IBM Hardware Management Console

Best Practices

Ron Barker
Minh Nguyen
Shamsundar Ashok

January 2007

Table of Contents

1 Overview

This paper is a compilation of useful information, recommendations and best practices for managing the IBM Hardware Management Console (HMC) for IBM System p5™ and IBM System i5™ servers. It includes suggestions on planning for an HMC, initial configuration, security, user management, problem determination, and code installation and maintenance.

IBM System p5 customers frequently ask about the security of the HMC. Addressing those questions was one of the motivations for writing this document. The security section includes a list of things IBM has done as of this writing to enhance the security of the HMC, as well as recommended customer practices. This is a constantly evolving area, however, and additional changes are to be expected over time.

It is assumed the reader has some knowledge of the HMC and its management, as well as reference to IBM documentation. This document is not intended to be an all-inclusive management handbook, but a supplement to existing documentation. Information about the HMC may be found online using the IBM Systems Hardware Information Center.

http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp

From that site, you can find files that can be downloaded and viewed in PDF format or printed. For a list of downloadable files as of this date, go to this link:

http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp?topic=/ipha8/icpdflist.htm

2 Planning

Planning should be done for any major change in any computing environment, including adding new servers, performing upgrades and implementing software changes. Careful planning involves creating a timeline and dividing the project into phases, each with a specific, stated outcome. In effective planning, you draw up a list of assignments and responsibilities, as well as document the current environment and the desired result.

As far as planning for the HMC is concerned, the customer needs to begin with some fairly simple questions: is an HMC needed? What models are available? Where are they set up? How do they connect to servers they manage? How are they maintained and by whom, and where can I find documentation?

2.1 Do I Need an HMC?

Mid-range and larger p5 and i5 servers need an HMC to create and manage logical partitions, dynamically reallocate resources, invoke Capacity on Demand, utilize Service Focal Point and facilitate hardware control. High-end servers with Bulk Power Controllers (BPC), such as the IBM System p5 model 590, p5-595 and p5-575 systems, require at least one HMC acting as a DHCP (Dynamic Host Configuration Protocol) server. Two HMCs are recommended for enhanced availability. Mission critical solutions, even those hosted on entry or mid-range servers, may benefit from having dual HMCs.

HMCs may not be cost-effective for distributed, entry level systems that nevertheless require the capabilities of Advanced POWER Virtualization. Entry level servers without an HMC can be configured with a hosting partition called the IBM Integrated Virtualization Manager (IVM). It provides a subset of HMC functions and a single point of control for small system virtualization. IVM does not offer the full range of management capabilities found on an HMC, but it may be sufficient for a small server with one to eight processors. IVM is a component of the Virtual I/O Server Version 1.2, which comes with purchase of the Advanced POWER Virtualization hardware feature code.

2.2 HMC Models

POWER5 systems use the IBM 7310 Hardware Management Console, which comes in rack-mounted and desktop configurations. The same HMC supports i5 and p5 servers. One HMC can support both i5 and p5 systems simultaneously. The rack-mounted models currently are the 7310-CR2 and 7310-CR3. The desk side units are 7310-C03 and 7310-C04.

The number of servers each HMC can manage varies by server size and complexity. Each server partition will have a connection to an Ethernet network, and the HMC will be logically connected to each partition via the network connection. In addition, there will be an Ethernet connection to the Flexible Service Processor (FSP) or Bulk Power Assembly (BPA) on each managed server. On dual-processor models that do not have a BPA, each FSP will be connected to the HMC, and the FSPs will be linked together over the second port on each FSP.

At the time of this writing, up to 48 entry and mid-range and 32 high-end servers can be managed by a single HMC, and up to a total of 254 logical partitions. HMC performance may vary, however, depending on the unique combination of servers and the number of partitions and I/O drawers implemented.

2.3 Physical Location of an HMC

An HMC should be located close to the servers it manages, nominally 50 feet. Service personnel use the HMC and its service applications to maintain systems and record service actions. For POWER4 servers, the distance between the HMC and the managed server was limited by the supported length of the serial cable connecting the two systems. On p5, the serial cable has been replaced by an Ethernet connection. The distance restriction is still necessary, however, to enable service personnel to go back and forth between an HMC and a managed server during a service call.

2.4 Planning for Network Connectivity

There are two types of networks described in p5 and HMC documentation. One is a called a “private” network and the other is called “open.”

An open network is the easiest to describe. It means any standard network connection, such as would be used to connect an HMC and a logical partition, or an HMC and a remote workstation.

The private network is a non-routable subnet. It may sometimes be referred to as a service network. In the context of the HMC, it is nearly always true that a single HMC will be the DHCP server for a private network. The only exception would be a configuration including a Cluster-Ready Hardware Server involving a Cluster 1600 High Performance Switch. That configuration is beyond the scope of this white paper. For the vast majority of HMC installations, a private network describes one HMC acting as a DHCP server connected over a non-routable subnet to one or more FSPs (FSP) or BPAs.

A server with dual HMCs would be connected to two private networks with each HMC acting as a DHCP server on a unique, non-routable subnet.

In order to attach multiple p5 servers to one or a pair of HMCs, network switches may be required. If you are planning to implement private networks over a switch that supports VLAN (Virtual Local Area Network) technology, be sure that a broadcast from the FSP will reach the HMC DHCP server quickly before the FSP port goes to its default IP address. For example, if the switch port must have spanning tree enabled, it should also have PortFast or the equivalent enabled.

As an additional step, determine whether the switch requires that the network interface on the HMC be set to a specific speed or whether auto-detect may be used. Hubs generally require the HMC to be set to a specific speed and duplex setting.

2.5 Private vs. Open

The network connection between the HMC and the FSP can be either private or open on low-end to mid-range servers. Private is preferred, and therefore a best practice. A private network is *required* for systems that have a BPA, such as the models 590, 595 and 575.

On a private network, the HMC acts as a DHCP server for the managed systems' FSPs and BPCs. The IP address is assigned from a range of non-routable addresses selected by the systems administrator when he configures DHCP on the HMC. There are 20 possible ranges to choose from, and all are non-routable subnets. The non-routable subnets isolate the HMC and the FSPs from other HMC network interfaces.

An HMC may also manage FSPs over an open network on low-end and mid-range systems. This scenario requires that the FSPs be network reachable from the HMC. All HMC-to-FSP communication is SSL encrypted, whether over a private or open network.

In an open configuration, the FSP's IP addresses must be set manually on each managed server. They cannot be DHCP clients of any server other than a managing HMC.

Addresses can be set using the Advanced System Management Interface (ASMI) on the FSP. This involves directly connecting a laptop to one of the ports on the FSP and using HTTPS to log into one of the two pre-defined IP addresses. The HMC1 port defaults to 192.168.2.147; HMC2 defaults to 192.168.3.147. The systems administrator can login as user "admin" using the default password "admin," which should be changed during the initial installation for security reasons. If no laptop is available, an ASCII terminal can be used on the native serial port to access the FSP menus in character mode.

Open networks are used for communications between a logical partition and the HMC. This connection is largely to facilitate traffic over the Resource Monitoring and Control (RMC) subsystem, which is the backbone of Service Focal Point (SFP) and required for dynamic resource allocation. The open network also is the means by which remote workstations may access the HMC, and it could be the path by which an HMC communicates with IBM Service through an Internet connection.

Regardless of which type of network is involved, customers must provide their own networking infrastructure, such as cables, switches or hubs. Switches that support virtual networks (VLAN) may be used to create one or more private or open networks, as conditions require.

2.6 Customer Setup

The HMC running in a System p5 environment is a customer setup machine. Customers can purchase an IBM Machine Control Program Remote Support Agreement (MCPRSA) that provides phone technical support for routine installation, configuration and how-to questions. Contracts can be purchased for one- and three-year periods. The agreement covers the cost of all Licensed Machine Code upgrades during the contract. Hardware service contracts for on-site hardware support are available beyond the initial warranty period.

Customers are responsible for installing and updating the Licensed Machine Code on all HMCs and System p5 managed servers. We will discuss an update strategy in detail later. Systems administrators should become familiar with the information and tools available on the HMC support home page:

http://www14.software.ibm.com/webapp/set2/sas/f/hmc/home.html

The systems administrator should use the link on the above URL to sign up for subscription services. This enables an administrator to receive email notification of new releases of software and firmware. Knowing what is current is essential to correctly and successfully managing System p5 servers and the HMC.

2.7 Documentation

Documentation for the IBM System p5 and i5 servers and the HMC is available through the IBM Systems Hardware Information Center. This is a browser-oriented information retrieval application that can be installed on servers and workstations, but which also is accessible over the Internet.

Updates can be downloaded over the Internet for workstation-based installations. The online version can be accessed at this URL:

http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp

One of the features provided by the online version is the ability to create checklists for planning and installation. The Information Center enables users to print documents or portions of them, as well as bookmark pages for easy reference in the future.

Careful attention should be paid to the interdependencies between HMC code levels and system firmware on System p5 servers. These relationships can be found on the firmware support web pages. This will be discussed in more detail later in this document, but it is important to understand that new system firmware may require an upgrade of HMC code. Since upgrading the HMC does not disrupt partition operation, and since dual or redundant HMCs are supported, keeping HMC code on the most current level is recommended as a best practice. The general rule is that the HMC must support the highest level of system firmware on any server that it manages.

3 Initial Configuration

HMCs come with Licensed Machine Code preinstalled, but you may need to reinstall if the code has been superseded or you have a disk failure. Customized setup and installation instructions come with all new machines or upgrades. Prior to receiving a new system, users can refer to the IBM Systems Hardware Information Center online to see what specific preparations may be needed.

http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp

3.1 Install and Configure the HMC First

If you are going to use a private network, you should install and configure your HMC before connecting it to a network or powering on the servers it will manage. This will enable you to configure the networks properly and bring up DHCP in an orderly manner and insure that the managed servers will connect correctly. The HMC should be located within 50 feet of the servers it manages and with easy access to networks and power. If a dial-up modem is going to be used to contact IBM Service, a telephone line also needs to be available.

The POWER5 HMC includes a Setup Wizard that can be used by a systems administrator to customize the system. The wizard starts the first time you login after a new installation. You are not required to use it. Administrators who are comfortable with the HMC Configuration menus are free to use them instead. The Setup Wizard can be run at any time from the main menu on the HMC console.

3.2 Changing Passwords

Whether you use the Setup Wizard or the HMC Configuration menus, the first task you should do is change the hscroot and root passwords from their default values. As shipped with the Licensed Machine Code, the hscroot password is “abc123.” It should be changed to a new, seven-character value. The default root password is “passw0rd,” and it needs to be changed as well. Logging in as root is disabled. Note that while you will not be using the root password for daily administration, you may need it from time to time when performing problem determination, usually with the assistance of IBM support or product engineering. Be sure to save all passwords in a secure location where they can be retrieved in an emergency.

3.3 Creating User IDs

Additional user IDs should be created on the HMC so that not every user is accessing the system with the same user ID and password, and not necessarily with the same level of authority.

Administrators with hmcsuperadmin authority, which is what hscroot has, should have their own user IDs and passwords. This will facilitate auditing administrative actions on the HMC. Other users may have other pre-defined roles with more restricted authority. As will be described later, the tasks and resources that can be accessed by users can be customized and is very granular.

A special, optional user ID hscpe can be created either initially or when needed. This is the user ID needed to gain root access to the system. The hscpe user can enter a password obtained from IBM that allows him to run the pesh command to override the restricted shell and switch-user to root. This also requires that the user know the root password. The password used to override the restricted shell is good for one day and must be obtained by contacting IBM Support and providing the HMC’s serial number.

3.4 Configuring Call-Home

For many years, IBM has offered a capability called “call-home” on its servers. This describes the ability for a server to automatically notify IBM Service in the event of a hardware problem, as well as the ability to transmit other service related information or vital product data as required.

On p5 HMCs, the configuration of call-home functionality has been simplified. The setup wizard automatically prompts for the necessary information. Using the HMC menus, look for the Service Applications folder and the Remote Support menu, and follow the prompts.

The Remote Support menu has several tasks, including configuring customer information and customizing outbound and inbound communications. If customer information hasn’t been configured correctly, the HMC will not be able to “call home” for support. Note that inbound communication is optional. When allowed, the customer has real-time control over the inbound session, and it can be terminated by the systems administrator.

Historically, call-home support involved dialing out by modem. That option is still available. Network support was added, if the customer was willing to configure a Virtual Private Network from the HMC through their corporate firewall to the Internet. In many cases, customers were not willing to allow that because although the traffic was secured by Internet Protocol Security (IPSec) and tunneling, the HMC’s IP address was exposed. At HMC V5R2, support was added for outbound, SSL socket based communications, and as of HMC V6R1, proxy support was added. With or without a proxy server, the SSL Internet option allows use of Network Address Translation (NAT) firewalls between an HMC and IBM Service, thus allowing the customer to hide the HMC’s true IP address behind their corporate firewall and send encrypted information to IBM.

3.5 Configure Electronic Service Agent

The final step during initial configuration is to authorize HMC administrators to access Electronic Service Agent (ESA) and to decide how events should be called to the attention of the customer.

The information that’s entered for the HMC on the Remote Support menus mentioned above includes the customer name, administrator name, email address, phone number, alternate phone number, fax number and alternate fax number. This enables IBM Service to contact the customer when a call-home service event occurs.

ESA may be configured to send email to customer accounts when service events are generated. The emails can be sent to distribution lists. Most customers will want to use the ESA filters to ensure that only serviceable events are sent via email, and not every message generated by the agent. Optionally, SNMP traps may also be configured to send notification to specific IP addresses when an event occurs. This could be used in conjunction with a network or system monitoring program.

After the configuration of Remote Support and ESA has been completed, the customer should test the process to make sure it works. The test option appears on the final configuration screen under Remote Support and Customize Outbound Connectivity.

4 Security

Physical security of the HMC is a customer responsibility. The HMC should be located in a secure room, if possible. Usually, because of its proximity to the servers it manages, the HMC will be located in a secured data center. However, when that is not possible, there are ways of providing additional protection against unauthorized physical access. These protections are mainly provided by changes in the BIOS settings on the Intel chip that powers the HMC:

- Change the startup device settings in BIOS to prevent the use of a Recovery CD or diskette to boot into single-user mode.
- Assign a power-on password in BIOS to prevent unauthorized changes to BIOS settings.
- Unattended start mode can be set in BIOS to allow the HMC to reboot without the power-on password following restoration of power after an unplanned outage. However, the keyboard and mouse at the local console will remain locked until the power-on password is entered.

4.1 Network Security

The HMC must be properly networked to perform its server management functions. The private or service network is used to communicate with FSPs, and the open network is used to collect serviceable events from managed servers and to dynamically reallocate resources. A network is also the means by which remote administrators access and manage the HMC itself.

Two versions of the Web-based System Manager (WebSM) client code (for Windows 2000 and later or Linux) reside on the HMC and are downloadable using a browser as follows:

http://<HMC_hostname>/remote_client.html

To download the client package from the HMC, the user is required to enter a valid HMC user ID and password. Once the WebSM client package has been installed, the user can connect to the HMC by:

- Entering `wsm <hmc_hostname>` if the user is on a remote Linux system, or
- Double clicking on the WebSM remote client icon on the Windows desktop

A login dialog is then displayed to prompt the user for an id and password.

4.2 Secure WebSM

A secure WebSM connection using Secure Socket Layer (SSL) code is also available on the HMC. The SSL code can also be downloaded as follows:

http://<HMC_hostname>/remote_client_security.html

The SSL protocol provides server authentication, data encryption and data integrity. The HMC can be configured to require all clients to connect via SSL or to give clients the option of connecting via SSL. The first option, required, is more secure and therefore is preferred as a best practice. The HMC Security Manager application, which can only be accessed by a system administrator with super user authority and which must be configured at the console, controls these options.

Secure WebSM uses the RSA public key cryptography algorithm. The user on the client is authenticated to the server by his login password. The user name and password are sent encrypted over the SSL socket. The SSL protocol protects against changes or substitutions to data transmissions between the server and client machines. All data transmissions between the server and client machines are encrypted by the SSL protocol using the RSA RC4 algorithm with 128 bits key used for bulk encryption, and 1024 bits key used in the key exchange of the SSL protocol.
It's important to understand that SSL encryption is not the default for WebSM clients. It must be configured.

4.3 WebSM Ports

On the HMC, a WebSM server runs under xinetd control and listens on port 9090. When a remote WebSM client connects to the HMC, the WebSM server first authenticates the user ID and password. Once the authentication is completed, an instance of a WebSM Server running a separate Java Virtual Machine will be created. A pair of ports in the range of 30000 and 30009 is used as the communications channel between this WebSM server and the remote WebSM client.

Unless SSL encryption has been enabled as described above, the packets sent between the client and the HMC are in clear text. Even when encryption is used in the main WebSM client, a virtual terminal session to a partition opened from the client is not encrypted because it uses a separate application.

Customers who choose not to use the remote management function can disable the remote WebSM and Apache servers through the HMC Configuration menus or by using the command ***chhmc***.

The following command disables all remote WebSM connections to the HMC:

```
$ chhmc –c websm –s disable
```

The following command disables the HTTP service on HMC:

```
$ chhmc –c http –s disable
```

4.4 Restricting Network Access

The POWER5 HMC enables a firewall to block all incoming network traffic, with the exception of a well known set of ports. Within these well known ports, further access restriction can be customized based on IP address or host name.

4.4.1 Power 5 HMC port information

Port	Protocol	Application	Enabled by Default	Modification allowed in Network Config	Security	NOTES
9090, 9940, 30000-30009	tcp	WebSM	no	yes	SSL	
22	tcp	ssh	no	yes	3DES	
80	tcp	http	yes	yes		
443	tcp	https	yes	yes	SSL	
657	udp/tcp	RMC	yes	yes		
9,920	tcp	FCS	yes	yes		Call Home
9,900	udp	FCS	yes	yes		Call Home
4,411	tcp	Web server	yes	yes		InfoCenter (V4)
4,412	tcp	Web server	yes	yes		InfoCenter (V5)
9,443	tcp	Secure Web Server	yes	yes	SSL	Remote ASM(V5)
9,735	tcp	Vtty	yes	yes		
2300, 2301	tcp	5250 console	yes	yes	2301 uses SSL	
6,000	tcp	X11	yes	no	Xhost -	
5,988	tcp	CIM	yes	yes	SSL	
9197, 9198	tcp	CIM Indication	yes	yes	SSL	
123	udp	NTP	no	yes		
1,701	udp	l2tp	yes	yes		
427	udp	SLP	no	yes		Used in Cluster
2,049	tcp	NFS	no	no		
69	tcp	TFTP	no	no		
n/a	icmp	Ping	yes	yes		
500, 4500	udp	IPSec	no	no		VPN
162	tcp/udp	SNMP	no	yes		
8899	tcp	CRHS	yes	yes		
12347-12348	udp	RSCT Peer Domain	yes	yes		

The firewall interface allows you to customize remote access to the HMC by IP address and network mask. An example of the customization screen is shown below.

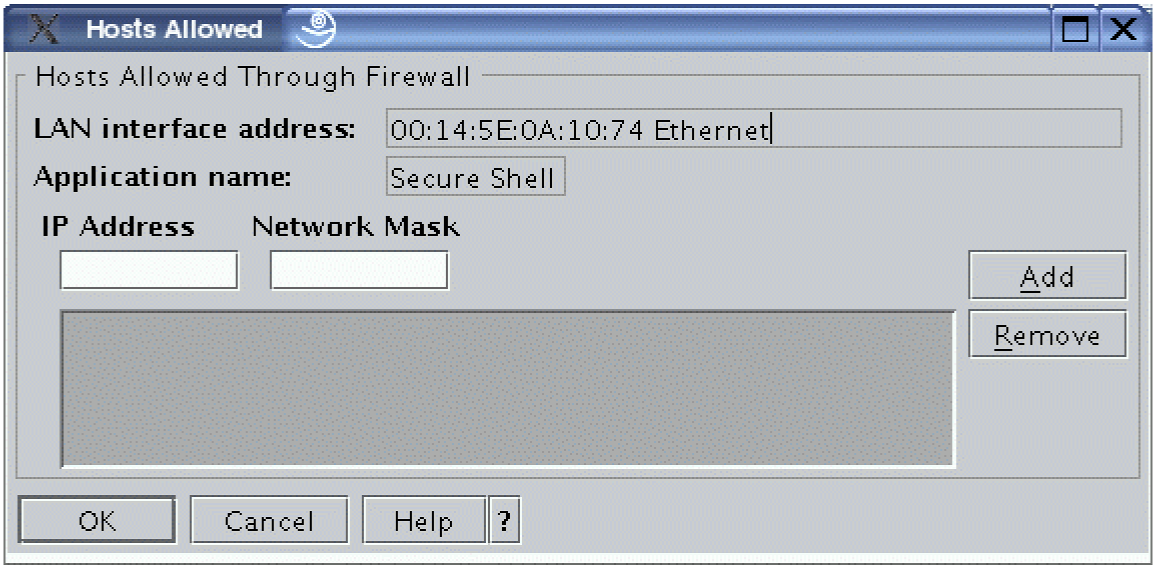

LAN interface address: 00:14:5E:0A:10:74 Ethernet

Application name: Secure Shell

4.5 Network Access between HMC and FSP

The HMC communicates with the FSP to perform its management functions. To do this, it establishes an SSL connection with port 30000 and 30001 of the FSP's Ethernet port. It is recommended that the network used for this communication channel be private, although an open network is supported except on systems with a Bulk Power Controller.

4.6 Restricted shell on the HMC

The HMC provides a rich set of commands that encompasses most of the tasks found in the graphical user interface. We choose to use SSH as a means to run these commands because it provides a secure way to perform remote command execution. However, by itself, SSH would provide an authenticated user full access to the shell. To protect the HMC from users trying to gain higher privileges by some means of exploiting the system, we are enforcing a restricted shell when remotely connecting to the HMC via SSH or when opening an rshterm on the HMC console. In the restricted shell environment, users will only have access to a small subset of operating system commands, along with all the HMC commands. Users will not be able to use the ***cd*** command, nor can they use re-direction.

Because the full list of HMC commands is published at the IBM HMC's web site, we will only reference a few security-related commands in this paper.

- **MKAUTHKEYS:** This command updates the caller's ***authorized_keys2*** file under the ***$HOME/.ssh/*** directory with a given DSA or RSA key generated from a client. Typically, on Linux and UNIX systems, the key can be generated using the ***ssh-keygen*** command. With the setup of the key in this file, a user can run HMC commands from a script without having to enter a password or passphrase.
- **MKHMCUSR:** This command creates a user on the HMC.
- **CHHMCUSR:** This command changes the properties of an HMC user. This command must be used to change hscroot's password. Using an operating system command to change the password will render the HMC unusable because hscroot's password is encrypted and safely saved away for communicating with various subsystems running on the HMC.
- **RMHMCUSR:** This command removes a user on the HMC. Root and hscroot users cannot be removed using this command.
- **LSHMCUSR:** This command lists users on the HMC.
- **CHHMC:** This command changes subsystems and network settings on the HMC, such as SSH, WebSM, syslog, http and network settings.
- **LSHMC:** This command displays various HMC's configuration settings, code version and Vital Product Data.
- **LSSVCEVENTS:** This command displays console events entries.
- **HMCSHUTDOWN:** This command can be used from a remote client to shut down or reboot the HMC. This command notifies the FSP on the managed server that it is gracefully going away. If this command is not used, the FSP will attempt to generate an error to indicate unexpected loss of communication with the HMC.
- **UPDHMC:** This command performs Software updates on the HMC. Software can be installed from a remote ftp server or locally from the DVD-RAM drive on the HMC.

4.7 Auditing capabilities on the HMC

A secure system also requires strong auditing capabilities. This section describes some of the logging/auditing functions on the HMC.

Most tasks performed on the HMC (either locally or remotely) are logged in a file ***iqyylog.log***. These entries can be viewed by using the View Console Events task, under the HMC Management—System Configuration application, or by using the ***lssvcevents*** command from the restricted shell. A log entry contains the timestamp, the user name and the task being performed. When a user logs in to the HMC locally or from a remote client, entries are also recorded in this file. For remote login, the client hostname or IP address are also captured. For example:

```
lssvcevents -t console
Earliest Timestamp    Description
10/16/03 07:27:51 PM  HSCE2175 User hscroot login failed
from remote host abcd.xyz.com with IP address
9.99.999.9999
```

Standard log entries that come from ***syslogd*** can be also be viewed on the HMC by viewing the file ***/var/hsc/log/secure***. This file can be read by users with System Administrator role. It is under ***logrotate*** control. A valid user can simply use the ***cat*** or ***tail*** command to view the file. A user with the System Administrator role could also use the ***scp*** command to securely copy the file to another system.

For customers who wish to copy ***syslogd*** entries to a remote system, the ***chhmc*** command can be used to change ***/etc/syslog.conf*** file on the HMC to specify a system to which to copy. For example, the following command line will cause syslog entries to be sent to the hostname *myremotesys.company.com*:

```
chhmc –c syslog –s add –h myremotesys.company.com
```

The systems administrator needs to make sure that the syslogd daemon running on the target system is setup to receive message from the network. On most Linux systems, this can be done by adding the –r option to the SYSLOGD_OPTIONS in file /etc/sysconfig/syslog.

In AIX, the /etc/syslog.conf file would be edited by uncommenting the appropriate lines at the bottom of the file, such as:

```
*.debug /tmp/syslog.out rotate size 100k files 4
*.crit /dev/console
```

Then the systems administrator would enter:

```
#  touch /tmp/syslog.out
# refresh –s syslogd
```

Another method of further limiting commands in the restricted shell along with logging the command(s) executed by users in syslog, is to use the *logssh* command. This command is introduced in HMC Version 5 Release 1.0, and can be added to the authorized_keys2 file to restrict a user from being able to open a pseudo-tty using SSH. To do this, the system administrator, hscroot, would enter:

```
     # ssh hscroot@stratus mkauthkeys -a -u minh
'"command=\"logssh \${SSH_ORIGINAL_COMMAND}\" ssh-
rsa
AAAAB3NzaC1yc2EAAAABIwAAAIEAzxTNjoXAvyZBw390oJ27uj
90PxZNtUWhYVN1/kaAfilSIr3z5Hhm7BdaaarUru94qhiM0xds
6cgQpNUQUy6GByoWDrNhdEIdAzXj3uaPscY6wKkNia0llTJPUo
BDBsadaa4oEc0/4poNG/X3uYrsdnbbMNkt/jmnEilSXIgOEmWk
= minh@somehost"'
```

This command sets up user *minh* on the *stratus* HMC so that he can login from *somehost* and run HMC commands using SSH. Each command executed by this user will be logged in *syslog*. In addition, user *minh* will not be able to open a pseudo-tty using SSH, cannot run the *mkauthkeys* command to undo this setup, nor can the *scp* command be used.

In this example, the string `‘ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAIEAzxTNjoXAvyZBw390oJ27uj90PxZNtUWhYVN1/kaAfilSIr3z5Hhm7BdaaarUru94qhiM0xds6cgQpNUQUy6GByoWDrNhdEIdAzXj3uaPscY6wKkNia0llTJPUoBDBsadaa4oEc0/4poNG/X3uYrsdnbbMNkt/jmnEilSXIgOEmWk= minh@somehost’` is generated by running the command ***ssh-keygen***, which is part of the OpenSSH package, on *somehost*, a remote workstation. When this command is run, it will ask for a passphrase, which should be left empty in order to eliminate prompting for passphrase when a script is running. Several SSH packages for various platforms and operating systems are available from a variety of sources.

4.8 Managing and Understanding Security Vulnerabilities on HMC

As stated in the HMC Code Update section, HMC users can subscribe to email notification of corrective service at the following web site:

http://www14.software.ibm.com/webapp/set2/sas/f/hmc/home.html

Whenever vulnerability is discovered on the HMC, a bulletin will be sent out to users on how to obtain the fix. In most cases, because of the closed nature of the HMC and the presence of the restricted shell, some vulnerability found on non-HMC systems will not apply. Each time a new release of HMC code is made available on the support web site, a list of security fixes included in the release is also published.

4.9 Resource Monitoring and Control

The Resource Monitoring and Control subsystem (RMC) is based on IBM’s Reliable, Scalable Cluster Technology (RSCT.) It is installed and used on the

HMC for establishing a trusted communication channel between the HMC and the partitions on the managed server. Some of the tasks performed through this channel include:

- Dynamic allocation of hardware resources on the partitions
- Graceful shutdown of the AIX operating systems running on the partitions
- Send hardware error log entries from the AIX partitions to the HMC to provide a single focal point for error collection

RMC uses port 657 for HMC-to-partition communication. Initially, the TCP protocol was used, but in recent releases of AIX and HMC code, the connectionless User Datagram Protocol (UDP) has been implemented. RMC employs access control lists to authenticate communication between the partitions and the HMC. The authentication is established during configuration steps on the HMC, thus, when transmitting messages over port 657, the HMC and the partition can be sure with whom they are communicating. For additional information regarding RMC and Reliable Cluster Scalable Technology refer to the Information Center.

4.10 CIM and Cluster System Management

The HMC uses Open CIMOM (Common Information Model Object Manager) to model the hardware resources of the managed servers. It is therefore CIM compliant and can provide information about its CIM objects to remote CIM clients. A CIM server runs on the HMC and listens on port 5988 for remote CIM requests. Only requests that supply a valid user ID and password on the HMC are honored. The Cluster System Management (CSM) managing server uses this facility on the HMC to perform various hardware control functions such as power on/off of partitions or servers in an IBM System Cluster 1600 environment. The same SSL protocol used by the WebSM client and server can be used to secure the communication between CIM clients and the HMC.

5 User Management

The HMC provides a User Management facility to create and manage users, as well as control their access to HMC tasks and managed system resources. User roles can be assigned to each HMC user along with the associated tasks the roles that can be performed. Before exploring how to use this facility, it’s helpful to establish some common terminology.

- **Resource –** Also called a managed resource, this is an entity to be managed in the system. Care should be taken to distinguish between ***resource type*** and ***resource instance***. Resource types represent all of the managed entities in the same category, such as logical partitions (LPARs). Resource instances mean individual managed entities, such as a specific LPAR.
- **Task –** In general, this is a function, typically a management operation performed on a resource.
- **Role -** A set of predefined configuration records for access management. When there is a need to assign a set of tasks for a certain resource, a role can be used to associate the tasks with the resources. An example of how roles can be of help is they allow separation of personnel management and work assignment in a customer environment. The Task and Managed Resource Roles section below contains more information on types of roles.

5.1 HMC Users and Access

The HMC comes with two predefined users: `hscroot` and `root`. Along with `hscpe`, these user IDs are special, i.e. reserved. The `hscroot` user ID and password are used to log in to the HMC when you configure the system after installation. The `root` user ID and password are used by a service provider to perform maintenance procedures, and cannot be used to log in to the HMC. The `hscpe` user should only be created for use by a service provider when performing problem determination. Afterward, it should be deleted, but recreated again when needed in the future. See the section on Problem Determination.

The `hscroot` and `root` are predefined user IDs and cannot be deleted from the HMC. They come with default passwords, but it is strongly recommended that they be changed during HMC setup and configuration. While user `root` cannot log in, `hscroot` can, and the `hscroot` user can change/set the password for the `root` user. `hscroot` is in effect a common user ID across all HMCs in a system environment. It has an `hmcsuperadmin` task role and can hence be dangerous if present across all HMCs with the common, default password. You will be prompted to change the password for both of these user IDs when using the HMC Guided Setup Wizard. The passwords can also be modified as described below.

To manage users employing the HMC WebSM console, a user with `hmcsuperadmin` authority would go to the **HMC Management** -> **HMC Users** plugin. There, select the **Manage HMC Users and Access** task. A window with user profiles for existing, defined users will be displayed. From the **User** menu you can **Add…**, **Modify…**, **Copy…**, or **Remove** a user. The following are a few notes on these tasks:

- When you add a user, the user IDs must not exceed 32 characters; passwords must exceed six characters; you can select a password expiration date by selecting the "Enable strict password rules option;" and you must choose a task role for this user. If you do not choose a managed resource role, 'AllSystemResources' will be selected by default. Please see the below section for more information on task and resource roles.
- When modifying user `hscroot` or `root`, you can only change the password, not the role.
- When logged in as a user without the `hmcsuperadmin` task role, it might be necessary to change your password. This can be accomplished by the **Change User Password** task in place of the **Manage HMC Users and Access** task mentioned above.
- When copying a user, the Add User dialog is effectively rendered with source user information (except password!) filled in. Pay careful attention to make necessary changes to create a new user ID at a minimum.

All the above tasks also can be accomplished through the HMC command line interface. Any user, regardless of task role, can use `lshmcusr` to view user profiles. Only a user with the `hmcsuperadmin` task role can run `mkhmcusr` or `rmhmcusr` to create or delete a user, respectively. While any user can run `chhmcusr` to modify an existing user, users who do not have `hmcsuperadmin` task role are limited to use of the `–t passwd` attribute to change their own password.

The `mkhmcusr` command only takes a task role as an argument. It does not allow for a managed resource role to be specified. As a result, it defaults to the AllSystemResources resource role. You can first create a user with this command, then use

```
chhmcusr –t assign –o a –r resourcerole | resource
```

Use chhmcusr with the `–v` attribute and value to assign a customized resource role (discussed below) or resource instance to the newly created user.

5.2 Task and Managed Resource Roles

As mentioned in the introduction to this section, a role is a way of grouping access privileges. On the HMC, roles are divided into two classes: task roles and managed resource roles.

A task role is a grouping of tasks, carried out either through WebSM or the Command Line Interface (CLI.) A managed resource role, which is also referred to as a resource role, is a grouping of resource types and/or resource instances, e.g. managed frames, managed systems and logical partitions. When specific partitions or servers are selected for a customized resource role, a user with this role can view and affect only those managed system.

On the HMC console view, the managed system must be visible under Server Management for specific partitions to be viewable to users with a specific role. Resources are viewed hierarchically. Thus, in this example the CLI would allow the managed system to be listed, but both the GUI and CLI would not allow any tasks outside of property views to be executed on the managed system unless they were authorized in the custom role definition. If a specific partition was selected in the resource role, partition views would be restricted to that logical partition.

As mentioned in the previous section, when creating a user you must specify a task role and one more resource roles. The HMC comes predefined with the following five task roles:

- `hmcsuperadmin` - The super administrator acts as the root user, or manager, of the HMC system. The super administrator has unrestricted authority to access and modify most of the HMC system. This should not be confused with user `root`.
- `hmcservicerep` - A service representative is generally someone physically at the managed system location to install, configure or repair managed systems.
- `hmcoperator` - An operator is responsible for daily system operation.
- `hmcpe` - A product engineer assists in support situations (for both the managed system and the HMC), but cannot access HMC user management functions. To provide support with access for your system, you must create and administer user IDs with the product engineer role; see section on Problem Determination.
- `hmcviewer` - A viewer can view HMC information, but cannot change any configuration information.

In addition, the HMC comes predefined with the AllSystemResources resource role. This is a dynamic resource role in that it is a container for all defined resources at any instance; it is not a static grouping of resources. This name is only defined for the graphical interface. It is not a recognized attribute by the command line interface:

```
hscroot@myhost:~> lshmcusr --filter names=hscroot
name=hscroot,taskrole=hmcsuperadmin,description=,pwage
=99999,resourcerole=
```

From the HMC console as a user with `hmcsuperadmin` authority, you can manage resource and task roles by selecting the **HMC Management** -> **HMC Users** plugin, then selecting **Manage Access Task Roles and Managed Resource Roles**. From the rendered dialog, you can choose to manage resource roles or task roles by selecting a role type. From the **Edit** menu, you can **Add…**, **Copy…**, **Remove**, or **Modify…** roles. Below are some notes on these tasks:

- You will be given an error if you try to modify or remove the predefined task roles, or the AllSystemResources resource role.
- You cannot copy the AllSystemResources resource role. You can only copy other, static managed resource roles.
- When adding or copying a resource role, ‘CEC Management’ and ‘All Logical Partitions’ are resource types for the managed system and LPARs, respectively. ‘CEC Management’ includes ‘All Logical Partitions,’ but not vice-versa. Note that LPAR profiles are implicit resources when a LPAR resource instance or type is selected. Similarly, system profiles resources are implicit with managed system resource instances or type.
- When adding or copying a task role, the available tasks to choose from are limited based on the parent role (‘Based on’ drop-down box).

Any user role can view task and resource roles from the restricted shell using the `lsaccfg` command. Below are some notes on this command:

- `lsaccfg -t resourcerole` won’t show AllSystemResources.
- `lsaccfg –t resource` does not show an equivalent to ‘CEC Management.’ It instead only shows all available managed system resources. It does have an ALL_PARTITIONS resource type however.
- `lsaccfg –t resource` shows managed resources in the form “<ResourceID=resource_type:resource_ID><UserDefinedName=name_value>”. The role commands typically are used in scripts, usually redirecting output to `mkaccfg` or `chaccfg` to make or change new roles. These commands won’t be input manually due to their complexity and verboseness. Therefore, there’s a `–script` flag that makes the output more manageable:
-
  ```
  hscroot@myhost:~> lsaccfg -t resource –script
  "resources=cec:root/ibmhscS1_0|9406-
  520*10007CA|IBMHSC_ComputerSystem,lpar:root/ibmhscS1_
  0|ALL_PARTITIONS*9406-
  520*10007CA|IBMHSC_Partition,lpar:root/ibmhscS1_0|7*9
  ```

```
406-
520*10007CA|IBMHSC_Partition,lpar:root/ibmhscS1_0|3*9
406-
520*10007CA|IBMHSC_Partition,lpar:root/ibmhscS1_0|6*9
406-
520*10007CA|IBMHSC_Partition,lpar:root/ibmhscS1_0|2*9
406-
520*10007CA|IBMHSC_Partition,lpar:root/ibmhscS1_0|10*
9406-
520*10007CA|IBMHSC_Partition,lpar:root/ibmhscS1_0|9*9
406-
520*10007CA|IBMHSC_Partition,lpar:root/ibmhscS1_0|5*9
406-
520*10007CA|IBMHSC_Partition,lpar:root/ibmhscS1_0|1*9
406-
520*10007CA|IBMHSC_Partition,lpar:root/ibmhscS1_0|8*9
406-
520*10007CA|IBMHSC_Partition,lpar:root/ibmhscS1_0|4*9
406-520*10007CA|IBMHSC_Partition"
```

- `lsaccfg -t taskrole` shows a list of accessible (by this role) tasks separated by a '+' following the resource type they operate on.

You can create customized HMC roles by modifying predefined HMC roles. Creating customized HMC roles is useful for restricting or granting specific task privileges to a certain user. Through the GUI this can be accomplished when you add or copy a managed resource or task role. Through the restricted shell, you can use the `mkaccfg` or `chaccfg` commands. These tasks can only be performed by a user with an `hmcsuperadmin` role. Note that this does not mean a user with the just the `hmcsuperadmin` task role, but any user whose task role is a customized one with `hmcsuperadmin` as the parent task role with the appropriate GUI and CLI tasks grouped in.

Below are some examples of using the mkaccfg command:

```
# create a resource role for LPAR ID 9
# note the '\' escape character preceding a '|'
hscroot@myhost:~> mkaccfg -t resourcerole -i
"name=hscroot_group_orig,resources=lpar:root/ibmhscS1_
0\|9*9406-520*10007CA\|IBMHSC_Partition"

# create a resource role for a managed system
hscroot@myhost:~> mkaccfg -t resourcerole -i
"name=nonhscroot_group,resources=cec:root/ibmhscS1_0\|
9406-520*10007CA\|IBMHSC_ComputerSystem
```

```
# create a resource role for a managed system and LPAR
# resource type; note that the managed system includes
# the LPARs, so adding ALL_PARTITIONS is redundant
hscroot@myhost:~> mkaccfg -t resourcerole -i
"name=groupies,resources=cec:root/ibmhscS1_0\|9406-
520*10007CA\|IBMHSC_ComputerSystem,lpar:root/ibmhscS1_
0\|ALL_PARTITIONS*9406-520*10007CA\|IBMHSC_Partition"

# create a customized task role based on hmcoperator
# only the lshwres and lssyscfg CLI commands are
accessible
hscroot@myhost:~> mkaccfg -t taskrole -i
name=buttered_role,parent=hmcoperator,resources=cec:ls
hwres+lssyscfg
```

Note that in the above example creating a task role, configuration *data* was given. A configuration file could have been passed as an attribute instead. The configuration file containing the configuration information creates the access control role. The format for this file must be in comma separated value (CSV) format. A line feed (<LF>) marks the end of a record. There can only be one configuration record in the file.

For each particular HMC version, the defined set of tasks can not change. From *version to version*, the set of tasks may change - the newer version of HMC may add new tasks or remove existing tasks. An installation from Recovery media wipes out all existing role data stored on the HMC, so it's important that upgrade data be saved before performing the upgrade.

In an HMC upgrade installation (from Recovery media or an upgrade zip file), predefined role definitions will be updated as appropriate. Customized roles will remain unchanged. While it's unlikely that all tasks supported by a customized role will be removed in an upgrade, should that happen the user with this task role will only be able to login and nothing more. The same follows if all tasks assigned to a task role were removed from that role by an `hmcsuperadmin` user. Similarly, if all task roles assigned to a user are deleted by an `hmcsuperadmin` user, that user ID will remain, but will be made ineffective. The GUI will show no task role when viewing the user's profile, and the CLI will give output like the following:

```
hscroot@myhost:~> lshmcusr --filter names=noTaskRole
                           -Ftaskrole
Undefined
```

Thus, while it is unlikely to occur, it would be a good practice to document customized roles and to pay attention to the REAME files that come with updates to the HMC code. This should help prevent unexpected results.

When deleting resources, managed resource roles that include these resources are not modified. The reason is that these resources might not be permanently removed. For example, if a managed system was removed from an HMC's management domain, it could very well be added back. Hence, the corresponding managed system resource will not be deleted. It is the user responsible for the management of the resource role to decide if the removed resource should remain part of the role.

For a script that can be used to automate HMC user management, please refer to Appendix I.

6 Problem Determination

This section is intended to cover issues and problems that may be encountered during operation of the HMC itself. While some mention will be made of managed systems in the context of server and frame management, service applications needed to enable service environments, e.g. Electronic Service Agent™, Service Focal Point, and the Remote Support Facility will not be discussed.

6.1 Problem Analysis

The HMC is composed of many subsystems and layers to its code stack. Every subsystem maintains a dynamic trace – a "flight recorder" of sorts. As of HMC Version 5, Release 1.0, the vital traces are stored in persistent files in the HMC filesystem. At periodic intervals, typically every hour, ***cron*** jobs are run, depending on the process, for those subsystems or applications that tend to generate heavier quantities of trace data. This ***cron*** job checks the respective trace file's size to see if it has exceeded a fixed, static threshold. If the file does exceed this threshold, the trace file is backed up and a new trace file's generation begins. Also, when the HMC is rebooted, the last-running trace file for some subsystems will be backed up; for others, it will be overwritten.

Up to eight backup copies are maintained at any given time. While this may seem sufficient, the size of the trace files and the number of backups maintained are dependent upon the HMC load. For example, in an environment where an HMC is managing two frames and 40 logical partitions, the amount of trace data generated can be voluminous over a short period. If IBM Support and Service will be needed to diagnose an HMC problem, it is vital that the trace files be extracted from the HMC as soon as possible after the problem has been observed.

The HMC provides a command ***`pedbg`*** to assist in this process. This command can only be run as user hscpe with the `hmcpe` task role. The man page should be consulted for the full list of options this command provides. The most important features it provides are:

- `-d` For pre-V5R1.0 HMC releases, this command can be issued with an attribute value of `on` to enable trace logging. Note that this does not mean trace logging will begin! While it is enabled, to initiate it the HMC must be rebooted.
- `-l` For errors given pointing to failures in the RMC subsystem (ex. DLPAR), this attribute can be used to list the status of various RMC Resource Managers.

- `-c` Collect trace files and (optional) stack traces of various subsystems. This initiates an interactive task which will prompt for verification if trace for certain subsystems should be collected. All files will be collected into a zip file, and the option to store to the HMC filesystem or DVD will be given. If the former is chosen then root access will be needed to extract the zip file from the HMC (see below) or the `sendfile` HMC CLI can be used. This file can later be deleted from /dump using the –r attribute once the log file has been handled appropriately.

One case where ***pedbg*** will not be as helpful for trace logging is for the remote HMC Console application. On Windows^{TM} platforms from where this Console can be launched, the following must first be typed at a command prompt prior to launching the Console (assuming `pedbg` was already run on the HMC):

```
dir > c:\program files\websm\config\wsmdebug
wsm.bat 2>\temp\stderr.log 1>\temp\stdout.log
```

It is the presence of wsmdebug in the correct directory that enables the tracing. Note that for the second command, the PATH environment variable will be set up for all users during WebSM installation such that this command can be executed from any directory. AIX and Linux users of `wsm` can use the following command:

```
# touch /tmp/wsmdebug
```

While ***pedbg*** should suffice for enabling trace collection, situations can arise where it will be helpful for support to gain root-level access on the HMC. One example was given above: extracting the zip file from /dump. The HMC has a restricted shell which can be accessed via SSH or by selecting an rshterm from the pop-up menu on the console. (This menu is accessed by right clicking in the grey area of the screen.)

The HMC provides a command to escape temporarily from the restricted shell: ***pesh***. This command may only be run by user `hscpe` with task role `hmcpe`. This command accepts one parameter – the HMC serial number – which can be obtained by issuing the command ***lshmc –v*** (see the "SE" tag attribute). You will then be prompted for a password, which you must obtain from IBM Support. See the following:

http://publib.boulder.ibm.com/infocenter/iseries/v1r2s/en_US/info/ipha5/contacting_support.htm

This password is valid until the end of the calendar day for which it was issued. Hence, a password obtained at 11 pm will expire in one hour, i.e. midnight. Once accepted, the password will give full shell access. Note that while the password will expire at the end of the calendar day, once logged in as user root you may remain logged in indefinitely. This is not recommended because it can be a security exposure. It is recommended that this user ID be deleted after use and recreated, temporarily, as needed.

6.2 Problem Logging and Tracking

More detailed information about either information or error messages received during command execution can be obtained by using the ***showLog*** command. This command can be run in an xterm on the HMC console by a user who has become root. (In order to become root, you will need to follow the steps listed above for obtaining a password to run in conjunction with the hscpe task role and the ***pesh*** command.) This document in Information Center provides more details:

http://publib.boulder.ibm.com/infocenter/iseries/v1r2s/en_US/info/iphau/viewhmclogs.htm#viewhmclogs

Any user except those with the hmcviewer task role can view system event information included in the console logs on the GUI via the HMC Management -> HMC Configuration -> View Console Events task. This will display system events logged during HMC operation, enumerating HMC activity in response to user-initiated tasks, whether the command succeeded or failed. This will not display all entries in the HMC Console logs, but a subset of them. This task can also be executed on the command line by using the ‘`lssvcevents`’ command with ‘`-t console`’ flag.

6.3 Problem Correction

As mentioned previously, all HMC tasks - user-initiated and otherwise - require interactions between various subsystems on the HMC. Failures in one or more subsystems may occur, and it is very useful if failures can be isolated. For example, usually when a task fails it’s a good idea to try another way to perform the same operation. Assume a task cannot be performed from the GUI; it was initiated and the GUI “hung.” Typically, this means the panels are grayed out, especially after minimizing and maximizing. Check whether it can be done through the command line. If that can be performed successfully, most likely the GUI is the problem. If both ways do work, then the backend is likely to be the culprit.

With that being said, let’s consider some possible scenarios in HMC system management. A common source of curiosity is HMC performance. On the HMC

7310-C03 desktop model, or 7310-CR2 / 7310-CR3 rack-mount models, make sure the hyperthreading option under the Advanced options in BIOS is disabled. On these models, this can be verified on the command line as user `root` by typing:

```
    dmesg | grep -i hyper
```

There will be no output if hyper threading is disabled.

Performance can suffer if trace and log files fill the HMC file system. Disk space usage can be checked by typing (only as user `root`!) ***‘df –h’***. If any filesystem partition is in 100% use, your service provider should be consulted.

The managed systems and frames can be in one of many states, as reported by the HMC. A description of managed system states can be found at:

http://publib.boulder.ibm.com/infocenter/iseries/v1r2s/en_US/info/ipha1/poweronstates.htm,

A description of frame states can be found at:

http://publib.boulder.ibm.com/infocenter/iseries/v1r2s/en_US/info/ipha1/framepoweronstates.htm.

Among the states that cause confusion:

- No Connection: The HMC cannot build a valid connection to the FSP / BPC. The reason will be displayed as an error code under the Op Panel value column on the GUI. If you believe that this state has been reached in error, you can reset the network connection between the HMC and FSP. By right-clicking on the managed system to bring up the pop-up menu, or selecting the managed system and going to the ‘Selected’ menu, select ‘Reset or Remove’ Connection on the HMC Console. You must have task role `hmcsuperadmin`, `hmcoperator`, or `hmcpe` to perform this operation. You can also initiate this task from the command line: `rmsysconn` command with `–o reset` flag.

- Definitions for No Connection codes, as well as corrective measures that can be taken, can be found at

  http://publib.boulder.ibm.com/infocenter/iseries/v1r2s/en_US/info/ipha5/hmc_codes_0.htm

.

- Incomplete: The HMC is unable to gather all system information from the managed system or frame. In some cases this could be due to a network error causing a temporary disruption to HMC–FSP interactions, or managed system hardware configuration changes being performed from a redundant HMC. To verify, an attempt can be made to recover from this state by using the Rebuild Managed System GUI task, or the `chsysstate` command with `–o rebuild –r sys` flags. Neither can be performed by task role `hmcviewer`. If this does not change the state, try resetting the HMC–FSP connection (see above – No Connection), then try rebooting the HMC if resetting does not help. If the problem persists, gather the trace files and logs for support. More information can be found at:

  http://publib.boulder.ibm.com/infocenter/iseries/v1r2s/en_US/info/ipha5/managedsystemstates.htm#managedsystemstates__incompletestate

- Recovery: The save area of the FSP, where partition profile and some partition information is kept, could be corrupted, cleared, or out-of-sync with the cached copy the HMC maintains in its filesystem. First, it would be good to know whether the managed system has been updated recently; firmware updates could clear NVRAM. If no system update has been performed recently, you can perform the Recovery Partition Data task on the managed system from the GUI. You will be presented with two options: Restore profile data from HMC backup data, and Initialize the managed system. The former (`chsysstate –o recover –r sys` on the CLI – not as `hmcviewer`) will restore the save area with the cached copy on the HMC. The latter (`rstprofdata –l 4` on the CLI – only as `hmcsuperadmin` or `hmcoperator`) will **clear all partition configuration information**! Don't use this unless you're willing to rebuild the partition from scratch. If neither approach works, gather trace files and logs. For more information see:

  http://publib.boulder.ibm.com/infocenter/iseries/v1r2s/en_US/info/ipha5/managedsystemstates.htm#managedsystemstates__recoverystate

A problem more severe than No Connection is the situation when no systems or frames appear where they had appeared before. While there can be many reasons for this, one common scenario observed is when a managed system or frame is removed from the HMC. This could have been through the Remove Connection task or `rmsysconn` in the CLI. When this system is then added back into the HMC's management domain, the HMC (as DHCP server) will not redetect it. If you remove a managed system, and have reason to believe this HMC might again manage it in the future, '`mksysconn -o auto`' should be run to purge the HMC of its management history and allow it to once again provide IP addresses to the managed server.

Another observed problem has been the GUI isn’t reflecting the true managed system or frame status or configuration. The CLI can be used to see if it gives up-to-date information. If so, this means the GUI has either stopped receiving indication data, or no indications are being propagated to it. A Reload (F5) can refresh the GUI in this situation. If the command line is also incorrect, the HMC has stopped receiving event notifications from the managed system or frame. To recovery from this situation, perform the Rebuild Managed System task.

The FSP provides the HMC with the capability to set locks on the platform. The FSP does not interpret the locks, but rather leaves it up to the HMC functions to do that. These locks are used for synchronization of operations from one or two (dual) HMCs. In a dual HMC environment, situations can arise where both HMCs perform tasks against the same managed system and require the same lock.

When HMC 2 needs to perform a task that requires a lock that HMC 1 is currently holding, HMC 2 will wait and retry to acquire the lock. If after a few attempts it is unsuccessful, the operation will fail and the user will be notified accordingly. Often this blocked state can be mistaken for a “hang,” when that in fact is not the case. However, it is possible for HMC 1 to acquire a lock and fail to release it. Should this happen, HMC 2 can be used to disconnect HMC 1. When an HMC is disconnected, all locks owned by the HMC are reset. To due this, any `hmcsuperadmin` user can run the ‘Disconnect Another HMC’ GUI task on HMC 2 against HMC1. This can only be done from the graphical interface. There is no corresponding command line version of this task.

7 Maintaining Licensed Internal Code

Make sure you keep track of new releases, updates and emergency fixes to HMC code. You can do this in one of two ways:

- Sign up for the technical support subscription service to receive emails when updates become available on the web
- Monitor the web manually on a regular basis: http://www14.software.ibm.com/webapp/set2/sas/f/hmc/home.html

Read the web site carefully. Select the appropriate platform, POWER4 or p5, whichever is appropriate. There are many additional resources on the site, such as links to additional technical information, hints and tips, and the latest Command Line Specification, where you'll find new commands that may have been added.

You can order Recovery CDs or download packages that contain the files needed to burn your own Recovery CD. The files used to create CDs have an .iso file extension. The CDs created from these packages are bootable. You can download updates to HMC code as well as emergency fixes, and you can order CDs containing the updates and fixes. The CDs containing updates and fixes are NOT bootable.

7.1 Critical Console Data Backups

It's important to maintain a current Critical Console Date (CCD) backup to use in recovering the HMC after the loss of a disk drive. Whenever you go to a new version level of HMC code, or use a Recovery CD to update the HMC, you should create a new CCD backup immediately following the installation. If you update HMC code between releases using the Corrective Service files downloadable from the web, and then create new CCD backups after the update, you can use those CCD backups and the last-used Recovery CD to rebuild the HMC to the level in use when the disk drive was lost.

Another example where a CCD would be useful is when replacing an FSP or BPC on a POWER5 server. A fresh CCD backup should be made before starting the replacement in order to preserve the DHCP lease file on the HMC that lists the starting FSP and BPC IP addresses. If for some reason things don't work after replacing the FSP or BPC, this backup can be used to restore the original information so the customer can go back to the original FSP and BPC. If the replacement is successful, a new IP address will be assigned to the new component and the lease file will be updated. At this point, a new CCD backup should be created capturing that freshly updated DHCP lease file.

7.2 HMC Code Installation/Upgrade/Update

HMC systems leave manufacturing pre-installed with the most recent level of code. However, there may be time when re-installation is needed at the customer’s location.

Beginning with Version 5 Release 1.0, POWER5 HMC code can be installed using DVD media, or over the network, using a server that accepts PXE (**P**re-Boot E**x**ecution **E**nvironment) requests.

7.3 Install/Recovery

Installation is the simplest form of applying code on the HMC. It is used by manufacturing to install code prior to shipping the HMC. When the HMC is at the customer’s location, an install should only be needed for disaster recovery or when the customer desires to reload the HMC from scratch. In disaster recovery, a systems administrator can use an appropriate Recovery CD and the Critical Console Data backup to get the HMC back to the state it was in prior to the failure. An old Critical Console Data backup should not be used after upgrading to newer version or release of HMC. A new CCD backup should be created, as mentioned above.

7.4 Upgrade

It’s important to distinguish between updating and upgrading a system. The terms are not synonymous. To upgrade is to bring the system to a higher version or release of HMC code. When the HMC’s version number is incremented, such as going from Version 4 to Version 5, the upgrade method must be used in order to apply the new version of HMC code. Prior to an upgrade, the systems administrator should perform a Save Upgrade Data to preserve configuration information on the HMC, i.e., network settings, users’ data and partition profiles. This data is saved in a special location on the HMC’s hard disk that will not be erased during the upgrade process. When the upgrade process completes, the data will be restored to the HMC’s file system.

Only perform a Save Upgrade Data when you are upgrading an HMC. Do not use it when performing service work on the POWER5 server. If you are planning to service or replace an FSP or BPC on a p5 server, do a Critical Console Data backup first.

7.5 Update (Corrective Service)

In between HMC releases, or between upgrades, there will be times when interim fixes or cumulative service packs need to be applied. Interim fixes consist of security fixes or fixes that are considered critical to be released immediately to customers. Service packs are generally larger in contents. Both can be installed on the HMC by using the Install Corrective Service task under HMC Code Update, or by using the ***updhmc*** command on the HMC.

7.6 HMC Code Update and Service Strategy

HMC uses the VRM (Version, Release, and Maintenance) nomenclature to describe operating system releases. HMC version and release information can be viewed locally at an HMC console by bringing up the Help menu, then clicking on About Hardware Management Console. You can also see it, either locally or from a remote WebSM client, by going to Licensed Internal Code Maintenance -> HMC Code Update and reading the status section. From the command line, you can get the current code level by issuing the command ***lshmc –V***.

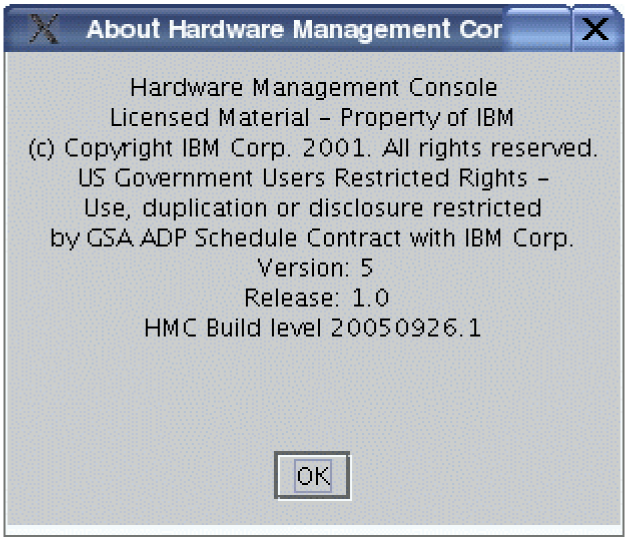

Hardware Management Console
Licensed Material - Property of IBM
(c) Copyright IBM Corp. 2001. All rights reserved.
US Government Users Restricted Rights -
Use, duplication or disclosure restricted
by GSA ADP Schedule Contract with IBM Corp.
Version: 5
Release: 1.0
HMC Build level 20050926.1

Above is a splash panel displayed by selecting the Help task at the HMC console.

Every version of HMC code is made available on bootable Recovery media. Within a version, releases are available as either Recovery CDs or downloadable Corrective Service files.

Corrective service is a cumulative maintenance release within a single version that customers can use to update the HMC from any previous releases within the same version. For example, a customer who is currently at Version 4 Release 1 or Version 4 Release 4 can update to Version 4 Release 5 using the same corrective service.

Corrective service is not provided whenever the version number is incremented, for example from Version 4 to Version 5. If this happens, only Recovery media may be used to perform the upgrade.

Corrective service is relatively easy to apply by using the Install Corrective Service task on the HMC console, either locally or remotely, or by running the ***updhmc*** command. As successive corrective service updates are installed, the size of the Critical Console Data increases. The Backup Critical Console Data task backs up both data and binary changes on the HMC. Over time, this can mean that the size of the CCD backup will be quite large. To shrink the size of the CCD, update with Recovery media after performing a Save Upgrade Data. The latter step only saves needed user and configuration data. After the update, a new Critical Console Data backup can be made, and it will be smaller.

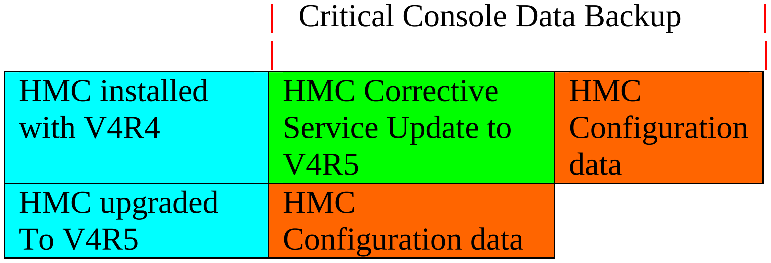

Critical Console Data Backup

HMC installed with V4R4

HMC Corrective Service Update to V4R5

HMC Configuration data

HMC upgraded To V4R5

HMC Configuration data

Corrective services zip files can be downloaded from the HMC support web site:

http://www14.software.ibm.com/webapp/set2/sas/f/hmc/home.html

It can also be ordered from IBM at that same web site.

If you are upgrading HMC machine type 7315 to support the POWER5 systems, contact your IBM Sales Representative or Business Partner, and order Hardware Feature Code 0961 for the HMC Recovery media.

Interim fixes or service packs are also treated as corrective service, and they are installed in the same manner as described above. The difference is primarily the size and how they are shown to users. Customers will see a PTF (Program

Temporary Fix) value associated with the interim fix or service pack when using a command such as ***lshmc –V***, for example:

MH00324: Maintenance Package for V4R5.0 (06-21-2005)

7.7 HMC Code Update on Multiple Machines

Some customers will have a large number of HMCs. Updating code on a large number of machines can be time consuming, especially if manual intervention or physical access to the local HMC is needed. Fortunately, there are methods to overcome this problem

7.7.1 Remote command

There is a rich set of commands on the HMC. These commands are available locally as well as remotely via OpenSSH. This allows a remote workstation installed with SSH client software to remotely execute commands on the HMC. The ***updhmc*** command previously mentioned is such a command that allows interim fixes, service packs and cumulative maintenance releases to be remotely installed on the HMC. The following example illustrates a scenario where an HMC code update is performed simultaneously on multiple machines from a remote workstation.

From the remote system installed with ssh, generate public key file with the *ssh-keygen* command using an empty passphrase and deploy the file to all the HMC. In this example, the HMCs host names are hmc1 through hmc7.

```
for i in 1 2 3 4 5 6 7
do
  scp hmc_update.zip hscroot@hmc$i:/home/hscroot
done
for i in 1 2 3 4 5 6 7
do
  ssh hscroot@hmc$i “updhmc –t l –f /home/hscroot/hmc_update.zip –c –r”
done
```

The first for loop in the code example above copies an interim fix whose file name is hmc_update.zip, to seven HMCs. The second for loop runs the ***updhmc*** command on each of the same seven HMCs. When the command finishes, it will remove the update file and reboot the HMC.

7.7.2 Network Installation/Update/Backup/Restore

Beginning with Version 5 Release 1.0, you can select the integrated network adapter in your HMC as a start-up or IPL device. This will allow the HMC to contact a remote system that supports PXE requests to perform installs, upgrades, backups or restore operations on the HMC. The remote system will need to have DHCP, TFTP and NFS servers running. To perform a secure backup/restore operation over the network, the remote system also needs to have an OpenSSH server running.

7.7.3 Server setup

The following describes the steps to set up a system to accept PXE boot requests from an HMC. In this example, the system is running SuSE 9.2.

- Make sure that tftp, dhcp and nfs servers are running on the system. Many gateways by default block dhcp and tftp access. To avoid these issues, it is recommended that the system and HMC be connected to the same switch, and that the switch permit dhcp and tftp.
- Create the directory /var/tftp. This directory is the default directory used by tftp daemon and is specified in the /etc/xinetd.d/tftp configuration file.
- Setup the /etc/dhcpd.conf with the **allow bootp** and **allow booting** options. An example of /etc/dhcpd.conf follows:

```
allow bootp;
allow booting;
ddns-update-style  none;
default-lease-time 14400;
max-lease-time     172800;
subnet 192.168.1.0 netmask 255.255.255.0 {
   range 192.168.1.100 192.168.1.200;
   option routers 192.168.1.1;
   option domain-name "mycompany.com";
   option domain-name-servers 192.168.1.1;
   filename "pxelinux.0";
}
```

In the above file, the system is setup to accept PXE requests and will serve IP address between the range of 192.168.1.100 and 192.168.1.200. The file pxelinux.0 will be the initial boot file loaded by the client, in this case the HMC.

This file is part of the *syslinux* package and should be copied from /usr/lib/syslinux/ to /var/tftp/ directory.

- Create the following directories /var/tftp/hmc and /var/tftp/pxelinux.cfg.
- Obtain the following files from HMC support web site or HMC Recovery media:

File Name	**Location on Recovery Media**	**Directory to copy to**
bzImage	Media 1: isolinux /	/var/tftp/hmc
initrd.gz	Media 1: images/	/var/tftp/hmc

- Create an NFS export the directory /home/hmc/images. If you want to use this directory to restore a backed up image of an HMC via NFS, you must export the directory for write access.

- Create a file *default* under /var/tftp/pxelinux.cfg/ directory. Below is a sample file:

```
default hmc

label hmc
        kernel hmc/bzImage
        append initrd=hmc/initrd.gz media=network server=192.168.1.1
        dir=/home/hmc/images mode=manual root=/dev/hda2 vga=0x317
```

In this example, we are telling the boot server to use the kernel file bzImage in the /var/tftp/hmc directory and start it with the parameters specified in the next line, beginning with the keyword **append**. The **server** argument should have the IP address of your server. The **dir** argument specifies where to get the image to install.

If the HMC has multiple network adapters, you will need to specify the interface that is used to initiate the PXE request. To do this, specify the argument xNIF or xNETWORK_INTERFACE, e.g:

```
        append initrd=hmc/initrd.gz media=network server=192.168.1.1
        dir=/home/hmc/images mode=manual  xNIF= eth1 vga=0x317
```

- Obtain the following files from IBM and copy them to /home/hmc/images. The following table shows the location of the file on the Recovery Media that comes with the HMC. The files also can be obtained from the support web site at http://www14.software.ibm.com/webapp/set2/sas/f/hmc/home.html .

File Name	**Location on Media**	**Directory to copy to**
disk1.img	Media 1: base/	/home/hmc/images/
disk2.img	Media 1: images/	/home/hmc/images/
disk3.img	Media 2: images/	/home/hmc/images/

7.8 Startup the HMC Using the Network

Once the server is setup, follow the steps below to start the HMC using the network adapter as the startup device. If the HMC is currently running, you must first shut down and restart it by using the command ***hmcshutdown*** or by exiting the console.

When the HMC restarts, and the following options will be displayed:

```
Press F1 for Setup
Press F12 for Startup Device
```

Press F12 key, and select the network adapter. This will be the integrated Ethernet adapter on your HMC.

If F12 option is not displayed:

- Press F1, and then specify the network interface as one of the startup devices via the BIOS Setup utility, as follows:

    1. From BIOS Setup, find and select Startup or Start Options, and then Startup Sequence to view the list of startup devices.
    2. Depending on the type of HMC (desktop or rack-mount), use the +/- key or arrow key to make the network interface an entry in the startup list, *after* the hard disk.
    3. If the HMC is a desktop, find and enable the Start-up Device Menu prompt. This action allows you to press the F12 key when the HMC powers on, and then select the network device in your startup list.
    4. For the CR2 or CR3 machine type, the Planar Ethernet PXE/DHCP entry should have Planar Ethernet 1 and Planar Ethernet 2. On desktop HMC machines, press the Esc key, and then select Devices -> Network Setup. Make sure that both PXE Boot Agent and PXE Base code are set to Enabled. Network installation does not occur if these two settings are not enabled.
    5. When you have finished, save the settings, then exit the BIOS setup utility to restart the boot process. At this point, you can press

F12 and select the network device as the startup device for this boot instance.

NOTE: In a future release of HMC, there will be enhancements made to the ***chhmc*** command to allow network boot settings to be changed prior to rebooting the HMC. This enhancement will only be for CR2 and CR3 HMC machine types. The following command would then enable the HMC to startup from network on the next boot.

```
chhmc –c netboot –s enable
```

The HMC will now broadcast a request to obtain an IP address from the DHCP server. It will then contact the TFTP server to proceed to load the two files **bzImage** and **initrd.gz** from the server. It will then start up.
Next the HMC will display a window that gives user a choice of tasks to perform: Backup, Upgrade, Restore and Install:

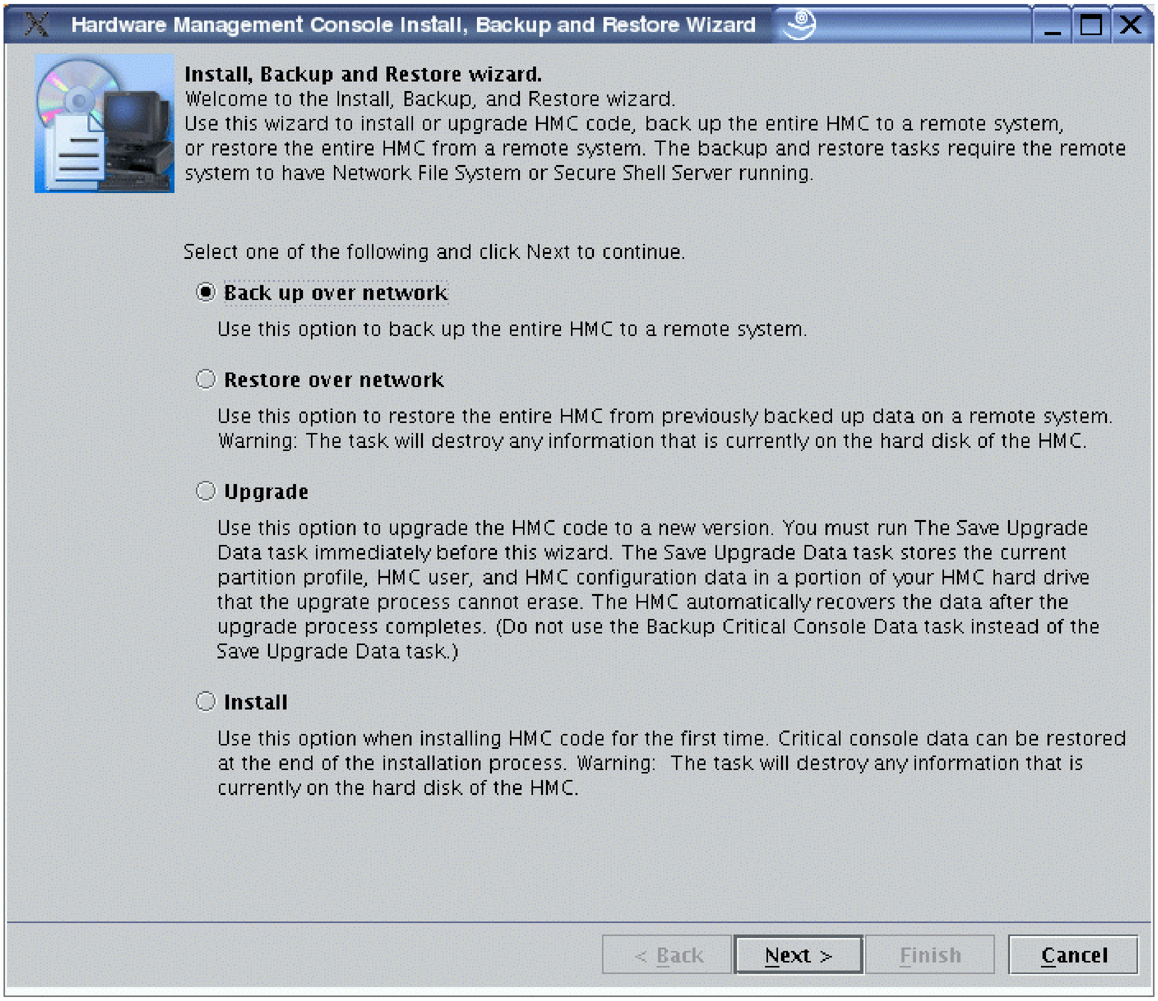

Hardware Management Console Install, Backup and Restore Wizard

Install, Backup and Restore wizard.
Welcome to the Install, Backup, and Restore wizard.
Use this wizard to install or upgrade HMC code, back up the entire HMC to a remote system, or restore the entire HMC from a remote system. The backup and restore tasks require the remote system to have Network File System or Secure Shell Server running.

Select one of the following and click Next to continue.

● **Back up over network**
Use this option to back up the entire HMC to a remote system.

○ **Restore over network**
Use this option to restore the entire HMC from previously backed up data on a remote system. Warning: The task will destroy any information that is currently on the hard disk of the HMC.

○ **Upgrade**
Use this option to upgrade the HMC code to a new version. You must run The Save Upgrade Data task immediately before this wizard. The Save Upgrade Data task stores the current partition profile, HMC user, and HMC configuration data in a portion of your HMC hard drive that the upgrate process cannot erase. The HMC automatically recovers the data after the upgrade process completes. (Do not use the Backup Critical Console Data task instead of the Save Upgrade Data task.)

○ **Install**
Use this option when installing HMC code for the first time. Critical console data can be restored at the end of the installation process. Warning: The task will destroy any information that is currently on the hard disk of the HMC.

- If you have more than one network card on your HMC, select Default if the install images reside on the same server to which the HMC sent PXE requests. If the install images reside on a different server, you can select another network adapter and configure it in order to obtain the install images. However, you must configure the second adapter using a different subnet.
- If a Restore operation is selected, the directory on the remote server must be accessible for writing.

7.9 Automating the Process

To perform these network tasks without having to be physically at the HMC console, you can follow the procedure below. (Note, you will still have to be physically at the HMC console to initially set the network interface as the startup device).

7.9.1 Server setup

There are two additional parameters that you will need to add to the append tag in the default PXE configuration file:

append initrd=hmc/initrd.gz media=network server=192.168.1.1
dir=/home/hmc/images mode=manual **mode=auto**
autocfg=/home/hmc/images/HmcInstall.cfg vga=0x317

The **mode** parameter needs to be set to auto to indicate an unattended mode is desired. The **autocfg** parameter specifies the configuration file that will specify other information. The file must be named HmcInstall.cfg and must be in the same directory as specified in the **dir** parameter.

Field Name	**Possible values**	**Description**
optype	*Install, Upgrade, Backup, Restore*	Operation to perform
media	*network, media*	Network access the image(s)
interface	*eth*x (where X is 0, 1, 2 or 3)	Network interface use to access image(s)
protocol	*dhcp, static*	Obtain IP address dynamically or statically
transtype	*nfs, ssh*	Non secure or secure transfer
host	IP address	IP address of server. Can be host name if DNS is

		specified, however the full hostname with domain should be specified.
xdir	Directory or File name	Directory or Full path of file to backup to or restore from.
restore	*yes* or *no*	This is optional. Only specify no if you are NOT restoring from a file that was backed up from the exact same machine.
mode	*auto* or *manual*	Indicate auto or manual mode.
ipaddr	IP address of the interface	Specify this ONLY if protocol is static.
gateway	IP address of gateway	Specify this ONLY if protocol is static.
dns	IP address of Dynamic Name Server	Specify this ONLY if protocol is static.
userid	User Name	Specify this ONLY if transtype is ssh. This user id MUST exist on the server specified by the host value.
passwd	Password	Specify this ONLY if transtype is ssh

HmcInstall.cfg file specification are shown in this table. Field names in bold are required. Values are indicated in italics, and they are case sensitive.

7.9.2 No PXE or TFTP or DHCP server

If no PXE, TFTP or DHCP server is available, there is still an option to startup the HMC from Recovery media and contact a remote server that has an NFS server running. To do this, insert the first Recovery media into the HMC DVD-RAM drive and reboot the HMC. By using the Recovery media, you don’t have to modify the startup sequence on the HMC to include a network interface.

7.9.3 Preparing the HMC

Once the server is set up, follow the steps below to prepare the HMC. If the HMC is currently running, then you must first shut down and restart it by using the command ***hmcshutdown*** or by exiting the console.

When the HMC restarts, the following options will be displayed:

```
Press F1 for Setup
Press F12 for Startup Device
```

Press F1 to go to Setup. Select Startup Sequence, and change the first startup device to be network. Save the settings and exit.

If you are not yet ready to start your network operation, you can reboot the HMC at this point using the hard disk as the start up device. To do this, you must override the current startup device (which is the network adapter that you have just selected), by pressing F12 key when powering on the HMC, and select the hard disk as current startup device. This will starts the HMC from hard disk. If you do not see F12 in the options, the Start up Device Menu prompt is disabled in BIOS. You will need to press F1, go into the setup menu, find and enable the Start up Device Menu prompt under Startup menu.

When you are ready to start the network operation, you can remotely login to the HMC via SSH and issue the command ***hmcshutdown –r –t now***. This will restart the HMC and use the network adapter on the HMC to send out PXE requests. After the HMC has started up, it will recognize that an unattended operation is desired, based on the **mode** and **autocfg** parameters in the default PXE configuration file.

Next it will obtain the HmcInstall.cfg file from the server and use it to proceed with the operation specified in the file. When the operation is completed, it will complete the booting process and come up to the HMC login panel. You will have to change the startup device list and move the network interface to the position after the hard disk later on.

7.9.4 Setup on Server when there are multiple HMCs

In an environment with multiple HMCs, each running at different code level or possessing different configuration, it is sometimes desirable to have a unique configuration on the server for each HMC to send PXE requests. For example, HMC A may choose to contact the server to do automatic backup to System X, while HMC B may choose to contact the same server with user intervention. In addition, HMC B has multiple Ethernet adapter cards, requiring it to specify the network interface with the **xNIF** parameter. In this scenario, the server can be setup to specify the PXE configuration files by using the MAC address or IP address of each HMC. If you wish to use the IP address as a file name, note that it is NOT the IP address of the HMC, but an IP address served by this server. This information can be obtained by looking up the DHCP lease file on the server.

For example, if HMC A is served IP address 192.168.1.21, then you will need to create a file C0A80115 under /tftpboot/pxelinux.cfg directory. C0A80115 is the hexadecimal value of the IP address without the dot. If you wish to use the MAC address instead, and the value is 00025557165A, then you will need to create a file 00-02-55-57-16-5a under /tftpboot/pxelinux.cfg directory. Below are the contents of two PXE configuration files:

File C0A80115 (HMC A):

```
default hmc

label hmc
  kernel hmc/bzImage
  append initrd=hmc/initrd.gz media=network server=192.168.1.1
dir=/home/hmc/images/SQ6 mode=auto
autocfg=/home/hmc/images/HmcInstall.cfg vga=0x317
```

File C0A80126 (HMC B):

```
default hmc

label hmc
  kernel hmc/bzImage
  append initrd=hmc/initrd.gz media=network server=192.168.1.1
dir=/home/hmc/images/mydir mode=manual xNIF=eth1 vga=0x317
```

8 Tips for Maintaining Licensed Internal Code

Make sure you keep track of new releases, updates and emergency fixes to HMC code. You can do this in one of two ways:

- Sign up for the technical support subscription service to receive emails when updates become available on the web
- Monitor the web manually on a regular basis at http://www14.software.ibm.com/webapp/set2/sas/f/hmc/home.html

8.1 Code and Resources on the Web

Read the web site carefully. Select the correct HMC version, whichever is appropriate. There are many resources on this site, such as:

- Links to additional technical information
- Hints and tips
- The latest Command Line Specification (command line reference)
- Order Recovery CDs/DVDs
- Download packages that contain the ISO files needed to burn your own Recovery media. (Note: this media is bootable.)
- Download updates to HMC code as well as emergency fixes, or order CDs containing the updates and fixes. (Note: this media is not bootable.)

8.2 Maintain Backups

Maintain a current Critical Console Data backup. If you use Recovery media to update your Licensed Internal Code to a new release level, make a new CCD backup after the upgrade process. A CCD backup created at V4R3.0 will **not** work on a system that was upgraded to V4R4.0 using a Recovery CD. However, if you are trying to recover after losing a disk on a system that was updated using the Corrective Service files, you *can* use a CCD backup (created at V4R4.0) with your V4R3.0 Recovery CD.

9 Tips for Maintaining System Firmware on the Managed System

The naming convention for system firmware updates is as follows:

01SFXXX_YYY_ZZZ

XXX is the release level
YYY is the service pack level
ZZZ is the last disruptive service pack level

Upgrades from one release level to another (xxx) are **always** disruptive, meaning you must re-IPL. Updates between service pack levels **may** be run concurrently, but you need to check.

You may need to upgrade existing HMC code to support a new server that is running the latest system firmware. As soon as you upgrade the HMC for the new server, you should plan on upgrading the existing managed servers to the new system firmware level. You should upgrade the HMC code before upgrading the system firmware on the existing servers or attaching new servers.

9.1 Understanding Update Packaging

New updates will be in a repository location. Supported repositories include the IBM service web site; IBM support system (RETAIN); CD or DVD media in the DVD drive on the HMC; or an FTP site. The HMC hard drive may be used as a repository location for updating other servers. The standard location (data source) for updates on an FTP server is /opt/ccfw/data, although any directory can be used. In other words, the customer can copy update files from the web to an accessible system within their network that can serve as a local FTP server.

It is important to keep in mind that a complete update consists of two files:

- a firmware code fix pack in RPM format, and
- a cover letter in XML format.

For example, the managed system firmware fix pack level 99 in code released in July 2005 consists of 01SF230_099_099.rpm and 01SF230_099_099.xml. If one of these files is missing, then that fix pack is *not* considered as a valid level by the LIC Update process on the HMC.

9.2 Concurrent versus Disruptive Updates

Starting with HMC V4R5, the behavior of the code update wizard on the HMC changed. The new behavior favors concurrent updates over those requiring a reboot. The wizard uses the name of the package to make the distinction. Concurrent-capable update filenames differ from disruptive update filenames in that the FFF and DDD fields are NOT equal. For example, a filename 01SF230_107_107.rpm is a disruptive code update package. An example of a concurrent-capable code update filename would be 01SF230_165_108.rpm, where the FFF field is 165 (the fix pack level) and the DDD field is 108 (the last disruptive code level).

When the fully concurrent-enabled HMC performs the code update wizard function, a filename of 01SF230_107_107.rpm and 01SF230_107_107.xml will cause the HMC to perform a concurrent install with deferred disruptive activation. This means that the system administrator will need to schedule a re-IPL at a convenient time in order to bring the new code into active status. It will be applied concurrently, but activated only after a reboot.

Since system release level 230, which came out in August 2005, power frame fix packs (which begin with 02BP in their file names) have been concurrent with respect to server operations. Although the Bulk Power Controllers reboot during the process, power to the frame remains up and no interruption is required.

You can perform a concurrent update if the disruptive base is already installed and active. For example, you can concurrently install a code update with the filename 01SF230_165_108.rpm if the target system is already activated at the 01SF230_YYY_108 code level, where "YYY" is a number equal to or greater than 108. This option does not apply to upgrades because upgrades are always disruptive.

Applying the “165” level files to a system running 01SF230_105_105.rpm would result in the need to perform a disruptive activate (requiring an outage) to get the “108” base onto the managed system simultaneously with the “165” concurrent update. This is why the file set is referred to as a concurrent-capable code update file set. The code update is only concurrent (no need to perform the re-IPL to have the code “online”) when its last disruptive level or another concurrent level using that base is already activated on the managed system.

File set names must follow naming rules as the HMC displays the fix pack level, which can be seen either on the “View system information” task menu or in the output of the **lslic** command entered in the restricted shell. Once a level value is assigned to a fix pack in the “FFF” field of the filename, it must not be re-issued to other fix packs within that release level.

Appendix I

The HMC does not have any Active Directory or LDAP support for automated user management. In large environments with many servers, there can be multiple HMCs (see Planning). It may not be desirable to manually manage user IDs across all HMCs, especially as users gain and lose access to the environment. As previously mentioned, the HMC CLI for user management is very useful for scripting, and hence can be used to address this situation. Using the CLI and SSH, we can create a script called `hmcusers` that can be used to add and remove users from a remote workstation as follows:

```
hmcusers: -o [ add | delete ] -h <list of hmc hosts>
          -l <login> -u <list of user names> -a <hmc
access>
   -o     The operation to perform, add or delete
users.
   -h     The list of HMC perform the operation.
   -l     The login to use on HMC to perform the
operation.
   -u     The list of user to create.
   -a     The access on HMC.
```

To prevent password prompting when using this script remotely, RSA / DSA keys can be deployed on HMCs to be managed, under a desired `hmcsuperadmin` ID. The `mkauthkeys` on the HMC can be used to do this; please see the section on Security for more details.

To add three users – deoli, tri, and bob -- with `hmcoperator` task roles on HMCs hmc1.austin.ibm.com and 9.53.188.188, using the login admin1, type

```
hmcusers -o add -h "hmc1.austin.ibm.com 9.53.188.188"
-l admin1 -u "deoli tri bob" -a "hmcoperator"
```

Note that admin1 would have to have hmcsuperuser authority to run this script.

To delete users salma and steve on two HMCs using login superadmin1 type

```
hmcusers -o delete -h "hmc2.austin.ibm.com
hmc3.austin.ibm.com" -l superadmin1 -u "salma steve"
```

Below is the code of such a script:

```
#!/bin/ksh
usage()
```

```
{
  echo "hmcusers: -o [ add | delete ] -h <list of hmc
hosts>"
  echo "          -l <login> -u <list of user names> -
a <hmc access>"
  echo "   -o     The operation to perform, add or
delete users."
  echo "   -h     The list of HMC perform the
operation."
  echo "   -l     The login to use on HMC to perform
the operation."
  echo "   -u     The list of user to create."
  echo "   -a     The access on HMC."
  exit 1
}

while getopts "o:h:l:u:a:" _arg; do
   case $_arg in
   o) op=$OPTARG
         ;;
   h) hl=$OPTARG
        ;;
   l) lid=$OPTARG
        ;;
   u) ul=$OPTARG
        ;;
   a) access=$OPTARG
        ;;
   *) usage
      ;;
   esac
done
if [ "$op" == "" ]; then
   echo "operation not specified"
   usage
fi
if [ "$hl" == "" ]; then
   echo "host list not specified"
   usage
fi
if [ "$lid" == "" ]; then
   echo "login id not specified"
   usage
fi
if [ "$ul" == "" ]; then
```

```
   echo "user names not specified"
   usage
fi

if [ "$op" == "add" ]; then
  # some validation here
  if [ "$access" == "" ]; then
     echo "missing HMC access role."
     usage
  fi
  for i in $hl
  do
    ping -c 1 $i >/dev/null 2>&1
    if [ $? -ne 0 ]; then
          echo "Unable to contact "$i
          echo "Skipping this host."
          continue
    fi
    for u in $ul
    do
       pass=$u$RANDOM
       ssh $lid@$i mkhmcusr -u $u -a $access -d "User
$u" --passwd $pass -M 90
       if [ $? -eq 0 ]; then
          echo "Created user "$u" with default
password "$pass" on "$i
       else
          echo "Failed to create user "$u" on "$i
       fi
    done
  done
fi
if [ "$op" == "delete" ]; then
  for i in $hl
  do
    ping -c 1 $i >/dev/null 2>&1
    if [ $? -ne 0 ]; then
          echo "Unable to contact "$i
          echo "Skipping this host."
          continue
    fi
    for u in $ul
    do
       ssh $lid@$i rmhmcusr -u $u
       echo "in delete"
       if [ $? -eq 0 ]; then
          echo "Deleted user "$u" on "$i
       else
          echo "Failed to delete user "$u" on "$i
       fi
    done
  done
fi
```

For more information about the tasks each HMC user role can perform and the commands associated with each task, see

http://publib.boulder.ibm.com/infocenter/iseries/v1r2s/en_US/info/ipha1/overviewoftasksandroles.htm.

IBM Corporation
Systems and Technology Group
Route 100
Somers, New York 10589

Produced in the United States of America
March 2006

This document was developed for products and/or services offered in the United States. IBM may not offer the products, features, or services discussed in this document in other countries.

The information may be subject to change without notice. Consult your local IBM business contact for information on the products, features and services available in your area.

All statements regarding IBM future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only.

IBM hardware products are manufactured from new parts, or new and used parts. In some cases, the hardware product may not be new and may have been previously installed. Regardless, our warranty terms apply.

This equipment is subject to FCC rules. It will comply with the appropriate FCC rules before final delivery to the buyer.

Information concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of the non-IBM products should be addressed with those suppliers.

When referring to storage capacity, 1TB equals total GB divided by 1000; accessible capacity may be less

Many of the features described in this document are operating system dependent and may not be available on Linux. For more information, please check: http://www.ibm.com/servers/eserver/pseries/linux/whitepapers/linux_pseries.html.

The IBM home page on the Internet can be found at: http://www.ibm.com.

The IBM System p5, @server p5 and pSeries home page on the Internet can be found at: http://www.ibm.com/servers/eserver/pseries.