant Capacity software to send encrypted email messages from your system to HP. The `sendmail` daemon, if running, can also be used to receive email. For purposes of this email configuration, only the ability to send email is required.

Mail applications invoke `sendmail` to send email. The configuration file, `/etc/mail/sendmail.cf`, offers tremendous flexibility.

### Overview of email routing across the internet

When the Instant Capacity software uses `sendmail` to send email to HP, `sendmail` determines where it should initially send the email (the first hop). Mail often goes through multiple systems (hops) before it reaches the final destination. To determine the first hop for the email, `sendmail` uses one of the following:

- The email is routed to a mail relay host if it is configured in the `/etc/mail/sendmail.cf` configuration file. This is the easiest implementation and can be done with just a one line change (DS) to the default `/etc/mail/sendmail.cf` file.

Note that the relay host must be configured to properly route (forward) the email to the final destination.

- DNS MX records — this method requires that the Instant Capacity system be in an environment (network) where DNS (Domain Name Server) is operating and properly configured. `sendmail` on the system queries a DNS server for the name of the email server to forward the email to (for the first hop) in order for the email to reach the final destination (`hp.com`).

In all cases, the following requirements must be met:

- HP mail servers that receive email expect the host (the mail server in the last hop before reaching HP) to be properly registered in DNS. If not, the HP mail server rejects, or “bounces”, the email.
- The From field (email address) in the email message must be known by the receiving mail server (that is, the hostname is registered in DNS and advertised on the internet). Otherwise, the receiving mail server at HP rejects the email. This field in the email can be configured with a simple one-line modification (DM) to the `/etc/mail/sendmail.cf` file.

In some DNS environments, changes to the default `/etc/mail/sendmail.cf` file might not be needed to properly route email from the Instant Capacity system to HP.

- In some environments, configuring your system to properly send email from the system to HP can require as little as a two-line edit (or none) to the `/etc/mail/sendmail.cf` file. In most organizations, configuring mail, including `sendmail` and DNS configurations, is usually handled by the IT team.

### Example A-1 Example Edit to Sendmail Configuration (`/etc/mail/sendmail.cf`)

---

```
DMmy_company.com
DSmailhub.my_company.com
```

---

In this example:

- The Instant Capacity system’s hostname is `myICAPsystem.my_site.my_company.com`.
- The From field of the email is set to `my_company.com` rather than to the exact hostname of the Instant Capacity system. This is because most organizations do not advertise the names of their internal servers to the internet; however, they do advertise a few (select) high-level domain names.

- The Instant Capacity system is not advertised to the internet, but hostname `mycompany.com` is advertised and reachable from the internet.
- Email is forwarded from the system to a mail relay host called `mailhub`. The mail server called `mailhub` can either be directly connected to the internet and send the email directly to HP, or it can forward the email to another mail server on its way to HP.




---

**NOTE:** Any bounced Instant Capacity email messages are sent to the `adm` mailbox.

---

## Steps to Confirm or Diagnose Email Configuration

After you configure your Instant Capacity system to send email over the internet, follow these steps to confirm the email configuration or to aid in debugging the configuration:

1. Send an email message from your system to an email address in the same domain (intranet) and confirm receipt of the email message.
2. Send an email message from your system to an email address outside of your domain (to the internet, for example, to a `yahoo` or `hotmail` email address) and confirm receipt of the email message.
3. Send an email message from your system to someone at HP (for example, an HP representative in a local account team) and confirm that the person at HP received the email message.
4. As `root`, execute the following command:  

```
/usr/sbin/icapnotify <reply_address>
```
5. If all the previous steps are successful, but asset reports are still not visible on the HP portal, examine your email configuration to determine whether outgoing messages are automatically being modified or appended, for example, to include something like a privacy notice. Additions or modifications to encrypted asset reports might cause them to be rejected by the portal.

The command in step 4 sends an email message to HP's audit application. HP sends a confirmation email message to the `reply_address`. Receipt of the confirmation email message confirms successful email configuration.

## Configuring Instant Capacity's From Email Address

One of the email requirements of the Instant Capacity program is that the From email address, on email messages sent by the Instant Capacity software from your system, must be DNS resolvable.

The Instant Capacity software uses `adm@localhost.domain` as the default From email address (where `localhost` is the hostname of your system and `domain` is its DNS domain). If the default From email address is undesirable, you can configure the Instant Capacity software to use a From address you specify.

### Configuring a Specified From Address

To configure your specified Instant Capacity From email address, execute the following command:

```
/usr/sbin/icapmodify -f from_address
```

You can verify the configured Instant Capacity From email address by using the `/usr/sbin/icapstatus` command.

After you configure a specified From email address, the Instant Capacity software uses it on all subsequent email messages sent from your system.

### Reverting to the Default From Address

If you specified an Instant Capacity From email address and you want to revert to the default From email address (`adm@localhost.domain`), execute the following command:

```
/usr/sbin/icapmodify -f ""
```

## Configuring Your Server to Send but Not Receive Email

For security reasons, some organizations do not allow incoming mail. If you want your Instant Capacity system to only send email but not receive email, complete the following configuration procedure:

1. To prevent the `sendmail` daemon from starting up again when your system reboots, edit the `/etc/rc.config.d/mailservs` file, and change the value of `SENDMAIL_SERVER` to 0:

```
vi /etc/rc.config.d/mailservs
```

```
#####  
# Mail configuration. See sendmail(1m) #  
#####  
#  
# BSD's popular message handling system  
#  
# SENDMAIL_SERVER: Set to 1 if this is a mail server  
# and should run the sendmail daemon.  
# SENDMAIL_SERVER_NAME: If this is not a mail server, but a  
# client being served by another  
# system, then set this variable to  
# the name of the mail server system  
# name so that site hiding can be  
# performed.  
#  
export SENDMAIL_SERVER=0  
export SENDMAIL_SERVER_NAME=
```

2. To immediately stop the server from receiving email, kill the active `sendmail` daemon by executing the following command:

```
/sbin/init.d/sendmail stop
```

## Testing Email Transmission of the Asset Report



---

**NOTE:** The following procedure assumes your Instant Capacity system is capable of sending internet email.

---

Execute the following command to send your asset report by email to HP:

```
/usr/sbin/icapnotify <reply_address>
```

The specified `reply_address` receives an acknowledgment email from HP confirming the receipt of your asset report. Use an email client to verify the acknowledgment email from HP to the `reply_address`.

## Measurement Software and Instant Capacity Systems

Systems with Instant Capacity components (and systems contributing usage rights to Instant Capacity systems) might have fewer active cores than the total number of cores in the system. This fundamental difference between the number of active cores and the total number of cores can cause some processor measurement products and utilities to report incorrect information. Additionally, when a core is dynamically activated, some software products must recognize the change in the number of active cores in order to report correct processing information.

### HP OpenView Measurement Products

HP OpenView measurement products, such as MeasureWare and GlancePlus, must be version C.02.60 or later to provide correct measurements. Earlier versions of the OpenView measurement products might not work correctly on Instant Capacity systems.

#### On HP Integrity Superdome Systems

On HP Integrity Superdome systems, HP recommends updating to the GlancePlus Pak version C.03.20 or later.

### Other Measurement Software

Please consult your measurement software vendor about whether their software works properly on Instant Capacity systems, and update the measurement software versions as needed.

## Dynamic Processor Resilience (HP-UX)

The LPMC monitor, within the Support Tools Manager (STM) diagnostics, generates Information events for all cache errors that are detected. After three errors (Threshold) are detected on a processor in 1440 minutes, or a 24-hour period of time (Period), the monitor deactivates that particular processor, marks it for deconfiguration on the next system reboot, and generates a Serious event. After the failed processor is deactivated, the LPMC monitor attempts to activate one of the inactive Instant Capacity cores, if any are available. This method ensures the processing power of the system is unchanged.

A default value of 3 is assigned to Threshold, except for the PCX-W+ family of processors, which have a value of 5 assigned. The default value assigned to a Period is 1440 minutes, or 24 hours, in all possible processor configurations.

An inactive processor under warranty or support automatically replaces a failed processor. HP also services and replaces any failed processor.

### Monarch Processors

For details about replacing failed monarch processors, see “Failed Monarch Processors (HP-UX)” (page 73).

## Security Issues

### Customer protections which iCAP assumes to be in place

Instant Capacity commands provide system status information and facilitate system configuration modification, and are therefore executable only by users with root level access. An assumption is made that there exist administrative policies which exercise the appropriate degree of control over root level access.

### Disabling the iCAP daemon (HP-UX)

On a system with full usage rights (no iCAP components), you can disable the iCAP daemon (`icapd`) by commenting out its entry in the `/etc/inittab` system file, resetting the init task (`init -q`), and killing `icapd` via `kill -9` or `kill -s SIGTERM`.

Note that disabling the daemon in this way on an iCAP or GiCAP system is a violation of the iCAP contract with HP. After 12 to 24 hours, the system goes out of compliance and an exception notification email is sent. Also, other partition management software cannot determine whether the system contains iCAP components and, as a result, refuses to manage any components that are present.

### Customer Security Requirements

The Instant Capacity software is designed to provide maximum protection for sensitive customer information. It follows these customer security requirements:

- Sensitive customer data (names, phone numbers, email addresses, hostnames, IP addresses) is not transmitted to HP.
- There are no transmissions of authentication credentials in clear (nonencrypted) text.
- Nonsuperuser access to iCAP commands and data is not allowed.
- Confidential information is encrypted when transmission is required.
- Appropriate protections are accorded to confidential data and authentication credentials.

### Security Tuning Options

Instant Capacity asset reporting (via email to HP) is optional and is turned off by default. Customers can enable asset reporting by executing the `icapnotify -a on` command.

---

## B Considerations for OpenVMS Systems

This appendix covers the following topics:

- “CLI Support on OpenVMS” (page 210)
- “DCL Commands” (page 212)
- “Special OpenVMS-Specific Features and Considerations” (page 220)
- “Restrictions” (page 221)

## CLI Support on OpenVMS

OpenVMS provides a CLI (command-line interface) to the Instant Capacity software. The HP-UX command syntax can be implemented using foreign command symbols. The DCL ICAP command provides DCL command support.

### HP-UX Style Commands

The HP-UX command syntax can be used on OpenVMS systems by defining foreign command symbols to the iCAP images. Add the following three symbol declarations to your LOGIN.COM file or to the SYLOGIN file to define commands that use the HP-UX syntax:

```
$ icapmodify === $ICAP_MODIFY
$ icapnotify === $ICAP_NOTIFY
$ icapstatus === $ICAP_STAT
```

Command options are specified as described in the HP-UX manpages for each command.

### OpenVMS Command Mapping

Table B-1 shows the HP-UX iCAP commands and their OpenVMS equivalents.

**Table B-1 HP-UX and OpenVMS Command Equivalents**

HP-UX Style	OpenVMS Style
icapstatus	icap show status
icapstatus -s	icap show status/snapshot
icapmodify -C <codeword>	icap apply "codeword"
icapmodify -c <address>	icap set email/contact="address"
icapmodify -f <address>	icap set email/from="address"
icapmodify -i <id>	icap set system_id "id"
icapmodify -r	icap reconcile
icapmodify -w <days>	icap set warning_days "days"
icapmodify -a	icap activate /cpu=/defer/ticap
icapmodify -d	icap deactivate /cpu/defer
icapmodify -s	icap set active_cpus n
icapnotify -a	icap set asset/state=on/off
icapnotify -n	icap set notification/state=on/off
icapmanage -i -u <rule_file>	icap manage install <rule_file>
icapmanage -C <codeword>	icap manage codeword <codeword>
icapmanage -a -g <group_name>	icap manage add group <group_name>
icapmanage -r -g <group_name>	icap manage remove group <group_name>
icapmanage -T <host>[,<host>]...[-g <group_name>]	icap manage test <host,host,...> [/group=<group_name>]
icapmanage -a -m <member_name>:<host>[,<host>]... -g <group_name>	icap manage add member <member_name> /host_list=(host,host,...) [/group=<group_name>]
icapmanage -r -m <member_name>	icap manage remove member <member_name>
icapmanage -s -g <group_name> [-b] [-v]	icap manage status <group_name> [/BRIEF] [/FULL]

**Table B-1 HP-UX and OpenVMS Command Equivalents (continued)**

HP-UX Style	OpenVMS Style
icapmanage -R [<host>[,<host>]...] [-U <rule_file>]	icap manage report <host,host,...> <rule_file>
icapmanage -x <host>	icap manage extract <host>

## OpenVMS iCAP Files

Table B-2 lists the files that are new to iCAP version 9.x on OpenVMS.

**Table B-2 OpenVMS iCAP Files**

OpenVMS File Location	HP-UX File Location	Supplied By
sys\$manager:GiCAP.log	/var/adm/GiCAP.log	Created at run time.
sys\$system:GiCAP.rules	/etc/opt/iCAP/GiCAP.rules	Supplied in kit.
sys\$system:GiCAP.rules_list		Supplied in kit.
sys\$system:GiCAP_db.dat	/etc/opt/iCAP/GiCAP_db	Created at run time.
sys\$manager:GiCAP_keygen.com	/etc/opt/iCAP/GiCAP_keygen	Supplied in kit.
sys\$sysroot: [SYSMGR.SYSTEM.SSH2] GICAPKEY.; /SYS\$SYSROOT/sysmgr/system/ssh2/ GICAPKEY	/etc/opt/iCAP/.GiCAPKey	Created by GiCAP_keygen.com.
sys%sysroot: [SYSMGR.SYSTEM.SSH2] GICAPKEY.PUB /SYS\$SYSROOT/sysmgr/system/ssh2/ GICAPKEY.PUB	/etc/opt/iCAP/.GiCAPKey.pub	Created by GiCAP_keygen.com.

## DCL ICAP Command

The `ICAP` command supports six command options to perform iCAP operations on OpenVMS systems.

## DCL Commands

### ICAP ACTIVATE

#### Name

`ICAP ACTIVATE` - Immediately activates additional cores on the system. (HP-UX equivalent: `icapmodify -a`)

#### Format

`ICAP ACTIVATE /CPU=n [/DEFER] [/TICAP]`

#### Qualifiers

- `/CPU=n` Specifies the number of additional cores to activate. This qualifier is required.
- `[/DEFER]` Defers the activation until the next reboot. (HP-UX equivalent: `-D` option)
- `[/TICAP]` Authorize the use of temporary capacity to satisfy this activation request. (HP-UX equivalent: `-t` option)

## ICAP APPLY

### Name

ICAP APPLY - Apply an iCAP codeword. (HP-UX equivalent: `icapmodify -C`)

### Format

ICAP APPLY "*codeword*"

### Parameter

*"codeword"* An iCAP codeword obtained from the HP Utility Pricing Solutions portal. Enclose the codeword in double quotation marks.

## ICAP DEACTIVATE

### Name

ICAP DEACTIVATE - Deactivates cores on the system. (HP-UX equivalent: `icapmodify -d`)

### Format

ICAP DEACTIVATE */CPU=n [qualifiers]*

### Qualifiers

*/CPU=n* Specifies the number of cores to deactivate. This qualifier is required.

*/DEFER* Defers the deactivation until the next shutdown. (HP-UX equivalent: `-D` option)

## ICAP RECONCILE

### Name

ICAP RECONCILE - Activates or deactivates cores (subject to compliance limits) to bring the system to a state where the intended active number of cores are active. (HP-UX equivalent: `icapmodify -r`)

### Format

ICAP RECONCILE

## ICAP SET

### Name

ICAP SET - Sets various iCAP management variables.

### Format

ICAP SET *parameter [qualifiers]*

### Parameters

**ACTIVE\_CPU** Sets the number of active cores and the number of intended active cores. (HP-UX equivalent: `icapmodify -s`)

#### Format

ICAP SET ACTIVE\_CPU *count*

#### Value

*count*: the number of cores to set active in the npartition.

**ASSET** Sets the asset reporting email on or off. (HP-UX equivalent: `icapnotify -a`)

#### Format

ICAP SET ASSET *[qualifier]*

#### Qualifiers

`/STATE=state`: specify ON or OFF for the state qualifier value.

**EMAIL** Sets the system contact email addresses. (HP-UX equivalent: `icapmodify -c`)

#### Format

ICAP SET EMAIL *qualifiers*

#### Qualifiers

`/CONTACT`: The email address that receives the configuration change notifications and exception reports.

`/FROM`: The From address for the email sent from the iCAP system.

**NOTIFICATION** Sets the iCAP change configuration email notifications on or off. (HP-UX equivalent: `icapnotify -n`)

#### Format

ICAP SET NOTIFICATION *[qualifier]*

#### Qualifiers

`/STATE=state`: specify ON or OFF for the state qualifier value.

**SYSTEM\_ID** Sets the system identification used for iCAP asset reporting. (HP-UX equivalent: `icapmodify -I`)

#### Format

ICAP SET SYSTEM\_ID *"id"*

### Value

*id*: A user-defined string to identify this system when tracking or reporting usage. Specify a null string (" ") to set the system ID to the default value. The default value is the local hostname.

WARNING\_DAYS Sets the temporary capacity warning period to the number of days specified. (HP-UX equivalent: `icapmodify -w`)

### Format

```
ICAP SET WARNING_DAYS days
```

### Value

*days*: the number of days of temporary capacity before temporary capacity expiration warning email is sent to the system contact.

## ICAP SHOW

### Name

ICAP SHOW - Show the status and settings of the iCAP software on the OpenVMS system. (HP-UX equivalent: `icapstatus`)

### Format

ICAP SHOW STATUS [*qualifiers*]

### Parameter

STATUS Show the iCAP status and system settings to the standard output device.

### Qualifiers

/SNAPSHOT Creates a string of snapshot information containing encrypted audit data and displays the string to the standard output device. (HP-UX equivalent: `icapstatus -s`)

## ICAP\_SERVER

### Name

ICAP\_SERVER - iCAP server process.

### Description

The ICAP\_SERVER process performs the same functions as the `icapd` daemon process on HP-UX systems. For more information, see the HP-UX `icapd` manpage. To ensure compliance, the ICAP\_SERVER is always running on OpenVMS systems in an iCAP complex.

# Special OpenVMS-Specific Features and Considerations

## Core Activation and Deactivation

Unlike HP-UX, the OpenVMS operating system provides a user interface to start and stop system processor resources. When you enter the `START /CPU` command on an OpenVMS system in a complex containing iCAP resources, the `ICAP_SERVER` process validates that the start operation does not take the complex out of compliance. When you enter the `STOP /CPU` command, the CPU might restart at a later time if the count of intended active cores on the system is greater than the actual active cores.



**IMPORTANT:** The `ICAP` command or the corresponding HP-UX foreign commands must be used on OpenVMS systems when stopping and starting in complexes containing iCAP components. Using the `START /CPU` command can result in unintended consequences, such as an unexpected usage of temporary capacity or the deactivation of cores on the system or on another system in the complex. Using the `STOP /CPU` command can result in an unexpected restart of the core or the unexpected start of a core in another system in the complex.

To start cores on OpenVMS, use the `ICAP ACTIVATE/CPU=` command. To stop cores on OpenVMS, use the `ICAP DEACTIVATE/CPU=` command.

---

## Email Considerations

The iCAP software requires that SMTP mail be configured on the OpenVMS system to send email to the system contact. For more information about setting up SMTP mail, see your IP provider's documentation.

## Restrictions

- Instant Capacity software on OpenVMS Version 8.4 does not support HP virtual partitioning (vPars).
- Global Instant Capacity features, including the use of the `icapmanage` command, are not supported on OpenVMS.
- Instant Capacity on OpenVMS does not support internationalization. Only English language support is provided.
- LPMC and HPMC are not available on OpenVMS systems.



---

# Glossary

<b>activate cell</b>	The process of changing an <i>inactive cell</i> into an <i>active cell</i> . A <i>cell</i> is added to a partition using the <code>parmodify</code> and <code>parcreate</code> commands, and is activated through a reboot or reconfig, or through <i>cell online activation</i> .
<b>activated core</b>	A <i>core</i> that has been turned on by the Instant Capacity software or during installation. Cores are activated with the <code>icapmodify</code> command (or the <code>vparmodify</code> command in an HP-UX virtual partition) while HP-UX or OpenVMS is running.
<b>active cell</b>	A <i>cell</i> that is available for use by the software running on the <i>nPartition</i> . This implies that the cell's processors and memory (and I/O, if the cell is attached to an I/O chassis) are all available for use by the operating system. An active cell has the following characteristics: <ul style="list-style-type: none"><li>• It is present and populated.</li><li>• It is powered on.</li><li>• It is assigned to an <i>nPartition</i>.</li><li>• It is released from <i>boot-is-blocked</i>.</li></ul>
<b>active nPartition</b>	An <i>nPartition</i> with at least one <i>active cell</i> . <i>See also</i> inactive <i>nPartition</i> .
<b>actual active</b>	The current number of active cores in the partition. In a virtual partition, the count represents the total number of cores assigned to all virtual partitions. The number of actual active cores for each partition is displayed using the <code>icapstatus</code> command. <i>See also</i> intended active.
<b>add-on system</b>	A system that has been converted to an Instant Capacity system. This process is performed by an HP service representative.
<b>BCH</b>	Boot console handler. The system firmware user interface that allows boot-related configuration changes and operations on PA-RISC systems. For example, BCH provides a way to specify boot options and the choice of boot devices. The <i>EFI</i> Boot Manager provides a similar function for Intel™ Itanium™-based systems.
<b>BIB</b>	Boot-is-blocked. The state of a <i>cell</i> that is powered on but is not allowed to boot. BIB exists as soon as power is enabled to a cell, although the system firmware completes its power-on self-test sequence before waiting for BIB to be cleared by the Service Processor. BIB is cleared when the Service Processor is told to boot an <i>nPartition</i> . BIB is also cleared when the system firmware determines that there is no active Service Processor in a complex.
<b>boot console handler</b>	<i>See</i> BCH.
<b>boot-is-blocked</b>	<i>See</i> BIB.
<b>bound core</b>	For <i>vPars</i> versions before A.04, a <i>core</i> that can process interrupts for a <i>virtual partition</i> . Bound cores cannot be migrated from one virtual partition to another if either of the partitions is running. Every virtual partition must have at least one bound core.
<b>cell</b>	A circuit board that contains processors and memory, controlled by a cell controller (CC). A cell is the basic building block of an <i>nPartition</i> in a complex.
<b>cell OLA</b>	Cell online activation. The process of changing an <i>inactive cell</i> to an <i>active cell</i> in a booted partition without requiring a reboot. If such a cell is attached to a powered-on I/O chassis, then the chassis is also activated as part of the cell OLA. Some operating systems do not support online activation of cells.
<b>cell online activation</b>	<i>See</i> cell OLA.
<b>cell-based server</b>	A server in which all <i>cores</i> and memory are contained in <i>cells</i> , each of which can be assigned for exclusive use by an <i>nPartition</i> . Each <i>nPartition</i> runs its own instance of an operating system.
<b>codeword</b>	The mechanism used with Instant Capacity software B.06.x and later for adjusting available <i>usage rights</i> for system components ( <i>RTU</i> codewords), for applying an amount of temporary capacity to a system ( <i>TiCAP</i> codewords), and for applying <i>sharing rights</i> to a GiCAP system to

	enable the creation of one or more groups (GiCAP codewords). Codewords are purchased from HP and retrieved from the <i>Utility Pricing Solutions Portal</i> . <i>See also</i> RTU, sharing rights, usage rights.
<b>configured processor core</b>	A processor that is configured at the boot console handler ( <i>BCH</i> or <i>EFI</i> ) and whose cores are now available for activation by the Instant Capacity software.
<b>core</b>	The actual data-processing engine within a processor. A single processor can have multiple cores, and a core can support multiple execution threads.
<b>deactivate cell</b>	The process of changing an <i>active cell</i> into an <i>inactive cell</i> . A <i>cell</i> becomes inactive when a <i>shutdown for reconfiguration</i> operation is performed on its <i>nPartition</i> . A cell can also be deactivated by setting its <i>use-on-next-boot</i> value to <i>No</i> and then performing a <i>reboot for reconfiguration</i> operation on the <i>nPartition</i> , or dynamically deactivated through a cell online deactivation (cell OLD) without rebooting.
<b>deactivated core</b>	<i>See</i> inactive core.
<b>deconfigured processor</b>	A processor that has not yet been configured at the boot console handler ( <i>BCH</i> or <i>EFI</i> ). The Instant Capacity software cannot activate a processor core that is deconfigured.
<b>EFI</b>	Extensible Firmware Interface. The system firmware user interface that allows boot-related configuration changes and operations on Intel™ Itanium™-based systems. For example, EFI provides ways to specify boot options and list boot devices. The <i>boot console handler (BCH)</i> provides a similar function for PA-RISC systems.
<b>Extensible Firmware Interface</b>	<i>See</i> EFI.
<b>failback</b>	The process of restoring a system in a failover state back to its original state.
<b>failover</b>	The operation that takes place when a primary service (network, storage, or CPU) fails, and the application continues operation on a secondary unit.
<b>GiCAP</b>	Global Instant Capacity. Software that enables you to move usage rights for Instant Capacity components within a group of servers.
<b>Global Instant Capacity</b>	<i>See</i> GiCAP.
<b>guest</b>	<i>See</i> virtual machine.
<b>guest OS</b>	The operating system that is running on a <i>virtual machine</i> .
<b>HA</b>	High availability. The ability of a server or partition to continue operating despite the failure of one or more components. High availability requires redundant resources, such as CPU resources and memory, in specific combinations.
<b>hard partition</b>	<i>See</i> nPartition.
<b>high availability</b>	<i>See</i> HA.
<b>host</b>	<ol style="list-style-type: none"> <li>1. A system or partition that is running an instance of an operating system.</li> <li>2. The physical machine that is the <i>VM Host</i> for one or more <i>virtual machines</i>.</li> </ol>
<b>host name</b>	The name of a system or partition that is running an OS instance.
<b>host OS</b>	The operating system that is running on the host machine.
<b>hyper-threading</b>	The logical CPU feature that allows multiple execution threads (logical CPUs) within a single core. Hyper-threading does not result in a true multicore processor, but it can add performance benefits. True multicore processors typically deliver much greater performance than equivalent hyper-threading technology.
<b>iCAP component</b>	<i>See</i> Instant Capacity component.
<b>iCAP core</b>	<i>See</i> Instant Capacity core.
<b>inactive cell</b>	A <i>cell</i> that is not available for use by software running on an <i>nPartition</i> . This term usually describes a cell that has the following status (though any cell that is not <i>active</i> is by definition inactive). <ul style="list-style-type: none"> <li>• The slot is present and is populated.</li> <li>• Power is enabled.</li> </ul>

- *Boot-is-blocked*.
- The cell is assigned to an nPartition.

See also active cell.

<b>inactive core</b>	A <i>core</i> that either has not yet been activated or that has been turned off by the Instant Capacity software and returned to the pool of inactive cores. Inactive cores are available for activation. New HP-UX or OpenVMS processes are not assigned to an inactive core, and all processes running on the inactive core are migrated to other cores (with the exception that interrupt handlers might not be migrated from inactive cores).
<b>inactive nPartition</b>	An <i>nPartition</i> in which all of its cells are <i>inactive</i> . See also active nPartition.
<b>Instant Access Capacity</b>	Also called IAC. The amount of <i>temporary capacity</i> included with the purchase of an Instant Capacity component.
<b>Instant Capacity</b>	Also called iCAP. The HP Utility Pricing Solutions product that allows you to purchase and install additional processing power through the use of a two-step purchase model. Initially, you purchase system components ( <i>cores</i> , <i>cell</i> boards, memory) at a fraction of the regular price because the usage rights are not included. These Instant Capacity components are inactive but are installed and ready for use. When extra capacity is needed, you pay the remainder of the regular price for the usage rights to activate the components. If the regular price for the component is reduced by the time the usage rights are purchased, the remainder price is proportionally reduced, providing additional savings.  Earlier versions of iCAP were referred to as Instant Capacity on Demand, or iCOD.
<b>Instant Capacity component</b>	Also called a component without usage rights. A <i>core</i> , <i>cell</i> board, or memory that is physically installed in an Instant Capacity system but is not authorized for use. Before it can be used, an <i>RTU</i> must be purchased and a <i>codeword</i> must be applied to the system.
<b>Instant Capacity core</b>	Also called a core without usage rights. A <i>core</i> that is physically installed in an Instant Capacity system but that does not have usage rights and is not activated. After obtaining usage rights, Instant Capacity cores can be turned on by the Instant Capacity software or during installation. Cores with usage rights are activated with the <code>icapmodify</code> command (or the <code>vparmodify</code> command in a virtual partition) while HP-UX or OpenVMS is running.
<b>Integrity Virtual Machines</b>	See Integrity VM.
<b>Integrity VM</b>	A soft partitioning virtualization product that allows you to install and run multiple systems ( <i>virtual machines</i> ) on the same physical host system (Integrity server or nPartition). The Integrity server or nPartition acts as a <i>VM Host</i> for the <i>virtual machines</i> (also referred to as guests). The virtual machines share a single set of physical hardware resources, yet each <i>virtual machine</i> is a complete environment in itself and runs its own instance of an operating system (referred to as a guest OS). See also virtual machine, VM Host.
<b>intended active</b>	The number of cores a user requests to be active for a partition by the Instant Capacity software at the next reconcile operation. A reconcile operation is normally a reboot, although other actions can trigger a reconcile operation, such as moving cores between virtual partitions. You adjust the number of intended active cores by using the <code>icapmodify -a</code> , <code>-d</code> , and <code>-s</code> options. Other commands, such as <code>parmodify</code> and <code>parcreate</code> , can also affect this value. The number of intended active cores for each partition is displayed using the <code>icapstatus</code> command. See also actual active.
<b>load balancing</b>	The distribution of processing activity across two or more servers in order to avoid overloading any one server. Load balancing is performed on Instant Capacity systems by activating or deactivating resources on partitions and spreading the load across the system (or across members of a GiCAP group) so that no partition or member of a group is overloaded.
<b>local nPartition</b>	When an <i>nPartition</i> command is being executed, the nPartition that is running the command.
<b>migrating processing cores</b>	The process of activating and deactivating cores across partitions or across members in a GiCAP group for load balancing.

<b>monarch processor</b>	Also known as the boot processor. The main controlling <i>core</i> of the operating system. This core is designated as CPU 0. The LPMC monitor does not deactivate or replace a failing monarch processor.
<b>nPartition</b>	Also known as a hard partition. A <i>partition</i> in a <i>cell-based server</i> that consists of one or more <i>cells</i> and one or more I/O chassis. Each nPartition operates independently of other nPartitions and either runs a single instance of an operating system or is further divided into <i>virtual partitions</i> .
<b>nPartition Provider</b>	The <i>WBEM</i> services provider for <i>nPartition</i> information about <i>cell-based servers</i> .
<b>online activation</b>	The ability to activate a <i>deactivated core</i> using Instant Capacity software while HP-UX or OpenVMS is running. No reboot is required. This is done with the <code>icapmodify</code> command, or, in a virtual partition, the <code>vparmodify</code> command. Online activation is the default behavior of the Instant Capacity software.
<b>partition</b>	A subset of server hardware that includes <i>core</i> , memory, and I/O resources on which an operating system (OS) can be run. A partition can be either a software partition ( <i>virtual partition</i> ) or a hard partition ( <i>nPartition</i> ). This type of partitioning allows a single server to run an OS independently in each partition with isolation from other partitions. <i>See also</i> nPartition, virtual partition.
<b>Pay per use</b>	Also called PPU. The HP software product that is a part of the HP Utility Pricing Solutions program. The customer acquires a specific hardware platform and number of cores and then is charged for usage of the cores based on system demand.
<b>processor</b>	The hardware component that plugs into a processor socket. Processors can contain more than one <i>core</i> .
<b>reboot for reconfiguration</b>	The process of rebooting an <i>nPartition</i> so that all <i>active cells</i> in the nPartition are reset with <i>boot-is-blocked</i> (BIB) set. When the operating system running on the nPartition has finished shutting down, these cells begin their power-on self-test sequence, then wait for BIB to be cleared by the <i>Service Processor</i> . When all of the cells in the nPartition complete self-test, the Service Processor boots the nPartition.  On the HP-UX operating system, reboot for reconfiguration is performed using the <code>reboot</code> or <code>shutdown</code> command with the <code>-R</code> option. Do not use the <code>-H</code> option to ensure that the nPartition automatically reboots after reconfiguration.  <i>See also</i> shutdown for reconfiguration.
<b>Right to Use</b>	<i>See</i> RTU.
<b>rights migration</b>	On a GiCAP system, the transfer of core, cell, or memory usage rights by deactivating cores, cells, or memory in an active partition and activating them elsewhere in a GiCAP group. Rights migration cannot be performed on an inaccessible partition. Migrated usage rights do not expire.  Rights migration is also used to describe the activation and deactivation sequence used for <i>load balancing</i> partitions of a single server.
<b>rights seizure</b>	The acquisition of core usage rights from a failed partition on a GiCAP group member, which might be off line because of some disaster, transferring them to the Group Manager to make them available to other group members. Seized usage rights from a fully unavailable member normally expire after a period of time and revert to the original member from which they were seized.
<b>RTU</b>	Right to Use. A type of <i>codeword</i> used to activate and adjust available usage rights for Instant Capacity components (memory, <i>cell</i> board, or <i>core</i> ). An RTU codeword can be applied only to the system for which it was purchased, and the application of an RTU codeword adjusts the number of component-specific usage rights on the system.  <i>See also</i> codeword, usage rights.
<b>Service Processor</b>	An independent support processor for HP servers that support <i>nPartitions</i> . The Service Processor provides a menu of service-level commands plus commands to reset and reboot nPartitions and configure various parameters.  In HP servers, the Service Processor is sometimes called the Management Processor (MP) or the Guardian Service Processor (GSP).

<b>sharing rights</b>	A type of <i>codeword</i> applied to a Group Manager to enable the addition of members with Instant Capacity components to groups. To share resources across groups, you must purchase <i>GiCAP</i> sharing rights, acquire the GiCAP codeword from the HP <i>Utility Pricing Solutions Portal</i> ( <a href="http://www.hp.com/go/icap/portal">http://www.hp.com/go/icap/portal</a> ), and apply the associated codeword to the Group Manager system. You purchase at least as many GiCAP sharing rights as the total number of cores without usage rights across all the potential group members. Members can be added to a GiCAP group as long as there are sufficient sharing rights available and as long as the grouping rules indicate hardware compatibility. <i>See also</i> codeword.
<b>shutdown for reconfiguration</b>	The process of shutting down an <i>nPartition</i> in such a way that all <i>active cells</i> in the <i>nPartition</i> are reset with the <i>boot-is-blocked (BIB)</i> attribute. When the operating system that is running on the <i>nPartition</i> has finished shutting down, these cells begin their power-on self-test sequence and then wait for BIB to be cleared by the <i>Service Processor</i> . As a result, the <i>nPartition</i> becomes <i>inactive</i> .  On the HP-UX operating system, shutdown for reconfiguration is performed using the <code>shut down</code> or <code>reboot</code> commands with the <code>-R</code> and <code>-H</code> (or <code>-RH</code> ) options. <i>See also</i> <code>reboot</code> for reconfiguration.
<b>system</b>	A server, <i>nPartition</i> , <i>virtual partition</i> , or <i>virtual machine</i> that is running an instance of an operating system.
<b>Temporary Instant Capacity</b>	An HP product that enables customers to purchase prepaid core activation usage rights, for a specified (temporary) period of time. Temporary capacity is sold in 30 processing-day increments. Temporary capacity was formerly known as Temporary Instant Capacity on Demand, or TiCOD.
<b>TiCAP</b>	<i>See</i> Temporary Instant Capacity.
<b>TiCOD</b>	<i>See</i> Temporary Instant Capacity.
<b>unused capacity</b>	The difference between actual active and intended active cores. Unused capacity can be a side effect of using the <code>vparmodify</code> command to deactivate cores without immediately activating the cores somewhere else. This is usually a transient state because a user typically migrates cores from one virtual partition to another with a deactivate command followed by an activate command. However, unused capacity can persist if, for example, a utility such as <code>gWLM</code> ignores an error status from an activation and leaves the previously deactivated cores in an unassigned state. The Instant Capacity software always takes unused capacity into account when a request is made to activate or deactivate cores and attempts to eliminate the discrepancy between intended active and actual active.
<b>usage rights</b>	Usage rights are used by Instant Capacity to activate system components (memory, <i>cell</i> boards, and <i>cores</i> ), and are freed by deactivating components. Usage rights for a complex are adjusted by the application of a Right to Use ( <i>RTU</i> ) codeword, and they can be shared between systems through the use of <i>Global Instant Capacity</i> . <i>See also</i> RTU.
<b>use-on-next-boot</b>	A per- <i>cell</i> flag in the Partition Configuration Data. This flag is used by system firmware during the process of booting an <i>nPartition</i> . If a cell is assigned to an <i>nPartition</i> and this flag is not set, then the cell is not activated the next time the <i>nPartition</i> is booted.
<b>Utility Pricing Solutions Portal</b>	An HP website ( <a href="http://www.hp.com/go/icap/portal">http://www.hp.com/go/icap/portal</a> ) that gives customers an interface to obtain <i>codewords</i> for Instant Capacity systems, view asset reports, view <i>Pay per use</i> system-utilization information and to check other information related to their Utility Pricing Solutions accounts.
<b>virtual machine</b>	A software entity provided by HP <i>Integrity Virtual Machines</i> , VMware ESX, or Microsoft Virtual Server. This technology allows a single server or (with Integrity Virtual Machines) <i>nPartition</i> to act as a <i>VM Host</i> for multiple individual virtual machines, each running its own instance of an operating system (referred to as a <i>guest OS</i> ). Virtual machines are servers in the HP Virtual Server Environment (VSE).
<b>virtual partition</b>	A software <i>partition</i> of a server, or of a single <i>nPartition</i> , where each virtual partition can run its own instance of an operating system. A virtual partition cannot span an <i>nPartition</i> boundary. <i>See also</i> <i>nPartition</i> , virtual machine.
<b>VM</b>	<i>See</i> virtual machine.

<b>VM Host</b>	A server running software such as <i>HP Integrity Virtual Machines</i> , VMware ESX, or Microsoft Virtual Server, that provides multiple <i>virtual machines</i> , each running its own instance of an operating system.
<b>vPars</b>	An HP software product that provides <i>virtual partitions</i> . <i>See also</i> virtual machine, virtual partition.
<b>WBEM</b>	Web-Based Enterprise Management. A set of web-based information services standards developed by the Distributed Management Task Force, Inc. A WBEM provider offers access to a resource. WBEM clients send requests to providers to get information about and access to the registered resources. <i>See also</i> nPartition Provider.

---

# Index

## A

- activating cores, 57
  - virtual partition environment, 64, 65
- Administration System, 22
- asset report
  - testing email transmission, 205
- assigning cell to partition, 67
- assumed values in icapstatus command, 192
- audit application, 22

## B

- boot time compliance, 66
- boot time enforcement, 38

## C

- cell
  - assigning to partition, 67
  - unassigning from partition, 69
- cell boards, 33
  - activating, 40, 92
- cell removal implications, 197
- cimconfig command, 46
- cimserver command, 46
- codewords, 36
  - applying, 31, 36, 55
  - TiCAP, 37
- compliance and enforcement, 38
- compliance exceptions, 137
- components, 33
  - activation, 33
- configuration change notification, 39
- configuring email, 202
- core activation, 40, 57
  - deconfigured cores, 57
  - example, 58
  - OpenVMS, 220
  - virtual partition environment, 64, 65
- core deactivation, 59
  - example, 59
  - virtual partition environment, 64, 65
- core test activation, 72
- cores, 33
  - activating, 40
  - replacement of failed, 73
- creating a new partition, 196

## D

- database, 22
- deactivating cores, 59
  - virtual partition environment, 64, 65
- deferred activation, 57
  - overriding, 61
- deferred deactivation, 59
- disaster recovery, 122
- dual-core support, 195

dynamic processor resilience, 207

## E

- email
  - asset report transmission, 205
  - configuration change notification, 39
  - configuring, 202
  - configuring From address, 204
  - OpenVMS, 220
  - requirements, 30, 202
  - sent by Instant Capacity software, 145
  - TiCAP expiration reminder, 83
- email configuration
  - troubleshooting, 142
- exceptions, 137

## F

- failed core replacement, 73
  - on OpenVMS, 73
- failed monarch processors, 73
- failed primary processor
  - on OpenVMS, 73
- frequently asked questions, 144

## G

- Genesis partitions, 199
- GiCAP, 101
  - adding new partitions, 121
  - benefits, 102
  - codeword, 102
  - creating groups, 109
  - disaster recovery, 122
  - effect of temporary capacity, 112
  - Group Manager availability, 119, 120, 128
  - Group Manager failover, 120
  - Group Manager requirements, 105
  - Group Managers, 105
  - grouping rules, 35, 107
  - location of members, 128
  - log file, 148
  - multiple groups, 127
  - overview, 35, 102
  - ports, 30
  - reinstalling a group member, 118
  - reinstalling HP-UX, 48
  - removing a group member, 117
  - requirements, 104
  - resource sharing, 112
  - rights seizure, 122
  - sharing rights, 35, 108
  - standby Group Managers, 105
  - status reporting, 113
  - temporary capacity, 115
  - TiCAP prefetch, 113, 116
  - transferring database, 105
  - upgrades, 121

WBEM requirements, 104, 128  
Global Instant Capacity (*see* GiCAP)  
grouping rules, 35

## H

high availability, 102, 124  
HP OpenView measurement software, 206  
HP-UX  
    reinstalling, 48  
HP-UX 11i v1 requirements, 28  
HP-UX 11i v2 requirements, 29  
HP-UX 11i v3 requirements, 30  
HPMC  
    core failure and replacement, 73

## I

IAC (*see* Instant Access Capacity)  
ICAP\_SERVER time zone, 44  
icapd daemon time zone, 44  
icapmanage command, 105  
    adding group members, 109  
    adding partitions, 121  
    creating groups, 109  
    installing grouping rules, 107  
    removing partitions, 121  
    rights seizure, 122  
    standby Group Managers, 105  
    testing hardware compatibility, 109  
icapmodify command, 39, 54, 57, 64  
    deferred activation, 57  
    deferred deactivation, 59  
    example, 55, 59  
    increasing TiCAP warning period, 85  
    instant activation, 57  
    instant deactivation, 59  
icapnotify command, 39  
icapstatus command, 43, 52  
    assumed values, 192  
    example GiCAP output, 110  
    example output, 53  
    TiCAP use, 83  
inactive cells, 33  
inactive cores, 33  
installing Instant Capacity software, 46  
Instant Access Capacity, 18, 78  
Instant Capacity audit application, 22  
Instant Capacity Cell Board, 91  
    accidental activation, 98  
    activating, 97  
    activation exception, 99  
    and TiCAP, 100  
    license and support, 94  
    ordering, 93  
    overview, 41, 92  
    usage rights, 95  
    usage rights examples, 96  
Instant Capacity compliance and enforcement, 38  
Instant Capacity components, 33  
Instant Capacity database, 22

Instant Capacity portal, 22  
Instant Capacity program requirements, 28  
Instant Capacity software  
    downloading, 46  
    installing from HP Software Depot, 47  
    installing from HP-UX media, 46  
    installing on OpenVMS, 47  
    reinstalling, 48  
    removing, 49  
    validation, 47  
Instant Capacity software overview, 21  
Instant Capacity software products, 20  
Instant Capacity software requirements, 28  
    HP-UX 11i v1, 28  
    HP-UX 11i v2, 29  
    HP-UX 11i v3, 30  
    OpenVMS, 30  
Instant Capacity software validation, 42  
Instant Capacity status reporting, 43  
Instant Capacity supported hardware platforms, 23  
Instant Capacity supported OS versions, 23  
Instant Capacity system hardware, 21  
Instant Capacity system overview, 20  
Instant Capacity version history, 25  
instant deactivation, 59  
intended active  
    understanding and managing, 63

## L

load balancing, 102  
load-balancing active cores, 62  
log file history, 144  
LPMC  
    core failure and replacement, 73  
    deactivations in virtual partitions, 73

## M

measurement software and Instant Capacity, 206  
memory, 33  
monarch processor  
    replacement of failed, 73

## N

network requirements, 30  
noncompliance, 38  
nPartition  
    minimum active cores, 59  
    reinitializing, 199

## O

OpenVMS  
    CLI support, 210  
    core activation and deactivation, 220  
    core deactivation, 220  
    DCL commands, 212  
    DCL ICAP command, 212  
    email considerations, 220  
    HP-UX Command Mapping, 210  
    HP-UX style commands, 210

- ICAP ACTIVATE command, 212
- ICAP APPLY command, 213
- ICAP DEACTIVATE command, 214
- ICAP RECONCILE command, 215
- ICAP SET command, 216
- ICAP SHOW command, 218
- ICAP\_SERVER process, 219
- restrictions, 221
- special features and considerations, 220
- system files, 211
- OpenVMS considerations, 209
- OpenVMS requirements, 30
- overriding deferred activation, 61
- overriding deferred deactivation, 61

**P**

- par commands with PC System Management Station, 200
- parmodify command, 57
- partitions
  - shutting down, 57, 198
- patches
  - for HP-UX 11i v1, 29
  - for HP-UX 11i v2, 29
- PC System Management Station and par commands, 200
- performance optimization, 71
- portal, 22
- processing capacity
  - increasing, 40
- processors, 33
- program requirements, 28
- psets
  - compatibility with Instant Capacity, 201

**R**

- reinstalling
  - preserving Instant Capacity information, 48
- removing Instant Capacity software, 49
- rights migration, 122
- rights seizure, 122
  - and temporary capacity, 123
  - effects, 122
  - summary, 124

**RTU**

- acquiring codewords, 40
- application order, 55
- applying codeword, 40
- applying codewords, 31, 36, 55
- purchasing, 40, 55

**S**

- security issues, 208
- sharing rights, 35, 108
  - applying, 108
- shutdown command, 57
- shutting down a partition with Instant Capacity cores, 198
- software application considerations, 71
- software overview, 21
- software products, 20
- software requirements, 28
- software validation, 42
- sslClientVerificationMode
  - setting, 46
- static virtual partitions, 64
- status
  - reporting, 43
  - viewing, 52
- supported hardware platforms, 23
- supported OS versions, 23
- system contact
  - role requirement, 31
  - setting, 54
- system hardware, 21
- system management commands, 22
- system overview, 20

**T**

- temporary capacity, 75
- temporary instant capacity (*see* TiCAP)
- TiCAP, 75
  - acquiring and configuring, 79
  - activating cores, 79
  - and rights seizure, 123
  - and system compliance, 38, 76
  - and virtual partitions, 81
  - applying codeword, 79
  - applying codeword example, 79
  - ceasing use, 81
  - codewords, 37
  - compliance enforcement, 86
  - core test activation, 72
  - deactivating, 80
  - depletion, 87
  - example of activating processor, 79
  - exceptions, 87
  - expiration, 86
  - expiration reminder, 83
  - Group Manager prefetch, 116
  - HP-UX licensing, 78
  - icapstatus example, 83
  - in GiCAP systems, 112, 115
  - low balance, 87
  - negative balance, 38, 87
  - OpenVMS licensing, 78
  - ordering, 78
  - overview, 37, 76
  - pooling, 102, 112
  - prefetch, 113
  - reducing consumption, 81
  - tracking usage, 83
  - transferring, 76
  - using, 79
  - warning period, 85
- time zone
  - for ICAP\_SERVER, 44
  - for icapd daemon, 44
- troubleshooting, 137
  - compliance exceptions, 137

- email configuration, 142
- Instant Capacity software, 140

## U

- unassigning cell from partition, 69
- unused capacity, 64
- upgrades
  - GiCAP, 121
  - Instant Capacity version B.06.x or later, 193
- usage rights
  - purchasing, 31, 55
  - requirements, 31
- use-on-next-boot flag, 41, 92
- Utility Pricing Solutions portal, 22

## V

- version history, 25
- viewing status, 52
  - example, 53
- virtual partitions
  - activating cores, 65
  - deactivating cores, 65
  - integration with Instant Capacity, 31
  - static, 64
- vparmodify command, 57, 64, 65
  - reducing TiCAP consumption, 81
- vPars
  - integration with Instant Capacity, 31
  - required versions, 31

## W

- WBEM
  - GiCAP requirements, 128
  - requirements, 46
  - requirements for GiCAP, 104