

User Guide

Pigeon Point Shelf Manager

Release 3.7.2

September 19, 2022

nVent

Schroff GmbH

hardware.management@nVent.com

www.pigeonpoint.com

schroff.nVent.com



This document is furnished under license and may be used or copied only in accordance with the terms of such license. The content of this manual is furnished for informational use only, is subject to change without notice, and should not be construed as a commitment by nVent. nVent assumes no responsibility or liability for any errors or inaccuracies that may appear in this book.

Except as permitted by such license, no part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, manual, recording, or otherwise, without the prior written permission of nVent.

The Pigeon Point Shelf Manager uses an implementation of the MD5 Message-Digest algorithm that is derived from the RSA Data Security, Inc. MD5 Message-Digest algorithm.

All nVent marks and logos are owned or licensed by nVent Services GmbH or its affiliates. All other trademarks are the property of their respective owners. nVent reserves the right to change specifications without notice.

Table of contents

1	About this document	9
1.1	Shelf Manager documentation	9
1.1.1	Conventions used in this document	9
1.2	Additional resources	10
2	Introduction	11
2.1	In this section	11
2.2	Intelligent Platform Management: an ATCA overview	11
2.3	Pigeon Point Board Management Reference: hardware and firmware	14
2.4	Pigeon Point Entry Level Shelf Manager and ShMM	15
2.4.1	Pigeon Point Shelf Manager features	15
2.4.2	Support for dual redundant operation	16
2.4.3	System Manager Interface	18
2.4.4	Pigeon Point ShMM Shelf Management Mezzanines	19
3	Configuration	22
3.1	In this section	22
3.2	Setting up U-Boot	22
3.2.1	U-Boot environment variables	24
3.2.2	Assigning values to environment variables	34
3.2.3	Configuring U-Boot environment variables for the Shelf Manager	35
3.2.4	Establishing the secondary RC script	38
3.3	Setting up Shelf Manager configuration file	38
3.3.1	Carrier-specific configuration file	96
3.3.2	Obtaining configuration variables from Shelf FRU Information	96
3.3.2.1	Using the FRU Info Compiler to place configuration variables in the Shelf FRU Information	96
3.3.2.2	Using the Shelf Manager Command Line Interface to place configuration variables in the Shelf FRU Information	97
3.3.3	Verbosity level description	97
3.3.4	Verbosity console level description	98
3.4	Setting up Ethernet	98
3.4.1	Usage of the first Ethernet interface	98
3.4.1.1	Assigning an additional IP address to the first network interface	98
3.4.1.2	RMCP address propagation	99
3.4.2	Usage of the second Ethernet interface	99
3.4.2.1	Active-standby usage of the two network interfaces	100
3.4.2.2	Active-active usage of the two network interfaces	101
3.4.2.3	Support for site-dependent ShMC cross-connects	102
3.4.2.4	Bonded usage of the two network interfaces	103
3.4.3	Setting up the HPM.2 client channel	104
3.4.4	Gigabit Ethernet Support on ShMM-700	107
3.4.5	Using the ShMM-500 and ShMM-700 alternate software redundancy interface	107
3.4.6	Using the ShMM-1500 alternate software redundancy interface	110

3.4.7	Changing the default ShMM network parameters	110
3.4.8	Assigning VLAN IDs	119
3.4.9	Assigning IP addresses to the Shelf Manager via DHCP (version 4)	120
3.4.10	Using IPv6 for shelf-external access to the Shelf Manager (ShMM-700 only)	126
3.4.11	Assigning IPv6 addresses to the Shelf Manager via DHCPv6 (ShMM-700 only)	129
3.4.11.1	Handling DHCPv6 options	130
3.5	Configuring the FRU Information	132
3.5.1	Accessing the Shelf FRU Information	132
3.5.2	Setting up the Shelf FRU Information	134
3.5.3	Setting up the Shelf FRU Information using the CLI	137
3.5.4	Other FRU Information repositories	137
3.6	Configuring carrier and shelf attributes using HPDL	138
3.6.1	Compiling HPDL definitions	138
3.6.2	Compiling SDRs	139
3.6.3	Deploying HPDL data to the ShMM file system	139
3.6.4	Deploying HPDL Data to FRU Information areas	140
3.6.4.1	Specifying the location of the Carrier FRU Information	140
3.6.4.2	Using the FRU Info Compiler to place HPDL and SDR data in FRU Information	141
3.6.4.3	Using the Shelf Manager Command Line Interface	141
3.6.4.4	Placing HPDL data and SDRs in the FRU Information for specific FRUs	142
3.7	Configuring the cooling management strategy	143
3.7.1	Default cooling management strategy	143
3.7.2	Configuring a specific cooling management strategy	145
3.8	Configuring local sensors	145
3.9	Setting the Auxiliary Firmware Revision	151
3.10	Setting up the clock	151
3.10.1	Obtaining date and time from a time server	152
3.11	Setting up and using ShMM Power On Self Tests	153
3.12	Configuring external event handling	155
3.12.1	Detailed steps to configure external event handling	155
3.12.2	External event handler operation	157
3.13	Configuring the Platform Event Trap format	158
3.13.1	Parsed example of SNMP Trap	159
3.14	Configuring the IntegralHPI interface	161
3.14.1	HPI domain support in IntegralHPI	161
3.14.2	HPI SNMP subagent support	161
3.14.3	IntegralHPI client configuration	161
3.15	Configuring system services	162
3.16	Configuring the speed of I ² C buses (ShMM-700 only)	166
4	Using the Shelf Manager	168
4.1	In this section	168
4.2	ShMM login	168
4.3	Starting the Shelf Manager	168

4.4	Redundant operation	177
4.4.1	Initialization of the redundant Shelf Managers	179
4.4.2	Redundancy and CPLD State sensor	180
4.4.3	Reboot Reason sensor	183
4.5	Operation in shelves with radial IPMB-0	184
4.5.1	Operation in ShMM-500-based shelves with radial IPMB-0	184
4.5.2	Operation in ShMM-1500-based shelves with radial IPMB-0	186
4.5.3	Operation in ShMM-700-based shelves with radial IPMB-0	190
4.5.3.1	Upgrading the routing element SmartFusion cSoC(s)	191
4.6	Automatic SEL truncation	192
4.7	Cooling state sensors	192
4.8	Deadlock detection	194
4.9	I ² C Bus fault isolation	198
4.9.1	I ² C bus fault isolation – original approach (single multiplexer or switch, not HPDL-based)	199
4.9.2	I ² C bus fault isolation – HPDL-based approach (multiple multiplexers or switches)	200
4.10	Assignment of LAN configuration parameters to boards and modules	201
4.10.1	Structure and composition of the supported parameters	202
4.10.2	Obtaining LAN configuration parameters on the Shelf Manager level	202
4.10.2.1	Using the FRU Info Compiler to place LAN configuration data in the Shelf FRU Information	204
4.10.2.2	Using the Shelf Manager Command Line Interface to place LAN configuration parameters in the Shelf FRU Information	204
4.10.2.3	Configuring a Linux DHCP server for obtaining LAN configuration parameters	204
4.10.3	Dispatching LAN configuration parameters to boards	206
4.10.4	Synchronous assignment of LAN configuration parameters to boards	207
4.11	HPI System Event sensor	208
4.12	ShMM tests available via the diagnostic infrastructure	208
5	Using the IPMI analysis tools	212
5.1	In this section	212
5.2	IPMB trace collection daemon (ipmb_traced)	212
5.3	IPMI trace analyzer	216
5.3.1	Introduction to the Wireshark GUI	216
5.3.2	Introduction to terminal-oriented Wireshark	219
5.3.3	Installing IPMI analyzer software	220
5.4	Collecting and analyzing a trace using the GUI and command line tools	220
5.4.1	Collecting IPMI traces on IPMB-0 in unattended mode	220
5.4.2	Collecting IPMI traces on IPMB-0 in controlled mode	225
5.4.3	Collecting an IPMI trace in board trace mode	228
5.4.4	Collecting traces of IPMI traffic over the network	228
5.4.5	Analyzing a trace using the GUI	231
5.4.6	Analyzing a trace using command-line tools	233
6	Re-initializing the ShMM	236
6.1	In this section	236
6.2	Re-initializing the U-Boot environment	236

6.3	Re-initializing the file system	237
6.4	Resetting the login password	238
7	Re-programming the ShMM-500/1500	239
7.1	In this section	239
7.2	Firmware reliable upgrade procedure overview	239
7.3	Flash partitioning	240
7.4	The /var/upgrade file system	246
7.5	Reliable upgrade procedure status file	246
7.6	Reliable upgrade utility	246
7.7	Reliable upgrade utility use scenarios	253
7.8	Reliable upgrade examples	254
7.8.1	Example 1	254
7.8.2	Example 2	257
7.8.3	Example 32	260
8	Re-programming the ShMM-700	264
8.1	In this section	264
8.2	Firmware reliable upgrade procedure overview	264
8.3	Flash partitioning and image layout	266
8.4	Changing the 'user' flash partition format to JFFS2	268
8.5	Reliable upgrade procedure status file	271
8.6	Reliable upgrade utility	272
8.7	Reliable upgrade of SmartFusion cSoC image	279
8.8	Reliable upgrade utility use scenarios	280
8.9	Finding out component versions via the imls utility	281
8.10	Reliable upgrade examples	283
8.10.1	Example 1	283
8.10.2	Example 2	290
8.10.3	Example 3	294
8.10.4	Example 4	301
9	HPI-based Shelf Manager upgrades	308
Appendix A	Converting to ShMM-700-based Shelf Managers in ShMM-500 shelves	309
Appendix B	Customer support	311
Appendix C	Revision history	312
C.1	Release 2.1.0	312
C.2	Release 2.2.0	312
C.3	Release 2.3.0	312
C.4	Release 2.4.0	313
C.5	Release 2.4.1	313
C.6	Release 2.4.2	314
C.7	Release 2.4.4	314
C.8	Release 2.5.0	314

C.9	Release 2.5.2	314
C.10	Release 2.5.3	315
C.11	Release 2.6.0	315
C.12	Release 2.6.1	315
C.13	Release 2.6.4	316
C.14	Release 2.6.4.2	316
C.15	Release 2.6.4.4	316
C.16	Release 2.7.0	316
C.17	Release 2.7.1	317
C.18	Release 2.7.2	317
C.19	Release 2.7.3	317
C.20	Release 2.7.4	318
C.21	Release 2.8.0	318
C.22	Release 2.8.1	318
C.23	Release 2.8.2	318
C.24	Release 3.0.0	318
C.25	Release 3.1.1	319
C.26	Release 3.2.0	319
C.27	Release 3.3.0	320
C.28	Release 3.4.0	320
C.29	Release 3.5.0	321
C.30	Release 3.5.0.1	321
C.31	Release 3.5.1	321
C.32	Release 3.6.0	321
C.33	Release 3.6.1	321
C.34	Release 3.6.1.2	322
C.35	Release 3.6.1.5	322
C.36	Release 3.6.1.6	322
C.37	Release 3.7.0	322
C.38	Release 3.7.1	322
C.39	Release 3.7.1.3	322
C.40	Release 3.7.1.6	323
C.41	Release 3.7.1.8	323
C.42	Release 3.7.2	323

List of figures

Figure 1: Management aspects and potential Pigeon Point product sites in an example AdvancedTCA shelf.....	13
Figure 2: Pigeon Point Shelf Manager architecture and context	17
Figure 3: Implementation options for ShMC cross-connects	103
Figure 4: Dual star radial IPMB-0 topology	187
Figure 5: Redundant dual star radial IPMB-0 topology.....	188

Figure 6: Main window (IPMB trace analysis)	217
Figure 7: Main window (IPMI over network trace analysis)	219
Figure 8: Open Capture File window	221
Figure 9: GUI main window (unattended mode)	222
Figure 10: Capture Options window	226
Figure 11: Collected trace (controlled mode)	227
Figure 12: Capture Options window (analyzing networked IPMI traffic)	229
Figure 13: Collected trace (networked IPMI traffic)	230
Figure 14: Message 244 is selected.....	231
Figure 15: IPMI protocol layers (IPMB trace)	232
Figure 16: Filtered trace	233

List of tables

Table 1: Shelf Manager documentation.....	9
Table 2: Conventions used in this document.....	9
Table 3: ShMM-500R, ShMM-1500R and ShMM-700R features and variants	19
Table 4: ShMM models.....	20
Table 5: U-Boot environment variables	25
Table 6: Configuration parameter types and descriptions	38
Table 7: Shelf Manager configuration parameters	40
Table 8: U-Boot environment variables and descriptions	95
Table 9: Verbosity configuration parameters and levels	97
Table 10: HPM.3 DHCP Client ID	121
Table 11: HPM.3 DHCP Unique Identifier format	122
Table 12: Slot Location record for AdvancedTCA.....	122
Table 13: Positional parameters	130
Table 14: Option values format.....	130
Table 15: Shelf FRU Information and configuration variable settings	134
Table 16: Network services	163
Table 17: OS level bus number mapping for ShMM-700.....	167
Table 18: IPMB Topology record	189
Table 19: Cooling state sensors.....	193
Table 20: ShMM tests implemented in the diagnostic infrastructure	209
Table 21: Options for the IPMB trace collection daemon	213
Table 22: IPMB trace collection packet field.....	216
Table 23: IPMB message list basic information	217
Table 24: IPMI filter primitives.....	232
Table 25: Flash partitioning for 16MB Flash ShMM-500s.....	240
Table 26: Flash partitioning for 32MB Flash ShMM-500s.....	241

Table 27: Flash partitioning for 64MB Flash ShMM-500s	242
Table 28: Flash partitioning for 32MB Flash ShMM-1500s	244
Table 29: Flash partitioning for 64MB Flash ShMM-1500s	245
Table 30: Standard source file names for upgrade components	248
Table 31: Flash partitioning for ShMM-700s (64MB Flash) in UBIFS mode	266
Table 32: File names for firmware images stored as UBIFS files	267
Table 33: Flash partitioning for ShMM-700s (64MB Flash) in JFFS2 mode	269
Table 34: Reliable upgrade stages	271
Table 35: Protocols recognized in the URLs by the reliable upgrade utility	275
Table 36: Standard source file names for upgrade components	277

1 About this document

This document describes the Pigeon Point Shelf Manager. This PDF file requires Adobe Acrobat Reader 7.0 or later for full functionality. For a free reader, go to <http://www.adobe.com>.

1.1 Shelf Manager documentation

This document is one of two documents in the Shelf Manager documentation set. These documents are available in PDF file format. The complete set includes:

Table 1: Shelf Manager documentation

DOCUMENT	DESCRIPTION
<i>Pigeon Point Shelf Manager User Guide</i>	This document describes the overall configuration and use of the Pigeon Point Shelf Manager.
<i>Pigeon Point Shelf Manager External Interface Reference</i>	This document describes how to use the Shelf Manager command line interface, web interface, Simple Network Management Protocol (SNMP) interface and Remote Management Control Protocol (RMCP) interface.

1.1.1 Conventions used in this document

This table describes the textual conventions used in this document.

Table 2: Conventions used in this document

CONVENTION SAMPLE	DESCRIPTION
setenv	This 10 point bold Courier font is used for text entered at keyboard in example dialogues, which typically occur as one or more separate lines.
U-Boot 1.0.2 (Apr 18 2006 - 14:58:54)	This 10 point normal Courier font is used for ShMM output in example dialogues.
<i>addmisc</i>	This 12 point italic Courier New font is used for special text within normal paragraphs. The types of such special text include command names, file names, configuration parameters and command parameters, plus other text that could be entered by or displayed to a Shelf Manager user. This font is also used for command syntax definitions.

"Get Device ID"	IPMI commands defined by the IPMI specification or as PICMG extensions are shown in the normal font, surrounded by double quotes. This matches the corresponding convention used in PICMG specifications.
-----------------	---

1.2 Additional resources

For more information about Pigeon Point products, please visit the following website:

<http://www.pigeonpoint.com/products.html>.

2 Introduction

This section provides an overview of the Pigeon Point Shelf Manager and Shelf Management Mezzanine (or ShMM, currently the ShMM-500R, ShMM-1500R and ShMM-700R) products. The Pigeon Point Shelf Manager is a shelf-level management solution for AdvancedTCA® (ATCA®) products.

The Pigeon Point ShMM, when coupled with a corresponding carrier board, provides the necessary hardware to run the Shelf Manager within an ATCA shelf. This document focuses on aspects of the Shelf Manager and ShMM that are common to any ShMM carrier used in an AdvancedTCA context. Carrier-specific and shelf-specific details are documented by shelf providers.

The ShMM-500R complies with the Restriction of Hazardous Substances (RoHS) directive, but is equivalent to its predecessor, the ShMM-500, from a software point of view. All references to the ShMM-500 in this document apply to the ShMM-500R, unless otherwise noted. Both the ShMM-1500R and ShMM-700R were designed to be RoHS-compliant from the start and there are no ShMM-1500 or ShMM-700 products, though references to both product names may be simplified as ShMM-1500 or ShMM-700, respectively.

The ShMM-1500R is an additional ShMM variant that is based on a PowerPC processor, unlike the ShMM-500R, which is based on a MIPS-32 processor. The ShMM-700R is the newest ShMM variant; it is based on the Freescale i.MX28 (ARM9-based) processor. Both the ShMM-1500R and the ShMM-700R have different physical dimensions and a different approach for connecting to the ShMM carrier. Therefore, distinct ShMM carrier board designs are required for Shelf Manager solutions based on each of the ShMM-500R, ShMM-1500R or ShMM-700R ShMM variants. However, the Pigeon Point Shelf Manager provides the same functionality and high level interfaces for all three ShMM variants.

The Pigeon Point Shelf Manager is adaptable to manage CompactPCI platforms as well. This document focuses primarily on AdvancedTCA contexts, but provides CompactPCI-specific comments where appropriate.

2.1 In this section

This section contains the topics listed below. Just click on a topic to go to it.

- Intelligent Platform Management: an ATCA overview
- Pigeon Point Board Management Reference: hardware and firmware
- Pigeon Point Shelf Manager and ShMM

2.2 Intelligent Platform Management: an ATCA overview

The Pigeon Point products are the first Intelligent Platform Management building blocks designed from the ground up for modular platforms like AdvancedTCA, in which there is a strong focus on a dynamic population of Field Replaceable Units (FRUs) and maximum service availability.

The Intelligent Platform Management Interface (IPMI) specification provides a solid foundation for the management of such platforms, but requires significant extension to support them well.

IPMI defines a management infrastructure that is widely used across the PC and server industry. PICMG 3.0, the AdvancedTCA specification, defines the necessary extensions to IPMI. PICMG specifications are based on revision 1.5 of IPMI. However, the Pigeon Point Shelf Manager conforms to revision 2.0 of the IPMI specification and implements numerous IPMI 2.0 features, including: RMCP+, Virtual LAN support, and firmware firewall commands.

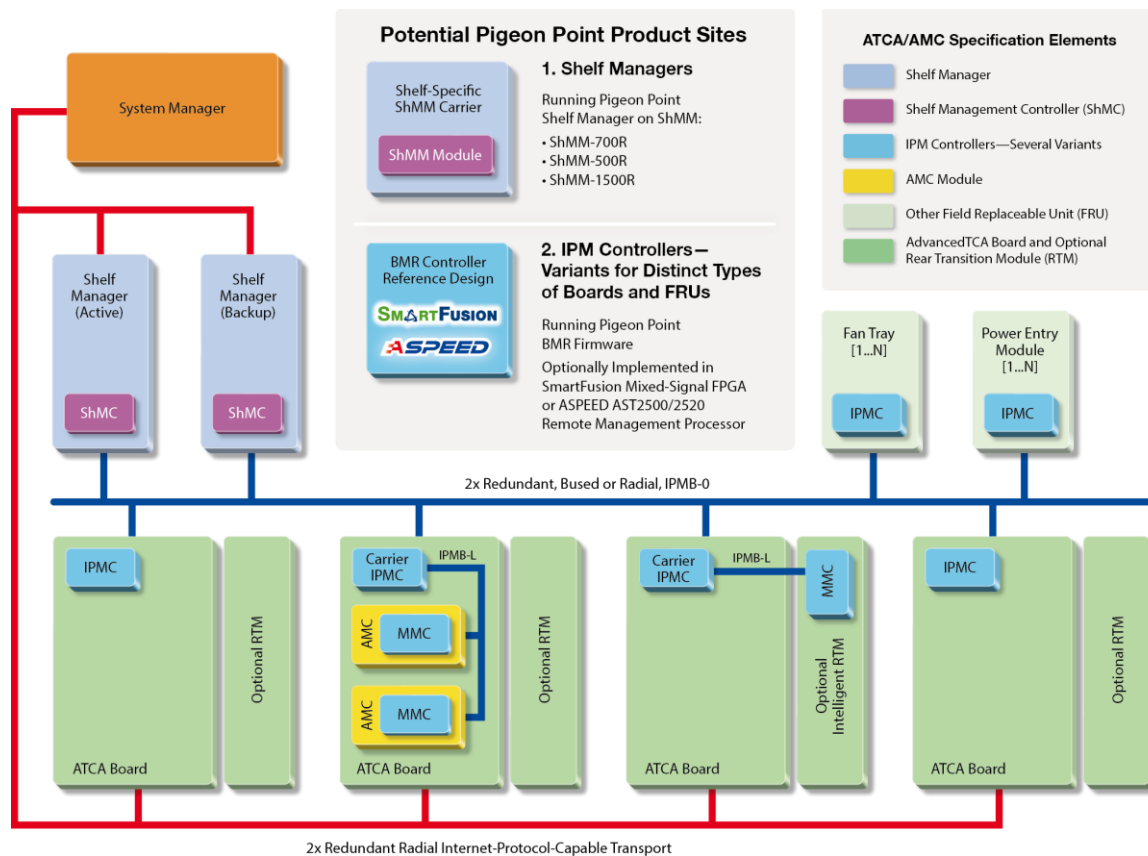
PICMG has significantly extended IPMI to cover the needs of open modular architectures. In fact, about 30% of the 656 pages of PICMG 3.0 are devoted to hardware platform management, including the definition of 38 new commands, ten new FRU Information data structures—several quite complex—and 3 new sensor types. There are hundreds of pages of additional management content in the AdvancedMC specification, the Hardware Platform Management (HPM.x) specifications and complementary subsidiary specifications.

The strategy for the Pigeon Point Shelf Manager is to fully support these extensions and also map them to other platform architectures such as CompactPCI.

Note: AdvancedTCA has adopted the term “shelf” for alignment with typical practice in telecommunications applications. Traditionally (for instance, in the CompactPCI specifications), the term “chassis” has been used with essentially the same meaning.

Figure 1 shows the logical elements of an example AdvancedTCA shelf, identified in terms of the ATCA specification, and potential sites for incorporation of Pigeon Point products.

Figure 1: Management aspects and potential Pigeon Point product sites in an example AdvancedTCA shelf



An AdvancedTCA Shelf Manager communicates inside the shelf with IPM Controllers, each of which is responsible for local management of one or more Field Replaceable Units (FRUs), such as boards, fan trays or power entry modules. Management communication within a shelf occurs primarily over the Intelligent Platform Management Bus (IPMB), which is implemented on a dual-redundant basis as IPMB-0 in AdvancedTCA.

The PICMG Advanced Mezzanine Card (AdvancedMC or AMC) specification, AMC.0, defines a hot-swappable mezzanine form factor designed to fit smoothly into the physical and management architecture of AdvancedTCA.

Figure 1 includes an AMC carrier with a Carrier IPMC and two installed AMC modules, each with a Module Management Controller (MMC). The figure also shows an instance of an MMC installed on a Rear Transition Module (RTM), which is then defined as an intelligent RTM. In both cases, on-carrier management communication between a Carrier IPMC and its MMCs occurs over IPMB-L (“L” for Local).

An overall System Manager (typically external to the shelf) can coordinate the activities of multiple shelves. A System Manager typically communicates with each Shelf Manager over Ethernet. The *Pigeon Point Shelf Manager External Interface Reference* document details the interfaces and protocols that are available for such interactions. Those interfaces include optional support for the SA Forum’s Hardware Platform Interface (HPI). There are two Pigeon Point HPI implementations,

one implemented internally and called IntegralHPI, and the other based on OpenHPI (www.openhpi.com). Both implementations are documented in the *Pigeon Point HPI User Guide*.

The next two sections address the board and shelf levels of management, highlighting the following Pigeon Point products and their capabilities as well as the relevant AdvancedTCA functionality:

- Pigeon Point Board Management Reference (BMR) firmware and corresponding hardware reference design, which together implement various types of management controllers.
- Pigeon Point Shelf Manager software and ShMM mezzanine module, which, together with an appropriate ShMM carrier board, implement an AdvancedTCA-compliant Shelf Manager and Shelf Management Controller (ShMC).

2.3 Pigeon Point Board Management Reference: hardware and firmware

This hardware and firmware level includes the local management of full-size 8U AdvancedTCA boards as well as other auxiliary FRUs, such as fan trays or power entry modules. Based on the interfaces specified by IPMI and extended by AdvancedTCA and AdvancedMC, any compliant Shelf Manager can work with any compliant IPM Controller and the FRUs that it represents, including AMCs.

This section focuses on Pigeon Point technology and specifically on controller solutions for AdvancedTCA and AdvancedMC.

The Pigeon Point BMR reference design can be implemented as part of any board or other FRU, and executes the corresponding firmware, thereby realizing a compliant IPM Controller. The BMR firmware represents one or more FRUs (via IPMB-0) to the Shelf Manager, including:

- Providing inventory information identifying each such FRU, including its manufacturer and other data.
- Describing and implementing a set of logical sensors (such as for temperature, state of IPMB-0, and operational state for each FRU (activated, deactivated, etc.)).
- Generating events (typically directed to the Shelf Manager) for exceptional conditions detected by any sensor, based on its configured event generation settings.
- Negotiating with the Shelf Manager for resources needed by the FRU(s), including power and interconnects.

BMR firmware running on a Carrier IPMC additionally represents its installed AMCs (or an intelligent RTM) to the Shelf Manager, including negotiating for power resources on their behalf. Furthermore, a Carrier IPMC negotiates with its modules and on-carrier switching resources regarding interconnect configurations.

The AMC management architecture is purposely designed to: 1) avoid impacting existing ATCA Shelf Managers and 2) minimize the resources required to implement a Module Management Controller, since board real estate and cost are at a premium on AMCs.

There are Pigeon Point BMR variants for AdvancedTCA IPM Controllers, as well as for Carrier IPMCs and MMCs. The principal BMR variants are based on the Microsemi SmartFusion intelligent mixed signal FPGA.

Pigeon Point Board Management Starter Kits for each of these BMR variants include all the materials necessary (documentation, schematics, bill of materials, firmware source code and development tools, etc.) for Intelligent FRU developers to integrate a reference design directly into their boards and take immediate advantage of the fully validated BMR firmware.

More details (including product briefs) on the available Pigeon Point BMR variants and corresponding Starter Kits are available at www.pigeonpoint.com/products.html.

2.4 Pigeon Point Shelf Manager and ShMM

The Pigeon Point Shelf Manager (consistent with AdvancedTCA Shelf Manager requirements) has two main responsibilities

- Manage/track the FRU population and common infrastructure of a shelf, especially the power, cooling and interconnect resources and their usage. Within the shelf, this management/tracking primarily occurs through interactions between the Shelf Manager and the IPM Controllers over IPMB-0.
- Enable the overall System Manager to join in that management/tracking through the System Manager Interface, which is typically implemented over Ethernet.

Much of the Pigeon Point Shelf Manager software is devoted to routine missions such as powering a shelf up or down and handling the arrival or departure of FRUs, including negotiating assignments of power and interconnect resources.

In addition, the Shelf Manager can take direct action when exceptions are raised in the shelf. For instance, in response to temperature exceptions the Shelf Manager can raise the fan levels or, if that step is not sufficient, even start powering down FRUs to reduce the heat load in the shelf.

2.4.1 Pigeon Point Shelf Manager features

The Pigeon Point Shelf Manager features are listed below:

- Executes on the ShMM, a compact mezzanine module, installed on a suitable carrier board for the shelf.
- Conforms to the AdvancedTCA specification.
- Monitors activities within the shelf via the ATCA-specified dual redundant Intelligent Platform Management Bus (IPMB).
- Accepts and logs events posted by any intelligent FRU in the shelf (reflecting exceptions in temperatures, voltages, etc.); posts alerts outside the shelf based on configurable IPMI Platform Event Filters.
- Supports hot swapping of Field Replaceable Units (FRUs), while maintaining full management visibility.
- Interfaces to standard “Telco Alarm” infrastructures, via ShMM carrier-implemented dry contact relays.
- Supports redundant Shelf Manager instances for high availability.
- Integrates a watchdog timer, which resets the ShMM if not periodically strobed; such resets automatically trigger a switchover to the backup ShMM, if configured.

- Includes battery-backed real-time clock for time-stamping events.
- Implements rich set of shelf-external interfaces accessible over Ethernet, including Remote Management Control Protocol (RMCP, required by AdvancedTCA), command line, web browser, Simple Network Management Protocol (SNMP).
- Includes a Shelf Adaptation Layer based on the Pigeon Point Hardware Platform Description Language (HPDL) that makes it straightforward to adapt the Shelf Manager for operation across a wide range of ATCA shelf architectures.
- Optionally supports an additional built-in shelf-external interface that complies with the Hardware Platform Interface (HPI), a set of application programming interfaces (APIs) for managing hardware platforms that is defined by the Service Availability Forum (www.saforum.org).

The Pigeon Point Shelf Manager can also be used in CompactPCI shelves.

2.4.2 Support for dual redundant operation

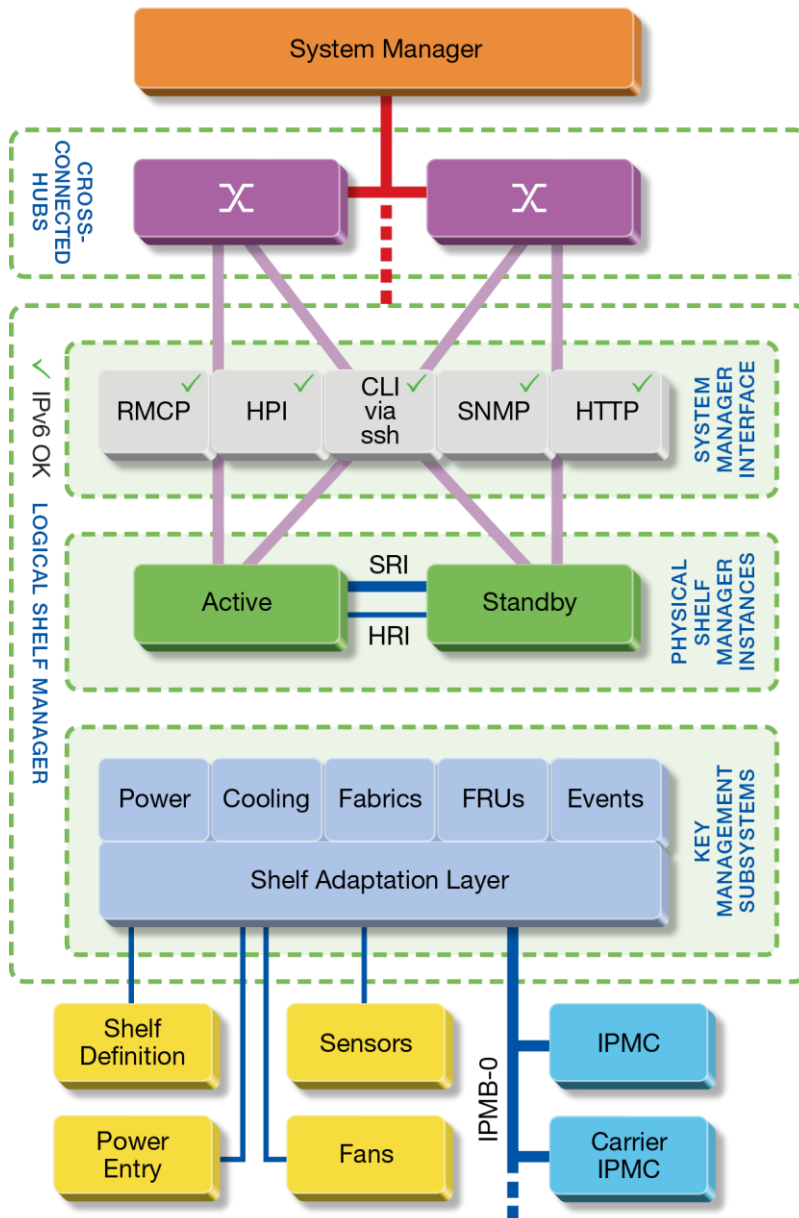
The Pigeon Point Shelf Manager can be configured with active/backup instances to maximize availability.

Figure 2 shows how both Shelf Manager instances are accessible to the System Manager, with primarily the active instance interacting at any given time. Only the active instance communicates over IPMB-0 with the IPM Controller population in the shelf.

The two Shelf Manager instances communicate over TCP/IP, with the active instance posting incremental software level state updates to the backup instance via a Software Redundancy Interface (SRI) implemented between them. As a result, the backup can quickly step into the active role if necessary. The SRI is a USB or a high speed UART-based serial link, depending on the ShMM variant.

Hardware level status updates are communicated between the Shelf Manager instances via the Hardware Redundancy Interface (HRI), which is also implemented differently in each of the three ShMM variants.

Figure 2: Pigeon Point Shelf Manager architecture and context



As shown in Figure 2, ShMC cross connects allow both ShMMs to be connected with both Base Interface Hubs. This improves system availability because either hubs or ShMMs can switchover independently, if necessary. ShMC cross connect support requires the use of both ShMM Ethernet links and is implemented in essentially all modern ShMM-based shelves that support Shelf Manager redundancy.

In the now-obsolete ShMM-300 family, one of the two ShMM Ethernet links was typically routed directly between the two ShMMs and used for the SRI. With this arrangement, ShMC cross connect support was not possible. A non-Ethernet link for SRI purposes was introduced with the ShMM-500R and carried forward to the ShMM-1500R and ShMM-700R. Given this history, a non-Ethernet SRI is often referenced elsewhere in this document as an Alternate Software Redundancy

Interface. This non-Ethernet SRI is implemented via a pair of cross-connected USB links on the ShMM-500R and ShMM-700R and with a high speed, UART-based link on the ShMM-1500R.

The HRI ensures that three key items of coordination status are available to each Shelf Manager instance:

- Presence: each Shelf Manager instance knows whether the other instance is present in the shelf.
- Health: each instance knows whether the other instance considers itself “healthy.”
- Switchover Request: the backup instance can force a switchover if necessary.

The HRI is implemented differently in each of the three ShMM families.

Figure 2 shows other aspects of the architecture and context of the Pigeon Point Shelf Manager, as well. These include the Shelf Adaptation Layer (and the shelf description that drives it) and the main Shelf Manager subsystems that implement the features summarized in section 2.4.1.

2.4.3 System Manager Interface

Another major subsystem of the Pigeon Point Shelf Manager implements the System Manager Interface. “System Manager” is a logical concept that may include software as well as human operators in the “swivel chairs” of an operations center. The Pigeon Point Shelf Manager provides a rich set of System Manager Interface options, which provide different mechanisms of access to similar kinds of information and control regarding a shelf. Figure 2 shows the System Manager Interface options supported by the Pigeon Point Shelf Manager.

One such mechanism is the IPMI LAN Interface. To maximize interoperability among independently implemented shelf products, this interface is required by the AdvancedTCA specification and supports IPMI messaging with the Shelf Manager via the IPMI Remote Management Control Protocol (RMCP).

A System Manager that uses RMCP to communicate with shelves should be able to interact with any ATCA-compliant Shelf Manager. This relatively low level interface provides essentially complete access to the IPMI aspects of a shelf, including the ability for the System Manager to issue IPMI commands to IPM Controllers in the shelf, using the Shelf Manager as a proxy.

The Pigeon Point Shelf Manager also supports Simple Network Management Protocol (SNMP) access to the shelf. This popular management protocol is supported with a custom Management Information Base (MIB) providing Get and Set access to a wide range of information and controls regarding the shelf.

In addition, the Pigeon Point Shelf Manager provides two interfaces oriented towards human users rather than programmatic ones:

- Command Line Interface (CLI): This interface provides a comprehensive set of textual commands that can be issued to the Shelf Manager via either a physical serial connection or a telnet connection.
- Web-based Interface: This interface enables essentially the same functionality as the CLI, with access to the Shelf Manager via a web browser.

Finally, the Pigeon Point Shelf Manager can optionally include an IntegralHPI subsystem, which provides access to the Shelf Manager via the Hardware Platform Interface (HPI). IntegralHPI

operates within the Shelf Manager, fully leveraging the Shelf Manager's facilities for managing the elements and events in the shelf. IntegralHPI also takes advantage of the mature redundancy framework of the Shelf Manager to deliver a fully redundant HPI service.

Using these mechanisms, the System Manager can access information about the current state of the shelf, including current FRU population, sensor values, threshold settings, recent events and overall shelf health.

These aspects of ATCA's System Manager Interface are considered to be part of the shelf-external interfaces. They are documented separately in the *Pigeon Point Shelf Manager External Interface Reference*.

2.4.4 Pigeon Point ShMM Shelf Management Mezzanines

The Pigeon Point Shelf Manager executes on the ShMM, a small Shelf Management Mezzanine with ShMM-500R, ShMM-1500R and ShMM-700R variants. ShMM-500R and ShMM-1500R mezzanines are available with: 1) 32 Mbytes of Flash and 64 Mbytes of SDRAM or 2) 64 Mbytes of Flash and 128 Mbytes of SDRAM. All ShMM-700Rs have 64 Mbytes of Flash and 128 Mbytes of SDRAM. The following table shows the main characteristics of the key ShMM variants.

Table 3: ShMM-500R, ShMM-1500R and ShMM-700R features and variants

FEATURE	SHMM-500R	SHMM-1500R	SHMM-700R
CPU	Broadcom Au1550	Freescale MPC8343	Freescale i.MX28
Processor core(s)	333 MHz MIPS-32	250 MHz PowerPC	297 MHz ARM9
SDRAM	64 or 128 Mbytes	128 Mbytes with ECC	128 Mbytes
Flash	32 or 64 Mbytes	32 or 64 Mbytes	64 Mbytes
Ethernet	Dual 10/100 Mbit	Dual 10/100 Mbit	Dual 10/100 Mbit
Serial	Two, one with modem controls	Two, one with modem controls	Three, one with modem controls and one with flow control.
Universal Serial Bus (USB)	Host and device ports	No	Host and device ports
PCI interface to carrier devices	No	Yes	No
Duplex IPMB-0	Yes	Yes	Yes
ATCA watchdog timer	Yes	Yes	Yes
Real-time clock, optionally battery backed on ShMM carrier	Yes	Yes	Yes

FEATURE	ShMM-500R	ShMM-1500R	ShMM-700R
General Purpose I/O signals	Nine	Nine	Ten
Shelf Manager hardware redundancy and hot swap interface, via on-board PLD, as indicated	Yes, via CPLD (Complex Programmable Logic Device)	Yes, via FPGA (Field Programmable Logic Array)	Yes, via FPGA (Field Programmable Logic Array)
High speed interface(s) to on-carrier devices	Multiple ports supporting either SPI or SMBus (with the latter used to implement IPMB-0, if so configured)	SPI	SPI
JTAG interface for processor debug and Flash programming	Yes	Yes	Yes
Physical dimensions	67.60mm X 50.80mm	92mm x 50.80mm	67.60mm x 41.75mm
Form factor definition	SO-DIMM-144 w/ proprietary pin assignments	Proprietary	DDR3 SO-DIMM-204 compatible w/ proprietary pin assignments

The following table provides a high level description of the full range of ShMM models, including some that are no longer shipping; for those in the latter category, the table shows the end of life (EOL) date.

Table 4: ShMM models

PART #	PROCESSOR SPEED	FLASH SIZE	RAM SIZE	COMMENTS
ShMM-300R (EOL June, 2007)	47MHz	16Mb	32Mb	Replaced with ShMM-500R
ShMM-500RE-333M16F32R (EOL July, 2016)	333MHz	16Mb	32Mb	Entry level variant of ShMM-500R; no support for ShMM redundancy, along with other differences
ShMM-500R-333M16F64R	333MHz	16Mb	64Mb	Replaced with ShMM-

(EOL August, 2008)				500R-333M32F64R
ShMM-500R-333M32F64R (EOL July, 2016)	333MHz	32Mb	64Mb	
ShMM-500R-333M64F128R (EOL July, 2016)	333MHz	64Mb	128Mb	
ShMM-1500R-250M64F128R	250MHz	64Mb	128Mb	Encryption code present
ShMM-1500R-250M64F128R-NE	250MHz	64Mb	128Mb	Encryption code removed
ShMM-700R	297MHz	64Mb	128Mb	

This edition of the *User Guide* focuses on the ShMM-500R, ShMM-1500R and ShMM-700R and uses the shorthand “ShMM” to refer to all three variants. A separate edition of this document covers the Entry Level Shelf Manager that is implemented on the ShMM-500RE.

The ShMM-1500R can be ordered with encryption code present or encryption code removed. The ShMM-500R and the ShMM-700R have an ECCN (United States Export Control Classification Number) of 5A992. The ShMM-1500R with or without the encryption code also has an ECCN of 5A992. For more information regarding these export classification topics, please contact nVent.

3 Configuration

The Shelf Manager application runs on a specialized distribution of Linux. For the ShMM-500 and ShMM-1500, the distribution is Monterey Linux. For the ShMM-700, that distribution is Pigeon Point Linux. The lowest layer in both cases is the firmware monitor, which is called U-Boot on all three ShMM variants.

3.1 In this section

This section contains the topics listed below. Just click on a topic to go to it.

- Setting up U-Boot
- Setting up Shelf Manager configuration file
- Setting up Ethernet
- Configuring the FRU Information
- Configuring carrier and shelf attributes using HPDL
- Configuring the cooling management strategy
- Configuring local sensors
- Setting the Auxiliary Firmware Revision
- Setting up the clock
- Setting up and using ShMM Power On Self Tests
- Configuring external event handling
- Configuring the Platform Event Trap format
- Configuring the IntegralHPI interface
- Configuring system services
- Configuring Speed of Master-Only I²C Buses (ShMM-700 only)

3.2 Setting up U-Boot

On a power-up/reset of the ShMM, the hardware starts executing the U-Boot firmware in Flash. The firmware performs basic initialization of the ShMM, and unless the user explicitly disables the Autoboot feature (thus forcing the firmware to switch to the maintenance user command interface), commences booting the Linux kernel. Linux is booted from the kernel and root file system images residing in Flash. U-Boot relocates the kernel image to RAM, sets up kernel parameters, and passes control to the kernel entry point.

For ShMM-500:

U-Boot 1.1.4 (Apr 27 2005 - 19:17:09)

CPU: Au1550 324 MHz, id: 0x02, rev: 0x00

Board: ShMM-500

S/N: 00 00 00 00 00 00 00 00 00 00 03 03 03

DRAM: 64 MB
Flash: 16 MB
In: serial
Out: serial
Err: serial
Net: AulX00 ETHERNET
Hit any key to stop autoboot: 0
shmm500

For ShMM-1500:

U-Boot 1.1.4 (Jun 15 2006 - 17:49:12) MPC83XX
Clock configuration:
Coherent System Bus: 99 MHz
Core: 249 MHz
Local Bus: 24 MHz
CPU: MPC83xx, Rev: 1.1 at 249.975 MHz
Board: ShMM-1500R
PCI1: 32 bit, 33 MHz
I2C: ready
DRAM: 128 MB
FLASH: 64 MB
PCI: Bus Dev VenId DevId Class Int
00 17 1172 0001 ff00 00
In: serial
Out: serial
Err: serial
FPGA: firmware version 1.12, carrier id 0
Net: TSEC0, TSEC1
Hit any key to stop autoboot: 0
shmm1500

For ShMM-700:

PowerPrep start initialize power...
Battery Voltage = 0.80V
No battery or bad battery detected!!!.Disabling battery voltage measurements.
Feb 24 201207:45:29
FRAC 0x92926152
memory type is DDR2
Wait for ddr ready 1
power 0x00820616
Frac 0x92926152
start change cpu freq
hbus 0x00000003
cpu 0x00010001

```

start test memory access
ddr2 0x40000000
finish simple test
U-Boot 2009.08 (Feb 15 2012 - 18:50:25)
Freescale i.MX28 family
CPU: 297 MHz
BUS: 99 MHz
EMI: 130 MHz
GPMI: 24 MHz
I2C: ready
DRAM: 128 MB
SFGEN: N25Q512A detected, total size 64 MB
A2F: SPICOMM protocol v1.6, M3 firmware v0.7, FPGA design v0.41.0.0
A2F: A2F firmware, version 0.7
A2F: Last reset cause: HARD
A2F: Device type: A2F060M3E-FG256
A2F: MSS clock frequency: 40 MHz
A2F: Fabric clock frequency: 20 MHz
A2F: Fast delay calibration: 1330 cycles per 100uS
A2F: eNVM: 128 KB (00000000 - 00020000)
A2F: eSRAM: 16 KB (20000000 - 20004000)
A2F: Extram start: 200022F0
RUPG: booting from image 1 (confirmed)
In: serial
Out: serial
Err: serial
Net: FEC0: 00:18:49:01:8f:26, FEC1: 00:18:49:01:8f:27
FEC0, FEC1
Hit any key to stop autoboot: 0
shmm700

```

Note: “shmm500”, “shmm1500” and “shmm700” are the U-Boot prompts on each ShMM variant for user commands. For a complete set of supported commands, type help. In this document, example dialogues that are applicable to all three ShMM variants, represent this prompt with “shmmxx00”. Example dialogues that are applicable to only one of the variants use the relevant prompt.

3.2.1 U-Boot environment variables

U-Boot includes a set of environment variables that should be configured prior to U-Boot use. The following table describes the default set of variables available. Some of these variables are set automatically by U-Boot and should not be modified by a user.

Table 5: U-Boot environment variables

ENVIRONMENT VARIABLE	DESCRIPTION
<code>a2f_upg_file</code>	<p>ShMM-700 only: This variable contains the name of the SmartFusion cSoC image file that U-Boot looks for on the boot volume. This variable is set automatically by U-Boot during bootstrap and must not be changed by a user; its value is used by U-Boot to pass information about the SmartFusion (A2F) cSoC upgrade image file to the Linux kernel. Only three values can be used: <code>a2f-upgrade.dat</code> (in the case of a normal A2F upgrade), <code>a2f-auto-rollback.dat</code> (in the case of an A2F upgrade rollback), or empty (if no A2F upgrade is in progress).</p> <p><i>Note: U-Boot expects the <code>a2f-auto-rollback.dat</code> file to be stored on the boot volume and to contain an image of the currently running firmware. The image is used to rollback an A2F upgrade if new candidate firmware is not programmed successfully or fails any of the post-upgrade tests.</i></p>
<code>addmisc</code>	<p>Appends <code>quiet</code> and <code>console</code> settings to <code>bootargs</code>. This variable is normally not modified.</p> <p><i>Note: not present on ShMM-700</i></p>
<code>alt_image</code>	<p>ShMM-700 only: This variable contains the rollback image set index. This variable is used by <code>load_ubifs</code> command to determine which of the Linux kernel and RFS images to boot if the current image set is broken and there is no reliable upgrade running. The variable is set automatically by U-Boot during bootstrap and should not be changed by the user; its value is always opposite to <code>image_sel</code>.</p>
<code>baudrate</code>	<p>Serial port baud rate, default is <code>115200</code>.</p>

ENVIRONMENT VARIABLE	DESCRIPTION
<i>bootargs</i>	<p>Command line to be passed to the Linux kernel. May contain references to other U-Boot environment variables, which is resolved at run-time. On both ShMM-500 and ShMM-1500, the default value is:</p> <pre>root=/dev/ram rw console=ttyS0,115200</pre> <p>On ShMM-700, the default value is:</p> <pre>console=ttySP0,115200n8 mtdparts=m25p80-1:512K(uboot0),512K(uboot1),256K(envvars0),256K(envvars1),256K(ubifs_cache),30976K(boot),32M(user) quiet fec_mac=\${ethaddr} root=/dev/ram</pre>
<i>bootargs_common</i>	ShMM-700 only: U-Boot command that sets up common kernel bootargs parameters. It's called by <i>bootargs_initrd</i> and <i>bootargs_nfs</i> commands to set up <i>console</i> , <i>mtdparts</i> , <i>quiet</i> , <i>fec_mac</i> kernel arguments.
<i>bootargs_initrd</i>	ShMM-700 only: U-Boot command that sets up kernel bootargs string so that the kernel shall be booted with RFS image stored in RAM. It executes <i>bootargs_common</i> command and then extends bootargs string with <i>root=/dev/ram</i> kernel parameter.
<i>bootargs_nfs</i>	ShMM-700 only: U-Boot command that sets up kernel bootargs string to prepare the system to be booted with NFS volume as RFS. It executes <i>bootargs_common</i> command and then extends bootargs string with the following NFS parameters: <pre>root=/dev/nfs ip=\${ipaddr}:::eth0:off nfsroot=\${serverip}:\${rootpath},v3,tcp.</pre> So, the kernel shall use <i>ipaddr</i> to configure <i>eth0</i> interface and access <i>serverip</i> host to use its exported <i>rootpath</i> share as root.
<i>bootcmd</i>	<p>U-Boot command executed to accomplish auto-booting.</p> <p>On both ShMM-500 and ShMM-1500, normally, this is something similar to <i>bootm BFB00000 BFC40000</i>, which starts the Linux image stored in Flash.</p> <p>On ShMM-700, the boot command is more complex because of the reliable A2F upgrade support and because boot images are stored in a file system, and, normally, this is something similar to:</p> <pre>run mount_ubifs; a2f upgrade; run ubifs</pre>
<i>bootdelay</i>	Autoboot delay value, in seconds. Default setting is 3.
<i>bootfile</i>	Parameter that specifies what kernel image should be used by the <i>net</i> and <i>nfs</i> boot options.

ENVIRONMENT VARIABLE	DESCRIPTION
<code>console</code>	Setting for the kernel and init script console port and baud rate. On ShMM-500 and ShMM-1500, the default is <code>ttys0,115200</code> . On ShMM-700, the default is <code>ttysp0,115200n8</code>
<code>corrupted_images</code>	This U-Boot-managed counter records the number of times that a corrupted boot image is discovered. When U-Boot fails to boot the kernel or the configured RFS image due to corruption of the boot image in Flash, U-Boot executes a recovery procedure that essentially designates the previously provisional (candidate) Flash device as the current persistent (confirmed) Flash device, increments this variable and proceeds with the boot process. The recovery procedure is not performed if the reliable upgrade WDT is active or if the corrupted boot image is not located in Flash. Support temporarily limited to ShMM-500/1500.
<code>ethact</code>	ShMM-700 only: Parameter that specifies which Ethernet controller should be used by U-Boot when handling <code>bootp</code> , <code>dhcp</code> , <code>nfs</code> , <code>ping</code> , <code>rarpboot</code> , <code>tftpboot</code> commands. Supported values are: <code>FEC0</code> , <code>FEC1</code> . <code>FEC0</code> selects <code>eth0</code> interface, <code>FEC1</code> selects <code>eth1</code> interface. Default value is <code>FEC0</code> .
<code>ethaddr</code>	MAC address of the primary on-chip Ethernet controller. The value of this variable is set automatically by U-Boot. This address is passed to the kernel Ethernet driver.
<code>eth1addr</code>	MAC address of the secondary Ethernet controller. The value of this variable is set automatically by U-Boot. On ShMM-500 and ShMM-1500, this address is passed to the kernel Ethernet driver. On ShMM-700, this address is not passed to the kernel Ethernet driver directly, but both the U-Boot and the Linux kernel FEC driver set up the MAC address for the second Ethernet controller as <code>ethaddr</code> incremented by 1, so both U-Boot and Linux use the same MAC for the second Ethernet controller.
<code>eth0_switch_bus</code>	This variable defines the logical I ² C bus that is connected to the GbE Switch for the ShMM-700 Ethernet 0 PHY interface. The default value is <code>0</code> , which corresponds to physical I ² C interface I2C4. <i>Note: this feature is supported starting in release 3.7.2.</i>
<code>eth0_switch_address</code>	This variable defines the 7-bit I ² C address of the GbE Switch for the ShMM-700 Ethernet 0 PHY interface. The default value is <code>0x5F</code> . <i>Note: this feature is supported starting in release 3.7.2.</i>

ENVIRONMENT VARIABLE	DESCRIPTION
<code>eth1_switch_bus</code>	This variable defines the logical I ² C bus that is connected to the GbE Switch for the ShMM-700 Ethernet 1 PHY interface. The default value is <code>5</code> , which corresponds to physical I ² C interface I2C3. <i>Note: this feature is supported starting in release 3.7.2.</i>
<code>eth1_switch_address</code>	This variable defines the 7-bit I ² C address of the GbE Switch for the ShMM-700 Ethernet 1 PHY interface. The default value is <code>0x5F</code> . <i>Note: this feature is supported starting in release 3.7.2.</i>
<code>flash_reset</code>	<p>Instructs Linux to erase the Flash filesystems (<code>/etc</code> and <code>/var</code>), restoring them from the contents of the current RFS image (<code>y/n</code>). The system startup script sets this variable back to <code>n</code> after the Flash erase. Default is <code>n</code>.</p> <p>On ShMM-700, this variable can also be set to <code>user</code> (<code>y/n/user</code>) to cause the system startup script to erase the entire 'user' partition (where <code>/etc</code> and <code>/var</code> are located). This mode can be used in case of unrecoverable 'user' partition corruption or together with the <code>user_partition</code> variable while converting the 'user' partition to the JFFS2 or UBIFS format.</p> <p><i>Note: this mode is supported starting in release 3.5.0.</i></p>
<code>gatewayip</code>	Default gateway IP address. This variable can be passed as a part of the kernel command line to automatically configure routing for the network interfaces. Default setting: <code>192.168.0.1</code> .
<code>hostname</code>	Network host name; default is <code>shmm500</code> for ShMM-500, <code>shmm1500</code> for ShMM-1500 and <code>shmm700</code> for ShMM-700.
<code>image_sel</code>	ShMM-700 only: This variable contains the current image set index. The variable is used in the <code>load_ubifs</code> command to determine which of the Linux kernel and RFS images to boot during bootstrap. The variable is set automatically by U-Boot during bootstrap.
<code>io_config</code>	ShMM-500 only: Determines if the Programmable Serial Controllers (PSCs – used for the IPMB-0 interface) on the ShMM-500 are configured for the dual-slave-address-configuration (<code>y/n</code>). Default setting: <code>y</code> . If setting is not <code>y</code> , the IPMB-0 interface does not work properly

ENVIRONMENT VARIABLE	DESCRIPTION
<i>ipaddr</i>	IP address used by the primary on-chip Ethernet interface. This variable configures the network interface specified by <i>ipdevice</i> automatically if the <i>rc_ifconfig</i> variable is set to <i>y</i> . Note that the system startup script sets the least significant bit of this variable to the least significant bit of the Hardware Address for the ShMM carrier; that is, if the Hardware Address is an even value, the last bit in the IP address is set to 0, otherwise it is set to 1. This is done in the startup script <i>/etc/netconfig</i> to support coordinated IP address configurations on redundant ShMMs. To disable this functionality, simply remove the <i>/etc/readhwaddr</i> file.
<i>ip1addr</i>	IP address used by the secondary Ethernet interface. This variable can be passed as a part of the kernel command line to automatically configure the corresponding kernel network interface.
<i>ip2addr</i>	IP address used by the third network interface, if present. This variable can be passed as a part of the kernel command line to automatically configure the corresponding kernel network interface. This variable is not assigned by default.
<i>ipdevice</i>	Device corresponding to <i>\$ipaddr</i> . Default value is <i>eth0</i> .
<i>ip1device</i>	Device corresponding to <i>\$ip1addr</i> . On ShMM-500 and ShMM-1500, the default is <i>eth1</i> . On ShMM-700, <i>usb0</i> is the default.
<i>ip2device</i>	Device corresponding to <i>\$ip2addr</i> . Not assigned by default, but can be assigned if more than two network interfaces are used on the ShMM (e.g. a USB- or serial-based network interface can be designated this way).
<i>kernel_start</i>	The absolute starting address of the kernel image in Flash. This variable is set automatically by U-Boot during bootstrap. <i>Note: on ShMM-700, this variable contains the absolute starting address of the kernel image in RAM. This variable is used in <code>ubifs, nfs, net, load ubifs, load ubifs alt</code> to specify the address in physical memory where the kernel image should be loaded and where it is expected to be during system boot. The default value is 0x42000000.</i>

ENVIRONMENT VARIABLE	DESCRIPTION
<i>load_ubifs</i>	ShMM-700 only: U-Boot command that loads current Linux kernel and RFS images from the Flash drive into RAM. Images are selected on the basis of the <i>image_sel</i> variable value. The kernel image is loaded at the <i>loadaddr</i> address, and the RFS image is loaded at the <i>rfsaddr</i> address. This command is invoked by <i>ubifs</i> , explicitly.
<i>load_ubifs_alt</i>	ShMM-700 only: U-Boot command that loads alternative Linux kernel and RFS images from the Flash drive into RAM. Images are selected on the basis of the <i>alt_image</i> variable value. The kernel image is loaded at the <i>loadaddr</i> address and the RFS image is loaded at the <i>rfsaddr</i> address. This command is invoked by <i>ubifs</i> command if one of the images from the current image set (selected by the <i>image_sel</i> value) is broken and there is no reliable upgrade running.
<i>log_max</i>	Specifies the size limit for the syslog file in bytes. Default is <i>250000</i> . When the log file size reaches the limit, the file is renamed as a backup file, and a new syslog file is started (replacing any existing backup), so that the maximum space used by the overall syslog is twice the <i>log_max</i> value.
<i>log_remote</i>	Specifies the IP address for the remote syslog facility. By default, this option is not set. The syslog daemon on Linux systems can be configured to receive messages from remote hosts, normally using the <i>-r</i> option, so using the <i>log_remote</i> setting it is possible to send the ShMM system log to a remote system.
<i>log_remote2</i>	Specifies the IP address for the secondary remote syslog facility. By default, this option is not set. The syslog daemon on Linux systems can be configured to receive messages from remote hosts, normally using the <i>-r</i> option. Note this option is only used when the <i>log_remote</i> option is set.
<i>logging</i>	Specifies if messages log file should be maintained in ram or Flash. Default is <i>ram</i> , which is the recommended option.

ENVIRONMENT VARIABLE	DESCRIPTION
<code>mac_override</code>	ShMM-700 only: This variable allows overriding MAC address configuration for both Ethernet controllers for the Linux kernel. This variable is used in <code>bootargs</code> generation and by default configures <code>ethaddr</code> to be passed as a configuration parameter. Default value is <code>set bootargs \${bootargs} fec_mac=\${ethaddr}</code> . If a different MAC address is specified and it is a valid MAC address, the FEC0 controller is configured with it, while FEC1 controller is configured with an incremented value. If no value is specified (i.e. <code>mac_override</code> is empty or doesn't set <code>fec_mac</code> value), the Linux kernel driver selects MAC addresses by its built-in algorithm. The MAC addresses generated by the kernel in this latter case may be the same as what the U-Boot reports via <code>ethaddr</code> and <code>eth1addr</code> variables, but this is not guaranteed.
<code>mount_ubifs</code>	ShMM-700 only: U-Boot command that invokes MTD configuration U-Boot scripts. This command is executed by U-Boot during bootstrap. Default value is <code>run set_mtdparts</code> .
<code>mtddids</code>	ShMM-700 only: This variable contains SPI interfaces configuration and is used privately by the U-Boot to configure the SPI subsystem to access Flash drives. Default value is <code>spi0=SPI-0</code> .
<code>mtdpart1</code>	ShMM-700 only: This variable contains the Flash drive partition map. The variable is used by the U-Boot and its value is also passed to the Linux kernel as part of <code>bootargs</code> string. The variable specifies how the Flash drive is split into partitions and consists of pairs of partition sizes and their names. Default value is <code>512K(uboot0),512K(uboot1),256K(envvars0),256K(envvars1),256K(ubifs_cache),30976K(boot),32M(user)</code> . See the ShMM-700R Hardware Architecture Specification for more information.
<code>net</code>	This variable can be used as a replacement for <code>bootcmd</code> as a means of booting a kernel and RFS image from TFTP. Use <code>run net</code> .
<code>netmask</code>	Network netmask, default value is <code>255.255.255.0</code>

ENVIRONMENT VARIABLE	DESCRIPTION
<code>nfs</code>	<p>This variable can be used as a replacement for <code>bootcmd</code> as a means of booting and running with an NFS mounted root filesystem. The <i>Monterey Linux User Guide</i> sample NFS project provides details for ShMM-500 and ShMM-1500 users.</p> <p><i>Note: the kernel image is fetched from a TFTP server by default.</i></p>
<code>password_reset</code>	Instructs Linux to restore the password from the current RFS image. The default password for user <code>root</code> in the standard provided RFS images is the empty password <code>""</code> . Default value for <code>password_reset</code> is <code>n</code> .
<code>post_normal</code>	Determines the list of POST tests that are executed on each boot-up. If not set, compile-time default settings are used. The test names listed in a value of this variable are separated by space characters.
<code>post_poweron</code>	Determines the list of POST tests that are executed after power-on reset only (vs. on each boot-up). If not set, compile-time default settings are used. The test names listed in a value of this variable are separated by space characters.
<code>serial#</code>	This variable contains the ShMM serial number assigned during manufacturing. The variable is set automatically by U-Boot during bootstrap.
<code>rootpath</code>	ShMM-700 only: This variable specifies the NFS root to be used by the Linux kernel when booting the ShMM via the <code>bootcmd_nfs</code> command. Default value is <code>/rootfs</code> .
<code>quiet</code>	Instructs the kernel upon bootup not to print progress messages to the serial console. Default is <code>quiet</code> .
<code>ramargs</code>	<p>Sets the kernel command line in the <code>bootargs</code> variable as appropriate for the root filesystem to be mounted from a ramdisk.</p> <p><i>Note: not present on ShMM-700.</i></p>
<code>ramdisk</code>	Specifies what <code>.rfs</code> image should be used by the <code>net</code> boot option.

ENVIRONMENT VARIABLE	DESCRIPTION
<code>ramsize</code>	Size of the system memory, in bytes. Default setting: calculated from the SDRAM configuration encoding in the build-time configuration block. <i>Note: not present on ShMM-700.</i>
<code>rc_ifconfig</code>	Allows the <code>/etc/rc</code> script to set up the IP address instead of <code>shelfman</code> . Default is <code>n</code> (allow <code>shelfman</code> to set up IP addresses).
<code>rc2</code>	Specifies secondary RC script that is to be invoked. This is the carrier-specific startup script. Default is <code>/etc/rc.shmm700-hpdl</code> on ShMM-700 and <code>/etc/rc.carrier3</code> on ShMM-500 and ShMM-1500 or other appropriate script for the relevant target platform. This variable must be set to a carrier-specific value matching the carrier on which the ShMM is installed.
<code>rfs_start</code>	The absolute starting address of the root filesystem image in Flash. This variable is set automatically by U-Boot during bootstrap. <i>Note: on ShMM-700, this variable contains the absolute starting address of the RFS image in RAM. This variable is used in the <code>ubifs</code>, <code>net</code>, <code>load_ubifs</code>, and <code>load_ubifs_alt</code> commands to specify the address in physical memory where the RFS image should be loaded and where it is expected to be during ShMM boot. The default value is 0x46000000.</i>
<code>rmcpaddr</code>	Default IP address for the RMCP service.
<code>serverip</code>	IP address of the TFTP server
<code>set_mtdparts</code>	ShMM-700 only: U-Boot command that is invoked to set up the MTD configuration in U-Boot to perform boot volume access. The command is invoked by <code>mount_ubifs</code> during bootstrap.
<code>start_rc2_daemons</code>	Instructs the secondary startup script to start or not start the <code>snmpd/boa</code> and <code>shelfman</code> daemons after bootup. Default is <code>y</code> .
<code>time_proto</code>	Protocol used to retrieve time from a network time server; possible values are <code>ntp</code> and <code>rdate</code> .
<code>time_server</code>	Time server for synchronization at runtime. If this variable is not specified, time is extracted from the hardware clock at system startup.

ENVIRONMENT VARIABLE	DESCRIPTION
<i>timezone</i>	Local time zone in CCCn format where n is the offset from GMT and optionally negative, while CCC identifies the time zone. The default is <i>UTC</i> .
<i>ubifs</i>	ShMM-700 only: U-Boot command executed to perform system boot from the Flash drive. The command performs setting <i>bootargs</i> variable, loading both kernel and RFS images into RAM and booting the system. If at least one of the images is broken and there is no reliable upgrade in progress, the command also outputs an error message and performs a single attempt to boot the ShMM using the alternative (rollback) kernel and RFS images.
<i>user_partition</i>	<p>ShMM-700 only: Specifies the format of the ‘user’ partition, which can be either <i>jffs2</i> or <i>ubifs</i>, where <i>ubifs</i> is the default value. This variable value must only be changed along with the <i>flash_reset=user</i> setting to change the ‘user’ partition format.</p> <p><i>Note: such partition format changes are supported starting with release 3.5.0.</i></p>
<i>ver</i>	<p>The U-Boot version string. The value is taken from the U-Boot image. The version string contains the base U-Boot version and the build timestamp, as well. For example, on the ShMM-700, the variable contains something like:</p> <p><i>U-Boot 2009.08 (Feb 24 2012 - 07:44:12)</i></p>

3.2.2 Assigning values to environment variables

To assign a value to an environment variable, on either ShMM variant, use the format:

```
shmmxx00 setenv <variable_name> <new_value>
```

For example:

```
shmmxx00 setenv bootdelay 1
```

Once all of the environment variables have been properly set, you need to save them back out to the Flash so that they remain after the ShMM is powered down. The *saveenv* command is used for this purpose.

```
shmmxx00 saveenv
```

The *setenv* functionality is also available as a Linux utility with the same usage. To display U-Boot variables at the shell prompt, use the additional *getenv* utility or issue the *setenv* command without parameters.

3.2.3 Configuring U-Boot environment variables for the Shelf Manager

When U-Boot is started for the first time, the following default environment variables are defined.

For ShMM-500:

```
addip=setenv bootargs $(bootargs)
ip=$(ipaddr):$(serverip):$(gatewayip):$(netmask):$(hostname):$(ipdevice)
addmisc=setenv bootargs $(bootargs) $(quiet) console=$(console),$(baudrate)
bootargs=root=/dev/ram rw console=ttyS0,115200
bootcmd=run ramargs addmisc; bootm $(kernel_start) $(rfs_start)
bootdelay=3
bootfile=sentry.shmm500.kernel
baudrate=115200
console=ttyS0
ethaddr= 00:00:1a:18:xx:yy
eth1addr= 00:00:1a:18:xx:zz
netmask=255.255.0.0
hostname=shmm500
gatewayip=192.168.0.1
ipdevice=eth0
ipladdr=192.168.1.2
ipldevice=eth1
rc2=/etc/rc.carrier3
ipaddr=192.168.0.22
start_rc2_daemons=y
flash_reset=n
password_reset=n
logging=ram
net= tftpboot 80400000 $(bootfile); tftpboot 81200000 $(ramdisk); run ramargs
addmisc; bootm 80400000 81200000
nfs=tftpboot 80800000 $(bootfile); run nfsargs addip addmisc; bootm
nfsargs=setenv bootargs root=/dev/nfs rw nfsroot=$(serverip):$(rootpath)
quiet=quiet
rc_ifconfig=n
ramargs=setenv bootargs root=/dev/ram rw
ramdisk=sentry.shmm500.rfs
rootpath=/rootfs
rmcpaddr=192.168.0.2
serverip=192.168.0.7
timezone=UTC
```

For ShMM-1500:

```
addip=setenv bootargs $(bootargs)
ip=$(ipaddr):$(serverip):$(gatewayip):$(netmask):$(hostname):$(ipdevice)
addmisc=setenv bootargs $(bootargs) $(quiet) console=$(console),$(baudrate)
bootargs=root=/dev/ram rw console=ttyS0,115200
```

```
bootcmd=run ramargs addmisc; bootm $(kernel_start) $(rfs_start)
bootdelay=3
bootfile=sentry.shmm1500.kernel
baudrate=115200
console=ttyS0
ethaddr= 00:50:c2:3f:xx:yy
eth1addr= 00:50:c2:3f:xx:zz
netmask=255.255.0.0
hostname=shmm1500
gatewayip=192.168.0.1
ipdevice=eth0
ipladdr=192.168.1.2
ipldevice=eth1
rc2=/etc/rc.carrier3
ipaddr=192.168.0.22
start_rc2_daemons=y
flash_reset=n
password_reset=n
logging=ram
net= tftpboot 400000 $(bootfile); tftpboot 1200000 $(ramdisk); run ramargs
addmisc; bootm 400000 1200000
nfs=tftpboot 400000 $(bootfile); run nfsargs addip addmisc; bootm
nfsargs=setenv bootargs root=/dev/nfs rw nfsroot=$(serverip):$(rootpath)
quiet=quiet
rc_ifconfig=n
ramargs=setenv bootargs root=/dev/ram rw
ramdisk=sentry.shmm1500.rfs
rootpath=/rootfs
rmcpaddr=192.168.0.2
serverip=192.168.0.7
timezone=UTC
```

For ShMM-700:

```
bootargs=console=${console}
bootcmd=run mount_ubifs; a2f upgrade; run ubifs
bootdelay=3
baudrate=115200
netmask=255.255.255.0
bootfile="sentry.shmm700.kernel"
loadaddr=0x42000000
console=ttySP0,115200n8
rfsaddr=0x46000000
mtdids=spi0=SPI-0
```

```

mtdparts1=512K(uboot0),512K(uboot1),256K(envvars0),256K(envvars1),256K(ubifs_cac
he),30976K(boot),32M(user)
bootfile=sentry.shmm700.kernel
ramdisk=sentry.shmm700.rfs
quiet=quiet
mac_override=set bootargs ${bootargs} fec_mac=${ethaddr}
rootpath=/rootfs
ipdevice=eth0
ipldevice=usb0
ip2device=eth1
ipaddr=192.168.0.22
ipladdr=192.168.1.2
gatewayip=192.168.0.1
rmcpaddr=192.168.0.2
netmask=255.255.255.0
hostname=shmm700
flash_reset=n
password_reset=n
logging=ram
timezone=UTC
rc_ifconfig=n
start_rc2_daemons=y
rc2=/etc/rc.shmm700-hpdl
set_mtdparts=set mtdparts mtdparts=SPI-0:${mtdparts1}
bootargs_common=set bootargs console=${console} mtdparts=m25p80-1:${mtdparts1}
${quiet}; run ${mac_override}
bootargs_initrd=run bootargs_common; set bootargs ${bootargs} root=/dev/ram
bootargs_nfs=run bootargs_common; set bootargs ${bootargs} root=/dev/nfs
ip=${ipaddr}:::eth0:off nfsroot=${serverip}:${rootpath},v3,tcp
mount_ubifs=run set_mtdparts
load_ubifs=ubifsload ${loadaddr} uImage.${image_sel}; ubifsload ${rfsaddr}
rfs.${image_sel}
load_ubifs_alt=ubifsload ${loadaddr} uImage.${alt_image}; ubifsload ${rfsaddr}
rfs.${alt_image}
ubifs=run bootargs_initrd; run load_ubifs; bootm ${loadaddr} ${rfsaddr}; echo
ERROR: booting active image failed, trying alternate image...; run
load_ubifs_alt; bootm ${loadaddr} ${rfsaddr}; echo ERROR:bad image(s), aborting
nfs=tftp ${loadaddr} ${bootfile}; run bootargs_nfs; bootm ${loadaddr}
net=run bootargs_initrd; tftp ${loadaddr} ${bootfile} ; tftp ${rfsaddr}
${ramdisk} ; bootm ${loadaddr} ${rfsaddr}

```

Several of these environment variables need to be reconfigured with values that are appropriate to the network context in which the ShMM is used.

3.2.4 Establishing the secondary RC script

The secondary RC script gets invoked when the system configuration is established during the boot process. It is called from the primary RC script `/etc/rc`. The secondary script is a carrier-specific startup script and is `/etc/rc.carrier3` by default on ShMM-500/1500 (and `/etc/rc.shmm700-hpdl` on ShMM-700) or some other script that is appropriate for that platform. A typical name for this script is `/etc/rc.<target_platform>`.

The name of this carrier-specific startup script is defined by the U-Boot environment variable `rc2`. The variable `rc2` is the one environment variable that must definitely be changed for a ShMM and its carrier to work properly in a shelf.

The RC2 script sets up environment variables `CARRIER` and `CARRIER_OPTIONS`. These variables inform the Shelf Manager about the carrier on which it is installed and define carrier-specific options as necessary for each supported carrier. By default, the values of these environment variables are propagated to the corresponding configuration variables (see Table 7: Shelf Manager configuration parameters) `CARRIER` and `CARRIER_OPTIONS`. The configuration variables, in their turn, are retrieved and used by the Shelf Manager.

The U-Boot variable `start_rc2_daemons` instructs the secondary startup script to start or not start the daemons `snmpd` (SNMP server), `boa` (HTTP server) and `shelfman` (the Shelf Manager) after Linux boots. If the U-Boot variable `start_rc2_daemons` is set to `y`, the secondary RC script should also define command-line options for automatic invocation of the `shelfman` daemon. It may also provide other configuration services.

3.3 Setting up Shelf Manager configuration file

The Shelf Manager configuration file (`shelfman.conf`) is located in the `/etc` directory. Each line in the file is either a comment line (starting with `#`) or a `<name> = <value>` pair, representing the assignment for the configuration parameter. The name and the value are separated with the equal sign `=`.

The configuration parameter name is case-insensitive. Each configuration parameter is one of the following types: Boolean, number, string, or IP-address. The values of string type of configuration parameters are case-sensitive. The format of the value conforms to the type of the configuration parameter as shown in the following table.

Table 6: Configuration parameter types and descriptions

CONFIGURATION PARAMETER TYPE	DESCRIPTION
Boolean	A Boolean can be represented by either the strings <code>FALSE</code> (<code>false</code>) or <code>TRUE</code> (<code>true</code>), or by their numerical representations of <code>0</code> and <code>1</code> , respectively.
Number	A whole (possibly signed) numeric value; hexadecimal notation " <code>0x...</code> " is also supported.
String	A string, quoted (always with double quotes <code>""</code>) or unquoted. Quoted strings may contain blanks; unquoted

CONFIGURATION PARAMETER TYPE	DESCRIPTION
	strings are terminated by the first blank. The maximum string size is specified separately for each string-oriented configuration parameter. These values are case-sensitive.
IP-address	An Internet Protocol address in decimal-dot ("xxx.xxx.xxx.xxx") notation.

It is possible to specify a value of an environment variable as a configuration parameter value, using the notation `$<envvar>`; in that case, the value of the variable `<envvar>` is substituted when the configuration file is read. Here is an example:

```
DEFAULT_RMCP_IP_ADDRESS = $IPADDR
```

After the Shelf Manager has been brought up for the first time, the IP addresses are stored with the IPMI LAN configuration parameters. The LAN configuration parameters can be accessed or modified via any of the RMCP, CLI, web, or SNMP external interfaces and take precedence over the *shelfman* configuration file when the Shelf Manager is restarted. This is to ensure the persistency of any modifications that are made to the LAN IP Addresses and gateway via those interfaces.

If the Shelf Manager IP Connection record in the Shelf FRU Information contains an IP address, it takes precedence over all other settings of the external or RMCP IP address. Preferably, the Shelf FRU Information should either not specify this address at all or set it to `0.0.0.0` to ensure that addresses can be controlled through the Shelf Manager configuration file and the IPMI LAN configuration parameters. The value of `0.0.0.0` for an IP-address type of configuration parameter is interpreted as 'undefined'.

The Shelf FRU Information should specify the RMCP IP address; the Shelf Manager uses it and propagates it to the LAN configuration parameters. RMCP is available in the absence of the Shelf FRU Information only if the configuration parameter `RMCP_WITHOUT_SHELF_FRU` is defined and set to `TRUE`. In this case, the Shelf Manager uses the IP address stored in the channel parameters. If there are no stored channel parameters, it uses the IP address specified in the boot parameters of the ShMM.

Some Shelf Manager configuration variables can get their values from the Shelf FRU Information. This may be useful to users who need to associate certain configuration parameters with a specific instance of a shelf. In that case, even if ShMMs or ShMM carriers are moved between shelves, the configuration parameters specified in the Shelf FRU Information stay with the shelf and override the parameter values specified in configuration files on the ShMM.

If the same variable is specified in a configuration file on the ShMM and in the Shelf FRU Information, the value from the Shelf FRU Information overrides the value from configuration files. Note that configuration files are parsed early during the Shelf Manager initialization, while the Shelf FRU Information can be found substantially later.

Not all variables can be specified in the Shelf FRU Information; the main reason is that many variables control the Shelf Manager initialization behavior, which happens before the Shelf FRU Information is found. For such variables, overriding the value when the Shelf FRU Information is found, could not have any effect, since the choice they control has typically already been made. If

such variables are specified in the Shelf FRU Information, they are parsed successfully but value assignments are silently ignored.

The configuration variable *SHELF_MANAGER_CONFIGURATION_IN_SHELF_FRU_INFO* controls whether this functionality of the Shelf Manager is enabled. By default, the value of this variable is *FALSE* and the functionality is disabled.

The following table lists the configuration parameters that are currently supported.

Table 7: Shelf Manager configuration parameters

NAME	TYPE	DEFAULT	DESCRIPTION	CAN BE OBTAINED FROM SHELF FRU INFO
<i>2_X_SYSTEM</i>	Boolean	None	If specified, this parameter explicitly designates the current shelf as CompactPCI (if <i>TRUE</i>) or AdvancedTCA (if <i>FALSE</i>). If not specified, the choice of the shelf type is made automatically. This parameter should not be specified unless it is necessary to override an incorrect hardware detection algorithm for the shelf type.	No
<i>ACTIVATE_LOCAL_WITHOUT_SHELF_FRU</i>	Boolean	FALSE	If set to <i>TRUE</i> , both IPM controllers exposed by the active Shelf Manager (representing the physical and the logical Shelf Managers) are activated even if the Shelf FRU Information cannot be found. This option should be used with caution, because the power consumption of the payload of the physical Shelf Manager IPM controllers may potentially exceed the power capability of the corresponding slot in the shelf.	No
<i>ALARM_CUTOFF_TIMEOUT</i>	Number	600	The alarm cutoff timeout (time after which the alarm cutoff is deactivated), in seconds.	Yes

<code>ALLOW_ALL_COMMANDS_FROM_IPMB</code>	Boolean	TRUE	If set to <i>TRUE</i> , most of the commands allowed from the RMCP interface are allowed from IPMB-0 as well (except for session-related commands). For example, “Cold Reset” and user management commands are accepted from IPMB-0 in this case. In this case, a malicious IPM controller can seriously jeopardize the functionality of the shelf.	Yes
<code>ALLOW_CHANGE_EVENT_RECEIVER</code>	Boolean	TRUE	If set to <i>TRUE</i> , the Event receiver address for the Shelf Manager can be set to an address other than 20h, LUN 0. If set to <i>FALSE</i> , any attempt to change event receiver address for the Shelf Manager is rejected.	Yes
<code>ALLOW_CLEARING_CRITICAL_ALARM</code>	Boolean	FALSE	If set to <i>TRUE</i> , the critical alarm condition can be cleared by the CLI command <i>clia alarm clear</i> .	Yes
<code>ALLOW_POWER_UNRELATED_FRU_IN_CRITICAL_STATE</code>	Boolean	FALSE	This variable affects the behavior of the Shelf Manager with respect to powering up FRUs that are in state M3 when the shelf is in Critical thermal alert state. If set to <i>TRUE</i> , the FRU can be powered on if the Critical alert state is caused by temperature sensors that belong to a different FRU. If set to <i>FALSE</i> , no FRU can be powered up when the shelf is in the Critical alert state. In the case of a Critical alert caused by a shelf-wide sensor, no FRU can be powered up, irrespective of the value of this variable.	Yes

<i>ALLOW_RESET_STANDALONE</i>	Boolean	FALSE	If set to <i>TRUE</i> , the command “Cold Reset” is accepted even if the Shelf Manager does not have an available backup, and reboots the Shelf Manager. By default, the command “Cold Reset” is accepted only in a dual redundant configuration and causes a switchover.	Yes
<i>ALLOWED_CIPHER_SUITES</i>	Number	0xFFFFFFFF (decimal -1)	The bit mask of the IPMI 2.0 cipher suite IDs that a LAN client is allowed to use when establishing a session to the Shelf Manager. By default, all supported cipher suites can be used. If security considerations require, the set of allowed cipher suites can be restricted using this parameter. A bit value of 1b in the bit position corresponding to a cipher suite ID (0 to 14) indicates that the corresponding cipher suite can be used. There must be at least one cipher suite enabled, so a value 0 of this parameter is treated as if all cipher suites are allowed (just as with the default value).	Yes

<i>ALTERNATE_CONTROLLER</i>	Boolean	TRUE	<p>Use alternate controller on the Shelf Manager with the address that is equal to the ShMM hardware address. If this variable is set to <i>TRUE</i>, the active Shelf Manager exposes two IPMB addresses: 20h and a second address based on its hardware address; the backup Shelf Manager exposes only the IPMB address based on its hardware address.</p> <p>After a switchover, the address 20h is exposed by the former backup Shelf Manager, which now exposes two IPMB addresses. If this variable is set to <i>FALSE</i>, the Shelf Manager exposes only the logical address 20h, this is allowed only for a non-redundant Shelf Manager.</p> <p>In the redundant configuration, this variable must be set to <i>TRUE</i>.</p>	No
<i>ATCA_TESTER_COMPATIBILITY</i>	Boolean	FALSE	<p>This variable, if set, turns off event handling optimizations in the Shelf Manager, so that the Shelf Manager behavior is compatible with the Polaris ATCA Tester; in particular, duplicate hot swap events are handled as separate events if they have different IPMI sequence numbers.</p>	No
<i>AUTO_SEND_MESSAGE</i>	Boolean	TRUE	<p>Automatically convert an RMCP request sent to a non-Shelf Manager IPMB address into a “Send Message” request directed to that address.</p>	No
<i>AXIE_NOT_READY_STARTUP_DELAY</i>	Number	1	<p>The AXIe startup delay time in milliseconds.</p>	No

<i>AXIE_POWER_MANAGEMENT</i>	Boolean	FALSE	Enables overall support of AXIe power management functionality in the Shelf Manager; must be turned on explicitly.	No
<i>AXIE_ROOT_PREFERENCE_CHANNELS</i>	String(256)	""	Specifies the order of fabric channels that are processed by the Shelf Manager during the E-keying of the System Module. The channel numbers are separated by commas. The number 0 designates the System Module itself.	No
<i>AXIE_SEQUENCING_ENUM_READY_TIMEOUT</i>	Number	-1	Timeout for the AXIe "Enumeration Not Ready" signal to get cleared. This timeout value is measured in milliseconds. The value 0 means to not wait and the value -1 means to wait indefinitely. Any other value specifies the number of milliseconds that the Shelf Manager waits for the AXIe "Enumeration Not Ready" signal to get cleared. If the timeout expires and the signal does not get cleared, the Shelf Manager terminates the wait, issues an error message and proceeds to the next state of power sequencing.	No

<i>AXIE_SEQUENC ING_M4_TIMEO UT</i>	Number	-1	Timeout for all AXIe modules to reach state M4. This timeout value is measured in milliseconds. The value 0 means to not wait and the value -1 means to wait indefinitely. Any other value specifies the number of milliseconds that the Shelf Manager waits for all AXIe modules to reach state M4. If the timeout expires and not all AXIe modules have reached state M4, the Shelf Manager terminates the wait, issues an error message and proceeds to the next state of power sequencing.	No
<i>AXIE_SEQUENC ING_SYSTEM_M ANAGER_TIMEO UT</i>	Number	-1	Timeout for the Shelf Manager to receive the nVent-specific command "AXIe Allow Enable PCIe Hosts" from the System Manager. This timeout value is measured in milliseconds. The value 0 means to not wait; the value -1 means to wait indefinitely. If the timeout expires and the nVent-specific command "AXIe Allow Enable PCIe Hosts" is not received from the System Manager, the Shelf Manager terminates the wait, issues an error message and proceeds to the next stage of power sequencing.	No
<i>AXIE_SEQUENC ING_WAIT_DEL AY</i>	Number	100	The delay in milliseconds between successive attempts by the Shelf Manager to poll the state of hardware and software signals and advance the current power sequencing state if it can be advanced.	No

<i>AXIE_SYSTEM_MODULE_IPMB_ADDRESS</i>	Number	0x82	The IPMB address of the AXIe system module.	No
<i>AXIE_TIMING_ASYMMETRIC_MATCH</i>	Boolean	FALSE	<p>This variable specifies how the E-Keying logic in the Shelf Manager matches AXIe timing links. If this variable is set to <i>TRUE</i>, asymmetric matching applies; that is, links tagged as “Instrument Slot Inputs” match links tagged as “System Slot Outputs” on the other side. If this variable is set to <i>FALSE</i>, regular matching rules apply to AXIe timing links.</p> <p>This configuration variable affects only AXIe configurations; in non-AXIe configurations its value is ignored.</p>	No
<i>BOARD_LAN_PARAMETERS_ASSIGN_FROM_SHM</i>	Boolean	FALSE	<p>This variable allows assigning LAN configuration (IP, Netmask, Gateway) to boards based on the IP address at LAN channel 1 plus geographical board number, e.g. If ShMM RMCP IP is set to <i>1.2.3.4</i> then board 3 will get IP <i>1.2.3.7</i> with the same network and gateway IP address. This only works if board has a LAN channel with IP Address Source value <i>3</i> (BIOS or system software). Note that it’s up to user to keep the assigned address range within logical network boundaries and not to have gateway IP address in this range.</p>	No

<i>BOARD_LAN_PARAMETERS_CHANNEL_LIST</i>	String(64)	""	Deprecated. The list of IPMI channel numbers on boards and modules that are available for assignment of LAN configuration parameters by the Shelf Manager. Channel numbers, each in the range 1 to 7, can be separated by commas, spaces or any other separators, for example: <i>3, 4, 5</i> . LAN parameters are assigned only to channels specified in this list and in the order in which they appear in this list.	Yes
<i>BOARD_LAN_PARAMETERS_SYNCHRONOUS</i>	String(256)	""	Deprecated. This string value represents the list of descriptors for FRUs that require synchronous assignment of LAN configuration parameters. It is meaningful only if the variable <i>BOARD_LAN_PARAMETERS_USE_DHCP</i> is <i>TRUE</i> . The detailed syntax for this value is defined in section 4.10.4.	Yes

<i>BOARD_LAN_PARAMETERS_USE_DHCP</i>	Boolean	FALSE	<p>If <i>TRUE</i>, use DHCP to retrieve board/module LAN configuration parameters.</p> <p>If this parameter is <i>FALSE</i>, the following options are available:</p> <ul style="list-style-type: none"> • If the Shelf FRU Information contains one or more Board/AMC LAN Configuration Parameters multirecords, use these to retrieve the board/module LAN configuration parameters. • Otherwise, if the LAN configuration parameters file exists on the ShMM, use this file to retrieve the board/module LAN configuration parameters. • Otherwise, Shelf Manager assignment of board/module LAN configuration parameters is not available. 	Yes
<i>CARRIER</i>	String(16)	"PPS"	The name of the specific carrier board on which the ShMM is installed.	No
<i>CARRIER_OPTIONS</i>	String(256)	Variable	The carrier-specific options; defined separately for each supported carrier. By default, this parameter is set from the environment variable <i>\$CARRIER_OPTIONS</i> .	No
<i>CONSOLE_LOGGING_ENABLED</i>	Boolean	FALSE	Output log messages to the console on which the Shelf Manager was started.	No

<i>COOLING_EXHAUST_SENSOR_LIST</i>	String(32)	"" (undefined)	This parameter is used in certain shelf-specific cooling algorithms and specifies the exhaust temperature sensors of the shelf. The format of the parameter is a comma-separated list of tuples in the form "<owner-address>/<sensor-number>". The sensor LUN must be 0. For example, the value <i>"0x20/120, 0x20/121"</i> indicates that sensors 120 and 121 on the logical Shelf Manager (IPMB address 20h) represent the exhaust temperature sensors of the shelf.	Yes
<i>COOLING_FAN_DECREASE_TIMEOUT</i>	Number	0	The minimum timeout between successive decrements of the fan speed during operation of the cooling algorithm in Normal state. Should be a multiple of <i>COOLING_POLL_TIMEOUT</i> ; if not, it is rounded up to the next multiple. If the parameter is omitted or set to 0, this timeout is equal to <i>COOLING_POLL_TIMEOUT</i> .	Yes
<i>COOLING_FAN_INCREASE_TIMEOUT</i>	Number	0	The minimum timeout between successive increments of the fan speed during operation of the cooling algorithm in Minor Alert state. Should be a multiple of <i>COOLING_POLL_TIMEOUT</i> ; if not, it is rounded up to the next multiple. If the parameter is omitted or set to 0, this timeout is equal to <i>COOLING_POLL_TIMEOUT</i> .	Yes

<i>COOLING_INTAKE_SENSOR_LIST</i>	String(32)	"" (undefined)	This parameter is used in certain shelf-specific cooling algorithms and specifies the intake temperature sensors of the shelf. The format of the parameter is a comma-separated list of tuples in the form "<owner-address>/<sensor-number>". The sensor LUN must be 0. For example, the value " <i>0x20/124,0x20/125</i> " indicates that sensors 124 and 125 on the logical Shelf Manager (IPMB address 20h) represent the intake temperature sensors of the shelf.	Yes
<i>COOLING_KEEP_POWERED_OFF_FRUS_IN_M1</i>	Boolean	FALSE	If set, the FRUs that are powered off by the Shelf Manager due to a critical temperature stay in state M1 and are not automatically activated and powered on when the temperature condition goes away; manual intervention is needed to reactivate each such FRU.	Yes
<i>COOLING_NO_ALL_FRUS_POWER_DOWN_IN_CRITICAL_ALERT</i>	Boolean	FALSE	Do not power down the whole shelf if a critical thermal alert is reported by a shelf-wide temperature sensor or a sensor that cannot be associated with a specific FRU. Specific FRUs can still be powered down in the case of a critical thermal alert in this mode. This mode should be used with caution.	Yes

<code>COOLING_NO_POWER_DOWN_IN_CRITICAL_ALERT</code>	Boolean	FALSE	Do not power down FRUs that experience critical thermal alerts. This mode is not compliant with the PICMG 3.0 (AdvancedTCA) specification and should be used with extreme caution.	Yes
<code>COOLING_MANAGEMENT</code>	String(64)	"" (undefined)	If specified, the name of the shared library that implements cooling management. The actual name of the library is <code>libcooling_<xxx>.so</code> , where <code><xxx></code> is the value of this configuration parameter. This library is dynamically loaded by the Shelf Manager and must be located in <code>/var/bin</code> or <code>/lib</code> .	No
<code>COOLING_POLL_TIMEOUT</code>	Number	30	The maximum time (in seconds) between successive invocations of the cooling monitoring and management thread.	Yes
<code>CPLD_ACTIVE_WORKAROUND</code>	Boolean	TRUE	This flag (when <code>TRUE</code>) enables a special workaround to detect the loss of the Active signal in the CPLD. (Registers in the CPLD (Complex Programmable Logic Device) on the ShMM expose the state of the hardware redundancy signals to the software.) This loss may happen on some platforms when hot inserting a ShMM carrier. The workaround can be turned off by setting this value to <code>FALSE</code> , if this problem does not exist for a specific platform. This parameter is not applicable to the ShMM-700.	No

<i>CTCA_FRU_RESET_TIMEOUT</i>	Number	500	CompactPCI shelves only. The time in milliseconds during which the Shelf Manager holds the BD_SEL# line low in order to reset a CompactPCI board.	No
<i>CTCA_HEALTHY_TIMEOUT</i>	Number	0	CompactPCI shelves only. The time in seconds during which the Shelf Manager waits for the HEALTHY# signal to appear when powering on a CompactPCI board. If the HEALTHY# signal does not appear within the specified time, the Shelf Manager deactivates the board. 0 (the default) stands for “infinity”.	No
<i>CTCA_INITIAL_FAN_LEVEL</i>	Number	15	CompactPCI shelves only. The initial fan speed (in the range 0..15) that Shelf Manager applies to fan trays in CompactPCI shelves. 0 corresponds to the slowest, and 15 to the fastest possible speed.	No
<i>DEFAULT_GATEWAY_IP_ADDRESS</i>	IP-address	None	The default IP address used for the gateway for shelf-external (RMCP-based) communication, if the corresponding parameter is set to 0.0.0.0 in the IPMI LAN Configuration Parameters for channel 1. If a non-zero gateway IP address is provided in the LAN Configuration Parameters, the value provided in the Shelf Manager configuration file is ignored.	No

<i>DEFAULT_GATEWAY_IP_ADDRESS2</i>	IP-address	None	The default IP address used for the gateway for shelf-external (RMCP-based) communication on the second network interface, if the corresponding parameter is set to <i>0.0.0.0</i> in the IPMI LAN Configuration Parameters for channel 2. If a non-zero gateway IP address is provided in the LAN Configuration Parameters, the value provided in the Shelf Manager configuration file is ignored.	No
<i>DEFAULT_HPM2_CLIENT_IP_ADDRESS</i>	IP-address	None	The default IP address used for HPM.2-based network communication with intelligent FRUs in the shelf. This IP address is used only if the corresponding parameter is set to <i>0.0.0.0</i> in the IPMI LAN Configuration Parameters for the HPM.2 channel (channel 8 on the Shelf Manager) and if the HPM.2 channel is not mirrored to another channel. If a non-zero IP address is provided in the LAN Configuration Parameters for the HPM.2 channel, the value provided in the Shelf Manager configuration file is ignored. If the HPM.2 channel is mirrored to another channel, the IP address is taken from the parameters of that other channel and this parameter is also ignored.	No

<i>DEFAULT_HPM2 _GATEWAY_IP_ ADDRESS</i>	IP-address	None	The default IP address used as the gateway IP address for HPM.2-based network communication with intelligent FRUs in the shelf. This IP address is used only if the corresponding parameter is set to 0.0.0.0 in the IPMI LAN Configuration Parameters for the HPM.2 channel (channel 8 on the Shelf Manager) and if the HPM.2 channel is not mirrored to another channel. If a non-zero gateway IP address is provided in the LAN Configuration Parameters for the HPM.2 channel, the value provided in the Shelf Manager configuration file is ignored. If the HPM.2 channel is mirrored to another channel, the gateway IP address is taken from the parameters of that other channel and this parameter is also ignored.	No
--	------------	------	--	----

<i>DEFAULT_HPM2_NETMASK</i>	IP-address	Variable	The network mask for the network adapter used for the HPM.2-based network communication with intelligent FRUs in the shelf. This mask is used only if the corresponding parameter is set to <i>0.0.0.0</i> in the IPMI LAN Configuration Parameters for the HPM.2 channel (channel 8 on the Shelf Manager) and if the HPM.2 channel is not mirrored to another channel. The default value depends on the class of the client IP address used for the HPM.2-based network communication. (see parameter <i>DEFAULT_HPM2_CLIENT_IP_ADDRESS</i>). For example, for an IP address of class C, this parameter is set to <i>255.255.255.0</i>	No
<i>DEFAULT_RMCP_IP_ADDRESS</i>	IP-address	None	The default IP address used for shelf-external (RMCP-based) communication; it is switched over between the redundant instances of the Shelf Manager. This IP address is used only if the corresponding parameter is set to <i>0.0.0.0</i> in the IPMI LAN Configuration Parameters for channel 1 and in the Shelf Manager IP Connection record in Shelf FRU Information. If a non-zero IP address is provided in the LAN Configuration Parameters and/or Shelf FRU Information, the value provided in the Shelf Manager configuration file is ignored.	No

<i>DEFAULT_RMCP_IP_ADDRESS2</i>	IP-address	None	The default IP address used for shelf-external (RMCP-based) communication on the second network interface; it is switched over between the redundant instances of the Shelf Manager. This IP address is used only if the corresponding parameter is set to <i>0.0.0.0</i> in the IPMI LAN Configuration Parameters for channel 2. If a non-zero IP address is provided in the LAN Configuration Parameters, the value provided in the Shelf Manager configuration file is ignored.	No
<i>DEFAULT_RMCP_NETMASK</i>	IP-address	Variable	The network mask for the network adapter used for RMCP communication. This mask is used only if the corresponding parameter is set to <i>0.0.0.0</i> in the IPMI LAN Configuration Parameters for channel 1 and in the Shelf Manager IP Connection record in Shelf FRU Information. The default value depends on the class of the default IP address used for shelf-external (RMCP-based) communication. (see parameter <i>DEFAULT_RMCP_IP_ADDRESS</i>). For example, for an IP address of class C, this parameter is set to <i>255.255.255.0</i>	No

<i>DEFAULT_RMCP_NETMASK2</i>	IP-address	Variable	The network mask for the second network adapter used for RMCP communication. This mask is used only if the corresponding parameter is set to <i>0.0.0.0</i> in the IPMI LAN Configuration Parameters for channel 2. The default value depends on the class of the default IP address used for shelf-external (RMCP-based) communication. (see parameter <i>DEFAULT_RMCP_IP_ADDRESS2</i>). For example, for an IP address of class C, this parameter is set to <i>255.255.255.0</i>	No
<i>DEFAULT_VLAN_ID</i>	Number	0	The default Virtual LAN ID used for the first LAN channel. This value is used only as the default value for the corresponding LAN configuration parameter; once LAN configuration parameters are defined and stored on the ShMM, the value provided in the Shelf Manager configuration file is ignored.	Yes
<i>DEFAULT_VLAN_ID2</i>	Number	0	The default Virtual LAN ID used for the second LAN channel. This value is used only as the default value for the corresponding LAN configuration parameter; once LAN configuration parameters are defined and stored on the ShMM, the value provided in the Shelf Manager configuration file is ignored.	Yes

<i>DETECT_DEADLOCKS</i>	Boolean	TRUE	This variable turns on the deadlock detection in CLI and RMCP server facilities in the Shelf Manager. The detection is based on an internal watchdog that must be periodically strobed by the threads serving CLI and RMCP requests. If one of the threads fails to strobe the internal watchdog, the actual watchdog does not get strobed, and ultimately the ShMM resets, initiating a failover to the backup Shelf Manager. In addition (currently on all platforms except ShMM-700), lock data structures are periodically checked directly for the presence of a deadlock. See section 4.8 for further background.	No
<i>DEVICE_POLL_TIMEOUT</i>	Number	15	The time (in seconds) between successive polls of the IPMB-0 devices by the Shelf Manager via sending the “Get Device ID” command to them.	Yes
<i>DHCP_FOR_RMC_P_ONLY</i>	Boolean	FALSE	If this variable is set, only the RMCP IP addresses are assigned via the DHCP mechanism; private IP addresses of the ShMMs are not touched by DHCP.	Yes
<i>DHCP_REQUEST_OPTION_LIST</i>	Boolean	FALSE	This variable causes the DHCP client to request an explicit list of options via parameter request list option 55. By default, the DHCP request does not specify desired options, so the DHCP server does not necessarily return such information.	Yes

<i>DHCP_SERVER_ADDRESS</i>	IP-address	None	This parameter is the IP address of the DHCP server; it applies only if the variable <i>USE_DHCP</i> is <i>TRUE</i> . If this parameter is omitted or set to <i>0.0.0.0</i> , and the <i>USE_DHCP</i> variable is <i>TRUE</i> , the Shelf Manager accepts address information from any DHCP server that responds to its broadcast discovery request.	Yes
<i>DHCP_SPECIAL_CLIENT_ID_FORMAT</i>	Boolean	FALSE	If this parameter is set to <i>TRUE</i> , the DHCP Client ID has a special format that is used by specific customers for Pigeon Point products. By default this variable is <i>FALSE</i> and a standard Pigeon Point Client ID format is used, as described in section 3.4.9 of this document.	Yes
<i>DHCPV6_OPTIONS_SCRIPT</i>	String	""	ShMM-700 only: This variable contains the path to an external executable (or a shell script) that the DHCPv6 client invokes when it receives values of requested DHCPv6 options from the server. The client passes the values of the options to this executable via command-line parameters. See section 3.4.11 for more information.	Yes

<i>DHCPV6_REQUESTED_OPTIONS</i>	String	""	ShMM-700 only: This variable contains the list of DHCPv6 options that the DHCPv6 client requests from the server. The client parses the values of these options that the server returns in the REPLY PDU and feeds them to an external script specified by the configuration parameter <i>DHCPV6_OPTIONS_SCRIPT</i> . No options are requested by default. See section 3.4.11 for more information.	Yes
<i>DHCPV6_SERVER_ADDRESSES</i>	String	""	ShMM-700 only: This string parameter is the list of IPv6 addresses of the allowed DHCPv6 servers, separated by blanks, commas or semicolons; it applies only if the variable <i>USE_DHCPV6</i> is <i>TRUE</i> . If this parameter is omitted and the <i>USE_DHCPV6</i> variable is <i>TRUE</i> , the Shelf Manager accepts address information from any DHCPv6 server that responds to its broadcast solicitation request.	Yes

<i>DISABLE_SESSION_AUDIT_EVENTS</i>	Number	0	<p>This bit mask controls whether assertion events are generated by the Session Audit sensor in response to specific state changes of the sensor. By default, assertion events are generated by the Session Audit sensor (sensor 139, IPMB address 20h) and logged to the SEL for all defined states. The bits of the bit mask have the following meaning:</p> <p>[0] - 1b disables generation of “Session Activated” events.</p> <p>[1] - 1b disables generation of “Session Deactivated” events.</p> <p>[2] - 1b disables generation of “Invalid username or password” events.</p> <p>[3] - 1b disables generation of “Invalid password disable” events.</p>	Yes
<i>ENABLE_DIRECT_SHELF_FRU_WRITE</i>	Boolean	FALSE	<p>This variable controls whether the Shelf Manager allows direct writes to the Shelf FRU Information (FRU Device ID #254 on the IPM controller at 20h on IPMB-0) via the IPMI command “Write FRU Data”. By default, direct writes are prohibited (as mandated by PICMG 3.0 R2.0 ECN-002, which requires that Shelf FRU Information writes use a special locking protocol so that only one writer is active at once). Setting this variable to <i>TRUE</i> enables direct writes for compatibility with System Manager applications that rely on the pre-ECN-002 behavior.</p>	No

<i>ENABLE_INTEGRALHPI</i>	Boolean	FALSE	This variable determines whether the Shelf Manager exposes the Hardware Platform Interface (HPI) defined by the Service Availability Forum (www.saforum.org). Setting this variable to <i>TRUE</i> enables the separately licensed IntegralHPI implementation within the Shelf Manager, and makes the Shelf Manager accessible via HPI.	No
<i>ENABLE_LOCKS_LOGGING</i>	Boolean	TRUE	If this variable is <i>TRUE</i> (default value), the Shelf Manager records all requests to acquire and release locks (mutexes) in an internal memory area. This information can be subsequently recovered and analyzed by the analysis tool <i>dumplog</i> in case of a deadlock or another synchronization violation. Setting this variable to <i>FALSE</i> makes post-mortem lock analysis impossible, but substantially improves the performance of the Shelf Manager. See section 4.8 for further background.	No
<i>ENABLE_NEW_DHCPV4_CLIENT</i>	Boolean	TRUE	This variable determines which DHCPv4 client is used by the Shelf Manager if <i>USE_DHCP</i> is set to <i>TRUE</i> . By default, the new DHCPv4 client is used. See section 3.4.9 for more information.	No

<i>ENABLE_RTC_TRICKLE_CHARGER</i>	Boolean	FALSE	If this variable is <i>TRUE</i> , the trickle charge feature of the RTC DS1339 device on the ShMM is enabled. This setting is required for ShMM carriers that connect a super cap to the RTC backup power supply input (VBAT or VBACKUP_RTC).	No
<i>EXIT_IF_HEALTHY_LOST_IN_STANDALONE_MODE</i>	Boolean	FALSE	This variable defines what to do if the Shelf Manager runs without backup and detects a loss of the local Healthy bit. If this variable is <i>TRUE</i> , the Shelf Manager exits, the ShMM reboots and the Shelf Manager is restarted. If this variable is <i>FALSE</i> , the Shelf Manager sets the Healthy bit and continues operation	Yes
<i>EXIT_IF_NO_SHELF_FRU</i>	Boolean	FALSE	If <i>TRUE</i> , the Shelf Manager exits (probably resetting the ShMM) if no Shelf FRU Information can be found.	No
<i>EXTERNAL_EVENT_HANDLER</i>	String(255)	""	This is the path to an executable file (or a script file) on the Shelf Manager that performs local handling of events via PEF.	No

<i>FAN_FULL_SPEED_DELAY</i>	Number	0	The delay in seconds after Shelf Manager startup or after a switchover, during which the cooling algorithm does not check the number of present fan trays, giving the existing fan trays enough time to activate. This applies to the carriers where cooling management raises the fan speed to maximum if the actual number of fan trays in the shelf is fewer than what is specified in the Shelf Address Table.	No
<i>FAN_LEVEL_STEP_DOWN</i>	Number	1	The number of fan steps by which the fan speed is decreased during operation of the cooling algorithm in the Normal state. This parameter may be overridden by a ShMM carrier-specific cooling algorithm.	Yes
<i>FAN_LEVEL_STEP_UP</i>	Number	1	The number of fan steps by which the fan speed is increased during operation of the cooling algorithm in the Minor Alert state. This parameter may be overridden by a carrier-specific cooling algorithm.	Yes
<i>HPDL</i>	Boolean	FALSE	Turns on HPDL support in the Shelf Manager. The carrier and chassis HPDL data and SDRs are taken from the FRU Information or from the files and are used to define the behavior of the platform, plus the number and types of managed FRUs and sensors.	No

<i>HPDL_ON_SUBSIDIARY_FRUS</i>	Boolean	TRUE	Turns on support of HPDL information stored on subsidiary FRUs. If <i>TRUE</i> , the Shelf Manager looks for HPDL data and SDRs in the FRU Information of its subsidiary FRUs. If these data are found for a specific FRU, they are used to substitute definitions for that FRU from the carrier or chassis HPDL data and/or SDRs.	Yes
<i>HPDL_NETWORK_ELEMENT_ID_EEPROM_OFFSET</i>	Number	0	If non-zero, specifies the offset to the Network Element ID data structure in Shelf FRU Information EEPROMs. This 33-byte data structure consists of the three separate 11-byte Network Element IDs and is transparent for the Shelf Manager but can be read and modified with CLI commands. A positive value of this parameter indicates the offset from the beginning of the EEPROM, a negative value means the offset from the end of the EEPROM. Zero value indicates the absence of the Network Element ID. This parameter makes sense only for HPDL-based systems and only if local EEPROMs are used to store Shelf FRU Information.	Yes

<i>HPM2_ENABLE_CLIENT_SESSIONS</i>	Boolean	TRUE	<p>Enables the Shelf Manager to establish and use HPM.2 client sessions to HPM.2-capable intelligent FRUs in the Shelf. If this parameter is <i>TRUE</i>, the Shelf Manager, when detecting an intelligent FRU (an IPMC or an AMC) that supports HPM.2 sessions, attempts to establish a client session to that FRU. Subsequently the Shelf Manager uses these sessions to retrieve SDRs and FRU information from that FRU (which can be significantly faster than using IPMB).</p> <p>If this parameter is <i>FALSE</i>, the Shelf Manager does not establish HPM.2 client sessions to intelligent FRUs in the shelf; this mode of operation is compatible with previous releases of the Shelf Manager that lacked this level of HPM.2 support.</p>	No
------------------------------------	---------	------	--	----

<i>HPM2_ENABLE_INACTIVE_STATE_MANAGEMENT</i>	Boolean	TRUE	<p>Enables HPM.2 Extended Inactive State Management in the Shelf Manager. If this parameter is <i>TRUE</i>, the Shelf Manager, when detecting an IPMC that supports Extended Inactive State Management, attempts to turn on the extended management power and enable corresponding communication ports on that IPMC. Subsequently the Shelf Manager maintains extended management power and corresponding communication ports while the IPMC is in state M1.</p> <p>If this parameter is <i>FALSE</i>, the Shelf Manager does not manage extended management power on IPMCs; this mode of operation is compatible with previous releases of the Shelf Manager that lacked HPM.2 support.</p>	No
<i>HPM2_ENABLE_SHMM_INACTIVE_STATE</i>	Boolean	FALSE	<p>Enables HPM.2 Extended Inactive State Management support in the physical Shelf Manager on non-HPDL systems and on HPDL-systems where extended management power or extended ports are not specified explicitly. In this mode, no additional extended management power is requested, but both cross-connect ports stay available in the inactive state (M1). With the default value of this parameter (<i>FALSE</i>), Shelf Manager operation is compatible in this area with previous releases.</p>	No

HPM2_CLIENT_MIRROR_CHANNEL	Number	1	<p>This configuration parameter specifies whether the HPM.2 client channel (channel 8 for the Shelf Manager) is mirrored to some other LAN channel, and the number of this channel.</p> <p>If this parameter has a non-zero value, the HPM.2 client channel uses the LAN configuration (IP address, subnet mask, gateway IP address, etc.) of the given channel. If a LAN configuration parameter is changed for the mirror channel, it is also automatically changed for the HPM.2 client channel. Changing LAN configuration parameters for the HPM.2 client channel is not allowed.</p> <p>If this parameter has value 0, mirroring is not in use, and the HPM.2 client channel shall be set up and configured explicitly, in the same way as the mandatory LAN channel 1 and the optional LAN channel 2.</p> <p>By default, the HPM.2 channel is mirrored to the LAN channel 1, utilizing the same LAN configuration parameters for the RMCP server and the HPM.2 RMCP client in the Shelf Manager.</p>	No
----------------------------	--------	---	---	----

<i>HPM2_NET_ADAPTER</i>	String(16)	"" (undefined)	The name of the network adapter used for HPM.2-based network communication with intelligent FRUs in the Shelf. This parameter is ignored if the HPM.2 client channel (channel 8 on the logical Shelf Manager) is mirrored to another network channel, but must be specified otherwise.	No
<i>HPM3_ENABLE</i>	Boolean	TRUE	Enables HPM.3-compliant format of DHCP client IDs when retrieving addressing information for the logical and physical Shelf Manager(s) over DHCP. If this parameter is <i>FALSE</i> , the Shelf Manager uses the old format of DHCP client IDs; this mode of operation is compatible with previous releases of the Shelf Manager that did not support HPM.3.	No

<i>IGNORE_FAILURE_DIRECTED_POWER_DOWN</i>	Boolean	TRUE	This parameter tells the Shelf Manager to ignore board power down failure(s) during a Critical Alert and not to power down all boards in the shelf in this case. This parameter also instructs the Shelf Manager to ignore board power level decrease failure(s) during a Major Alert and not to decrease power level for all boards in the shelf in this case. If this parameter is <i>FALSE</i> , all boards in the shelf are powered down, if board(s) that caused the Critical Alert cooling state fail to power off; also, the power level for all boards in the shelf is decreased, if decreasing power level fails for board(s) that caused the Major Alert cooling state.	Yes
<i>INITIALIZATION_SCRIPT</i>	String(256)	""	This parameter specifies a command line that is executed by the active Shelf Manager during initialization. This initialization script can perform additional platform-specific initialization of the Shelf Manager; when it is started, the Local Healthy and Active bits in CPLD are already set. The Shelf Manager waits for the completion of the initialization script. The initialization script is run only during the startup of the Shelf Manager; it is not run after a switchover.	No

<i>INITIAL_FAN_LEVEL</i>	Number	15	This parameter specifies the initial fan level that the Shelf Manager applies to fan trays. Usually fan level values are in 0..15 range, where 0 is the slowest, and 15 is the fastest possible fan speed. This parameter has an alias <i>CTCA_INITIAL_FAN_LEVEL</i> for CompactPCI shelves.	No
<i>INITIAL_SLOW_LINK_DELAY</i>	Number	0	The initial delay, in seconds, before the Shelf Manager starts testing the integrity of the physical network link between the Shelf Manager and the System Manager (the RMCP link; see the description of the configuration parameter <i>SWITCHOVER_TIMEOUT_ON_BROKEN_LINK</i>). A non-zero delay can be used to accommodate slow network links that need significant time to initialize after shelf power up.	Yes
<i>INNER_SEQUENCE_NUMBER_IN_SEND_MSG_RESPONSE</i>	Boolean	TRUE	This variable controls which sequence number is used in the response to a “Send Message” command bridged from LAN to IPMB. If <i>TRUE</i> , the sequence number of the command encapsulated in the Send Message request is used. If <i>FALSE</i> , the sequence number of the Send Message request itself is used. According to a clarification being proposed for the latest version of the PICMG 3.0 specification, the first variant is correct, while the Shelf Manager historically used the second variant.	No

<i>INTEGRALHPI_ DEFAULT_AUTO_ _EXTRACT_TIM EOUT</i>	Number	0	This variable defines the seconds portion of the default auto-extraction timeout value for all exposed HPI resources. A negative value causes the timeout to be set to <i>SAHPI_TIMEOUT_BLOCK</i> .	Yes
<i>INTEGRALHPI_ DEFAULT_AUTO_ _EXTRACT_TIM EOUT_MSEC</i>	Number	1	This variable defines the milliseconds portion of the default auto-extraction timeout value for all exposed HPI resources.	Yes
<i>INTEGRALHPI_ DEFAULT_AUTO_ _INSERT_TIME OUT</i>	Number	0	This variable defines the seconds portion of the default auto-insertion timeout value for this IntegralHPI domain. A negative value causes the timeout to be set to <i>SAHPI_TIMEOUT_BLOCK</i> .	Yes
<i>INTEGRALHPI_ DEFAULT_AUTO_ _INSERT_TIME OUT_MSEC</i>	Number	1	This variable defines the milliseconds portion of the default auto-extraction timeout value for this IntegralHPI domain.	Yes
<i>INTEGRALHPI_ DOMAIN_ID</i>	Number	1	The HPI domain ID of the current shelf; used to differentiate, on the HPI client level, among multiple shelves whose Shelf Managers support HPI.	No
<i>INTEGRALHPI_ KEEPALIVE_IN TERVAL</i>	Number	2	This configuration parameter applies to TCP-based IntegralHPI sessions; it is used to set the interval (in seconds) between successive keepalive probes.	Yes

<i>INTEGRALHPI_KEEPA LIVE_PROBES</i>	Number	5	This configuration parameter applies to TCP-based IntegralHPI sessions; it is used to set the number of unacknowledged probes to send before considering the TCP session to be dead.	Yes
<i>INTEGRALHPI_KEEPA LIVE_TIME</i>	Number	10	This configuration parameter applies to TCP-based IntegralHPI sessions; it is used to set the interval (in seconds) between the last data packet sent and the first keepalive probe.	Yes
<i>INTEGRALHPI_SESSION_ EVENTS_CAPACITY</i>	Number	100	The maximum number of entries in the event queue for each IntegralHPI session. When choosing this number, the memory capacity of the ShMM should be considered. Each event queue entry occupies approximately 300 bytes and 64 simultaneous IntegralHPI sessions are supported. Therefore, increasing this number by 1 causes the Shelf Manager to use 19.2 Kbytes of additional ShMM memory.	Yes
<i>INTEGRALHPI_SHELF_ ENTITY_LOCATION</i>	Number	1	The entity location for the AdvancedTCA Shelf HPI Resource; used to differentiate, on the HPI client level, among multiple shelves whose Shelf Managers support HPI.	Yes

<i>IPMB_ADDRESS</i>	Number	0	The IPMB address of the Shelf Manager, overriding the hardware address. The value of 0 causes the Shelf Manager to read the hardware address from hardware and set IPMB address to hardware address * 2.	No
<i>IPMB_LINK_ISOLATION_TIMEOUT</i>	Number	-1	In radial shelves, if an IPMB link is disabled due to the isolation algorithm, the link is automatically enabled after this time interval (in seconds). -1 (the default) stands for “forever”.	Yes
<i>IPMB_LOCAL_FAILURE_RETRIES</i>	Number	10	ShMM-700 only: in the case of a local IPMB failure (as reported by the IPMB driver), the number of times the IPMB operation is retried. After the specified number of retries is exhausted on IPMB-A and is also exhausted on IPMB-B, a switchover to the backup Shelf Manager takes place.	No
<i>IPMB_RETRIES</i>	Number	3	The number of attempts to send an IPMB request before finally giving up, if no response is received to this request.	No
<i>IPMB_RETRY_TIMEOUT</i>	Number	0	The amount of time (in seconds) that the Shelf Manager waits for a response after sending an IPMB request, before retrying the request.	No

<i>IPMB_RETRY_TIMEOUT_MSEC</i>	Number	500	The millisecond part of the retry timeout value that combines with the <i>IPMB_RETRY_TIMEOUT</i> value to form the full timeout value. This variable allows retry timeout values less than a second, given a zero value for the configuration variable <i>IPMB_RETRY_TIMEOUT</i> . Using an overall IPMB timeout value less than 250 msec is not recommended.	No
<i>IPMC_PRESERVE_ON_REVISION_CHANGE</i>	Boolean	TRUE	Setting this variable to <i>TRUE</i> preserves the Shelf Manager's identification of an IPM controller after a firmware upgrade if only the Firmware Revision and/or Auxiliary Firmware Revision information is changed. If the variable is set <i>FALSE</i> , any change in "Get Device ID" response data during the Shelf Manager's regular polls is considered to signal the presence of a different IPM controller.	Yes
<i>ISOLATE_MUX_ADDRESS</i>	Number	0x70	The 7-bit I ² C multiplexer address (on platforms where SHMM_GPIO8 is used to control access from the Shelf Manager to the multiplexer on the master-only I ² C bus, via the original mechanism that is not based on HPDL).	No

<i>ISOLATE_MUX_BUS</i>	Number	Platform-dependent (5 on ShMM-700, 0 on other platforms)	The number of the logical I ² C bus on which the I ² C switch/multiplexer resides (on platforms where SHMM_GPIO8 is used by the Shelf Manager to reset the I ² C switch/multiplexer via the original mechanism that is not based on HPDL).	No
<i>ISOLATE_MUX_IGNORE_COUNT</i>	Number	10	This parameter instructs the isolation algorithm to skip this number of accesses to a faulty bus before trying to enable it (on the platforms where SHMM_GPIO8 is used to control access from the Shelf Manager to the multiplexer on the master-only I ² C bus, via the original mechanism that is not based on HPDL).	No
<i>ISOLATE_MUX_IGNORE_TIME</i>	Number	10	This parameter instructs the isolation algorithm to avoid accessing the faulty bus and immediately return a failure for specified number of seconds after a failure detection (on the platforms where SHMM GPIOs are used to control access from the Shelf Manager to multiplexers on the master-only I ² C bus, via an HPDL-based mechanism).	No
<i>ISOLATE_MUX_ON_GPIO8</i>	Boolean	FALSE	Must be set to <i>TRUE</i> for platforms where SHMM_GPIO8 is used by the Shelf Manager to reset an I ² C switch/multiplexer on an I ² C bus (but must be <i>FALSE</i> if the HPDL-based reset mechanism is used).	No

<i>ISOLATE_MUX_SWITCHOVER_THRESHOLD</i>	Number	0	If this parameter is non-zero and <i>ISOLATE_MUX_ON_GPIO8</i> is <i>TRUE</i> , the I ² C isolation algorithm will perform a switchover to the backup Shelf Manager if a fault exists on the main I ² C bus after the I ² C switch is reset. In this case, this parameter specifies the number of consecutive failed I ² C switch resets that are attempted, after which, if the main I ² C bus remains inaccessible, a switchover takes place and the active Shelf Manager reboots. This parameter is only applicable on platforms where SHMM GPIO8 is used by the Shelf Manager to reset an I ² C switch via the original mechanism that is not based on HPDL.	No
<i>LOCAL_SHELF_FRU</i>	Boolean	TRUE	Create a local FRU #1 on the Shelf Manager that exposes the Shelf FRU Information (obtained from the file <i>/var/nvdata/shelf_fru_info</i>).	No
<i>M7_TIMEOUT</i>	Number	-1	The maximum time (in seconds) for a FRU to stay in M7 state; after the expiration of this time the FRU is automatically transitioned to M0. -1 (the default) stands for "forever". Setting this parameter to 0 completely prevents FRUs from going into state M7.	Yes
<i>MAX_ALERT_POLICIES</i>	Number	64	The maximum number of PEF Alert Policies available. The maximum allowed value for this parameter is 128.	No

<i>MAX_ALERT_STRINGS</i>	Number	64	The maximum number of PEF Alert Strings available.	No
<i>MAX_DEFERRED_ALERTS</i>	Number	32	The maximum number of outstanding PEF alerts.	No
<i>MAX_EVENT_FILTERS</i>	Number	64	The maximum number of PEF event filters available.	No
<i>MAX_EVENT_SUBSCRIBERS</i>	Number	64	The maximum number of entities that can simultaneously subscribe to receive event notifications from the Shelf Manager.	No
<i>MAX_EVENT_SUBSCRIBER_IDLE_TIME</i>	Number	60	The maximum timeout for an event subscriber, in seconds, between the moment when an event arrives and the moment when the subscriber retrieves this event from the Shelf Manager. If this timeout is exceeded, the subscriber is considered dead and is automatically unregistered.	Yes
<i>MAX_INCOMING_IPMB_REQUESTS</i>	Number	128	The size of the internal Shelf Manager queue for incoming IPMB requests. Incoming IPMB requests are stored in this queue before processing.	No
<i>MAX_NODE_BUSY_TRANSMISSIONS</i>	Number	255	The maximum number of transmissions of an IPMB command if the receiver always returns the completion code Node Busy in response.	Yes
<i>MAX_OEM_FILTERS</i>	Number	16	The maximum number of PEF OEM event filters available.	No
<i>MAX_PENDING_EVENT_NOTIFICATIONS</i>	Number	1024	The maximum number of outstanding event notifications for each active subscriber.	No
<i>MAX_PENDING_IPMB_REQUESTS</i>	Number	64	The maximum number of pending IPMB requests awaiting response.	No

<i>MAX_SEL_ENTRIES</i>	Number	1024	The maximum number of entries in the System Event Log. CAUTION: Large values of this parameter can significantly degrade Shelf Manager performance.	No
<i>MAX_SESSIONS</i>	Number	32	The maximum number of simultaneous IPMI sessions. The maximum allowed value for this parameter is 63.	No
<i>MAX_USERS</i>	Number	32	The maximum number of IPMI users. The maximum allowed value for this parameter is 63.	No
<i>MICRO_TCA</i>	Boolean	FALSE	If <i>TRUE</i> , the Shelf Manager operates as a MicroTCA Shelf Manager (the second RMCP channel is used for interaction with Carrier Managers).	No
<i>MIN_FAN_LEVEL</i>	Number	1	The minimum fan level; the cooling management code does not reduce the fan level of any fan below this value when controlling the fan level automatically	Yes
<i>MIN_SHELF_FRUS</i>	Number	1	The minimum number of Shelf FRU Information instances in the shelf that the Shelf Manager must detect to start up successfully.	No

<i>NORMAL_STABLE_TIME</i>	Number	3600	The time in seconds for which the Shelf Manager preserves the minimum fan level dynamically found in Normal mode (that is, the minimum fan level that does not cause thermal alerts). After this time expires, the cooling algorithm decreases the minimum fan level, if possible, to allow the shelf to decrease the fan level if the thermal load in it has also decreased.	Yes
<i>NO_M0_M1_EVENT_AT_STARTUP</i>	Boolean	FALSE	If <i>TRUE</i> , the physical Shelf Manager does not send an M0->M1 hot swap event when the Shelf Manager application is restarted on the ShMM; the event is sent only the first time after the ShMM is plugged into a shelf. The default value is <i>FALSE</i> for backward compatibility.	No
<i>PEF_USE_KEYED_ALARMS</i>	Boolean	FALSE	If <i>TRUE</i> , the Shelf Manager uses special methods of managing TELCO alarms raised via PEF as OEM actions. Each event condition that causes a TELCO alarm is counted separately in this case, so that if a specific alarm is raised due to multiple event conditions, it stays raised while any of these conditions stay asserted.	No

<i>PET_FORMAT</i>	Number	0	<p>Specifies the format of the Platform Event Traps that are sent by the Shelf Manager as the Alert action initiated by event processing in the Platform Event Filtering facility. The values are defined as follows:</p> <ul style="list-style-type: none"> • 0 = the default IPMI format defined by IPMI Platform Event Trap Format v1.0 specification. • 1 = plain text format; all the event details are sent as plain ASCII text in a single variable. • 2 = multi-variable format; each event field is encoded as a separate variable defined in PPS-PET-MIB. 	Yes
<i>PET_OEM_WITH_SEVERITY_STRING</i>	Boolean	FALSE	If <i>TRUE</i> , the event severity and alert string is added to a trap message in the nVent OEM (Plain Text and Multi-Variable) PET formats. The default IPMI PET format always includes these fields from the corresponding event filter.	Yes
<i>PET_PORT</i>	Number	162	Specifies destination port number for the Platform Event Trap datagrams. The default UDP port number is 162.	Yes
<i>PET_SNMP_VERSION</i>	Number	1	Specifies the SNMP version for traps generated by the Shelf Manager. Currently supported versions are 1 (the default) and 2.	Yes
<i>POWER_UNLISTED_FRUS</i>	Boolean	TRUE	Allow the FRUs not listed in the power management table in the Shelf FRU Information to be activated and powered up.	No

<i>PREFER_CHASSIS_HPD_L_SDRS_FROM_FILE</i>	Boolean	FALSE	If <i>TRUE</i> , the preferred source of shelf HPDL configuration data is files on the ShMM, not records in Shelf FRU Information. Default is <i>FALSE</i> .	No
<i>PROPAGATE_RMCP_ADDRESS</i>	Boolean	FALSE	If <i>TRUE</i> , the active Shelf Manager propagates the RMCP IP address to the backup Shelf Manager, which configures the network interface specified by the <i>RMCP_NET_ADAPTER</i> variable using that IP address, but with the least significant bit inverted.	Yes
<i>REAPPLY_POWER_MAX_COOLING_STATE</i>	String(16)	"NORMAL"	Sets the most severe shelf thermal state in which the Shelf Manager can power a FRU up after that FRU has been powered down due to a Critical thermal alert. Possible values of this variable are <i>NORMAL</i> , <i>MINOR_ALERT</i> , <i>MAJOR_ALERT</i> . The <i>CRITICAL_ALERT</i> value is not supported, since it could lead to oscillations caused by immediate reactivation of the problematic FRU.	No

<i>REDUNDANCY_COMPRESSION_THRESHOLD</i>	Number	-1	<p>This parameter controls compression of messages sent over the redundancy interface. A message is sent in compressed format if its size exceeds the parameter value. The value <i>-1</i> turns compression off.</p> <p>Using message compression can improve the speed of redundancy information transfer on relatively slow redundancy links (like the serial interface used on ShMM-1500).</p>	No
<i>REDUNDANCY_ENABLED</i>	Boolean	TRUE	Run the Shelf Manager in redundant mode.	No
<i>REDUNDANCY_NET_ADAPTER</i>	String(16)	"eth0"	The name of the network adapter used for communication between redundant instances of the Shelf Manager. " <i>usb0</i> " is the recommended value if USB network interfaces are used for redundancy.	No
<i>REDUNDANCY_NET_ADAPTER2</i>	String(16)	"" (undefined)	<p>The name of the second network adapter used for communication between redundant instances of the Shelf Manager (if the dual USB network interface is used for this purpose). "<i>usb1</i>" is the recommended value if USB network interfaces are used for redundancy. If the configuration variable <i>REDUNDANCY_NET_ADAPTER</i> is set to "<i>usb0</i>", this variable defaults to "<i>usb1</i>".</p>	No

<i>REDUNDANCY_NETMASK</i>	IP-address	0.0.0.0	The netmask to assign to the redundancy IP address; by default (if 0), the netmask is determined automatically from the class of the IP address. <i>255.255.255.128</i> is the recommended value if USB network interfaces are used for redundancy.	No
<i>REDUNDANCY_PORT</i>	Number	1040	The TCP port used for interactions between redundant instances of the Shelf Manager.	No
<i>REDUNDANT_IP_ADDRESS</i>	IP-address	None	The IP address used for redundant communications. This address actually specifies a pair of IP addresses that differ only in the least significant bit. They are assigned to redundant Shelf Managers according to their hardware addresses.	No
<i>RESERVATION_RETRIES</i>	Number	10	The maximum number of times the Shelf Manager retries the “Reserve Device SDR” command.	Yes
<i>RMCP_NET_ADAPTER</i>	String(16)	“eth0:1”	The name of the network adapter used for RMCP-based communication.	Yes
<i>RMCP_NET_ADAPTER2</i>	String(16)	None	The name of the alternate network adapter used for RMCP-based communications, if cross-connect links are supported by the hardware. <i>eth1</i> is the recommended value if the parameter is specified.	Yes

<i>RMCP_PLUS_ONLY</i>	Boolean	FALSE	If this variable is <i>TRUE</i> , support of the regular RMCP protocol (from IPMI 1.5) is disabled on the Shelf Manager and only the RMCP+ protocol (from IPMI 2.0) is supported. (This mode may be needed for security reasons.)	No
<i>RMCP_WITHOUT_SHELF_FRU</i>	Boolean	FALSE	RMCP is available in the absence of the Shelf FRU Information only if this configuration parameter is defined and set to <i>TRUE</i> .	No
<i>RUN_HPI_SNMP_SUBAGENT</i>	Boolean	FALSE	Specifies whether the HPI SNMP subagent should be started automatically.	No
<i>SAVE_CRASH_INFORMATION</i>	Boolean	FALSE	If <i>TRUE</i> , the Shelf Manager saves crash information in the <i>/var/crash</i> directory in addition to setting a corresponding reboot reason. The following system signals are treated as indicators of a Shelf Manager crash: SIGSEGV, SIGABRT, SIGILL, SIGFPE, SIGBUS. The saved crash information includes a compressed system log and core file.	No
<i>SDR_READ_RETRIES</i>	Number	3	The maximum number of times the Shelf Manager retries the “Read Device SDR” command.	Yes
<i>SEL_FILE_COMPRESSION_ENABLED</i>	Boolean	TRUE	If <i>TRUE</i> , the Shelf Manager compresses the SEL before writing it to the ShMM Flash; this reduces the size of the file that is periodically written to the Flash and thus increases the lifetime of the Flash.	No

<i>SEL_FILE_JOURNALING_ENABLED</i>	Boolean	TRUE	If <i>TRUE</i> , the Shelf Manager uses journaling when writing the SEL to the ShMM Flash; this reduces the size of the data periodically written to the Flash and thus increases the lifetime of the Flash. The reduction is even greater if SEL compression is simultaneously enabled.	No
<i>SEL_FILE_WRITE_DELAY</i>	Number	3	The minimum number of seconds between subsequent updates of the SEL file on the ShMM Flash. Setting this number to a higher value causes SEL updates to be less frequent; this reduces the total amount of data written to the Flash and thus increases the lifetime of the Flash, but increases the probability of event loss if both redundant Shelf Managers go out of service.	No
<i>SEL_HIGH_WATERMARK</i>	Number	0	This value is the “high watermark” for the algorithm that controls automatic purging of the SEL; if the actual percentage of free entries in the SEL falls below this value, or the SEL overflows, the Shelf Manager starts a thread that purges old records from the SEL in order of decreasing age.	Yes
<i>SEL_LOW_WATERMARK</i>	Number	0	This value is the “low watermark” for the algorithm that controls automatic purging of the SEL; if the thread that purges old records from the SEL starts, it purges records until the percentage of occupied entries in the SEL falls below this value.	Yes

<i>SENSOR_POLL_INTERVAL</i>	Number	1	The time (in seconds) between successive polls of local Shelf Manager sensors by the Shelf Manager.	No
<i>SESSION_SEQUENCE_WINDOW</i>	Number	128	This is the window of acceptable RMCP sequence numbers; the wider this window is, the more tolerant is the Shelf Manager to RMCP packets being dropped during transfer. If the difference in the sequence numbers of a received packet and the previous packet exceeds the window size, the Shelf Manager closes the RMCP session as a corrupted one.	No
<i>SHELF_FRU_IN_EEPROM</i>	Boolean	TRUE	If <i>TRUE</i> , the Shelf FRU Information is retrieved from EEPROMs on the backplane in a carrier-specific way; if <i>FALSE</i> , the Shelf FRU Information is obtained from a file on the Flash file system.	No
<i>SHELF_FRU_IPMB_SOURCE1</i>	Number	0	If defined (non-zero), specifies the IPMB address of the first designated source of Shelf FRU Information in the shelf. (Shelf FRU Information is located at FRU 1.) If this value is defined, the search for the Shelf FRU Information on the IPMB is limited to the designated sources only.	No

<i>SHELF_FRU_IP MB_SOURCE2</i>	Number	0	If defined (non-zero), specifies the IPMB address of the second designated source of Shelf FRU Information in the shelf. (Shelf FRU Information is located at FRU 1.) If this value is defined, the search for the Shelf FRU Information on the IPMB is limited to the designated sources only.	No
<i>SHELF_FRU_TI MEOUT</i>	Number	5 or 20	The time interval (in seconds) during initialization that the Shelf Manager waits for Shelf FRU Information devices to be detected. If the configuration variable <i>SHELF_FRU_IN_EEPROM</i> is <i>TRUE</i> , the default value is 20. The actual time is set to 20 if the value of this parameter is less than 20.	No
<i>SHELF_MANAGER_ CONFIGURATION_IN_ SHELF_FRU_INFO</i>	Boolean	FALSE	If <i>TRUE</i> , instructs the Shelf Manager to retrieve configuration parameters from a special set of records in the Shelf FRU Information. If these parameters are successfully retrieved, they override the values from configuration files. Not all configuration variables can be overridden from the Shelf FRU Information; for example, this variable cannot be overridden (for obvious reasons) and, to be effective, must be specified in the configuration file.	No

<i>SHORT_SEND_MESSAGE_RESPONSE</i>	Boolean	TRUE	Determines the type of the “Send Message” response provided by the Shelf Manager: required by the current PICMG 3.0 (if <i>TRUE</i>) or compatible with the previous versions of the Shelf Manager (if <i>FALSE</i>).	No
<i>SPECIAL_HPI_ENTITY_FOR_AMC_CARRIER</i>	Boolean	TRUE	If <i>TRUE</i> , IntegralHPI uses the entity type <i>SAHPI_ENT_SUBBOARD_CARRIER_BLADE</i> for AMC carrier boards; otherwise the entity type <i>SAHPI_ENT_PICMG_FRONT_BLADE</i> is used.	No
<i>SWAPPED_CROSS_CONNECTS</i>	Boolean	FALSE	Swaps the names of network adapters used for cross-connects on the ShMC with the odd hardware address.	Yes
<i>SWITCHOVER_ON_BROKEN_LINK_BACKUP_DELAY</i>	Number	0	This parameter specifies the delay on the backup side during a switchover caused by a broken physical network link. When the active Shelf Manager requests the switchover, the backup Shelf Manager agrees to it (if its physical network link is present) only after this number of seconds passes, following the first request from the active Shelf Manager. A non-zero value for this parameter can be used to prevent switchovers in scenarios where a hub board is inserted into the shelf that activates the active Shelf Manager link several seconds after the backup Shelf Manager link.	Yes

<i>SWITCHOVER_ON_HANDLE_OPEN</i>	Boolean	FALSE	<p>If <i>TRUE</i>, switchover-related behavior of the active Shelf Manager is affected by the Hot Swap state of its physical IPM controller, as follows:</p> <ul style="list-style-type: none"> - If the physical IPM controller on the active Shelf Manager goes to state M1 due to its Hot Swap Handle being open or due to programmatic deactivation, a switchover to the backup Shelf Manager is initiated; - If the physical IPM controller on the active Shelf Manager goes to state M5 and there is no available backup Shelf Manager, the physical IPM controller is not deactivated and stays in M5 indefinitely. 	Yes
<i>SWITCHOVER_TIMEOUT_ON_BROKEN_LINK</i>	Number	-1	<p>This parameter affects when or whether the Shelf Manager initiates a switchover when the physical network link between the Shelf Manager and the System Manager (the RMCP link) is broken. If the link remains broken for at least the number of seconds given in this parameter, a switchover takes place; if the link is restored during this timeout period, no switchover takes place. If the value of this parameter is <i>-1</i>, no automatic switchovers take place on broken RMCP links.</p>	Yes
<i>SYSLOG_LOGGING_ENABLED</i>	Boolean	TRUE	Output log messages to the system log.	No

<i>SYSTEM_MANAGER_TRUNCATES_SEL</i>	Boolean	FALSE	If <i>TRUE</i> , the Shelf Manager algorithm for truncating the SEL automatically is disabled; the System Manager is responsible for truncating the SEL by monitoring the value of the sensor SEL State of the type Event Logging Disabled on the Shelf Manager, and removing events from the SEL by sending the “Delete SEL Entry” command to the Shelf Manager.	No
<i>TACHOMETER_THRESHOLD_UPDATE_DELAY</i>	Number	15	This parameter applies to fan trays managed by the Shelf Manager and implementing adaptive setting of tachometer thresholds depending on the current fan level. This parameter specifies the delay in seconds between setting a new fan level that is higher than the current fan level and increasing the tachometer threshold value; the delay is needed to allow the fans to reach the higher speed.	Yes
<i>TASKLET_RETRIES</i>	Number	3	This parameter specifies the total number of Shelf Manager retries to execute a tasklet (activation, deactivation, getting information, etc.) if it encounters a transient failure when doing that. If the Shelf manager fails to perform a tasklet after the <i>TASKLET_RETRIES</i> retries, it gives up. Increasing this value may improve the robustness of the Shelf Manager operation, but degrade its performance.	Yes

<i>TIMEPROTO</i>	String(16)	""	Protocol used to retrieve time from a network time server; possible values are <i>ntp</i> and <i>rdate</i> . The default value is an empty string which makes the script to use <i>rdate</i> .	Yes
<i>TIMESERVER</i>	IP-address	None	IP address of the time server for synchronization at runtime. If this variable is not specified, time is extracted from the hardware clock at shelf startup.	Yes
<i>TURBO_MODE_M IN_FAN_FAILU RES</i>	Number	1	This parameter specifies the number of tachometer underspeed faults (crossing Major or Critical thresholds) cause the remaining fans to go to TURBO mode (that is, run at full speed) in the default HPDL cooling management algorithm. If the shelf implements zoned cooling, the specified number of tachometer faults is counted relative to each zone and setting TURBO mode affects only the fan trays that participate in cooling of the corresponding zone.	Yes

<i>TURBO_MODE_MIN_MISSING_FAN_TRAYS</i>	Number	1	<p>This parameter specifies the minimum number of missing/inactive fan trays that triggers the remaining fan trays to go to TURBO mode (that is, run at full speed) in the default HPDL cooling management algorithm. If this parameter is set to 0, this feature of setting fans to their maximum speed (due to missing/inactive fan trays) is disabled, entirely.</p> <p>If the shelf implements zoned cooling, the specified number of missing/inactive fan trays is counted relative to each zone and setting TURBO mode affects only the fan trays that participate in cooling of the corresponding zone</p>	Yes
<i>UNCONDITIONAL_SDR_REREAD_ON_VERSION_CHANGE</i>	Boolean	FALSE	<p>If the value of this parameter is <i>TRUE</i> the Shelf Manager unconditionally re-reads SDRs from an IPM controller when it receives a Version Change event from that controller, even if the Sensor Population Change Indicator in the “Get Device SDR Info” response does not change. This is done for the benefit of ATCA boards that are not fully IPMI-compliant in this respect.</p>	Yes
<i>USE_DHCP</i>	Boolean	FALSE	<p>Requests assignment of RMCP-accessible and private IP addresses for the Shelf Manager from a DHCP server; the configuration parameter <i>DHCP_SERVER_ADDRESS</i> can be used to specify the IP address of the DHCP server.</p>	Yes

<i>USE_DHCPV6</i>	Boolean	FALSE	ShMM-700 only: Requests assignment of IPv6 addresses for the Shelf Manager from a DHCPv6 server; the configuration parameter <i>DHCPV6_SERVER_ADDRESSES</i> can be used to specify the IPv6 addresses of the allowed DHCPv6 servers.	Yes
<i>USE_IPMB_A_ONLY</i>	Boolean	FALSE	This parameter allows using <i>IPMB-A</i> only at systems with one management bus. By default, both <i>IPMB-A</i> and <i>IPMB-B</i> are used.	No
<i>USE_SECOND_CHANNEL</i>	Boolean	FALSE	This parameter applies only if two network interfaces on the ShMM are used for RMCP communication. If <i>TRUE</i> , the two network interfaces on the ShMM are used in parallel mode; if <i>FALSE</i> , they are used in redundant mode.	Yes
<i>VERBOSITY</i>	Number	7	The Shelf Manager verbosity level.	No
<i>VERBOSITY_CONSOLE</i>	Number	N/A	The Shelf Manager verbosity level for console output.	No
<i>VERIFY_SHELF_FRU_CHECKSUM</i>	Boolean	TRUE	Enable verification of checksums in Shelf FRU Information records; if set to <i>FALSE</i> , Shelf Manager ignores checksums	No
<i>WATCHDOG_ENABLED</i>	Boolean	TRUE	This configuration parameter is used to enable or disable the hardware watchdog timer. When enabled, the watchdog will cause the Shelf Manager to reboot if the watchdog expires.	No

<i>WATCHDOG_TIMEOUT</i>	Number	2	<p>The hardware watchdog timeout, in seconds. If the software does not strobe the watchdog at least once during this period, the hardware resets the ShMM. The supported range for this parameter is 2..30 seconds, which is valid for all supported ShMM platforms.</p> <p>Increasing the default value may be necessary if additional software besides the Shelf Manager application runs on the ShMM, increasing the CPU load. Also see section 4.4.3 for the possible reasons to increase the watchdog timeout.</p> <p>The downside of increasing the watchdog timeout is that the length of time to switch over to the backup ShMM will be also be increased in the event of a watchdog expiration.</p>	No
-------------------------	--------	---	--	----

By default, the configuration file variables are used automatically when the ShMM is brought up for the first time. The default configuration file imports several environment variables set by U-Boot.

Table 8: U-Boot environment variables and descriptions

U-BOOT ENVIRONMENT VARIABLE	DESCRIPTION
<i>\$IPADDR</i>	Default RMCP IP Address
<i>\$IPDEVICE</i>	Default RMCP network adapter
<i>\$IP1ADDR</i>	Default Redundant IP Address
<i>\$IP1DEVICE</i>	Default Redundant network adapter
<i>\$GATEWAY</i>	Default gateway used for RMCP communication

The environment variables *\$CARRIER* and *\$CARRIER_OPTIONS* are set by the secondary RC script. The name of this carrier-specific startup script is defined by the U-Boot environment variable *rc2*. If needed, the U-Boot environment variables can be reset to default values, refer to section 6.2.

3.3.1 Carrier-specific configuration file

After reading the common configuration file `/etc/shelfman.conf`, the Shelf Manager reads the carrier-specific configuration file `/etc/shelfman.conf.<carrier-name>`, where `<carrier-name>` is the name of the ShMM carrier used in the relevant shelf, in lower case characters. Settings in the carrier-specific configuration file override settings for the same variable in the common configuration file.

This mechanism allows redefinition of common settings on a carrier-specific basis. Typically, only a few critical configuration variables are defined in the carrier-specific file. For instance, the appropriate value for the `MIN_FAN_LEVEL` parameter may well be determined by the shelf architecture and the fan facilities that it implements. This mechanism allows such shelf-specific constraints to be enforced.

One result of this mechanism is that to change the effective value of a configuration parameter that is specified in the carrier-specific configuration file, the change must be made in that configuration file. A change for such a variable in the common configuration file will not have any effect.

3.3.2 Obtaining configuration variables from Shelf FRU Information

Besides carrier-specific redefinition of configuration variables, there is another mechanism for retrieving and applying configuration variables from the Shelf FRU Information. This mechanism allows redefining certain configuration variables for a specific instance or specific type of shelf.

Configuration variables are stored in the FRU Information as a sequence of OEM-type multirecords, similar to the multirecords defined in the xTCA specifications. These multirecords have the following nVent, Schroff GmbH (formerly Pentair Technical Products) Manufacturer ID (00400Ah) and the manufacturer-specific record type is 10. The concatenated contents of these multirecords represent the configuration variables in the same text format as in configuration files, but optionally compressed using *gzip*.

Some restrictions apply to the configuration variables in the Shelf FRU Information. First, not all variables can be redefined from the Shelf FRU Information. Table 7 indicates, for each variable, whether it can be redefined from the Shelf FRU Information. Second, retrieval of configuration variables from the Shelf FRU Information must be explicitly enabled for the Shelf Manager via the configuration variable `SHELF_MANAGER_CONFIGURATION_IN_SHELF_FRU_INFO`. By default, this variable is set to `FALSE`. This is done to allow the Shelf Manager to start in the case of a bad configuration stored in the Shelf FRU Information, and for compatibility with previous releases of the Shelf Manager.

3.3.2.1 Using the FRU Info Compiler to place configuration variables in the Shelf FRU Information

The FRU Information Compiler is a Pigeon Point tool that generates a FRU Information image from its textual description. The image can then be stored in an appropriate EEPROM device using various techniques. The FRU Information Compiler implements special syntax that can be used to incorporate a configuration file into a FRU Information image that is being created.

For example, the following lines add a compressed configuration file `shelfman.conf.acb3.gz` to the current FRU Information image:

...

[nVent.Shelf Manager Configuration]

File = shelfman.conf.acb3.gz

3.3.2.2 Using the Shelf Manager Command Line Interface to place configuration variables in the Shelf FRU Information

An extended version of the CLI command *frudataw* can be used to dynamically update Shelf FRU Information with configuration information. The corresponding configuration file (optionally compressed) must be first downloaded to the ShMM.

The corresponding command line has the following syntax:

```
clia frudataw -p <ipmb-address> <fru-id> <in-file>
```

```
clia frudataw -p -c <ipmb-address> <fru-id>
```

The parameters *<ipmb-address>* and *<fru-id>* specify the IPMB address and FRU device ID of the FRU Information to update. For the Shelf FRU Information, use IPMB address 20h and FRU device ID 254.

The parameter *<in-file>* indicates the configuration file (which is possibly compressed) to be stored in the target FRU Information.

The option *-c* removes (clears) the configuration data from the target FRU Information.

For example, the following command can be used to update the Shelf FRU Information with a new version of the configuration file:

```
#clia frudataw -p 20 254 shelfman.conf.acb3.gz
```

3.3.3 Verbosity level description

This parameter controls the verbosity of output sent by the Shelf Manager to the system log. This parameter also controls the verbosity of output to the console, if the configuration parameter *VERBOSITY_CONSOLE* is omitted. The *VERBOSITY* configuration parameter is a hexadecimal bit mask, with each bit enabling output of a specific type of message:

Table 9: Verbosity configuration parameters and levels

VERBOSITY CONFIGURATION PARAMETER	VERBOSITY LEVEL
0x01	Error messages
0x02	Warning messages
0x04	Informational messages
0x08	Verbose informational messages
0x10	Trace messages (not recommended)
0x20	Verbose trace messages (not recommended)
0x40	Messages displayed for important commands sent to the IPM Controllers during their initialization (not

VERBOSITY CONFIGURATION PARAMETER	VERBOSITY LEVEL
	recommended)
0x80	Verbose messages about acquiring and releasing internal locks (not recommended)

The default debug level is 0x07, enabling error, warning and informational messages.

3.3.4 Verbosity console level description

This parameter controls the verbosity of output sent by the Shelf Manager to the console. The *VERBOSITY_CONSOLE* configuration parameter is a hexadecimal bit mask, with each bit enabling output of a specific type of message, according to Table 9.

If this parameter is omitted, the value of the *VERBOSITY* configuration parameter is used to control the verbosity of output to both the system log and the console.

3.4 Setting up Ethernet

The ShMM uses two Ethernet ports, one of them being used for shelf-external access. Since RMCP is the only shelf-external interface that is required by ATCA, the shelf-external Ethernet port is referenced as the RMCP port, though the other shelf-external interfaces (HTTP, TELNET, FTP) are accessible via this port as well. The other Ethernet port can be used for communication between redundant Shelf Managers, or potentially for other purposes if the redundancy link is provided by another mechanism.

Note that manual changes in the ShMM network configuration that are initiated separately from the Shelf Manager (e.g. via the *ifconfig* command) can disrupt Shelf Manager operations. The implications of any such changes must be thoroughly understood before they are undertaken.

3.4.1 Usage of the first Ethernet interface

Since the RMCP Ethernet port is directly connected to the site network, the IP address should be set up appropriately for that network.

For example, if the site uses the IP address range 192.168.0.x, the RMCP Ethernet port should be set to a unique IP address within that range such as 192.168.0.2. In a redundant ShMM setup, it is important to note that only one ShMM (the active ShMM) has the RMCP IP address enabled on the RMCP Ethernet port. The backup ShMM assigns the same IP address to the RMCP Ethernet port, but only enables it when that ShMM assumes the active role. This way, the RMCP IP address maintains availability in a failover situation. If the configuration variable *PROPAGATE_RMCP_ADDRESS* is set to *TRUE*; in that case the backup ShMM assigns a conjugate to the RMCP IP address to its RMCP Net Adapter (that is, it assigns an address that is different in the least significant bit of the fourth byte).

3.4.1.1 Assigning an additional IP address to the first network interface

In the default configuration, no IP address is assigned to the first network interface (and the ShMM is not accessible over the network) until the Shelf Manager starts and the RMCP IP address is assigned. However, it may be useful in some cases to assign an IP address to the RMCP network interface and have the ShMM accessible over the network as soon as the operating system is

booted. In that case, it is also desirable that when the Shelf Manager is started, for the RMCP IP address to coexist with the originally assigned IP address rather than replacing it.

To achieve this configuration, it is necessary to instruct the Shelf Manager to assign the RMCP IP address not to the first network adapter itself (*eth0*) but to its first alias (*eth0:1*). The initial IP address is assigned in that case to the network adapter itself (*eth0*) during the start of the operating system. This initial assignment happens in the initialization script */etc/rc*; it is accomplished by: 1) enabling the U-Boot variable *rc_ifconfig* (setting it to *y*), 2) assigning the original IP address to the U-Boot variable *ipaddr*, for example:

```
shmmx500 setenv rc_ifconfig y
shmmx500 setenv ipaddr 192.168.1.240
```

and 3) changing the value of the Shelf Manager configuration parameter *RMCP_NET_ADAPTER* to *eth0:1*.

In a redundant configuration, the U-Boot variable *ipaddr* is allowed to have the same value on both ShMMs. The actual initial IP address assigned to each of the two redundant ShMMs is based on the value of *ipaddr* but is modified depending on the hardware address of the ShMM. The least significant bit of the IP address is set to the least significant bit of the hardware address. In the example above, the IP address is 192.168.1.240 for the ShMM with an even hardware address, and is 192.168.1.241 for the ShMM with an odd hardware address. This modification of the IP address can be turned off by removing the file */etc/readhwaddr*.

3.4.1.2 RMCP address propagation

There is an optional feature of the Shelf Manager that allows the backup ShMM also to be exposed on the external network with an IP address that is different from the RMCP IP address only in the least significant bit. The netmask and default gateway on the backup ShMM is the same as on the active ShMM. For example, if the RMCP IP address is 192.168.0.2, the backup ShMM has the corresponding IP address 192.168.0.3, with the same netmask and default gateway. To enable this feature, it is necessary to define the Shelf Manager configuration parameter *PROPAGATE_RMCP_ADDRESS* as *TRUE* in the Shelf Manager configuration file.

3.4.2 Usage of the second Ethernet interface

The second Ethernet interface can be dedicated for use as a private network between redundant ShMMs and used to synchronize state information between the active and backup ShMMs.

Unlike the RMCP Ethernet port, this second Ethernet interface is always enabled on both the active and backup ShMM, but with a small twist – both the active and backup ShMM specify the same IP address for the redundancy interface, but software assigns the next logical IP address to the ShMM with an odd hardware address. For instance, the default setting for the redundancy Ethernet port is 192.168.1.2. The odd-addressed ShMM assumes the address 192.168.1.3. This way the active and backup ShMM can be identically configured but still assume unique IP addresses for the redundancy Ethernet link.

When the ShMMs use a non-Ethernet network interface for communication between the redundant Shelf Managers (see section 3.4.5), the second network interface can potentially be used for other purposes, such as for support in the Shelf Manager for ShMC cross connects, in

accordance with the PICMG ECN 3.0-2.0-001. In that case, the second network interface connects the Shelf Manager with one of the ATCA network hub boards.

To configure the Shelf Manager to support cross-connects, it is necessary to define the configuration parameter *RMCP_NET_ADAPTER2*. In that case, if the backplane and hub boards also support cross-connects, the Shelf Manager uses the two network adapters (*RMCP_NET_ADAPTER* and *RMCP_NET_ADAPTER2*) for RMCP communication. Three usage models are available for this feature:

- active-standby,
- active-active, and
- bonded.

3.4.2.1 Active-standby usage of the two network interfaces

The two network interfaces can be used in an active-standby way. In this mode, at any given time, RMCP communication passes through only one adapter (initially this is *RMCP_NET_ADAPTER*). However, if the Shelf Manager detects that the adapter currently used for RMCP communication becomes physically disconnected from the network (link broken), it automatically switches to the other (alternate) adapter. The first adapter is turned off, the RMCP IP address is transparently moved to the other adapter, and three ARP notifications are broadcast to notify other systems about the MAC address change. This change is transparent for the System Manager and does not break existing RMCP connections.

However, if the Shelf Manager detects that the other network adapter is also physically disconnected from the network, it does not perform the IP address switchover described above, but performs a full switchover to the backup Shelf Manager. A full switchover is also performed if the Shelf Manager detects physical disconnection of the network adapter used for RMCP communication, in non cross-connect configurations.

Detection of physical disconnection of the RMCP network adapter is controlled by the Shelf Manager configuration parameter *SWITCHOVER_TIMEOUT_ON_BROKEN_LINK*. The value of this parameter is the time interval in seconds during which the adapter stays physically disconnected, before the Shelf Manager performs an IP address switchover or full switchover. Detection and switchover is disabled if the value of this parameter is equal to *-1*.

Another configuration parameter, *INITIAL_SLOW_LINK_DELAY*, specifies the time interval from the start of the Shelf Manager, during which detection of physical disconnection is not performed. This allows using Ethernet links that are slow to start and need some time after shelf power up to establish the physical connection.

In the case of a Shelf Manager switchover (as opposed to the IP address switchover mentioned above) caused by a broken link, a special situation occurs when the first hub board is inserted into a shelf with no hub boards and cross-connect Shelf Manager links.

From the physical connection perspective, this is a transition from a situation where both physical network links were broken on both Shelf Managers, to a situation where at least one physical network link is present on both Shelf Managers. However, if in this case, due to the specifics of the hub board, the link to the backup Shelf Manager comes up earlier than the link to the active Shelf Manager, an undesirable switchover may occur. To prevent a switchover in this situation, the

backup Shelf Manager can delay (for a specified number of seconds) doing a switchover after the original switchover request from the active Shelf Manager, if the physical link to the backup Shelf Manager is present. If, during this time, the physical link to the active Shelf Manager also comes up, the active Shelf Manager will stop sending switchover requests, and the switchover won't happen. Otherwise, if the physical link to the active Shelf Manager remains broken after the expiration of this delay, the switchover will finally take place. The value of the delay is specified in the configuration variable *SWITCHOVER_ON_BROKEN_LINK_BACKUP_DELAY*.

Usually the parameter *RMCP_NET_ADAPTER2* is assigned the value *eth1*, while the value of the configuration parameter *RMCP_NET_ADAPTER* (the main network adapter used for RMCP communication) is *eth0*, as in the following sample.

```
RMCP_NET_ADAPTER = "eth0"  
RMCP_NET_ADAPTER2 = "eth1"
```

However, other configurations are possible. For example, values *eth0:1* and *eth1:1* can be used if additional (permanent) IP addresses need to be assigned to both network interfaces.

The two network interfaces are used in active-standby mode if the configuration parameter *USE_SECOND_CHANNEL* is set to *FALSE*.

If the Shelf Manager configuration parameter *PROPAGATE_RMCP_ADDRESS* is defined as *TRUE*, active-standby management also applies to the network interface on the backup Shelf Manager, which is assigned the RMCP-derived IP address on the backup Shelf Manager. That is, if the backup Shelf Manager detects that the currently used adapter becomes physically disconnected from the network (link broken), it automatically switches to the other (alternate) adapter. The first adapter is turned off and the RMCP-derived IP address of the backup Shelf Manager is transparently moved to the other adapter.

3.4.2.2 Active-active usage of the two network interfaces

The approach to network interface redundancy outlined in the previous section seems to be insufficient for some configurations where the network connection between the Shelf Manager and the System Manager goes through several switches and may break on an Ethernet segment that is not adjacent to the ShMM.

This type of failure cannot be immediately recognized by the Shelf Manager. Checking the accessibility of the System Manager from the Shelf Manager does not seem practical in this case, since the architecture of the System Manager and its usage of IP addresses is not defined in the ATCA specification and the Shelf Manager design should not artificially limit it. The solution in this case should be implemented at the System Manager level.

Some new features introduced in ECN-002 to the ATCA specification (PICMG 3.0 R2.0) facilitate a solution. The command “Get Shelf Manager IP Addresses” allows the System Manager to retrieve the IP addresses exposed by the Shelf Managers, both active and backup. Using this information, the System Manager can check the availability of the Shelf Managers and the current distribution of responsibilities between them (active vs. backup).

In active-active mode, instead of having a single RMCP network address that is switched between the two network interfaces (see description of active-standby mode in 3.4.2.1), the Shelf Manager

supports RMCP on both interfaces with different IP addresses, as two separate IPMI channels (channels 1 and 2).

Each channel has its own set of LAN configuration parameters that includes the IP address, network mask, default gateway, etc. Both addresses are available via the “Get Shelf Manager IP Addresses” command.

The System Manager, in this case, is responsible for switching over to a different IP address if the currently used IP address becomes unavailable, or maintaining two parallel RMCP sessions to both addresses. If IP addresses on both interfaces become unavailable, the System Manager can access one of the IP addresses on the backup Shelf Manager and initiate a switchover in a non-standard way (for example, log in to the backup ShMM via *telnet* and issue the command *clia switchover*).

The active-active mode is enabled by setting the configuration parameter `USE_SECOND_CHANNEL` to `TRUE`. In this case, the configuration parameter `DEFAULT_RMCP_IP_ADDRESS2` should be set to a non-zero IP address. This IP address becomes the default IP address for the channel 2, and the configuration variables `DEFAULT_RMCP_NETMASK2` and `DEFAULT_GATEWAY_IP_ADDRESS2` specify the network mask and default gateway IP address for the channel 2, respectively.

The values of the above-mentioned configuration parameters have no effect if the corresponding values in the channel configuration parameter file are non-zero. The channel configuration parameters for the channel 2 are stored in the files `/var/nvdata/ch2_param` and `/var/nvdata/ch2_param_v2` on the ShMM (the second file is used only with the Shelf Manager release 2.6 and higher).

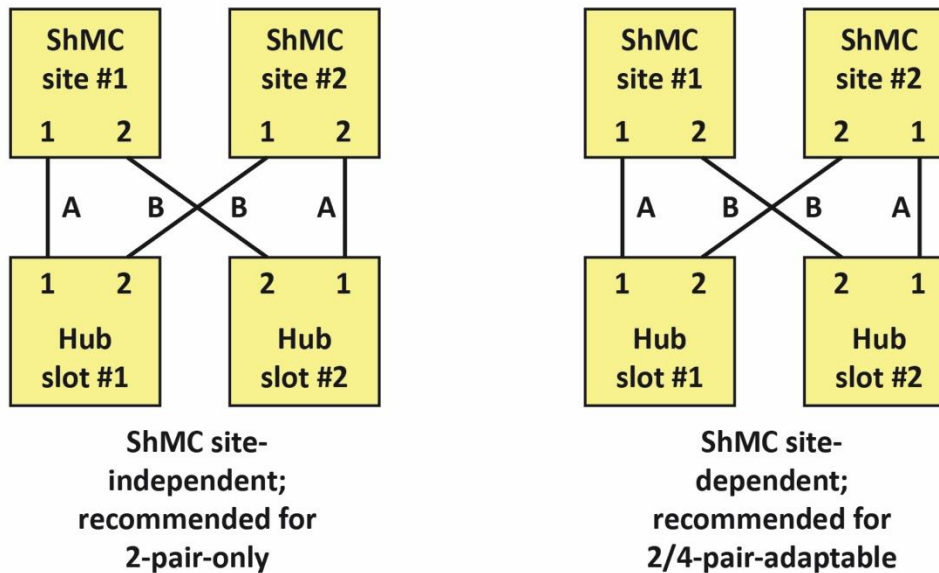
Starting with release 2.5.2, the Shelf Manager recognizes a second instance of the Shelf Manager IP Connection record in the Shelf FRU Information as the source of the IP address, netmask and default gateway for the second RMCP network interface. Therefore, it is possible to specify RMCP addressing information for both interfaces in the Shelf FRU Information.

In active-active mode, the IP addresses on both network interfaces are switched over to the backup ShMM as a result of a switchover.

3.4.2.3 Support for site-dependent ShMC cross-connects

According to the ECN 3.0-2.0-002 for revision 2.0 of the ATCA specification, ShMC cross-connects can be implemented with either a site independent or a site dependent configuration, as shown in Figure 3.

Figure 3: Implementation options for ShMC cross-connects



In the site-dependent variant, after a switchover in the default Shelf Manager configuration, each specific RMCP address is propagated to the corresponding network interface (eth0 or eth1) on the other ShMC, and thus becomes associated with a different hub board. In this case, from the perspective of the hub board, the ShMC switchover becomes non-transparent, which may be undesirable if, for example, each hub board hosts a System Manager. The configuration parameter *SWAPPED_CROSS_CONNECTS* can be used to avoid this non-transparency.

If the configuration parameter *SWAPPED_CROSS_CONNECTS* is defined and set to *TRUE* and if the hardware address of the ShMC is odd, the Shelf Manager swaps the values of the configuration parameters *RMCP_NET_ADAPTER* and *RMCP_NET_ADAPTER2*. As a result, after a switchover, the RMCP addresses on the newly active Shelf Manager correspond to the same hub boards as before the switchover.

3.4.2.4 Bonded usage of the two network interfaces

The two physical network interfaces (*eth0* and *eth1*) can be used in a bonded fashion using the kernel bonding driver that joins several physical devices into one logical device and selects an active device for data transfers. In this mode, both physical network interfaces that can be used for outside connections are enslaved by the bonding driver at system startup and the Shelf Manager uses the single bonding interface for RMCP-based communication. This mode is useful when both devices are connected to the same network switch/hub and it does not matter which device is used at any given moment, as long as the Shelf Manager remains available via the network.

The bonding driver in the kernel is configured to use an active-backup policy, which means that only one slave device in the bond is active. A different slave device becomes active only if the active slave device fails. The bond's MAC address is externally visible on only one network adapter port to avoid problems with switch/hub hardware. Link status of the slave devices in the bond is polled every 100 milliseconds, which is the recommended value. Also, in the case of an active device switch, the bonding driver sends out a gratuitous ARP packet to notify switch hardware about the port change, since bonding slave devices share the same MAC address taken from the

first device in the slave device list. When a network link loss is detected, active device switching is performed only after 500 milliseconds to avoid potential problems with short-term link loss. The idea is similar to the Active-Standby usage of two network interfaces except that device link state monitoring and switching is performed at low level by the kernel instead of the Shelf Manager.

The bonding mode is activated by changing the *ipdevice* variable value in U-Boot to *bond0*. This has to be done at the U-Boot level so that the system startup scripts perform the device initialization for *bond0* before the Shelf Manager starts. By default, this value is passed to the Shelf Manager as the default RMCP network adapter via the setting in */etc/shelfman.conf*:

```
RMCP_NET_ADAPTER = $IPDEVICE
```

It is possible to use an alias interface of *bond0* if it is necessary to have a permanent IP address on the ShMM while the RMCP address floats between active and standby ShMMs i.e.

```
RMCP_NET_ADAPTER = "bond0:1"
```

The parameter *RMCP_NET_ADAPTER2* must not be set while using the bonding interface as an RMCP adapter. The bonding mode assumes single RMCP network adapter usage from the Shelf Manager perspective, so the parameter *USE_SECOND_CHANNEL* must be set to *FALSE*.

3.4.3 Setting up the HPM.2 client channel

The logical Shelf Manager supports RMCP client connectivity with intelligent FRUs in the shelf (IPMCs, AMCs and intelligent RTMs), as described in the HPM.2 specification. A separate IPMI LAN-based channel (channel #8) is allocated on the logical Shelf Manager for that purpose.

This channel can be used in one of the two modes: mirrored or standalone. The mode is specified by the value of the Shelf Manager configuration parameter *HPM2_CLIENT_MIRROR_CHANNEL*. The value *1* or *2* specifies the mirrored mode; the value *0* specifies the standalone mode.

In the mirrored mode, the channel is “mirrored” on an existing RMCP server-mode LAN channel (the mandatory channel #1 or the optional channel #2). This means that the HPM.2 client channel shares the network adapter and LAN configuration parameters (the IP address, the subnet mask, etc.) with the original channel. These parameters can be read via the command “Get LAN Configuration Parameters” through the HPM.2 client channel or the original channel. They can be changed by sending the command “Set LAN Configuration Parameters” to the original channel. An attempt to change them through the HPM.2 client channel will fail. This mode can be used for shelves in which both the outside network and the HPM.2 shelf-internal network are available through the same network interface. The default mode for the HPM.2 client channel is mirroring it to channel #1.

In the standalone mode, a separate set of LAN configuration parameters and a separate network interface are assigned to the HPM.2 client channel. The configuration parameters can be managed via the commands “Get LAN Configuration Parameters” and “Set LAN Configuration Parameters” directed to the HPM.2 client channel. A set of Shelf Manager configuration parameters is defined that specify the network adapter name and initial values of certain LAN configuration parameters. In the standalone mode, values must be assigned to all the parameters in the following list:

- *HPM2_NET_ADAPTER*
- *DEFAULT_HPM2_CLIENT_IP_ADDRESS*

- `DEFAULT_HPM2_NETMASK`
- `DEFAULT_HPM2_GATEWAY_IP_ADDRESS`

This mode can be used for shelves in which the outside network and the HPM.2 shelf-internal network are available through different network interfaces.

If the HPM.2 client channel is configured properly, the Shelf Manager attempts to establish an RMCP session with each HPM.2-capable IPMC, when it first encounters it in the shelf, and with each HPM.2-capable AMC or intelligent RTM, when the corresponding FRU is activated. The session is maintained by the Shelf Manager while the corresponding FRU is present in the shelf, and is attempted to be reestablished if dropped by the controller representing the FRU. The CLI commands `ipmc` and `fru` in verbose mode can be used to determine if an HPM.2 client session to a specific module currently exists. For example, the following output indicates that HPM.2 client sessions are established to the IPMC at 84h and to MMC on AMC #2, represented by that same IPMC:

```
# clia ipmc -v 84
```

```
Pigeon Point Shelf Manager Command Line Interpreter
```

```
84: Entity: (0xa0, 0x60) Maximum FRU device ID: 0x04
```

```
    PICMG Version 2.3
```

```
    Hot Swap State: M4 (Active), Previous: M3 (Activation In Process), Last
    State Change Cause: Normal State Change (0x0)
```

```
    Device ID: 0x12, Revision: 0, Firmware: 1.20 (ver 1.2.0), IPMI ver 2.0
```

```
    Manufacturer ID: 00400a, Product ID: bada, Auxiliary Rev: 00000000
```

```
    Device ID String: "BMR-A2F-ATCA-BTR"
```

```
    Global Initialization: 0xc, Power State Notification: 0xc, Device
    Capabilities: 0x2d
```

```
    Controller provides Device SDRs
```

```
    Supported features: 0x2d
```

```
    "Sensor Device" "SEL Device" "FRU Inventory Device" "IPMB Event
    Generator"
```

```
    HPM.1 Components:
```

```
        0: "A2F-AMCc F/W": Version 1.20, Aux 00000000
```

```
        1: "A2F-AMCc B/L": Version 1.20, Aux 00000000
```

```
        2: "A2F-AMCc F/I": Version 1.20, Aux 00000000
```

```
        4: "A2F-AMCc IAP": Version 2.30, Aux 00000000
```

```
    HPM.2 Session: 2; Channel: 5; State: Established
```

```
    HPM.2 Capabilities:
```

```
        SOL Extensions
```

```
        Extended Inactive State Management: 20W, ON
```

```
        Version Change Sensor for LAN Configuration
```

```
        Dynamic Sessions
```

```
    HPM.2 Channels: 5
```

```
    HPM.2 LAN Parameters: start at 192 (0xc0), revision 1
```

```
    HPM.2 SOL Parameters: start at 192 (0xc0), revision 1
```

```
    HPM.3 Capable: Parameters: start at 205 (0xcd), revision 1
```

Links:

```
84: Base Interface (0x00), Channel: 1
    Link: Disabled Ports: 1
84: Base Interface (0x00), Channel: 2
    Link: Disabled Ports: 1
84: Fabric Interface (0x01), Channel: 1
    Link: Disabled Ports: 1
84: Fabric Interface (0x01), Channel: 2
    Link: Disabled Ports: 1
84: Update Channel Interface (0x02), Channel: 1
    Link: Disabled Ports: 1
```

clia fru -v 84 4

Pigeon Point Shelf Manager Command Line Interpreter

84: FRU # 4 (AMC # 2)

Entity: (0xc1, 0x62)

Hot Swap State: M4 (Active), Previous: M3 (Activation In Process), Last State Change Cause: Normal State Change (0x0)

Device Type: "FRU Inventory Device behind management controller" (0x10), Modifier 0x0

Device ID String: "BMR-A2F-AMCm BTR"

Current Power Level: 0x01, Maximum Power Level: 0x01, Current Power Allocation: 42.4 Watts

Device ID: 0x34, Revision: 0, Firmware: 1.20 (ver 1.2.0), IPMI ver 2.0

Manufacturer ID: 00400a, Product ID: beef, Auxiliary Rev: 00000000

HPM.2 Session: 2; Channel: 5; State: Established

HPM.2 Capabilities:

SOL Extensions

Extended Inactive State Management

Version Change Sensor for LAN Configuration

Dynamic Sessions

HPM.2 Channels: 5

HPM.2 LAN Parameters: start at 192 (0xc0), revision 1

HPM.2 SOL Parameters: start at 192 (0xc0), revision 1

HPM.3 Capable: Parameters: start at 205 (0xcd), revision 1

Established HPM.2 sessions are used by the Shelf Manager in the following scenarios:

- CLI commands *sendcmd* and *sendamc* with the option *-x* are used to send IPMI commands to the corresponding targets over an HPM.2 session, if one exists; the requests and responses in this case can be significantly longer than when the command is sent over IPMB. The length of a request can be up to 150 bytes and a response can be up to 248 bytes in *sendcmd* and up to 241 bytes in *sendamc*.

- In IntegralHPI, populating an Inventory from the FRU Information uses an HPM.2 session if it exists for the target FRU; in that case, the chunk size for a single read operation is 128 bytes; otherwise, it is 16 bytes.
- The CLI command `fruinfo` uses an HPM.2 session if it exists for the target FRU; in that case, the chunk size for a single read operation is 128 bytes; otherwise, it is 16 bytes.
- Some (not all) internal Shelf Manager requests to read a FRU Information record use an HPM.2 session if it exists for the target FRU; in that case, the chunk size for a single read operation is 128 bytes; otherwise, it is 16 bytes.
- The Shelf Manager code that reads Device SDRs from IPMCs uses an HPM.2 session if it exists for the target IPMC. In that case, each SDR is read in a single read operation; by default, each SDR is read by several partial read operations.

3.4.4 Gigabit Ethernet Support on ShMM-700

Starting from release 3.7.2, the Shelf Manager supports Gigabit Ethernet connections using one or two Microchip KSZ9563RNX devices that can be implemented on the carrier board. By default when a Microchip KSZ9563R switch is connected to the ShMM-700R Ethernet 0 PHY interface it is assumed to be connected to ShMM-700R I²C Bus I2C4 (logical I²C Bus 0, refer to Table 17) at 7-bit I²C address 5Fh. Likewise when a Microchip KSZ9563R switch is connected to the ShMM-700R Ethernet 1 PHY interface it is assumed to be connected to ShMM-700R I²C Bus I2C3 (logical I²C Bus 5) at 7-bit I²C address 5Fh. If needed the I²C bus and address values can be changed through four dedicated U-Boot environment variables `eth0_switch_bus`, `eth0_switch_address`, `eth1_switch_bus`, `eth1_switch_address`, refer to Table 5. Note if these U-Boot environment variables are not defined the default values are used, thus these U-Boot variables only need to be defined if the default value needs to be changed.

WARNING: Upon boot-up the U-Boot and kernel image will attempt to access the GbE Switch (Microchip KSZ9563RNX) at I²C address 5Fh on physical I²C buses I2C4 and I2C3. Customers that have ShMM-700R Carrier board designs that implement an I²C device with a 7-bit address of 5Fh on physical I²C bus I2C4 or I2C3, that is not the Microchip KSZ9563RNX, must contact the Pigeon Point Support (support.pigeonpoint@nVent.com).

3.4.5 Using the ShMM-500 and ShMM-700 alternate software redundancy interface

In ShMM-500s and ShMM-700s configured for ShMC Cross-connect operation, two additional network interfaces are implemented over the two USB connections. In this configuration, they always connect the two redundant Shelf Managers. These interfaces are named `usb0` and `usb1`. The interface `usb0` always exists, while the interface `usb1` exists only if the interface `usb0` is active on the peer Shelf Manager (which means that the peer Shelf Manager is physically installed and running). Also, the interfaces are cross-connected: `usb0` on the first Shelf Manager is connected to `usb1` on the second Shelf Manager, and vice versa.

The Shelf Manager supports usage of the USB network interfaces for communication between the redundant Shelf Managers. To use this feature, it is necessary to define two redundancy network adapters in the Shelf Manager configuration file `/etc/shelfman.conf`, as follows:

```
REDUNDANCY_NET_ADAPTER = usb0
```

```
REDUNDANCY_NET_ADAPTER2 = usb1
```

One additional consideration relates to the definition of the subnet mask for the redundancy network interfaces. In the legacy case, when only one redundant network adapter is used, two different IP addresses are derived from the redundancy IP address specified in */etc/shelfman.conf*. They are assigned to the two endpoints of the redundancy connection and differ only in the least significant bit.

However, when two redundancy network adapters are used, four different IP addresses are used, one for each of the endpoints (two endpoints on each of the two redundant Shelf Managers). To ensure proper operation, the two endpoints on the same Shelf Manager (usb0 and usb1) must belong to different logical networks, while usb0 on one Shelf Manager and usb1 on the other Shelf Manager must belong to the same logical network. This is achieved by dividing the IP address space into two ranges.

These ranges (logical networks) are defined by the subnet mask given by the parameter *REDUNDANCY_NETMASK* from the configuration file */etc/shelfman.conf*. If the network mask is 255.255.255.128 then the first range is 192.168.1.0 - 192.168.1.127 and the other is 192.168.1.128 – 192.168.1.255.

The usb0 endpoint on the first Shelf Manager and the usb0 endpoint on the other Shelf Manager are in the first range. The usb1 endpoint on the first Shelf Manager and the usb1 endpoint on the other Shelf Manager are in the second range.

The 4 IP addresses in question can be derived from one IP address (for example, the IP address assigned to usb0 on the Shelf Manager with the even hardware address) and the network mask (for which the recommended value is 255.255.255.128). The rules are as follows:

- To compute the IP address for usb0 on the Shelf Manager with the even hardware address when the file */etc/readhwaddr* is present you should set the least significant bit of *REDUNDANT_IP_ADDRESS* to 0.
- To compute the IP address for usb1 on the other Shelf Manager you should toggle the least significant bit in the usb0 IP address on the first Shelf Manager. This guarantees that usb0 on Shelf Manager with the even hardware address and usb1 on the Shelf Manager with the odd hardware address are in the same logical network and are not equal to each other.
- To compute the IP address for usb1 on the Shelf Manager with even hardware address you should take the IP address for usb0 on the same Shelf Manager and toggle the least non-zero bit of the network mask. This guarantees that the IP addresses for usb0 and usb1 on the Shelf Manager with the even hardware address are in different logical networks.
- The last step is to compute the IP address for usb0 on the Shelf Manager with the odd hardware address. You should either toggle the least significant bit in the IP address for usb1 on the Shelf Manager with the even hardware address or toggle the least non-zero bit of the network mask in the IP address for usb0 on the Shelf Manager with odd hardware address. The result is the same. This guarantees that usb0 on Shelf Manager with the odd hardware address and usb1 on Shelf Manager with the even hardware address are in the same logical network and are not equal.

Here is an example of deriving IP addresses for the USB network interfaces, under the assumption that the following definitions are in */etc/shelfman.conf*:

```
REDUNDANT_IP_ADDRESS = 192.168.1.2
REDUNDANCY_NETMASK = 255.255.255.128
```

The least significant non-zero bit in the network mask is the 7th bit (where smaller bit numbers are less significant). To toggle this bit in an IP address it is sufficient to add 128 (if this bit is set to zero in the IP address) or subtract 128 (if this bit is set to 1 in the IP address).

To toggle the least significant bit in an IP address it is sufficient to add 1 if the IP address is even or subtract 1 if the IP address is odd. Since *REDUNDANT_IP_ADDRESS* is even, the computations are the same whether the file */etc/readhwaddr* is present or not.

On the ShMM with the even hardware address the assignment of IP addresses looks like this:

- usb0: 192.168.1.2 (no changes)
- usb1: 192.168.1.130 (toggling the least significant non-zero bit of the netmask)

On the ShMM with the odd hardware address the assignment of IP addresses looks like this:

- usb0: 192.168.1.131 (toggling the least significant bit of the IP address and the least non-zero bit of the netmask)
- usb1: 192.168.1.3 (toggling the least significant bit of the IP address)

Here is another example of deriving IP addresses for the USB network interfaces, under the assumption that the following definitions are in */etc/shelfman.conf*:

```
REDUNDANT_IP_ADDRESS = 192.168.1.13
REDUNDANCY_NETMASK = 255.255.255.128
```

Suppose also that the file */etc/readhwaddr* is present.

The least significant non-zero bit in the network mask is 7th bit. To toggle this bit in an IP address it is sufficient to add 128 (if this bit is set to zero in the IP address) or subtract 128 (if this bit is set to 1 in the IP address). To toggle the least significant bit in an IP address it is sufficient to add 1 if the IP address is even or subtract 1 if the IP address is odd.

On the ShMM with the even hardware address the assignment of IP addresses looks like this:

- usb0: 192.168.1.12 (since the file */etc/readhwaddr* is present, the least significant bit should be set to zero)
- usb1: 192.168.1.140 (toggling the least significant non-zero bit of the netmask)

On the ShMM with the odd hardware address the assignment of IP addresses looks like this:

- usb0: 192.168.1.141 (toggling the least significant bit of the IP address and the least significant non-zero bit of the netmask)
- usb1: 192.168.1.13 (toggling the least significant bit of the IP address)

Starting from Shelf Manager release 2.8.0, data transferred over the Software Redundancy Interface can be compressed. A configuration variable

REDUNDANCY_COMPRESSION_THRESHOLD specifies the threshold for message size; messages that are bigger than this threshold are compressed by the *gzip* algorithm before sending and uncompressed after receiving. Smaller messages are sent in uncompressed form. The default value

of the threshold is -1 , which effectively turns compression off for compatibility with previous Shelf Manager releases. Since previous releases of the Shelf Manager do not understand compressed data, compression should be turned on only if both redundant Shelf Managers are release 2.8.0 or higher of the Shelf Manager.

3.4.6 Using the ShMM-1500 alternate software redundancy interface

For ShMM-1500s that are configured for ShMC Cross-connect operation, software redundancy communication takes place over an FPGA-implemented serial interface with SLIP (Serial Line IP support) configured on it, thereby enabling TCP/IP to be used by the Shelf Manager for this link, just as it is on the ShMM-500. The corresponding Linux serial interface is named `sl0`.

To use this interface, it is necessary to define the redundancy network adapter in the Shelf Manager configuration file `/etc/shelfman.conf`, as follows:

```
REDUNDANCY_NET_ADAPTER = sl0
```

In all other aspects, configuration of the serial interface `sl0` for redundant communication is the same as in the case of the second Ethernet adapter (see 3.4.2). In particular, both the active and backup ShMM specify the same IP address for the redundancy interface, but software assigns the next logical IP address to the ShMM with an odd hardware address.

Compression of data transferred over the Software Redundancy Interface applies also to ShMM-1500. Compression may be even more beneficial for performance on ShMM-1500 since the SLIP interface is somewhat slower than the USB-based SRI used on the ShMM-500. However, as with the ShMM-500, compression is turned off by default.

3.4.7 Changing the default ShMM network parameters

Configuring a ShMM to work in a specific network environment requires changing the following network parameters:

- RMCP IP address (*rmcpaddress*)
- RMCP gateway address (*gatewayip*)
- RMCP netmask (*netmask*)

If the second ShMM Ethernet is used for software redundancy communication between a pair of ShMMs, the network parameters above would typically not need to be adapted to a different network environment, since this is a dedicated private network between the redundant ShMMs. However, if the second network interface is used for RMCP communication, please use steps 10 and 11 below to change the settings; U-Boot environment variables do not apply to the second network interface.

Changing the RMCP network parameters is a two step process. First, the U-Boot network environment variables need to be updated, then the booted active ShMM module network settings need to be updated using the Shelf Manager command line interface (CLIA). Specific steps are shown below:

1. Attach a serial port console connection to the ShMM module.
This typically is 115200 Baud, N/8/1 for ShMM-500, ShMM-1500 and ShMM-700.
2. Reset the ShMM carrier and press the space bar to interrupt the automatic boot-up process.
You should see on ShMM-500:

```
U-Boot 1.1.2 (Apr 27 2005 - 19:17:09)
Board: ShMM-500
S/N: 00 00 00 00 00 00 00 00 00 03 03 03
DRAM: 64 MB
Flash: 16 MB
In: serial
Out: serial
Err: serial
Net: AulX00 ETHERNET
Hit any key to stop autoboot: 0
shmm500
```

On ShMM-1500, you should see the following:

```
U-Boot 1.1.4 (Jun 15 2006 - 17:49:12) MPC83XX
Clock configuration:
  Coherent System Bus: 99 MHz
  Core: 249 MHz
  Local Bus: 24 MHz
CPU: MPC83xx, Rev: 1.1 at 249.975 MHz
Board: ShMM-1500R
PCI1: 32 bit, 33 MHz
I2C: ready
DRAM: 128 MB
FLASH: 64 MB
PCI: Bus Dev VenId DevId Class Int
      00 17 1172 0001 ff00 00
In: serial
Out: serial
Err: serial
FPGA: firmware version 1.12, carrier id 0
Net: TSEC0, TSEC1
Hit any key to stop autoboot: 0
shmm1500
```

On ShMM-700, you should see the following:

```
PowerPrep start initialize power...
Battery Voltage = 1.44V
No battery or bad battery detected!!!.Disabling battery voltage measurements.
Feb 24 201207:45:29
FRAC 0x92926152
memory type is DDR2
Wait for ddr ready 1
power 0x00820616
Frac 0x92926152
```

```
start change cpu freq
hbus 0x00000003
cpu 0x00010001
start test memory access
ddr2 0x40000000
finish simple test
U-Boot 2009.08 (Feb 24 2012 - 07:44:12)
Freescale i.MX28 family
CPU: 297 MHz
BUS: 99 MHz
EMI: 130 MHz
GPMI: 24 MHz
I2C: ready
DRAM: 128 MB
SFGEN: N25Q512A detected, total size 64 MB
A2F: SPICOMM protocol v1.6, M3 firmware v0.8, FPGA design v0.42.0.0
A2F: A2F firmware, version 0.8
A2F: Last reset cause: HARD
A2F: Device type: A2F060M3E-FG256
A2F: MSS clock frequency: 40 MHz
A2F: Fabric clock frequency: 20 MHz
A2F: Fast delay calibration: 1330 cycles per 100uS
A2F: eNVM: 128 KB (00000000 - 00020000)
A2F: eSRAM: 16 KB (20000000 - 20004000)
A2F: Extram start: 200022D0
RUPG: booting from image 0 (confirmed)
In: serial
Out: serial
Err: serial
Net: FEC0: 00:18:49:01:8f:26, FEC1: 00:18:49:01:8f:27
FEC0, FEC1
Hit any key to stop autoboot: 0
shmm700
```

3. Echo current network settings:

```
shmmxx00 printenv rmcpaddr netmask gatewayip
rmcpaddr=192.168.0.44
netmask=255.255.255.0
gatewayip=192.168.0.1
shmmxx00
```

4. Change settings and commit to non-volatile storage:

```
shmmxx00 setenv rmcpaddr 10.1.1.10
shmmxx00 setenv netmask 255.255.0.0
```



```
shmmxx00 setenv gatewayip 10.1.1.1
shmmxx00 saveenv
Saving Environment to EEPROM...
shmmxx00
```

Note: saveenv output is slightly different on the ShMM-700.

5. Boot the ShMM up to full operational state and log in as user “root”.

On ShMM-500:

```
shmm500 reset
U-Boot 1.1.2 (Apr 27 2005 - 19:17:09)
CPU: Au1550 324 MHz, id: 0x02, rev: 0x00
Board: ShMM-500
S/N: 00 00 00 00 00 00 00 00 00 03 03 03
DRAM: 64 MB
Flash: 16 MB
In: serial
Out: serial
Err: serial
Net: AulX00 ETHERNET
Hit any key to stop autoboot: 0
## Booting image at bfb00000 ...
Image Name: MIPS Linux-2.4.26
Created: 2005-05-07 17:35:21 UTC
Image Type: MIPS Linux Kernel Image (gzip compressed)
Data Size: 843144 Bytes = 823.4 kB
Load Address: 80100000
Entry Point: 802bc040
Verifying Checksum ... OK
Uncompressing Kernel Image... OK
## Loading Ramdisk Image at bfc40000 ...
Image Name: sentry RFS Ramdisk Image
...
shmm500 login: root
BusyBox v0.60.5 (2005.05.07-17:27+0000) Built-in shell (msh)
#
```

On ShMM-1500:

```
shmm1500 reset
Resetting the board.
U-Boot 1.1.4 (Jun 15 2006 - 17:49:12) MPC83XX
Clock configuration:
Coherent System Bus: 99 MHz
```

```

Core:                249 MHz
Local Bus:           24 MHz
CPU:   MPC83xx, Rev: 1.1 at 249.975 MHz
Board: ShMM-1500R
PCI1:  32 bit, 33 MHz
I2C:   ready
DRAM:  128 MB
FLASH: 64 MB
PCI:   Bus Dev VenId DevId Class Int
       00  17  1172  0001  ff00  00
In:    serial
Out:   serial
Err:   serial
FPGA:  firmware version 1.12, carrier id 0
Net:   TSEC0, TSEC1
Hit any key to stop autoboot:  0
## Booting image at f00c0000 ...
   Image Name:   Linux-2.4.25
   Created:      2006-06-29  14:51:42 UTC
   Image Type:   PowerPC Linux Kernel Image (gzip compressed)
   Data Size:    786177 Bytes = 767.8 kB
   Load Address: 00000000
   Entry Point:  00000000
   Verifying Checksum ... OK
   Uncompressing Kernel Image ... OK
## Loading RAMDisk Image at f04c0000 ...
   Image Name:   sentry.shmm1500 RFS Ramdisk Imag
   Created:      2008-02-15  16:35:56 UTC
   Image Type:   PowerPC Linux RAMDisk Image (gzip compressed)
   Data Size:    3826312 Bytes = 3.6 MB
   Load Address: 00000000
   Entry Point:  00000000
   Verifying Checksum ... OK
   Loading Ramdisk to 07bfc000, end 07fa2288 ... OK
...
shmm1500 login: root
On ShMM-700:
shmm700 reset
resetting ...
PowerPrep start initialize power...
Battery Voltage = 1.44V
No battery or bad battery detected!!!.Disabling battery voltage measurements.
Feb 24 201207:45:29

```

```
FRAC 0x92926152
memory type is DDR2
Wait for ddr ready 1
power 0x00820616
Frac 0x92926152
start change cpu freq
hbus 0x00000003
cpu 0x00010001
start test memory access
ddr2 0x40000000
finish simple test
U-Boot 2009.08 (Feb 24 2012 - 07:44:12)
Freescale i.MX28 family
CPU: 297 MHz
BUS: 99 MHz
EMI: 130 MHz
GPMI: 24 MHz
I2C: ready
DRAM: 128 MB
SFGEN: N25Q512A detected, total size 64 MB
A2F: SPICOMM protocol v1.6, M3 firmware v0.8, FPGA design v0.42.0.0
A2F: A2F firmware, version 0.8
A2F: Last reset cause: HARD
A2F: Device type: A2F060M3E-FG256
A2F: MSS clock frequency: 40 MHz
A2F: Fabric clock frequency: 20 MHz
A2F: Fast delay calibration: 1330 cycles per 100uS
A2F: eNVM: 128 KB (00000000 - 00020000)
A2F: eSRAM: 16 KB (20000000 - 20004000)
A2F: Extram start: 200022D0
RUPG: booting from image 0 (confirmed)
In: serial
Out: serial
Err: serial
Net: FEC0: 00:18:49:01:8f:26, FEC1: 00:18:49:01:8f:27
FEC0, FEC1
Hit any key to stop autoboot: 0
RUPG: upgrade file is not present, skipping
Loading file 'uImage.1' to addr 0x42000000 with size 1399152 (0x00155970)...
Loaded from UBIFS cache
Loading file 'rfs.1' to addr 0x46000000 with size 2805717 (0x002acfd5)...
Loaded from UBIFS cache
## Booting kernel from Legacy Image at 42000000 ...
```

```

Image Name:      Linux-2.6.34.8-shmm700
Image Type:      ARM Linux Kernel Image (uncompressed)
Data Size:       1399088 Bytes =  1.3 MB
Load Address:    40008000
Entry Point:     40008000
Verifying Checksum ... OK
## Loading init Ramdisk from Legacy Image at 46000000 ...
Image Name:      ShMM700 RFS 3.7.0
Image Type:      ARM Linux RAMDisk Image (gzip compressed)
Data Size:       2277767 Bytes =  2.2 MB
Load Address:    00000000
Entry Point:     00000000
Verifying Checksum ... OK
Loading Kernel Image ... OK
OK
Starting kernel ...
...
shmm700 login: root
#

```

6. Allow the ShMM to start up.

Please note that the settings that were changed in the U-Boot firmware are not necessarily propagated to the Linux environment. The reason for this is that the Shelf Manager needs to maintain its own copy of the network configuration data in order to manage failover situations. If this is the first time the Shelf Manager has been booted, or if the Flash filesystems has been reset prior to bootup, then the Shelf Manager uses the network settings provided by U-Boot to set up this networking context (and thus the changes you made in U-Boot are propagated forward). Otherwise, the network configuration data is stored on the ShMM Flash file system in the following files:

- `/var/nvdata/ch1_param` - parameters for channel 1
- `/var/nvdata/ch1_param_v2` - additional parameters for channel 1 (Shelf Manager 2.6 and higher)
- `/var/nvdata/ch1_param_ipv6_v2` - additional IPv6-related parameters for channel 1 (ShMM-700 only, starting with Shelf Manager release 3.5)
- `/var/nvdata/ch2_param` - parameters for channel 2
- `/var/nvdata/ch2_param_v2` - additional parameters for channel 2 (Shelf Manager release 2.6 and higher)
- `/var/nvdata/ch2_param_ipv6_v2` - additional IPv6-related parameters for channel 2 (ShMM-700 only, starting with ShelfManager release 3.5)

In this latter case, the following steps are required to configure the network settings in the Shelf Manager context.

7. First, check to see if you are interacting with the active Shelf Manager.

You only need to make changes on the active Shelf Manager as it updates the backup with the network configuration changes via the redundancy interface. Use the *hwri* command, and look for “Active” in the output of the command (see the underlined fragment of the output below). The *hwri* command reports the current state of the hardware redundancy signals on the ShMM, exposed through the CPLD on the ShMM-500 and through the FPGA on the ShMM-1500 and ShMM-700. The flags reported by the *hwri* command are identical for all three variants, though the actual register interface and bit layout is different among those variants. If you are not using the active ShMM, connect to the other ShMM device and repeat step 7:

```
# hwri
```

```
HWRI word: 3307
    1000h - Local Presence
    0100h - Remote Presence
    0002h - Local Healthy
    0200h - Remote Healthy
    0004h - Local Switchover Request
    2000h - Active
```

8. Get the current IP settings (for channel 1):

```
# clia getlanconfig 1
```

```
Pigeon Point Shelf Manager Command Line Interpreter
```

```
Authentication Type Support: 0x15 ( None MD5 Straight Password/Key )
```

```
Authentication Type Enables:
```

```
    Callback level: 0x00
```

```
    User level: 0x15 ( "None" "MD5" "Straight Password/Key" )
```

```
    Operator level: 0x15 ( "None" "MD5" "Straight Password/Key" )
```

```
    Administrator level: 0x15 ( "None" "MD5" "Straight Password/Key" )
```

```
    OEM level: 0x00
```

```
IP Address: 206.25.139.28
```

```
IP Address Source: Static Address (Manually Configured) (0x01)
```

```
MAC Address: 00:50:c2:22:50:30
```

```
Subnet Mask: 0.0.0.0
```

```
IPv4 Header Parameters: 0x40:0x40:0x10
```

```
Primary RMCP Port Number: 0x026f
```

```
Secondary RMCP Port Number: 0x0298
```

```
BMC-generated ARP Control: 0x02
```

```
    Enable BMC-generated ARP Response
```

```
Gratuitous ARP Interval: 2.0 seconds
```

```
Default Gateway Address: 206.25.139.3
```

```
Default Gateway MAC Address: 00:00:00:00:00:00
```

```
Backup Gateway Address: 0.0.0.0
```

Backup Gateway MAC Address: N/A
Community String: "public"
Number of Destinations: 16
Destination Type:
 N/A
Destination Address:
 N/A
802.1q VLAN ID: 0 (disabled)
VLAN priority: 0
Cipher Suite Entry count: 15
Supported Cipher Suite IDs: 0h, 1h, 2h, 3h, 4h, 5h, 6h, 7h, 8h, 9h, Ah, Bh, Ch, Dh, Eh
Cipher Suite Privilege Levels:
 ID 00h, Priv.Level 'User' (2); ID 01h, Priv.Level 'User' (2);
 ID 02h, Priv.Level 'Administrator' (4); ID 03h, Priv.Level 'OEM Proprietary' (5);
 ID 04h, Priv.Level 'OEM Proprietary' (5); ID 05h, Priv.Level 'OEM Proprietary' (5);
 ID 06h, Priv.Level 'User' (2); ID 07h, Priv.Level 'Administrator' (4);
 ID 08h, Priv.Level 'OEM Proprietary' (5); ID 09h, Priv.Level 'OEM Proprietary' (5);
 ID 0Ah, Priv.Level 'OEM Proprietary' (5); ID 0Bh, Priv.Level 'Administrator' (4);
 ID 0Ch, Priv.Level 'OEM Proprietary' (5); ID 0Dh, Priv.Level 'OEM Proprietary' (5);
 ID 0Eh, Priv.Level 'OEM Proprietary' (5);
Destination Address VLAN TAGs:
 N/A
Bad password threshold: disabled

9. Change the IP settings (for channel 1):

```
# clia setlanconfig 1 ip 10.1.1.10
Pigeon Point Shelf Manager Command Line Interpreter
IP set successfully
# clia setlanconfig 1 subnet_mask 255.255.0.0
Pigeon Point Shelf Manager Command Line Interpreter
Subnet Mask set successfully
# clia setlanconfig 1 dft_gw_ip 10.1.1.1
Pigeon Point Shelf Manager Command Line Interpreter
Default Gateway Address set successfully
#
```

10. If the second network interface is used for RMCP, get the current IP settings for channel 2:

```
# clia getlanconfig 2
```

```
Pigeon Point Shelf Manager Command Line Interpreter
```

```
...
```

11. Change the IP settings (for channel 2):

```
# clia setlanconfig 2 ip 10.2.1.10
```

```
Pigeon Point Shelf Manager Command Line Interpreter
```

```
IP set successfully
```

```
# clia setlanconfig 2 subnet_mask 255.255.0.0
```

```
Pigeon Point Shelf Manager Command Line Interpreter
```

```
Subnet Mask set successfully
```

```
# clia setlanconfig 2 dft_gw_ip 10.2.1.1
```

```
Pigeon Point Shelf Manager Command Line Interpreter
```

```
Default Gateway Address set successfully
```

3.4.8 Assigning VLAN IDs

Starting from release 2.6.0, the Shelf Manager supports Virtual LANs (VLANs), in accordance with version 2.0 of the IPMI specification. That is, RMCP/RMCP+ traffic to and from the Shelf Manager can be passed over a specific virtual LAN; the ID of that virtual LAN is stored in the LAN configuration parameters and is automatically added to the outgoing RMCP/RMCP+ packets (and stripped from the incoming packets) in accordance with the IEEE 802.1q specification. Separate VLAN IDs are supported for the two LAN channels (1 and 2) that can carry RMCP/RMCP+ traffic.

To assign a VLAN ID to a channel, use the IPMI “Set LAN Configuration Parameters” command, the CLI *setlanconfig* command or the CLI analog in the SNMP or WEB interface. Changing the VLAN ID may disrupt current RMCP/RMCP+ communication on the affected channel, so it should be used with caution. The VLAN ID is a 12-bit unsigned number, the value 0 indicates that virtual LANs are not used on the channel. Once set, the VLAN ID is stored persistently in the LAN configuration parameters on the ShMM Flash file system.

Example: Get the current VLAN ID settings for channel 1 (where no VLANs are currently in use):

```
# clia getlanconfig 1 vlan_id
```

```
Pigeon Point Shelf Manager Command Line Interpreter
```

```
802.1q VLAN ID: 0 (disabled)
```

Set the VLAN ID for channel 1 to a different value:

```
# clia setlanconfig 1 vlan_id enabled 5
```

```
Pigeon Point Shelf Manager Command Line Interpreter
```

```
VLAN ID set successfully
```

The IPMI 2.0 specification allows distinct VLAN IDs to be used for the outgoing SNMP traps (LAN alerts) generated by the Shelf Manager as a result of PEF processing. As allowed by the IPMI specification, the Shelf Manager does not support this extension and always uses the main VLAN ID that is assigned to the corresponding channel for such LAN alerts. Each LAN configuration parameter that defines a VLAN ID for a LAN alert destination is read-only in the Shelf Manager and identical to the VLAN ID for the corresponding LAN channel.

The default VLAN ID for a channel is usually 0 (VLANs disabled), but can be configured via the configuration parameters `DEFAULT_VLAN_ID` (for channel 1) and `DEFAULT_VLAN_ID2` (for channel 2). These defaults apply in the absence of the LAN configuration parameters on a ShMM, on a fresh ShMM or after a Shelf Manager upgrade from a pre-2.6 version to the version 2.6. Once the LAN configuration parameters are defined and stored, the “Set LAN Configuration Parameters” mechanism should be used to change the VLAN ID.

VLAN support is available on all ShMM variants.

3.4.9 Assigning IP addresses to the Shelf Manager via DHCP (version 4)

DHCP (Dynamic Host Configuration Protocol) and DHCP servers can be used to assign IP addresses to the Shelf Manager. This section focuses on using version 4 of this protocol (as opposed to DHCPv6) that deals with IPv4 addresses (as opposed to IPv6 addresses). The abbreviations DHCP and DHCPv4 are used in this section to designate DHCP version 4, and the abbreviations IP, IPv4 are used in this section to designate IP version 4.

The following types of IP addresses can be assigned via DHCP:

- RMCP accessible addresses (for one or both network interfaces)
- Private Shelf Manager addresses (for one or both network interfaces, for both Shelf Managers).

A total of six Shelf Manager-related IP addresses can be assigned via DHCP.

Two different DHCPv4 clients are currently available in the Shelf Manager. The new DHCPv4 client is available starting in release 3.6.0 and has the following differences from the legacy DHCPv4 client:

- The new client supports finite lease times, and performs renewal and rebinding of IP addresses on the DHCP server as defined by the DHCP protocol (unlike the legacy client, where the lease time is assumed to be always infinite).
- The new client operates also on the backup Shelf Manager; it is responsible for obtaining the IP addresses for the backup physical Shelf Manager instance (unlike the legacy client, where the active Shelf Manager operates as a DHCP proxy client for the backup Shelf Manager). This is to ensure that a request comes to the DHCP server from the same network interface on the client for which the address is requested (which is a DHCPv4 protocol requirement).
- The new client assigns an IP address to each specific interface as soon as it receives the address for that interface from the DHCP server (unlike the legacy client, where the active Shelf Manager collects all six IP addresses from the DHCP server before assigning any of them). The old behavior is not applicable in the new architecture because the DHCP client on the backup Shelf Manager operates independently of the DHCP client on the active Shelf Manager.
- The state of the client is passed over the redundancy interface; therefore, after a switchover, the newly active Shelf Manager knows that state, including the remaining timeout periods for all leases associated with the logical Shelf Manager. The state of the DHCP client for the logical Shelf Manager is therefore transparent for the DHCP server; the client does not restart the DHCP protocol state machine after the switchover. The IP address for the logical Shelf Manager moves over to the other Shelf Manager during a switchover, as in the static case.

- The IP address for the physical Shelf Manager that becomes active stays unchanged during the switchover. The physical Shelf Manager that lost the active role reboots and, after its activation as the backup, asks the DHCP server again for an address. It depends on the server whether it provides the same or a different IP address; the Client Identifier that the Shelf Manager sends to the DHCP server is the same, however.
- The new client does not support the DHCP proxy mechanisms for IPM controllers that are described in section 4.10.4 of this document and which have been obsoleted by the HPM.3 specification.

The choice of which client to use for DHCPv4 client functionality (legacy vs. new) depends on the Shelf Manager configuration parameter `ENABLE_NEW_DHCPV4_CLIENT`; the default value is for this parameter is `TRUE`.

In some future release, multiple releases after release 3.6.0, support for the legacy DHCP client is likely to be removed from the Shelf Manager, together with the differentiating configuration parameter. Thereafter, only the new DHCP client will be included with the Shelf Manager in subsequent releases.

A specific IP address is designated by a particular value of the Client Identifier that is passed from the Shelf Manager to the DHCP server. All formats of Client Identifiers described below are supported by both legacy and new clients.

If the HPM.3 Client Identifier format is enabled (which is the default), the Client Identifier has the following standard format, defined in the HPM.3 specification:

Table 10: HPM.3 DHCP Client ID

FIELD	LENGTH	FORMAT	CONTENT
<i>Header</i>	1	Binary	61 (3Dh)
<i>Data Length</i>	1	Binary	Length of remaining content (44, 2Ch)
<i>Type</i>	1	Binary	255 (FFh)
<i>IAID</i>	4	Binary	<i>IPMI LAN Channel Index</i> with range 0 through Ah and all other values reserved; see section 2.5.1 of HPM.3, DHCP-Assigned Platform Management Parameters Specification.
<i>DUID</i>	39	Binary	<i>DHCP Unique Identifier</i> . See Table 11: HPM.3 DHCP Unique Identifier format

The DHCP Unique Identifier (DUID) has the following format:

Table 11: HPM.3 DHCP Unique Identifier format

FIELD	LENGTH	FORMAT	CONTENT
<i>DUID format</i>	2	Binary	DUID-EN identifier (2)
<i>DUID header part 1</i>	4	Binary	PICMG IANA ID (12634, 00315Ah)
<i>DUID header part 2</i>	5	ASCII	Text “HPM.3” identifier (48504D2E33h)
<i>DUID header part 3</i>	2	ASCII	HPM.3 DUID version (“-1”, 2D31h)
<i>Shelf address type/length</i>	1	Binary	<i>Shelf Address Type/Length</i> Byte field from the Shelf FRU Information (AdvancedTCA Address Table or MicroTCA Shelf Information record)
<i>Shelf address</i>	20	Binary	<i>Shelf Address</i> field from the Shelf FRU Information (AdvancedTCA Address Table or MicroTCA Shelf Information record)
<i>Shelf type</i>	1	Binary	0 for AdvancedTCA shelves, 1 for MicroTCA shelves
<i>Slot location</i>	4	Binary	See Table 12: Slot Location record for AdvancedTCA

Table 12: Slot Location record for AdvancedTCA

OFFSET	LENGTH	DEFINITION
0	1	<i>Primary Site Type</i> . The Site Type of the Slot, for a Front Board or other Intelligent FRU that attaches directly to IPMB-0.
1	1	<i>Primary Site Number</i> . The Site Number of the Slot, for a Front Board or other Intelligent FRU that attaches directly to IPMB-0.

2	1	<i>Secondary Site Type</i> . Applicable when the Slot is a subsidiary site on the FRU designated by the <i>Primary Site Type</i> and <i>Primary Site Number</i> ; identifies the Site Type of that Slot.
3	1	<i>Secondary Site Number</i> . Applicable when the Slot is a subsidiary site on the FRU designated by the <i>Primary Site Type</i> and <i>Primary Site Number</i> ; identifies the Site Number of that Slot.

For the Shelf Manager addresses, the fields are assigned as follows:

- Primary Site Type = 3 (Dedicated Shelf Manager)
- Primary Site Number = 0 to request RMCP accessible addresses, 1 or 2 to request Private Shelf Manager addresses (the value 1 or 2 is the site number for the corresponding Shelf Manager in the Shelf Address Table)
- Secondary Site Type and Secondary Site Number = 0
- IPMI LAN Channel Index = 0 for the physical interface used by the first LAN channel and =1 for the physical interface used by the second LAN channel. The first physical interface is normally eth0 and the second physical interface is normally eth1, but this can be changed by the Shelf Manager configuration variables *RMCP_NET_ADAPTER* and *RMCP_NET_ADAPTER2*.

If the HPM.3 Client Identifier format is disabled, then in the default implementation, the Client Identifier is based on the Shelf Address string that is stored in the Shelf FRU Information, plus the Request Identifier (Request ID) byte at the end.

The Request ID byte has the following format:

- Request ID bits 7..4 – Shelf Manager number (0, 1, 2, with 0 for logical Shelf Manager)
- Request ID bits 3..0 – Interface number (0 for eth0, 1 for eth1)

If the Shelf Address string in the Shelf FRU Information is empty, a hard-coded string of 7 zero bytes is used. In that case, the Client Identifier values are implemented as follows:

Currently the Client Identifier values are hardcoded as follows:

```
00:00:00:00:00:00:00:10 - Shelf Manager 1, eth0
00:00:00:00:00:00:00:11 - Shelf Manager 1, eth1
00:00:00:00:00:00:00:20 - Shelf Manager 2, eth0
00:00:00:00:00:00:00:21 - Shelf Manager 2, eth1
00:00:00:00:00:00:00:00 - Logical Shelf Manager eth0
00:00:00:00:00:00:00:01 - Logical Shelf Manager eth1
```

However, these Client IDs can be redefined for a specific carrier inside the corresponding ShMM carrier-specific module (if the HPM.3 Client Identifier format is disabled).

To use this feature in the Shelf Manager, it is necessary to define the configuration parameter *USE_DHCP* in the Shelf Manager configuration file */etc/shelfman.conf*, as follows:

```
USE_DHCP = TRUE
```

The configuration parameter *ENABLE_NEW_DHCPV4_CLIENT* determines whether the legacy or the new DHCP client is used. By default, this variable is set to *TRUE* which means the new client is

used. To use the legacy DHCP client, set this variable explicitly to *FALSE* in the Shelf Manager configuration file, as follows:

```
ENABLE_NEW_DHCPV4_CLIENT = FALSE
```

By default, the Shelf Manager uses the first DHCP server that answers the *DHCPDISCOVER* request to assign the IP addresses. If more than one DHCP server is present in the network, the configuration parameter *DHCP_SERVER_ADDRESS* can be used in the Shelf Manager configuration file */etc/shelfman.conf* to specify the IP address of the DHCP server to be used, as follows:

```
DHCP_SERVER_ADDRESS = 192.168.1.50
```

In that case, only the specified DHCP server is used.

In addition, if the configuration parameter *DHCP_FOR_RMCP_ONLY* is set to *TRUE*, only the RMCP-accessible (Logical Shelf Manager) IP addresses are assigned via DHCP. In this case, the private ShMM IP addresses, if any, are left untouched.

The DHCP server should be configured to provide a unique IP address for each Client Identifier. Also, when the legacy client is in use, the lease time of each address must be set as “infinite” (time value 0xFFFFFFFF) to avoid IP address expiration. The new client correctly handles finite lease times.

If the configuration of the DHCP server changes over time, it may be necessary to restart the procedure of IP address assignment on the DHCP client for the Shelf Manager. The CLI command *dhcp restart* can be used for this purpose. Another CLI command, *dhcp status*, can be used to find out the current DHCP client status on the Shelf Manager (which IP addresses have already been retrieved and assigned). Please refer to the *Shelf Manager External Interface Reference* for the detailed information about these CLI commands.

In addition to providing an IP address related information, a DHCP server can be configured, via the TFTP Server Name and Bootfile Name options, to provide the name of a TFTP server and identify a bootfile on that server. For the benefit of other applications that may run on the ShMM in parallel with the Shelf Manager, the DHCP client, if it receives this additional information, stores it in the ShMM file system. A separate file is used for each of the 6 network interfaces processed by the DHCP client, according to the following rules:

/tmp/tftp_info_0 for the first network interface on the ShMM with the lower Hardware Address

/tmp/tftp_info_1 for the second network interface on the ShMM with the lower Hardware Address

/tmp/tftp_info_2 for the first network interface on the ShMM with the higher Hardware Address

/tmp/tftp_info_3 for the second network interface on the ShMM with the higher Hardware Address

/tmp/tftp_info_4 for the first RMCP-accessible interface

/tmp/tftp_info_5 for the second RMCP-accessible interface

The first four files above are stored on the corresponding ShMM only, the last two are stored on both redundant ShMMs.

Each file consists of two lines. The first line represents the TFTP server name and the second line represents the boot file name; for example:

```
# cat /tmp/tftp_info_2
TFTPSERVER="192.168.1.253"
BOOTFILE="/tftpboot/ssh_script_eth0_ShM2"
```

The example configuration file below shows how to configure the Linux DHCP server (DHCPD) to provide IP addresses, as well as TFTP server and bootfile information to the Shelf Manager. Fixed predefined addresses are used for that purpose.

This file should be located as */etc/dhcpd.conf* on the system hosting the DHCP server. Other DHCP servers (such as those on non-Linux operating systems) are configured differently.

```
allow booting;
allow bootp;

option domain-name "tst";
option subnet-mask 255.255.255.0;
option domain-name-servers 192.168.1.100;
option ntp-servers 192.168.1.50;
option routers 192.168.1.253;
option vendor-class-identifier "nVent";

min-lease-time 4294967295;
default-lease-time 4294967295;

use-host-decl-names on;
ddns-update-style ad-hoc;

subnet 192.168.1.0 netmask 255.255.255.0
{
    host client00 {
        option dhcp-client-identifier 0:0:0:0:0:0:0:10;
        fixed-address 192.168.1.140;
        option tftp-server-name "192.168.1.253";
        option bootfile-name "/tftpboot/ssh_script_eth0_ShM1";
    }
    host client01 {
        option dhcp-client-identifier 0:0:0:0:0:0:0:11;
        fixed-address 192.168.1.141;
        option tftp-server-name "192.168.1.253";
        option bootfile-name "/tftpboot/ssh_script_eth1_ShM1";
    }
    host client02 {
        option dhcp-client-identifier 0:0:0:0:0:0:0:20;
```

```

        fixed-address 192.168.1.142;
        option tftp-server-name "192.168.1.253";
        option bootfile-name "/tftpboot/ssh_script_eth0_ShM2";
    }
    host client03 {
        option dhcp-client-identifier 0:0:0:0:0:0:0:21;
        fixed-address 192.168.1.143;
        option tftp-server-name "192.168.1.253";
        option bootfile-name "/tftpboot/ssh_script_eth1_ShM2";
    }
    host client04 {
        option dhcp-client-identifier 0:0:0:0:0:0:0:0;
        fixed-address 192.168.1.144;
    }
    host client05 {
        option dhcp-client-identifier 0:0:0:0:0:0:0:1;
        fixed-address 192.168.1.145;
    }
}

```

3.4.10 Using IPv6 for shelf-external access to the Shelf Manager (ShMM-700 only)

Starting with release 3.5, the Shelf Manager supports IPv6 in parallel with IPv4 for communication with the System Manager and other shelf-external entities, using RMCP or other application level protocols. This support exists only on ShMM-700, since the Linux kernel on ShMM-500/1500 does not support IPv6.

Starting with release 3.5.1, the Shelf Manager also supports the PICMG 3.0 R3.0 amendments defined by ECN 3.0-3.0-001. In particular, Shelf Manager IPv6 Connection records in the Shelf FRU Information and the new format of the command “Get Shelf Manager IP Addresses” are supported.

For configuration of IPv6-based services, the Shelf Manager supports new IPv6-related LAN configuration parameters, which are defined in the IPMI 2.0, rev 1.1 specification (parameters 50-80). These LAN configuration parameters, as well as the previously supported parameters, are stored persistently and are redundant (that is, changes to the parameters are passed from the active to the backup Shelf Manager over the redundancy interface and stored there). The parameters are stored in the files `/var/nvdata/ch1_param_ipv6_v2`, `/var/nvdata/ch2_param_ipv6_v2` (for LAN channels 1 and 2, respectively). From the user perspective, configuration can be done via the CLI, using the commands `setlanconfig` and `getlanconfig`.

The IPv6-based services in the Shelf Manager are designed as an addition, not replacement, of the existing IPv4-based services. So, it is not possible to make the Shelf Manager accessible only over IPv6 but not over IPv4.

IPv6-based services are always available for LAN channel 1, and for LAN channel 2, if the second network interface is used in the Active-Active mode. Use of IPv6 over a VLAN-based interface or over a bonded interface is also supported.

Use of IPv6 for HPM.2-based communication with FRUs in the shelf is not currently supported, since there are no provisions yet for the use of IPv6 in the HPM.2 specification.

Also, IPv6 is not supported for inter-ShMM communication over the software redundancy interface; only IPv4 can be used for that communication.

It is possible to use IPv6 as the underlying protocol for LAN alert messages sent to remote destinations over the network. For that, the Shelf Manager supports an extended format for LAN configuration parameter 19 (Destination Addresses), which allows specifying an IPv6 address for a LAN alert destination. LAN alerts sent to such destinations use the IPv6 protocol and packet format.

The following IPv6-related LAN configuration parameters are essential for IPv6 operation on a LAN channel:

- IPv6/IPv4 Addressing Enables (51). This parameter must be set to the value 02h (Enabled IPv6 and IPv4 addressing simultaneously) to allow IPv6-based communication with the Shelf Manager over the channel. Setting this parameter to 00h (IPv6 addressing disabled) causes the Shelf Manager to stop listening to IPv6 addresses and sending IPv6-based LAN alerts. The parameter value 01h (Enabled IPv6 addressing only, IPv4 addressing is disabled) is not supported by the Shelf Manager, since it cannot operate in IPv6-only mode.
- IPv6 Static Addresses (57). For IPv6-based services, at least one static IPv6 address must be configured on the channel. Up to 4 addresses can be configured. RMCP service is available on any of these addresses. The Shelf Manager makes sure that an RMCP request received on a specific address is responded to from the same address. However, RMCP service is not available on the link-local IPv6 addresses that are automatically configured by the kernel (starting with the `fe80:` prefix). When static IPv6 addresses are assigned or removed, the Shelf Manager assigns the address to or removes it from the physical network interface of the ShMM via the corresponding `ioctl()` call.
- IPv6 Router Address Configuration Control (64). This parameter controls whether static IPv6 router addresses, assigned via corresponding LAN configuration parameters, are taken into consideration by the Shelf Manager. The value 03h of this parameter turns on the use of static routers, and the value 02h turns it off. Other values are not valid for the Shelf Manager; the Shelf Manager always honors dynamic router assignment performed via the IPv6 Router Discovery protocol.
- Static Router 1 IP Address (65), Prefix Length (67), Prefix Value (68); Static Router 2 IP Address (69), Prefix Length (71), Prefix Value (72). These are the configuration parameters for the two static routers that Shelf Manager can use if the parameter IPv6 Router Address Configuration Control turns on the use of static routers.

The following are the default values of the IPv6-related LAN parameters in the factory-provided RFS images:

- IPv6/IPv4 Addressing Enables = 02h (Enabled IPv6 and IPv4 addressing simultaneously).

- IPv6 Static Addresses (57) = not configured.
- IPv6 Router Address Configuration Control (64) = 03h (enable static and dynamic router address configuration).
- Static Router 1 IP Address (65), Prefix Length (67), Prefix Value (68); Static Router 2 IP Address (69), Prefix Length (71), Prefix Value (72) = not configured.

Note: The above parameters are reset to the default values, as assigned in the current RFS image, if the /var/nvdata files are erased, as would occur after re-initializing the file system or possibly after a reliable upgrade with, for example, the --erase-all mode enabled.

To configure the Shelf Manager with the default IPv6-related LAN parameter values for responding to RMCP traffic over IPv6, it is necessary at least to assign a static IPv6 address to the corresponding channel. For example, to assign the IPv6 address 3ffe:302:11:2::26, with the network part of 64 bits, as the first address (the address index starts from 0) to the LAN channel 1, the following command can be used:

```
# clia setlanconfig 1 ipv6_addresses 0 3ffe:302:11:2::26/64
```

Configuration of static routers is optional, depending on whether appropriate IPv6 routers exist and can be discovered dynamically (via the Router Discovery protocol) on the target local network. For example, if such dynamic discovery is not feasible or desirable, the following commands can be used to configure statically the default router with the address 3ffe:302:11:2::7f on the Shelf Manager:

```
# clia setlanconfig 1 ipv6_router1_ip 3ffe:302:11:2::7f
# clia setlanconfig 1 ipv6_router1_prefix_length 0
# clia setlanconfig 1 ipv6_router1_prefix_value ::
```

The values assigned via the `clia setlanconfig` command are stored persistently and survive a Shelf Manager reboot or switchover. In the case of a switchover, the static IPv6 addresses and other configuration information are failed over to the former backup Shelf Manager, which now becomes active, and RMCP communication can continue without interruption (as occurs with IPv4).

Static Shelf Manager IPv6 addresses and static router IPv6 addresses can also be configured in the Shelf Manager IPv6 Connection records in the Shelf FRU Information, as defined in PICMG 3.0 R3.0 as amended by ECN 3.0-3.0-001. If a Shelf Manager IPv6 Connection record in the Shelf FRU Information contains one or two IPv6 addresses or static router addresses, such definition(s) take(s) precedence over all other settings for the corresponding channel and is/are propagated to the LAN configuration parameters for the relevant channel. (Note that, unlike the IPv4 case, multiple IPv6 addresses can be assigned to a single network interface.)

According to the specification, multiple records can be specified to define settings separately for different LAN channels. In the Shelf Manager case, multiple Shelf Manager IPv6 Connection records are handled similarly to multiple Shelf Manager IPv4 Connection records (see section 3.4.2.2 for details). In particular, the first Shelf Manager IPv6 Connection record defines IPv6 settings for the first network interface (or for both network interfaces in the Active-Standby mode). The second record, if present, defines settings for the second network interface in Active-Active mode.

In addition, the command “Get Shelf Manager IP Addresses” can be used by a System Manager to retrieve information about IPv6 addresses (in addition to IPv4 addresses) of both the active and backup Shelf Managers.

See the *Shelf Manager External Interfaces Reference* for further details on the IPv6-related LAN configuration parameters.

3.4.11 Assigning IPv6 addresses to the Shelf Manager via DHCPv6 (ShMM-700 only)

Starting with release 3.6.0, the Shelf Manager supports DHCPv6 for automatic assignment of IPv6 addresses to the Shelf Manager. This support exists only on ShMM-700, since the Linux kernel on ShMM-500/1500 does not support IPv6.

The following types of IP addresses can be assigned via DHCP:

- RMCP accessible addresses (for one or both network interfaces)
- Private Shelf Manager addresses (for one or both network interfaces, for both Shelf Managers).

In DHCPv6, the Client Identifier is actually a DHCP Unique Identifier (DUID) for the client shelf, and the specific network interface is identified by a separate identifier (IAID). For the Shelf Manager, the DUID has the format shown in Table 11: HPM.3 DHCP Unique Identifier format. The *IPMI LAN Channel Index* for a network interface is used as the IAID for that interface, with range 0 through Ah and all other values reserved; see section 2.5.1 of HPM.3, DHCP-Assigned Platform Management Parameters Specification for more information.

DHCPv6 client functionality in the Shelf Manager is disabled by default. To enable it, set the configuration variable `USE_DHCPV6` to `TRUE` in the Shelf Manager configuration file.

As with DHCPv4, by default, the Shelf Manager uses the first DHCPv6 server that answers the *SOLICIT* request to assign the IPv6 addresses. If more than one DHCPv6 server is present in the network, the configuration parameter `DHCPV6_SERVER_ADDRESSES` can be used in the Shelf Manager configuration file `/etc/shelfman.conf` to specify the IPv6 addresses of the DHCP server to be used, as follows:

```
DHCPV6_SERVER_ADDRESSES = fe80::215:f2ff:fe67:81c6;fe80::218:49ff:fe01:8f8c
```

In that case, only DHCPv6 servers that have one of the specified addresses are used. Multiple DHCP server addresses can be specified in the value of this parameter; they must be separated by blanks, commas or semicolons.

In addition, also as in the DHCPv4 case, if the configuration parameter `DHCP_FOR_RMCP_ONLY` is set to `TRUE`, only the RMCP-accessible (Logical Shelf Manager) IP addresses are assigned via DHCPv6. In this case, the private ShMM IP addresses, if any, are left untouched.

One significant difference of DHCPv6 from DHCPv4 is that IPv6 addresses, assigned via DHCPv6, do not replace statically-assigned addresses but continue to coexist with them. Also, more than one IPv6 address can be assigned to a single network interface via DHCPv6. The total number of IPv6 addresses (both static and dynamic) assigned to a single physical network interface on ShMM-700 must not exceed 16, however. (This is a Linux kernel limitation.)

IPv6 addresses assigned via DHCPv6 are reported by the Shelf Manager as the values of the LAN configuration parameter IPv6 Dynamic (SLAAC/DHCPv6) Address (59), with “source/type” = 2 (DHCPv6). All IPv6 addresses assigned via DHCPv6 are RMCP-capable.

3.4.11.1 Handling DHCPv6 options

The Shelf Manager application does not currently explicitly request any DHCPv6 options from the server. The client sends and parses the options needed for the proper functionality of the DHCPv6 protocol (Client Identifier, Server Identifier, IA_NA, Status Code, Elapsed Time). A third-party application running on the ShMM-700 can, however, use the Shelf Manager DHCPv6 client to request options from the DHCPv6 server and handle their values.

The mechanism supporting this functionality involves the configuration parameter *DHCPV6_REQUESTED_OPTIONS*, which contains a comma-separated list of option numbers, and the configuration parameter *DHCPV6_OPTIONS_SCRIPT*, which provides the path to an executable file or a shell script that is spawned by the Shelf Manager executable when the values of the requested options arrive from the DHCPv6 server. The script receives these values, as well as the corresponding network interface name, DUID and IAID, in its parameters.

These configuration parameters can be specified, for example, as follows:

```
DHCPV6_REQUESTED_OPTIONS = "31,59,23,24,21,22,27,28,29,30,33,34,17"
```

```
DHCPV6_OPTIONS_SCRIPT = "/etc/dhcpv6cfg.sh"
```

The positional parameters to the script invoked by the Shelf Manager have the following meaning:

Table 13: Positional parameters

POSITION	MEANING
1	Path to a temporary file that contains received options and values in the format "name=value" on each line.
2	Network interface name (e.g. "eth0")
3	IAID, as a string of hexadecimal bytes (in network byte order)
4	DUID, as a string of hexadecimal bytes

The names of options in the temporary file are constructed as follows. For options that consist of several sub-values, several lines of output are included; a numeric suffix "_1", "_2", etc is appended to the option name on each subsequent line.

The format of the option values in the file depends on the option number and is shown in the following table.

Table 14: Option values format

OPTION NUMBER	MEANING	OPTION NAME IN THE FILE	VALUE FORMAT
21	SIP server names	SIP_DOMAIN_n	Text string
22	SIP server addresses	SIP_SERVER_n	IPv6 address
23	DNS server names	DNS_SERVER_n	Text string
24	Local domain names	DOMAIN_n	Text string
27	NIS server addresses	NIS_SERVER_n	IPv6 address
28	NISP server addresses	NISP_SERVER_n	IPv6 address

29	NIS server names	NIS_DOMAIN_n	Text string
30	NISP server names	NISP_DOMAIN_n	Text string
31	SNTP server addresses	SNTP_SERVER_n	IPv6 address
33	BCMS server names	BCMCS_DOMAIN_n	Text string
34	BCMS server addresses	BCMCS_SERVER_n	IPv6 address
59	Boot file URL	BOOTFILE_URL	Text string
	Other options	OPTION_n, where "n" is the decimal option number	String of hexadecimal bytes

Below is an example of a script that echoes the values of certain DHCPv6 options:

```
#!/bin/sh
echo ===== DHCPv6 Configuration script start
=====

cat $1
echo -----
. $1
echo Interface:      $2
echo IAID:           $3
echo DUID:           $4
echo BOOTFILE_URL:   $BOOTFILE_URL
echo SNTP_SERVER:    $SNTP_SERVER_1
echo DNS_SERVER:     $DNS_SERVER_1
echo DOMAIN_NAME_1:  $DOMAIN_1
echo DOMAIN_NAME_2:  $DOMAIN_2
echo SIP_SERVER_1:    $SIP_SERVER_1
echo SIP_SERVER_2:    $SIP_SERVER_2
echo SIP_DOMAIN_1:    $SIP_DOMAIN_1
echo SIP_DOMAIN_2:    $SIP_DOMAIN_2
echo NIS_DOMAIN:      $NIS_DOMAIN_1
echo NIS_SERVER:      $NIS_SERVER_1
echo NISP_DOMAIN:     $NISP_DOMAIN_1
echo NISP_SERVER:     $NISP_SERVER_1
echo BCMCS_DOMAIN:    $BCMCS_DOMAIN_1
echo BCMCS_SERVER:    $BCMCS_SERVER_1
echo OPTION_17:       $OPTION_17
echo ===== DHCPv6 Configuration script end
=====
```

And here is some example output of this script:

```
===== DHCPv6 Configuration script start =====
SNTP_SERVER_1="3ffe:501:ffff:100:200:ff:fe00:3f3e"
DNS_SERVER_1="3ffe:501:ffff:100:200:ff:fe00:3f3e"
```

```
DOMAIN_1="example.com"
DOMAIN_2="test.example.com"
SIP_DOMAIN_1="win.sip2.test.example.com"
SIP_DOMAIN_2="win.sip1.test.example.com"
SIP_SERVER_1="3ffe:501:ffff:100:200:ff:fe00:3f3d"
SIP_SERVER_2="3ffe:501:ffff:100:200:ff:fe00:3f3e"
NIS_SERVER_1="3ffe:501:ffff:100:200:ff:fe00:3f4e"
NISP_SERVER_1="3ffe:501:ffff:100:200:ff:fe00:3f4f"
NIS_DOMAIN_1="nis.test.example.com"
NISP_DOMAIN_1="nisp.test.example.com"
-----
Interface: eth1.8
IAID: 00000001
DUID:
00020000315A48504D2E332D31C454455354000000000000000000000000000000003000000
BOOTFILE_URL:
SNTP_SERVER: 3ffe:501:ffff:100:200:ff:fe00:3f3e
DNS_SERVER: 3ffe:501:ffff:100:200:ff:fe00:3f3e
DOMAIN_NAME_1: example.com
DOMAIN_NAME_2: test.example.com
SIP_SERVER_1: 3ffe:501:ffff:100:200:ff:fe00:3f3d
SIP_SERVER_2: 3ffe:501:ffff:100:200:ff:fe00:3f3e
SIP_DOMAIN_1: win.sip2.test.example.com
SIP_DOMAIN_2: win.sip1.test.example.com
NIS_DOMAIN: nis.test.example.com
NIS_SERVER: 3ffe:501:ffff:100:200:ff:fe00:3f4e
NISP_DOMAIN: nisp.test.example.com
NISP_SERVER: 3ffe:501:ffff:100:200:ff:fe00:3f4f
BCMCS_DOMAIN:
BCMCS_SERVER:
OPTION_17:
===== DHCPv6 Configuration script end =====
```

3.5 Configuring the FRU Information

This section describes configuring the Field Replaceable Unit (FRU) Information.

3.5.1 Accessing the Shelf FRU Information

According to the ATCA specification, the Shelf FRU Information should be redundant (at least two copies per shelf) and each copy may be represented by separate IPM controllers, as FRU #1. Some ATCA shelves adopt this approach fully. For such shelves, the default configuration file must be changed: the variable `LOCAL_SHELF_FRU` must be set to `FALSE`. This enables a shelf-wide search for potential sources of Shelf FRU Information on IPMB-0.

On most ATCA shelves implementing this approach, the Shelf FRU Information is accessed via two IPM controllers with known IPMB addresses. In this case, it is possible to limit the search for Shelf FRU Information to the two pre-determined IPMB locations.

To do this, two configuration variables `SHELF_FRU_IPMB_SOURCE1` and `SHELF_FRU_IPMB_SOURCE2` must be defined in the configuration file `/etc/shelfman.conf`. These variables are of the Number type and contain the IPMB addresses of the two designated IPM controllers that represent Shelf FRU Information. For example, to limit the search for the Shelf FRU Information to IPM controllers at 66h and 68h respectively, these variables should be defined as follows (note the use of “0x” prefix for hexadecimal addresses):

```
SHELF_FRU_IPMB_SOURCE1 = 0x66
SHELF_FRU_IPMB_SOURCE2 = 0x68
```

However, most known ShMM-based shelves provide two SEEPROMs that are connected to the Shelf Manager via the master-only I²C interface, usually with each SEEPROM residing on its own bus behind an I²C multiplexer.

Each of these two SEEPROMs stores a copy of the Shelf FRU Information, providing the needed redundancy. Some shelves do not provide even this type of storage; in that case it is possible to store the Shelf FRU Information on the ShMM itself, as a single Flash file `/var/nvdata/shelf_fru_info`.

The redundant Shelf Managers each have their own copy of that Flash file, and synchronize them using the redundancy protocol, so even in that case some degree of redundancy is still preserved.

The Shelf Manager configuration (as represented in `shelfman.conf`) must be aligned with the mechanisms for accessing Shelf FRU Information that are provided by the shelf. The default configuration supports the approach where redundant Shelf FRU Information instances are represented by separate IPM controllers as FRU #1 as well as the “two SEEPROMs” approach. The following key configuration variables are set as follows for that case:

```
LOCAL_SHELF_FRU = TRUE
SHELF_FRU_IN_EEPROM = TRUE
```

In this case, however, support for SEEPROMs must also be provided by the carrier-specific module in the Shelf Manager.

If none of the above approaches is supported by the shelf, and there are no sources of Shelf FRU Information represented by separate IPM controllers, the system integrator must resort to the Flash file as storage for Shelf FRU Information. In that case, the following changes to the default configuration should be done:

- set the variable `SHELF_FRU_IN_EEPROM` to `FALSE`
- set the variable `MIN_SHELF_FRUS` to `1`

The last change is necessary because there is only one copy of the Shelf FRU Information on each ShMM. The variable `LOCAL_SHELF_FRU` must retain its default value of `TRUE`.

The following table summarizes the configuration variable settings that correspond to the various Shelf FRU Information source possibilities described above.

Table 15: Shelf FRU Information and configuration variable settings

SOURCE OF THE SHELF FRU INFORMATION	SETTINGS OF CONFIGURATION VARIABLES
Non-volatile storage (likely EEPROMs), accessed via IPM controllers on IPMB-0	<p>LOCAL_SHELF_FRU = FALSE</p> <p>SHELF_FRU_IN_EEPROM = does not matter</p> <p>MIN_SHELF_FRUS = minimum number of IPM controllers on IPMB-0 providing the Shelf FRU Information (usually 2)</p> <p>Optionally:</p> <p>SHELF_FRU_IPMB_SOURCE1 = IPMB address of the first designated source</p> <p>SHELF_FRU_IPMB_SOURCE2 = IPMB address of the second designated source</p>
EEPROMs, accessed locally by the Shelf Manager	<p>LOCAL_SHELF_FRU = TRUE</p> <p>SHELF_FRU_IN_EEPROM = TRUE</p> <p>MIN_SHELF_FRUS = the number of EEPROMs providing the Shelf FRU Information (usually 2)</p>
Flash file	<p>LOCAL_SHELF_FRU = TRUE</p> <p>SHELF_FRU_IN_EEPROM = FALSE</p> <p>MIN_SHELF_FRUS = 1</p>

3.5.2 Setting up the Shelf FRU Information

Since the contents of the Shelf FRU Information is crucial for successful management of the shelf, it is necessary to set up the Shelf FRU Information on a fresh shelf before starting the Shelf Manager on it. This procedure consists of the following steps:

- Creating a description of the shelf in a formalized text format (“INF format”).
- Compiling the text description using the FRU Information Compiler.
- Placing the binary image of the FRU Information into the appropriate storage.

The first two steps are documented separately in the user manual for the FRU Information Compiler. The last step is documented here and depends on where the Shelf FRU Information is stored.

The simplest case is if the Shelf FRU Information is stored on a Flash file on the ShMM. In that case, the binary image file should be downloaded on the ShMM via FTP and copied to the location `/var/nvdata/shelf_fru_info`.

The following log represents an example of the above process. It should be performed on an x86 Linux machine (NOT on the ShMM, itself!).

```
# mv shelf_fru.bin shelf_fru_info
# ftp 192.168.1.230
Connected to 192.168.1.230.
```

```

220 shmm-230 FTP server (Version wu-2.6.2(1) Sun Dec 15 17:40:37 GMT 2002)
ready.
Name (192.168.1.230:serjio): ftp
331 Guest login ok, send your complete e-mail address as password.
Password:
230 Guest login ok, access restrictions apply.
Remote system type is UNIX.
Using binary mode to transfer files.
ftp> cd /var/nvdata
250 CWD command successful.
ftp> del shelf_fru_info
250 DELE command successful.
ftp> put shelf_fru_info
local: shelf_fru_info remote: shelf_fru_info
227 Entering Passive Mode (192,168,1,230,107,162)
150 Opening BINARY mode data connection for shelf_fru_info.
226 Transfer complete.
129 bytes sent in 5.1e-04 secs (2.3e+02 Kbytes/sec)
ftp> quit
221-You have transferred 129 bytes in 1 files.
221-Total traffic for this session was 640 bytes in 1 transfers.
221-Thank you for using the FTP service on shmm-230.
221 Goodbye.
#

```

Another case is when the Shelf FRU Information is stored in SEEPROMs that are accessible to the Shelf Manager via the master-only I²C bus. In that case, the binary file should be copied onto the ShMM and then written into the SEEPROMs. The utility *eeepromw* can be used for that purpose.

The exact location of SEEPROMs on the master I²C bus is carrier-specific, but typically they are located at address 0xA4 on channels 1 and 2 of the I²C multiplexer, and the multiplexer itself resides at address 0xE0. In that case, the following commands (issued on the ShMM!) can be used to download the file `/var/nvdata/shelf_fru_info` to the SEEPROMs. It is assumed that the file was compiled by the FRU Information Compiler and downloaded onto the ShMM, as in the previous example:

```

#eeepromw -c 1 A4 /var/nvdata/shelf_fru_info
#eeepromw -c 2 A4 /var/nvdata/shelf_fru_info

```

The general syntax for the *eeepromw* utility is as follows:

```

eeepromw [-b <multiplexer>] [-c <channel>] <eeeprom-address> <file>
[<count>]

```

where:

- `<multiplexer>` is the address of the I²C multiplexer on the I²C bus (default=0xE0),
- `<channel>` is the channel on the multiplexer to use (default=0),
- `<eeeprom-address>` is the address of the target SEEPROM,

- `<file>` is the path to the file to write onto the target SEEPROM.
- `<count>` is how many bytes to write.

The optional parameter `<count>` specifies how many bytes to write. If this parameter is not specified, all the contents of the file are written into the target SEEPROM. The total number of bytes written to SEEPROM does not exceed the size of the SEEPROM.

The reciprocal `eeepromr` utility allows the user to read the contents of an SEEPROM into a file on the Flash, and has the following parameters:

```
eeepromr [-b <multiplexer>] [-c <channel>] <eeeprom-address> <file>
<count>
```

where:

- `<file>` is the path to the file where data are written from SEEPROM,
- `<count>` is how many bytes to read.

The remaining parameters are the same as for `eeepromw`.

There is one caveat on using the utility `eeepromw`. The Shelf Manager must not be running when the SEEPROM is updated, because the Shelf Manager can change the I²C multiplexer settings trying to access an I²C device behind the multiplexer while the update is in progress. This will cause the update to fail with high probability.

There are several ways to ensure that the Shelf Manager is not running on the ShMM. One is to turn off automatic start of the Shelf Manager by setting the U-Boot variable `start_rc2_daemons` to `n`. The other way is to terminate the Shelf Manager using the command `clia terminate`.

On some shelves, access to Shelf FRU Information SEEPROMs is also guarded by the ACTIVE# signal on the ShMM. In that case, make sure before running `eeepromw` that the Hardware Redundancy Interface flags “Local Healthy” and “Active” are set for the local ShMM.

For the ShMM-500 and ShMM-1500, make sure that the CPLD control word has the following bits set:

- 0x02: Local Healthy
- 0x04: Local Switchover Request
- 0x20: Independent IPMB/Watchdog.

To turn all these bits on, use the following `cpld` command:

```
#cpld 26 26
```

For the ShMM-700, use the `hwri` command to set these flags and the “Local Presence” flag:

```
# hwri set presence healthy switchover
```

After that, make sure that the output of `hwri` reports the “Active” flag set.

The remaining case, where the Shelf FRU Information resides on separate IPM controllers inside the shelf, is completely shelf-specific and is beyond the scope of this document.

3.5.3 Setting up the Shelf FRU Information using the CLI

Because of the limitations and system dependencies of the approaches listed in the previous sections, another method is recommended to update the Shelf FRU Information. This method involves using the CLI command `frudataw` and it is applicable to all of the three cases considered in the previous section. This method requires that the Shelf FRU Information file to be placed on the active ShMM. After that, the following command can be used (possibly several times) to update the locations that contain the Shelf FRU Information:

```
clia frudataw <ipmc-address> <fru-id> <file>
```

where:

- `<ipmc-address>` is the address of the IPM controller that contains the Shelf FRU Information;
- `<fru-id>` is the FRU device ID of that location;
- `<file>` is the name of the file on the ShMM that contains the new Shelf FRU Information image.

The Shelf Manager must be running when this command is issued. After the Shelf FRU Information is updated, the shelf must be completely restarted to accommodate the changes.

To update the Shelf FRU Information stored in a Flash file with the contents of the file `newdata.bin`, update a single location with the IPMB address 20h and FRU device ID 1, like this:

```
#clia frudataw 20 1 newdata.bin
```

To update the Shelf FRU Information stored in SEEPROMs with the contents of the file `newdata.bin`, use two commands to update locations with the IPMB address 20h and FRU device IDs 1 and 2, like this:

```
#clia frudataw 20 1 newdata.bin
```

```
#clia frudataw 20 2 newdata.bin
```

To update the Shelf FRU Information stored on separate IPM controllers in the shelf, use several (usually two) commands with the corresponding IPMB addresses and FRU IDs. FRU ID is usually 1 in this case. Here is an example of updating the Shelf FRU Information in a shelf, where the Shelf FRU Information is stored on the IPMCs at 62h and 64h:

```
#clia frudataw 62 1 newdata.bin
```

```
#clia frudataw 64 1 newdata.bin
```

This method can be also used to update any other FRU Information repository in the shelf, if the IPMB address and the FRU device ID of that repository are known.

Note: The Shelf Manager doesn't support dynamic reconfiguration of the Shelf FRU Information. After any modifications of the Shelf FRU Information performed via CLI, RMCP, SNMP and other interfaces the whole shelf must be restarted to accommodate the changes.

3.5.4 Other FRU Information repositories

The Shelf Manager itself exposes at least one IPM controller (the ShMC at IPMB address 20h). For most carriers, the Shelf Manager also exposes a “physical” IPM controller that represents the

resources of the carrier board and has an IPMB address derived from the physical address of the carrier. While the ShMC is exposed only by the active Shelf Manager and is subject to switchover, the “physical” IPM controller is exposed separately by both active and backup Shelf Managers.

For both of these IPM controllers, FRU Information is stored in Flash files on the ShMM:

`/var/nvdata/bmc-fru-information` for the ShMC

`/var/nvdata/shelfman-fru-information` for the “physical” IPM controller

For some carriers, there may be additional FRUs represented by the ShMC.

The location of the FRU Information for these FRUs is carrier-specific.

Reading and writing these FRU Information repositories can be done via the IPMI commands “Read FRU Data”, “Write FRU Data”, addressed to the appropriate FRUs.

3.6 Configuring carrier and shelf attributes using HPDL

HPDL is a special language used to describe the structure and device population of specific AdvancedTCA shelves, ShMM carriers and non-intelligent subsidiary FRUs managed by the Shelf Manager (such as Fan trays, PEMs, alarm panels). This language is called HPDL (which is short for Pigeon Point Hardware Platform Description Language).

A full description of the HPDL is available to shelf developers in a separate document (the *Pigeon Point HPDL Reference*).

HPDL descriptions are organized in units called HPDL modules. Each module corresponds to a single text file. One module can describe one or more carriers and/or shelves. The HPDL modules are compiled by a special compiler (*hpdlc*) and stored as binary data. The encoding of the binary data conforms to the ASN.1 Basic Encoding standard (ISO 8825). The Shelf Manager reads and interprets this data during initialization and obtains the following information from it:

- Population of the master-only I²C bus on the carrier and shelf;
- Types and attributes of non-intelligent FRUs managed by the Shelf Manager;
- Information about IPMI sensors exposed by the Shelf Manager;
- Associations between IPMI sensors and physical signals exposed by master-only I²C devices.

Based on this information, the Shelf Manager configures master-only I²C devices, creates FRU descriptors for non-intelligent FRUs, creates IPMI sensors and associates them with the corresponding FRUs.

In addition to HPDL modules, collections of SDRs are used to configure IPMI sensors on the carrier and in the shelf. These collections of SDRs are compiled with the SDR compiler, stored as binary data and retrieved by the Shelf Manager during initialization. The Shelf Manager uses these saved SDRs for the sensors described in HPDL that it creates during initialization.

3.6.1 Compiling HPDL definitions

HPDL definitions in text format are compiled by the HPDL compiler into a binary format that conforms to the ASN.1 basic encoding rules (ISO 8825) and stored in an output file. The HPDL compiler executable is available both for Linux (x86) and Windows operating systems (with the

names *hpdlc* and *hpdlc.exe* respectively). The Linux version of the compiler can be run on various Linux distributions, including Red Hat, Mandriva Linux and Ubuntu.

The command line for the HPDL compiler has the following synopsis:

```
hpdlc [-d] [-v] <in-file> [ <out-file> ]
```

The parameter *<in-file>* specifies the input file that contains the HPDL descriptions in text format.

The optional parameter *<out-file>* specifies the output file; if this parameter is omitted, the name of the output file is based on the input file name with the suffix *.bin* appended at the end.

Option *-d* causes the HPDL compiler to output additional debugging information when compiling the input file.

Option *-v* causes the HPDL compiler to identify its version.

Error messages, if any, are output on the standard error stream of the HPDL compiler.

For example, the following command line compiles the input file *acb3.hpd1* into the binary file *acb3.hpd1.bin*:

```
>hpdlc acb3.hpd1
```

3.6.2 Compiling SDRs

Most of the attributes of sensors referenced in HPDL files are actually defined through Sensor Data Records (SDRs). The details of these SDRs are specified in text files and compiled to binary by a separate Python-based compiler. This compiler is available for both Linux and Windows operating systems; there is a separate *User Guide* for this compiler.

For example, the following command line compiles the input SDR file *acb3.sdr* into the binary SDR file *acb3.sdr.bin*:

```
>python sdrc.py acb3.sdr
```

3.6.3 Deploying HPDL data to the ShMM file system

The Shelf Manager first looks for HPDL data and SDRs in FRU Information areas (Shelf FRU Information for the shelf level data and carrier FRU Information for the ShMM carrier data). If the data is not found there, it is retrieved from the following files on the ShMM file system:

<i>/var/nvdata/carrier_data</i>	Binary HPDL definitions for the carrier
<i>/var/nvdata/carrier_sdrs</i>	Binary SDRs for the carrier sensors
<i>/var/nvdata/chassis_data</i>	Binary HPDL definitions for the shelf
<i>/var/nvdata/chassis_sdrs</i>	Binary SDRs for the shelf sensors

If the Shelf Manager cannot retrieve HPDL data from these locations, it uses the values of the following environment variables to locate the information:

<i>CARRIER_HPDL</i>	Binary HPDL definitions for the carrier
<i>CARRIER_SDRS</i>	Binary SDRs for the carrier sensors
<i>CHASSIS_HPDL</i>	Binary HPDL definitions for the shelf
<i>CHASSIS_SDRS</i>	Binary SDRs for the shelf sensors

Therefore, when HPDL definitions cannot be placed in some part of a FRU Information EEPROM, they are placed in the RFS and the environment variables listed above provide the locations inside the RFS (somewhere below the */etc* directory).

To override the default HPDL definitions stored in the RFS, a user should copy the files generated by the HPDL and SDR compilers to the directory */var/nvdata* on the ShMM file system by using, for example, FTP. The directory */var/nvdata* is in the persistent part of the ShMM file system; that is, the file content persists across ShMM reboots.

3.6.4 Deploying HPDL Data to FRU Information areas

FRU Information areas are the recommended deployment location for HPDL data and SDRs for the carrier and shelf/chassis (in carrier FRU Information and shelf FRU Information, respectively). FRU Information for a specific FRU (Fan tray, Power Entry Module) can also be used; it provides FRU-specific HPDL data that amends definitions for that FRU in the shelf/chassis HPDL description. FRU Information is usually stored in EEPROM-type devices.

HPDL data and SDRs are stored in the FRU Information as a sequence of OEM-type multirecords, similar to the multirecords defined in the AdvancedTCA specification. These multirecords have an nVent, Schroff GmbH (formerly Pentair Technical Products) Manufacturer ID (00400Ah) and the manufacturer-specific record type is 7 for HPDL data multirecords and 8 for SDR multirecords. The effective binary image for HPDL and SDR data is a concatenation of the contents of all multirecords of the corresponding type that are present in a given FRU Information.

Since in many cases the capacity of EEPROM devices used to store the FRU Information is small, HPDL data and SDRs can be placed into the FRU Information in a compressed form. The GZIP compression algorithm is supported. Compressed data is detected automatically when retrieved by the Shelf Manager from the FRU Information and the data is automatically uncompressed in that case.

To support this deployment approach, several new features have been added to the Pigeon Point FRU Compiler and to the Pigeon Point Shelf Manager CLI (Command Line Interface); these features enable the user to incorporate compiled HPDL data and SDRs into the FRU Information.

3.6.4.1 Specifying the location of the Carrier FRU Information

The preferred location for the Carrier FRU Information is an EEPROM on the ShMM carrier. In addition, however, the Shelf Manager supports getting the Carrier FRU Information from a file on the ShMM file system; this can be useful if the ShMM carrier lacks an appropriate EEPROM.

However, if the Carrier FRU Information is stored in an EEPROM, the location and type of this EEPROM needs to be communicated to the Shelf Manager. Carrier HPDL data cannot be used for that because of a “chicken and egg” problem – in order to read carrier HPDL data, the Shelf Manager has to already know the location or the Carrier FRU Information.

So, the location of the Carrier FRU Information (if it is stored in an EEPROM) is specified as a carrier option *CARRIER_FRU_LOCATION* for the HPDL carrier (part of the string value of the Shelf Manager configuration parameter *CARRIER_OPTIONS*).

If the Carrier FRU Information is located in a file on the ShMM file system, the carrier option *CARRIER_FRU_LOCATION* must not be present.

The following syntax governs this carrier option:

```
CARRIER_FRU_LOCATION = <device-type>:<bus>:<address>[,<size>]
```

where

- *<device-type>* is the EEPROM device type (See Section 3 of the *Pigeon Point HPDL Reference* for the list of supported EEPROM devices). Typical values are *AT24C16*, *AT24LC256*, *ADM1026*.
- *<bus>* is the number of the master-only I²C bus where the EEPROM is located (typically 0)
- *<address>* is the 7-bit I²C address of the EEPROM device on that bus, in hexadecimal
- *<size>* is an optional argument that specifies the size of the occupied portion of the EEPROM, that is, the number of bytes which the Shelf Manager should read to get the entire Carrier FRU Information. If this argument is not specified, the Shelf Manager reads the entire EEPROM.

For example, the following definition:

```
CARRIER_OPTIONS = "CARRIER_FRU_LOCATION = ADM1026:0:2E,2048"
```

specifies that the Carrier FRU Information should be retrieved from the EEPROM of the ADM1026 device that is located on bus 0 at address 2Eh (7-bit), and that only the first 2048 bytes of this EEPROM (which has the capacity of 8192 bytes) should be read.

3.6.4.2 Using the FRU Info Compiler to place HPDL and SDR data in FRU Information

The FRU Information Compiler is a Pigeon Point tool that generates a FRU Information image from its textual description. The image can then be stored in an appropriate EEPROM device using various techniques. The FRU Information Compiler implements special syntax that can be used to incorporate binary HPDL and SDR data into a FRU Information image that is being created.

For example, the following instructions add binary files *acb3.hpdl.bin.gz* and *acb3.sdr.bin.gz* (which are compressed HPDL and SDR files, respectively) to the current FRU Information image:

```
...
[nVent.HPDL Data]
File = acb3.hpdl.bin.gz

[nVent.HPDL SDR Data]
File = acb3.sdr.bin.gz
...
```

3.6.4.3 Using the Shelf Manager Command Line Interface

An extended version of the CLI command *frudataw* can be used to dynamically update any FRU Information in a shelf with HPDL data and SDRs. The corresponding binary files should be first downloaded on the ShMM.

The corresponding command line has the following syntax:

```
clia frudataw -d <ipmb-address> <fru-id> <in-file>
clia frudataw -s <ipmb-address> <fru-id> <in-file>
```

```
clia frudataw -d -c <ipmb-address> <fru-id>
clia frudataw -s -c <ipmb-address> <fru-id>
```

The parameters *<ipmb-address>* and *<fru-id>* specify the IPMB address and FRU device ID of the FRU Information to update. When updating HPDL data or SDRs in the Shelf FRU Information, use IPMB address 20h and FRU device ID 254, or the addresses of actual locations of the Shelf FRU Information.

The parameter *<in-file>* indicates the binary HPDL data or SDR file (which is possibly compressed) to be stored in the target FRU Information.

The option *-d* indicates that HPDL data is to be updated; the option *-s* indicates that SDRs data is to be updated.

The option *-c* removes (clears) the HPDL data or SDRs, respectively, from the target FRU Information.

For example, the following commands can be used to update the Shelf FRU Information with new versions of HPDL data and SDRs on a shelf where the Shelf FRU Information is stored in two redundant EEPROMs on the backplane, accessed locally by the Shelf Manager:

```
#clia frudataw -d 20 1 std14slot.hpdl.bin.gz
#clia frudataw -s 20 1 std14slot.sdr.bin.gz
#clia frudataw -d 20 2 std14slot.hpdl.bin.gz
#clia frudataw -s 20 2 std14slot.sdr.bin.gz
```

3.6.4.4 Placing HPDL data and SDRs in the FRU Information for specific FRUs

This section addresses placing HPDL data and SDRs into the FRU Information for a specific FRU (e.g. a fan tray). This information is used by the Shelf Manager to amend the attributes of that FRU from the attributes for that FRU that are specified in the HPDL chassis definition.

The following information about the FRU is supplied in the chassis definition and cannot be redefined by the HPDL data in the FRU:

- FRU site type and number
- Presence signal definition.

Device selection definitions can be redefined by the HPDL data in the FRU; however a problem arises when several slots exist in the chassis for the same type of the FRU, and addresses of devices of the FRU are different between different slots. In that case, it is not possible to define device location information on the FRU level; it is possible to do that only on the chassis level. This consideration limits the possibilities of redefining device information on the FRU level.

It is not necessary to have both HPDL data and SDRs in a given FRU Information area; only HPDL data or only SDRs may be present. For example, to change the power consumption of the FRU, it is sufficient to have only HPDL data. To change certain sensor attributes (e.g. the conversion formula, default threshold values or default hysteresis), it is sufficient to have only SDRs.

One problem with using HPDL data in a specific FRU is how to correlate the sensor numbers in the description to the actual sensor numbers. The sensor numbers are assigned within an IPM controller, not within a specific FRU, so the same sensors on different FRUs of the same type have different numbers; and it is typically not known in advance in which slot a specific FRU is going to

be placed. This problem is solved in the following way: the sensor numbers in the HPDL/SDR data for a specific FRU have relative numbers, and these numbers are correlated to the actual numbers in ascending order.

For example, assume that the replacement SDRs describe sensors 1, 2, 3 and 4, and the sensor numbers for the corresponding FRU in the HPDL chassis definition are 124, 208, 209 and 210. In that case, the replacement SDR for sensor 1 applies to sensor 124; the replacement SDR for sensor 2 applies to sensor 208, and so on.

3.7 Configuring the cooling management strategy

Different shelf vendors have different requirements for the cooling management strategy of the Shelf Manager. The PICMG 3.0 specification contains several requirements related to cooling management that must be satisfied by all vendors (e.g. the specification requires the Shelf Manager to deactivate a FRU that is in a critical thermal condition). However, some vendors have their own vendor-specific requirements that need to be satisfied for specific shelves (e.g. “turn on the Critical TELCO alarm when the thermal condition in the shelf becomes critical”). Until now, vendor-specific cooling management was implemented inside the Shelf Manager. With the advent of HPDL as a means to describe the architecture of carriers and shelves without additional coding, it became desirable to make cooling management modular, so that cooling management strategies can be added to the Shelf Manager without modification of the Shelf Manager code, and the proper cooling management strategy can be chosen dynamically.

To satisfy this goal, the following approach is supported in the Shelf Manager:

- Cooling management modules are implemented as Linux shared libraries (`*.so` files).
- A cooling management API is defined; this API fully defines the interface between the Shelf Manager and the cooling management module in the form of function calls.
- Support has been added to the Shelf Manager for dynamically loading the cooling management modules and binding to the functions that they export.

This approach is utilized only if HPDL is used to describe the carrier and the shelf. For non-HPDL-defined carriers and shelves, the cooling management strategy remains bound to carrier-specific code in the Shelf Manager.

3.7.1 Default cooling management strategy

With HPDL, a reasonable default cooling management strategy is provided. It is not necessary to configure the cooling management strategy if the default is satisfactory for a given carrier/shelf (which is likely the case for many carriers and shelves).

The default cooling management has the following features:

- zoned cooling is supported; the Shelf Fan Geography record from the Shelf FRU Information is used to define cooling zones.
- in the normal cooling state (that is, when no thermal thresholds on any sensors are crossed), the Shelf Manager attempts to minimize the fan level, but at the same time prevent thermal alerts. It does that by adaptively choosing the lowest possible fan level that allows the shelf to avoid thermal alerts, for each fan, taking cooling zones into account.

- In the minor alert cooling state (non-critical thermal thresholds are crossed for one or more sensors) the Shelf Manager periodically increases the fan level for the fans that serve the cooling zone(s) where those thresholds have been crossed, until the fan level reaches its maximum or the thermal condition goes away.
- In the major alert cooling state (critical thermal thresholds are crossed for one or more sensors) the Shelf Manager sets the fan level to the maximum for the fans that serve the cooling zone(s) where those thresholds have been crossed. In addition, if the thermal condition is caused by a specific FRU, and the FRU supports power levels lower than the current one, the Shelf Manager reduces power consumption of the FRU by assigning it the next lower power level.
- In the critical alert cooling state (non-recoverable thermal thresholds are crossed for one or more sensors) the Shelf Manager sets the fan level to maximum for the fans that serve the cooling zone(s) where those thresholds have been crossed. In addition, if the thermal condition is caused by a specific FRU, the FRU is powered down. If the thermal alert is caused by a shelf-wide temperature sensor, all FRUs are powered down, as prescribed by the PICMG 3.0 specification. After a FRU is powered down, its further handling by the Shelf Manager depends on the value of the configuration variable `COOLING_KEEP_POWERED_OFF_FRUS_IN_M1`.
- If this variable is set to `FALSE` (which is the default), the FRU is activated but kept in state M3 until the thermal alert goes away. The cooling state in which the FRU can be powered back is specified by the value of the configuration variable `REAPPLY_POWER_MAX_COOLING_STATE`. By default, this state is “Normal”; when the shelf cools down sufficiently and the shelf cooling state reaches the specified state, the Shelf Manager powers the FRU back on.
- Otherwise, if the configuration variable is set to `TRUE`, the FRU stays in the state M1 indefinitely and is not automatically powered back on when the thermal alert goes away; in that case an intervention by the System Manager or by an operator is needed to activate and power on the FRU.
- Also, when a new FRU is installed in the shelf when the shelf is in the critical alert cooling state, by default it is not powered up and stays in M3 until the critical alert goes away. This behavior can be overridden by the configuration variable `ALLOW_POWER_UNRELATED_FRU_IN_CRITICAL_STATE`. If this variable is set to `TRUE`, a new FRU can be powered up if: 1) the critical alert state is caused by a temperature sensor that belongs to a different FRU and 2) the critical alert state is not caused by a shelf-wide temperature sensor).
- In addition, a fan management strategy is implemented that sets the fan level to the maximum for all fans in the cooling zone in the following cases:
 - If some fans are missing in the shelf, based on the fan population specified in the Address Table in the Shelf FRU Information, all fan trays that cool the same zone(s) as the missing ones, are set to maximum.
 - If one or more of the fan tachometer sensors have a major or critical threshold crossed (a fan is stopped or rotates too slowly), all fan trays that cool the same zone(s) as the underspeed fan(s) are set to maximum speed. The number of simultaneously underspeed fans that cause the remaining fans to go full speed is configurable via the configuration parameter

`TURBO_MODE_MIN_FAN_FAILURES` (and the default value is 1). Setting this parameter to 0 turns off this feature.

- The default cooling management strategy optionally supports leaving cold-sensitive FRUs that are too cold at the time of shelf startup unpowered until they warm up. This may be useful for shelves installed in severe climate conditions. A special Pigeon Point OEM-specific record must be present in the Shelf FRU Information to activate this feature. This record contains the list of IPMB addresses and FRU IDs of FRUs considered cold-sensitive. A cold-sensitive FRU is considered too cold if any temperature sensor associated with it generates an event at startup that indicates that the temperature is below the lower non-recoverable threshold. Such a FRU is not immediately powered on, but kept in the state M3 until all of its temperature sensors are above their lower non-recoverable thresholds; then it is powered on.

3.7.2 Configuring a specific cooling management strategy

To configure a specific cooling management strategy, the following steps need to be taken:

1. Create a new cooling management strategy module (shared library) or choose one of the several libraries supplied with the Shelf Manager (in the RFS directory `/lib`), that are based on cooling strategies from certain vendors.
2. If a new module is created, perform the following steps:
 - 2.1. Implement the specific cooling management functionality, using the Shelf Manager Cooling Management API (which is described in a separate document).
 - 2.2. Choose a name for the shared library that conforms to the following pattern: `libcooling_<xxx>.so`, where `<xxx>` is the name of the strategy, typically the vendor name.
 - 2.3. Place the shared library either in the directory `/lib` or in the directory `/var/bin` on the ShMM.
3. Set the value of the configuration variable `COOLING_MANAGEMENT` in `/etc/shelfman.conf` to the name of the strategy (`<xxx>` above) and restart the Shelf Manager; thereafter, whenever it starts, the Shelf Manager loads and uses the designated shared library for cooling management.

3.8 Configuring local sensors

Note: This section applies only to shelves where HPDL is not used. With HPDL, sensor SDRs are defined as part of the carrier and chassis definition and do not require additional configuration from a file on the ShMM.

Local sensors on the ShMM can be configured when the ShMM is started. (This capability applies to sensors that are associated with either: 1) the Shelf Manager or 2) the physical IPM controller that takes its IPMB-0 address from the hardware address of the ShMM carrier slot.) Only sensor attributes that are defined in Sensor Data Records (such as thresholds, hysteresis values, sensor name, linearization parameters, etc.) can be configured at this time. The Sensor Device Records (SDRs) defining these sensors are read from the file `/var/nvdata/user_sdr`.

This file must contain an array of binary SDRs that are compliant with the IPMI specification. However, these SDRs can contain only partial sensor definition, if only a subset of the attributes of the sensor need to be redefined (see below).

The Pigeon Point SDR compiler utility can be used to produce the binary SDRs from plain text human-readable text files. The SDR compiler can also decode binary SDR data and produce a human-readable text file from it. This utility is described in the *SDR Compiler User Guide*. The current version of the SDR compiler (as distributed with release 2.2 or later of the Shelf Manager) must be used.

In order to take advantage of this sensor configuration facility, you should install the SDR Compiler on either a Linux or a MS Windows system. Then you should create an SDR definition text file according to the format described in a SDR Compiler User Guide.

You can use standard text editors such as *vi* for Linux or Notepad for MS Windows. The standard extension for SDR definition text files is *.inf*. The SDR Compiler is a command-line utility. To produce a binary SDR definition file *user_sdr* from a text SDR definition file you should use the utility in the compilation mode. For example:

```
>python sdrc.py test.inf user_sdr
```

user_sdr is an optional name of an output file. If it is not indicated in the command line a binary file *test.bin* will be created by the SDR Compiler.

The SDR compiler is written in Python, so you also need a Python interpreter (version 2.3 or later) to run it. Python is available for free downloading, with support for both Windows and Linux at www.python.org.

After producing the binary SDR definition file, you should place it on your ShMM in the directory */var/nvdata* under the name *user_sdr*. For example, you can use FTP:

```
>ftp 192.168.1.191
```

```
Connected to 192.168.1.191 (192.168.1.191).
```

```
220 shmm+191 FTP server (Version wu-2.6.2(1) Wed Oct 5 21:30:04 GMT 2005) ready.
```

```
Name (192.168.1.191:username) :anonymous
```

```
331 Guest login ok, send your complete e-mail address as password.
```

```
Password:
```

```
230 Guest login ok, access restrictions apply.
```

```
Remote system type is UNIX.
```

```
Using binary mode to transfer files.
```

```
ftp> cd /var/nvdata
```

```
250 CWD command successful.
```

```
ftp> put user_sdr
```

```
local: user_sdr remote: user_sdr
```

```
227 Entering Passive Mode (192,168,1,191,246,195)
```

```
150 Opening BINARY mode data connection for user_sdr.
```

```
226 Transfer complete.
```

```
124 bytes sent in 8.5e-05 secs (7.6e+03 Kbytes/sec)
```

```
ftp> quit
```

```
221-You have transferred 124 bytes in 1 files.
```

221-Total traffic for this session was 1165 bytes in 1 transfers.
 221-Thank you for using the FTP service on shmm+191.
 221 Goodbye.

A newly installed SDR definition file `/var/nvdata/user_sdr` takes effect only after a restart of the Shelf Manager; the Shelf Manager log during that restart could look like the following example:

```
# daemon -f shelfman -lcs
# <I> 03:45:09.041 [152] Pigeon Point Shelf Manager ver. 3.7.0. Built on Aug
27 2017 11:03:27
<*> 03:45:09.055 [152] Limits: code=(10af4040:10ba95c0), end_data=10c355d4,
start_stack=10c36e64, esp=10c36874, eip=10ba1838
<*> 03:45:09.063 [152] Stack limits: curr=800000, max=ffffffff
<*> 03:45:09.067 [152] Data limits: curr=ffffffff, max=ffffffff
<*> 03:45:09.090 [152] *** Lock log print buffer at 10c10e90 ***
<*> 03:45:09.094 [152] *** Pthread lock log print buffer at 10c14ed0 ***
<I> 03:45:09.380 [152] Enabling the CPLD Active bit workaround
<I> 03:45:09.435 [152] User SDR size = 124
<I> 03:45:09.453 [152] 20 # 120, Type = 1
<I> 03:45:09.467 [152] a0 # 2, Type = 1
...
```

If the SDR definition file `/var/nvdata/user_sdr` is not present, the configurations of the local sensors are unmodified from the default established in the Shelf Manager, and the log looks like the following example:

```
# daemon -f shelfman -lcs
# <I> 17:11:28.004 [522] Pigeon Point Shelf Manager ver. 3.7.0. Built on Aug
27 2017 19:03:33
<*> 17:11:28.009 [522] Limits: code=(400000:529030), end_data=10062000,
start_stack=7fff7df0, esp=7fff7758, eip=2ab0d2e4
<*> 17:11:28.010 [522] Stack limits: curr=1ff000, max=7fffffff
<*> 17:11:28.010 [522] Data limits: curr=7fffffff, max=7fffffff
<*> 17:11:28.014 [522] *** Lock log print buffer at 1003c910 ***
<*> 17:11:28.014 [522] *** Pthread lock log print buffer at 10040940 ***
<W> 17:11:28.027 [522] Custom SDR initialization file is absent
...
```

It is important to understand that replacement SDRs must be closely coordinated with SDRs that are defined within the Shelf Manager. Therefore, additions or modifications to the set of replacement SDRs in a shelf should only be undertaken in close cooperation with the shelf supplier.

The following rules apply to SDRs that are used to configure local sensors (referenced as replacement SDRs, below):

- Every replacement SDR must be a Full Sensor Record (type 01). Compact Sensor Records (type 02) are not supported and are never used by the Shelf Manager for local sensors.
- The following fields are mandatory in every replacement SDR (note that this list is smaller than the list of mandatory fields for normal SDRs processed by the SDR compiler):

- Sensor Owner ID
- Sensor Number
- Sensor Initialization
- If a field from a replacement SDR is used (as a result of the operation of other rules) to replace an attribute for a target sensor, and the field is not specified in the replacement SDR, a value zero (0) is used for this field.
- The tuple (Sensor Owner ID, Sensor Number) identifies the sensor that is to be configured. The Sensor Owner ID is a literal IPMB address; so for sensors on the physical IPM controller two instances of the corresponding replacement SDR must be present, one each for the IPMB addresses associated with each of the two redundant dedicated ShMC slots.
- The fields Entity ID and Entity Instance, if specified in the replacement SDRs and different from (0,0), replace the corresponding attributes of the sensor.
- The Sensor Initialization field identifies the specific attributes to be redefined and indicates what fields in the rest of the replacement SDR are applicable. It is specified as a list of symbolic constants; the following constants are defined:
 - THRESHOLDS: indicates that the sensor thresholds are to be replaced with the corresponding threshold values specified in the replacement SDR
 - HYSTERESIS: indicates that the sensor hysteresis values are to be replaced with the corresponding hysteresis values specified in the replacement SDR
 - SENSOR_TYPE: indicates that the sensor type is to be replaced with the corresponding fields Sensor Type and Event/Reading Type from the replacement SDR
 - EVENTS: indicates that the event masks are to be replaced with the corresponding fields from the replacement SDR
- The field Sensor Capabilities cannot be replaced.
- If and only if the symbolic constant SENSOR_TYPE has been specified in the Sensor Initialization field, the fields Sensor Type and Event/Reading Type from the replacement SDR replace the corresponding attribute of the sensor (even if the values of these fields are 0).
- If and only if the symbolic constant EVENTS has been specified in the Sensor Initialization field, the following fields from the replacement SDR replace the corresponding attributes of the target sensor:
 - For Threshold-Based sensors:
 - Lower Threshold Reading Mask
 - Upper Threshold Reading Mask
 - Threshold Assertion Event Mask
 - Threshold Deassertion Event Mask
 - Settable Threshold Mask
 - Readable Threshold Mask
 - For Discrete sensors:

- Assertion Event Mask
- Deassertion Event Mask
- Discrete Reading Mask
- The fields Sensor Units 1, Base Unit, Modifier Unit cannot currently be replaced.
- The following sensor attributes are replaced from the replacement SDR fields, if at least one of these fields is specified in the replacement SDR with a non-zero value:
 - Linearization
 - M
 - Tolerance
 - B
 - Accuracy
 - Accuracy exp
 - R exp
 - B exp
 - Analog Characteristic Flags
 - Nominal Reading
 - Normal Maximum
 - Normal Minimum
 - Sensor Maximum Reading
 - Sensor Minimum Reading
- If and only if the symbolic constant THRESHOLDS has been specified in the Sensor Initialization field, the fields listed below from replacement SDRs specify replacement threshold values for the target sensor. Note, however, that only thresholds supported by the sensor implementation can be redefined. (Some local sensors do not support all possible thresholds.) Here is a list of replaceable threshold types, all of which must be specified in the raw format (with the “0x” prefix):
 - Upper Non-Recoverable Threshold
 - Upper Critical Threshold
 - Lower Non-Critical Threshold
 - Lower Non-Recoverable Threshold
 - Lower Critical Threshold
 - Lower Non-Critical Threshold

If and only if the symbolic constant HYSTERESIS has been specified in the Sensor Initialization field, the following fields from the replacement SDRs specify the replacement hysteresis values for the target sensor:

 - Positive Hysteresis

- **Negative Hysteresis**
Hysteresis values must be specified in the raw format (with the “0x” prefix).
- The field OEM from the replacement SDR always replaces the corresponding attribute of the target sensor (even if not specified).
- The field Id String replaces the target sensor name only if specified in the replacement SDR.

The example below illustrates a typical local sensor configuration text definition. It redefines thresholds for the two temperature sensors on the physical IPM controller. It is assumed that the physical IPM controllers have IPMB-0 addresses 10h and 12h in the target shelf.

```
[Full Sensor Record]
Owner Id = 0x10
Sensor Number = 2
Sensor Initialization = THRESHOLDS
Lower Non-Critical Threshold = 0xb0
Lower Critical Threshold = 0xc0
Lower Non-Recoverable Threshold = 0xd0
Upper Non-Critical Threshold = 0x40
Upper Critical Threshold = 0x48
Upper Non-Recoverable Threshold = 0x50
[Full Sensor Record]
Owner Id = 0x10
Sensor Number = 3
Sensor Initialization = THRESHOLDS
Lower Non-Critical Threshold = 0xb0
Lower Critical Threshold = 0xc0
Lower Non-Recoverable Threshold = 0xd0
Upper Non-Critical Threshold = 0x40
Upper Critical Threshold = 0x48
Upper Non-Recoverable Threshold = 0x50
[Full Sensor Record]
Owner Id = 0x12
Sensor Number = 2
Sensor Initialization = THRESHOLDS
Lower Non-Critical Threshold = 0xb0
Lower Critical Threshold = 0xc0
Lower Non-Recoverable Threshold = 0xd0
Upper Non-Critical Threshold = 0x40
Upper Critical Threshold = 0x48
Upper Non-Recoverable Threshold = 0x50
[Full Sensor Record]
Owner Id = 0x12
Sensor Number = 3
Sensor Initialization = THRESHOLDS
```

```
Lower Non-Critical Threshold = 0xb0
Lower Critical Threshold = 0xc0
Lower Non-Recoverable Threshold = 0xd0
Upper Non-Critical Threshold = 0x40
Upper Critical Threshold = 0x48
Upper Non-Recoverable Threshold = 0x50
```

3.9 Setting the Auxiliary Firmware Revision

The Auxiliary Firmware Revision can be set when the Shelf Manager is started. The Auxiliary Firmware Revision is reported by the “Get Device ID” command targeted to a physical ShMC (at the hardware-specified IPMB-0 address, versus the logical Shelf Manager at IPMB-0 address 0x20) and is stored in a single Flash file `/var/nvdata/aux-fw-revision`. If the file `/var/nvdata/aux-fw-revision` is absent, the Auxiliary Firmware Revision is not defined.

According to IPMI v2.0 R1.0, Section 20.1 Get Device ID Command, the Auxiliary Firmware Revision is a 4-byte data item. The file `/var/nvdata/aux-fw-revision` should contain the eight hexadecimal digits that represent the 4 bytes in question. No separators are allowed. The first two hexadecimal digits represent the most significant byte of the Auxiliary Firmware Revision.

For example, assume that the file `/var/nvdata/aux-fw-revision` contains the string ‘a0b1efcd’. When the Shelf Manager is started and the RMCPTA connection is established, it is possible to obtain the Auxiliary Firmware Revision via RMCP. The parameter `<IPMB-address>`, below, represents the IPMB-address of the alternative controller. Here is a dialogue with the Pigeon Point internal tool RMCPTA, but any RMCP client can be used to make this query:

```
RMCPTA{1}-> TargetFwd <IPMB-address>
RMCPTA{1}-> GetDeviceId
Completion Code = 0x00 (OK)
Device ID      = 0x00
Device Revision = 0x0
Device Mode    = normal operation ; Device SDR present
Firmware Rev.  = 2.30
IPMI Version   = 1.5
Device Support = IPMB Req.Gen; FRU; Sensor;
Manufacturer ID = 0x0400A
Product ID     = 0x0000
AUX FW Rev.    = 0xA0B1EFCD (A0 B1 EF CD)
```

3.10 Setting up the clock

When the system is brought up for the first time, the clock is not set and must be initialized. Initially the clock is set to January 1, 1970. The date can be accessed via the serial console.

```
# date
Thu Jan 1 03:16:30 UTC 1970
```

In order to change the date, you should type in the correct date using the date application. The format for the date command is `MMDDhhmm[[CC]YY][.ss]`, where:

- MM - Month
- DD - Day
- hh - Hour (using 24 hour notation)
- mm - Minute
- ss - Second
- YY - Year (2-digit)
- CC - Century

For example:

```
# date 042916282006
Sat Apr 29 16:28:00 UTC 2006
```

To make the date persistent, you need to store it using the *hwclock* application.

```
# hwclock -w
```

In some cases, you might get the error message:

```
mktime: cannot convert RTC time to UNIX time
```

This error can be ignored. It is due to the original date being in an uninitialized state.

3.10.1 Obtaining date and time from a time server

It is possible to obtain the system date and time from a time server during system startup and synchronize it periodically thereafter. (This facility is critical if the ShMM carrier does not have an RTC battery.) There are two network time protocols that can be used for that purpose: NTP and RFC 868 (rdate). The specific protocol to be used is selected when configuring the ShMM.

To enable obtaining the network time via the NTP protocol, it is necessary to define the U-Boot variables *time_proto*, *time_server* and optionally the additional variable *timezone*. The variable *time_proto* determines the adjust time protocol (if this variable is undefined, by default the RFC 868 (rdate) protocol is used). This variable should be set to *ntp* to enable the NTP protocol. The usage of the other variables is identical their usage with RFC 868 (rdate) and is described below.

To enable obtaining the network time via the RFC 868 protocol (rdate) over TCP, it is necessary to define the U-Boot variable *time_server* and optionally the additional variable *timezone*. The variable *time_proto* should be left undefined or set to *rdate*.

The variable *time_server* contains the IP address of the time server that the Shelf Manager queries for the system time after startup. This server should support RFC 868 over TCP as required by the *rdate* utility or support NTP as required by the *ntpdate* utility. This variable is propagated to the Linux level as the environment variable *TIMESERVER*.

If this variable is set, the startup script */etc/netconfig* starts the script */etc/timesync* as a daemon, which runs in an endless loop and queries the time server with a default interval of 300 seconds. To change this interval, edit the script */etc/timesync* and change the value of the variable *INTERVAL*. If a query fails, the script starts polling the server more frequently (with the default interval of 30 seconds); this new interval is the value of the script variable *INTERVAL2*. When a subsequent query succeeds, the polling rate reverts back to *INTERVAL*. The

TIMESERVER variable can be changed by the Shelf Manager if the Shelf Manager is configured to receive its RMCP IP address via DHCP. In that case, depending on the configuration of the DHCP server, the Shelf Manager can also receive the IP address of the NTP server via DHCP. It stores this address in the */tmp/timeserverip* file, which is used by the script */etc/timesync* to override the current value of the *TIMESERVER* variable.

Note that during each reboot of the Shelf Manager the clock is reset to January 1, 1970. Therefore, the system time stays incorrect until the script */etc/timesync* succeeds in obtaining correct time from the time server.

The variable *timezone* contains the name of the current time zone followed by its offset from Greenwich Meridian Time (GMT). The offset is positive for time zones to the west of Greenwich and negative for time zones to the east of Greenwich. This variable is propagated to the Linux level as the environment variable *TZ*. The default value of this variable is *UTC*, i.e. Universal Coordinated Time which matches Greenwich time. Note that textual timestamps generated on a ShMM correspond to the local time zone. That is, these timestamps include the offset defined by the environment variable *TZ*.

The time sent by time servers is GMT time; if the time zone on the Shelf Manager is not set or not set correctly, the time obtained from the time server is interpreted incorrectly. The three-letter name of the time zone is not used by the Shelf Manager, but is propagated to set the Linux time zone. (For instance, if the time zone name XXX0 is used, the date command produces output like the following: "Thu Sep 9 21:24:24 XXX 2004".) Daylight saving time is not supported.

Here is an example of a time zone definition for US Eastern Time:

```
timezone = EST5
```

Here the digit 5 specifies that the time zone is 5 hours west of GMT. Any three letters can replace *EST*; they are used to identify the time zone in (for example) the Linux date command output.

For timezones with half-hour or quarter-hour differences from an integer offset (e.g. *UTC+05:30*), it is possible to use the */etc/localtime* file. Due to limited RFS space, there is no time zone database on ShMM, but such file can be copied from a Linux system (e.g. */usr/share/zoneinfo/Asia/Calcutta*) and installed on a ShMM as */etc/localtime*. If this file is detected at ShMM startup, the *TZ* variable is not set, so date-related library calls use the time zone specified in */etc/localtime*.

3.11 Setting up and using ShMM Power On Self Tests

The available Power On Self Test (POST) tests are built into U-Boot. The choice of the available POST tests that are actually executed is controlled by the dedicated U-Boot environment variables *post_normal* and *post_poweron*.

The value of the environment variable *post_normal* contains names of tests that are executed on each boot-up. These names are separated by space characters. These tests do not take much time and can be run on a regular basis.

The value of the environment variable *post_poweron* contains names of tests that are executed after power-on reset only (vs. on each boot-up). These names are separated by space characters.

As the POST tests are executed, the results are logged in a textual form in a dedicated area in SDRAM. Results for each particular test have the following form:

```
<4>POST <test name> [<test-specific output>] [PASSED|FAILED]
```

The POST framework provides interfaces for accessing the results of the POST tests in U-Boot and in Linux.

U-Boot supports a `log show` command, which can be used to access the POST test results. This command outputs the contents of the POST log buffer onto the serial console.

```
shmmxx00 log show
<4>POST uart UART 0 test failed
<4>FAILED
<4>POST crc PASSED
```

In a POST-enabled configuration, the Linux kernel shares its internal message log buffer with the POST log buffer. This causes POST results to be automatically displayed on the serial console during the kernel bootstrap. Interactively, the user can access the kernel log buffer (and thus the POST test results) using the `dmesg` command implemented by *busybox*.

The defined names for POST tests are:

- *memory* (SDRAM tests, recommended for execution on the first boot-up after power-on)
- *crc* (Flash checksum verification, recommended for execution on each boot-up; not available on ShMM-700)
- *uart* (UART verification, specific to the Au1550 processor, recommended for execution on each boot-up)
- *ethernet* (Ethernet MACs test, specific to the Au1550 processor, recommended for execution on each boot-up)
- *i2c* (Master-Only I²C test, recommended for execution on each boot-up)

The names of the tests can be used in values of the `post_poweron` and `post_normal` environment variables.

Starting from release 2.4.1, the IPMI command “Get Self Test Results” directed to the logical Shelf Manager (IPMB address 20h) returns the results of the POST performed by U-Boot at the startup of the ShMM. If all tests have passed, the status code 0x55 is returned. If any tests have failed, the device-specific failure code 0x59 is returned. The third byte of the response contains the following bit mask in that case:

- [7:5] Reserved
- [4] 1b = Ethernet test failed
- [3] 1b = UART test failed
- [2] 1b = U-Boot CRC test failed
- [1] 1b = I²C test failed
- [0] 1b = Memory test failed.

3.12 Configuring external event handling

Sometimes there is a need for specific IPMI events to be processed by a user-defined application or script that is external to the Shelf Manager, but still executing on the ShMM. This is called “external event handling”. This feature allows extensions of Shelf Manager functionality by user programs or scripts executed on the ShMM.

External event handling is implemented via the PEF alerting mechanism. As an extension of this mechanism, all alerts with a destination channel set to the System Interface Channel (Fh) are delivered by the Shelf Manager to the standard input of a designated external event handler program. This program is specified via the configuration parameter

`EXTERNAL_EVENT_HANDLER`.

3.12.1 Detailed steps to configure external event handling

To use external event handling the user should take the following steps:

1. Prepare a script or executable for external event handling and place it on the Shelf Manager Flash filesystem (for example, into the directory `/var/bin`).

2. Specify the location of the external event handler in the Shelf Manager configuration file (`/etc/shelfman.conf`) as the value of the configuration parameter

`EXTERNAL_EVENT_HANDLER`, for example:

```
EXTERNAL_EVENT_HANDLER = /var/bin/ext_handler.sh
```

If the configuration parameter `EXTERNAL_EVENT_HANDLER` is not defined or is defined as an empty value, external event handling is disabled.

3. Start the Shelf Manager.

4. Enable PEF via the CLI, if not enabled yet:

```
# clia setpefconfig control 1
```

5. Enable alerting, if not enabled yet:

```
# clia setpefconfig action_control 1
```

6. Set alert strings if necessary:

```
# clia setpefconfig alert_string 1 "test string"
```

In the example above, `1` is the number of the string; to set multiple strings, use sequential string numbers `1`, `2`, etc. for subsequent strings.

7. Set at least one alert policy that sends alerts to the destination channel Fh (the IPMI System Interface):

```
# clia setpefconfig alert_policy 1 1 8 F 1 1
```

Where the fields in the example have the following meanings (see the *Pigeon Point Shelf Manager External Interface Reference* document for detailed descriptions of the `setpefconfig` and other CLI commands):

- `1` - number of alert policy table entry
- `1` - policy number
- `8` - policy (enabled, always send alert to this destination)

- *F* - destination channel
- *1* - destination selector
- *1* - alert string selector

8. Set event filters for relevant events. For example, the following event filter can be used to capture hot swap transitions to state M0 of FRU #0 on IPMC 9Ch:

```
# clia setpefconfig event_filter 1 80 1 1 0 9C FF F0 FF FF FF FF 0F FF 0 0 0 0
FF FF 0
```

where:

- *1* - filter number
- *80* - filter configuration (enabled, software configurable filter)
- *1* - filter action - alert
- *1* - alert policy number
- *0* - event severity (unspecified)
- *9C* - Slave Address or Software ID from Event Message.
- *FF* - channel Number / LUN to match
- *F0* - type of sensor
- *FF* - sensor
- *FF* - event trigger
- *FF* - Event Data 1 Event Offset Mask
- *FF* - Event Data 1 Event Offset Mask
- *0F* - Event Data 1 AND Mask
- *FF* - Event Data 1 Compare 1
- *0* - Event Data 1 Compare 2
- *0* - Event Data 2 AND Mask
- *0* - Event Data 2 Compare 1
- *0* - Event Data 2 Compare 2
- *FF* - Event Data 3 AND Mask
- *FF* - Event Data 3 Compare 1
- *0* - Event Data 3 Compare 2

Note: when an event handler is placed in the directory /var/bin as assumed in the example above, care must be taken to preserve the state of that directory across reliable firmware upgrades (see section 7.6 for details). Also note that the /var/bin directory is erased when the U-Boot variable flash_reset is set to y (see section 6.3 for details).

3.12.2 External event handler operation

The external event handler can be a user-provided binary executable program or a shell script. The handler is invoked by the Shelf Manager when an event matching a relevant PEF filter occurs. During that invocation, a pipe is created between the Shelf Manager and the external event handler. This pipe serves as standard input for the external event handler.

For each event, the Shelf Manager submits a single line of text into the pipe in the following format:

```
raw="raw data" alert_string="alert string data"
```

The raw data are the 16 bytes of the SEL record, with the characters representing each byte separated by a colon (":") from the next byte. For example:

```
raw="36:3:2:27:df:49:0:20:0:4:f1:18:6f:a1:0:78" alert_string="test"
```

The external event handler should read its standard input line by line and process each line as a separate event.

There are possible two approaches to event processing in the external event handler:

- Read the standard input line by line and process incoming data. In that case, one instance of the external event handler is spawned in the beginning and handles all subsequent events.
- Read one line from the standard input, process the data on that line and exit. In that case, the Shelf Manager spawns a new instance of the external event handler for the next event.

Here is an example of a simple external event handler script:

```
#!/bin/sh

while read line; do
    raw=""
    eval $line;
    if [ ! -z "$raw" ]; then
        # do event processing here
        echo "event $raw!"
    fi
done
```

The following external event handler script example is more sophisticated. It powers down either IPMC or the entire chassis on Upper Non-Recoverable threshold crossing event from any voltage sensor, based on event source address. For sensors at BMC 20h the entire chassis is powered down, otherwise only the event source IPMC is powered down. The "raw" variable contains 16 bytes of a SEL entry defined in IPMI 2.0 specification and the script extracts source IPMB address, sensor type (02h for voltage), event direction and reading type (01h for asserted threshold events) and threshold type (lower half of the event data byte#1) from it.

```
#!/bin/sh

LOGGER="logger -t EVENTHAND -p 6 -s"

while read line; do
    raw=""
```

```
eval $line;
if [ "$raw" ]; then
    # $LOGGER "raw=$raw"
    source=`echo $raw | cut -d: -f 8`
    sentype=`echo $raw | cut -d: -f 11`
    direction=`echo $raw | cut -d: -f 13`
    threshold=`echo $raw | cut -d: -f 14 | cut -c2`
    $LOGGER "source=$source sentype=$sentype direction=$direction
threshold=$threshold"
    # voltage sensor type & asserted event & UNR threshold
    if [ "$sentype" = "2" -a "$direction" = "1" -a "$threshold" = "b" ]; then
        if [ "$source" = "20" ]; then
            $LOGGER "Powering down the chassis via IPMI:ChassisControl(power-down)"
            clia sendcmd 20 0 2 0
        else
            $LOGGER "Deactivating IPMC 0x$source via
PICMG:SetFruActivate(deactivate)"
            clia deactivate $source 0
        fi
    fi
fi
done
```

3.13 Configuring the Platform Event Trap format

The Shelf Manager sends SNMP Traps as a result of event processing when an Alert action is initiated or as a result of the “Alert Immediate” command which is used to test Alert destinations. The Platform Event Trap format is controlled by the *PET_FORMAT* configuration variable. Currently there are three supported formats. All formats use the same enterprise OID = iso(1).org(3).dod(6).internet(1).private(4).enterprises(1).wired_for_management(3183).PET(1).version(1), but have different OIDs for the embedded variable(s). A separate MIB file for supported SNMP traps is available.

By default or with *PET_FORMAT* set to 0, the Shelf Manager generates trap messages according to IPMI Platform Event Trap Format v1.0 specification. In this format, event data is packed as binary in a single variable with OID=.1.3.6.1.4.1.3183.1.1.1. A sample application *snmptc* is available for parsing this format.

If *PET_FORMAT* is set to 1, the Shelf Manager generates trap messages in plain text format, packing event into single variable with OID=.1.3.6.1.4.1.3183.1.1.2 as an ASCII string. This format is convenient for visual monitoring without specialized tools.

If *PET_FORMAT* is set to 2, Shelf Manager generates trap messages in multi-variable format, packing each field of the event (event number, timestamp, generator etc.) as a separate variable with OIDs in the .1.3.6.1.4.1.3183.1.1.3.[1..13] range. Processing traps in this format requires the MIB file.

It is possible to change the Platform Event Trap format dynamically using the *clia setpefconfig pet_format <N>* command with *<N> = 0..2*. The command *clia getpefconfig pet_format* shows the currently used PET format.

Here is some sample output of the standard Linux *snmptrapd* tool for the supported formats:

PET_FORMAT = 0

```
Received 111 bytes from 192.168.0.2
192.168.0.2: Enterprise Specific Trap (2453248) Uptime: 191 days,
4:23:00.02,
PPS-PET-MIB::ipmi-trap-data = Hex-STRING: BF CE 7A 60 AE 50 11 DC 00 80
00 50 C2 3F CD 24 00 01 12 BE 74 62 00 00 20 20 02 20 C1 00 00 00 FF FF 00 00 00
00 00 19 0A 40 00 00 00 00 C1
```

PET_FORMAT = 1

```
Received 251 bytes from 192.168.0.2
192.168.0.2: Enterprise Specific Trap (2453248) Uptime: 0:00:41.04,
PPS-PET-MIB::ipmi-trap-text = STRING: "RecordID:0004;
RecordType:Platform Event; Time:Dec 19 17:11:43 2007; Generator:20 LUN0 Chan0;
SensorType:25; SensorNumber:C1; EventType:6F; EventDirection:asserted;
EventData: 00 FF FF"
```

PET_FORMAT = 2

```
Received 296 bytes from 192.168.0.2
192.168.0.2: Enterprise Specific Trap (2453248) Uptime: 489 days,
11:31:17.15,
PPS-PET-MIB::ipmi-trap-record-id = INTEGER: 4,
PPS-PET-MIB::ipmi-trap-record-type = INTEGER:
platformEvent(2),
PPS-PET-MIB::ipmi-trap-timestamp = INTEGER: 1198084292,
PPS-PET-MIB::ipmi-trap-generator-address = INTEGER: bmc(32),
PPS-PET-MIB::ipmi-trap-generator-lun = INTEGER: lun0(0),
PPS-PET-MIB::ipmi-trap-generator-channel = INTEGER: ipmb(0),
PPS-PET-MIB::ipmi-trap-sensor-type = INTEGER:
entityPresence(37),
PPS-PET-MIB::ipmi-trap-sensor-number = INTEGER: 193,
PPS-PET-MIB::ipmi-trap-event-type = INTEGER:
sensorSpecific(111),
PPS-PET-MIB::ipmi-trap-event-direction = INTEGER: asserted(0),
PPS-PET-MIB::ipmi-trap-event-data = Hex-STRING: 00 FF FF ,
PPS-PET-MIB::ipmi-trap-entire-record = Hex-STRING: 04 00 02 C4 50 69
47 20 00 04 25 C1 6F 00 FF FF
```

3.13.1 Parsed example of SNMP Trap

The following is a parsed version of sample SNMP Trap in the default (IPMI PET v1.0) format. The original event which is reported by this trap follows:

```
Event: at Tue Aug 5 16:29:21 2008; from:(0x20,0,0); sensor:(0x12,0x85);
event:0x6f(asserted): 0xC0 0x00 0x00
```

This trap describes a system-reconfigure event, which is generated during Shelf Manager startup and switchover.

In the example below, an italicized bold hexadecimal represents the Tag of a field and an underlined hexadecimal represents the length of a field.

```
---message length 125-----
30 7B 02 01 00 04 06 70 75 62 6C 69 63 A4 6E 06
09 2B 06 01 04 01 98 6F 01 01 40 04 C0 A8 00 02
02 01 06 02 03 12 6F 00 43 04 61 A3 ED 13 30 4D
30 4B 06 0A 2B 06 01 04 01 98 6F 01 01 01 04 3D
1D 3A 30 76 62 F2 11 DD 00 80 00 50 C2 3F FB 56
00 07 13 ED A3 61 00 00 20 20 01 20 85 00 00 C0
00 00 00 00 00 00 00 19 00 00 40 0A 00 00 80 4B
01 54 65 73 74 53 74 72 69 6E 67 00 C1
-----
```

Here is a detailed description of the content of this trap:

30 7B - ASN.1 tag and remaining length;
02 01 00 - SNMP version field, len = 1, 0=SNMPv1;
04 06 70 75 62 6C 69 63 - community field, "public";
A4 6E - PDU Trap command tag and remaining length;
06 09 2B 06 01 04 01 98 6F 01 01 - Agent ID, Enterprise OID(.1.3.6.1.4.1.3183.1.1);
40 04 C0 A8 00 02 - Agent IPv4 (192.168.0.2);
02 01 06 - Generic Status, len=1, EnterpriseSpecific;
02 03 12 6F 00 - Specific Status, len=3, (SenType=12h, EventType=6Fh, EventOffset=00h, Direction=Asserted);
43 04 61 A3 ED 13 - Time in ticks (0.01sec), obviously wrong...;
30 4D - ASN.1 tag at variable bindings start and remaining length;
30 4B - First variable header and remaining length;
06 0A 2B 06 01 04 01 98 6F 01 01 01 - Variable OID(.1.3.6.1.4.1.3183.1.1.1);
04 3D - Variable data length, data follows (see Table 17-7 "PET Variable Bindings Field" of IPMI 2.0 specification);
1D 3A 30 76 62 F2 11 DD 00 80 00 50 C2 3F FB 56 - GUID;
00 07 - trap sequence;
13 ED A3 61 00 00 - Local Timestamp (Tue Aug 5 16:29:21 2008);
20 20 - trap source = event source = IPMI;
01 - Event Severity (Monitor);
20 - Sensor Device (BMC)
85 - Sensor Number;
00 - Entity (00h=unspecified)
00 - Entity Instance (00h=unspecified_
C0 00 00 00 00 00 00 00 - Event Data;
19 - Language Code (English);

00 00 40 0A - Manufacturer ID, MSB first (00400Ah[†] assigned to nVent, Schroff GmbH (formerly Pentair Technical Products));
 00 00 - System ID, MSB first;
 80 4B 01 - PET multirecord, ASCII, len=Bh, Text Alert String;
 54 65 73 74 53 74 72 69 6E 67 00 - Null terminated Alert String "TestString";
 c1 - No more fields;

3.14 Configuring the IntegralHPI interface

The Pigeon Point Shelf Manager optionally includes IntegralHPI, an implementation of the Service Availability Forum (SAF, www.saforum.org) Hardware Platform Interface (HPI), operating as a subsystem within the Shelf Manager. On the ShMM-500 and ShMM-1500, IntegralHPI requires a variant of the ShMM that has additional memory: specifically, 128 megabytes of RAM and 64 megabytes of Flash. These memory resources needed for IntegralHPI are present on all ShMM-700s, so no special ShMM configuration requirements exist for IntegralHPI.

IntegralHPI is implemented as a shared library within the Shelf Manager and is turned off by default. To turn it on, it is necessary to set the configuration variable `ENABLE_INTEGRALHPI` to `TRUE`, and ensure that the corresponding shared library `libintegralhpi.so` (and several other libraries that IntegralHPI depends on) is present on the ShMM. A special RFS image containing those libraries is installed on your shelf if it is enabled for IntegralHPI support.

3.14.1 HPI domain support in IntegralHPI

IntegralHPI exposes the ATCA shelf managed by the (potentially redundant) Shelf Manager as a single HPI domain. This domain has Domain ID=1 by default. In order to allow a single HPI client to manage multiple ATCA shelves, the domain ID for a given shelf can be configured via the Shelf Manager configuration variable `INTEGRALHPI_DOMAIN_ID`.

3.14.2 HPI SNMP subagent support

The IntegralHPI-enabled RFS includes a public domain HPI SNMP subagent that allows a remote client to communicate with IntegralHPI using the SNMP protocol. The mapping between the HPI interface and the SNMP protocol was specified by the Service Availability Forum (SAF, <http://www.saforum.org>) some time ago. However, currently neither the mapping definition nor the subagent software itself is actively supported. In order to automatically start the HPI SNMP subagent at system startup, the Shelf Manager configuration variable `RUN_HPI_SNMP_SUBAGENT` should be set to `TRUE`.

3.14.3 IntegralHPI client configuration

The HPI interface is a C language function call interface. To use it across a network boundary, a remote procedure call (RPC) protocol must be used. IntegralHPI is compatible with the RPC protocol implemented in OpenHPI (www.openhpi.org). This means that OpenHPI client utilities and libraries can be used to communicate with IntegralHPI in the same way as they can communicate with an OpenHPI daemon. IntegralHPI accepts an OpenHPI library connection to the TCP port 4743, which is the default for an OpenHPI daemon.

[†]This IANA Enterprise Number was originally assigned to Pigeon Point Systems, which has since been acquired by Pentair.

Currently, the recommended approach for client configuration is the following:

- Download OpenHPI distribution (release 2.10, 2.12 ,2.14 or 2.16) from www.openhpi.org, configure it, build and install on the client system. If you have access to Pigeon Point OpenHPI (a special distribution of OpenHPI) for either of the above OpenHPI releases, you can use it instead of the original distribution.
- Set the environment variable `OPENHPI_DAEMON_HOST` to the RMCP IP address of the target Shelf Manager.
- Use the OpenHPI client library (which is located by default in `/usr/local/lib/libopenhpi.so`) to link to client applications.

Note: The OpenHPI library `libopenhpi.so.2` from the Pigeon Point OpenHPI 2.12.0.1 release or from the Pigeon Point OpenHPI 2.14.0.0 release and higher is strongly recommended for use with client HPI applications since it contains the redundancy framework improvement to avoid the potential loss of an HPI call that is being processed during a Shelf Manager switchover. The OpenHPI library `libopenhpi.so.2` in the Shelf Manager RFS already contains this improvement. Also note that the posted open source OpenHPI versions 2.16 and higher already include the necessary library updates.

3.15 Configuring system services

The Shelf Manager file system contains several network services that can be used for debug and upgrade purposes. The following table contains a list of ports used by the Shelf Manager with a brief description of the service configuration with each port:

Table 16: Network services

PORT NUMBER/TYPE	NETWORK SERVICE DESCRIPTION
21/tcp	<p><i>FTP server</i></p> <p>Activated via the <code>/etc/inetd.conf</code> file; to disable the FTP server, comment out the 'ftp' line in that file. This service is based on <i>WU FTP</i> daemon, so the configuration is located in the <code>/etc/ftp*</code> files, mostly in <code>/etc/ftpaccess</code>.</p> <p><i>Note: On ShMM-700 only, the FTP server can be configured to listen to IPv6 sockets. To achieve that, it is necessary to replace 'tcp' with 'tcp6' in the corresponding line in the <code>/etc/inetd.conf</code> file:</i></p> <pre>ftp streamtcp6 nowaitroot /bin/in.ftpd -a</pre> <p><i>This option is turned off by default, because the FTP server does not correctly show the peer IP address in the system logs in the case of an IPv6 connection.</i></p>
22/tcp 22/tcp6	<p><i>SSH server</i></p> <p>Activated via the <code>/etc/inetd.conf</code> file; to disable the SSH server, comment out the 'ssh' line in that file. Run <code>sshd --help</code> to see the list of supported arguments that can be specified.</p> <p>On ShMM-700 only, the SSH server listens to both IPv6 and IPv4 sockets, so SSH sessions to a ShMM-700 can be established over both IPv6 and IPv4.</p>

PORT NUMBER/TYPE	NETWORK SERVICE DESCRIPTION
23/tcp	<p><i>Telnet server</i></p> <p>Activated via the <code>/etc/inetd.conf</code> file; to disable the Telnet server, comment out the <code>'telnet'</code> line in that file. Run <code>telnetd --help</code> to see the list of supported arguments that can be specified.</p> <p><i>Note: On ShMM-700 only, the telnet server can be configured to listen to IPv6 sockets. To achieve that, it is necessary to replace 'tcp' with 'tcp6' in the corresponding line in the <code>/etc/inetd.conf</code> file:</i></p> <pre><i>telnetstreamtcp6 nowaitroot /bin/telnetd</i></pre> <p><i>This option is turned off by default, because the telnet server does not correctly show the peer IP address in the system logs in the case of an IPv6 connection.</i></p>
80/tcp	<p><i>HTTP server with the ShM Web interface</i></p> <p>Activated by the <code>/etc/rc.common</code> script if the U-Boot variable <code>start_rc2_daemons</code> is set to <code>y</code>; to disable the HTTP server, comment out the corresponding line in the script. The HTTP server supports the configuration file <code>/etc/httpd.conf</code> where it is possible to specify allowed and denied IP ranges and the access login and password, for example:</p> <pre>A:192.16 # Allow any address that begins with 192.16. D:* # Deny from other IP connections /cgi-bin:username:password</pre> <p>where password is generated via the <code>httpd -m password</code> command.</p>

PORT NUMBER/TYPE	NETWORK SERVICE DESCRIPTION
161/udp 161/udp6	<p><i>SNMP server with ShM SNMP interface</i></p> <p>Activated by the <code>/etc/rc.common</code> script if the U-Boot variable <code>start_rc2_daemons</code> is set to <code>y</code>; to disable the SNMP server, comment out the corresponding line in the script. The SNMP server supports the configuration file <code>/etc/snmpd.conf</code>. The validation criteria for incoming SNMP requests are hardcoded in this configuration file. For SNMP v1/v2c, the community string in an incoming request must be <code>public</code>. For SNMP version 3 requests, the user name in an incoming request must be <code>overlord</code>. These parameters can, however, be changed by editing the configuration file <code>/etc/snmpd.conf</code>. On ShMM-700, the SNMP server handles requests over both IPv6 and IPv4. To disable IPv6 support, comment out the <code>agentaddress</code> and <code>com2sec6</code> directives in the configuration file.</p>
514/tcp	<p><i>RSH server</i></p> <p>Activated by the <code>/etc/inetd.conf</code> file; to disable the RSH server, comment out the <code>'shell'</code> line in that file. Run <code>in.rshd --help</code> to see list of supported arguments that can be specified.</p>
623/udp	<p><i>ShM RMCP interface</i></p> <p>Activated by the <code>/etc/rc.common</code> script if the U-Boot variable <code>start_rc2_daemons</code> is set to <code>y</code>; to disable the RMCP interface, comment out the corresponding line in the script. This interface is available only on the active ShMM and only if the Shelf Manager is operational. The Shelf Manager only accepts connections targeted to configured RMCP address(es).</p>
1040/tcp	<p><i>ShM Redundancy interface</i></p> <p>This is the inter-ShMM communication interface needed in redundant configuration. It only accepts connection targeted to the IP address specified as the configuration parameter <code>REDUNDANT_IP_ADDRESS</code> in the Shelf Manager configuration file.</p>

PORT NUMBER/TYPE	NETWORK SERVICE DESCRIPTION
4743/tcp 4743/tcp6	<p><i>HPI server</i></p> <p>This service is available when an instance of an on-ShMM OpenHPI or IntegralHPI running on the ShMM. IntegralHPI is activated as a part of the Shelf Manager if the configuration variable <code>ENABLE_INTEGRALHPI</code> is set to <code>TRUE</code> and the HPI dynamic library is present. OpenHPI (on-ShMM version) is started by the <code>/etc/rc.common</code> script, so it is possible to comment out corresponding lines in that script to disable it. On ShMM-700 only, the HPI server listens to both IPv6 and IPv4 sockets, so HPI sessions with a ShMM-700 can be established over both IPv6 and IPv4.</p>

3.16 Configuring the speed of I²C buses (ShMM-700 only)

Two of the ShMM-700 I²C buses (with logical numbers 2 and 3) support both high speed (400KHz) and normal speed (100KHz) operation. By default, both buses operate in high-speed mode, but this can be changed by means of a Linux kernel startup parameter `i2c-mxs.speed`, which has the following syntax:

```
i2c-mxs.speed=<speed1>,<speed2>
```

where `<speed1>` and `<speed2>` can have values `100` or `400` and represent the speed of buses 2 and 3, respectively.

According to normal Linux kernel naming conventions, the name of the parameter is composed of two parts separated by the dot symbol: `i2c-mxs` corresponds to the I²C driver name, and `speed` is the parameter name within that driver.

For example, to set both buses to normal speed, the following parameter should be passed to the Linux kernel:

```
i2c-mxs.speed=100,100
```

The recommended method to pass this parameter to the kernel is to append it to the value of the U-Boot variable `bootargs_initrd`, for example, by typing the following at the U-Boot prompt:

```
shmm700 setenv bootargs_initrd $(bootargs_initrd) i2c-mxs.speed=100,100
shmm700 saveenv
```

Note: On ShMM-700, to provide compatibility with designs based on previous generation of ShMMs, bus mapping takes place on the OS level. For example, logical I²C bus 0 (represented by the device node `/dev/i2c0`) that goes to the ShMM-700R carrier and accommodates carrier-based I²C devices, is actually physical I²C bus 4. Logical I²C bus 1 (represented by the device node `/dev/i2c1`) that accommodates on-ShMM I²C devices, is actually physical I²C bus 5. The details of the mapping are shown in the table below.

Table 17: OS level bus number mapping for ShMM-700

LOGICAL BUS NUMBER (/DEV/I2C<N>)	PHYSICAL BUS NUMBER (ACCORDING TO THE ShMM-700R HARDWARE SPECIFICATION)
0	4
1	5
2	0
3	1
4	2
5	3

4 Using the Shelf Manager

This section introduces the overall operation of the Shelf Manager, including operation in a redundant (active/standby) pair.

4.1 In this section

This section contains the topics listed below. Just click on a topic to go to it.

- ShMM login
- Starting the Shelf Manager
- Redundant operation
- Operation in shelves with radial IPMB-0
- Automatic SEL truncation
- Cooling state sensors
- Deadlock detection
- I²C Bus fault isolation
- Assignment of LAN configuration parameters to boards and modules
- HPI System Event sensor
- ShMM tests available via the diagnostic infrastructure

4.2 ShMM login

Once the ShMM has been fully booted, you are prompted to login. You can log in as user root; the default password, as set in the factory-provided RFS images, is the empty password "". We highly recommend that you change passwords during the configuration of the ShMM. If the password is forgotten, it can be reset as described in section 6.4. The following is a typical log-in session for ShMM-500 (which would look similar on ShMM-1500 and ShMM-700):

```
shmm500 login: root
Password:
# ls /
bin          dev          etc          lib          mnt
proc         sbin        tmp          usr          var
##
```

4.3 Starting the Shelf Manager

The Shelf Manager software is implemented in the executable file *shelfman* in the directory */bin*. During normal use of the Shelf Manager it is invoked automatically by startup scripts. Typical users never need to invoke it manually as described here. Nevertheless, for the unusual circumstances in which manual invocation may be necessary, the details are described here.

The syntax of a Shelf Manager command line invocation is defined as follows:

```
shelfman [<options>] &
```


or

```
daemon -f shelfman [<options>]
```

The following options are recognized:

```
-h <address>
-v <verbosity>
-c <path>
-cs[-]
-w[-]
-wt
-l[c][s]
-g <ip_address>
-sf[-]
-port <port>
-ph[=<IPMB-addr1>,<IPMB-addr2>]
-i <command_line>.
```

A detailed description of these options is given below.

-h <address>

This option overrides the hardware address of the FRU site where the Shelf Manager resides. This hardware address is used as the IPMB address for the Shelf Manager (the address that is based on its hardware address, separate from 20h) and is treated as a hexadecimal number. This option can be used if no hardware address is automatically available for the Shelf Manager. If this option is not used, the Shelf Manager obtains the hardware address in a carrier-specific way; the IPMB address is the hardware address multiplied by 2.

-v <verbosity>

Set the initial debug verbosity mask (see section 3.3.2), both for the system log and for the console. If this option is not present, the values of the parameters `VERBOSITY` and `VERBOSITY_CONSOLE` in the file `/etc/shelfman.conf` determine the initial debug verbosity mask.

-c <path>

Path to non-volatile configuration files. The default path is `/var/nvdata`.

-cs[-]

Enforce (`-cs`) / disable (`-cs-`) checking Shelf FRU Information checksums. If a checksum is invalid, an error message is produced and the Shelf FRU Information is not used. If a checksum is valid, no message is produced. If the option `-cs` is not present in the command line, the value of the parameter `VERIFY_SHELF_FRU_CHECKSUM` from the file `/etc/shelfman.conf` determines Shelf Manager behavior in this area.

-w[-] | -wt

Enable (*-w*) / disable (*-w-*) the watchdog timer. If neither of the *-w-* or *-wt* options is present in the command line, the value of the parameter *WATCHDOG_ENABLED* from the file */etc/shelfman.conf* determines Shelf Manager actions in this area.

The *-wt* option puts the watchdog timer in test mode: that is, the actual timer is not turned on or strobed, but the strobing thread in the Shelf Manager still runs and a warning message is printed if the interval between subsequent strobes exceeds 500 ms. (In normal operation, if this interval exceeds one second, the Shelf Manager is reset.)

-l[c][s]

Set logging destination: use *s* to choose syslog and *c* to choose console. If this option is not specified, the values of the parameters *CONSOLE_LOGGING_ENABLED* and *SYSLOG_LOGGING_ENABLED* from the file */etc/shelfman.conf* determine Shelf Manager actions in this area.

-g <ip_address>

Set the default gateway IP address. This address is only used if no gateway address is set in the LAN Configuration Parameters for channel 1.

-sf[-]

The option *-sf* forces the Shelf Manager to use EEPROMs for Shelf FRU Information storage. The option *-sf-* forces the Shelf Manager not to use EEPROMs for Shelf FRU Information storage. If neither of these options is present in the command line, the value of the parameter *SHELF_FRU_IN_EEPROM* from the file */etc/shelfman.conf* determines Shelf Manager actions in this area.

-p <port>

This option sets the redundancy communication port. If this option is not present, the value of the parameter *REDUNDANCY_PORT* from the file */etc/shelfman.conf* determines the Shelf Manager actions in this area.

-ph[=<IPMB-addr1>,<IPMB-addr2>]

This option defines IPMB addresses of pseudo-hubs (by default, 0x82 and 0x84). Pseudo-hubs are virtual IPM controllers that are created and emulated by the Shelf Manager and behave like IPM controllers for Base Interface hub boards. This option can be used for testing and when non-intelligent hub boards (non-compliant boards that do not implement an IPM controller) are installed in the shelf.

-i <command_line>

This option defines the command line for the initialization script that is invoked by the active Shelf Manager during initialization. This initialization script can be used for platform-specific initialization. The CPLD bits Local Healthy and Active are set when the command line is invoked and the Shelf Manager waits for completion of the initialization script before progressing further. Note that the initialization script is run only at the startup of the active Shelf Manager, and is not called after a switchover. The corresponding command line can also be specified as the value of the configuration parameter *INITIALIZATION_SCRIPT* in the Shelf Manager configuration file.

To run the Shelf Manager from the command line, type the following:

daemon -f shelfman [<options>]

For example, here is a typical log of the Shelf Manager starting with options that cause it to route log information both to the console and to the log file and to use EEPROMs for Shelf FRU Information:

```
# daemon -f shelfman -lcs -sf
```

```
<*> 14:35:17.338 [199] Pigeon Point Shelf Manager ver. 3.7.0. Built on Aug 27
2017 22:00:07
<*> 14:35:17.348 [199] *** Lock log print buffer at 1006c720 ***
<*> 14:35:17.350 [199] *** Pthread lock log print buffer at 10070f70 ***
<I> 14:35:17.362 [199] Reading configuration file: /etc/shelfman.conf
<I> 14:35:17.384 [199] Shelfman: Using EEPROM for Shelf FRU
<I> 14:35:17.392 [199] Set watchdog timeout to 2 seconds, error=0
<W> 14:35:18.810 [199] Custom SDR initialization file is absent
<I> 14:35:18.812 [199] Carrier set to "GENERIC"
<I> 14:35:18.818 [199] [shm_crypto_init] initialization succeeded
<I> 14:35:18.824 [199] Device GUID: {B4A841B0-5EBD-11DE-0080-001849000EEC}
<I> 14:35:18.829 [199] Netmask not specified in two-adapter redundancy
configuration: defaulting to 255.255.255.128
<I> 14:35:18.834 [199] Interface usb1 activated successfully
<I> 14:35:18.836 [199] Input redundancy socket successfully bound to:
192.168.0.192:1040
<I> 14:35:18.838 [199] Output redundancy socket successfully bound to:
192.168.0.64:1041
<I> 14:35:18.841 [199] Active listening thread created successfully
<I> 14:35:18.843 [199] Connecting to: 192.168.0.65:1040
<I> 14:35:21.844 [199] *** Running in Active mode (connection to backup
failed: -148) ***
<I> 14:35:21.846 [199] Redundancy protocol initialized successfully
<I> 14:35:21.883 [199] Reboot reason sensor was registered successfully, value
8
<I> 14:35:21.951 [199] Shelfman: Registered "Local Temp." sensor { 0:3 } at
0x2C ...
<E> 14:35:21.995 [199] ADM1026: Ext 1 temp sensor not implemented: -128,
ignored
<E> 14:35:22.045 [199] ADM1026: Ext 2 temp sensor not implemented: -128,
ignored
<I> 14:35:22.049 [199] Shelfman: Registered "3.3STBY voltage" sensor { 0:4 }
at 0x2C ...
<I> 14:35:22.054 [199] Shelfman: Registered "3.3MAIN voltage" sensor { 0:5 }
at 0x2C ...
<I> 14:35:22.058 [199] Shelfman: Registered "+5V voltage" sensor { 0:6 } at
0x2C ...
<I> 14:35:22.062 [199] Shelfman: Registered "Vcpp voltage" sensor { 0:7 } at
0x2C ...
<I> 14:35:22.067 [199] Shelfman: Registered "+12V voltage" sensor { 0:8 } at
0x2C ...
<I> 14:35:22.071 [199] Shelfman: Registered "-12V voltage" sensor { 0:9 } at
0x2C ...
```

```

<I> 14:35:22.188 [199] FT 0 FRU successfully registered as FRU 02
<I> 14:35:22.190 [199] Initializing Fans
<I> 14:35:22.192 [199] Registering fan RD facility
<I> 14:35:22.194 [199] Activating fan tray 0, fan level=5, power level=0
<I> 14:35:22.197 [199] Controller 20, FRU 2: ATCA state set to M1, prev=M0,
cause=0, locked=0
<I> 14:35:22.203 [199] Controller 20, FRU 2: updating blue LED for M1
<I> 14:35:22.206 [199] Carrier version set to 82:64:01:00
<I> 14:35:22.243 [199] Alarm sensor registered
<I> 14:35:22.245 [199] Alarm Input Sensor: reading=0, initial update 0
<I> 14:35:22.247 [199] Alarm input sensor registered
<I> 14:35:22.253 [199] SEL: initialized SEL compression resources,
max_sel_entries = 1024
<I> 14:35:22.256 [199] SEL: compression is ON, journaling is ON
<I> 14:35:22.285 [205] SEL dedicated write thread: started
<I> 14:35:22.291 [206] SEL truncation thread pid started successfully
<I> 14:35:22.294 [199] ShM SEL Activation complete
<I> 14:35:22.352 [208] PEF thread: starting
<W> 14:35:22.355 [199] External event handler is not defined, so PEF external
handler thread is not started
<I> 14:35:22.357 [199] PEF activated successfully
<I> 14:35:22.359 [199] PEF initialized successfully, System Event Sensor #133
<I> 14:35:22.363 [199] Chassis facility activated successfully
<I> 14:35:22.366 [199] Chassis facility initialized successfully
<I> 14:35:22.371 [199] Registering shelf FRU notification: 0x2abd05f8
<I> 14:35:22.416 [199] Session: session descriptor size = 996 bytes
<I> 14:35:22.490 [215] Started incoming message thread, hc=1
<I> 14:35:22.493 [215] Status change on IPMB 0: flags = 186
(ONL,ATT,ONL_CH,ATT_CH)
<I> 14:35:22.495 [215] IPMB 0 went online
<I> 14:35:22.498 [215] Status change on IPMB 1: flags = 186
(ONL,ATT,ONL_CH,ATT_CH)
<I> 14:35:22.500 [215] IPMB 1 went online
<I> 14:35:22.524 [216] IPMC Stored write thread started
<I> 14:35:22.663 [199] Starting power thread
<I> 14:35:22.681 [199] Controller FC, FRU 0: ATCA state set to M1, prev=M0,
cause=0, locked=0
<I> 14:35:22.689 [199] Created IPMC for 84, exists=0
<I> 14:35:22.691 [199] Get Stored Device ID OK for 84
<I> 14:35:22.693 [199] Operational state for SA 84, FRU 0 is set to M7 from
stored SDR!
<I> 14:35:22.701 [199] SDR Repository: broadcasting Get Device ID
<I> 14:35:22.763 [224] CLI subscribing for deadlocks, pid=224
<I> 14:35:22.784 [224] Registering 224 facility to the triggering watchdog
<I> 14:35:22.806 [225] FRU data write thread started successfully

```

```

<I> 14:35:22.834    [199] Generating list of the trusted Shelf FRU Information
sources
<I> 14:35:22.836    [199] FRU 1 on 0x20 IPMC was added to the Shelf FRU trusted
list
<I> 14:35:22.842    [205] Creating new SEL Journal...
<I> 14:35:22.850    [199] Controller 20, FRU 0: ATCA state set to M1, prev=M0,
cause=0, locked=0
<I> 14:35:22.852    [199] Before_init_fru_hs_state!
<I> 14:35:22.854    [199] set the FRU 1 initial state
<I> 14:35:22.856    [199] Controller 20, FRU 1: ATCA state set to M1, prev=M0,
cause=0, locked=0
<I> 14:35:22.859    [199] Registering shelf FRU notification: 0x4eb18c
<I> 14:35:22.861    [199] Registering shelf FRU notification: 0x2ab44fdc
<I> 14:35:22.863    [199] Shelfman: Running.
<I> 14:35:22.865    [226] Shelf FRU: registering for SDR Repository
notifications!
<I> 14:35:22.873    [226] Added 20#1 to the fru list, treated as Shelf FRU Info
storage, size = 2048 (data size = 529)
<I> 14:35:22.875    [226] 1 Equal Shelf FRUs have been detected
<I> 14:35:22.877    [226] Calling shelf FRU notifications, count=1, op=1
<I> 14:35:22.879    [226] Calling shelf FRU notification: 0x2abd05f8
<I> 14:35:22.881    [226] Calling shelf FRU notification: 0x4eb18c
<I> 14:35:22.883    [226] Lan apply cfg parameters: HW addr 0xFC, cross 0,
use_second 0, adapter eth0, adapter2
<I> 14:35:22.886    [226] Lan apply cfg parameters: IP C601A8C0, GW FD01A8C0,
MASK 00FFFFFF
<W> 14:35:22.888    [226] Shelf FRU Info found: failed to get IP Connection
record from Shelf FRU, err=-61
<I> 14:35:22.891    [226] LAN: channel 1 address 192.168.1.198:623
<I> 14:35:22.893    [226] LAN: setting RMCP IP address to 192.168.1.198
<I> 14:35:22.914    [212] Controller 20, FRU 0: ATCA state set to M2, prev=M1,
cause=2, locked=0
<I> 14:35:22.916    [212] Controller FC, FRU 0: ATCA state set to M2, prev=M1,
cause=2, locked=0
<I> 14:35:22.935    [226] LAN: setting RMCP subnet mask to 255.255.255.0
<I> 14:35:22.939    [226] [channel_param_start_channel_on_backup] channel 1,
propagate 0
<I> 14:35:22.941    [226] DHCP Client is disabled
<I> 14:35:22.943    [226] Calling shelf FRU notification: 0x2ab44fdc
<I> 14:35:22.945    [226] LAN Config Parameters data in the Shelf FRU 0 bytes
<I> 14:35:22.948    [226] Found Power Distribution record instance=0, size=49
<I> 14:35:22.950    [226] Allowance for FRU Activation Readiness = 10 seconds,
start=23
<I> 14:35:22.952    [226] Found Power Management record, format version=1
<I> 14:35:22.954    [226] Retrieving initial power budget
<I> 14:35:22.975    [227] RMCP: starting server thread for 192.168.1.198:623

```

```

<I> 14:35:22.977 [227] RMCP: started server thread for 192.168.1.198:623
<I> 14:35:22.979 [227] Registering 227 facility to the triggering watchdog
<I> 14:35:23.585 [214] Hot Swap event: SA=0x20 FRU 2 sensor_number = 0x03, M0-
>M1, cause 0
<I> 14:35:23.586 [214] State M1, sa=20, FRU=2, op_state=M0, power_cycle=0
<I> 14:35:23.647 [214] Hot Swap event: SA=0xfc FRU 0 sensor_number = 0x00, M0-
>M1, cause 0
<I> 14:35:23.649 [214] State M1, sa=FC, FRU=0, op_state=M0, power_cycle=0
<I> 14:35:23.659 [214] Hot Swap event: SA=0x20 FRU 2 sensor_number = 0x03, M0-
>M1, cause 0
<I> 14:35:23.661 [214] State M1, sa=20, FRU=2, op_state=M1, power_cycle=0
<I> 14:35:23.667 [214] Hot Swap event: SA=0xfc FRU 0 sensor_number = 0x00, M0-
>M1, cause 0
<I> 14:35:23.669 [214] State M1, sa=FC, FRU=0, op_state=M1, power_cycle=0
<I> 14:35:23.729 [214] Hot Swap event: SA=0x20 FRU 0 sensor_number = 0x00, M0-
>M1, cause 0
<I> 14:35:23.731 [214] State M1, sa=20, FRU=0, op_state=M0, power_cycle=0
<I> 14:35:23.743 [214] Hot Swap event: SA=0x20 FRU 1 sensor_number = 0x02, M0-
>M1, cause 0
<I> 14:35:23.745 [214] State M1, sa=20, FRU=1, op_state=M0, power_cycle=0
<I> 14:35:23.755 [214] Hot Swap event: SA=0x20 FRU 0 sensor_number = 0x00, M1-
>M2, cause 2
<I> 14:35:23.763 [217] [ipmc_cooling_scan_sensors] Reread sensors on the SA
0x20
<I> 14:35:23.767 [228] SA 0x20 FRU 0 is ACTIVATING
<I> 14:35:23.779 [214] Hot Swap event: SA=0xfc FRU 0 sensor_number = 0x00, M1-
>M2, cause 2
<I> 14:35:23.787 [217] [ipmc_cooling_scan_sensors] Reread sensors on the SA
0xFC
<I> 14:35:23.801 [228] Tasklet ACTIVATE (20, 0): is_local_address 1
<I> 14:35:23.811 [221] Initial indicator for SA 0x20: 27.(state new)
<I> 14:35:23.814 [221] Initial indicator for SA 0xfc: 33.(state new)
<I> 14:35:23.817 [221] sdrrep: fetched new SDRs, time=6
<I> 14:35:23.821 [213] Controller 20, FRU 0: ATCA state set to M3, prev=M2,
cause=1, locked=0
<I> 14:35:23.824 [213] Controller 20, FRU 0: ATCA state set to M4, prev=M3,
cause=0, locked=0
<I> 14:35:23.828 [214] Hot Swap event: SA=0x20 FRU 0 sensor_number = 0x00, M2-
>M3, cause 1
<I> 14:35:23.835 [217] [ipmc_cooling_scan_sensors] Reread sensors on the SA
0x20
<I> 14:35:23.842 [217] [ipmc_cooling_scan_sensors] Reread sensors on the SA
0x20
<I> 14:35:23.847 [214] Hot Swap event: SA=0x20 FRU 0 sensor_number = 0x00, M3-
>M4, cause 0
<I> 14:35:23.853 [217] [ipmc_cooling_scan_sensors] Reread sensors on the SA
0x20

```

```

<I> 14:35:23.864 [229] SA 0xfc FRU 0 is ACTIVATING
<I> 14:35:23.881 [218] Adjust power for 20, 00: from 0 to 5000
<I> 14:35:23.882 [218] Sending sync point
<I> 14:35:23.891 [233] SA 0x20 FRU 0 is OPERATIONAL
<I> 14:35:23.905 [229] Tasklet ACTIVATE (fc, 0): is_local_address 1
<W> 14:35:23.907 [229] Local FRU 0 at SA 0xFC is not listed in Shelf FRU Info;
still activate
<I> 14:35:23.918 [212] Controller 20, FRU 1: ATCA state set to M2, prev=M1,
cause=2, locked=0
<I> 14:35:23.920 [212] Controller 20, FRU 2: ATCA state set to M2, prev=M1,
cause=2, locked=0
<I> 14:35:23.942 [221] sdrrep: Skipping 20, read already in progress
<I> 14:35:23.944 [221] sdrrep: Skipping FC, read already in progress
<I> 14:35:23.947 [214] Hot Swap event: SA=0x20 FRU 1 sensor_number = 0x02, M1-
>M2, cause 2
<I> 14:35:23.963 [214] Hot Swap event: SA=0x20 FRU 2 sensor_number = 0x03, M1-
>M2, cause 2
<I> 14:35:23.973 [234] SA 0x20 FRU 1 is ACTIVATING
<I> 14:35:23.978 [230] sdrrep_full_ipmc_sdrs_update_thread for 20
<I> 14:35:23.979 [213] Controller FC, FRU 0: ATCA state set to M3, prev=M2,
cause=1, locked=0
<I> 14:35:23.984 [230] sdrrep_full_ipmc_sdrs_update_thread for 20: attempt 0
<I> 14:35:23.992 [230] Registered potential fan 0x100eb2b8, SA 0x20 FRU 2
<I> 14:35:23.997 [230] sdrrep: finished reading SDRs for SA=20, added: 15,
preserved: 0, deleted: 0; time=7 seconds
<I> 14:35:24.003 [214] Hot Swap event: SA=0xfc FRU 0 sensor_number = 0x00, M2-
>M3, cause 1
<I> 14:35:24.021 [231] sdrrep_full_ipmc_sdrs_update_thread for FC
<I> 14:35:24.024 [235] SA 0x20 FRU 2 is ACTIVATING
<I> 14:35:24.031 [217] Move fan 0x100eb2b8, cnt 2 0->1
<I> 14:35:24.031 [217] [ipmc_cooling_scan_sensors] Reread sensors on the SA
0xFC
<I> 14:35:24.023 [231] sdrrep_full_ipmc_sdrs_update_thread for FC: attempt 0
<I> 14:35:24.071 [218] Adjust power for FC, 00: from 0 to 20000
<I> 14:35:24.073 [218] Sending sync point
<I> 14:35:24.077 [213] Controller FC, FRU 0: ATCA state set to M4, prev=M3,
cause=0, locked=0
<I> 14:35:24.081 [214] Hot Swap event: SA=0xfc FRU 0 sensor_number = 0x00, M3-
>M4, cause 0
<I> 14:35:24.097 [236] Set Port State (disable): ipmc=FC, chan=1, it=0,
ports=1, lt=1, ext=0, group=0
<I> 14:35:24.106 [222] IPMC 20 # 1 is in the list already. Exiting.
<I> 14:35:24.109 [234] Tasklet ACTIVATE (20, 1): is_local_address 1
<W> 14:35:24.111 [234] Local FRU 1 at SA 0x20 is not listed in Shelf FRU Info;
still activate

```

```
<I> 14:35:24.115    [236] Set Port State (disable): ipmc=FC, chan=2, it=0,
ports=1, lt=1, ext=0, group=0
<I> 14:35:24.130    [217] [ipmc_cooling_scan_sensors] Reread sensors on the SA
0xFC
<I> 14:35:24.134    [213] Controller 20, FRU 1: ATCA state set to M3, prev=M2,
cause=1, locked=0
<I> 14:35:24.137    [213] Controller 20, FRU 1: ATCA state set to M4, prev=M3,
cause=0, locked=0
<I> 14:35:24.147    [231] sdrrep: finished reading SDRs for SA=FC, added: 29,
preserved: 0, deleted: 0; time=7 seconds
<I> 14:35:24.152    [214] Hot Swap event: SA=0x20 FRU 1 sensor_number = 0x02, M2-
>M3, cause 1
<I> 14:35:24.167    [236] SA 0xfc FRU 0 is OPERATIONAL
<I> 14:35:24.179    [214] Hot Swap event: SA=0x20 FRU 1 sensor_number = 0x02, M3-
>M4, cause 0
<I> 14:35:24.211    [235] Tasklet ACTIVATE (20, 2): is_local_address 1
<W> 14:35:24.213    [235] Local FRU 2 at SA 0x20 is not listed in Shelf FRU Info;
still activate
<I> 14:35:24.236    [213] Controller 20, FRU 2: ATCA state set to M3, prev=M2,
cause=1, locked=0
<I> 14:35:24.241    [214] Hot Swap event: SA=0x20 FRU 2 sensor_number = 0x03, M2-
>M3, cause 1
<I> 14:35:24.257    [218] Adjust power for 20, 01: from 0 to 5000
<I> 14:35:24.258    [218] Sending sync point
<I> 14:35:24.269    [237] SA 0x20 FRU 1 is OPERATIONAL
<I> 14:35:24.301    [218] Adjust power for 20, 02: from 0 to 20000
<I> 14:35:24.302    [218] Sending sync point
<I> 14:35:24.311    [213] Controller 20, FRU 2: ATCA state set to M4, prev=M3,
cause=0, locked=0
<I> 14:35:24.315    [214] Hot Swap event: SA=0x20 FRU 2 sensor_number = 0x03, M3-
>M4, cause 0
<I> 14:35:24.326    [238] SA 0x20 FRU 2 is OPERATIONAL
<I> 14:35:24.393    [217] IPMC Cooling Management: cooling state switched from
Unknown to Normal
```

The Shelf Manager software on the ShMM-1500 is optionally delivered without encryption-related code, for reasons having to do with export regulations. If present, the encryption-related code is located in the shared library `/lib/libpps_encryption.so`. If the encryption library is absent, the Shelf Manager does not allow opening RMCP+ sessions with encryption support. When ordering ShMM-1500s, you must specify which version is needed:

- ShMM-1500R-250M32F64R, ShMM-1500R-250M64F128R: Encryption enabled
- ShMM-1500R-250M32F64R-NE, ShMM-1500R-250M64F128R-NE: Encryption removed

The following line in the log above:

```
<I> 14:35:18.818    [199] [shm_crypto_init] initialization succeeded
```

indicates that the encryption library has been found and loaded successfully. Otherwise, the following messages appear:


```
<W> 14:34:04.480    [242] [shm_crypto_init] failed to load libpps_encryption.so
library, err libpps_encryption.so: cannot open shared object file: No such file
or directory
<I> 14:34:04.482    [242] [shm_crypto_init] the Shelf Manager will not support
encryption
```

4.4 Redundant operation

The active Shelf Manager exposes the ShMC device (address 20h) on IPMB, manages IPMB and the IPM controllers and interacts with the System Manager over RMCP and other shelf-external interfaces. It maintains an open TCP connection with the backup Shelf Manager. It communicates all changes in the state of the managed objects to the backup Shelf Manager.

The backup Shelf Manager does not expose the ShMC on IPMB, does not actively manage IPMB and IPM controllers, nor interact with the System Manager via the shelf-external interfaces (with one exception noted below). Instead, it maintains the state of the managed objects in its own memory (volatile and non-volatile) and updates the state as directed by the active Shelf Manager. These state updates are not applied immediately, but are cached on the backup Shelf Manager.

The backup Shelf Manager may become active as the result of a switchover. Two types of switchover are defined:

- Cooperative switchover: the active and backup Shelf Managers negotiate the transfer of responsibilities from the active to the backup Shelf Manager; this mode is supported via the CLI *switchover* command issued on the active or backup Shelf Manager.
- Forced switchover: the backup Shelf Manager determines that the active Shelf Manager is no longer alive or healthy, and forcefully takes on the responsibilities of the active Shelf Manager. Forced switchover can also be initiated via the CLI *switchover* command with the option *-force*, addressed to the backup Shelf Manager.

As discussed briefly in section 2.4.2, redundant ShMMs share three Hardware Redundancy Interface (HRI) data items between them (for six items, total); these items together indicate, for each ShMM, the health, presence and switchover request state of the ShMM itself (referenced as the local state) and its peer (referenced as remote state). Each ShMM variant has a different hardware implementation of these signals.

At the software level, each of the Shelf Managers can access a register that provides a view of these data items. The register is implemented in a PLD (a CPLD – Complex Programmable Logic Device – on the ShMM-500 or an FPGA – Field Programmable Logic Array – on the ShMM-1500 and ShMM-700) that is part of each ShMM. For details of the HRI implementation for each architecture, please see the appropriate ShMM Hardware Architecture specification. In the remainder of this section, these HRI data items are referenced informally as the Healthy, Presence and Switchover Request bits, each with Local and Remote variants. In addition, the set of HRI data items also includes an Active bit, which is 1 on the active Shelf Manager.

This section and most other Shelf Manager user documentation refers to the device that implements the HRI functionality as a CPLD, though that is not strictly true on either the ShMM-1500 or the ShMM-700.

The backup Shelf Manager recognizes the departure of the active Shelf Manager when the Remote Healthy or Remote Presence bit becomes 0. The Remote Presence bit monitors the presence of the

peer Shelf Manager; this bit changing to 0 on the ShMM-500 and ShMM-1500 means that the board hosting the peer Shelf Manager has been removed from the shelf. For the ShMM-700R this bit changing to 0 indicates that the peer Shelf Manager has been reset or removed from the shelf.

WARNING: Only the ShMM carrier board hosting the backup Shelf Manager should be removed from a shelf. If the ShMM carrier board hosting the active Shelf Manager needs to be removed the user should first issue the CLI *switchover* command on the active or backup Shelf Manager, and then remove the ShMM carrier board after the switchover has completed.

The Remote Healthy bit is set by the peer Shelf Manager during initialization; this bit changing to 0 means that the remote Shelf Manager has become not healthy (typically, has been powered off or reset).

Another situation that needs some action from the backup Shelf Manager is when the TCP connection between the Shelf Managers gets closed. This happens when either the communication link between the two Shelf Managers is broken, or the *shelfman* process on the active Shelf Manager terminates, in a voluntary or involuntary way, or due to a software exception.

Also, since the keepalive option is enabled on the TCP connection, it closes shortly after the active ShMM is switched off or reset. In the case of Shelf Manager termination, it is possible that the TCP connection is closed before the Remote Healthy bit becomes 0. So, in order to determine why the TCP connection closed, the backup Shelf Manager samples the state of the Remote Healthy bit immediately and, if it is still 1, samples it again after some delay. If the Remote Healthy bit becomes 0, the backup Shelf Manager concludes that the active Shelf Manager is dead. In that case it initiates a switchover and assumes the active Shelf Manager role.

Otherwise, if the Remote Healthy bit retains the value of 1, the backup Shelf Manager concludes that the communication link between the Shelf Managers is broken. In that case, no switchover is initiated; instead the backup Shelf Manager repeatedly reinitializes itself and tries to establish a connection with the active Shelf Manager, until the communication link is restored. Re-initialization is achieved by rebooting the ShMM and automatically restarting the Shelf Manager after the reboot. Special logic in the Shelf Manager guarantees that it does not try to become active at startup if the peer Shelf Manager is already active.

The Shelf Manager uses a watchdog timer to protect against becoming unresponsive due to infinite loops or other software bugs. In the event the watchdog timer on the active Shelf Manager triggers, that ShMM is reset, causing the Remote Healthy bit on the backup ShMM to become 0 and triggering a switchover.

In addition to the scenarios described above, the active Shelf Manager monitors the state of the Local Healthy and Active bits. If either of these bits becomes 0 on the active Shelf Manager, and the backup Shelf Manager is present, the active Shelf Manager initiates a reboot of the ShMM, in this way initiating a switchover to the backup Shelf Manager. This allows third-party applications on the ShMM to participate in monitoring health of the whole ShMM and initiate a switchover if necessary. Also, this behavior allows the former active ShMM to reinitialize itself as a backup in the case of a forced switchover from the backup Shelf Manager.

However, the behavior of a standalone Shelf Manager (one without a backup) in the case of the loss of either Local Healthy or Active bit is different.

In the case of the Local Healthy bit, the behavior is configurable and depends on the value of the configuration variable `EXIT_IF_HEALTHY_LOST_IN_STANDALONE_MODE`. By default, the value is `FALSE`, and in that case the Shelf Manager restores the Local Healthy bit and continues the operation. However if the value is `TRUE`, the standalone Shelf Manager reinitializes itself by rebooting the ShMM.

On ShMM-500, when the Active bit changes to 0 in the absence of a backup Shelf Manager, it is likely the result of a known hardware issue that exists on some ShMM carriers. In that case, the Shelf Manager attempts to restore the Active bit by requesting the active state again via setting the Switchover Request bit in the CPLD. If the recovery of the Active bit fails, the standalone Shelf Manager reinitializes itself by rebooting the ShMM. No check of the Active bit takes place if the configuration variable `CPLD_ACTIVE_WORKAROUND` is set to `FALSE` (the default value is `TRUE` for all carriers and changing it is not recommended). The workaround described in this paragraph does not apply to the ShMM-700 and the configuration variable `CPLD_ACTIVE_WORKAROUND` has no effect on the ShMM-700.

After a switchover, the former backup, now active, Shelf Manager performs additional initialization, applies cached state changes and collects any necessary further information from the IPM Controllers on IPMB. The new active Shelf Manager then exposes the ShMC device (address 20h) on IPMB, and assumes the IP address that was used for RMCP and other shelf-external interactions between the formerly active Shelf Manager and the System Manager. Since the RMCP session information is propagated from the active Shelf Manager to the backup Shelf Manager, RMCP sessions survive the switchover. For the System Manager using RMCP, the switchover is transparent.

The switchover is also transparent for the Web interface and for the SNMP interface (provided that they use the redundancy IP address that is switched over). Command-line interface sessions, since they are initiated locally on the target Shelf Manager, do not survive a switchover and need to be re-established again on the newly active Shelf Manager. The command-line interface support on the backup Shelf Manager is limited, but it does allow the backup Shelf Manager to request a switchover using the `switchover` command.

The formerly active Shelf Manager after the switchover can cease to exist or reinitialize itself as the backup Shelf Manager. To reinitialize as the backup Shelf Manager the formerly active ShMM must reboot the operating system.

4.4.1 Initialization of the redundant Shelf Managers

When starting, the Shelf Manager tries to determine its role (active or backup) and establish a TCP connection with its peer Shelf Manager. It does this using the following algorithm:

- The Shelf Manager starts listening for incoming TCP connection requests from its peer Shelf Manager.
- If either of the Remote Healthy and Remote Present bits is 0, the Shelf Manager assumes that there is no peer Shelf Manager and assumes the active role.
- Otherwise, the Shelf Manager issues a connection request, trying to establish a TCP connection to its peer. If the connection request fails, the Shelf Manager assumes the active role.

- If the connection request succeeds and there is no pending incoming connection request, the Shelf Manager assumes that an active Shelf Manager already exists in the shelf, and assumes the backup role.
- If both: 1) the outgoing connection request succeeds and 2) an incoming connection request is pending, the situation is symmetric between the two Shelf Managers and one of the Shelf Managers should assume the active role and the other should assume the backup role. The decision is based on the value of the Shelf Manager geographic address.

In PICMG 3.0 systems where full hardware addresses are available, hardware addresses assigned to the slots where peer Shelf Managers reside must have different least significant bits. In the case of a tie, the Shelf Manager with the even hardware address becomes the active Shelf Manager and the Shelf Manager with the odd hardware address becomes the backup. In PICMG 2.x systems, where the Shelf Manager geographic address is represented by the GEOGRAPHIC_POSITION bit, the Shelf Manager with GEOGRAPHIC_POSITION=0 becomes the active Shelf Manager, and the other Shelf Manager becomes the backup.

After the connection is established, the active Shelf Manager performs the initial synchronization with the backup Shelf Manager, sending it all the appropriate redundant data. During the initial synchronization, the backup Shelf Manager verifies that the active Shelf Manager uses the same carrier-specific module. If the carrier-specific modules are different, further data propagation and collaboration between the active and the backup Shelf Managers is not possible. Therefore, in that case, the backup Shelf Manager does not proceed with initialization, but reboots and reestablishes the connection in a loop, expecting the active Shelf Manager with the correct carrier-specific module to appear. The only exception to this is that in order to facilitate upgrades from non-HPDL to HPDL-based configurations, the HPDL-based backup Shelf Manager can coexist with a non-HPDL active Shelf Manager, but only when a reliable upgrade operation is in progress.

4.4.2 Redundancy and CPLD State sensor

Both active and standby Shelf Managers expose a sensor that indicates the high-level redundancy state of the ShMM, along with the state of the low-level redundancy bits exposed by the CPLD, and redundancy-related exceptional conditions in the CPLD, if any. This discrete sensor is associated with the physical Shelf Manager, has number 128, name “CPLD State” and OEM event/reading type code DEh. On the ShMM-700, where there is no CPLD and the programmatic interface to HRI is different, this sensor has the name “HWRI state”, though it is fully state-compatible with its counterpart on the ShMM-500/1500.

On the ShMM-500/1500, the sensor reads the second byte of the CPLD CSR0 register (bits 8-15) which contains various redundancy-related data, such as the Remote Presence, Remote Healthy, Remote Switchover Request, Local Presence and Active bits. For a detailed description of the CPLD register values, please see the ShMM Hardware Architecture specification. On the ShMM-700, the sensor uses the ShMM-700 HRI programmatic interface to obtain the values of the HRI bits.

The state mask reflects the high-level redundancy state (normal state) of the ShMM, as well as various low-level redundancy-related exceptions. Several bits of the exception group can be set in the mask, along with one of the normal state bits. The meanings of the bits in the sensor state mask are given below:

High-level redundancy state bits (normal state):

[0] - The current Shelf Manager is Active with no Backup.

[1] - The current Shelf Manager is Active with a Backup.

[2] - The current Shelf Manager is a Backup.

Low-level exception state bits (should be treated as low-level errors if set):

[3] - The Shelf Manager is Active with a Backup but Backup is unavailable for switchover.

[4] - The Shelf Manager is a Backup but the remote presence bit is not set.

[5] - The Shelf Manager is a Backup but the remote switchover request bit (the remote active bit on ShMM-700) is not set.

[6] - The Shelf Manager is a Backup but the HRI Active bit is set.

[7] - The Shelf Manager is Active with a Backup but the remote presence bit is not set.

[8] - The Shelf Manager is Active with a Backup but the remote healthy bit is not set.

[9] - The Shelf Manager is Active with a Backup but the HRI Active bit is not set.

[10] - The local presence bit is not set for the current Shelf Manager.

[11] - The Shelf Manager is Active with no Backup, but the remote healthy bit is set.

[12] - The Shelf Manager is Active with no Backup but the remote switchover request bit (the remote active bit on ShMM-700) is set.

[13] - The Shelf Manager is Active but senses the HRI Active bit set on other ShMM.

The current reading and state mask of the sensor can be retrieved in many ways, e.g. with the CLI command *sensordata*.

The sensor also generates events when the state mask changes. The format of the event message is given in the following table:

BYTE	DATA FIELD
1	<i>Event Message Rev</i> = 04h (IPMI 1.5 and IPMI 2.0)
2	<i>Sensor Type</i> = DEh (Redundancy and CPLD state)
3	<i>Sensor Number</i> = 80h
4	<i>Event Direction</i> [7] = 0b (Assertion) <i>Event Type</i> [6:0] = 6Fh (General Availability)

5	<p><i>Event Data 1:</i></p> <p>[7:4] = 7h (previous state and severity in Event Data 2, OEM code in Event Data 3)</p> <p>[3:0] = The highest new offset:</p> <p>00h - The current Shelf Manager is Active with no Backup.</p> <p>01h - The current Shelf Manager is Active with a Backup.</p> <p>02h - The current Shelf Manager is a Backup.</p> <p>03h - The Shelf Manager is Active with a Backup but Backup is unavailable for switchover.</p> <p>04h - The Shelf Manager is a Backup but the remote presence bit is not set.</p> <p>05h - The Shelf Manager is a Backup but the remote switchover request bit is not set.</p> <p>06h - The Shelf Manager is a Backup but the CPLD Active bit is set.</p> <p>07h - The Shelf Manager is Active with a Backup but the remote presence bit is not set.</p> <p>08h - The Shelf Manager is Active with a Backup but the remote healthy bit is not set.</p> <p>09h - The Shelf Manager is Active with a Backup but the CPLD Active bit is not set.</p> <p>0Ah - The local presence bit is not set for the current Shelf Manager.</p> <p>0Bh - The Shelf Manager is Active with no Backup, but the remote healthy bit is set.</p> <p>0Ch - The Shelf Manager is Active with no Backup but the remote switchover request bit is set.</p> <p>0Dh - The Shelf Manager is Active but senses the CPLD Active bit set on other ShMM.</p>
6	<p><i>Event Data 2:</i></p> <p>[7:4] = Event Severity</p> <p>[3:0] = The highest previous offset</p>
7	<p><i>Event Data 3 (bits 8-15 of the ShMM CPLD state on the ShMM-500/1500):</i></p> <p>[7] 1b = Reboot was caused by watchdog expiration</p> <p>[6] 1b = Interrupt happened</p> <p>[5] 1b = Active</p> <p>[4] 1b = Local presence</p> <p>[3] 1b = Hot Swap latch open</p> <p>[2] 1b = Remote switchover request</p>

	<p>[1] 1b = Remote healthy</p> <p>[0] 1b = Remote presence</p> <p>On the ShMM-700, this byte should be ignored; its value is always 00h</p>
--	---

4.4.3 Reboot Reason sensor

Both active and standby Shelf Managers also expose a sensor that indicates the reason for the last reboot. This discrete sensor is associated with the physical Shelf Manager, has number 129, name “Reboot Reason” and OEM event/reading type code DDh.

The state mask for the sensor indicates the cause of the last reboot; its bits have the following meanings:

- [0] - The reboot reason is either a Shelf Manager crash or unknown.
- [1] - The reboot was caused by a switchover operation.
- [2] - The reboot was caused by a forced switchover operation.
- [3] - The reboot was caused by the CLI command *terminate*.
- [4] - The reboot was caused by loss of the HEALTHY bit.
- [5] - The reboot was caused by loss of the ACTIVE bit.
- [6] - The reboot of the Backup ShMM happened because the redundancy connection was broken but the Active ShMM was still alive.
- [7] - The reboot happened due to an error during the Shelf Manager startup.
- [8] - The reboot was caused by the ShMM hardware watchdog.
- [9] – The reboot was initiated by software (a *reboot()* system call).
- [10] – The ShMM has been power cycled.
- [11] – The reboot happened because of an SDRAM error.
- [12] – The reboot was caused by the IPMI “Cold Reset” command
- [13] – The reboot was caused by I²C Isolation switchover
- [14] – The reboot was caused by a local IPMB-0 failure (ShMM-700 only)

The sensor reading byte is always 0 and does not have any meaning.

Note: in case of a Shelf Manager reboot initiated by software, this sensor may report two different reasons based on system shutdown speed. If the system shuts down within the default setting of 2 seconds (i.e. while the watchdog timer has not expired), the reboot reason will be a reboot() call (state mask = 9). Otherwise, the reboot reason will be the hardware watchdog. To avoid this non-deterministic result, it is possible to increase the WATCHDOG_TIMEOUT configuration variable value, for example WATCHDOG_TIMEOUT could be increased to 5 seconds.

In addition to the state mask reported by the reboot reason sensor, information about recent reboots is stored in the file `/var/nvdata/reboot_reason.log` on the ShMM. The typical log file is shown below:

```
04/10/2014 04:55:35 The reboot was caused by power operation
```



```
04/10/2014 04:41:04 The reboot was caused by power operation
04/09/2014 23:32:47 The reboot was caused by power operation
04/08/2014 05:14:29 The reboot was caused by WATCHDOG
04/07/2014 00:10:43 The reboot was caused by Shelf Manager abnormal termination:
signal = 11(Segmentation fault), pid = 139
04/06/2014 02:11:59 The reboot was caused by power operation
```

Each line in the file contains a timestamp and text description of a ShMM reboot. The text description normally corresponds to the description of the reboot reason sensor state shown above. However, if the corresponding reboot reason is 0 (“The reboot reason is either a Shelf Manager crash or unknown”) and it is known that a crash has occurred, the log message notes that and adds information about the signal that caused the crash, along with the Shelf Manager process identifier.

4.5 Operation in shelves with radial IPMB-0

Some shelves implement radial control of IPMB-0 links in the shelf. In that case, the segment of IPMB-0 leading to each IPM controller in the shelf can be turned on and off individually by the Shelf Manager. This applies individually to both the IPMB-A and IPMB-B portions of IPMB-0. The operation of the Shelf Manager in such shelves is different from the shelves with a simple bused IPMB-0.

The Shelf FRU Information in radial shelves must contain special records that identify the shelf as implementing a radial IPMB-0 topology and specify the routing of the IPMB links to the IPM controllers. These records must conform to the ATCA specification, section 3.8.1.1.

Radial IPMB-0 is implemented differently for ShMM-500s and for ShMM-1500s. For ShMM-500s, there is a single logical bus for each of the two constituent buses of IPMB-0 (that is, IPMB-A and IPMB-B) and radial IPMB is actually implemented in a pseudo-radial, in which the radial segments are only implemented at the physical level. For ShMM-1500s, there are distinct logical buses for each radial segment. Support for radial IPMB is not yet implemented in the Linux kernel for the ShMM-700, though on the hardware level, the same approach as on the ShMM-500 should work fine.

4.5.1 Operation in ShMM-500-based shelves with radial IPMB-0

For pseudo-radial IPMB-0 implementations on the ShMM-500, the Shelf Manager uses a single logical IPMB interface on each of IPMB-A and IPMB-B to communicate with all the IPM controllers, but there is a separate GPIO control on the carrier for each of the radial segments that can be used to turn that segment on and off. Only some carriers provide this support.

If both: 1) the Shelf FRU Information contains the radial IPMB-0 description record and 2) the carrier supports radial operation, the Shelf Manager on ShMM-500s implements isolation of faulty bus segments in the case of a persistent error on IPMB-0. This prevents such an error on one or several specific IPM controllers from causing a failure of an entire bus. Isolation is performed independently for IPMB-A and IPMB-B. Isolated segments are turned off (which means that the corresponding IPM controller(s) is(are) isolated from the corresponding bus).

Such segments are kept in the isolated state for `IPMB_LINK_ISOLATION_TIMEOUT` seconds, after which they are automatically turned on; if the fault still exists, the persistent IPMB-0 error

occurs again and the faulty links is isolated again. (By default, however, the value of this configuration parameter is `-1`, which means “never re-enable the link automatically”).

In addition, isolated links can be turned on in one of the following ways:

- manually, via the CLI command `setipmbstate` with parameters indicating the ShMC target address (0x20) and the corresponding link number and corresponding IPM controller is removed from the shelf (that is, when it goes to the state M0).
- automatically, when the corresponding IPM controller is removed from the shelf (that is, when it goes to the state M0).

To find out the current isolation state of radial IPMB-0 links, the CLI command `getipmbstate` can be used. When applied to the ShMC target address (0x20), this command shows the state of all radial IPMB-0 links. For example, the command:

```
# clia getipmbstate 20
```

can yield a report like this:

```
Pigeon Point Shelf Manager Command Line Interpreter
```

```
20: Link: 1, LUN: 0, Sensor # 10 ("IPMB LINK 1")
    Bus Status: 0x8 (IPMB-A Enabled, IPMB-B Enabled)
    IPMB A State: 0x08 (LocalControl, No failure)
    IPMB B State: 0x08 (LocalControl, No failure)
20: Link: 2, LUN: 0, Sensor # 11 ("IPMB LINK 2")
    Bus Status: 0x8 (IPMB-A Enabled, IPMB-B Enabled)
    IPMB A State: 0x08 (LocalControl, No failure)
    IPMB B State: 0x08 (LocalControl, No failure)
20: Link: 3, LUN: 0, Sensor # 15 ("IPMB LINK 3")
    Bus Status: 0x8 (IPMB-A Enabled, IPMB-B Enabled)
    IPMB A State: 0x08 (LocalControl, No failure)
    IPMB B State: 0x08 (LocalControl, No failure)
20: Link: 4, LUN: 0, Sensor # 16 ("IPMB LINK 4")
    Bus Status: 0x4 (IPMB-A Disabled, IPMB-B Enabled)
    IPMB A State: 0x07 (Isolated, Undiagnosed communication failure)
    IPMB B State: 0x08 (LocalControl, No failure)
20: Link: 5, LUN: 0, Sensor # 17 ("IPMB LINK 5")
    Bus Status: 0x8 (IPMB-A Enabled, IPMB-B Enabled)
    IPMB A State: 0x08 (LocalControl, No failure)
    IPMB B State: 0x08 (LocalControl, No failure)
20: Link: 6, LUN: 0, Sensor # 18 ("IPMB LINK 6")
    Bus Status: 0x8 (IPMB-A Enabled, IPMB-B Enabled)
    IPMB A State: 0x08 (LocalControl, No failure)
    IPMB B State: 0x08 (LocalControl, No failure)
20: Link: 7, LUN: 0, Sensor # 19 ("IPMB LINK 7")
    Bus Status: 0x8 (IPMB-A Enabled, IPMB-B Enabled)
    IPMB A State: 0x08 (LocalControl, No failure)
    IPMB B State: 0x08 (LocalControl, No failure)
```

```

20: Link: 8, LUN: 0, Sensor # 20 ("IPMB LINK 8")
    Bus Status: 0x8 (IPMB-A Enabled, IPMB-B Enabled)
    IPMB A State: 0x08 (LocalControl, No failure)
    IPMB B State: 0x08 (LocalControl, No failure)
20: Link: 9, LUN: 0, Sensor # 21 ("IPMB LINK 9")
    Bus Status: 0x8 (IPMB-A Enabled, IPMB-B Enabled)
    IPMB A State: 0x08 (LocalControl, No failure)
    IPMB B State: 0x08 (LocalControl, No failure)
20: Link: 10, LUN: 0, Sensor # 22 ("IPMB LINK 10")
    Bus Status: 0x8 (IPMB-A Enabled, IPMB-B Enabled)
    IPMB A State: 0x08 (LocalControl, No failure)
    IPMB B State: 0x08 (LocalControl, No failure)
20: Link: 11, LUN: 0, Sensor # 23 ("IPMB LINK 11")
    Bus Status: 0x8 (IPMB-A Enabled, IPMB-B Enabled)
    IPMB A State: 0x08 (LocalControl, No failure)
    IPMB B State: 0x08 (LocalControl, No failure)
20: Link: 12, LUN: 0, Sensor # 24 ("IPMB LINK 12")
    Bus Status: 0x8 (IPMB-A Enabled, IPMB-B Enabled)
    IPMB A State: 0x08 (LocalControl, No failure)
    IPMB B State: 0x08 (LocalControl, No failure)
20: Link: 13, LUN: 0, Sensor # 25 ("IPMB LINK 13")
    Bus Status: 0x8 (IPMB-A Enabled, IPMB-B Enabled)
    IPMB A State: 0x08 (LocalControl, No failure)
    IPMB B State: 0x08 (LocalControl, No failure)
20: Link: 14, LUN: 0, Sensor # 26 ("IPMB LINK 14")
    Bus Status: 0x8 (IPMB-A Enabled, IPMB-B Enabled)
    IPMB A State: 0x08 (LocalControl, No failure)
    IPMB B State: 0x08 (LocalControl, No failure)
20: Link: 15, LUN: 0, Sensor # 27 ("IPMB LINK 15")
    Bus Status: 0x8 (IPMB-A Enabled, IPMB-B Enabled)
    IPMB A State: 0x08 (LocalControl, No failure)
    IPMB B State: 0x08 (LocalControl, No failure)

```

The output above indicates that IPMB-0 link 4 for IPMB-A has been isolated as a result of a persistent IPMB failure. To manually re-enable the link, the following CLI command can be used:

```
# clia setipmbstate 20 A 4 1
```

4.5.2 Operation in ShMM-1500-based shelves with radial IPMB-0

On ShMM-1500-based shelves with radial IPMB-0 support, there is a separate logical IPMB interface for each of the two portions (IPMB-A and IPMB-B) of each radial IPMB segment. There is no need for faulty link isolation in this case, since each segment is separately accessible and permanently isolated from its peer segments on a logical basis.

Radial operation of ShMM-1500 requires all of the following: 1) radial support on the shelf level, 2) the presence of a radial IPMB description record in the Shelf FRU Information and 3) radial support


on the ShMM carrier. For non-HPDL systems, a special Pigeon Point defined record (the IPMB topology record) must be present in the ShMM carrier FRU information to specify the type of the IPMB topology supported by the carrier. In HPDL-based systems, the IPMB topology type can be specified in the *IPMB_TOPOLOGY* clause in the HPDL Carrier definition; in that case, the IPMB topology record is not needed.

Two different topologies of radial IPMB-0 are supported by the ShMM-1500: dual star and redundant dual star. The implemented topology is determined by the design of the ShMM carrier.

With the dual star topology, only one IPMB hub exists on each ShMM carrier. This hub connects the ShMM to radial links on either IPMB-A or IPMB-B, depending on the physical position of the ShMM carrier in the shelf. This topology has the following properties:

- Provides equivalent connectivity to the dual star topology that ATCA defines for the Gigabit Ethernet Base Interface
- Includes two cross-links that connect an IPMB (either A or B) port on one ShMM to the radial hub on the other ShMM carrier
- Requires two installed Shelf Managers for dual redundant operation of IPMB-0
- Establishes two distinct and unconnected fault zones for the two Shelf Managers
- Defined by the ATCA architecture only after the amendments made by ECN-002 for PICMG 3.0 R2.0

Figure 4: Dual star radial IPMB-0 topology

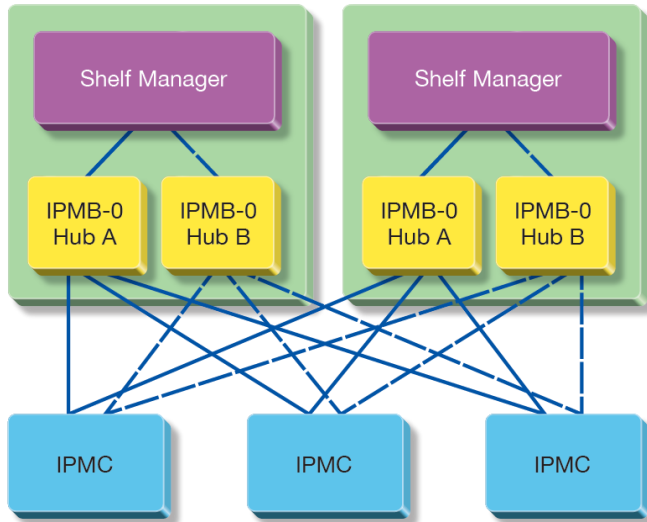


With the redundant dual star topology, each ShMM carrier has two IPMB-0 hubs, with each hub connecting one of the ShMM IPMB ports (A or B) to corresponding radial links on the backplane. This topology has the following properties:

- Defined in PICMG 3.0 R1.0 and supported as the only IPMB-0 radial topology until the ECN-002 amendments to PICMG 3.0 R2.0
- Supports dual redundant operation of IPMB-0 even with the only one Shelf Manager installed

- Does not need or include IPMB-0 cross-links between the two ShMM carriers

Figure 5: Redundant dual star radial IPMB-0 topology



The IPMB Topology record has the following format (which follows the regular FRU Info multirecord format):

Table 18: IPMB Topology record

OFFSET	LENGTH	DEFINITION
0	1	<i>Record Type ID</i> . The value of C0h (OEM) is used,
1	1	<i>End of List/Version</i> [7] – End of List. Set to one for the last record. [6:4] – Reserved, write as 0h. [3:0] – Record format version (=2h for this definition).
2	1	<i>Record length</i>
3	1	<i>Record Checksum</i> . Holds the zero checksum of the record.
4	1	<i>Header Checksum</i> . Holds the zero checksum of the header.
5	3	<i>Manufacturer ID</i> . LS byte first. Write as this three byte ID assigned to nVent (16394 decimal, or 0Ah 40h 00h).
8	1	<i>nVent Record ID</i> . For this record, the value 5h is used.
9	1	<i>Record Format Version</i> . For this record, the value 0h is used.
8	1	<i>IPMB Topology Type</i> . The following values are defined: 0h – unspecified 1h – Bused IPMB-0 carrier 2h – Dual star radial IPMB-0 carrier 3h – Redundant dual star radial IPMB-0 carrier 4h – Universal redundant dual star radial IPMB-0 carrier (not currently supported) Other values are reserved.

In a FRU Info Compiler input file, the following approach is used to define the IPMB Topology record (for a redundant dual star radial configuration, in this example):

```
[nVent.IPMB Topology]
Carrier Type = 3
```

In an HPDL carrier definition file, the IPMB topology can be specified in a *IPMB_TOPOLOGY* clause that has the following syntax:

```
<ipmb_topology_clause> ::= "IPMB_TOPOLOGY" <bus_type> ";"
<bus_type> ::= "BUSED" | "DUAL_STAR" | "REDUNDANT_DUAL_STAR"
```

For example, a redundant dual star radial configuration, the following definition can be used:

```
IPMB_TOPOLOGY REDUNDANT_DUAL_STAR;
```

The radial functionality is turned on if the following two prerequisites are satisfied:

- the carrier is designated as a dual star or redundant dual star type, either in the HPDL carrier definition or in an IPMB Topology record in the carrier FRU Information with IPMB Topology Type = 2h or 3h;
- the Shelf FRU Information contains an IPMB radial mapping record.

If the above prerequisites are satisfied, the logical Shelf Manager creates and hosts multiple IPMB link sensors, each of which reports the state of the corresponding radial segment. As in the case of ShMM-500, the CLI command *getipmbstate* can be used to show the state of all segments or of a specific IPMB segment; the CLI command *setipmbstate* can be used to selectively enable or disable a specific IPMB segment.

4.5.3 Operation in ShMM-700-based shelves with radial IPMB-0

On ShMM-700-based shelves, radial IPMB-0 support is implemented very similar to the ShMM-1500 case. As on the ShMM-1500, there is a separate logical IPMB interface for each of the two portions (IPMB-A and IPMB-B) of each radial IPMB segment, and the supported topologies include dual star and redundant dual star.

Since all ShMM-700-based shelves are HPDL-based, IPMB topology in those systems is always specified in HPDL and no IPMB Topology record in the Carrier FRU Information is needed or supported. The IPMB topology clause has the following syntax (which is backwards-compatible with ShMM-1500, but with an extension specifying the position of the IPMB cross link).

```
<ipmb_topology_clause> ::= "IPMB_TOPOLOGY" <bus_type>
                        [", " <cross_link_num>] ";"
<bus_type> ::= "BUSED" | "DUAL_STAR" | "REDUNDANT_DUAL_STAR"
<cross_link_num> ::= 0 | 1 [", " "NOBUS"]
```

The value `BUSED` of the `<bus_type>` token indicates a bused IPMB-0 carrier. The default value is `BUSED`. For hybrid IPMB-0 designs, i.e. carrier designs that support both a radial IPMB-0 and a bused IPMB-0, the radial IPMB-0 topology is specified. For carrier designs that only support a bused IPMB topology the IPMB Topology clause is optional and does not need to be included.

The value `DUAL_STAR` corresponds to the “non-redundant dual star” radial topology (one Shelf Manager implements all the IPMB-A radial links, the other Shelf Manager implements all the IPMB-B radial links).

The value `REDUNDANT_DUAL_STAR` corresponds to the “redundant dual-star” radial topology where each Shelf Manager implements both IPMB-A and IPMB-B radial links.

The `<cross_link_num>` specifies the link number used for communication with the peer Shelf Manager in a radial configuration and is used only with the Redundant Dual Star topology. The default value is 0, which means that no cross link exists and the Shelf Managers communicate with each other over the bused IPMB-0. Currently, only link number 1 can be specified as the cross link; as a result, the only legal values for `<cross_link_num>` are 0 and 1.

If a cross link exists (i.e. `<cross_link_num>` is 1), communication between redundant Shelf Managers goes over the cross link, even if a bused IPMB-0 also exists, as would be the case for a hybrid IPMB-0 topology. The keyword `NOBUS` can be specified if `<cross_link_num>` is 1. This indicates a bused IPMB-0 is not physically implemented in the shelf. In this case, the active Shelf Manager does not attempt to access the bused IPMB-0 but emulates the bused IPMB logically over the cross link. The bused IPMB link sensor (“IPMB LINK”) in this case monitors the state of the cross link.

The remaining aspects of radial IPMB-0 support in ShMM-700-based shelves are the same as in ShMM-1500-based-shelves, except for programming/upgrading the routing elements (i.e. the SmartFusion Customizable System-on-Chip (cSoC) devices), which is described in the next section.

4.5.3.1 Upgrading the routing element SmartFusion cSoC(s)

The routing element SmartFusion cSoC(s) can be programmed/upgraded using the *ripmb_prog* utility, which is included as part of the standard ShMM-700R RFS image. In addition to programming the SmartFusion cSoC(s), this utility can be used to obtain the version of the currently-installed image.

IMPORTANT: before running *ripmb_prog* to program/upgrade the routing element SmartFusion cSoC(s), terminate the Shelf Manager by issuing a *clia terminate* command. Upon completion of the SmartFusion cSoC programming the Shelf Manager must be re-initialized by issuing a *reboot* command. Failing to do so may lead to SmartFusion cSoC access errors and unpredictable behavior. If *ripmb_prog* is invoked only to obtain the version of the currently-installed image, terminating the Shelf Manager is not necessary.

To program a new SmartFusion cSoC image to the routing element SmartFusion cSoC devices, issue the following command:

```
ripmb_prog <image_name>
```

where *<image_name>* is the name of the upgrade image (with a *.dat* extension). The SmartFusion cSoC upgrade image file includes a SmartFusion A2F060 image concatenated with a SmartFusion A2F200 image. Each individual image includes all the programmable content needed for the A2F060 device or the A2F200 device (including FPGA fabric and Cortex-M3 firmware). The *ripmb_prog* utility will determine the SmartFusion device type (A2F060 or A2F200) as well as the number of SmartFusion devices that are connected to the JTAG chain and program the appropriate image from the concatenated SmartFusion cSoC image file to each device.

IMPORTANT: all routing element SmartFusion cSoCs on the ShMM-700R carrier must be programmed with the same version of the SmartFusion cSoC image file. If the upgrade procedure is interrupted at any point, it must be re-started to ensure all SmartFusion cSoCs are completely programmed. When the programming is successful, the following message(s) are displayed by the *ripmb_prog* utility (depending on the number of SmartFusion devices to be programmed):

```
Image file: radial.dat
```

```
Determining the number of devices in the JTAG scan chain: 2
```

```
Programming device #1:.....SUCCESS
```

```
Programming device #2: .....SUCCESS
```

To obtain the version of the currently-installed routing element image, issue the following command:

```
ripmb_prog -V
```

This command displays the version of the routing element SmartFusion cSoC firmware and FPGA design in the following format: *<firmware_version>/<fpga_version>*, for example: 1.3/1.0.1.1. If this command displays "**VERSION MISMATCH**", the SmartFusion cSoCs must be re-programmed with the desired image.

4.6 Automatic SEL truncation

The System Event Log (SEL) on the Shelf Manager stores events from all IPM controllers and FRUs in the shelf and can easily exceed its maximum capacity. To prevent overflows, the Shelf Manager can automatically truncate the SEL, removing the oldest entries as SEL approaches its limits.

The automatic truncation algorithm works as follows. The two configuration parameters are defined: *SEL_HIGH_WATERMARK* and *SEL_LOW_WATERMARK*. The first parameter relates to the percentage of free entries in the SEL, and the second one relates to the percentage of occupied entries in the SEL. When the percentage of free SEL entries falls below *SEL_HIGH_WATERMARK*, the truncation thread wakes up and starts deleting the oldest entries from the SEL using the IPMI command “Delete SEL Entry”. The thread works until the percentage of occupied entries in the SEL falls below *SEL_LOW_WATERMARK*; then it stops and stays dormant until the percentage of free SEL entries falls below the *SEL_HIGH_WATERMARK* again.

This algorithm is enabled by default, but can be turned off by setting the configuration variable *SYSTEM_MANAGER_TRUNCATES_SEL* to *TRUE*. In that case, automatic truncation is turned off, but the Shelf Manager assists the System Manager by indicating the current state of the SEL via the sensor SEL State of the type Event Logging Enabled, as follows (in accordance with the IPMI 2.0 specification):

- The reading returned by that sensor is equal to the percentage of the SEL that is occupied (0 to 100);
- The sensor assumes the state SEL Almost Full (5h) when the percentage of free entries in the SEL falls below the high watermark (the value of the configuration parameter *SEL_HIGH_WATERMARK*). An event is generated in that case and placed in the SEL; Event Data Byte 3 of the event contains the percentage of the SEL that is occupied. No event is generated for this state until the occupancy of the SEL falls below the low watermark; this is done to prevent multiple events if SEL occupancy oscillates around the high watermark.
- A custom state (6h) is defined for the state when the SEL occupancy is below the low watermark (the percentage of occupied entries falls below the value of the configuration parameter *SEL_LOW_WATERMARK*). There is no event associated with this state; the System Manager can detect this state by polling the sensor value and its asserted states.

Using this information, the System Manager performs truncation of the SEL, presumably using the IPMI command “Delete SEL Entry”, sending it over the RMCP session.

4.7 Cooling state sensors

The cooling management strategy in the Shelf Manager is configurable and also depends on the currently chosen carrier. Cooling management strategy configuration is described in section 3.7. Irrespective of the current cooling management strategy, however, several sensors are defined that expose the current cooling state of the shelf to the System Manager. These sensors belong to the logical Shelf Manager (IPMB address 20h) and are defined in the following table:

Table 19: Cooling state sensors

SENSOR NUMBER	SENSOR NAME	EVENT/READING TYPE	SENSOR TYPE	OFFSETS
135	FT Oper. Status	Redundancy (Bh)	Management Subsystem Health (28h)	<p>00h = Full Redundancy = all fan trays defined in the Address Table are operational</p> <p>01h = Redundancy Lost = some of the fan trays defined in the Address Table are missing or non-operational. With the HPDL default cooling management strategy this causes the fan level for all remaining fan trays to be set to their maximum</p>
136	Cooling State	Severity (7)	Management Subsystem Health (28h)	<p>00h = transition to OK. The cooling state is Normal</p> <p>01h = transition to Non-Critical from OK. The cooling state is now Minor Alert, the previous cooling state was Normal.</p> <p>02h = transition to Critical from less severe. The cooling state is now Major Alert, the previous cooling state was either Normal or Minor Alert.</p> <p>04h = transition to Non-Critical from more severe. The cooling state is now Minor Alert, the previous cooling state was either Major or Critical Alert.</p> <p>05h = transition to Critical from Non-recoverable. The current cooling state is Major Alert, the previous cooling state was Critical Alert.</p> <p>06h = transition to Non-recoverable. The current cooling state is now Critical Alert.</p>

137	Fans State	Severity (7)	Management Subsystem Health (28h)	<p>00h = transition to OK. The fans state is Normal (no thresholds are crossed on fan tachometer sensors)</p> <p>01h = transition to Non-Critical from OK. The fans state is now Minor Alert (non-critical thresholds are crossed for some tachometer sensors), the previous fans state was Normal.</p> <p>02h = transition to Critical from less severe. The fans state is now Major Alert (critical thresholds are crossed for some tachometer sensors), the previous fans state was either Normal or Minor Alert.</p> <p>04h = transition to Non-Critical from more severe. The fans state is now Minor Alert, the previous fans state was either Major Alert or Critical Alert (non-recoverable thresholds are crossed for some tachometer sensors).</p> <p>05h = transition to Critical from Non-recoverable. The current fans state is Major Alert, the previous fans state was Critical Alert.</p> <p>06h = transition to Non-recoverable. The current fans state is now Critical Alert.</p>
-----	------------	--------------	-----------------------------------	---

4.8 Deadlock detection

The Shelf Manager contains built-in facilities for detecting internal deadlocks. A deadlock happens if two or more threads in the Shelf Manager application acquire locks in an incorrect order, so that one thread owns the lock requested by the other thread, and vice versa. As a result, the Shelf Manager becomes unresponsive on its external interfaces, though it still continues to run. Any deadlock is a result of a serious programming bug and should not happen during normal operation. The deadlock detection facilities, however, add an additional layer of protection and ensure that

should a deadlock happen during normal operation of the Shelf Manager, a switchover to the backup Shelf Manager takes place and the Shelf Manager continues operation without interruption.

There are two deadlock detection facilities: one that is triggered by watchdog timer expiration and one that directly detects incorrect locking patterns in the Shelf Manager.

The watchdog-based deadlock detection facility monitors the operation of the most important interface-oriented subsystems of the Shelf Manager, which now include the CLI processing thread and the RMCP processing threads. Normally, if a deadlock takes place anywhere in the Shelf Manager, one or more of these threads is affected immediately or very soon as a result of deadlock propagation.

The threads monitored by the watchdog-based deadlock detection facility must periodically strobe the internal watchdog timer. The strobe period is set separately for each type of monitored thread and is 30 seconds for the CLI thread and 3 seconds for the RMCP threads. If one of the monitored threads fails to strobe the internal watchdog timer on timely basis, the Shelf Manager does the following:

- issues a warning message to the console and/or to the syslog file.
- collects information about current usage of locks in the Shelf Manager.
- writes this information to a Flash file on the ShMM (*/var/nvdata/dumplog.bin*), where it survives reboots and can be subsequently analyzed by the analysis tool *dumplog*.
- if one of the monitored threads fails to do a timely strobe of the internal watchdog timer three times in a row, stops strobing the hardware watchdog timer, which causes a reset of the ShMM and a switchover to the backup Shelf Manager.

In this scenario, output similar to the following appears on the console:

```
<W> 17:25:49.368    [218] [ipmf_watchdog_thread_proc] 239 facility has failed,
don't confirm the watchdog
<I> 17:25:49.760    [218] Write 132232 bytes to the '/var/nvdata/dumplog.bin', 0
seconds, lock offset 131848
<W> 17:25:50.068    [218] [ipmf_watchdog_thread_proc] 239 facility has failed,
don't confirm the watchdog
<W> 17:25:50.378    [218] [ipmf_watchdog_thread_proc] 239 facility has failed,
don't confirm the watchdog
<W> 17:25:50.688    [218] [ipmf_watchdog_thread_proc] 239 facility has failed,
don't confirm the watchdog
<W> 17:25:50.998    [218] [ipmf_watchdog_thread_proc] 239 facility has failed,
don't confirm the watchdog
```

The direct deadlock detection mechanism is supported in the Shelf Manager, starting with release 2.5.2. This mechanism checks the state of the currently claimed locks every 5 seconds by looking at the internal lock structures. It analyzes that state information for circular dependencies that indicate a deadlock. If a deadlock is detected by this mechanism, the Shelf Manager does the following:

- issues a warning message to the console and/or to the syslog file.
- collects information about current usage of locks in the Shelf Manager.

- writes this information to a Flash file on the ShMM (`/var/nvdata/dumplog.bin`), where it survives reboots and can be subsequently analyzed by the analysis tool `dumplog`.
- stops strobing the hardware watchdog timer, which causes a reset of the ShMM and a switchover to the backup Shelf Manager.

Note: This second (direct) deadlock detection mechanism is not implemented on the ShMM-700. The reason is that with the new version of the Linux kernel (2.6.x) used on the ShMM-700, the detection of deadlocks by direct analysis of internal lock (mutex) structures is not reliable and generates substantially larger performance overhead. The first mechanism, watchdog-based monitoring of the CLI and RMCP processing threads, should be sufficient for detection of deadlocks on this platform.

In this scenario, an output similar to the following appears on the console:

```
lock: 0x2acd3940, owner 208, waiting 209
lock: 0x2ac432f0, owner 223, waiting 217 214 220 208
lock: 0x2ac43650, owner 220, waiting 223
<W> 14:09:43.866 [199] [ipmf_watchdog_thread_proc] a potential deadlock was
detected
<I> 14:09:44.307 [199] Write 148724 bytes to the '/var/nvdata/dumplog.bin', 1
seconds, lock offset 148232
lock: 0x2acd3940, owner 208, waiting 209
lock: 0x2ac432f0, owner 223, waiting 217 214 220 208
lock: 0x2ac43650, owner 220, waiting 223
<W> 14:09:44.616 [199] [ipmf_watchdog_thread_proc] a potential deadlock was
detected
lock: 0x2acd3940, owner 208, waiting 209
lock: 0x2ac432f0, owner 223, waiting 217 214 220 208
lock: 0x2ac43650, owner 220, waiting 223
<W> 14:09:44.926 [199] [ipmf_watchdog_thread_proc] a potential deadlock was
detected
lock: 0x2acd3940, owner 208, waiting 209
lock: 0x2ac432f0, owner 223, waiting 217 214 220 208
lock: 0x2ac43650, owner 220, waiting 223
<W> 14:09:45.236 [199] [ipmf_watchdog_thread_proc] a potential deadlock was
detected
lock: 0x2acd3940, owner 208, waiting 209
lock: 0x2ac432f0, owner 223, waiting 217 214 220 208
lock: 0x2ac43650, owner 220, waiting 223
<W> 14:09:45.546 [199] [ipmf_watchdog_thread_proc] a potential deadlock was
detected
```

For postmortem analysis of a deadlock, the utility `dumplog` (supplied in the Shelf Manager RFS) can be used. The following command line invokes this utility:

```
# dumplog -f /var/nvdata/dumplog.bin
```

The result is produced on standard output in a human-readable format. It includes information about which locks are owned by which threads, and a per-thread history of lock acquisition and release, for each currently active Shelf Manager thread. The output of the command looks something like this:

```
*** 24, Log for pthread with pid = 263, 18:17 ***
Jun 22 17:23:27 2007: Succeeded 2ac34530; pid=263, file="sdrrep.c", line=4288
Jun 22 17:23:27 2007: Released 2ac34530; pid=263, file="sdrrep.c", line=3715
Jun 22 17:23:27 2007: Acquiring 2ac34ad0; pid=263, file="ipmc.c", line=387
Jun 22 17:23:27 2007: Succeeded 2ac34ad0; pid=263, file="ipmc.c", line=387
Jun 22 17:23:27 2007: Released 2ac34ad0; pid=263, file="ipmc.c", line=484
Jun 22 17:23:27 2007: Acquiring 7d5ff740; pid=263, file="msg.c", line=775
Jun 22 17:23:27 2007: Succeeded 7d5ff740; pid=263, file="msg.c", line=775
Jun 22 17:23:27 2007: Acquiring 2acc1f80; pid=263, file="ipmf.c", line=1360
Jun 22 17:23:27 2007: Succeeded 2acc1f80; pid=263, file="ipmf.c", line=1360
Jun 22 17:23:27 2007: Acquiring 2ac344a0; pid=263, file="msg.c", line=924
Jun 22 17:23:27 2007: Succeeded 2ac344a0; pid=263, file="msg.c", line=924
Jun 22 17:23:27 2007: Released 2ac344a0; pid=263, file="msg.c", line=932
Jun 22 17:23:27 2007: Acquiring 2ac344a0; pid=263, file="msg.c", line=867
Jun 22 17:23:27 2007: Succeeded 2ac344a0; pid=263, file="msg.c", line=867
Jun 22 17:23:27 2007: Released 2ac344a0; pid=263, file="msg.c", line=876
Jun 22 17:23:27 2007: Acquiring 2acbe5ec; pid=263, file="ipmf_ipmb.c", line=1649
Jun 22 17:23:27 2007: Succeeded 2acbe5ec; pid=263, file="ipmf_ipmb.c", line=1649
Jun 22 17:23:27 2007: Released 2acbe5ec; pid=263, file="ipmf_ipmb.c", line=1656
Jun 22 17:23:27 2007: Acquiring 2acc1de0; pid=263, file="ipmf_ipmb_buffe",
line=98
Jun 22 17:23:27 2007: Succeeded 2acc1de0; pid=263, file="ipmf_ipmb_buffe",
line=98
Jun 22 17:23:27 2007: Released 2acc1de0; pid=263, file="ipmf_ipmb_buffe",
line=129
Jun 22 17:23:27 2007: Released 2acc1f80; pid=263, file="ipmf.c", line=1533
Jun 22 17:23:27 2007: CondWait Before 7d5ff740; pid=263, file="msg.c", line=790
Jun 22 17:23:27 2007: CondWait After 7d5ff740; pid=263, file="msg.c", line=790
.....
Jun 22 17:23:18 2007: Succeeded 2ac34970; pid=216, file="shelffru.c", line=3171
Jun 22 17:23:18 2007: Released 2ac34970; pid=216, file="shelffru.c", line=3196
Jun 22 17:23:18 2007: Acquiring 2ac34970; pid=216, file="shelffru.c", line=3171
Jun 22 17:23:18 2007: Succeeded 2ac34970; pid=216, file="shelffru.c", line=3171
Jun 22 17:23:18 2007: Released 2ac34970; pid=216, file="shelffru.c", line=3196
Total of 32 locks
2ac34530: owner=0, waiting=0
2ac34ad0: owner=0, waiting=0
2acc1f80: owner=0, waiting=0
2ac344a0: owner=0, waiting=0
2acbe5ec: owner=0, waiting=0
```

```
2acc1de0: owner=0, waiting=0
2ac34950: owner=0, waiting=0
2ac36230: owner=0, waiting=0
2ac34b90: owner=0, waiting=0
2acbe620: owner=0, waiting=0
2ac34df0: owner=0, waiting=0
2ac34970: owner=0, waiting=0
1004bd60: owner=0, waiting=0
2ac36350: owner=0, waiting=0
10051070: owner=0, waiting=0
2ac34450: owner=0, waiting=0
2acc1e28: owner=0, waiting=0
2acc1e40: owner=0, waiting=0
10090890: owner=0, waiting=0
2ac344f0: owner=0, waiting=0
2acbdb98: owner=0, waiting=0
10091fa8: owner=0, waiting=0
10091c10: owner=0, waiting=0
100923e8: owner=0, waiting=0
10092840: owner=0, waiting=0
1008d7a4: owner=0, waiting=0
7efffa80: owner=0, waiting=0
2ac34430: owner=0, waiting=0
2ac36310: owner=0, waiting=0
2ac33d98: owner=0, waiting=0
1008d81c: owner=0, waiting=0
10094c08: owner=0, waiting=0
```

For backwards compatibility with older versions of the Shelf Manager, deadlock detection is turned off if the value of the configuration variable `DETECT_DEADLOCKS` is `FALSE`. By default, this configuration variable has the value `TRUE`. This variable controls both the watchdog-based and the direct deadlock detection mechanisms.

4.9 I²C Bus fault isolation

On some carriers, ShMM GPIOs can be used to reset an I²C multiplexer or I²C switch that is implemented on an I²C bus. By resetting the I²C multiplexer or switch, it is possible to implement an algorithm in the Shelf Manager to isolate faulty devices on subordinate I²C buses that are “behind” that multiplexer/switch, so that access to devices on other I²C buses can still occur successfully.

In the original implementation, this feature was applicable only to one multiplexer or switch connected to the Shelf Manager. GPIO E8 (GPIO_08 on the ShMM-700R) was used to reset it. On ShMM-700R carriers that have multiple master-only I²C buses, the number of the bus to which the multiplexer or switch was connected, could be configured. By default, however, logical I²C bus 5 (physical bus 3) on the ShMM-700R was used for this functionality.

Starting with release 3.2.0, this feature has been re-implemented to allow more than one I²C multiplexer or switch to be resettable. These multiplexers or switches can also be located (in the case of the ShMM-700) on multiple I²C buses. The new implementation is not backwards-compatible and applies only to HPDL-based shelves. The information about reset support and the number of the ShMM GPIO used to reset the device must be contained in the HPDL device description of the corresponding multiplexer or switch.

The following subsections describe the usage of this feature in the original and HPDL-based mode, respectively.

4.9.1 I²C bus fault isolation – original approach (single multiplexer or switch, not HPDL-based)

In this approach, the isolation algorithm is controlled by the value of the configuration parameter *ISOLATE_MUX_ON_GPIO8* and works as follows:

- The multiplexer/switch 7-bit I²C address can be specified by the configuration parameter *ISOLATE_MUX_ADDRESS*. The default is *0x70*.
- When the error ETIMEDOUT is detected by the Linux I²C subsystem while trying to access a device on a subordinate bus, this bus is recorded as faulty in the Shelf Manager and the multiplexer/switch is held in reset via GPIO E8. This operation causes all the channels to be deselected on the multiplexer/switch, which allows the Shelf Manager to access devices on the main I²C bus even if one of the subordinate buses has a fault.
- The multiplexer/switch is re-enabled via GPIO E8 the next time some parts of the Shelf Manager tries to access it. This allows the Shelf Manager to access devices on the subordinate buses that are still functional.
- When an attempt is made to access the faulty subordinate bus (by writing the channel number or mask for the faulty bus into the multiplexer/switch), this operation is not permitted for the “ignore count” subsequent attempts. The faulty bus is not enabled; an error is instead returned for the corresponding write operation. The parameter “ignore count” is specified by the configuration parameter *ISOLATE_MUX_IGNORE_COUNT*.
- The algorithm knows which value in the multiplexer/switch register corresponds to the faulty bus, because it monitors all write accesses to the multiplexer/switch registers and remembers the last value before the occurrence of the fault.
- When the “ignore count” is exhausted, an attempt is made to access the faulty bus and check whether the fault is still present. If the attempt fails, the “ignore count” is reset to its original value. If the attempt succeeds, then the recorded fault condition for that bus is removed and operation continues normally.
- With this approach, it is also possible to initiate a switchover to the backup Shelf Manager if the main I²C bus remains inaccessible after the resetting the I²C switch/multiplexer. If the configuration parameter *ISOLATE_MUX_SWITCHOVER_THRESHOLD* is defined as a non-zero value, it specifies the maximum consecutive number of attempts to reset the switch/multiplexer. If after that number of attempts the main I²C bus is still inaccessible, a switchover to the backup Shelf Manager is initiated and the active Shelf Manager is rebooted.

Warning: The `ISOLATE_MUX_SWITCHOVER_THRESHOLD` configuration parameter can ONLY be used with I²C switch/multiplexor devices that disable, or isolate, all I²C channels after a reset. The default logical I²C bus number on which the I2C multiplexer resides can be changed with the configuration parameter `ISOLATE_MUX_BUS`.

4.9.2 I²C bus fault isolation – HPDL-based approach (multiple multiplexers or switches)

In this mode, the isolation algorithm is controlled by the HPDL description of the carrier and the shelf. The original isolation approach in this case must be turned off (that is, the configuration variable `ISOLATE_MUX_ON_GPIO8` must be set to `FALSE` in the configuration files).

The typical HPDL description of an I²C multiplexer supporting reset looks as follows:

```
DEVICE Multiplexor {
    PCA9545: 5: 0x70;
    SIGNALS {
        Channels: Channels, "6", "7", "8", "9", "Reset=8", "SwitchoverThreshold=10";
        #MuxIsolated: IsolatedChannelMask;
        MuxIslCh1: Channel1_Isolated;
        MuxIslCh2: Channel2_Isolated;
        MuxIslCh3: Channel3_Isolated;
        MuxIslCh4: Channel4_Isolated;
    };
};
```

A configuration parameter for the signal “Channels” in the form “Reset=<n>” specifies the number of the ShMM GPIO used to reset the device. This GPIO is chosen by the designers of the ShMM carrier. For distinct multiplexers, different GPIOs must be used.

Signals with internal names in the form “Channel<n>_Isolated” report the isolation state of a specific channel of the multiplexer or switch (starting from 1). The value of the signal is 1 if the channel is currently isolated by the isolation algorithm and 0 otherwise. These signals can be exposed as IPMI sensors using normal HPDL mechanisms.

The signal “IsolatedChannelMask” reports the bitmask that has 1s for currently isolated channels (bit 0 of the bitmask corresponds to the first channel, bit 1 – to the second channel, etc.).

In this mode, the isolation algorithm operates as follows:

- When the error ETIMEDOUT is detected by the Linux I²C subsystem while trying to access a device on a subordinate bus, this bus is recorded as faulty in the Shelf Manager and the multiplexer/switch is held in reset via the corresponding GPIO. This operation causes all the channels to be deselected on the multiplexer/switch, which allows the Shelf Manager to access devices on the main master-only I²C bus even if one of the subordinate buses has a fault.
- The multiplexer/switch channel that caused the error is noted and values of the signals “Channel<n>_Isolated” and “IsolatedChannelMask” are modified accordingly.

- The multiplexer/switch is re-enabled via the corresponding GPIO the next time some part of the Shelf Manager tries to access it. This allows the Shelf Manager to access devices on the subordinate buses that are still functional.
- When an attempt is made to access the faulty subordinate bus (by writing the channel number or mask for the faulty bus into the multiplexer/switch), this operation is not permitted for a certain (configurable) period of time after the failure. During this period, the faulty bus is not enabled; an error is instead returned for the corresponding write operation. This time period is specified by the configuration parameter `ISOLATE_MUX_IGNORE_TIME`.
- The algorithm knows which value in the multiplexer/switch register corresponds to the faulty bus, because it monitors all write accesses to the multiplexer/switch registers and remembers the last value before the occurrence of the fault.
- When the “ignore time” expires, an attempt is made to access the faulty bus and check whether the fault is still present. If the attempt fails, the “ignore time” is reset to its original value. If the attempt succeeds, then the recorded fault condition for that bus is removed and operation continues normally.
- With this approach, it is also possible to initiate a switchover to the backup Shelf Manager if the main I²C bus remains inaccessible after the resetting the I²C switch/multiplexer. If the configuration parameter "SwitchoverThreshold=<N>" on the Channel signal is defined as a non-zero value, it specifies the maximum consecutive number of attempts to reset the switch/multiplexer. If after that number of attempts the main I²C bus is still inaccessible, a switchover to the backup Shelf Manager is initiated and the active Shelf Manager is rebooted.

4.10 Assignment of LAN configuration parameters to boards and modules

In shelves where some of the boards and intelligent modules (currently only AMC modules are considered) implement LAN IPMI channels, there may be a need for a centralized repository of LAN configuration parameters (mainly the IP addresses, subnet masks and gateway addresses) for these boards and modules. This repository can be maintained on the Pigeon Point Shelf Manager. The parameters are stored statically on the ShMM or in the Shelf FRU Information. Also the Shelf Manager supports obtaining parameters from a DHCP server using the DHCP protocol.

One use for such LAN IPMI channels, once configured, is to support serial over LAN sessions via the management controllers on the boards or modules. This use involves a significant number of additional IPMI parameters. In the current Shelf Manager, only the following three parameters are supported: IP address, subnet mask and default gateway address. Other parameters (for instance, MAC addresses) may be added in future releases of the Shelf Manager.

Note: the functionality described in this section has been made obsolete by the HPM.3 specification and is not recommended to be used in new systems. It is described here for compatibility purposes. The Shelf Manager fully supports the mechanisms defined in the HPM.3 specifications to retrieve network configuration parameters for boards and modules from a DHCP server using the DHCP Proxy approach and assigns them to the target boards/modules with a series of “Set LAN Configuration Parameters” IPMI commands. See the HPM.3 specification for more information.

4.10.1 Structure and composition of the supported parameters

The parameters are specified on a per-slot basis. Slots are defined for both main boards and modules. For main boards, the slot number is the board physical slot number (site number). For AMC modules, the slot number is the combination of the AMC carrier physical slot number and AMC site number. For each slot, several parameter sets may be defined. Each parameter set contains parameters (currently the IP address, subnet mask and default gateway address) for one LAN channel. Parameter sets for a specific slot are ordered; they are assigned to a board/module in the order in which the Shelf Manager polls LAN channels of a board/module (see the definition of the configuration variable `BOARD_LAN_PARAMETERS_CHANNEL_LIST`).

In general, parameters sets are retrieved by the Shelf Manager dynamically on demand. During initialization, after the Shelf FRU Information is found and the shelf Address Table is known, the Shelf Manager retrieves parameter sets for all currently active boards and modules (in states M3 and M4) and performs the assignment. After initialization, when a new board or a module is installed in the shelf and reaches state M3, the Shelf Manager retrieves the corresponding parameter sets for it and performs the assignment. This assignment is performed synchronously, before powering up the payload of the corresponding board or module.

4.10.2 Obtaining LAN configuration parameters on the Shelf Manager level

There are three ways to obtain LAN configuration parameters on the Shelf Manager level.

1. Retrieved from a DHCP server

In this case (when the configuration variable `BOARD_LAN_PARAMETERS_USE_DHCP` is set to `TRUE`), the Shelf Manager retrieves parameter sets from a DHCP server. The same algorithm for finding the appropriate DHCP server is used, as in the case of obtaining Shelf Manager IP addresses via DHCP.

The queries are based on Client Identifiers, not on MAC addresses. In the default implementation, the Client Identifier is based on the Shelf Address string that is stored in the Shelf FRU Information, plus three endpoint Request Identifier (Request ID) bytes in the following format:

- Identifier byte 1 indicates the board physical slot address; the identifier byte value is slot address + 2Fh).
- Identifier byte 2 is the AMC site number on the board; for boards, this byte has value 0.
- Identifier byte 3 provides a parameter set number within the specific slot. For boards/modules that require the configuration of more than one LAN channel, the value of this byte is 0 for the first parameter set, 1 for the second parameter set and so on. For boards/modules that require configuration of a single LAN channel, this byte is 0.

Parameters for different slots are retrieved in parallel. The same sequence of retransmission timeouts (1, 2, 4, 8, 16, 32, 64 ... 64 seconds) applies if the DHCP server does not respond, as in the case of retrieval of the Shelf Manager parameters. The lease time is infinite.

The Shelf Manager maintains internally a record of assigned parameter sets. After a specific parameter set is no longer needed (e.g. when the board or the corresponding module is extracted), the Shelf Manager removes the corresponding parameter set from its records and sends a DHCP Release request for the corresponding Client ID to the DHCP server. This is done to prevent

exhaustion of the pool of IP addresses available to the DHCP server for allocation, in configurations where this pool is relatively small.

2. Read from a flash file on the ShMM

After the Shelf Manager starts up, it retrieves the LAN configuration parameters from a file on the ShMM `/var/nvdata/lan_config_params`. This is the simplest approach to retrieve the parameters.

The file has a text format. Each line defines one parameter set and has the following format:

```
<physical-slot> [", " <amc-slot> ] "=" <ip-address> ", " <subnet-mask> ", "
<default-gateway-ip-address>
```

Comment lines starting with `#` are allowed and ignored. Also, the rest of the line after a `#` character is treated as a comment and ignored. Empty lines are allowed and ignored, too.

For each slot, the file can define several parameter sets. All parameter sets defined for the same slot must be placed in the file together and have the same values to the left of the equal sign. They are assigned to LAN channels on the corresponding boards or modules that require such assignments.

3. Read from the Shelf FRU Information

In this case, there is a set of custom multirecords with the following Manufacturer ID (IANA Enterprise Number) assigned to nVent, Schroff GmbH (formerly Pentair Technical Products) (00400Ah)[†] and Record ID = 0Ch. The concatenation of the data in these multirecords (optionally compressed with *gzip*) represents the parameter data in the same format as for the text file above. Management of parameter sets is also performed in the same way as for the previous case.

An example of a LAN parameters configuration file is given below:

```
# Parameters for a board in physical slot 1
1 = 192.168.1.101, 255.255.255.0, 192.168.1.1
1 = 192.168.1.102, 255.255.255.0, 192.168.1.1
1 = 192.168.1.103, 255.255.255.0, 192.168.1.1
# Parameters for an AMC in the first AMC slot of a board in slot 1
1,1 = 192.168.1.104, 255.255.255.0, 192.168.1.1
1,1 = 192.168.1.105, 255.255.255.0, 192.168.1.1
1,1 = 192.168.1.106, 255.255.255.0, 192.168.1.1
# Parameters for an AMC in the second AMC slot of a board in slot 1
1,2 = 192.168.1.107, 255.255.255.0, 192.168.1.1
1,2 = 192.168.1.108, 255.255.255.0, 192.168.1.1
...
# Parameters for a board in physical slot 2
2 = 192.168.1.121, 255.255.255.0, 192.168.1.1
2 = 192.168.1.122, 255.255.255.0, 192.168.1.1
2 = 192.168.1.123, 255.255.255.0, 192.168.1.1
...
```

[†]This IANA Enterprise Number was originally assigned to Pigeon Point Systems, which has since been acquired by Pentair.

4.10.2.1 Using the FRU Info Compiler to place LAN configuration data in the Shelf FRU Information

The FRU Information Compiler implements special syntax that can be used to incorporate a LAN configuration parameters file into a FRU Information image that is being created.

For example, the following instructions add a compressed LAN configuration parameters file `lan_config_params.gz` to the current FRU Information image:

```
...
[nVent.Board/AMC LAN Configuration]
File = lan_config_params.gz
```

4.10.2.2 Using the Shelf Manager Command Line Interface to place LAN configuration parameters in the Shelf FRU Information

An extended version of the CLI command `frudataw` can be used to dynamically update Shelf FRU Information with LAN configuration information. The corresponding configuration file (optionally compressed) should first be downloaded to the ShMM.

The corresponding command line has the following syntax:

```
clia frudataw -l <ipmb-address> <fru-id> <in-file>
clia frudataw -l -c <ipmb-address> <fru-id>
```

The parameters `<ipmb-address>` and `<fru-id>` specify the IPMB address and FRU device ID of the FRU Information to update. For the Shelf FRU Information, use IPMB address 20h and FRU device ID 254.

The parameter `<in-file>` indicates the LAN configuration parameters file (which is possibly compressed) to be stored in the target FRU Information.

Option `-c` removes (clears) the LAN configuration data from the target FRU Information.

For example, the following command can be used to update the Shelf FRU Information with a new version of the LAN configuration parameters file:

```
#clia frudataw -l 20 254 lan_config_params.gz
```

4.10.2.3 Configuring a Linux DHCP server for obtaining LAN configuration parameters

The `/etc/dhcpd.conf` file on the DHCP server should have a record with a unique DHCP client identifier string for each channel of each FRU for which the LAN configuration parameters need to be assigned. Unless otherwise documented explicitly for certain specific types of ATCA shelves, the format of the client identifier is as follows:

```
<dhcp-client-identifier> ::= <shelfaddress in bytes> : <2F + board physical number> : <FRU ID> : <channel number>
```

The example below is for a shelf with 2 board slots (with IPMB addresses 82h, 84h) and the shelf address "XXX" (in ASCII):

```
##### Board LAN Configuration Params #####

# 0x82,0 0
host board1_fru0_ch0 {
```

```
    option dhcp-client-identifier 58:58:58:30:0:0;
    fixed-address 192.168.1.202;
    option routers 192.168.1.253;
    option subnet-mask 255.255.255.0;
}

host board1_fru0_ch1 {
    option dhcp-client-identifier 58:58:58:30:0:1;
    fixed-address 192.168.1.202;
    option routers 192.168.1.253;
    option subnet-mask 255.255.255.0;
}
host board1_fru0_ch2 {
    option dhcp-client-identifier 58:58:58:30:0:2;
    fixed-address 192.168.1.202;
    option routers 192.168.1.253;
    option subnet-mask 255.255.255.0;
}
host board1_fru0_ch3 {
    option dhcp-client-identifier 58:58:58:30:0:3;
    fixed-address 192.168.1.202;
    option routers 192.168.1.253;
    option subnet-mask 255.255.255.0;
}
host board1_fru0_ch4 {
    option dhcp-client-identifier 58:58:58:30:0:4;
    fixed-address 192.168.1.202;
    option routers 192.168.1.253;
    option subnet-mask 255.255.255.0;
}

# 0x84,0 0
host board2_fru0_ch0 {
    option dhcp-client-identifier 58:58:58:31:0:0;
    fixed-address 192.168.1.203;
    option routers 192.168.1.253;
    option subnet-mask 255.255.255.0;
}
host board2_fru0_ch1 {
    option dhcp-client-identifier 58:58:58:31:0:1;
    fixed-address 192.168.1.203;
    option routers 192.168.1.253;
    option subnet-mask 255.255.255.0;
}
```

```

host board2_fru0_ch2 {
    option dhcp-client-identifier 58:58:58:31:0:2;
    fixed-address 192.168.1.203;
    option routers 192.168.1.253;
    option subnet-mask 255.255.255.0;
}
host board2_fru0_ch3 {
    option dhcp-client-identifier 58:58:58:31:0:3;
    fixed-address 192.168.1.203;
    option routers 192.168.1.253;
    option subnet-mask 255.255.255.0;
}
host board2_fru0_ch4 {
    option dhcp-client-identifier 58:58:58:31:0:4;
    fixed-address 192.168.1.203;
    option routers 192.168.1.253;
    option subnet-mask 255.255.255.0;
}

```

4.10.3 Dispatching LAN configuration parameters to boards

For each board or module, the Shelf Manager detects whether it supports LAN channels and expects the configuration parameters for them to be assigned by the Shelf Manager. If this is the case for a particular board or module, the Shelf Manager sends it the configuration parameters from the next parameter set for the corresponding physical slot. Shelf Manager does this for all boards or modules in states M3-M5 during startup (after the Shelf FRU Information is found). After the initialization, Shelf Manager performs this procedure during board or module activation (when the board or module reaches state M3). The Shelf Manager maintains internal records to avoid sending configuration parameters to the same board or module twice.

To detect what LAN channels a specific board or module has, the Shelf Manager sends it the command “Get Channel Info”, with a channel number picked sequentially from the *BOARD_LAN_PARAMETERS_CHANNEL_LIST*. If a board or module responds with the Channel Medium Type = 802.3 LAN for some channel, the Shelf Manager sends the command “Get LAN Configuration Parameters (IP Address Source)” for that channel. If the IP Address source = 3 (address loaded by BIOS or system software), the Shelf Manager assigns configuration parameters for this channel. Commands are sent on IPMB-0; for modules, commands are directed to the module address on IPMB-L and are wrapped with “Send Message” commands directed to the IPMB-0 address of the carrier board.

Parameter sets are obtained in order as they are needed for a specific slot. If a parameter set cannot be obtained (for instance, if parameter sets are exhausted for a specific slot), configuration parameters are not assigned.

The Shelf Manager sends the configuration parameters to a board or module as a set of “Set LAN Configuration Parameters” commands, over IPMB-0; however the channel parameter in these

commands is set to channel N. For modules, commands are wrapped in “Send Message” commands. The following batch of “Set LAN Configuration Parameters” commands is sent:

- Channel = N, Parameter = Set In Progress (0), data = set in progress
- Channel = N, Parameter = IP Address (3), data = board IP address
- Channel = N, Parameter = Subnet Mask (6), data = board subnet mask
- Channel = N, Parameter = Default Gateway (12), data = board default gateway IP address
- Channel = N, Parameter = Set In Progress (0), data = commit write
- Channel = N, Parameter = Set In Progress (0), data = set complete

4.10.4 Synchronous assignment of LAN configuration parameters to boards

In the case of obtaining parameters from a DHCP server (if the configuration variable `BOARD_LAN_PARAMETERS_USE_DHCP` is `TRUE`), activation of LAN-enabled boards and modules (transitioning from M3 to M4) is not affected by the process of retrieving LAN configuration parameters for them; a board may reach M4 before its parameters are retrieved and assigned.

However, it is possible to instruct the Shelf Manager to assign LAN configuration parameters to designated boards/modules in a synchronous way before allowing them transition to M4. This is done by delaying power assignment to these FRUs until LAN configuration parameters are retrieved and assigned to them. This feature is controlled by a Shelf Manager configuration variable `BOARD_LAN_PARAMETERS_SYNCHRONOUS` that contains the list of all FRUs that are to be treated this way. Using synchronous assignment of LAN configuration parameters can indefinitely hold a target FRU in state M3, so it should be used with caution.

The string value of the configuration parameter `BOARD_LAN_PARAMETERS_SYNCHRONOUS` represents the list of descriptors for FRUs that require synchronous assignment of LAN configuration parameters. This list is meaningful only if the configuration parameter `BOARD_LAN_PARAMETERS_USE_DHCP` is set to `TRUE`. The following grammar applies to this value:

```
<descriptor-list> ::= <descriptor> | <descriptor-list> "," <descriptor>
<descriptor> ::= <ipmb-address> ["." <fru-designator>] | "*"
<fru-designator> ::= <fru-device-id>[:<fru-device-id>] | <site-type> "." <site-number-or-wildcard> | "*"
<site-number-or-wildcard> ::= <site-number>[:<site-number>] | "*"

```

Descriptors in the list are comma-separated. Each descriptor consists of the IPMB address and the FRU designator (either the FRU device ID or the site type and site number of the target subsidiary FRU). The IPMB address can be specified alone, or the FRU designator can be specified as a wildcard (*); in either case, synchronous assignment applies to all FRUs associated with that IPMB address. A wildcard descriptor (*) applies to all FRUs in the shelf.

If a FRU designator is specified, it is separated from the IPMB address by a dot. Site type and site number are also each separated by a dot. An additional FRU device ID or site number for this IPMB address can be designated, separated by semicolon. A site number can be specified as a wildcard (*), which makes the descriptor apply to all FRUs with the specified site type on a given IPMC. An

IPMB address is hexadecimal by default and FRU device ID, site type and site number are decimal by default.

For example, the value `82.0,86,8A.7.2:5,90.1:3,A0.7.*` applies synchronous assignment to the IPMC @ 82h, FRU 0; all FRUs associated with the IPMC @ 86h, AMC #2 and #5 associated with the IPMC @ 8Ah, FRU 1 and FRU 3 associated with the IPMC @ 90h and all AMC modules associated with the IPMC @ A0h.

4.11 HPI System Event sensor

This discrete sensor is associated with the logical Shelf Manager (IPMB address 20h), has number 0 on LUN 3, the name "HPI Sys Event" and OEM event/reading type code DBh. Its purpose is to enhance interaction between the Shelf Manager and Pigeon Point HPI implementations: IntegralHPI and Pigeon Point OpenHPI. This sensor sends IPMI events in a special format to signal HPI implementations that changes have occurred within the Shelf Manager, for example some fields within a Multi-record in the Shelf FRU Information have been updated. The HPI implementation is expected to process such an event and perform corresponding actions to react to the new conditions, for example to reread the Shelf FRU Information. Since this sensor is intended to be used for internal communication between the Shelf Manager and Pigeon Point HPI implementation, its behavior, set of states and event format are subject to change and it is not recommended for an ordinary user to make any direct use of it.

4.12 ShMM tests available via the diagnostic infrastructure

A Diagnostic Initiators (DI) facility has been implemented in the Shelf Manager, starting with release 2.8.0. Currently, nVent OEM commands are used for access to the diagnostic initiators; they are documented in the *Pigeon Point Shelf Manager External Interface Reference*. The overall infrastructure is documented in the IPM Controller Diagnostic Initiator Architecture document.

The Diagnostic Initiators facility is implemented on the physical Shelf Manager and provides access to several run-time tests for the ShMM. All tests are associated with Diagnostic Initiator 0. Cancellation is not supported for any of the tests. Parameters for the tests are described in the following table. The results string for most of these tests has the form "Status: XX", where XX is a two-digit hexadecimal test-specific reason code ("00" if the test completes successfully).

For the I²C test, the results string in the case of success is a list of 2-character hexadecimal addresses for I²C devices present on the bus. For the IPMB test, the results string (in the case of both success and failure) is the list of extended IPMB state and error flags reported by the driver.

Some of the tests are not implemented on ShMM-700 because they don't make sense for that platform; this is noted explicitly in the table below. However, to ensure compatible allocation of test identifiers across all platforms, the corresponding tests are still described as present on the ShMM-700; when invoked, such tests report failure, with the textual result string: "This test is not implemented".

Table 20: ShMM tests implemented in the diagnostic infrastructure

TEST NAME	TEST ID	PARAMETERS (NAME, TYPE, DESCRIPTION)	TEST DESCRIPTION
“UART INT” (Internal UART Loopback Test). Not implemented on ShMM-700.	0	“uart”: INT32: UART ID to test Acceptable values are 0 and 1. Default value is 0.	The UART RX and TX channels are internally looped back while the test is in progress, which may result in data loss for external communication on the affected UART. For example, if UART #0 is the system console, some console output may be lost or corrupted.
“UART EXT” (External UART Loopback Test). Not implemented on ShMM-700.	1	“uart”: INT32: UART ID to test Acceptable values are 0 and 1. Default value is 0.	This test requires manual interaction with the ShMM carrier to loop back the UART RX and TX pins.
“ETHER INT” (Internal Ethernet Loopback Test)	2	“mac”: INT32: MAC ID to test Acceptable values are 0 and 1. Default value is 0.	The RX and TX channels are internally looped back while the test is in progress, which may result in data loss and/or a broken network connection.
“ETHER PHY” (Ethernet PHY Loopback Test)	3	“mac”: INT32: MAC ID to test Acceptable values are 0 and 1. Default value is 0.	The RX and TX channels are internally looped back within the PHY while the test is in progress, which may result in data loss and/or a broken network connection.

TEST NAME	TEST ID	PARAMETERS (NAME, TYPE, DESCRIPTION)	TEST DESCRIPTION
"ETHER ARP" (Ethernet ARP Test)	4	<p>"mac": INT32: MAC ID to test</p> <p>Acceptable values are 0 and 1.</p> <p>Default value is 0.</p> <p>"ip": INT32: IP Address to resolve during the test.</p> <p>"timeout": INT32: Timeout value (in seconds) for the test.</p>	The test sends an ARP request with the specified IP address and waits for a reply in the specified timeout range.
"FLASH CRC" (Flash Checksum Test)	5	<p>"mtd": INT32: Software image number (MTD partition number for ShMM-500/1500) for the test.</p> <p>Acceptable values are:</p> <ul style="list-style-type: none"> • 2 – Active Linux kernel • 3 – Active U-Boot • 4 – Active RFS • 7 – Inactive Linux kernel • 8 – Inactive U-Boot • 9 – Inactive RFS • 10 – Active application image (ShMM-700 only) • 11 - Inactive application image (ShMM-700 only) 	The test calculates a checksum for the software image (for ShMM-500/1500, located in the corresponding MTD partition) and compares it with the checksum in the image header.

TEST NAME	TEST ID	PARAMETERS (NAME, TYPE, DESCRIPTION)	TEST DESCRIPTION
"I2C TEST" (I ² C Bus Test)	6	<p>"bus": INT32: I²C bus number to test.</p> <p>Acceptable values are 0 and higher (bus 0 represents the main master-only I²C bus going to the ShMM carrier (and often beyond the carrier to the shelf); other bus numbers are platform-specific).</p> <p>Default value is 0.</p>	The test tries to access the specified I ² C adapter device and enumerate all I ² C devices present on that bus by reading from all possible I ² C addresses on the bus and checking which addresses are acknowledged.
"IPMB TEST" (IPMB Bus Test)	7	<p>"bus": INT32: IPMB bus number to test.</p> <p>Acceptable values are 0 and higher (0 represents IPMB-A, 1 represents IPMB-B, higher values represent radial IPMB links).</p> <p>Default value is 0.</p>	The test queries the IPMB bus driver for extended information about the specified bus, which includes state flags and error flags. This information is then converted to a textual format and reported in the results string.

5 Using the IPMI analysis tools

The Shelf Manager software includes tools for analyzing the IPMI messaging traffic within a shelf. Categories of such traffic include IPMB messages going between the ShMM and other IPM controllers in the shelf as well as IPMI traffic between the ShMM and external management applications via Ethernet. The PICMG HPM.2 specification enables other traffic to be collected and analyzed as well, including IPMB-L traffic among a Carrier IPMC and the Module Management Controllers (MMCs) it represents, as well as IPMI messaging between an IPMC and its payload via the Payload Interface (which is called the System Interface in IPMI). These tools consist of the following two parts:

- The IPMB trace collection daemon. This application runs on the ShMM and interfaces with the IPMB-specific extensions in the Monterey Linux / Pigeon Point Linux kernel on the ShMM to collect IPMB trace data from IPMB-0. This daemon is also able to collect trace data from HPM.2-compliant management controllers in the shelf, forwarded via RMCP+ connections between the daemon and those management controllers. Even though this trace data is not necessarily only from IPMB, the name “IPMB trace daemon” is still used. The collected data is either stored into a file in a local or NFS-mounted file system (available on all ShMM variants), or transferred directly to the GUI client on a separate platform (available on ShMM-500/1500 only).
- The IPMI trace analyzer. This role is typically played by a GUI application that supports viewing and analyzing an IPMB trace collected by the daemon, either from a file or directly from a ShMM via a TCP/IP connection. Wireshark and its terminal-oriented variant *tshark* are two specific applications that can play this role. In addition, Wireshark and *tshark* can also collect and analyze IPMI packets that are transferred over an Internet Protocol (IP) based network via RMCP or RMCP+, both of which are layered above UDP.

5.1 In this section

This section contains the topics listed below. Just click on a topic to go to it.

- IPMB trace collection daemon (*ipmb_traced*)
- IPMI trace analyzer
- Collecting and analyzing a trace using the GUI and command line tools

5.2 IPMB trace collection daemon (*ipmb_traced*)

The IPMB trace daemon can operate in unattended trace mode or controlled trace mode.

In the unattended trace mode, the trace data is collected into a file in a ShMM-local or NFS-mounted file system. When the trace collection has been completed, the file can be transferred to the development host, then read and reviewed via the trace analyzer.

In the controlled trace mode, the trace data is transferred directly to the trace analyzer via a TCP/IP connection. The trace analyzer can display and analyze the collected data on a real time basis.

The IPMB trace collection daemon (*ipmb_traced*) runs on the ShMM and takes the role of IPMB trace collector.

The daemon interacts with the IPMB-specific extensions in the Linux kernel to capture IPMB traffic on IPMB-0. Specifically, the “promiscuous mode” of I²C adapter operation is used to capture all messages on the IPMB. (Promiscuous mode configures the adapter to receive all IPMB messages on the bus, not just the messages that are addressed to or issued by the ShMM, itself. However promiscuous mode is supported only on the ShMM-1500. On the ShMM-500 and ShMM-700, due to hardware limitations, adapters can only receive messages that are issued by or directed to a ShMM, itself, so promiscuous mode does not have its intended effect.)

In addition to IPMB-0 traffic, the daemon can also collect IPMI messaging traces from boards or modules in the shelf, via RMCP+ sessions with the management controllers on those boards or modules. To begin collecting a trace, the daemon establishes an RMCP+ session with the target board using a special type of payload: the IPMI Trace Payload that is defined in the HPM.2 LAN Attached IPM Controller Specification. Once the RMCP+ session with the IPMI Trace Payload is established, the IPM Controller begins sending trace data to the daemon in the IPMI Trace Payload format, and continues to do so until the session is terminated.

A specific instance of the daemon can collect either local or remote IPMI messaging trace, but not both. Also, when collecting a remote trace, it can do so only from one IPM Controller in the shelf. To collect a remote trace from each of multiple IPM Controllers, multiple instance of the daemon can be started on the ShMM.

When collecting a local or remote trace, the trace collection daemon runs in either the unattended trace or controlled trace modes.

When initiated in unattended trace mode, the daemon initiates trace collection and starts writing collected data into a file. This continues until the daemon is terminated by the user via *^C* or the SIGINT signal is sent to the daemon process operating as a background Linux application. The resulting trace file can then be transferred to another host for analysis using the IPMB trace analyzer.

When started in controlled trace mode, the daemon begins to listen for an incoming TCP/IP connection from the IPMB analyzer running on another host. Once the connection is established, the daemon initiates trace collection and starts transferring collected data to the analyzer over TCP/IP. This continues until the GUI terminates the connection. The daemon can be terminated by the user via *^C* or the SIGINT signal sent to the daemon process operating as a background Linux application.

The daemon command line has the following syntax:

```
ipmb_traced [OPTIONS]
```

The options are described in the following table:

Table 21: Options for the IPMB trace collection daemon

OPTION NAME	DESCRIPTION
<i>-u</i> <file_name> <i>--unattended</i> =<file_name>	Specify a file for collecting trace data in the unattended trace mode. By default the file name is

	<code>/var/ipmb_traced/ipmb_traced.log</code>
<code>-c </code> <code>--controlled</code>	Run daemon in the controlled trace mode
<code>-e </code> <code>--events</code>	Trace IPMB state change events. Applies to local trace mode only.
<code>-d <usec> </code> <code>--delay=<usec></code>	Specify a delay in microseconds for polling the IPMB-A/IPMB-B interfaces. This option has an effect only if the <code>-e (--events)</code> option is enabled. If there are IPMB state change events, they are extracted and collected at intervals that are no longer than the defined delay, but always whenever an IPMB message is received. The delay value is specified in a range of 0 to 1000000 microseconds (1 second). If the <code>-d</code> option is omitted, the delay value defaults to 100000 microseconds (0.1 second). Applies to local trace mode only.
<code>-p <port> </code> <code>--port=<port></code>	In controlled trace mode, specify the TCP port number to listen to; the client application connects to this port. Default port = 19000.
<code>-q <size> </code> <code>--quota=<size></code>	In unattended trace mode, the maximum trace file size (in KB). After this size is reached, the existing file is closed and renamed with the extension <code>.old</code> and a new trace file is created.
<code>-w </code> <code>--wireshark</code>	Produce output in the Wireshark format (otherwise, <code>ipmb_traced</code> produces output in the Ethereal format). On ShMM-700, only Wireshark output format is supported, so this option is ignored on ShMM-700.
<code>-b <ip_address> </code> <code>--board=<ip_address></code>	Specify board trace mode, in which IPMI messaging is traced on another IPM controller that is connected to the network. The trace daemon obtains the trace using an RMCP+ session via a HPM.2 IPMI Trace Payload with the target IPMC. The parameter <code><ip_address></code> specifies the IP address of the target IPMC.
<code>-m <mask> </code> <code>--mask=<mask></code>	In board trace mode, specify the bit mask of IPMI channels that are to be traced on the target IPMC. This bit mask is represented by

	<p>the numeric parameter <i><mask></i> (with prefix 0x used to specify a hexadecimal number).</p> <p>The following channel numbers can be used:</p> <ul style="list-style-type: none"> Channel 0 (bit mask 1) = IPMB-0. Channel 7 (bit mask 0x80) = IPMB-L. Channel 15 (bit mask 0x8000) = System (Payload) Interface. <p>The default value is 1 (trace IPMB-0 only).</p>
<pre>-i <number> --instance=<number></pre>	<p>In board trace mode, specify the payload instance number to use for establishing the RMCP+ session. This number must match the IPMI Trace Payload instance number used by the target IPMC for IPMB tracing; refer to the board documentation for this number. The default value is 1.</p>
<pre>-U <username> --username=<username></pre>	<p>In board trace mode, specify the RMCP+ user name for establishing the RMCP+ session with the board. The default user name is <i>rmcp</i>. Passwords are not used.</p>
<pre>-r [<number>] -- radial[=<number>]</pre>	<p>Enable radial IPMB link monitoring. Without this option, <i>ipmb_traced</i> doesn't monitor any radial IPMB links in the shelf. Optional <i><number></i> parameter allows limiting the number of device nodes to open. The default is 17, which is a cross-link plus 16 links for each board in the chassis.</p>
<pre>-h --help</pre>	<p>Display help text and exit</p>

If no options are provided, *ipmb_traced* runs in unattended trace mode, produces trace data in Ethernet format (Wireshark format on the ShMM-700) and puts them in the */var/ipmb_traced/ipmb_traced.log* file, but does not trace IPMB state change events, monitor radial IPMB links or collect any HPM.2-based IPM controller's trace data. This scenario assumes that a separate ShMM application, such as the Pigeon Point Shelf Manager, assigns IPMB addresses for the ShMM and attaches to the ShMM IPMB interface so that there can be IPMB traffic to trace.

The trace collection daemon stores the collected data in a file (or sends it directly to the analyzer) in the form of packets. Each packet describes a single IPMB transaction or event and contains the following fields. (The detailed layout of the fields depends on whether the Ethernet or Wireshark format has been selected for the trace collection daemon.)

Table 22: IPMB trace collection packet field

FIELD NAME	DESCRIPTION
Seconds	Transaction completion time in seconds
Micro seconds	Transaction completion time in microseconds (calculated from the start of the last second)
Event	IPMB interface state change information
Flags	Special flags returned by the Linux kernel driver
Bus	IPMB interface (IPMB-A or IPMB-B) on which the transaction was captured
Address	Destination I ² C address of the transaction
Size	Size of the transaction
Data	Transaction raw data

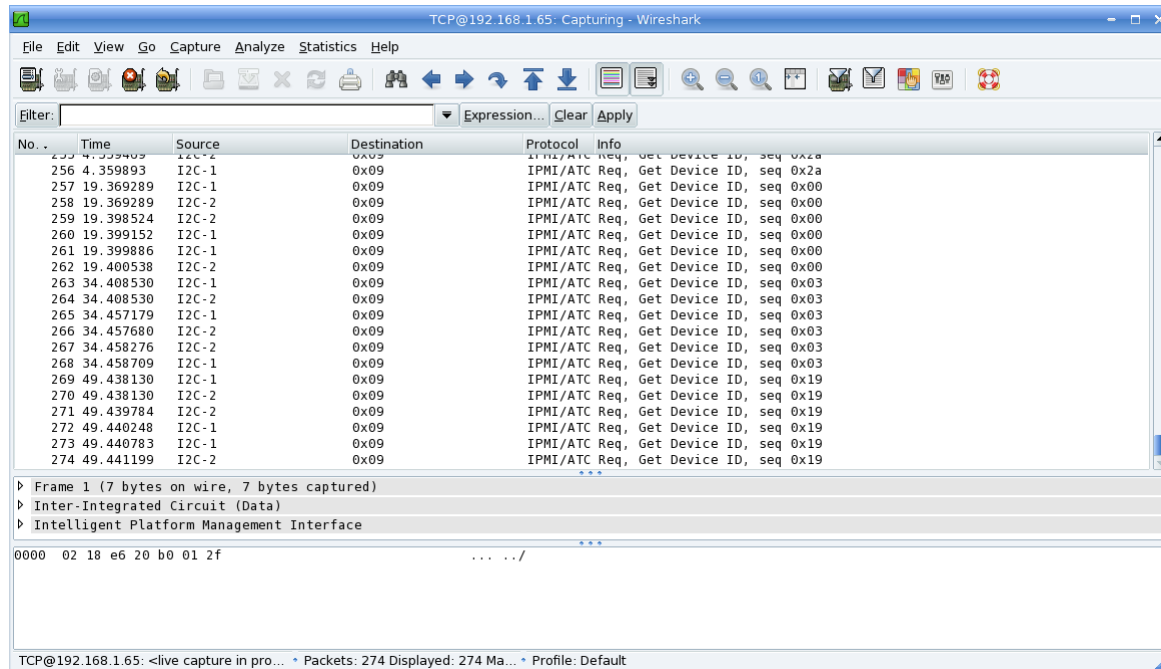
5.3 IPMI trace analyzer

The recommended IPMI trace analyzer for Pigeon Point Shelf Manager applications is a network traffic analyzer program called Wireshark (see www.wireshark.org), with its terminal-oriented variant, *tshark*. Wireshark provides a convenient interface for a complete analysis of layered protocol stack data. The Wireshark design allows developers to easily add support for new protocols, entire protocol stacks, and trace collection data sources in a collaborative manner. There are versions of Wireshark for both Windows and Linux environments. All these considerations combine to make Wireshark a great framework for IPMI trace analyzer applications. The following sections describe the key user interfaces of Wireshark that are relevant to IPMI trace analysis. Please consult the extensive user documentation for Wireshark at www.wireshark.org for more details.

5.3.1 Introduction to the Wireshark GUI

This section provides a brief introduction to the GUI; section 5.4 provides a step by step example of using it. Figure 6 provides a view of the GUI main window, which is composed of three frames.

Figure 6: Main window (IPMB trace analysis)



The upper frame shows a list of captured messages that provides basic information about each IPMB message. The list can be sorted by the contents of each field as a key.

In an IPMB trace analyzer usage, this frame displays the full list of collected IPMB messages. For each message, the following basic information is provided:

Table 23: IPMB message list basic information

GUI FRAME FIELD	INTERPRETATION
No.	The index (number) of the message in the trace
Time	Message date and time stamp, displayed as an offset from the first packet. The time stamp obtained on the target is available in the detailed packet information in the bottom frame.
Source	The IPMB the message was captured on. This field is either <i>I2C-1</i> for <i>IPMB-A</i> or <i>I2C-2</i> for <i>IPMB-B</i>
Destination	Destination I ² C address of the message. (in 7-bit format). This field is "–" for IPMB events.
Protocol	The type of the captured message. This field can be one of the following: <i>IPMI/ATCA</i> for an IPMI/ATCA formatted message

GUI FRAME FIELD	INTERPRETATION
	<i>EVENT</i> for an IPMB event <i>I2C READ</i> or <i>I2C WRITE</i> for a raw I ² C transaction
Info	Basic information about the captured message or event

The middle frame provides the protocol break down and data interpretation for a message selected in the message list provided in the upper frame.

In an IPMB trace analyzer usage, this frame provides interpretation of the layered contents of the message. Three protocol layers are supported:

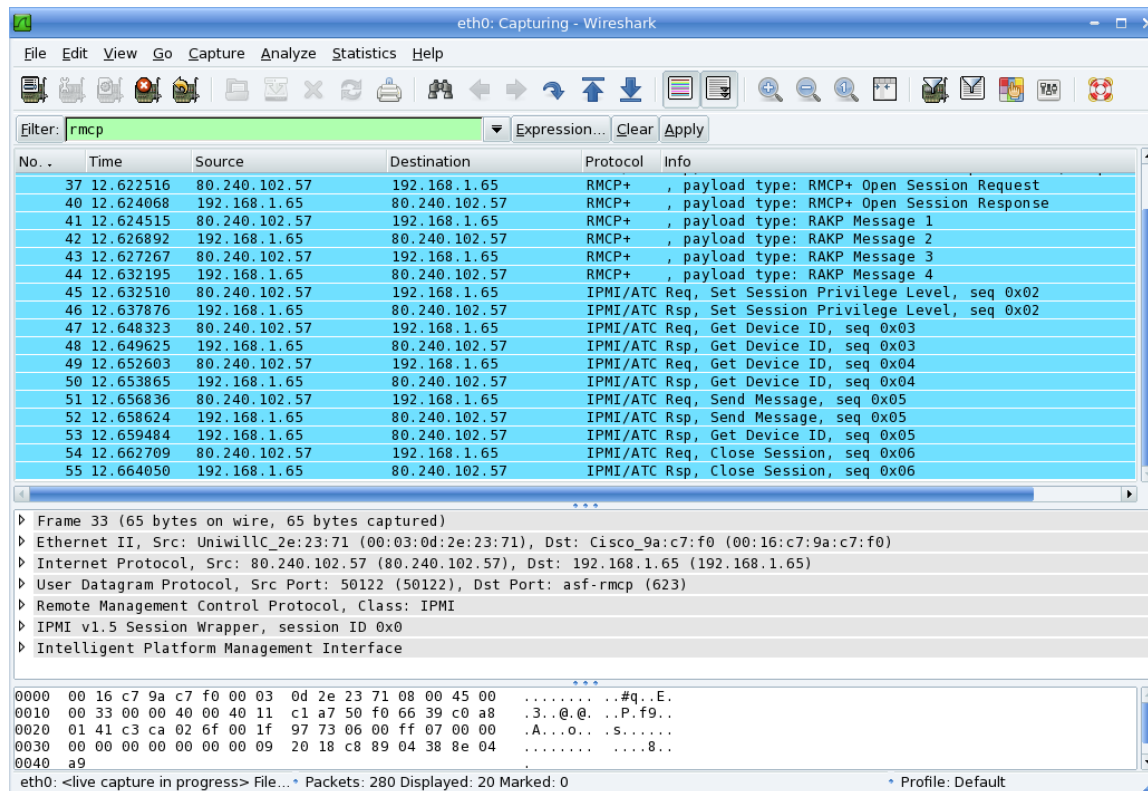
- Layer 1 provides low-level information about a message, such as the frame size, arrival time relative to the start of the trace, and the frame protocol (i2c:ipmi for IPMI messages).
- Layer 2 includes physical IPMB-0 and I²C information, such as the bus the message was captured on (IPMB-A or IPMB-B) and the destination I²C address of the frame.
- Layer 3 is available only for IPMI/ATCA messages, and provides message header information and information about the request/response data. Specifically, the message is identified as either a request or a reply, and the request/reply header fields are interpreted. The data fields are interpreted as defined by IPMI and PICMG specs for a particular request or response.

The bottom frame provides interpreted, hexadecimal and textual representations of a message selected in the message list provided in the upper frame.

The GUI provides a number of other GUI elements (dialogs, menu, windows, etc.) that are used to exercise various features of the IPMI Analyzer as well as customize its behavior. For instance, there is a trace collection dialog, trace filter window, and others. These facilities are described in detail in the Wireshark user documentation and man pages. Section 5.4, later in this document, provides a step by step procedure to use selected GUI elements for tracing and analyzing IPMB data.

When analyzing RMCP or RMCP+ based IPMI network traffic, the main GUI window looks similar, but with some small differences (as shown in Figure 7 below).

Figure 7: Main window (IPMI over network trace analysis)



Instead of I²C-related information, the Source and Destination columns show the source and destination IP addresses of the captured message. Instead of the three protocol layers in the case of an IPMB trace, seven protocol layers are shown:

- Layer 1 provides low-level information about a message, such as the frame size, arrival time relative to the start of the trace, and the list of protocols used in the frame.
- Layer 2 provides Ethernet-specific information about the frame (source and destination MAC addresses and the protocol type).
- Layer 3 provides IP-specific information from the frame.
- Layer 4 provides UDP-specific information from the frame.
- Layer 5 provides RMCP header information from the frame (version, sequence number and class).
- Layer 6 includes additional information about the RMCP or RMCP+ specific fields in the message.
- Layer 7 is available only for IPMI/ATCA messages, and provides message header information and information about the request/response data. Specifically, the message is identified as either a request or a reply, and the request/reply header fields are interpreted. The data fields are interpreted as defined by IPMI and PICMG specs for a particular request or response.

5.3.2 Introduction to terminal-oriented Wireshark

The Wireshark tool suite also includes a command line-oriented utility called *tshark*. This utility can be used to analyze traffic using only a command line approach, when running a GUI application

is not practical. The *tshark* tool has the same overall functionality as the GUI, but lacks its interactive capabilities.

Section 5.4 also includes step by step examples of using *tshark*, which is fully documented in the corresponding manual pages (*man tshark*). To get online help, use *tshark -?*.

5.3.3 Installing IPMI analyzer software

The IPMB trace daemon is included in the software image on the ShMM by default and does not require any specific installation.

The Wireshark client applications (and much more regarding Wireshark) can be obtained from www.wireshark.org.

The *netcat* (*nc*) utility is needed as a mediator between the IPMB trace daemon and Wireshark/tshark for real-time IPMI analysis (with IPMB trace daemon operating in controlled mode). Linux distributions typically include this utility. Its Windows variant (*nc.exe*) can be downloaded from <http://www.securityfocus.com/tools/139>. The source code of the utility is available at <http://netcat.sourceforge.net/download.php>.

5.4 Collecting and analyzing a trace using the GUI and command line tools

An IPMB trace can be collected in either of two ways: 1) unattended mode (where it is stored in a file during collection and analyzed later) or 2) controlled mode (where the trace is forwarded to directly to the analyzer as it is collected). The following sections provide guidance for each of these modes, as well as for collecting and analyzing an IPMI over network trace.

5.4.1 Collecting IPMI traces on IPMB-0 in unattended mode

In unattended mode, a trace is collected into a target-local file, transferred to the analyzer host, and then read and analyzed using the trace analyzer.

The following stepwise procedure shows how to collect an unattended trace and open it for analysis in the GUI. The GUI-specific steps are followed by corresponding guidance for opening the trace for analysis in *tshark*.

1. Start *ipmb_traced* on the ShMM with the option *-w*, optionally specifying the file to put the trace to:

```
# ipmb_traced -w -u /var/ipmb_traced.log &
```

If no *-u* option is provided, *ipmb_traced* puts the trace in the following file:

/var/ipmb_traced/ipmb_traced.log. In this case, the */var/ipmb_traced* directory must be created manually using the *mkdir* command before *ipmb_traced* is started:

```
# mkdir /var/ipmb_traced
```

2. Execute actions that create the IPMI traffic on IPMB-0 to be analyzed. For example, if insertion of a board generates the sequence of events and commands that need to be analyzed, do such a board insertion.
3. When trace collection is complete, stop *ipmb_traced* by sending the SIGINT signal to *ipmb_traced*:
 - Obtain the *ipmb_traced* process ID using the *ps* command:

```
# ps
```

```
...
```

```
39 ttyS0 root 0 S ipmb_traced -w -u /var/ipmb_traced.log
```

```
...
```

- Send the SIGINT signal to the *ipmb_traced* process:

```
# kill -2 39
```

4. Copy the trace file to the analyzer host:

```
# cd /var/
```

```
# ftp 172.16.1.24
```

```
Connected to 172.16.1.24.
```

```
220 ts FTP server (Version wu-2.6.2-5) ready.
```

```
...
```

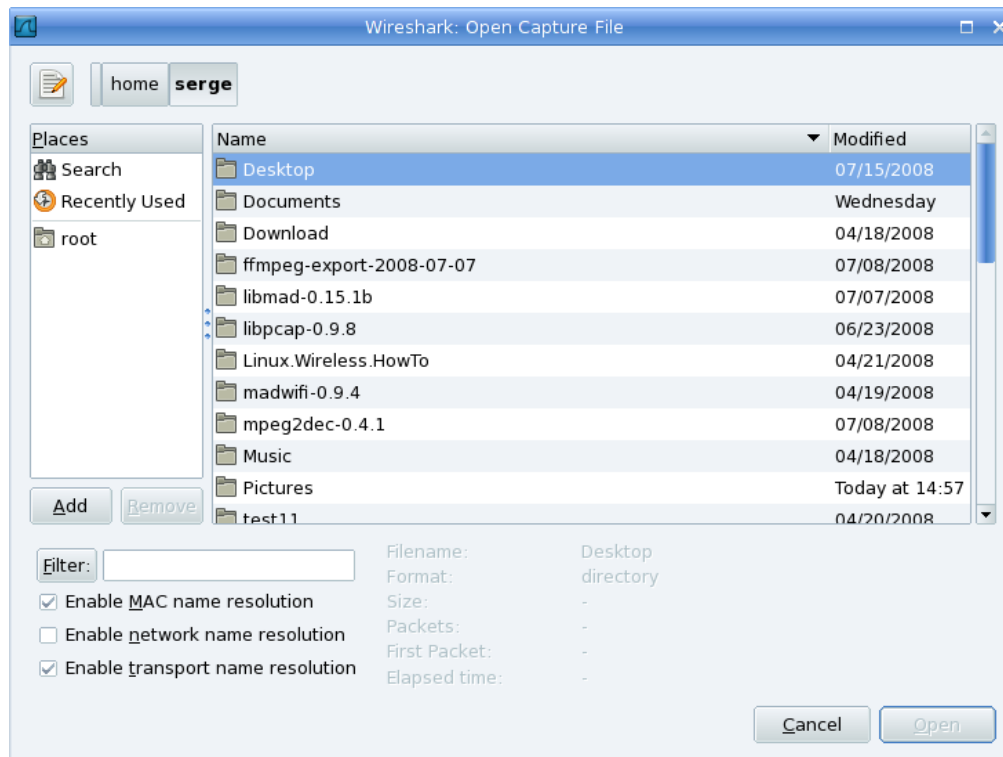
```
ftp> put ipmb_traced.log
```

5. On the analyzer host, start the GUI:

```
C:\> wireshark
```

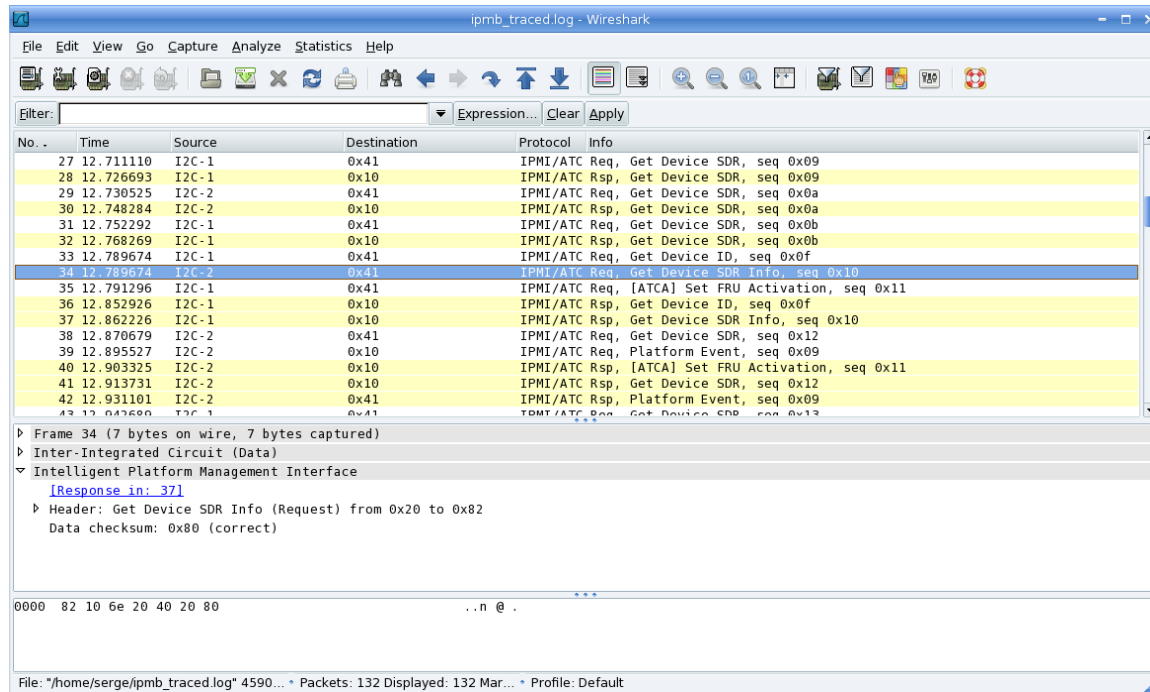
6. In the main menu, go to *File*, then *Open* and browse to the trace file. Open the trace file:

Figure 8: Open Capture File window



This will open the trace in the GUI main window:

Figure 9: GUI main window (unattended mode)



In a command-line context, at step 5, instead of starting the GUI utility, run the *tshark* utility (option *-r* specifies the trace data file name):

```
# tshark -r ipmb_traced.log
```

```
1 0.000000 I2C-1 -> 0x41 IPMI/ATCA Req, Get Sensor Reading,
seq 0x28
2 0.079507 I2C-1 -> 0x10 IPMI/ATCA Rsp, Get Sensor Reading,
seq 0x28
3 0.579642 I2C-2 -> 0x41 IPMI/ATCA Req, Get Sensor Reading,
seq 0x29
4 0.585695 I2C-2 -> 0x10 IPMI/ATCA Rsp, Get Sensor Reading,
seq 0x29
5 0.679566 I2C-1 -> 0x41 IPMI/ATCA Req, Get Sensor Reading,
seq 0x2a
6 0.686105 I2C-1 -> 0x10 IPMI/ATCA Rsp, Get Sensor Reading,
seq 0x2a
7 0.698566 I2C-2 -> 0x41 IPMI/ATCA Req, Get Sensor Reading,
seq 0x2b
8 0.713453 I2C-2 -> 0x10 IPMI/ATCA Rsp, Get Sensor Reading,
seq 0x2b
9 0.724908 I2C-1 -> 0x41 IPMI/ATCA Req, Get Sensor Reading,
seq 0x2c
10 0.739889 I2C-1 -> 0x10 IPMI/ATCA Rsp, Get Sensor Reading,
seq 0x2c
11 0.751980 I2C-2 -> 0x41 IPMI/ATCA Req, Get Sensor Reading,
seq 0x2d
12 0.767020 I2C-2 -> 0x10 IPMI/ATCA Rsp, Get Sensor Reading,
seq 0x2d
```

...

Use the `-V` option to obtain a detailed dump of the messages; use `-x` option to obtain a hexadecimal dump of messages, as well:

```
# tshark -r ipmb_traced.log -x -V
```

```
Frame 1 (8 bytes on wire, 8 bytes captured)
```

```
Arrival Time: Jun 24, 2008 20:58:03.238567000
```

```
[Time delta from previous captured frame: 0.000000000 seconds]
```

```
[Time delta from previous displayed frame: 0.000000000 seconds]
```

```
[Time since reference or first frame: 0.000000000 seconds]
```

```
Frame Number: 1
```

```
Frame Length: 8 bytes
```

```
Capture Length: 8 bytes
```

```
[Frame is marked: False]
```

```
[Protocols in frame: i2c:ipmi]
```

```
Inter-Integrated Circuit (Data)
```

```
Bus: I2C-1
```

```
Target address: 0x41
```

```
Flags: 0x00000000
```

```
Intelligent Platform Management Interface
```

```
[No corresponding response]
```

```
Header: Get Sensor Reading (Request) from 0x20 to 0x82
```

```
Target Address: 0x82
```

```
Target LUN: 0x00, NetFN: Sensor/Event Request (0x04)
```

```
.... ..00 = Target LUN: 0x00
```

```
0001 00.. = NetFn: Sensor/Event Request (0x04)
```

```
Header checksum: 0x6e (correct)
```

```
Source Address: 0x20
```

```
Source LUN: 0x00, SeqNo: 0x28
```

```
.... ..00 = Source LUN: 0x00
```

```
1010 00.. = Sequence Number: 0x28
```

```
Command: Get Sensor Reading (0x2d)
```

```
Data
```

```
Sensor Number: 2
```

```
Data checksum: 0x11 (correct)
```

```
0000 82 10 6e 20 a0 2d 02 11 ..n .-..
```

```
Frame 2 (11 bytes on wire, 11 bytes captured)
```

```
Arrival Time: Jun 24, 2008 20:58:03.318074000
```

```
[Time delta from previous captured frame: 0.079507000 seconds]
```

```
[Time delta from previous displayed frame: 0.079507000 seconds]
```

```
[Time since reference or first frame: 0.079507000 seconds]
```

```
Frame Number: 2
```

```
Frame Length: 11 bytes
```

```
Capture Length: 11 bytes
```

```

[Frame is marked: False]
[Protocols in frame: i2c:ipmi]
Inter-Integrated Circuit (Data)
  Bus: I2C-1
  Target address: 0x10
  Flags: 0x00000000
Intelligent Platform Management Interface
  [Response to: 1]
  [Responded in: 0.079507000 seconds]
  Header: Get Sensor Reading (Response) from 0x82 to 0x20
    Target Address: 0x20
    Target LUN: 0x00, NetFN: Sensor/Event Response (0x05)
      .... ..00 = Target LUN: 0x00
      0001 01.. = NetFn: Sensor/Event Response (0x05)
    Header checksum: 0xcc (correct)
    Source Address: 0x82
    Source LUN: 0x00, SeqNo: 0x28
      .... ..00 = Source LUN: 0x00
      1010 00.. = Sequence Number: 0x28
    Command: Get Sensor Reading (0x2d)
    Completion code: Command Completed Normally (0x00)
  Data
    Sensor Reading: 50
    Event Messages: Enabled, Sensor scanning: Enabled
      1... .... = Event Messages: Enabled
      .1.. .... = Sensor scanning: Enabled
      ..0. .... = Reading/status unavailable: False
    Threshold comparisons/assertions (byte 0)
      0... .... = Reserved / State 7 asserted: False
      .0.. .... = Reserved / State 6 asserted: False
      ..0. .... = At or above UNR threshold / State 5 asserted: False
      ...0 .... = At or above UC threshold / State 4 asserted: False
      .... 0... = At or above UNC threshold / State 3 asserted: False
      .... .0.. = At or below LNR threshold / State 2 asserted: False
      .... ..0. = At or below LC threshold / State 1 asserted: False
      .... ...0 = At or below LNC threshold / State 0 asserted: False
    Data checksum: 0xbf (correct)
0000 20 14 cc 82 a0 2d 00 32 c0 00 bf          ....-.2...
Frame 3 (8 bytes on wire, 8 bytes captured)
  Arrival Time: Jun 24, 2008 20:58:03.818209000
  [Time delta from previous captured frame: 0.500135000 seconds]
  [Time delta from previous displayed frame: 0.500135000 seconds]
  [Time since reference or first frame: 0.579642000 seconds]

```



```

Frame Number: 3
Frame Length: 8 bytes
Capture Length: 8 bytes
[Frame is marked: False]
[Protocols in frame: i2c:ipmi]
Inter-Integrated Circuit (Data)
  Bus: I2C-2
  Target address: 0x41
  Flags: 0x00000000
Intelligent Platform Management Interface
  [No corresponding response]
  Header: Get Sensor Reading (Request) from 0x20 to 0x82
    Target Address: 0x82
    Target LUN: 0x00, NetFN: Sensor/Event Request (0x04)
      .... ..00 = Target LUN: 0x00
      0001 00.. = NetFn: Sensor/Event Request (0x04)
    Header checksum: 0x6e (correct)
    Source Address: 0x20
    Source LUN: 0x00, SeqNo: 0x29
      .... ..00 = Source LUN: 0x00
      1010 01.. = Sequence Number: 0x29
    Command: Get Sensor Reading (0x2d)
  Data
    Sensor Number: 4
    Data checksum: 0x0b (correct)
0000 82 10 6e 20 a4 2d 04 0b                ..n .-..
...

```

5.4.2 Collecting IPMI traces on IPMB-0 in controlled mode

When the trace daemon is run in the controlled mode, the trace is transferred directly to the trace analyzer via a TCP/IP connection. The following stepwise procedure shows how to collect a controlled trace and open it for analysis in the GUI. The GUI-specific steps are followed by corresponding guidance for collecting a controlled trace in *tshark*.

1. On the ShMM, start *ipmb_traced* with the *-c* and *-w* flags:

```
# daemon -f ipmb_traced -c -w
```

On the analyzer host, use *netcat* to feed the data to Wireshark over the TCP/IP network (specify the IP address of the ShMM and the port on which the trace daemon listens for incoming connections, 19000 by default):

```
C:\> nc.exe 192.168.1.62 19000 | "c:\Program Files\Wireshark\wireshark.exe" -k -i -
```

The *-k* option instructs Wireshark to start capturing data immediately. The standard input is


specified as the interface that provides the trace data. To avoid repeatedly typing the long command, a Windows batch script can be created instead:

[rtwireshark.bat]

```
C:\distrib\nc.exe %1 19000 | "%ProgramFiles%\Wireshark\wireshark.exe" -k -i -
& exit
```


Given that this script is defined, invoke it from the command prompt or create a shortcut on the desktop and specify the IP address as an input parameter:

Figure 10: Capture Options window



2. Execute actions that create the IPMI traffic on IPMB-0 to be analyzed. For example, if insertion of a board generates the sequence of events and commands that need to be analyzed, do such a board insertion.
3. The GUI will update the main window showing the collected trace in real-time:

Figure 11: Collected trace (controlled mode)



4. To stop the capture session, go to the *Capture* submenu and select *Stop*.

In a command-line context, redirect the output of *netcat* to the *tshark* utility instead of the Wireshark GUI application. Like Wireshark, *tshark* can take the standard input option, *-*, to indicate the source of the data. Specify the IP address of the ShMM and the port on which the trace daemon listens for incoming connections (19000 by default) as input to *netcat*. Options *-x* and *-V*, as in the case of unattended mode, can be used to request a hexadecimal and/or detailed dumps of the messages. Any additional options needed for filtering and analysis of the message stream (see section 5.4.6) should also be specified at this time.

Next, execute the actions from step 5; the trace is then captured by *tshark* and sent to the command output. For example:

```
# nc 192.168.1.62 19000 | tshark -i -
```

Capturing on Standard input

```
1214325501.948358      I2C-1 -> 0x41      IPMI/ATCA Req, Get Sensor
Reading, seq 0x27
1214325501.999629      I2C-1 -> 0x10      IPMI/ATCA Rsp, Get Sensor
Reading, seq 0x27
1214325502.002796      I2C-2 -> 0x41      IPMI/ATCA Req, Get Sensor
Reading, seq 0x28
1214325502.017845      I2C-2 -> 0x10      IPMI/ATCA Rsp, Get Sensor
Reading, seq 0x28
1214325502.038082      I2C-1 -> 0x41      IPMI/ATCA Req, Get Sensor
Reading, seq 0x29
1214325502.038245      I2C-1 -> 0x10      IPMI/ATCA Rsp, Get Sensor
Reading, seq 0x29
1214325502.041232      I2C-2 -> 0x41      IPMI/ATCA Req, Get Sensor
Reading, seq 0x2a
```

...

To stop the capture session, terminate the *tshark* process (e.g. by pressing ^C).

5.4.3 Collecting an IPMI trace in board trace mode

If some board in the shelf supports an HPM.2 IPMI trace facility, the trace daemon can be used as a bridge between the board and the trace analyzer application. To support this mode, the target IPMI controller must be specifically configured to begin tracing the specified IPMI messaging interfaces when an RMCP+ session with the HPM.2 IPMI Trace Payload is established to a certain payload instance number and send IPMI trace packets in the appropriate format over the RMCP+ session while the session is open; when the session closes, the controller stops tracing IPMI messaging.

In this mode, the trace daemon, instead of collecting the IPMI trace locally on the Shelf Manager, establishes a RMCP+ session (with the HPM.2 IPMI Trace Payload) with the target board and forwards the data received over this session, to the trace analyzer (in controlled mode), or stores it in a file (in unattended mode).

For example, the following command line can be used to start the trace daemon in the board trace mode with a board that has the IP address 192.168.1.149. The specified channel mask instructs the IPMI controller to trace the IPMB-O and IPMB-L buses. The trace daemon operates in controlled mode.

```
# daemon -f ipmb_traced -c -w -b 192.168.1.149 -m 0x81
```

5.4.4 Collecting traces of IPMI traffic over the network

When collecting the traces of IPMI traffic transferred over the network RMCP or RMCP+ sessions, the trace daemon is not used. Wireshark or *tshark* collects the trace directly from one of the local network interfaces on the system on which that tool runs (that is, the analyzer host). Usually, the analyzer host is also the one that participates in the IPMI-oriented communication. However, if the analyzer host's local network interface supports promiscuous mode, Wireshark or *tshark* can collect a trace of IPMI communication between two separate systems that are not the analyzer host, if at least one of them is on the same LAN with the analyzer host.

The following stepwise procedure shows how to collect an IPMI-over-network trace and open it for analysis in the GUI. The GUI-specific steps are followed by corresponding guidance for collecting a similar trace in *tshark*.

1. On the analyzer host, start the GUI:

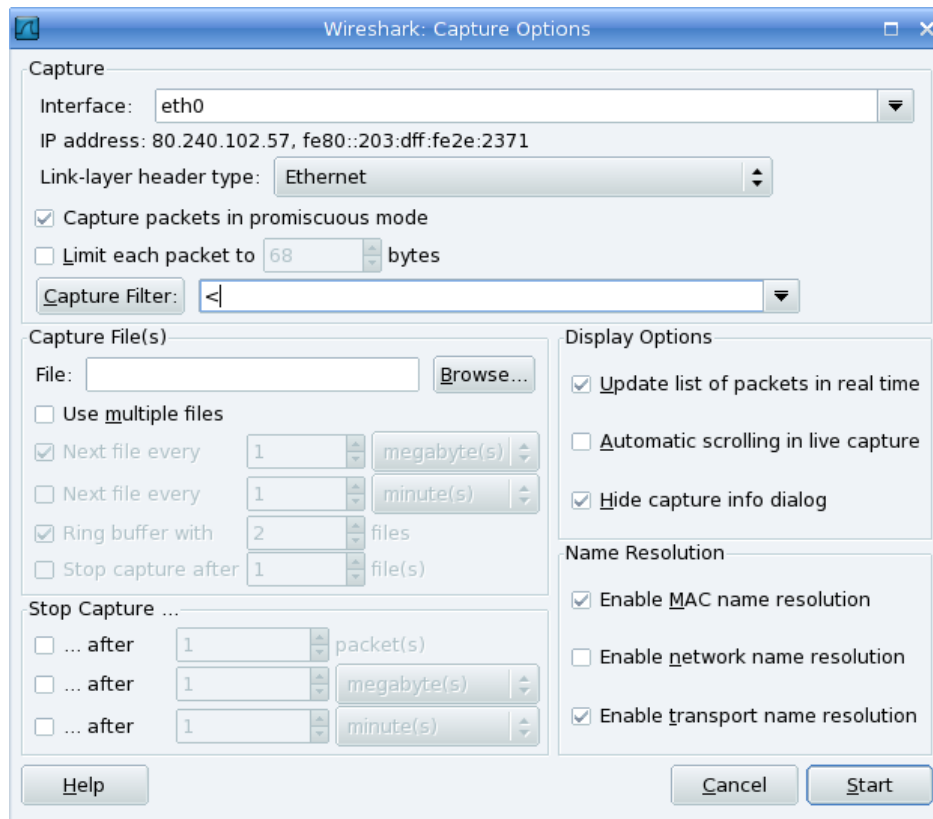
```
C:\> wireshark
```

2. In the main menu, go to *Capture* submenu and select *Options...*. This opens the *Capture Options* dialog. In the dialog, use the *Capture/Interface* window to specify the target. Choose the name of the network interface through which the IPMI communication takes place (e.g. *eth0*).

Additionally, you can use the *Update list of packets in real time* checkbox to specify that the GUI should show the trace in real-time.

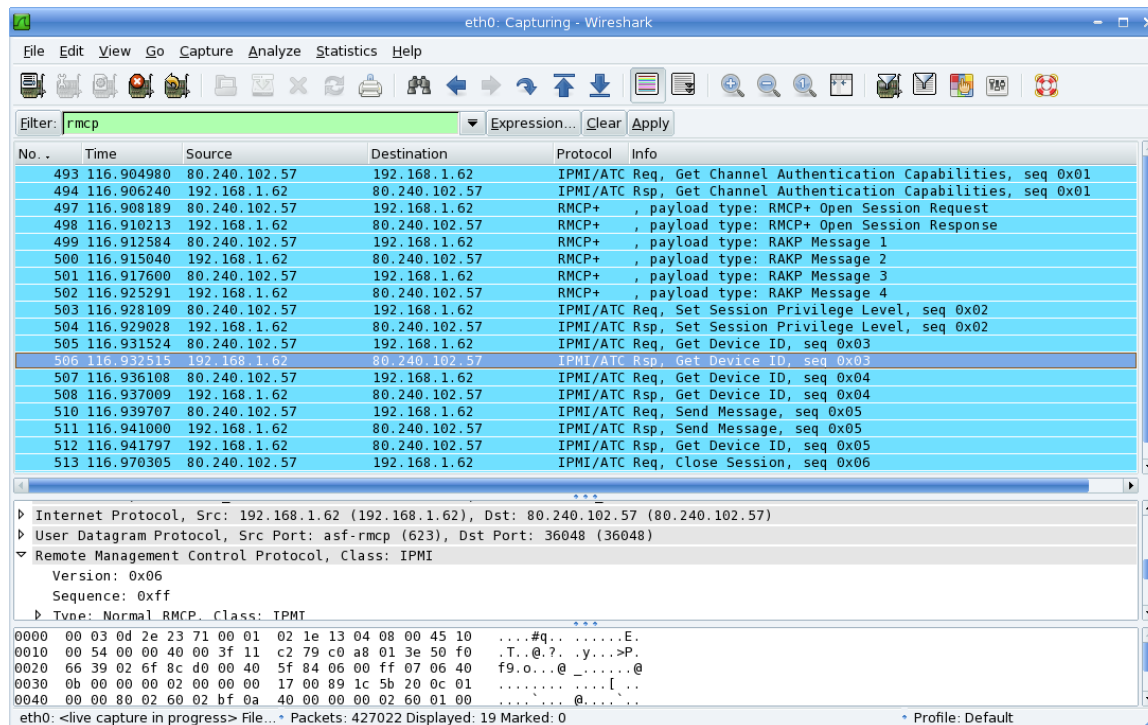
For instance:

Figure 12: Capture Options window (analyzing networked IPMI traffic)



3. Press the *Start* button to begin collecting the trace.
4. Set up the filter “rmcp” to include only RMCP-related traffic into the trace. Otherwise, all network traffic going through the specified adapter will be included into the trace. To do this, type *rmcp* in the Filter window in the top-left corner of the main screen and press the *Apply* button.
5. Execute actions that create the IPMI traffic on the network to be analyzed. For example, run a system management application that uses RMCP/RMCP+ to talk to the Shelf Manager.
6. The GUI will update the main window showing the collected trace in real-time:

Figure 13: Collected trace (networked IPMI traffic)



7. To stop the capture session, go to the *Capture* submenu and select *Stop*.

In a command-line context, instead of steps 1-3, run the *tshark* utility (specifying the target in the form of the network interface name (e.g. *eth0*) using option *-i*). Specify the filtering option *-R rmcp* to include only RMCP traffic in the trace. Options *-x* and *-V*, as in the case of unattended IPMB trace mode, can be used to request a hexadecimal and/or detailed dumps of the messages.

Any additional options needed for filtering and analysis of the message stream (see section 5.4.6) should also be specified at this time. Use the *&&* ("logical and") operator to combine additional filtering options with the filtering option *rmcp*.

Next, execute the actions from step 4; the trace is then captured by *tshark* and sent to the command output. For example:

```
# tshark -i eth0 -R rmcp
```

```
71.989290 80.240.102.57 -> 192.168.1.199 IPMI/ATCA Req, Get Channel
Authentication Capabilities, seq 0x0
71.990809 192.168.1.199 -> 80.240.102.57 IPMI/ATCA Rsp, Get Channel
Authentication Capabilities, seq 0x01
71.996590 80.240.102.57 -> 192.168.1.199 RMCP+ , payload type: RMCP+ Open
Session Reques
71.998471 192.168.1.199 -> 80.240.102.57 RMCP+ , payload type: RMCP+ Open
Session Respons
72.000184 80.240.102.57 -> 192.168.1.199 RMCP+ , payload type: RAKP Message 1
72.005487 192.168.1.199 -> 80.240.102.57 RMCP+ , payload type: RAKP Message
72.008539 80.240.102.57 -> 192.168.1.199 RMCP+ , payload type: RAKP Message 3
72.019149 192.168.1.199 -> 80.240.102.57 RMCP+ , payload type: RAKP Message 4
```

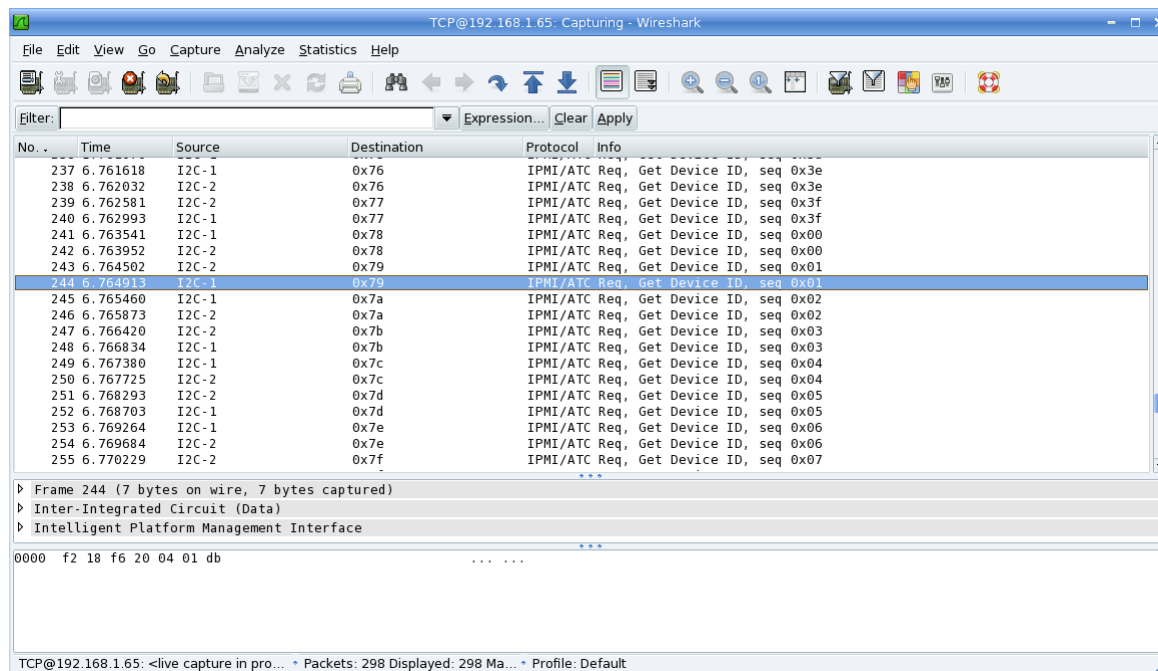
```
72.022416 80.240.102.57 -> 192.168.1.199 IPMI/ATCA Req, Set Session Privilege
Level, seq 0x02
72.033355 192.168.1.199 -> 80.240.102.57 IPMI/ATCA Rsp, Set Session Privilege
Level, seq 0x02
72.057400 80.240.102.57 -> 192.168.1.199 IPMI/ATCA Req, Get Device ID, seq 0x03
72.059134 192.168.1.199 -> 80.240.102.57 IPMI/ATCA Rsp, Get Device ID, seq 0x03
72.059568 80.240.102.57 -> 192.168.1.199 IPMI/ATCA Req, Get Device ID, seq 0x04
72.062264 192.168.1.199 -> 80.240.102.57 IPMI/ATCA Rsp, Get Device ID, seq 0x04
72.062680 80.240.102.57 -> 192.168.1.199 IPMI/ATCA Req, Send Message, seq 0x05
72.066428 192.168.1.199 -> 80.240.102.57 IPMI/ATCA Rsp, Send Message, seq 0x05
72.068235 192.168.1.199 -> 80.240.102.57 IPMI/ATCA Rsp, Get Device ID, seq 0x05
72.101030 80.240.102.57 -> 192.168.1.199 IPMI/ATCA Req, Close Session, seq 0x06
72.102654 192.168.1.199 -> 80.240.102.57 IPMI/ATCA Rsp, Close Session, seq 0x06
```

To stop the capture session, terminate the *tshark* process (e.g. by pressing `^C`).

5.4.5 Analyzing a trace using the GUI

To select a message in the Wireshark GUI main window, simply scroll to the message that interests you in the upper frame of the main window:


Figure 14: Message 244 is selected



After a message has been selected, its raw contents are shown in the bottom frame of the main window, both in hex and ASCII interpretations. The middle frame of the main window provides protocol interpretations of the message data.

By clicking on the right arrow/down arrow (or `+/-`) GUI elements, you can choose a particular protocol layer you wish to analyze:

Figure 15: IPMI protocol layers (IPMB trace)



A trace can be filtered using the filter expressions. Specifically, for IPMI trace filtering, the following expression primitives are available:

Table 24: IPMI filter primitives

FILTER PRIMITIVE NAME	DESCRIPTION
<i>ipmi.header.broadcast</i>	Broadcast destination address (only for Broadcast “Get Device ID” Request messages)
<i>ipmi.header.target</i>	Destination device (responder or requester) slave address
<i>ipmi.header.netfn</i>	Network Function (NetFn) of the message
<i>ipmi.header.trg_lun</i>	Destination (responder or requester) Logical Unit Number (LUN)
<i>ipmi.header.crc</i>	Header checksum
<i>ipmi.header.source</i>	Source device (requester or responder) slave address
<i>ipmi.header.sequence</i>	Request sequence
<i>ipmi.header.src_lun</i>	Source (requester or responder) LUN
<i>ipmi.header.command</i>	Command to execute
<i>ipmi.data.completion</i>	Completion code (only for response messages)
<i>ipmi.data.crc</i>	Data checksum


The filtering expression primitives above are applicable for both IPMB traffic and RMCP/RMCP+ IPMI traffic over the network. In the latter case, several RMCP filtering expression primitives are also available, but these are not widely used.

To specify a filter, use the *Filter* text entry box in the upper left corner of the main window. For instance, to select the messages with the destination address 0x20 and the source address 0x82, type the string:

```
ipmi.header.target == 0x20 && ipmi.header.source == 0x82
```

in the *Filter* text entry box:

Figure 16: Filtered trace



Note: Since the broadcast destination address is always zero, the `ipmi.header.broadcast` filter expression should be used by specifying the following filter string in the `Filter` entry of the main window:

```
ipmi.header.broadcast == 0x00
```

5.4.6 Analyzing a trace using command-line tools

GUI tools are much more powerful for analyzing a trace than command-line tools, so you should use GUI tools for trace analysis, if possible. When using command-line tools, you should first capture the data into a file in unattended mode and then analyze the captured data file using *tshark*. In this usage style, *tshark* can be run multiple times over the same data, thus allowing some degree of interactivity in the analysis. In controlled mode, *tshark* is run only once over the captured data; as a result, all applicable options must be specified in the command line at startup.

Another limitation of *tshark*, compared to Wireshark, is that *tshark* can only provide back references for the packets (the links *Response to* and *Response in* that allow the user to navigate across packets quickly). This limitation results from the fact that *tshark* is optimized for speed and thus calls the packet dissector only once per packet. The verbose *tshark* capture in section 5.4.1 shows an example of this limitation. The description of the first packet indicates *No corresponding message*, while the description of the second packet shows *Response to: 1*.

To filter messages in *tshark* (both in controlled and unattended mode), use the *-R* option, followed by the text of the filter expression (in quotes). For example, to select and dump in details the messages with the destination address 0x20 and the source address 0x82 (as in the previous section), when reading from a trace file, use the following command line:

```
# tshark -r ipmb_traced.log -V -R "ipmi.header.target==0x20 &&
ipmi.header.source==0x82"
```

The *-T* option can be used to specify the format of the output. The following formats are supported: plain text, PostScript, PDML (Packet Details Markup Language, based on XML), PSML (Packet Summary Markup Language, based on XML) and so called “field format”. For example, the following command line can be used to dump message details in the PostScript format:

```
# tshark -r ipmb_traced.log -V -T ps
```

The *-e* option can be used to select specific packet fields to be dumped; if this option is used, the “field format” must be also specified (*-T fields*). Several *-e* options can be used to specify multiple fields. In conjunction with the filtering option *-R*, this option provides a powerful tool for queries on the message stream, similar to relational database queries. For example, the following query dumps the values of net functions and command codes for IPMB messages sent from the source address 0x84 to the Shelf Manager (destination address 0x20):

```
# tshark -r ipmb_traced.log -e ipmi.header.netfn -e ipmi.header.command
"ipmi.header.target==0x84 && ipmi.header.source==0x20" -T fields
```

```
0x04      0x2d
0x04      0x2d
0x04      0x2d
0x04      0x2d
0x04      0x2d
0x04      0x2d
0x04      0x2d
0x04      0x2d
0x04      0x2d
0x04      0x2d
0x06      0x01
0x04      0x2d
0x04      0x20
0x04      0x2d
0x04      0x2d
0x04      0x2d
```

```
0x04    0x2d
0x04    0x2d
0x04    0x2d
0x04    0x2d
0x04    0x2d
0x04    0x2d
```

Full information about *tshark* options and syntax can be found in the *tshark* manual page (*man tshark*).

6 Re-initializing the ShMM

6.1 In this section

This section applies to the ShMM-500, ShMM-1500 and ShMM-700 and contains the topics listed below. Just click on a topic to go to it.

- Re-initializing the U-Boot environment
- Re-initializing the file system
- Resetting the login password

6.2 Re-initializing the U-Boot environment

The U-Boot environment variables are stored in the ShMM EEPROM on the ShMM-500 and ShMM-1500, and in two redundant dedicated Flash partitions on the ShMM-700. If you would like to restore the default values for the U-Boot environment variables, you must first erase the environment variables stored in EEPROM/Flash and reset (or power cycle) the ShMM.

To erase the EEPROM on a ShMM-500, you need to enter the following command from the U-Boot prompt:

```
shmm500 eeprom write 80400000 0 1800
EEPROM @0x50 write: addr 80400000 off 0000 count 6144 ... done
shmm500
```

To erase the EEPROM on a ShMM-1500, you need to enter the following command from the U-Boot prompt:

```
shmm1500 eeprom write 80400000 0 1000
EEPROM @0x50 write: addr 80400000 off 0000 count 4096 ... done
shmm1500
```

On a ShMM-700, the U-Boot environment variables are stored in two redundant dedicated Flash partitions. To erase the U-Boot environment variables on a ShMM-700, you need to enter the following commands from the U-Boot prompt:

```
shmm700 sf probe 2:0
65536 KiB N25Q512A at 2:0 is now current device
shmm700 sf erase 100000 80000
Erasing Flash... done!
shmm700
```

Then you need to reset the ShMM using the *reset* command, as shown below with resulting console output (for the ShMM-500; corresponding output for the ShMM-1500 and ShMM-700 is slightly different).

```
shmm500 reset
U-Boot 1.1.2 (Apr 27 2005 - 19:17:09)
CPU: Au1550 324 MHz, id: 0x02, rev: 0x00
Board: ShMM-500
S/N: 8000041
DRAM: 64 MB
```

```
Flash: 16 MB
*** Warning - bad CRC, using default environment
In:      serial
Out:     serial
Err:     serial
Net:     AulX00 ETHERNET
Hit any key to stop autoboot:  0
shmm500
```

Then save these environment settings. Since U-Boot configuration storage contains 2 copies of environment variables data and the *saveenv* command changes them in round-robin order, you need to invoke the *saveenv* command twice to update both copies.

```
shmmxx00 saveenv; saveenv; reset
```

6.3 Re-initializing the file system

The filesystem (*/var* and */etc*) is stored within the Flash and can be restored quite easily from the filesystem in the current RFS image. U-Boot has an environment variable called *flash_reset*. By setting this variable to *y* and then booting up the system, the */var* and */etc* files are erased and then restored from the filesystem in the current RFS image.

```
shmmxx00 setenv flash_reset y
shmmxx00 saveenv
shmmxx00 boot
```

The *flash_reset* variable is automatically set to *n* at system startup after re-initializing the Flash. The boot command (the second command above) begins booting the Linux kernel. It is during this process that the file system is re-initialized. The following output is shown on the console.

```
/etc/rc: Mounted /dev/pts
/etc/rc: Flash erase requested via U-BOOT var
/etc/rc: erasing mtdchar1 -> /etc
Erased 1024 Kibyte @ 0 -- 100% complete.
/etc/rc: erasing mtdchar0 -> /var
Erased 1536 Kibyte @ 0 -- 100% complete.
/etc/rc: Mounted /dev/mtdblock3 to /var
/etc/rc: /var/log mounted as FLASH disk
/etc/rc: Started syslogd and klogd
/etc/rc: /var/tmp mounted as RAM disk
/etc/rc: hostname demo
/etc/rc: /dev/mtdblock2 appears to be empty ... restoring from RFS /etc...
```

For the ShMM-700, the console output will look differently:

```
/etc/rc: Mount ubi0:user
/etc/rc: Flash erase requested via U-Boot variable
/etc/rc: erasing /etc
/etc/rc: erasing /var
```

```
/etc/rc: Selecting carrier from the environment:
/etc/rc:
/etc/rc: Extracting shelfman rootfs patch from /1/sentry.shmm700.app
/etc/rc: Extraction result : 0
/etc/rc: Placed /var/tmp to ram disk
/etc/rc: Setting hostname shmm~121
/etc/rc: Started syslogd and klogd
/etc/rc: Strobing the reliable upgrade WDT
/etc/rc: /etc appears to be empty ... restoring from RFS /etc...
```

6.4 Resetting the login password

The default login for the ShMM, as set in the factory-provided RFS images, is a user ID of “root” and the empty password “”. A shelf supplier can change the default password. We highly encourage users to change the password when configuring the Shelf Manager. In the event that the new password is forgotten, the password can be reset to the default value, as assigned in the RFS image, via the `password_reset` U-Boot variable. By setting this variable to `y` and then booting up the system, the root password is removed.

```
shmmxx00 setenv password_reset y
shmmxx00 saveenv
shmmxx00 boot
```

The following output is shown on the console, during boot up.

```
/etc/rc: hostname demo
```

7 Re-programming the ShMM-500/1500

This section describes how to update the firmware on the ShMM-500, or the ShMM-1500, using the reliable upgrade facilities that are built into Monterey Linux.

On the ShMM-700, a similar facility exists for reliable upgrade of ShMM firmware, but it is substantially different in implementation details and therefore is described in a separate section of this document (see section 8).

7.1 In this section

This section contains the topics listed below. Just click on a topic to go to it.

- Firmware reliable upgrade procedure overview
- Flash partitioning
- The `/var/upgrade` file system
- Reliable upgrade procedure status file
- Reliable upgrade utility
- Reliable upgrade utility use scenarios
- Reliable upgrade examples

7.2 Firmware reliable upgrade procedure overview

Monterey Linux provides a reliable upgrade procedure for the firmware images on a running and functioning ShMM. The procedure supports upgrade of the U-Boot firmware, the Linux kernel and the Linux root file system (or an arbitrary combination of these three images). If a software upgrade attempt fails (for instance, due to installation of a faulty U-Boot firmware image that is not capable of booting the ShMM or a Shelf Manager that can't start) the reliable upgrade procedure automatically falls back to the previous version of the firmware in persistent Flash.

Flash storage on either the ShMM-500 or the ShMM-1500 is divided into two areas. When a stable set of firmware is established in one of these areas, it is designated the persistent area. When new firmware is installed, it goes in the other area, which is initially designated provisional. Once a new set of firmware in the provisional area is validated, that area is designated the persistent area and continues in use until a future upgrade cycle starts the process over.

The reliable upgrade hardware mechanisms ensure that no matter what is installed to the provisional Flash, the ShMM always manages to boot from a software copy that is either fully functional or sufficiently sane to determine that there has been a failure in the upgrade session and consequently take appropriate corrective actions to revert to the safe software copy in persistent Flash.

At a higher level, the reliable upgrade hardware mechanisms are assisted by a software protocol based on logging of the status of the upgrade session to a non-volatile file in `/var/upgrade/status`. The software protocol ensures that the reliable upgrade does not finish until all the required actions, including those defined by custom "hook" scripts that may be needed for a specific application, have all completed successfully.

The Monterey Linux reliable upgrade procedure is described in full detail in Chapter 7 of the Monterey Linux User’s Guide: Au1550 Edition (for the ShMM-500) or MPC83xx Edition (for the ShMM-1500). Users of the reliable upgrade procedure should review that material as background. That material explains the underlying mechanisms, which provide substantial opportunities for customization to particular requirements. Any users who wish to extend or customize the workings of the reliable upgrade process must become thoroughly familiar with the Monterey Linux User’s Guide’s coverage of this topic.

The reliable upgrade procedure implemented in Monterey Linux is neutral to the embedded application running on the ShMM. The procedure provides a sufficient set of “hooks” allowing a specific application running on the ShMM to ensure that custom actions are carried out at appropriate points of the reliable upgrade. The remainder of this section focuses on provisions for reliable upgrade of the Pigeon Point Shelf Manager firmware that have been implemented using these hooks.

7.3 Flash partitioning

Both the ShMM-500 and the ShMM-1500 provide a hardware mechanism that allows swapping of the lower and upper halves of the Flash in the system memory map under control of software running on the CPU. This capability is implemented in support of the reliable upgrade procedure for software images in Flash. The reliable software upgrade procedure assumes that the Flash device contains two copies of the software, located in the lower and upper halves of Flash. All ShMMs are shipped with this partitioning, in which the Flash device is divided into two equal parts, each dedicated to holding one copy of the ShMM software.

This section assumes that the ShMM is configured to support reliable upgrade, including the two separate Flash regions. The tables below provide summaries of the Flash partitions maintained on the ShMM in this configuration. The exact layout of Flash partitions provided by the FOSL depends on the type of the ShMM, and the size of the Flash device installed.

The following tables provide a summary of the Flash partitions maintained by Linux for each of the ShMM-500 variants.

Table 25: Flash partitioning for 16MB Flash ShMM-500s

OFFSET IN FLASH (IN MBYTES)	SIZE (IN MBYTES)	DEVICE NODE	MOUNTED AS (ON STARTUP)	CONTENT
0	0.5	/dev/mtdchar1 0, /dev/mtdblock1 0	/var/upgrade	The second half of the /var/upgrade JFFS2 file system
0.5	1.5	/dev/mtdchar5, /dev/mtdblock5	Not mounted	The “other” /var JFFS2 file system
2	1	/dev/mtdchar6, /dev/mtdblock6	Not mounted	The “other” /etc JFFS2 file system
3	1	/dev/mtdchar7	Not mounted	The “other” Linux kernel image

OFFSET IN FLASH (IN MBYTES)	SIZE (IN MBYTES)	DEVICE NODE	MOUNTED AS (ON STARTUP)	CONTENT
4	0.25	/dev/mtdchar8	Not mounted	The “other” U-Boot firmware image
4.25	3.75	/dev/mtdchar9	Not mounted	The “other” Linux root file system (rfs) image
8	0.5	/dev/mtdchar10, /dev/mtdblock10	/var/upgrade	The first half of the /var/upgrade JFFS2 file system
8.5	1.5	/dev/mtdchar0, /dev/mtdblock0	/var	The /var JFFS2 file system
10	1	/dev/mtdchar1, /dev/mtdblock1	/etc	The /etc JFFS2 file system
11	1	/dev/mtdchar2	Not mounted	The Linux kernel image
12	0.25	/dev/mtdchar3	Not mounted	The U-Boot firmware image
12.25	3.75	/dev/mtdchar4	Not mounted	The Linux root file system (rfs) image

Table 26: Flash partitioning for 32MB Flash ShMM-500s

OFFSET IN FLASH (IN MBYTES)	SIZE (IN MBYTES)	DEVICE NODE	MOUNTED AS (ON STARTUP)	CONTENT
0	0.5	/dev/mtdchar10, /dev/mtdblock10	/var/upgrade	The first half of the /var/upgrade JFFS2 filesystem
0.5	1	/dev/mtdchar2, /dev/mtdblock2	Not mounted	The Linux kernel image
1.5	1	/dev/mtdchar1, /dev/mtdblock1	/etc	The /etc JFFS2 filesystem
2.5	1.75	/dev/mtdchar0, /dev/mtdblock0	/var	The /var JFFS2 filesystem
4.25	7.75	/dev/mtdchar4, /dev/mtdblock4	Not mounted	The Linux root file system (rfs) image
12	0.25	/dev/mtdchar3	Not mounted	The U-Boot firmware image

OFFSET IN FLASH (IN MBYTES)	SIZE (IN MBYTES)	DEVICE NODE	MOUNTED AS (ON STARTUP)	CONTENT
12.25	3.75	/dev/mtdchar1 1, /dev/mtdblock1 1	Not mounted	The first half of the app_jffs application-specific JFFS2 partition
16	0.5	/dev/mtdchar1 0, /dev/mtdblock1 0	/var/upgrade	The second half of the /var/upgrade JFFS2 file system
16.5	1	/dev/mtdchar7, /dev/mtdblock7	Not mounted	The “other” Linux kernel image
17.5	1	/dev/mtdchar6, /dev/mtdblock6	Not mounted	The “other” /etc JFFS2 file system
18.5	1.75	/dev/mtdchar5, /dev/mtdblock5	Not mounted	The “other” /var JFFS2 file system
20.25	7.75	/dev/mtdchar9, /dev/mtdblock9	Not mounted	The “other” Linux root file system (rfs) image
28	0.25	/dev/mtdchar8	Not mounted	The “other” U-Boot firmware image
28.25	3.75	/dev/mtdchar1 1, /dev/mtdblock1 1	Not mounted	The second half of the app_jffs application-specific JFFS2 partition

Table 27: Flash partitioning for 64MB Flash ShMM-500s

OFFSET IN FLASH (IN MBYTES)	SIZE (IN MBYTES)	DEVICE NODE	MOUNTED AS (ON STARTUP)	CONTENT
0	0.5	/dev/mtdchar1 0, /dev/mtdblock1 0	/var/upgrade	The second half of the /var/upgrade JFFS2 file system
0.5	1	/dev/mtdchar7	Not mounted	The “other” Linux kernel image
1.5	1	/dev/mtdchar6, /dev/mtdblock6	Not mounted	The “other” /etc JFFS2 file system
2.5	1.75	/dev/mtdchar5, /dev/mtdblock5	Not mounted	The “other” /var JFFS2 file system

OFFSET IN FLASH (IN MBYTES)	SIZE (IN MBYTES)	DEVICE NODE	MOUNTED AS (ON STARTUP)	CONTENT
4.25	15.75	/dev/mtdchar9	Not mounted	The “other” Linux root file system (rfs) image
20	8	/dev/mtdchar12, /dev/mtdblock12	Not mounted	The second half of the app1_jffs application-specific JFFS2 partition
28	0.25	/dev/mtdchar8	Not mounted	The “other” U-Boot firmware image
28.25	3.75	/dev/mtdchar11, /dev/mtdblock11	Not mounted	The second half of the app_jffs application-specific JFFS2 partition
32	0.5	/dev/mtdchar10, /dev/mtdblock10	/var/upgrade	The first half of the /var/upgrade JFFS2 file system
32.5	1	/dev/mtdchar2	Not mounted	The Linux kernel image
33.5	1	/dev/mtdchar1, /dev/mtdblock1	/etc	The /etc JFFS2 file system
34.5	1.75	/dev/mtdchar0, /dev/mtdblock0	/var	The /var JFFS2 file system
36.25	15.75	/dev/mtdchar4	Not mounted	The Linux root file system (rfs) image
52	8	/dev/mtdchar12, /dev/mtdblock12	Not mounted	The first half of the app1_jffs application-specific JFFS2 partition
60	0.25	/dev/mtdchar3	Not mounted	The U-Boot firmware image
60.25	3.75	/dev/mtdchar11, /dev/mtdblock11	Not mounted	The second half of the app_jffs application-specific JFFS2 partition

The following tables provide a summary of the Flash partitions maintained by Linux for each of the ShMM-1500 variants.

Table 28: Flash partitioning for 32MB Flash ShMM-1500s

OFFSET IN FLASH (IN MBYTES)	SIZE (IN MBYTES)	DEVICE NODE	MOUNTED AS (ON STARTUP)	CONTENT
0	0.25	/dev/mtdchar3	Not mounted	The U-Boot firmware image
0.25	0.5	/dev/mtdchar1 0, /dev/mtdblock1 0	/var/upgrade	The first half of the /var/upgrade JFFS2 filesystem
0.75	1	/dev/mtdchar2, /dev/mtdblock2	Not mounted	The Linux kernel image
1.75	1	/dev/mtdchar1, /dev/mtdblock1	/etc	The /etc JFFS2 filesystem
2.75	1.75	/dev/mtdchar0, /dev/mtdblock0	/var	The /var JFFS2 filesystem
4.5	7.75	/dev/mtdchar4, /dev/mtdblock4	Not mounted	The Linux root file system (rfs) image
12.25	3.75	/dev/mtdchar1 1, /dev/mtdblock1 1	Not mounted	The first half of the app_jffs application-specific JFFS2 partition
16	0.25	/dev/mtdchar8	Not mounted	The “other” U-Boot firmware image
16.25	0.5	/dev/mtdchar1 0, /dev/mtdblock1 0	/var/upgrade	The second half of the /var/upgrade JFFS2 file system
16.75	1	/dev/mtdchar7, /dev/mtdblock7	Not mounted	The “other” Linux kernel image
17.75	1	/dev/mtdchar6, /dev/mtdblock6	Not mounted	The “other” /etc JFFS2 file system
18.75	1.75	/dev/mtdchar5, /dev/mtdblock5	Not mounted	The “other” /var JFFS2 file system
20.5	7.75	/dev/mtdchar9, /dev/mtdblock9	Not mounted	The “other” Linux root file system (rfs) image
28.25	3.75	/dev/mtdchar1 1, /dev/mtdblock1 1	Not mounted	The second half of the app_jffs application-specific JFFS2 partition

Table 29: Flash partitioning for 64MB Flash ShMM-1500s

OFFSET IN FLASH (IN MBYTES)	SIZE (IN MBYTES)	DEVICE NODE	MOUNTED AS (ON STARTUP)	CONTENT
0	0.25	/dev/mtdchar3	Not mounted	The U-Boot firmware image
0.25	0.5	/dev/mtdchar1 0, /dev/mtdblock1 0	/var/upgrade	The first half of the /var/upgrade JFFS2 filesystem
0.75	1	/dev/mtdchar2, /dev/mtdblock2	Not mounted	The Linux kernel image
1.75	1	/dev/mtdchar1, /dev/mtdblock1	/etc	The /etc JFFS2 filesystem
2.75	2	/dev/mtdchar0, /dev/mtdblock0	/var	The /var JFFS2 filesystem
4.75	16	/dev/mtdchar4, /dev/mtdblock4	Not mounted	The Linux root file system (rfs) image
20.75	11.25	/dev/mtdchar1 1, /dev/mtdblock1 1	Not mounted	The first half of the app_jffs application-specific JFFS2 partition
32	0.25	/dev/mtdchar8	Not mounted	The “other” U-Boot firmware image
32.25	0.5	/dev/mtdchar1 0, /dev/mtdblock1 0	/var/upgrade	The second half of the /var/upgrade JFFS2 file system
32.75	1	/dev/mtdchar7, /dev/mtdblock7	Not mounted	The “other” Linux kernel image
33.75	1	/dev/mtdchar6, /dev/mtdblock6	Not mounted	The “other” /etc JFFS2 file system
34.75	2	/dev/mtdchar5, /dev/mtdblock5	Not mounted	The “other” /var JFFS2 file system
36.75	16	/dev/mtdchar9, /dev/mtdblock9	Not mounted	The “other” Linux root file system (rfs) image
52.75	11.25	/dev/mtdchar1 1, /dev/mtdblock1 1	Not mounted	The second half of the app_jffs application-specific JFFS2 partition

7.4 The `/var/upgrade` file system

Monterey Linux mounts a 1 MByte partition as a JFFS2 file system at `/var/upgrade`. This file system hosts the reliable upgrade procedure status file.

It is important to note that the `/var/upgrade` JFFS2 partition is composed of two non-contiguous Flash blocks (0.5 MBytes each) residing in both the lower and upper halves of the Flash device. Monterey Linux takes advantage of the ability of the Linux MTD and JFFS2 layers to support a file system in non-contiguous Flash sectors in order to implement `/var/upgrade` this way.

Another feature of the JFFS2 file system that makes `/var/upgrade` work for purposes of the reliable upgrade procedure is that the JFFS2 internal structures do not create any dependencies (such as linked lists) based on Flash sector numbers or absolute offsets in Flash. Instead, when mounting a file system on a partition, the JFFS2 scans all the Flash sectors comprising the partition and recreates the logical content of a file system in an internal in-RAM representation. This feature ensures that regardless of which half of the Flash the ShMM-500 has booted from, Linux is able to mount `/var/upgrade` as a JFFS2 file system and make use of the previous content of the file system.

7.5 Reliable upgrade procedure status file

The software reliable upgrade procedure maintains the status of the most recent upgrade procedure session in the file `/var/upgrade/status` residing in a dedicated file system (`/var/upgrade`), which is mounted by Linux regardless of which Flash the ShMM has booted from. If the file exists, it contains the status of an upgrade procedure session that either is in progress presently or has recently completed.

`/var/upgrade/status` is an ASCII format file that contains one or more new line-terminated records, each describing the status of a particular step in the upgrade procedure. The format of a record line is as follows:

```
<step>: <status>
```

where `<step>` is an integer ranging from 1 to 14 (with Step 14 corresponding to a completed upgrade session) and `status` is a human-readable string describing status of the current step of the upgrade procedure session. Refer to the *ML User Guide* for a list and explanation of these steps.

The status file is used by the reliable upgrade utility to maintain a software protocol atop the reliable upgrade procedure hardware mechanisms to reliably determine the status of the upgrade procedure and proceed as appropriate.

7.6 Reliable upgrade utility

A special user-space utility is provided that is used for carrying out the reliable upgrade procedure as well as checking the status of the most recent upgrade.

The utility can be called only from the superuser (root) account. Any attempt to run the utility from a non-superuser account is rejected.

If called with any of the `-s`, `-c`, or `-f` options, the utility is being used to carry out the reliable upgrade procedure. While in the upgrade procedure, the utility logs to `/var/upgrade/status` the status of each action it performs as it proceeds through the steps of the upgrade procedure. If

the utility detects a failure, the reliable upgrade procedure is terminated by adding to `/var/upgrade/status` a record indicating an unsuccessful completion of the upgrade procedure and exiting with an appropriate error code.

If the utility is called with the `-C` (“capital C”) option, it copies the images to Flash partitions according to other command-line parameters, but does not perform the final upgrade step of switching active partitions and rebooting. This final step can be done later using the command `rupgrade_tool -s -R`.

The utility prints any informational messages to `stdout`. Providing a `-v` flag to any option that supports it increases the verbosity of the informational messages. The utility prints any error messages to `stderr`.

The utility has the following synopsis:

```
rupgrade_tool -s {--<dst>=<src>}... [--proto=<protocol>] [-d]
[--hook=<args>] [-v] [-a] [-R]
rupgrade_tool -c [-v]
rupgrade_tool -f [--hook=<args>] [-v]
rupgrade_tool -w [-f]
rupgrade_tool -S [-v]
rupgrade_tool -h
rupgrade_tool -C {--<dst>=<src>}... [--proto=<protocol>] [-d] [-v]
[-a]
```

where the parameters are defined as follows.

```
-s {--<dst>=<src>}... [--proto=<protocol>] [--hook=<args>] [-v]
[-a] [-R]
```

Initiate the reliable upgrade procedure. As delivered with Shelf Manager support, this step includes the following actions:

- obtaining the images to copy: locally or via the network;
- copying the images to the provisional Flash;
- terminating the Shelf Manager instance running on the ShMM-500, if any;
- copying non-volatile data to the provisional Flash;
- resetting the ShMM and instructing it to boot from the provisional Flash.

However, if the specifier `-R` has been used in the command line, only the last action in the list above occurs (that is, no image copying takes place).

Because of the last step, an invocation of `rupgrade_tool -s` typically does not return and instead resets the ShMM. If `rupgrade_tool -s` does return, it indicates that the reliable upgrade procedure has failed and was terminated before proceeding to reset the ShMM in order to boot from the provisional Flash.

Before the first step of the upgrade procedure is initiated by the utility, it removes the `/var/upgrade/status` file. In other words, the status of the previous upgrade procedure

session (if any) is lost and overwritten by the status of the new upgrade procedure session as soon as `rupgrade_tool -s` is called.

There can be one or more `--<dst>=<src>` specifiers in a call to `rupgrade_tool -s`. Each such specifier defines the name of a to-be-installed upgrade image file and where the file is to be installed in the Flash of the ShMM.

`<dst>` defines the destination of a newly installed upgrade image and can be any of the following:

- `u` - Upgrade the U-Boot image in the provisional U-Boot firmware image partition (/dev/mtdchar3).
- `k` - Upgrade the Linux kernel image in the provisional Linux kernel image partition (/dev/mtdchar2).
- `r` - Upgrade the root file system image in the provisional root file system image partition (/dev/mtdchar4).

`<src>` specifies an upgrade image file to be copied to the provisional Flash partition specified by `<dst>`.

The specifier `-a`, if used, indicates that all three components (U-Boot, kernel and root file system) are to be upgraded and the source upgrade image files must have standard names. These standard names are defined separately for ShMM-500 and ShMM-1500, according to the following table:

Table 30: Standard source file names for upgrade components

COMPONENT NAME	SOURCE FILE NAME (ShMM-500)	SOURCE FILE NAME (ShMM-1500)
U-Boot	sentry.shmm500.u-boot	sentry.shmm1500.u-boot
Linux kernel	sentry.shmm500.kernel	sentry.shmm1500.kernel
RFS	sentry.shmm500.rfs	sentry.shmm1500.rfs

The image upgrade works as follows. For each of the specified source images, the image is copied to the ShMM using the specified copy protocol. If no `-d` specifier is supplied, the image is first copied to the RAM file system of the ShMM (specifically, the copy is to the `/tmp` directory) and then moved to Flash (that is, copied to the destination partition in Flash and then removed from the RAM file system).

If there is a `-d` specifier supplied in the call to `rupgrade_tool -s`, the intermediate copy to the `/tmp` directory is skipped and the image is copied directly to its destination in the Flash. Use of this specifier is intended for a scenario where there is insufficient run-time memory on the ShMM for an intermediate copy to the RAM file system.

If no `-d` specifier is supplied, the reliable upgrade procedure invokes a special script, the main purpose of which is to validate images after they are copied to the `/tmp` directory. If `-d` specifier is present, no such validation is performed.

Currently, the script `/etc/upgrade/step4vshm` supplied with the Shelf Manager does not perform specific image validation steps, but does take responsibility for filling in the Flash partitions for which no images are provided in the current call to `rupgrade_tool` (as would happen in a partial upgrade scenario). These partitions are copied from the current persistent Flash

to the provisional Flash. For example, if the current partial upgrade provides only a new RFS image, the script copies the U-Boot and kernel partitions from the persistent Flash to the provisional Flash. As soon as the first image has been installed to its destination, the utility proceeds to the second image (if there is one), and so on, until all the supplied image files have been successfully installed to Flash. A failure to successfully install an image immediately terminates the upgrade procedure (vs. skipping a failing image and proceeding to the next one).

This approach enables the user to separately upgrade the three parts of ShMM firmware (U-Boot, kernel and RFS image). However, the user should bear in mind that the parts that are not explicitly updated is copied from persistent Flash. nVent recommends use of one of the following upgrade approaches:

- Explicitly upgrading all three partitions.
- When fewer than 3 partitions are explicitly upgraded, omitting the `-d` specifier; in that case, the special script mentioned above automatically ensures that every upgrade is effectively a full upgrade covering all three partitions.

`<protocol>` specifies a file copy protocol used to pull each of the specified `<src>` files to the ShMM and can be any of the following:

- `no` - No copy is performed. This protocol assumes that all of the specified `<src>` files were pushed to the `/tmp` directory prior to start of the reliable upgrade procedure. This protocol choice cannot be used in conjunction with the `-d` option.
- `cp:<dir>` - Simple copy. This protocol assumes that all of the specified `<src>` files are to be copied from the specified directory in the ShMM local file system by the `cp` command. This protocol can be useful, for instance, for installation of upgrade images from an NFS mounted file system or even from a JFFS2 file system.
- `ftp:<server>:<dir>:<user>[:<pwd>]` - Copy from a remote FTP server. This protocol assumes that all of the specified `<src>` files are to be copied to the ShMM from the FTP server host specified as the host name or the IP address by `<server>`. All the images must reside in the directory specified by `<dir>` on the remote FTP server. The FTP connection is made using the account specified by the `<user>` parameter, with the password specified by the optional `<pwd>` parameter. If no `<pwd>` is supplied, the utility prompts for a password.
- `scp:<server>:<dir>:<user>[:<id_file>[:<port>]]` - Copy from a remote SSH server using the SCP protocol. This protocol assumes that all of the specified `<src>` files are to be copied to the ShMM from the secure shell server host specified as the host name or the IP address by `<server>`. All the images must reside in the directory specified by `<dir>` on the remote SSH server. (The path is absolute.) The user name for the SSH connection is specified by the `<user>` parameter. It is not possible to supply a password in this mode, but the `<id_file>` parameter allows specifying the secure shell identity file to be used for password-less authorization. The `<port>` parameter can be used to specify a non-standard port number for the remote SSH server. If the `<port>` parameter is not supplied, the default SSH service port number 22 is used.

A failure in copying an image to the ShMM causes the utility to terminate the upgrade procedure (vs. skipping a failing image and proceeding to the next one).

For each provisional Flash partition upgraded by the `-s` option, the to-be-upgraded partition is given write permissions after the validity of the image has been checked and right before the `src` image is about to be moved to the Flash. Write permissions are removed from the partition immediately after the full image has been moved to Flash. Combined with the fact that all the partitions containing the U-Boot, Linux kernel, and root file system images are read-only on boot-up of the ShMM, this ensures that applications cannot accidentally erase the critical boot-up images.

If the specifier `-R` specifier has been used in the command line, no copy operation is performed and it is assumed that corresponding upgrade images have been already copied into Flash. None of the previously described steps is executed in this case; execution of the reliable upgrade procedure begins from this point.

After all the specified images have been installed to their respective destinations in Flash, the utility invokes a “hook” script that enables custom actions required by an application at the point where the upgrade images have been already installed in Flash but the upgrade procedure has not yet initiated the hardware mechanisms of the reliable upgrade procedure by enabling the ShMM’s upgrade watchdog timer (WDT). (Refer to *ML User’s Guide* Chapter 7 for background and details on the upgrade WDT.)

The hook script `/etc/upgrade/step4hshm` is supplied with the Shelf Manager. It performs the following actions:

- terminates the Shelf Manager, performs a switchover to the backup ShMM without restarting the shelf and stops the ATCA watchdog timer,
- mounts the provisional `/etc` and `/var` Flash partitions and erases all files on them,
- optionally copies the current contents of the `/etc` directory to the provisional `/etc` Flash partition,
- optionally copies the current non-volatile Shelf Manager information from the directory `/var/nvdata` to the provisional `/var` file system; or optionally copies the whole `/var` directory to the provisional `/var` Flash partition,
- temporarily (just for the next boot), sets the boot delay to 0; this is done to minimize the time of the next boot and prevent the reliable upgrade watchdog timer from premature expiration.

This script is invoked as a sub-shell and given a single parameter, which is either the string specified by `<args>` or an empty string if no `args` is supplied. The parameter defines the mode of operation of the script with respect to copying non-volatile information from the persistent Flash partitions to the provisional Flash partitions, and can take the following values, each of which triggers a corresponding set of actions:

- no parameter supplied – the script erases both the provisional `/etc` and provisional `/var` directories, then copies Shelf Manager non-volatile information from the directory `/var/nvdata` to the provisional `/var` partition. The script also copies the contents of `/etc/ssh` directory to the provisional `/etc` partition to preserve SSH server keys.

This is the default mode of operation; in this case, the non-volatile data is preserved but the Shelf Manager configuration file is taken from the new RFS image.

- *conf* – the script adopts the default mode of operation, except that Shelf Manager configuration files are preserved via */var/upgrade* filesystem. During the next boot, configuration files are copied from */var/upgrade* to the */etc* partition.
- *erase* – the script erases both the provisional */etc* and provisional */var* directories; they are restored from the RFS default values during the next boot; the current Shelf Manager non-volatile data and configuration are not preserved. The only preserved configuration in this mode is the */etc/ssh* directory, which contains SSH server keys that are copied to the provisional */etc* partition. The current Shelf Manager non-volatile data and configuration are not preserved.
- *etc_copy* – the script erases both the provisional */etc* and provisional */var* directories; then it copies the contents of */etc* and the non-volatile information from the directory */var/nvdata* to the provisional Flash partitions. In this case, both the non-volatile data and the Shelf Manager configuration file are preserved.
- *copy* – the script erases both the provisional */etc* and provisional */var* directories, then copies the full contents of the */etc* and */var* directories onto the provisional partition. In this case, not only the configuration, but also the executable files placed to */var/bin* is copied and overrides executable files with the same name from the RFS image. This mode of operation is useful if the directory */var/bin* contains some special executables (e.g. a special version of the Shelf Manager or other utilities) that must be preserved across the upgrade.
- *flip* – the script overrides all other hook options, causing just the termination of the ShMM, flipping of the Flash banks, and a reboot of the ShMM with the upgrade WDT active.

*Note: When upgrading ShMM software from an encryption-present RFS to an encryption-removed RFS, always use the *erase* option to repopulate the directory */etc* from the RFS. This is because some files in the */etc* directory in an encryption-present RFS are encrypted and are not readable in a non-encrypted RFS; during the reliable upgrade, they should be replaced from the new encryption-removed RFS, where they are not encrypted.*

The script returns 0 on success and non-zero for failure. If a non-zero value is returned, the upgrade procedure is terminated.

The utility starts the upgrade WDT with a 12.8 sec timeout period. This timeout period is considered sufficient for any software that boots after the reset to proceed to the point where it is able to call *upgrade_tool -c* (which strobes the upgrade WDT in case it is active) without having to strobe the upgrade WDT in the interim. The utility performs a strobe of the upgrade WDT just before resetting the ShMM.

-c [-v]

Proceed with the reliable upgrade procedure after the ShMM is booted from the provisional Flash. The invocation of *upgrade_tool -c* is done from the */etc/rc* script. As described below, certain situations discovered by *upgrade_tool -c* imply a failure in the upgrade procedure and require corrective actions, including those resulting in the need to soft reset the ShMM. This means that an invocation of *upgrade_tool -c* may not return and instead may result in a

reset of the ShMM. If a reset takes place, it reverts the ShMM to the software installed in the persistent Flash.

If the upgrade WDT is active and has fired at any step prior to invocation of `rupgrade -c`, this means that the ShMM already reverted to the software in the persistent Flash. In this scenario, the utility disables the upgrade WDT and returns to the use of persistent Flash and terminates the upgrade procedure.

If the upgrade WDT is active but has not fired, this means that the ShMM successfully booted (up to this point) from the provisional Flash. The utility strobes the upgrade WDT and exits with the return code of 0 indicating that there is an upgrade procedure session in progress.

If the upgrade WDT is not active but the content of the `/var/upgrade/status` file indicates that the upgrade procedure is still in progress, this means that the ShMM rebooted due to a power-cycle at one of the steps of the upgrade procedure. In this scenario the utility performs the same corrective actions as for the situation when the upgrade WDT is active and has fired (see above).

Finally, if the upgrade WDT is not active and `/var/upgrade/status` is either not present or indicates that the upgrade procedure has finished (either successfully or unsuccessfully), the utility exits with the return value of 1, indicating that there is no upgrade procedure in progress.

`-f [--hook=args] [-v]`

Complete the upgrade procedure. The invocation of `rupgrade_tool -f` is done from inside the Shelf Manager after the Shelf Manager successfully completes its initialization. If the Shelf Manager is not started automatically, that invocation is done at the end of the `/etc/rc` script.

As soon as invoked, `rupgrade_tool -f` strobes the upgrade WDT and proceeds with establishing the new persistent Flash and disabling the upgrade WDT.

After disabling the upgrade WDT, the upgrade procedure can invoke a special script to take any actions necessary after the hardware resources associated with a reliable upgrade have been returned to their normal state, but before a completion record has been added to the status file. This feature is not currently used during reliable upgrades of the Shelf Manager.

As the last step, the utility updates `/var/upgrade/status` with a record indicating a successful completion of the upgrade procedure and exits with a value of 0.

`-w [-f]`

Print the current status of the most recent upgrade procedure. Essentially, this option dumps the content of the `/var/upgrade/status` file to `stdout`.

`rupgrade_tool -w` returns a value of 0 if the upgrade procedure has completed successfully, 1 if the upgrade procedure has completed unsuccessfully, an appropriate error code if there is no `/var/upgrade/status` file to be found.

If the `-f` specifier is supplied, `rupgrade_tool -w` removes the `/var/upgrade/status` file before exiting.

`-S [-v]`

Strobe the upgrade WDT. `rupgrade_tool -S` is intended as a shell-level interface to strobe the upgrade WDT, for use by newly installed software that is validating its sanity.

`rupgrade_tool -S` returns a value of 0.

`-u`

This option was previously used to undo a successful upgrade session. As of the 2.7.3 release, this option has been removed and is no longer supported. To undo an upgrade, use the options `-sR` `--hook=flip` or `-sR`.

`-C {--<dst>=<src>}... [--proto=<protocol>] [-v] [-d] [-a]`

Copy images to the Flash, but do not execute the final upgrade step. This usage scenario is complementary to the “upgrade without copying” usage scenario. With the option `-C`, `rupgrade_tool` performs the same operations as described above for the option `-s`, and terminates after copying images into the Flash (where the `rupgrade_tool` run with the options `-s -R` starts its operation).

`-h`

Show help to `stdout`.

7.7 Reliable upgrade utility use scenarios

It is intended that the reliable upgrade utility is used in the following sequence in order to carry out an upgrade of the ShMM:

- The user makes a call to `rupgrade_tool -s` to initiate the upgrade procedure. The call can be made either locally from the ShMM serial console or remotely over the network via `telnet`, `rsh`, `ssh`, or any equivalent.
- The user waits for the `rupgrade_tool -s` to reboot the ShMM. If the user is connected to the serial console locally, the status of the reboot is obvious from the messages printed by the U-Boot firmware and Linux to the serial console. If the connection to the ShMM is remote, the status of the reboot is less obvious.
For instance, a telnet connection times out on the reboot of the ShMM. The user can either assume that the upgrade procedure has been carried out successfully or wait for a certain amount of time required for the upgrade session to complete and then make a call to `rupgrade_tool -w` (again, remotely, over any of the remote shell tools mentioned above) in order to find out the status of the upgrade session. The amount of time to wait depends on the size of the upgrade images and the copy protocol used to pull the images to the ShMM as well as actions performed by the image validation script.
- On the ShMM, the startup script `/etc/rc` unconditionally makes a call to `rupgrade_tool -c`. If the call returns a value of 1, indicating that there is no upgrade in progress or an error code value indicating that the upgrade session has failed, the startup scripts proceed with the normal mode boot-up sequence. If however, a value of 0 is returned, indicating that there is an upgrade session in progress, the startup scripts proceed with validation of the sanity of the newly installed software, calling `rupgrade_tool -S` in the middle of operation to strobe the upgrade WDT in case the validation takes longer than the upgrade WDT timeout period, and finally start the Shelf Manager to perform final validation. The watchdog timer interval is set to 12.8 seconds; so the processing time in the `/etc/rc` script between the call to `rupgrade_tool -c` and strobing the WDT and between strobing the WDT and starting the Shelf Manager must not exceed 12.8 seconds each.

- During initialization, the Shelf Manager strobes the upgrade WDT once again, before trying to establishing a network connection with the peer Shelf Manager. Establishing a network connection may take up to 6 seconds. After that, and after successfully finishing the initialization (which indicates validity of the new configuration), the Shelf Manager makes a call to `rupgrade_tool -f`, which completes the upgrade procedure.
- The user optionally calls `rupgrade_tool -w` in order to find out the status of the upgrade session. As explained above, this option may be especially useful for a remote upgrade session where the progress of the upgrade cannot be observed directly from the messages printed to the serial console, as is the case for a local upgrade.
- After the completion of the reliable upgrade, the user can revert to the original images if he/she detects that the new images are not acceptable for any reason. To do this, the user calls `rupgrade_tool -sR -hook=flip`.

If necessary, the above sequence can be easily automated by developing a simple script designed to run on a remote network host. Alternatively, an operator can carry out the reliable upgrade manually, either locally from the serial console or remotely over the network.

7.8 Reliable upgrade examples

This section provides reliable upgrade examples for both ShMM-500 and ShMM-1500. The reliable upgrade procedure works similarly on these two types of ShMMs, but there may be minor differences in the appearance of the command output.

7.8.1 Example 1

This example shows a reliable upgrade of all three components (U-Boot, kernel and RFS image), copying `/etc` and `/var/nvdata` non-volatile directories to the provisional Flash. All images are taken from the local `/tmp` (which implies that they have already been copied there in some unspecified way). The U-Boot image is taken from `/tmp/u-boot.bin`, the kernel image is taken from `/tmp/sentry.kernel`, the RFS image is taken from `/tmp/sentry.rfs`. The upgrade procedure is started from the serial console. Comments are interspersed in the console log to provide additional background on the steps of the upgrade procedure.

First, the `rupgrade_tool` is started from the command prompt. The parameters show that all three Flash images are to be updated, with the Shelf Manager non-volatile data and configuration file preserved as well.

```
# rupgrade_tool -s --k=sentry.kernel --r=sentry.rfs --u=u-boot.bin
--hook=etc_copy -v
```

```
rupgrade_tool: PLB is 5
rupgrade_tool: EEPROM page saved
rupgrade_tool: persistent flash is 0
rupgrade_tool: provisional flash is 1
rupgrade_tool: copying image(s)
```

The upgrade utility attempts to invoke a validation script to check the images in `/tmp` currently supplied. If any of the specified file designators is not found in `/tmp`, the utility stops and a message like the following is produced.

```
rupgrade_tool: cannot open /tmp/u-boot.bin for reading.
```

```
rupgrade_tool: failed to copy images to flash
```

The utility proceeds to copy the images to the specified destinations in provisional Flash.

```
rupgrade_tool: invoking scripts (step4v*) [--u=u-boot.bin --k=sentry.kernel --r=sentry.rfs --hook=etc_copy]
```

```
rupgrade_tool: copying u-boot.bin from /tmp to /dev/mtdchar8 using 'cp' protocol
```

```
rupgrade_tool: copying sentry.kernel from /tmp to /dev/mtdchar7 using 'cp' protocol
```

```
rupgrade_tool: copying sentry.rfs from /tmp to /dev/mtdchar9 using 'cp' protocol
```

```
rupgrade_tool: invoking scripts (step4h*) [etc_copy]
```

At this point, the *step4hshm* hook script is invoked; it stops the Shelf Manager and copies non-volatile information to the provisional Flash.

```
/etc/upgrade/step4hshm: Stopping Shelf Manager...
```

```
/etc/upgrade/step4hshm: Erasing /var and /etc, copying /var/nvdata...
```

```
/etc/upgrade/step4hshm: Operation: copy /etc and /var/nvdata.
```

```
/etc/upgrade/step4hshm: Copying completed.
```

```
rupgrade_tool: image(s) copy OK
```

```
rupgrade_tool: watchdog started
```

```
rupgrade_tool: selected provisional flash
```

```
rupgrade_tool: reboot
```

```
Restarting system.
```

Here, the reliable upgrade procedure resets the ShMM. This causes U-Boot to start from the provisional Flash.

```
** Resetting Integrated Peripherals
```

```
U-Boot 1.1.2 (May 12 2005 - 21:27:13)
```

```
CPU: Au1550 324 MHz, id: 0x02, rev: 0x00
```

```
Board: ShMM-500
```

```
S/N: 08000412
```

```
DRAM: 64 MB
```

```
Flash: 16 MB
```

```
In: serial
```

```
Out: serial
```

```
Err: serial
```

```
Net: Au1X00 ETHERNET
```

```
Hit any key to stop autoboot: 0
```

```
## Booting image at bfb00000 ...
```

```
Image Name: MIPS Linux-2.4.26
```

```
Created: 2017-08-24 13:29:50 UTC
```

```
Image Type: MIPS Linux Kernel Image (gzip compressed)
```

```
Data Size: 844843 Bytes = 825 kB
```

```
Load Address: 80100000
```

```
Entry Point: 802bc040
```

```
Verifying Checksum ... OK
```

```
Uncompressing Kernel Image ... OK
```



```
## Loading Ramdisk Image at bfc40000 ...
Image Name:   sentry RFS Ramdisk Image
Created:      2017-08-27   10:51:03 UTC
Image Type:   MIPS Linux RAMDisk Image (gzip compressed)
Data Size:    2465924 Bytes =   2.4 MB
Load Address: 00000000
Entry Point:  00000000
Verifying Checksum ... OK
Starting kernel ...
init started: BusyBox v0.60.5 (2005.06.15-14:45+0000) multi-call binary
/etc/rc: Mounted /proc
/etc/rc: Mounting filesystems...
/etc/rc: Mounted /dev/pts
/etc/rc: Mounted /dev/mtdblock0 to /var
/etc/rc: Mounted /dev/mtdblock10 to /var/upgrade
```

At this point in the execution of the *rc* script, it invokes *rupgrade_tool -c* to check whether a reliable upgrade is in progress. The tool returns 0, confirming that an upgrade is in progress. Given that result, the *rc* script continues with the startup process.

```
/etc/rc: Checking the reliable upgrade watchdog timer...activated
/etc/rc: Mounted ram disk to /var/log
/etc/rc: Started syslogd and klogd
/etc/rc: Mounted ram disk to /var/tmp
/etc/rc: Setting hostname shmm+193
```

Since a reliable upgrade is in progress, the watchdog timer is strobed once more in the *rc* script.

```
/etc/rc: Strobing the reliable upgrade watchdog timer
/etc/rc: Mounted /dev/mtdblock1 to /etc
/etc/rc: Calling /etc/rc.carrier3
Board Hardware Address: 0xFE
/etc/netconfig: /etc/hosts has valid 192.168.1.193 entry
/etc/netconfig: Updating /etc/profile.sentry with IP settings
/etc/netconfig: ifconfig eth0 192.168.1.193
/etc/netconfig: ifconfig eth1 192.168.0.193
/etc/netconfig: route add default gw 192.168.1.253
/etc/netconfig: Starting inetd...
/etc/rc.carrier3: Starting up IPMBs ...
/etc/rc.carrier3: Updating /etc/profile.sentry with specific settings
/etc/rc.carrier3: Starting snmpd...
/etc/rc.carrier3: Starting httpd...
/etc/rc.carrier3: Starting Shelf Manager ...
<I> 02:48:08.463   [171] Pigeon Point Shelf Manager ver. 3.7.0. Built on Aug 27
2017 14:48:57
<*> 02:48:08.469   [171] Limits: code=(400000:5076f0), end_data=10062000,
start_stack=7fff7e30, esp=7fff78a0, eip=2ab0d2e4
```


The Shelf Manager starts and finalizes the reliable upgrade, calling `rupgrade_tool -f`.

```
eth0: link up
eth1: link up
eth1: going to full duplex
shmm+193 login:root
BusyBox v0.60.5 (2005.05.12-22:46+0000) Built-in shell (msh)
```

Finally, the user checks the status of the reliable upgrade, by calling `rupgrade_tool -w`.

```
# rupgrade_tool -w
Recent upgrade status:
1:PLB is 5
1:EEPROM page saved
2:persistent flash is 0
3:provisional flash is 1
4:copying image(s)
4:invoking scripts (step4v*) [--u=u-boot.bin --k=sentry.kernel --r=sentry.rfs -
-hook=etc_copy]
4:copying u-boot.bin from /tmp to /dev/mtdchar8 using 'cp' protocol
4:copying sentry.kernel from /tmp to /dev/mtdchar7 using 'cp' protocol
4:copying sentry.rfs from /tmp to /dev/mtdchar9 using 'cp' protocol
4:invoking scripts (step4h*) [etc_copy]
4:image(s) copy OK
5:watchdog started
6:selected provisional flash
7:reboot
9:WDT not fired, upgrade in progress.
11:provisional flash 1, updating EEPROM
12:EEPROM updated
13:upgrade WDT disabled
13:invoking scripts (step13h*) []
14:upgrade completed successfully
#
```

7.8.2 Example 2

This example shows a reliable upgrade of the RFS image only, copying `/etc` and `/var/nvdata` non-volatile directories to provisional Flash. The RFS image is taken from an FTP server at the IP address 192.168.1.253. The path to the RFS image on the FTP server is `/tftpboot/ru-ppc/sentry.sentry1500.rfs`. The upgrade procedure is started from a telnet session.

Note: Since only the RFS image is explicitly updated, the U-Boot and kernel images are automatically copied from the persistent Flash partition to the provisional Flash partition.

The local system must be able to access the FTP server over the network (that is, its network adapter must be up and configured and a route must exist from the ShMM to the FTP server). In

the example below, the ShMM is configured with the network address 192.168.1.174 (which is in the same network with the FTP server):

```
# telnet 192.168.1.174
Trying 192.168.1.174...
Connected to 192.168.1.174.
Escape character is '^]'.
BusyBox on shmm+174 login: root
BusyBox v0.60.5 (2005.05.07-17:27+0000) Built-in shell (msh)
```

The parameters to `rupgrade_tool -s` indicate that only the RFS is being upgraded and that the copy protocol is FTP, accessing a specified IP address and file, with user admin and no password supplied.

```
# rupgrade_tool -s --r=sentry.shmm1500.rfs
--proto=ftp:192.168.1.253:/tftpboot/ru-ppc:admin --hook=etc_copy -v
rupgrade_tool: PLB is 9
rupgrade_tool: EEPROM page saved
rupgrade_tool: persistent flash is 0
rupgrade_tool: provisional flash is 1
rupgrade_tool: copying image(s)
rupgrade_tool: copying sentry.shmm1500.rfs from 192.168.1.253:/tftpboot/ru-ppc
to /tmp using 'ftp' protocol 220 (vsFTPD 1.1.3) USER admin
```

The user is asked here for a password to the FTP site; that password is entered manually.

```
331 Please specify the password.
PASS *****
230 Login successful. Have fun.
TYPE I
200 Switching to Binary mode.
PASV
227 Entering Passive Mode (192,168,1,253,112,206) RETR /tftpboot/ru-
ppc/sentry.shmm1500.rfs
150 Opening BINARY mode data connection for /tftpboot/ru-ppc/sentry.shmm1500.rfs
(4069968 bytes).
226 File send OK.
QUIT
221 Goodbye.
```

In the next step, a special script `step4vshm` is invoked, that copies the U-Boot and kernel images from the persistent Flash to the provisional Flash. After that, the upgrade utility proceeds to copy the RFS image to its designated position in provisional Flash.

```
rupgrade_tool: invoking scripts (step4v*) [--r=sentry.shmm1500.rfs --
proto=ftp:192.168.1.253:/tftpboot/ru-ppc:admin:assword --hook=etc_copy]
/etc/upgrade/step4vshm: Erasing /dev/mtdchar7...Done
/etc/upgrade/step4vshm: Copying Kernel from /dev/mtdchar2 to
/dev/mtdchar7...Done
/etc/upgrade/step4vshm: Erasing /dev/mtdchar8...Done
```

```
/etc/upgrade/step4vsh: Copying U-Boot from /dev/mtdchar3 to
/dev/mtdchar8...Done
rupgrade_tool: copying sentry.shmm1500.rfs from /tmp to /dev/mtdchar9 using 'cp'
protocol
rupgrade_tool: invoking scripts (step4h*) [etc_copy]
/etc/upgrade/step4hsh: Stopping Shelf Manager... Done
/etc/upgrade/step4hsh: Cleaning new /var partition.... Done
/etc/upgrade/step4hsh: Cleaning new /etc partition.... Done
/etc/upgrade/step4hsh: Copying /etc contents..... Done
/etc/upgrade/step4hsh: Copying /var/nvdata contents... Done
/etc/upgrade/step4hsh: Upgrade complete.
rupgrade_tool: image(s) copy OK
rupgrade_tool: watchdog started
rupgrade_tool: selected provisional flash
rupgrade_tool: reboot
Restarting system.
```

The *step4hsh* hook script is invoked, which stops the Shelf Manager and preserves the non-volatile data. The utility then starts the upgrade WDT and reboots.

```
rupgrade_tool: invoking scripts (step4h*) [etc_copy]
/etc/upgrade/step4hsh: Stopping Shelf Manager... Done
/etc/upgrade/step4hsh: Cleaning new /var partition.... Done
/etc/upgrade/step4hsh: Cleaning new /etc partition.... Done
/etc/upgrade/step4hsh: Copying /etc contents..... Done
/etc/upgrade/step4hsh: Copying /var/nvdata contents... Done
/etc/upgrade/step4hsh: Upgrade complete.
rupgrade_tool: image(s) copy OK
rupgrade_tool: watchdog started
rupgrade_tool: selected provisional flash
rupgrade_tool: reboot
Restarting system.
```

At this point, the telnet session is closed after a certain inactivity period; after several seconds, it is possible to reconnect to the target again and check the status of the reliable upgrade by invoking *rupgrade_tool -w*.

```
# telnet 192.168.1.174
```

```
Trying 192.168.1.174...
```

```
Connected to 192.168.1.174.
```

```
Escape character is '^]'.
```

```
BusyBox on shmm+174 login: root
```

```
BusyBox v0.60.5 (2005.05.07-17:27+0000) Built-in shell (msh)
```

```
#
```

```
# rupgrade_tool -w
```

```
Recent upgrade status:
```

```
1:PLB is 9
```

```

1:EEPROM page saved
2:persistent flash is 0
3:provisional flash is 1
4:copying image(s)
4:copying sentry.shmm1500.rfs from 192.168.1.253:/tftpboot/ru-ppc to /tmp using
'ftp' protocol 4:invoking scripts (step4v*) [--r=sentry.shmm1500.rfs --
proto=ftp:192.168.1.253:/tftpboot/ru-ppc:admin --hook=etc_copy] 4:copying
sentry.shmm1500.rfs from /tmp to /dev/mtdchar9 using 'cp' protocol 4:invoking
scripts (step4h*) [etc_copy]
4:image(s) copy OK
5:watchdog started
6:selected provisional flash
7:reboot
9:WDT not fired, upgrade in progress.
10:provisional flash 1, disabling watchdog 12:upgrade WDT disabled 13:EEPROM
updated 13:invoking scripts (step13h*) [] 14:upgrade completed successfully
#

```

7.8.3 Example 3

This example shows an unsuccessful reliable upgrade. Power is turned off after the boot from the provisional Flash, but before the reliable upgrade is finalized. After turning the power back on, the rollback to the persistent Flash occurs. This reliable upgrade is initiated from the serial console. All three images are assumed to be already in */tmp*.

```

# rupgrade_tool -s --k=sentry.kernel --r=sentry.rfs --u=u-boot.bin
--hook=etc_copy -v
rupgrade_tool: PLB is 5
rupgrade_tool: EEPROM page saved
rupgrade_tool: persistent flash is 0
rupgrade_tool: provisional flash is 1
rupgrade_tool: copying image(s)
rupgrade_tool: invoking scripts (step4v*) [--u=u-boot.bin --k=sentry.kernel --
r=sentry.rfs --hook=etc_copy]
rupgrade_tool: copying u-boot.bin from /tmp to /dev/mtdchar8 using 'cp' protocol
rupgrade_tool: copying sentry.kernel from /tmp to /dev/mtdchar7 using 'cp'
protocol
rupgrade_tool: copying sentry.rfs from /tmp to /dev/mtdchar9 using 'cp' protocol
rupgrade_tool: invoking scripts (step4h*) [etc_copy]
Stopping Shelf Manager...
Pigeon Point Shelf Manager Command Line Interpreter
Terminating the Shelf Manager
Erasing /var and /etc, copying /var/nvdata...
Operation: copy /etc and /var/nvdata.
Copying completed.
rupgrade_tool: image(s) copy OK
rupgrade_tool: watchdog started

```

```
rupgrade_tool: selected provisional flash
rupgrade_tool: reboot
Restarting system.
```

The reliable upgrade procedure resets the ShMM here and starts U-Boot from the provisional Flash.

```
** Resetting Integrated Peripherals
U-Boot 1.1.2 (Apr 11 2005 - 15:16:25)
CPU: Au1550 324 MHz, id: 0x02, rev: 0x00
Board: ShMM-500
S/N: 8000044
DRAM: 64 MB
Flash: 16 MB
In: serial
Out: serial
Err: serial
Net: AulX00 ETHERNET
Hit any key to stop autoboot: 0
## Booting image at bfb00000 ...
   Image Name: MIPS Linux-2.4.26
   Created:    2005-04-11 10:35:08 UTC
   Image Type: MIPS Linux Kernel Image (gzip compressed)
   Data Size: 843129 Bytes = 823.4 kB
   Load Address: 80100000
   Entry Point: 802bc040
   Verifying Checksum ... OK
   Uncompressing Kernel Image ... OK
## Loading Ramdisk Image at bfc40000 ...
   Image Name: sentry RFS Ramdisk Image
   Created:    2005-04-22 9:10:41 UTC
   Image Type: MIPS Linux RAMDisk Image (gzip compressed)
   Data Size: 2400736 Bytes = 2.3 MB
   Load Address: 00000000
   Entry Point: 00000000
   Verifying Checksum ... OK
```

Power is turned off here. After some time, power is turned back on. Assignment of provisional Flash has been lost because of the power loss, so the system reverts back to the persistent Flash.

```
U-Boot 1.1.2 (Apr 11 2005 - 15:16:25)
CPU: Au1550 324 MHz, id: 0x02, rev: 0x00
Board: ShMM-500
S/N: 8000048
DRAM: 64 MB
Flash: 16 MB
```

```

In:    serial
Out:   serial
Err:   serial
Net:   AulX00 ETHERNET
Hit any key to stop autoboot:  0
## Booting image at bfb00000 ...
   Image Name:   MIPS Linux-2.4.26
   Created:      2005-04-11  10:35:08 UTC
   Image Type:   MIPS Linux Kernel Image (gzip compressed)
   Data Size:    843129 Bytes = 823.4 kB
   Load Address: 80100000
   Entry Point:  802bc040
   Verifying Checksum ... OK
   Uncompressing Kernel Image ... OK
## Loading Ramdisk Image at bfc40000 ...
   Image Name:   sentry RFS Ramdisk Image
   Created:      2005-04-11  18:27:17 UTC
   Image Type:   MIPS Linux RAMDisk Image (gzip compressed)
   Data Size:    2372311 Bytes = 2.3 MB
   Load Address: 00000000
   Entry Point:  00000000
   Verifying Checksum ... OK
Starting kernel ...
init started: BusyBox v0.60.5 (2005.02.07-16:45+0000) multi-call binary
hub.c: new USB device AU1550-1, assigned address 2
usb0: ? speed config #1: Ethernet Gadget
usb1: register usbnet usb-AU1550-1, Linux Device
serial#=8000048: not found
/etc/rc: Mounted /proc
/etc/rc: Mounting filesystems...
/etc/rc: Mounted /dev/pts
/etc/rc: Mounted /dev/mtdblock0 to /var
/etc/rc: Mounted /dev/mtdblock10 to /var/upgrade

```

The next step in the *rc* script is to call *rupgrade_tool -c* to check whether a reliable upgrade is in progress. The check determines that an attempted reliable upgrade failed. The message “restoring ADM1060 EEPROM to RAM” refers to the ShMM system supervisory device (an ADM1060), which supervises the ShMM boot up process and implements some of the hardware aspects of the reliable upgrade support. This message indicates that key variables affecting the boot process are being reverted to their state before the reliable upgrade was attempted.

```

/etc/rc: Checking the reliable upgrade watchdog timer
rupgrade_tool: Watchdog not active.
rupgrade_tool: restoring ADM1060 EEPROM to RAM
rupgrade_tool: upgrade failed

```

```
/etc/rc: Rupgrade -c Ret: 255
/etc/rc: Mounted ram disk to /var/log
/etc/rc: Started syslogd and klogd
/etc/rc: Mounted ram disk to /var/tmp
/etc/rc: Setting hostname shmm+193
/etc/rc: Mounted /dev/mtdblock1 to /etc
/etc/rc: Calling /etc/rc.carrier3
Board Hardware Address: 0xFE
/etc/netconfig: /etc/hosts has valid 192.168.1.193 entry
/etc/netconfig: Updating /etc/profile.sentry with IP settings
/etc/netconfig: Starting inetd...
/etc/rc.carrier3: Starting up IPMBs ...
/etc/rc.carrier3: Updating /etc/profile.sentry with specific settings
/etc/rc.carrier3: RC2 daemons not started by request
```

8 Re-programming the ShMM-700

This section describes how to update the firmware on the ShMM-700, using the reliable upgrade facilities. Comparisons and contrasts with the ShMM-500/ShMM-1500 style of firmware upgrades are included, but this section is also intended to be usable for readers without prior background on the ShMM-500/ShMM-1500 firmware upgrade facilities.

8.1 In this section

This section contains the topics listed below. Just click on a topic to go to it.

- Firmware reliable upgrade procedure overview
- Flash partitioning and image layout
- Changing the ‘user’ flash partition format to JFFS2
- Reliable upgrade procedure status file
- Reliable upgrade utility
- Reliable upgrade of SmartFusion cSoC image
- Reliable upgrade utility use scenarios
- Finding out component versions via the imls utility
- Reliable upgrade examples

8.2 Firmware reliable upgrade procedure overview

WARNING: Before downgrading the ShMM-700R Linux kernel, RFS and U-Boot images from a 3.5.0 or later Shelf Manager release to a pre-3.5.0 Shelf Manager release, users **MUST** verify that the Flash ‘user’ partition is formatted to UBIFS, for further details refer to section 8.4.

For the ShMM-700, a reliable upgrade procedure for the firmware images on a running and functioning ShMM is provided. An arbitrary combination of the following four software images can be upgraded reliably: the U-Boot firmware, the Linux kernel, the Linux root file system (RFS) and the application package. If an attempted reliable upgrade fails, an automatic rollback to the previous version of the firmware takes place. An upgrade failure can be caused, for example, by installation of a faulty kernel image that cannot properly boot up the ShMM, or a faulty Shelf Manager application that fails to start.

In contrast to the ShMM-500/1500, the ShMM-700 firmware includes a fourth component that contains just the application binaries for the Shelf Manager and accompanying applications; in many cases, it will be sufficient to upgrade only this component in order to upgrade to the next version of the Shelf Manager, which should result in faster upgrades. On the other hand, the RFS image on the ShMM-700 contains only system files and executables, and does not contain application binaries and application-specific files, which makes it smaller and less prone to change in successive versions of the ShMM-700 firmware.

Similar to the ShMM-500/1500, two sets of upgrade component images exist on the ShMM-700. The storage and layout of upgrade components on ShMM-700 is, however, different from ShMM-500/1500. Only the U-Boot images are assigned to dedicated Flash partitions. Kernel and RFS

images are stored as files on a single partition, formatted under the UBIFS file system. U-Boot on ShMM-700 is able to read files from UBIFS partitions and therefore is able to locate and load the proper kernel and RFS images. Also, there exist two instances of the application component package, one of which represents the candidate firmware and the other of which represents the confirmed firmware. They are stored on another (user-accessible) UBIFS Flash partition, in a compressed form. When the RFS is loaded and the initial startup script (*/etc/rc*) is executed, it picks up the appropriate instance of the application package, and unpacks it into certain directories in the root RAMFS file system, thus creating application binaries in the RAMFS file system on the fly.

Also, there are two instances of writable directories */etc* and */var* on the ShMM-700, so that each of the two image sets has its own instance of these directories. These directories contain writable application data. When the corresponding image set is loaded, the corresponding instance of */etc* and */var* directory is found on the Flash and mounted as */etc* and */var* in the current file tree.

When a stable firmware configuration is established in one of these image sets, it is designated as the confirmed image set. When a new firmware configuration is installed, it goes in the other image set, which is initially designated candidate. Once a new firmware configuration in the candidate set is validated, that image set is designated the confirmed image set and continues to be used until a future upgrade cycle starts the process over again.

ShMM reliable upgrades are enabled by hardware mechanisms which ensure that if defective firmware is present in the candidate image set and the candidate image set fails to boot or initialize, an automatic rollback takes place to the safe software copy in the confirmed image set. Even if the candidate image set boots and initializes successfully, but embedded software integrity checks run at the application level during initialization indicate that the new software image is not fully functional, hardware mechanisms can be invoked to reset the ShMM and revert to the safe confirmed image set.

At a higher level, the reliable upgrade utility uses these hardware mechanisms and also implements software mechanisms for carrying out reliable upgrades properly. For the benefit of the user, the progress of the current reliable upgrade session and its final outcome (success or failure) is logged in a special file (*/var/upgrade/status*) located on the User UBIFS partition.

The reliable upgrade procedure implemented on ShMM-700 was designed primarily to handle reliable upgrades of the Pigeon Point Shelf Manager firmware. In addition, however, the procedure can be applied to arbitrary embedded applications running on the ShMM. To achieve that, the procedure supports “hook” scripts that can be invoked at certain stages of the reliable upgrade procedure and can be used for customization of the procedure. This support for hook scripts is briefly discussed in section 8.6; however, the rest of this section 8 mainly concentrates on reliable upgrades in the context of Shelf Manager firmware, where hook scripts are not used.

Also, a special procedure is defined for the reliable upgrade of the SmartFusion cSoC image. The Microsemi SmartFusion cSoC is a hardware component of the ShMM-700 that includes an FPGA fabric and an ARM Cortex-M3 subsystem, combined on a single chip. Correspondingly, the image for the SmartFusion cSoC consists of an FPGA fabric image and the program code for the Cortex-M3 processor, combined in a single composite image. Update of the SmartFusion cSoC image requires special precautions and is implemented as a special case of the reliable upgrade

procedure. This special case is documented in detail in section 8.7. It is expected, however, that the SmartFusion cSoC image will be relatively stable, and that most Shelf Manager releases will not require an upgrade of the the SmartFusion cSoC.

8.3 Flash partitioning and image layout

The ShMM-700 provides a hardware mechanism that allows swapping of the two Flash partitions at certain addresses under control of software running on the i.MX287. This capability is implemented in support of the reliable upgrade procedure and applies to U-Boot images, stored on dedicated Flash partitions. Other firmware images (the Linux kernel, RFS and application package) are stored in files on UBIFS volumes, and there are two sets of these images (numbered 0 and 1). The choice of the set to use is made in U-Boot based on the chosen instance number evaluated by the U-Boot. Therefore, a hardware mechanism is needed only for choosing the U-Boot instance and exposing the chosen instance number to the software.

To accelerate loading of images from a UBIFS partition, U-Boot implements a special cache where it stores the map of the last booted Linux kernel and RFS images. The map describes how the images are laid out in terms of physical Flash sectors, and using this map allows U-Boot, during the second and subsequent boots of the same image set, to load images using direct Flash read operations and to avoid expensive mounting of UBIFS volumes. This information is stored on a dedicated Flash partition. It is invalidated explicitly by the reliable upgrade utility before writing a new image set during a reliable upgrade.

The U-Boot instance number and the number of the chosen set of images is calculated by the following formula:

$$\text{Number} = \text{UBOOT_IMAGE_SEL} \wedge \text{Candidate}$$

Where:

“UBOOT_IMAGE_SEL” is the ShMM-700 soft jumper, the value of which is the number of the current confirmed U-Boot image and other firmware images (0/1)

“Candidate” is a flag that has a value of 1 if a reliable upgrade is currently in progress and the candidate set of images should be loaded and 0 otherwise.

“^” indicates the Exclusive-OR logical operation

The following table provides a summary of the Flash partitions maintained by Pigeon Point Linux for the ShMM-700.

Table 31: Flash partitioning for ShMM-700s (64MB Flash) in UBIFS mode

OFFSET IN FLASH (IN MBYTES)	SIZE (IN MBYTES)	DEVICE NODE	MOUNTED AS (ON STARTUP)	CONTENT
0	0.5	/dev/mtd0	Not mounted	U-Boot image 0
0.5	0.5	/dev/mtd2	Not mounted	U-Boot image 1
1	0.25	/dev/mtd4	Not mounted	U-Boot environment variables, first copy

OFFSET IN FLASH (IN MBYTES)	SIZE (IN MBYTES)	DEVICE NODE	MOUNTED AS (ON STARTUP)	CONTENT
1.25	0.25	/dev/mtd6	Not mounted	U-Boot environment variables, second copy
1.5	0.25	/dev/mtd8	Not mounted	UBIFS cache (used for boot acceleration)
1.75	30.25	/dev/mtd10	Not mounted (temporarily mounted by U-Boot to load images and by the reliable upgrade utility to write images)	Boot partition: UBI with one UBIFS volume (named 'boot'); contains Linux kernel and RFS images
32	32	/dev/mtd12	/0, /1	User partition: UBI with one UBIFS volume (named 'user'); contains application package images, Shelf Manager files, other application files and user data

The following table lists names of the files and directories on UBIFS volumes that contain other firmware images. The file names for the SmartFusion cSoC upgrade files are also listed; the SmartFusion cSoC upgrade process is described in section 8.7.

Table 32: File names for firmware images stored as UBIFS files

FILE NAME	FLASH PARTITION	DESCRIPTION
ulimage.0	Boot	Linux kernel image 0
ulimage.1	Boot	Linux kernel image 1
rfs.0	Boot	Root file system (RFS) image 0
rfs.1	Boot	Root file system (RFS) image 1
a2f-upgrade.dat	Boot	SmartFusion (A2F) cSoC upgrade image (present only during an upgrade of the SmartFusion cSoC)

FILE NAME	FLASH PARTITION	DESCRIPTION
a2f-auto-rollback.dat	Boot	Rollback SmartFusion cSoC upgrade image (for use if a SmartFusion cSoC upgrade fails for any reason)
a2f-manual-rollback.dat	Boot	Previous SmartFusion cSoC upgrade image for a manual rollback
/0/etc	User	/etc directory: instance 0; mounted as /etc if image set 0 is active
/1/etc	User	/etc directory: instance 1; mounted as /etc if image set 1 is active
/0/var	User	/var directory: instance 0; mounted as /var if image set 0 is active
/1/var	User	/var directory: instance 1; mounted as /var if image set 1 is active
/0/sentry.shmm700.app	User	Application package image 0. Starting with release 3.5.0 this file, if image set 0 is active, is a symbolic link to the app.0 file in the Boot partition.
/1/sentry.shmm700.app	User	Application package image 1. Starting with release 3.5.0 this file, if image set 1 is active, is a symbolic link to the app.1 file in the Boot partition.

8.4 Changing the 'user' flash partition format to JFFS2

Starting with release 3.5.0, it is possible to change the format of the 'user' partition from UBIFS to JFFS2 and back. This enhancement was added due to a report of very rare cases of the UBIFS partition becoming unrecoverable after a power cycle during a write operation on Flash. There is no known solution to this reported problem in the UBIFS driver, so the ability to switch to a more reliable filesystem for the 'user' partition has been added. The risk of a 'boot' partition corruption is much lower, since it is mounted in read-only mode except for a short upgrade time that takes about a minute.

Corruption of UBIFS-formatted 'user' partition has the following implications:

- The Shelf Manager Application images for both firmware image sets are located in the 'user' partition, so even if the partition is re-formatted after corruption, the Shelf Manager Application images are lost and have to be restored manually.
- The upgrade status files are located in the 'user' partition, so in case of corruption they are lost.
- The 'user' partition filesystem is not separated between firmware image sets, so corruption leads to the loss of */var* and */etc* files for both image sets. Also, it is possible to fill the entire 'user' partition while working with one image set (for example when the *ipmb_traced* utility is configured to save IPMB trace in a file in flash without the size restriction). In that case, an

upgrade that attempts to preserve all */var* and */etc* contents will not work, due to the lack of file system space.

All these issues are resolved with the JFFS2 conversion enhancement. The Shelf Manager Application images, which are included in the application package, are now located in the safe 'boot' partition; the old locations remain as symbolic links. The 'user' partition is split into 3 parts in JFFS2 mode, so that there are separate partitions for each image set. In addition, there is a partition for the upgrade status files with the minimum possible size. The following table provides a summary of the Flash partitions for the ShMM-700 working in the JFFS2 mode.

IMPORTANT NOTICE	
<p>JFFS2 support for the ShMM-700R is available starting with release 3.5.0 and affects U-Boot, Linux kernel and RFS images. If the ShMM-700R is configured to use JFFS2 and it is necessary to downgrade it to a pre-3.5.0 release, there are two possible ways:</p> <p>Downgrade only the APP image, preserving JFFS2-capable U-Boot, Linux kernel and RFS images from release 3.5.0 or above.</p> <p>Change the 'user' partition format from JFFS2 to UBIFS before the downgrade, so that it is possible to downgrade any of the firmware components.</p>	

Table 33: Flash partitioning for ShMM-700s (64MB Flash) in JFFS2 mode

OFFSET IN FLASH (IN MBYTES)	SIZE (IN MBYTES)	DEVICE NODE	MOUNTED AS (ON ATARTUP)	CONTENT
0	0.5	/dev/mtd0	Not mounted	U-Boot image 0
0.5	0.5	/dev/mtd2	Not mounted	U-Boot image 1
1	0.25	/dev/mtd4	Not mounted	U-Boot environment variables, first copy
1.25	0.25	/dev/mtd6	Not mounted	U-Boot environment variables, second copy
1.5	0.25	/dev/mtd8	Not mounted	UBIFS cache (used for boot acceleration)

OFFSET IN FLASH (IN MBYTES)	SIZE (IN MBYTES)	DEVICE NODE	MOUNTED AS (ON ATARTUP)	CONTENT
1.75	30.25	/dev/mtd10	Not mounted (temporarily mounted by U- Boot to load images and by the reliable upgrade utility to write images)	Boot partition: UBI with one UBIFS volume (named 'boot'); contains Linux kernel and RFS images
32	15.25	/dev/mtdblock6	/0	JFFS2 User partition for image set 0; contains shelf manager files and user data.
47.25	15.25	/dev/mtdblock7	/1	JFFS2 User partition for image set 1; contains shelf manager files and user data.
62.5	1.5	/dev/mtdblock8	/var/upgrade	Partition for the upgrade status files. The minimum possible size for a JFFS2 partition is 5 erase blocks.

Before changing the 'user' partition format to JFFS2 it is necessary to verify that the ShMM-700 is running at least the Linux kernel, RFS and U-Boot images from release 3.5.0, or a later release, which can be done via the `imls` command. If the currently running Shelf Manager release includes the Linux kernel image, the RFS image or the U-Boot image from a release older than 3.5.0 (i.e. the Created date is prior to March 2015), it is necessary to upgrade the image(s) to release 3.5.0 or later, before changing the 'user' partition format to JFFS2.

To change the 'user' partition format to JFFS2, it is necessary to run the following commands.

```
# setenv user_partition jffs2 ; setenv flash_reset user ; reboot
```

WARNING: All files stored in the 'user' partition will be erased when the format is changed to JFFS2 (or from JFFS2 to UBIFS). The default files from the currently loaded RFS image will be copied to the 'user' partition in, for example, the `/var` and `/etc` directories.

The result can be verified in the `mount` command output at the ShMM shell prompt e.g.

```
# mount
```

```
...
```

```
ubi0:boot on /boot type ubifs (ro, sync, relatime, bulk_read)
/dev/mtdblock7 on /var type jffs2 (rw, relatime)
```

```
/dev/mtdblock8 on /var/upgrade type jffs2 (rw,relatime)
/dev/mtdblock7 on /etc type jffs2 (rw,relatime)
```

In case it is necessary to downgrade the Linux kernel image, the RFS image or the U-Boot image to a release older than 3.5.0 that does not support JFFS2, the 'user' partition format must be changed to UBIFS before the downgrade. This can be done via the following commands.

```
# setenv user_partition ubifs ; setenv flash_reset user ; reboot
```

The result can be verified in the `mount` command output at the ShMM shell prompt e.g.

```
# mount
...
ub1:boot on /boot type ubifs (ro, sync, relatime, bulk_read)
ubi0:user on /var type ubifs (rw, sync, relatime, bulk_read)
ubi0:user on /var/upgrade type ubifs (rw, sync, relatime, bulk_read)
ubi0:user on /etc type ubifs (rw, sync, relatime, bulk_read)
```

8.5 Reliable upgrade procedure status file

The special file `/var/upgrade/status` that is located on the User Flash partition contains the status of the most recent reliable upgrade session, that is, the session that is currently in progress or most recently completed.

`/var/upgrade/status` is a text file and consists of multiple lines; each line describes the status of a particular stage in the upgrade procedure. The format of a record line is as follows:
<stage>: <status> : <description>

where

- <stage> is an integer in the range 1 to 14 (where stage 14 is the final stage of a reliable upgrade session, meaning that the procedure is complete)
- <status> is a human-readable string that describes the status ("passed" or "failed") of the corresponding stage
- <description> is a description of the corresponding stage. The list of stages and the corresponding description of each stage is summarized in Table 34 below.

The status file can be used to track the progress of the reliable upgrade procedure; it is also used by the reliable upgrade utility itself to synchronize the state of the upgrade procedure across ShMM reboots.

Table 34: Reliable upgrade stages

STAGE NUMBER	DESCRIPTION
1	Images have been downloaded to /tmp and verified
2	Boot volume mount is about to begin
3	Boot volume mount is complete

STAGE NUMBER	DESCRIPTION
4	Candidate has been marked as non-valid
5	U-Boot image write is complete
6	Kernel image write is complete
7	RFS image write is complete
8	Boot volume umount is complete
9	Application package has been removed
10	New application package write is complete
11	Sync() is complete
12	/var and /etc have been processed
13	First strobe after reboot has been done
14	Upgrade is confirmed

8.6 Reliable upgrade utility

A special user-space utility *rupgrade* is provided on the ShMM. This utility has several modes of operation and allows the user to perform multiple tasks, including reliable upgrades of various ShMM software components, status checks for the most recent upgrade, and manual rollbacks to the previous software versions.

Another utility *backend* is provided that is not intended to be invoked directly by an interactive user. This utility performs more specific actions needed for reliable upgrade, like starting, strobing and stopping the reliable upgrade watchdog timer, getting and setting the number of the candidate and confirmed image set, activation of the candidate image set, etc. It can be invoked from system scripts, like the system startup script */etc/rc*.

Also, the *rupgrade* utility uses the services provided by the *backend* utility during its operation. The *rupgrade* utility is fully functional only when called from the superuser (root) account. If run from a non-root account, the utility can be used for read-only operations (like getting the upgrade status).

As the *rupgrade* utility operates in upgrade mode (see below), it logs the status of each action as it progresses through the stages of the procedure to the status file */var/upgrade/status*. In the case of a failure, the reliable upgrade utility adds a record to this file indicating that the reliable upgrade has completed unsuccessfully and then terminates with an appropriate error code.

The *rupgrade* utility outputs any informational messages to *stdout* and any error messages to *stderr*. All messages are also recorded to the */var/upgrade/log.txt* file.

The *rupgrade* utility can operate in several modes depending on the command-line parameters:

Upgrade mode:

```
rupgrade [-u <URL> | --uboot <URL>] [-k <URL> | --kernel <URL>] [-r
<URL> | --rootfs <URL>] [-a <URL> | --application <URL>] [--base
<string>] [--copy-config | --copy-home | --copy-etc | --copy-var |
--copy-all | --erase-all] [--clone] [--install-only]
[--skip-checksums] [-v <verbose level> | --verbose <verbose level>]
```

The main command-line parameters have the following meanings (and at least one of the following parameters must be specified):

- u, --uboot <URL>* – specify the location of the new U-Boot image
- k, --kernel <URL>* – specify the location of the new Linux kernel image
- r, --rootfs <URL>* – specify the location of the new Linux root file system image
- a, --application <URL>* – specify the location of the new application package image.

The following additional parameters are optional:

- verbose <verbose level>* – specify the verbosity level (see below)
- base <string>* – specify the string used as a prefix for all URLs in the command
- copy-config* – copy the SSH keys, the Shelf Manager configuration files and the subdirectory */var/nvdata* from the current image set to the new image set
- copy-home* – copy the entire */home* directory from the current image set to the new image set
- copy-etc* – copy the entire */etc* directory and the subdirectory */var/nvdata* from the current image set to the new image set
- copy-var* – copy the SSH keys and the entire */var* directory from the current image set to the new image set
- copy-all* – copy the entire */var* directory and the entire */etc* directory from the current image set to the new image set
- erase-all* – erase both the */var* and */etc* directories for the new image set; copy only the SSH keys from the current image set
- skip-checksums* – skip the verification of checksums in the images used for upgrade
- install-only* – don't activate the new image set after the installation; continue running from the current image set. The command *backend --activate* can be used later to activate the new image set.

The following verbosity levels are supported:

- *none* – do not output verbose messages to *stdout*; do not write them to the log file
- *file* – do not output verbose messages to *stdout*; do write them to the log file
- *all* – output verbose messages to *stdout* and write them to the log file

The following command variants invoke the HPM.1 upgrade mode:

```
rupgrade -p <URL> [--copy-config | --copy-home | --copy-etc |
--copy-var | --copy-all | --erase-all] [--skip-checksums]
[--install-only]
```

```
rupgrade --package <URL> [--copy-config | --copy-home |
--copy-etc | --copy-var | --copy-all | --erase-all]
[--skip-checksums] [--install-only]
```

Here, `-p <URL>` or `--package <URL>` specifies the location of the new image in HPM.1 format (which contains some combination of the new U-Boot, kernel, RFS and application package images). Other optional parameters have the same meaning as for a non-HPM.1 upgrade.

The following command variant invokes the SmartFusion (A2F) cSoC upgrade mode (see section 8.7):

```
rupgrade --upgrade-a2f <URL> [--skip-checksums]
```

The following command invokes rollback mode:

```
rupgrade --rollback
```

In this mode, the rollback takes place to the other image set (which is assumed to be the image set used before the last upgrade). Rollback doesn't affect the state of the SmartFusion cSoC image because the SmartFusion cSoC has its own auto-rollback mechanism that is utilized during an A2F reliable upgrade procedure.

Additional command variants invoke other modes of the reliable upgrade utility.

Get Upgrade Status mode (output current upgrade status on *stdout*):

```
rupgrade -s
```

```
rupgrade --status
```

Help mode (output a brief help summary on *stdout*):

```
rupgrade -h
```

```
rupgrade --help
```

Get Version mode (output the version of the utility to *stdout*):

```
rupgrade --version
```

Add Hook Mode:

```
rupgrade --add-hook <URL>
```

In this mode, the utility adds a new hook script (specifically, the script file designated by the `<URL>`) to the list of scripts executed during the upgrade procedure. Scripts are stored in the directory `/etc/hook`; to remove a hook script, just delete the corresponding file from this directory.

```
rupgrade --clone
```

In this mode the utility copies the current image set into the new image set. This new image set is not activated.

The following table lists the protocols that are recognized in the URLs by the reliable upgrade utility:

Table 35: Protocols recognized in the URLs by the reliable upgrade utility

PROTOCOL NAME	DESCRIPTION	URL SYNTAX	EXAMPLE URL
File	Get image from a local ShMM file.	file://<abs_path> <abs_path> - absolute path to the image file on the ShMM (starts with "/")	file:///tmp/rootfs.image
HTTP	Download the image from an HTTP server	http://<host>/<path> <host> - host name or IP address <path> - relative path to the image file (path origin depends on HTTP server configuration)	http://192.168.1.253/shmm700/rootfs.image
TFTP	Download the image using the TFTP protocol	tftp://<host>/<path> <host> - host name or IP address <path> - relative path to the image file (path origin depends on TFTP server configuration)	tftp://192.168.1.253/shmm700/rootfs.image
FTP	Download the image using the FTP protocol	ftp://[<user>[:<pass>]@]<host>/<path> <user> - user name <pass> - password <host> - host name or IP address <path> - relative path to the image file (path origin depends on FTP server configuration)	ftp://user:password@192.168.1.253/shmm700/rootfs.image

PROTOCOL NAME	DESCRIPTION	URL SYNTAX	EXAMPLE URL
SCP	Download the image using the SSH/SCP protocol	<code>scp://[<user>@]<host>/<path></code> <user> - user name <host> - host name or ip address <path> - relative path to the image file (path origin is usually the user's home directory). An absolute path can be specified by beginning it with <code>"/"</code> . The password must be entered interactively by the user.	<code>scp://user@192.168.1.253/shmm700/rootfs.image</code> <code>scp://user@192.168.1.253//home/user/shmm700/rootfs.image</code>

In upgrade mode (both regular and HPM.1), the reliable upgrade utility performs the following actions:

- obtains the images to copy: locally or via the network;
- copies the images to the candidate locations;
- copies non-volatile data to the candidate locations or erases the corresponding candidate locations (*/etc*, */var*), depending on the presence of command-line parameters *--copy-config*, *--copy-home*, *--copy-etc*, *--copy-var*, *--copy-all* and *--erase-all*.
- if the option *--install-only* is not specified, resets the ShMM and instructs it to boot the candidate U-Boot (which in turn loads images from the candidate image set).

Because of the last stage, an invocation of *rupgrade* in upgrade mode typically does not return and instead resets the ShMM (unless the option *--install-only* is used). If *rupgrade* in upgrade mode does return, it indicates that the reliable upgrade procedure has failed and was terminated before proceeding to reset the ShMM in order to boot the candidate images.

Not all images need to be specified to *rupgrade* in upgrade mode (or, in the HPM.1 upgrade mode, the HPM.1 image may not include all possible components). If any components are not specified during the reliable upgrade, the missing components are copied from current confirmed component set. This allows reliable upgrades to be done when only certain firmware components (e.g. the application package) change between releases while other components are left intact; in that case, upgrade images can be made smaller and the upgrade itself will be faster. An example of such a partial upgrade is shown as Example 2 at the end of this section.

The *backend* utility has the following synopsis. Only the command-line parameters that are intended to be used by external programs are shown here; the remaining parameters, not shown here, are for the exclusive use of the *rupgrade* utility.

Output the current image set index (0 or 1) to *stdout*:

```
backend -c
```

```
backend --get-current
```

Activate the new image set previously installed via *rupgrade* invocation with the option *--install-only* (that is, start the upgrade watchdog timer, set the other image set as candidate and reset the ShMM instructing it to boot the candidate U-Boot):

```
backend --activate
```

Confirm and finalize the current upgrade:

```
backend --confirm
```

Cancel (abort) the current upgrade:

```
backend --cancel
```

Strobe the reliable upgrade watchdog timer:

```
backend --strobe
```

Help mode (output a brief help to *stdout*):

```
backend -h
```

```
backend --help
```

Get Version mode (output the version of the utility to *stdout*):

```
backend --version
```

Table 36: Standard source file names for upgrade components

COMPONENT NAME	SOURCE FILE NAME
U-Boot	sentry.shmm700.u-boot
Linux kernel	sentry.shmm700.kernel
RFS	sentry.shmm700.rfs
Application Package	sentry.shmm700.app

The image upgrade works as follows:

For each of the specified source images, the image is copied to the ShMM via the protocol specified in the URL (FTP, TFTP, etc.). The image is copied to the RAM file system of the ShMM (specifically, the copy is to the */tmp* directory which is cleaned up after each reboot) and after that it is validated. Image validation involves verification of the image checksum and is performed unless the command-line parameter *--skip-checksums* is specified. This procedure is repeated for all images specified in the command line.

After all images have been successfully obtained and validated, they are finally, one by one, dispatched to their destination locations; that is, U-Boot is copied to the destination Flash partition and other images are copied to appropriate locations on the UBIFS Boot and User volumes. If the reliable upgrade utility does not succeed in copying or validating at least one image, it immediately terminates the overall upgrade. Subsequent images, if any, are ignored; they are not even copied to the ShMM.

Before replacing kernel and RFS images in the Boot volume, the reliable upgrade utility cleans up the UBIFS cache partition to tell U-Boot that it is not valid any more. The UBIFS cache is re-created by U-Boot automatically. When copying the kernel and RFS images to the Boot UBIFS Flash partition, the reliable upgrade utility mounts the partition before copying the image and then unmounts it shortly after the completion of the copy operation. Normally, the Boot Flash partition stays unmounted. This ensures that applications cannot accidentally erase the critical boot-up images and allows U-Boot to cache the location of boot images for faster loading.

After all the specified images have been installed at their respective destinations, the utility mounts the candidate */etc* and */var* directories and either erases all files in them or copies the current contents of these directories to the candidate locations. This behavior is controlled by several command-line parameters (all of them mutually exclusive):

- *--copy-config*: if specified, the SSH keys (*/etc/ssh*), the Shelf Manager configuration files (*/etc/shelfman.conf**) and the subdirectory */var/nvdata* are copied to the candidate directories */etc* and */var* from the current location; any other files in the candidate */etc* and */var* directories are erased
- *--copy-home*: if specified, the entire */home* directory is copied to the candidate directory */home* from the current location; any other files in the candidate */etc* and */var* directories are erased
- *--copy-etc*: if specified, the entire */etc* directory and the subdirectory */var/nvdata* are copied to the candidate directories */etc* and */var* from the current location; any other files in the candidate */var* directory are erased
- *--copy-var*: if specified, the SSH keys (*/etc/ssh*) and the entire */var* directory are copied to the candidate directories */etc* and */var* from the current location; any other files in the candidate */etc* directory are erased
- *--copy-all*: if specified, the entire */etc* directory and the entire */var* directory are copied to the candidate directories */etc* and */var* from the current location
- *--erase-all*: if specified, the candidate directories */var* and */etc* are erased, only the SSH keys (*/etc/ssh*) are copied from the current location.
- If none of the options above is specified (default mode), then the SSH keys (*/etc/ssh*) and the subdirectory */var/nvdata* are copied to the candidate directories */etc* and */var* from the current location; any other files in the candidate */etc* and */var* directories are erased.

The behavior defined by these parameters is compatible with the behavior of the reliable upgrade tool on the ShMM-500 and ShMM-1500.

As the final stage of the upgrade (which takes place if the option *--install-only* is not specified), the utility starts the reliable upgrade watchdog timer (WDT) and resets the ShMM using the FPGA capabilities for ShMM reset. (Refer to ShMM-700 Hardware Architecture Specification for background and details on the reliable upgrade WDT.) The reliable upgrade WDT is started with a 60 second timeout period. This period should be sufficient for the ShMM software to reach (after the reset and without strobing the reliable upgrade WDT) the point in time where the ShMM software can detect whether a reliable upgrade is in progress (by invoking *rupgrade --status*) and if so, strobe the reliable upgrade WDT via *backend --strobe*.

The reliable upgrade utility can invoke an optional “hook” script that enables custom actions required by an application. After each stage of the upgrade procedure, the hook script is invoked as a sub-shell and given a single parameter: the stage number. Table 34 above describes all stages of the reliable upgrade that are visible to the hook script.

The script returns 0 on success and a non-zero value in the case of a failure (in which case the upgrade procedure is terminated).

No hook script is set up by default. The user can run the reliable upgrade utility in the Add Hook mode to install a hook script. The script is copied to the directory `/etc/hook`; so if the directory `/etc` is copied to the candidate location, this installation is permanent; however, if the command-line parameter `--reset-etc` is specified, all hook scripts are erased after the upgrade.

8.7 Reliable upgrade of SmartFusion cSoC image

To upgrade the SmartFusion cSoC image, run the reliable upgrade utility in the SmartFusion (A2F) cSoC upgrade mode:

```
rupgrade --upgrade-a2f <URL> [--skip-checksums]
```

Here `<URL>` designates the file that contains the SmartFusion cSoC upgrade image, the option `--skip-checksums`, if specified, requests the utility not to verify the checksum of the image file. The SmartFusion cSoC image file includes a SmartFusion A2F060 image concatenated with a SmartFusion A2F200 image. Each individual image file includes all the programmable content needed for the A2F060 device or the A2F200 device (including FPGA fabric and Cortex-M3 firmware). The reliable upgrade utility when run in the SmartFusion cSoC upgrade mode will determine if the SmartFusion device on the ShMM-700 is an A2F060 or an A2F200 and program the appropriate image from the concatenated SmartFusion cSoC image file.

A reliable upgrade of SmartFusion cSoC image is done separately from a reliable upgrade of any other ShMM-700 firmware components and involves resetting the ShMM. Also, if the current image set has number 1, this procedure has a side effect that U-Boot image 0 is overwritten by the current U-Boot image; if the current image set has number 0, there are no such side effects.

The reliable upgrade of the SmartFusion cSoC image includes the following steps:

- The upgrade utility copies the upgrade image file onto the Flash Boot partition as `a2f-upgrade.dat`.
- If the current image set has number 1, the upgrade utility saves the contents of U-Boot image 0 to a file, and copies the U-Boot image from partition 1 to partition 0. This is done to ensure that if the upgrade process is interrupted, the ShMM-700 can still recover by booting from partition 0 (see below).
- The upgrade utility reboots the ShMM-700.
- When U-Boot sees the `a2f-upgrade.dat` file in the boot partition at startup, it initiates an A2F upgrade procedure via the SmartFusion cSoC JTAG pins. This procedure is performed with Flash remapping logic turned off (so that U-Boot has direct access to the Flash and U-Boot image 0 is always the active one). Also, the built-in i.MX287 watchdog timer is enabled during the upgrade so that if the upgrade process gets stuck, the ShMM-700 still has a potential to recover. As part of the A2F upgrade procedure the SmartFusion device type (A2F060 or

A2F200) implemented on the ShMM-700 is determined so the appropriate image from the concatenated SmartFusion cSoC image file can be programmed onto the device.

- If the upgrade is successful, U-Boot boots Linux from the current image set. When Linux boots, it performs some basic tests, and if everything works as expected, it disables the built-in i.MX287 watchdog timer and turns on the Flash remapping logic. It then restores the contents of the first U-Boot partition (if the backup file is present) and renames the SmartFusion cSoC upgrade file on the boot partition to `a2f-auto-rollback.dat` to complete the upgrade process.
- If an `a2f-auto-rollback.dat` file was already present, it is renamed as `a2f-manual-rollback.dat` before overwriting it with the current upgrade file. This allows a manual rollback procedure to be initiated by invoking the `rupgrade` utility running in the `--upgrade-a2f` mode with the file `a2f-manual-rollback.dat` as the parameter. The UBIFS Boot volume may need to be explicitly mounted first.
- If the upgrade process is interrupted or an error is detected, the ShMM-700 will get reset with the Flash remapping logic turned off. As a result, it will boot from partition 0, and U-Boot will detect that Flash remapping logic is turned off. In this case, it will try to revert to the previous image by repeating the above steps with the auto-rollback image file name (`a2f-auto-rollback.dat`).

Note: the SmartFusion cSoC upgrade images provided by nVent are encrypted with a private AES key. This key is pre-programmed in the SmartFusion cSoC device on the ShMM-700R during manufacturing. As a consequence, the SmartFusion cSoC upgrade images that are not built by nVent cannot be installed onto the ShMM-700R SmartFusion cSoC device. Attempting to do so will result in a failed upgrade.

8.8 Reliable upgrade utility use scenarios

It is intended that the reliable upgrade utility is used in the following sequence in order to carry out an upgrade of the ShMM:

- The user makes a call to `rupgrade` in upgrade mode (or in HPM.1 upgrade mode) to initiate the upgrade procedure. The call can be made either locally from the ShMM serial console or remotely over the network via `telnet`, `rsh`, `ssh`, or any equivalent.
- The user waits for `rupgrade` to reboot the ShMM. A user who has access to the ShMM serial console can see the messages produced there by U-Boot firmware and the Linux kernel, and therefore can easily monitor the status of the reboot. However, in the case of a remote connection to the ShMM, the status of the reboot can be more difficult to determine.

For example, a telnet- or ssh-based virtual terminal connection usually times out on the reboot of the ShMM. The user can either assume that the upgrade was successful or wait for some time, reconnect to the ShMM and remotely invoke `rupgrade --status` in order to find out the current upgrade session status. The amount of time to wait is determined by the initialization time of the ShMM (the time needed to boot the Linux kernel and start the proper

network daemons) and by the time needed to validate the functionality of the newly installed firmware image set.

- On the ShMM, the startup script `/etc/rc` unconditionally makes a call to `rupgrade --status`. The returned value indicates whether there is an upgrade session in progress (a value `0` is returned), no upgrade session is in progress (a positive value is returned) or the current upgrade session has failed (a negative error code is returned). If there is no upgrade session in progress (the returned value is non-zero), the startup script continues with the normal mode boot-up sequence. If however, there is an upgrade session in progress (the returned value is `0`), the startup script proceeds with sanity validation of the newly installed software. In the middle of validation, it calls `backend --strobe` to strobe the upgrade WDT in case the validation takes longer than the upgrade WDT timeout period. Finally, the startup script starts the Shelf Manager to perform final validation. The watchdog timer interval is set to 60 seconds; so the processing time in the `/etc/rc` script between the call to `rupgrade --status` and the call to strobe the WDT and between the call to strobe the WDT and the startup of the Shelf Manager must not exceed 60 seconds each.
- During initialization, the Shelf Manager strobes the upgrade WDT once again, before trying to establish a network connection with the peer Shelf Manager. Establishing a network connection may take up to 6 seconds. After that, and after successfully finishing the initialization (which indicates validity of the new configuration), the Shelf Manager makes a call to `backend --confirm`, which completes the upgrade procedure.
- At this point, the user can invoke `rupgrade --status` in order to find out the status of the current upgrade session. This may be useful if a reliable upgrade procedure is started remotely over the network via a virtual terminal (`telnet`, `rsh` or `ssh`) connection, because in that case the messages issued to the local serial console are not directly available.
- After the completion of the reliable upgrade, the user can revert to the original images if he/she detects that the new images are not acceptable for any reason. To do this, the user invokes `rupgrade --rollback`.

This sequence can be performed manually by an operator. However, if necessary, one can easily automate the above sequence via a script that runs on a host machine that is connected to the ShMM over the network, and remotely issues appropriate commands to the `rupgrade` and `backend` utilities on the ShMM.

8.9 Finding out component versions via the `imls` utility

A user can run the `imls` utility without parameters on the ShMM-700 to find out information (including version information) about the software components (U-Boot, kernel, RFS and the application package) that comprise both the current and the other image sets. The current image set number and the current version of the SmartFusion cSoC image is also shown. For example:

```
# imls
Current flash bank: 1

Current U-Boot:
Image Name:    U-Boot 9.0.8 for shmm700 board
```

Created: Wed Jul 20 18:57:24 2016
Image Type: ARM U-Boot Firmware (uncompressed)
Data Size: 524224 Bytes = 511.94 kB = 0.50 MB
Load Address: 00000000
Entry Point: 00000000

Current Kernel:

Image Name: Linux-2.6.34.8-shmm700
Created: Wed Jul 20 18:56:12 2016
Image Type: ARM Linux Kernel Image (uncompressed)
Data Size: 1531388 Bytes = 1495.50 kB = 1.46 MB
Load Address: 40008000
Entry Point: 40008000

Current RFS:

Image Name: ShMM700 RFS 3.7.0
Created: Thu Aug 31 17:41:29 2017
Image Type: ARM Linux RAMDisk Image (gzip compressed)
Data Size: 2951487 Bytes = 2882.31 kB = 2.81 MB
Load Address: 00000000
Entry Point: 00000000

Current Application:

Image Name: ShMM700 APP 3.7.0
Created: Thu Aug 31 18:56:06 2017
Image Type: ARM Linux Filesystem Image (gzip compressed)
Data Size: 2504410 Bytes = 2445.71 kB = 2.39 MB
Load Address: 00000000
Entry Point: 00000000

Other U-Boot:

Image Name: U-Boot 9.0.8 for shmm700 board
Created: Wed Jul 20 18:57:24 2016
Image Type: ARM U-Boot Firmware (uncompressed)
Data Size: 524224 Bytes = 511.94 kB = 0.50 MB
Load Address: 00000000
Entry Point: 00000000

Other Kernel:

Image Name: Linux-2.6.34.8-shmm700
Created: Wed Jul 20 18:56:12 2016
Image Type: ARM Linux Kernel Image (uncompressed)

Data Size: 1531388 Bytes = 1495.50 kB = 1.46 MB
 Load Address: 40008000
 Entry Point: 40008000

Other RFS:

Image Name: ShMM700 RFS 3.7.0
 Created: Thu Aug 31 18:39:21 2017
 Image Type: ARM Linux RAMDisk Image (gzip compressed)
 Data Size: 2951352 Bytes = 2882.18 kB = 2.81 MB
 Load Address: 00000000
 Entry Point: 00000000

Other Application:

Image Name: ShMM700 APP 3.7.0
 Created: Thu Aug 31 18:56:06 2017
 Image Type: ARM Linux Filesystem Image (gzip compressed)
 Data Size: 2504410 Bytes = 2445.71 kB = 2.39 MB
 Load Address: 00000000
 Entry Point: 00000000

A2F: SPICOMM protocol v1.7, M3 firmware v1.0, FPGA design v1.0.0.0

8.10 Reliable upgrade examples

This section provides reliable upgrade examples for ShMM-700.

8.10.1 Example 1

This example shows a reliable upgrade of all four components (U-Boot, kernel, RFS image and application package), copying */etc* and */var/nvdata* non-volatile directories to the candidate locations. All images are taken from the local */tmp* (which implies that they have already been copied there in some unspecified way). The U-Boot image is taken from */tmp/u-boot.bin*, the kernel image is taken from */tmp/sentry.kernel*, the RFS image is taken from */tmp/sentry.rfs*, the application image is taken from */tmp/sentry.app.tgz*. The upgrade procedure is started from the serial console. Comments are interspersed in the console log to provide additional background on the steps of the upgrade procedure.

First, the *rupgrade* utility is started from the command prompt in upgrade mode. The parameters show that all four images are to be updated, with the Shelf Manager non-volatile data and configuration file preserved as well.

```
# rupgrade --skip-checksums --copy-etc -k file:///tmp/sentry.kernel -r
file:///tmp/sentry.rfs -u file:///tmp/u-boot.bin -a file:///tmp/sentry.app.tgz
```

The upgrade utility attempts to invoke a validation script to check the images in */tmp* currently supplied. If any of the specified file designators is not found in */tmp*, the utility stops and a message like the following is produced.

```
<INFO> frontend.c:449: Downloading file "file:///tmp/u-boot.bin"
```

```
<ERR> download_easy.c:51: Failed to download file "file:///tmp/u-boot.bin":
Error
```

The utility proceeds to copy the images to the specified destinations in candidate Flash.

```
<INFO> frontend.c:449: Downloading file "file:///tmp/u-boot.bin"
<INFO> frontend.c:449: Downloading file "file:///tmp/sentry.kernel"
<INFO> frontend.c:449: Downloading file "file:///tmp/sentry.rfs"
<INFO> frontend.c:449: Downloading file "file:///tmp/sentry.app"
<INFO> frontend.c:535: Initiating partial upgrade:
<INFO> frontend.c:735:  * uboot "/tmp/_tmp_uboot"
<INFO> frontend.c:745:  * kernel "/tmp/_tmp_kernel"
<INFO> frontend.c:755:  * rfs "/tmp/_tmp_rfs"
<INFO> frontend.c:765:  * app "/tmp/_tmp_app"
<INFO> frontend.c:820: Calling backend to handle partial upgrade "/sbin/backend
--uboot /tmp/_tmp_uboot --kernel /tmp/_tmp_kernel --rootfs /tmp/_tmp_rfs --
application /tmp/_tmp_app -copy-etc"
<INFO> shmm700_hal.c:1285: Write storage state: not available -> available
[ALLOW]
UBI device number 1, total 484 LEBs (31657472 bytes, 30.2 MiB), available 0 LEBs
(0 bytes), LEB size 65408 bytes (63.9 KiB)
<INFO> shmm700_hal.c:369: Resetting UBIFS cache
Erasing 64 Kibyte @ 30000 -- 75 % complete.
<INFO> shmm700_hal.c:1913: Clean on defice 1[0]. Leaving file "/dev/mtd0"
untouched.
<INFO> backend.c:730: Burning image:.....
<INFO> backend.c:764: 393216 bytes was written from file "/tmp/_tmp_uboot" to
dev 1[0]
```

Here and below, *rupgrade* reports that a new image has been written to the storage device. The number in brackets is the candidate image set number. In this example, the confirmed image number is 1 and the candidate number is 0.

```
Erasing 64 Kibyte @ 30000 -- 75 % complete.
512+0 records in
512+0 records out
262144 bytes (256.0KB) copied, 1.567781 seconds, 163.3KB/s
<INFO> backend.c:730: Burning
image:.....
.....
.....
.....
.....
<INFO> backend.c:764: 1402348 bytes was written from file "/tmp/_tmp_kernel" to
dev 2[0]
<INFO> backend.c:730: Burning
image:.....
.....
.....
.....
.....
```

```
.....
.....
.....
<INFO> backend.c:764: 2283039 bytes was written from file "/tmp/_tmp_rfs" to dev
4[0]
<INFO> shmm700_hal.c:1285: Write storage state: available -> not available
[ALLOW]
<INFO> backend.c:730: Burning
image:.....
.....
.....
.....
.....
.....
<INFO> backend.c:764: 2110743 bytes was written from file "/tmp/_tmp_app" to dev
8[0]
<INFO> backend.c:808: Erasing user data on candidate bank
<INFO> backend.c:818: Copy /etc.
<INFO> shmm700_hal.c:1189: Write confirmed: 1 -> 1 [ALLOW]
<INFO> shmm700_hal.c:1194: Write candidate: 0 -> 0 [ALLOW]
<INFO> shmm700_hal.c:1265: Write upgrade watchdog: 0 -> 1 [ALLOW]
```

Here, the reliable upgrade procedure resets the ShMM. This causes U-Boot to start from the candidate Flash.

```
PowerPrep start initialize power...
Battery Voltage = 1.50V
No battery or bad battery detected!!!.Disabling battery voltage measurements.
Mar 21 201223:23:40
FRAC 0x92926152
memory type is DDR2
Wait for ddr ready 1
power 0x00820616
Frac 0x92926152
start change cpu freq
hbus 0x00000003
cpu 0x00010001
start test memory access
ddr2 0x40000000
finish simple test
U-Boot 2009.08 (Mar 21 2012 - 23:22:30)
Freescale i.MX28 family
CPU: 297 MHz
BUS: 99 MHz
EMI: 130 MHz
GPMI: 24 MHz
I2C: ready
```

```

DRAM: 128 MB
SFGEN: N25Q512A detected, total size 64 MB
A2F: SPICOMM protocol v1.6, M3 firmware v0.9, FPGA design v0.44.0.0
A2F: A2F firmware, version 0.9
A2F: Last reset cause: HARD
A2F: Device type: A2F060M3E-FG256
A2F: MSS clock frequency: 40 MHz
A2F: Fabric clock frequency: 20 MHz
A2F: Fast delay calibration: 1330 cycles per 100uS
A2F: eNVM: 128 KB (00000000 - 00020000)
A2F: eSRAM: 16 KB (20000000 - 20004000)
A2F: Extram start: 200022D0
RUPG: booting from image 0 (candidate)

```

Here, U-Boot reports that reliable upgrade is in progress and the candidate image set is being used. The candidate image set index is 0 and the confirmed image set is 1.

```

In: serial
Out: serial
Err: serial
Net: FEC0: 00:18:49:01:8f:26, FEC1: 00:18:49:01:8f:27
FEC0, FEC1
Hit any key to stop autoboot: 0
RUPG: reliable upgrade is in progress, skipping

```

Here, U-Boot reports that a reliable upgrade is in progress and therefore the SmartFusion cSoC upgrade stage is skipped; it doesn't matter if the SmartFusion cSoC upgrade file is present or not.

```

Creating 1 MTD partitions on "spi0":
0x00000001c0000-0x000002000000 : "mtd=5"

```

Below, U-Boot attaches the UBI partition and mounts the UBIFS volume to load the new kernel and RFS images. This stage is skipped and the messages listed below do not appear during regular boots because a UBIFS cache is used.

```

UBI: attaching mtd1 to ubi0
UBI: MTD device 1 is write-protected, attach in read-only mode
UBI: physical eraseblock size: 65536 bytes (64 KiB)
UBI: logical eraseblock size: 65408 bytes
UBI: smallest flash I/O unit: 1
UBI: VID header offset: 64 (aligned 64)
UBI: data offset: 128
UBI: attached mtd1 to ubi0
UBI: MTD device name: "mtd=5"
UBI: MTD device size: 30 MiB
UBI: number of good PEBs: 484
UBI: number of bad PEBs: 0
UBI: max. allowed volumes: 128

```

```
UBI: wear-leveling threshold:      4096
UBI: number of internal volumes: 1
UBI: number of user volumes:      1
UBI: available PEBs:              0
UBI: total number of reserved PEBs: 484
UBI: number of PEBs reserved for bad PEB handling: 0
UBI: max/mean erase counter: 11/9
UBIFS: read-only UBI device
UBIFS: mounted UBI device 0, volume 0, name "boot"
UBIFS: mounted read-only
UBIFS: file system size:   30807168 bytes (30085 KiB, 29 MiB, 471 LEBs)
UBIFS: journal size:      1569792 bytes (1533 KiB, 1 MiB, 24 LEBs)
UBIFS: media format:      w4/r0 (latest is w4/r0)
UBIFS: default compressor: LZO
UBIFS: reserved for root: 1522516 bytes (1486 KiB)
Loading file 'uImage.0' to addr 0x42000000 with size 1604076 (0x001879ec)...
Done
Loading file 'rfs.0' to addr 0x46000000 with size 2283039 (0x0022d61f)...
Done
## Booting kernel from Legacy Image at 42000000 ...
   Image Name:   Linux-2.6.34.8-shmm700
   Created:      2012-03-26  17:40:37 UTC
   Image Type:   ARM Linux Kernel Image (uncompressed)
   Data Size:    1402284 Bytes =  1.3 MB
   Load Address: 40008000
   Entry Point:  40008000
   Verifying Checksum ... OK
## Loading init Ramdisk from Legacy Image at 46000000 ...
   Image Name:   ShMM700 RFS 3.7.0
   Created:      2017-08-31  18:21:50 UTC
   Image Type:   ARM Linux RAMDisk Image (gzip compressed)
   Data Size:    2282975 Bytes =  2.2 MB
   Load Address: 00000000
   Entry Point:  00000000
   Verifying Checksum ... OK
   Loading Kernel Image ... OK
OK
Starting kernel ...
init started: BusyBox v1.16.2 (2011-11-03 18:19:59 MSK)
/etc/rc: Mounting filesystems...
/etc/rc: Mounted /proc
/etc/rc: Mounted /sys
/etc/rc: Mounted /dev/pts
```

```
/etc/rc: Boot using 0 image set
/etc/rc: Fetching evnironment settings:
/etc/rc: Extracted following params:
/etc/rc: rc2: /etc/rc.shmm700-hpdl
/etc/rc: ipaddr: 192.168.0.22
/etc/rc: ipdevice: eth0
/etc/rc: ipladdr: 192.168.1.2
/etc/rc: ipldevice: usb0
/etc/rc: ip2addr:
/etc/rc: ip2device:
/etc/rc: carrier:
/etc/rc: Checking the reliable upgrade watchdog timer...
<INFO> shmm700_hal.c:1265: Write upgrade watchdog: 1 -> 2 [ALLOW]
activated
```

Here, *rupgrade* is invoked for the first time after the reboot to strobe the reliable upgrade watchdog timer.

```
/etc/rc: Attaching UBI0 (user)
UBI device number 0, total 512 LEBs (33488896 bytes, 31.9 MiB), available 0 LEBs
(0 bytes), LEB size 65408 bytes (63.9 KiB)
/etc/rc: Attach result : 0
/etc/rc: Mount ubi0:user
/etc/rc: Attaching UBI1 (boot)
UBI device number 1, total 484 LEBs (31657472 bytes, 30.2 MiB), available 0 LEBs
(0 bytes), LEB size 65408 bytes (63.9 KiB)
/etc/rc: Attach result : 0
/etc/rc: Mount ubi1:boot
/etc/rc: Selecting carrier from the environment:
/etc/rc:
/etc/rc: Extracting shelfman rootfs patch from /0/sentry.shmm700.app
32979+1 records in
4122+1 records out
2110464 bytes (2.0MB) copied, 3.647968 seconds, 565.0KB/s
/etc/rc: Extraction result : 0
/etc/rc: Placed /var/tmp to ram disk
/etc/rc: Setting hostname shmm700
/etc/rc: Started syslogd and klogd
<INFO> shmm700_hal.c:1265: Write upgrade watchdog: 1 -> 2 [ALLOW]
/etc/rc: Strobing the reliable upgrade WDT
```

Here, *rupgrade* is invoked for the second time to strobe the reliable upgrade WDT.

```
/etc/rc: Mounted /0/etc to /etc
/etc/rc: Calling /etc/rc.shmm700-hpdl
/etc/readhwaddr: Board Hardware Address: 0x12
/etc/netconfig: /etc/hosts updated with shmm700 192.168.0.23 entry
```



```
/etc/netconfig: Updating /etc/profile.sentry with IP settings
/etc/netconfig: Starting /bin/inetd...
/etc/netconfig: Starting time synchronization script
/etc/rc.shmm700-hpdl: Updating /etc/profile.sentry with specific settings
/etc/rc.shmm700-hpdl: Starting snmpd...
/etc/rc.shmm700-hpdl: Starting httpd...
/etc/rc.shmm700-hpdl: Starting Shelf Manager ... for carrier type PPS700
/etc/rc.shmm700-hpdl: Command line: shelfman -sf-
/etc/rc.shmm700-hpdl: Strobing the reliable upgrade WDT
<INFO> shmm700_hal.c:1265: Write upgrade watchdog: 1 -> 2 [ALLOW]
```

rupgrade is called to strobe the reliable upgrade WDT once again, just before starting the Shelf Manager (or confirming the upgrade if automatic start is disabled via the environment variable *start_rc2_daemons*).

```
<*> 09:28:46.973 [134] Pigeon Point Shelf Manager ver. 3.7.0. Built on Aug 31
2017 22:20:24
<*> 09:28:47.012 [134] *** Lock log print buffer at 0x1a2d74 ***
<*> 09:28:47.013 [134] *** Pthread lock log print buffer at 0x1a7594 ***
<I> 09:28:47.044 [134] Reading configuration file: /etc/shelfman.conf
shmm700 login: <INFO> shmm700_hal.c:1265: Write upgrade watchdog: 1 -> 2 [ALLOW]
```

The reliable upgrade WDT is strobed once again, just to make sure it won't fire during the confirm procedure.

```
<INFO> shmm700_hal.c:1189: Write confirmed: 1 -> 0 [ALLOW]
```

rupgrade changes the confirmed image set index to the candidate image set index.

```
<INFO> shmm700_hal.c:1202: Write upgrade_state: "in progress" (2) -> "confirmed"
(4) [ALLOW]
<INFO> shmm700_hal.c:1265: Write upgrade watchdog: 1 -> 0 [ALLOW]
```

And finally, *rupgrade* turns the reliable upgrade WDT off. The reliable upgrade is complete.

Finally, the user checks the status of the reliable upgrade, by calling *rupgrade -s*.

```
# rupgrade -s
```

```
<INFO> frontend.c:37: Last upgrade log:
<INFO> shmm700_hal.c:1285: Write storage state: not available -> available
[ALLOW]
<INFO> shmm700_hal.c:369: Resetting UBIFS cache
<INFO> shmm700_hal.c:1913: Clean on defice 1[1]. Leaving file "/dev/mtd2"
untouched.
<INFO> backend.c:730: Burning image:
<INFO> backend.c:764: 393216 bytes was written from file "/tmp/_tmp_uboot" to
dev 1[1]
<INFO> backend.c:730: Burning image:
<INFO> backend.c:764: 1402348 bytes was written from file "/tmp/_tmp_kernel " to
dev 2[1]
<INFO> backend.c:730: Burning image:
<INFO> backend.c:764: 2283039 bytes was written from file "/tmp/_tmp_rfs" to dev
4[1]
```

```
<INFO> shmm700_hal.c:1285: Write storage state: available -> not available
[ALLOW]
<INFO> backend.c:730: Burning image:
<INFO> backend.c:764: 2110743 bytes was written from file "/tmp/_tmp_app" to dev
8[1]
<INFO> backend.c:808: Erasing user data on candidate bank
<INFO> backend.c:821: Copy /etc.
<INFO> shmm700_hal.c:1189: Write confirmed: 0 -> 0 [ALLOW]
<INFO> shmm700_hal.c:1194: Write candidate: 1 -> 1 [ALLOW]
<INFO> shmm700_hal.c:1265: Write upgrade watchdog: 0 -> 1 [ALLOW]
<INFO> shmm700_hal.c:1265: Write upgrade watchdog: 0 -> 1 [ALLOW]
<INFO> shmm700_hal.c:1265: Write upgrade watchdog: 1 -> 2 [ALLOW]
<INFO> shmm700_hal.c:1265: Write upgrade watchdog: 1 -> 2 [ALLOW]
<INFO> shmm700_hal.c:1265: Write upgrade watchdog: 1 -> 2 [ALLOW]
<INFO> shmm700_hal.c:1265: Write upgrade watchdog: 1 -> 2 [ALLOW]
<INFO> shmm700_hal.c:1265: Write upgrade watchdog: 1 -> 2 [ALLOW]
<INFO> shmm700_hal.c:1265: Write upgrade watchdog: 1 -> 2 [ALLOW]
<INFO> shmm700_hal.c:1200: Write confirmed: 1 -> 0 [ALLOW]
<INFO> shmm700_hal.c:1217: Write upgrade_state: "in progress" (2) -> "confirmed"
(4) [ALLOW]
<INFO> shmm700_hal.c:1280: Write upgrade watchdog: 1 -> 0 [ALLOW]
<INFO> frontend.c:46: Last upgrade stages log:
0 (initial): PASSED
1 (before storage mount): PASSED
2 (after storage mount): PASSED
3 (after marking image set invalid): PASSED
4 (after u-boot burn): PASSED
5 (after kernel burn): PASSED
6 (after rootfs burn): PASSED
7 (after storage umount): PASSED
8 (after cleaning previous apps set): PASSED
9 (after writting new apps set): PASSED
10 (after sync): PASSED
11 (after processing var and etc): PASSED
12 (after boot select): PASSED
13 (after reboot): PASSED
14 (after confirm): PASSED
#
```

8.10.2 Example 2

This example shows a reliable upgrade of the application package image only, copying */etc* and */var/nvdata* non-volatile directories from confirmed to candidate locations. The application package image is taken from an FTP server at the IP address 192.168.1.253. The path to the

application package image on the FTP server is `shmm700/sentry.shmm700.app`. The upgrade procedure is started from a telnet session.

Note: Since only the application package image is explicitly updated, the U-Boot, kernel and RFS images are automatically copied from the confirmed locations to the candidate locations.

The ShMM must be able to access the FTP server over the network (that is, its network adapter must be up and configured and a route must exist from the ShMM to the FTP server). In the example below, the ShMM is configured with the network address 192.168.1.123 (which is in the same network with the FTP server):

```
# telnet 192.168.1.123
```

```
Trying 192.168.1.123...
```

```
Connected to 192.168.1.123.
```

```
Escape character is '^]'.
```

```
shmm700 login: root
```

```
Password:
```

```
#
```

The parameters to `rupgrade` indicate that only the application package is being upgraded and that the copy protocol is FTP, accessing a specified IP address and file, with the user name `admin` and password `ADMINPWD`.

```
# rupgrade --skip-checksums --copy-etc -a
```

```
ftp://admin:ADMINPWD@192.168.1.253/shmm700/sentry.shmm700.app
```

```
<INFO> frontend.c:449: Downloading file "ftp://admin:ADMINPWD
@192.168.1.253/shmm700/sentry.shmm700.app"
```

```
<INFO> frontend.c:732: Initiating partial upgrade:
```

```
<INFO> frontend.c:765:  * app "/tmp/_shmm700_sentry.shmm700.app"
```

```
<INFO> frontend.c:784:  + copy etc
```

```
<INFO> frontend.c:820: Calling backend to handle partial upgrade "/sbin/backend
--application /tmp/_shmm700_sentry.shmm700.app --copy-etc"
```

```
<WARN> shmm700_hal.c:322: Device "ubil:boot" is mounted to folder "/boot".
```

```
<INFO> shmm700_hal.c:1255: Write storage state: not available -> available
[ALLOW]
```

```
<INFO> shmm700_hal.c:368: Resetting UBIFS cache
```

```
Erasing 64 Kibyte @ 30000 -- 75 % complete.
```

```
<INFO> shmm700_hal.c:1863: Clone:.....
```

```
<INFO> shmm700_hal.c:1898: File "/dev/mtd0" of device 1[0] cloned to file
"/dev/mtd2" of device 1[1].
```

```
<INFO> shmm700_hal.c:1642: Cleaning image "/var/upgrade/boot/uImage.1"
```

```
<INFO> shmm700_hal.c:1863:
```

```
Clone:.....
.....
.....
.....
.....
```

[illegible]

At this point, the telnet session is closed after a certain inactivity period; after several seconds, it is possible to reconnect to the target again and check the status of the reliable upgrade by invoking `rupgrade -s`.

```
# telnet 192.168.1.123
```

Trying 192.168.1.123...

Connected to 192.168.1.123.

Escape character is '^|'.

```
shmm700 login: root
```

Password:

```
# rupgrade -s
```

```
<INFO> frontend.c:87: Last upgrade log:
```

```
<WARN> shmm700 hal.c:322: Device "ubi1:boot" is mounted to folder "/boot".
```

```
<INFO> shmm700_hal.c:1255: Write storage state: not available -> available
[ALLOW]
```

```
<INFO> shmm700_hal.c:368: Resetting UBIFS cache
<INFO> shmm700_hal.c:1863: Clone:.....
<INFO> shmm700_hal.c:1898: File "/dev/mtd0" of device 1[0] cloned to file
"/dev/mtd2" of device 1[1].
<INFO> shmm700_hal.c:1642: Cleaning image "/var/upgrade/boot/uImage.1"
<INFO> shmm700_hal.c:1863:
Clone:.....
.....
.....
.....
.....
<INFO> shmm700_hal.c:1898: File "/var/upgrade/boot/uImage.0" of device 2[0]
cloned to file "/var/upgrade/boot/uImage.1" of device 2[1].
<INFO> shmm700_hal.c:1642: Cleaning image "/var/upgrade/boot/rfs.1"
<INFO> shmm700_hal.c:1863:
Clone:.....
.....
.....
.....
.....
.....
.....
.....
.....
.....
<INFO> shmm700_hal.c:1898: File "/var/upgrade/boot/rfs.0" of device 4[0] cloned
to file "/var/upgrade/boot/rfs.1" of device 4[1].
<INFO> shmm700_hal.c:1255: Write storage state: available -> not available
[ALLOW]
<INFO> shmm700_hal.c:1642: Cleaning image "/1/sentry.shmm700.app"
<INFO> backend.c:772: Burning image:<INFO> backend.c:806: 2377014 bytes were
written from file "/tmp/_shmm700_sentry.shmm700.app" to dev 8[1]
<INFO> backend.c:850: Erasing user data on candidate bank
<INFO> backend.c:860: Copy /etc.
<INFO> shmm700_hal.c:1161: Write confirmed: 0 -> 0 [ALLOW]
<INFO> shmm700_hal.c:1166: Write candidate: 1 -> 1 [ALLOW]
<INFO> shmm700_hal.c:1235: Write upgrade watchdog: 0 -> 1 [ALLOW]
<INFO> shmm700_hal.c:1235: Write upgrade watchdog: 1 -> 2 [ALLOW]
<INFO> shmm700_hal.c:1161: Write confirmed: 0 -> 1 [ALLOW]
<INFO> shmm700_hal.c:1174: Write upgrade_state: "in progress" (2) -> "confirmed"
(4) [ALLOW]
<INFO> shmm700_hal.c:1235: Write upgrade watchdog: 1 -> 0 [ALLOW]
<INFO> frontend.c:96: Last upgrade stages log:
0 (initial): PASSED
1 (before storage mount): PASSED
2 (after storage mount): PASSED
3 (after marking image set invalid): PASSED
4 (after u-boot burn): PASSED
5 (after kernel burn): PASSED
```

```

6 (after rootfs burn): PASSED
7 (after storage umount): PASSED
8 (after cleaning previous apps set): PASSED
9 (after writting new apps set): PASSED
10 (after sync): PASSED
11 (after processing var and etc): PASSED
12 (after boot select): PASSED
13 (after reboot): PASSED
14 (after confirm): PASSED

```

8.10.3 Example 3

This example shows an unsuccessful reliable upgrade. Power is turned off after the boot from the candidate locations, but before the reliable upgrade is finalized. After turning the power back on, the rollback to the confirmed images occurs. This reliable upgrade is initiated from the serial console. The default copy mode is used (copy SSH keys and `/var/nvdata`). All four images are assumed to be already in `/tmp`.

```

# rupgrade --skip-checksums -k file:///tmp/sentry.kernel -r
file:///tmp/sentry.rfs -u file:///tmp/u-boot.bin -a file:///tmp/sentry.app
<INFO> frontend.c:449: Downloading file "file:///tmp/u-boot.bin"
<INFO> frontend.c:449: Downloading file "file:///tmp/sentry.kernel"
<INFO> frontend.c:449: Downloading file "file:///tmp/sentry.rfs"
<INFO> frontend.c:449: Downloading file "file:///tmp/sentry.app"
<INFO> frontend.c:535: Initiating partial upgrade:
<INFO> frontend.c:735:  * uboot "/tmp/_tmp_uboot"
<INFO> frontend.c:745:  * kernel "/tmp/_tmp_kernel"
<INFO> frontend.c:755:  * rfs "/tmp/_tmp_rfs"
<INFO> frontend.c:765:  * app "/tmp/_tmp_app"
<INFO> frontend.c:820: Calling backend to handle partial upgrade "/sbin/backend
--uboot /tmp/_tmp_uboot --kernel /tmp/_tmp_kernel --rootfs /tmp/_tmp_rfs --
application /tmp/_tmp_app"
<INFO> shmm700_hal.c:1285: Write storage state: not available -> available
[ALLOW]
UBI device number 1, total 484 LEBs (31657472 bytes, 30.2 MiB), available 0 LEBs
(0 bytes), LEB size 65408 bytes (63.9 KiB)
<INFO> shmm700_hal.c:369: Resetting UBIFS cache
Erasing 64 Kibyte @ 30000 -- 75 % complete.
<INFO> shmm700_hal.c:1913: Clean on defice 1[1]. Leaving file "/dev/mtd2"
untouched.
<INFO> backend.c:730: Burning image:.....
<INFO> backend.c:764: 393216 bytes was written from file "/tmp/_tmp_uboot" to
dev 1[1]
Erasing 64 Kibyte @ 30000 -- 75 % complete.
512+0 records in
512+0 records out
262144 bytes (256.0KB) copied, 1.571406 seconds, 162.9KB/s

```

[illegible]

The reliable upgrade procedure resets the ShMM here and starts U-Boot from the candidate Flash.

```
PowerPrep start initialize power...
```

Battery Voltage = 1.46V

No battery or bad battery detected!!!.Disabling battery voltage measurements.

Mar 21 201223:23:40

```
FRAC  0x92926152
```

```
memory type is DDR2
```

```
Wait for ddr ready 1
```

```
power 0x00820616
```

Frac 0x92926152

```
start change cpu freq
```

```
hbus 0x00000003
```

```
cpu 0x00010001
```

```
start test memory access
ddr2 0x40000000
finish simple test
U-Boot 2009.08 (Mar 21 2012 - 23:22:30)
Freescale i.MX28 family
CPU: 297 MHz
BUS: 99 MHz
EMI: 130 MHz
GPMI: 24 MHz
I2C: ready
DRAM: 128 MB
SFGEN: N25Q512A detected, total size 64 MB
A2F: SPICOMM protocol v1.6, M3 firmware v0.9, FPGA design v0.44.0.0
A2F: A2F firmware, version 0.9
A2F: Last reset cause: HARD
A2F: Device type: A2F060M3E-FG256
A2F: MSS clock frequency: 40 MHz
A2F: Fabric clock frequency: 20 MHz
A2F: Fast delay calibration: 1330 cycles per 100uS
A2F: eNVM: 128 KB (00000000 - 00020000)
A2F: eSRAM: 16 KB (20000000 - 20004000)
A2F: Extram start: 200022D0
RUPG: booting from image 1 (candidate)
```

Here, U-Boot reports that a reliable upgrade is in progress and the candidate image set is being used. The candidate image set index is 1 and the confirmed image set is 0.

```
In: serial
Out: serial
Err: serial
Net: FEC0: 00:18:49:01:8f:26, FEC1: 00:18:49:01:8f:27
FEC0, FEC1
Hit any key to stop autoboot: 2
```

Power is turned off here. After some time, power is turned back on. The candidate Flash state has been lost because of the power loss, so the ShMM reverts back to the confirmed Flash.

```
PowerPrep start initialize power...
Battery Voltage = 1.32V
No battery or bad battery detected!!!.Disabling battery voltage measurements.
Mar 21 201223:23:40
FRAC 0x92926152
memory type is DDR2
Wait for ddr ready 1
power 0x00820616
Frac 0x92926152
```



```
start change cpu freq
hbus 0x00000003
cpu 0x00010001
start test memory access
ddr2 0x40000000
finish simple test
U-Boot 2009.08 (Mar 21 2012 - 23:22:30)
Freescale i.MX28 family
CPU: 297 MHz
BUS: 99 MHz
EMI: 130 MHz
GPMI: 24 MHz
I2C: ready
DRAM: 128 MB
SFGEN: N25Q512A detected, total size 64 MB
A2F: SPICOMM protocol v1.6, M3 firmware v0.9, FPGA design v0.44.0.0
A2F: A2F firmware, version 0.9
A2F: Last reset cause: HARD
A2F: Device type: A2F060M3E-FG256
A2F: MSS clock frequency: 40 MHz
A2F: Fabric clock frequency: 20 MHz
A2F: Fast delay calibration: 1330 cycles per 100uS
A2F: eNVM: 128 KB (00000000 - 00020000)
A2F: eSRAM: 16 KB (20000000 - 20004000)
A2F: Extram start: 200022D0
RUPG: booting from image 0 (confirmed)
```

After a power cycle, U-Boot reports that there is no reliable upgrade in progress and the confirmed image set is being used. The confirmed image set index is 0.

```
In: serial
Out: serial
Err: serial
Net: FEC0: 00:18:49:01:8f:26, FEC1: 00:18:49:01:8f:27
FEC0, FEC1
Hit any key to stop autoboot: 0
Creating 1 MTD partitions on "spi0":
0x00000001c0000-0x000002000000 : "mtd=5"
UBI: attaching mtd1 to ubi0
UBI: MTD device 1 is write-protected, attach in read-only mode
UBI: physical eraseblock size: 65536 bytes (64 KiB)
UBI: logical eraseblock size: 65408 bytes
UBI: smallest flash I/O unit: 1
UBI: VID header offset: 64 (aligned 64)
UBI: data offset: 128
```

```

UBI: attached mtd1 to ubi0
UBI: MTD device name:          "mtd=5"
UBI: MTD device size:          30 MiB
UBI: number of good PEBs:      484
UBI: number of bad PEBs:       0
UBI: max. allowed volumes:     128
UBI: wear-leveling threshold:  4096
UBI: number of internal volumes: 1
UBI: number of user volumes:    1
UBI: available PEBs:           0
UBI: total number of reserved PEBs: 484
UBI: number of PEBs reserved for bad PEB handling: 0
UBI: max/mean erase counter: 12/10
UBIFS: read-only UBI device
UBIFS: mounted UBI device 0, volume 0, name "boot"
UBIFS: mounted read-only
UBIFS: file system size:      30807168 bytes (30085 KiB, 29 MiB, 471 LEBs)
UBIFS: journal size:          1569792 bytes (1533 KiB, 1 MiB, 24 LEBs)
UBIFS: media format:          w4/r0 (latest is w4/r0)
UBIFS: default compressor: LZO
UBIFS: reserved for root: 1522516 bytes (1486 KiB)
RUPG: upgrade file is not present, skipping
Loading file 'uImage.0' to addr 0x42000000 with size 1604076 (0x001879ec)...
Done
Loading file 'rfs.0' to addr 0x46000000 with size 2283039 (0x0022d61f)...
Done
## Booting kernel from Legacy Image at 42000000 ...
   Image Name:   Linux-2.6.34.8-shmm700
   Created:      2012-03-26  17:40:37 UTC
   Image Type:   ARM Linux Kernel Image (uncompressed)
   Data Size:    1402284 Bytes =  1.3 MB
   Load Address: 40008000
   Entry Point:  40008000
   Verifying Checksum ... OK
## Loading init Ramdisk from Legacy Image at 46000000 ...
   Image Name:   ShMM700 RFS 3.7.0
   Created:      2017-08-31  18:21:50 UTC
   Image Type:   ARM Linux RAMDisk Image (gzip compressed)
   Data Size:    2282975 Bytes =  2.2 MB
   Load Address: 00000000
   Entry Point:  00000000
   Verifying Checksum ... OK
   Loading Kernel Image ... OK

```

OK

Starting kernel ...

init started: BusyBox v1.16.2 (2011-11-03 18:19:59 MSK)

/etc/rc: Mounting filesystems...

/etc/rc: Mounted /proc

/etc/rc: Mounted /sys

/etc/rc: Mounted /dev/pts

/etc/rc: Boot using 0 image set

/etc/rc: Fetching environment settings:

/etc/rc: Extracted following params:

/etc/rc: rc2: /etc/rc.shmm700-hpdl

/etc/rc: ipaddr: 192.168.0.22

/etc/rc: ipdevice: eth0

/etc/rc: ipladdr: 192.168.1.2

/etc/rc: ipldevice: usb0

/etc/rc: ip2addr:

/etc/rc: ip2device:

/etc/rc: carrier:

/etc/rc: Checking the reliable upgrade watchdog timer...activated

/etc/rc: Attaching UBI0 (user)

UBI device number 0, total 512 LEBs (33488896 bytes, 31.9 MiB), available 0 LEBs (0 bytes), LEB size 65408 bytes (63.9 KiB)

/etc/rc: Attach result : 0

/etc/rc: Mount ubi0:user

/etc/rc: Attaching UBI1 (boot)

UBI device number 1, total 484 LEBs (31657472 bytes, 30.2 MiB), available 0 LEBs (0 bytes), LEB size 65408 bytes (63.9 KiB)

/etc/rc: Attach result : 0

/etc/rc: Mount ubi1:boot

/etc/rc: Selecting carrier from the environment:

/etc/rc:

/etc/rc: Extracting shelfman rootfs patch from /0/sentry.shmm700.app

32979+1 records in

4122+1 records out

2110464 bytes (2.0MB) copied, 12.323593 seconds, 167.2KB/s

/etc/rc: Extraction result : 0

/etc/rc: Placed /var/tmp to ram disk

/etc/rc: Setting hostname shmm700

/etc/rc: Started syslogd and klogd

/etc/rc: Strobing the reliable upgrade WDT

/etc/rc: Mounted /0/etc to /etc

/etc/rc: Calling /etc/rc.shmm700-hpdl

/etc/readhwaddr: Board Hardware Address: 0x12

```
/etc/netconfig: /etc/hosts has valid shmm700 192.168.0.23 entry
/etc/netconfig: Updating /etc/profile.sentry with IP settings
/etc/netconfig: Starting /bin/inetd...
/etc/netconfig: Starting time synchronization script
/etc/rc.shmm700-hpdl: Updating /etc/profile.sentry with specific settings
/etc/rc.shmm700-hpdl: Starting snmpd...
/etc/rc.shmm700-hpdl: Starting httpd...
/etc/rc.shmm700-hpdl: Starting Shelf Manager ... for carrier type PPS700
/etc/rc.shmm700-hpdl: Command line: shelfman -sf-
/etc/rc.shmm700-hpdl: Strobing the reliable upgrade WDT
<*> 11:08:11.872 [122] Pigeon Point Shelf Manager ver. 3.7.0. Built on Aug 31
2017 22:20:24
<*> 11:08:11.913 [122] *** Lock log print buffer at 0x1a2d74 ***
<*> 11:08:11.915 [122] *** Pthread lock log print buffer at 0x1a7594 ***
<I> 11:08:11.963 [122] Reading configuration file: /etc/shelfman.conf
```

Finally, the user checks the status of the reliable upgrade, by calling *rupgrade -s*.

rupgrade -s

```
<INFO> frontend.c:37: Last upgrade log:
<INFO> shmm700_hal.c:1285: Write storage state: not available -> available
[ALLOW]
<INFO> shmm700_hal.c:369: Resetting UBIFS cache
<INFO> shmm700_hal.c:1913: Clean on defice 1[1]. Leaving file "/dev/mtd2"
untouched.
<INFO> backend.c:730: Burning image:
<INFO> backend.c:764: 393216 bytes was written from file "/tmp/file____tmp_u-
boot.bin" to dev 1[1]
<INFO> backend.c:730: Burning image:
<INFO> backend.c:764: 1402348 bytes was written from file
"/tmp/file____tmp_sentry.kernel" to dev 2[1]
<INFO> backend.c:730: Burning image:
<INFO> backend.c:764: 2283039 bytes was written from file
"/tmp/file____tmp_sentry.rfs" to dev 4[1]
<INFO> shmm700_hal.c:1285: Write storage state: available -> not available
[ALLOW]
<INFO> backend.c:730: Burning image:
<INFO> backend.c:764: 2110743 bytes was written from file
"/tmp/file____tmp_sentry.app" to dev 8[1]
<INFO> backend.c:808: Erasing user data on candidate bank
<INFO> backend.c:821: Copy ssh keys
<INFO> shmm700_hal.c:1189: Write confirmed: 0 -> 0 [ALLOW]
<INFO> shmm700_hal.c:1194: Write candidate: 1 -> 1 [ALLOW]
```

The log contains no 'confirm' message; the confirmed and candidate image set numbers are the same they were before reliable upgrade initiation.

```
<INFO> shmm700_hal.c:1265: Write upgrade watchdog: 0 -> 1 [ALLOW]
```

```
<INFO> frontend.c:46: Last upgrade stages log:
```

```
0 (initial): PASSED
1 (before storage mount): PASSED
2 (after storage mount): PASSED
3 (after marking image set invalid): PASSED
4 (after u-boot burn): PASSED
5 (after kernel burn): PASSED
6 (after rootfs burn): PASSED
7 (after storage umount): PASSED
8 (after cleaning previous apps set): PASSED
9 (after writting new apps set): PASSED
10 (after sync): PASSED
11 (after processing var and etc): PASSED
12 (after boot select): PASSED
13 (after reboot): FAILED
14 (after confirm): FAILED
```

The last stages are marked as 'FAILED' in the reliable upgrade status. The ShMM has automatically rolled back to the confirmed firmware because the upgrade failed.

8.10.4 Example 4

This example shows a reliable upgrade where image installation and activation are separated. Initially, the new images are installed into the other image set via *rupgrade* with the option *--install-only*. Later, the other image set is activated via the command *backend --activate*.

The default copy mode is used (copy SSH keys and */var/nvdata*). All four images are assumed to be already in */tmp*.

```
# rupgrade --install-only --skip-checksums -k file:///tmp/sentry.kernel -r
file:///tmp/sentry.rfs -u file:///tmp/u-boot.bin -a file:///tmp/sentry.app
144: Executing: /sbin/backend --get-wdt-state
144: System returned 0, exec result 0.
<INFO> Downloading file "file:///tmp/u-boot.bin"
<INFO> Downloading file "file:///tmp/sentry.kernel"
<INFO> Downloading file "file:///tmp/sentry.rfs"
<INFO> Downloading file "file:///tmp/sentry.app"
<INFO> Initiating partial upgrade:
<INFO> * uboot "/tmp/u-boot.bin"
<INFO> * kernel "/tmp/sentry.kernel"
<INFO> * rfs "/tmp/sentry.rfs"
<INFO> * app "/tmp/sentry.app"
<INFO> + install only
<INFO> Calling backend to handle partial upgrade "/sbin/backend --uboot /tmp/u-
boot.bin --kernel /tmp/sentry.kernel --rootfs /tmp/sentry.rfs --application
/tmp/sentry.app --install-only"
```

```

144: Executing: /sbin/backend --uboot /tmp/u-boot.bin --kernel
/tmp/sentry.kernel --rootfs /tmp/sentry.rfs --application /tmp/sentry.app --
install-only
148: Executing: /etc/hook/base_script.sh --stage 1 --silent
148: System returned 0, exec result 0.
<INFO> Write storage state: not available -> available [ALLOW]
148: Executing: mkdir -p /var/upgrade/boot
148: System returned 0, exec result 0.
<INFO> Resetting UBIFS cache
148: Executing: /sbin/eraseall /dev/mtd8
Erasing 256 Kibyte @ 0 -- 0 % complete.
148: System returned 0, exec result 0.
148: Executing: /etc/hook/base_script.sh --stage 2 --silent
148: System returned 0, exec result 0.
148: Executing: /etc/hook/base_script.sh --stage 3 --silent
148: System returned 0, exec result 0.
<INFO> Burning image:...
<INFO> 524288 bytes were written from file "/tmp/u-boot.bin" to dev 1[1]
<INFO> U-Boot checksum test after writing image: succeeded
148: Executing: eraseall /dev/mtd6 >/dev/null 2>/dev/null
148: System returned 0, exec result 0.
148: Executing: dd if=/dev/mtd4 of=/dev/mtd6 >/dev/null 2>/dev/null
148: System returned 0, exec result 0.
148: Executing: /etc/hook/base_script.sh --stage 4 --silent
148: System returned 0, exec result 0.
<INFO> Cleaning image "/var/upgrade/boot/uImage.1"
<INFO> Burning
image:.....
.....
.....
.....
<INFO> 1531460 bytes were written from file "/tmp/sentry.kernel" to dev 2[1]
<INFO> Kernel checksum test after writing image: succeeded
148: Executing: /etc/hook/base_script.sh --stage 5 --silent
148: System returned 0, exec result 0.
<INFO> Cleaning image "/var/upgrade/boot/rfs.1"
<INFO> Burning
image:.....
.....
.....
.....

```



```
228: System returned 0, exec result 0.
228: Executing: dd if=/dev/mtd8 of=/dev/null >/dev/null 2>/dev/null
228: System returned 0, exec result 0.
<INFO> Write confirmed: 0 -> 0 [ALLOW]
<INFO> Write candidate: 1 -> 1 [ALLOW]
<INFO> Write upgrade watchdog: 0 -> 1 [ALLЯ
```

```
U-Boot 2009.08 (Sep 28 2016 - 14:45:44)
Freescale i.MX28 family
CPU: 297 MHz
BUS: 99 MHz
EMI: 130 MHz
GPMI: 24 MHz
I2C: ready
DRAM: 128 MB
SFGEN: S25FL512S detected, total size 64 MB
A2F: SPICOMM protocol v1.7, M3 firmware v1.5, FPGA design v1.0.1.0
A2F: A2F firmware, version 1.5
A2F: Device type: A2F060M3E-FG256
A2F: MSS clock frequency: 40 MHz
A2F: Fabric clock frequency: 20 MHz
A2F: Fast delay calibration: 1330 cycles per 100uS
A2F: eNVM: 128 KB (00000000 - 00020000)
A2F: eSRAM: 16 KB (20000000 - 20004000)
A2F: Extram start: 20002750
A2F: Last reset cause: RUWD
A2F: VMON flags: 00000000
RUPG: booting from image 1 (candidate)
In: serial
Out: serial
Err: serial
Net: FEC0: 00:18:49:01:8f:78, FEC1: 00:18:49:01:8f:79
FEC0, FEC1
Hit any key to stop autoboot: 0
RUPG: reliable upgrade is in progress, skipping
Creating 1 MTD partitions on "spi0":
0x00000001c0000-0x000002000000 : "mtd=5"
UBI: attaching mtd1 to ubi0
UBI: MTD device 1 is write-protected, attach in read-only mode
UBI: physical eraseblock size: 262144 bytes (256 KiB)
UBI: logical eraseblock size: 262016 bytes
UBI: smallest flash I/O unit: 1
UBI: VID header offset: 64 (aligned 64)
```



```

UBI: data offset:                128
UBI: attached mtd1 to ubi0
UBI: MTD device name:            "mtd=5"
UBI: MTD device size:            30 MiB
UBI: number of good PEBs:        121
UBI: number of bad PEBs:         0
UBI: max. allowed volumes:       128
UBI: wear-leveling threshold:    4096
UBI: number of internal volumes: 1
UBI: number of user volumes:     1
UBI: available PEBs:             0
UBI: total number of reserved PEBs: 121
UBI: number of PEBs reserved for bad PEB handling: 0
UBI: max/mean erase counter: 598/527
UBIFS: read-only UBI device
UBIFS: mounted UBI device 0, volume 0, name "boot"
UBIFS: mounted read-only
UBIFS: file system size: 28297728 bytes (27634 KiB, 26 MiB, 108 LEBs)
UBIFS: journal size: 2096129 bytes (2047 KiB, 1 MiB, 6 LEBs)
UBIFS: media format: w4/r0 (latest is w4/r0)
UBIFS: default compressor: LZO
UBIFS: reserved for root: 1398497 bytes (1365 KiB)
Loading file 'uImage.1' to addr 0x42000000 with size 1531460 (0x00175e44)...
Done
Loading file 'rfs.1' to addr 0x46000000 with size 2684957 (0x0028f81d)...
Done
## Booting kernel from Legacy Image at 42000000 ...
   Image Name:   Linux-2.6.34.8-shmm700
   Created:      2016-09-28  19:45:42 UTC
   Image Type:   ARM Linux Kernel Image (uncompressed)
   Data Size:    1531396 Bytes =  1.5 MB
   Load Address: 40008000
   Entry Point:  40008000
   Verifying Checksum ... OK
## Loading init Ramdisk from Legacy Image at 46000000 ...
   Image Name:   ShMM700 RFS 3.7.0
   Created:      2017-08-27  19:40:32 UTC
   Image Type:   ARM Linux RAMDisk Image (gzip compressed)
   Data Size:    2684893 Bytes =  2.6 MB
   Load Address: 00000000
   Entry Point:  00000000
   Verifying Checksum ... OK
   Loading Kernel Image ... OK

```

OK

Starting kernel ...

```
init started: BusyBox v1.19.4 (2016-09-28 14:39:49 CDT)
/etc/rc: Current image set is 1
/etc/rc: Mounting filesystems...
/etc/rc: Mounted /proc
/etc/rc: Mounted /dev/pts
/etc/rc: Mounted /tmp
/etc/rc: Mounted /sys
/etc/rc: Checking the reliable upgrade watchdog timer...activated
<INFO> Write upgrade watchdog: 1 -> 2 [ALLOW]
/etc/rc: Attaching ubi0:user...done
/etc/rc: Mounting ubi0:user...done
/etc/rc: Attaching ubi1:boot...done
/etc/rc: Mounting ubi1:boot...done
<INFO> Write upgrade watchdog: 1 -> 2 [ALLOW]
/etc/rc: Extracting application package /boot/app.1...done
/etc/rc: Mounted /1/var to /var
/etc/rc: Strobing the reliable upgrade WDT
<INFO> Write upgrade watchdog: 1 -> 2 [ALLOW]
92: Executing: /etc/hook/base_script.sh --stage 13 --silent
92: System returned 0, exec result 0.
/etc/rc: /etc appears to be empty ... restoring from /etc in RFS...
/etc/rc: Preserved ssh_host_dsa_key in /etc/ssh
/etc/rc: Preserved ssh_host_rsa_key in /etc/ssh
/etc/rc: Mounted /1/etc to /etc
/etc/rc: Setting timezone to MSK-3
/etc/rc: Placed /var/tmp to ram disk
/etc/rc: Setting hostname shmm~120
/etc/rc: Started syslogd and klogd
/etc/rc: Calling /etc/rc.carrier7-hpdl
/etc/readhwaddr: Board Hardware Address: 0x10
/etc/netconfig: /etc/hosts updated with shmm~120 192.168.1.120 entry
/etc/netconfig: Updating /etc/profile.sentry with IP settings
/etc/netconfig: ifconfig eth0 192.168.1.120 netmask 255.255.255.0
/etc/netconfig: ifconfig usb0 192.168.0.120 netmask 255.255.255.0
/etc/netconfig: ifconfig eth1 172.16.2.120 netmask 255.255.255.0
/etc/netconfig: route add default gw 192.168.1.253
/etc/netconfig: Starting /bin/inetd...
/etc/netconfig: Starting time synchronization script
/etc/rc.carrier7-hpdl: Updating /etc/profile.sentry with specific settings
```

```
/etc/rc.carrier7-hpdl: RC2 daemons not started by request
/etc/rc: start_rc2_daemons=n - finishing reliable upgrade process now.
<INFO> Write confirmed: 0 -> 1 [ALLOW]
<INFO> Write upgrade_state: "in progress" (2) -> "confirmed" (4) [ALLOW]
<INFO> Write upgrade watchdog: 1 -> 0 [ALLOW]

shmm~120 login:
```

9 HPI-based Shelf Manager upgrades

In addition to the reliable upgrade facility, a Shelf Manager with IntegralHPI also supports upgrades over the HPI interface. HPI defines a special type of management instruments used for upgrades and called Firmware Upgrade Management Instruments (FUMIs). IntegralHPI creates a FUMI for each resource that represents a physical Shelf Manager. In a redundant configuration, there are two Shelf Manager-related FUMIs: one for the active and one for the backup Shelf Manager.

The HPI paradigm for FUMIs is defined in the HPI specification, and consists of a series of function calls to select the upgrade image, download the image to the target resource and activate the new image. Also the FUMI API includes optional support for manual and automatic rollbacks to the previous state. Shelf Manager FUMIs implement all functions defined for the “logical bank” model (except for `saHpiFumiAutoRollbackDisableSet()`); the “explicit bank” model is not supported for these FUMIs.

An image used for HPI-based Shelf Manager upgrades must be in HPM.1 format. This format allows creating a single image that consists of multiple components, with a special header. An HPM.1 Shelf Manager image includes, as HPM.1 components, the normal U-Boot, kernel and RFS images used for reliable upgrades. On ShMM-500/1500, the HPM.1 image consists only of these three components. On ShMM-700, the HPM.1 image also includes the application package and therefore, consists of four components. Starting with release 2.8.0, Shelf Manager images in HPM.1 format are available to ShMM-based shelf developers from a secure partner page, in addition to the traditional U-Boot, kernel and RFS images.

The implementation of these Shelf Manager FUMIs is done on top of the reliable upgrade facility. Once the image is downloaded to a physical ShMM, a call to activate the image invokes the reliable upgrade utility.

On ShMM-500/1500, the reliable upgrade utility places the embedded U-Boot, kernel and RFS images into the provisional Flash bank and makes the provisional Flash bank active. The utility is called with no `hook` option, so that the default mode of copying non-volatile data applies (only the directory `/var/nvdata` and SSH key files are copied).

On ShMM-700, the reliable upgrade utility places the embedded U-Boot, kernel, RFS images and application package into the candidate locations and makes the candidate image set active. The default mode of copying non-volatile data applies (only the directory `/var/nvdata` and SSH key files are copied).

As with the reliable upgrade facility, nVent recommends upgrading the backup Shelf Manager, then the active Shelf Manager. After the upgrade of the active Shelf Manager, a switchover occurs and the roles of the physical Shelf Managers (active vs. backup) become reversed. If it is desired to preserve the original role assignment, an additional switchover is needed; this can be done via the HPI Shelf Manager Failover control associated with the Virtual Shelf Manager resource.

See the *Pigeon Point HPI User's Guide* for more information and examples of the HPI-based Shelf Manager upgrade procedure.

Appendix A Converting to ShMM-700-based Shelf Managers in ShMM-500 shelves

The ShMM-700 supports converting a live shelf from ShMM-500-based management to ShMM-700-based management, without losing manageability of the shelf. This support minimizes the risk of upgrading shelves in the field to the ShMM-700.

It is important to understand the special aspects of this support:

- This conversion is only supported in one direction: going to ShMM-700-based management.
- Concurrent presence in the shelf of both types of ShMMs is only supported for transition purposes. Stable dual redundant management requires a pair of ShMMs from the same family.
- This functionality is only supported if the ShMM-500 uses a USB-based Software Redundancy Interface.
- The *CARRIER* Shelf Manager configuration parameter must be set to the same value on the ShMM-500 and ShMM-700 carriers.
- The IPMB-0 topology must be bused.

Please contact nVent if any of the following situations apply in a prospective scenario for converting a live shelf from ShMM-500-based management to ShMM-700-based management:

- An Ethernet-based Software Redundancy Interface is used.
- The *CARRIER* configuration parameter cannot be set to the same value on the ShMM-500 and ShMM-700 carriers.
- The shelf supports a Radial IPMB-0 topology.

The switchover procedure is not much different from a standard ShMM-500 to ShMM-500 or ShMM-700 to ShMM-700 switchover procedure, and is described below (where ShMM-500/700 carriers are referenced for physical board removals and insertions):

1. Copy (e.g. FTP) the new Shelf FRU Info file (for the ShMM-700 based shelf) into the directory */var/nvdata*.
2. Remove the backup ShMM-500 carrier.
3. Insert a ShMM-700 carrier into the vacant slot.
4. Wait until it comes up as a backup Shelf Manager.
5. Install the new Shelf FRU Info file via *clia frudataw 20 254 /var/nvdata/<FRU file>*
6. Issue a *clia switchover -force* command on the ShMM-700.
7. The ShMM-700 should immediately become active, and the ShMM-500R will automatically reboot. The ShMM-500 will keep rebooting without becoming a backup Shelf Manager, and can be removed at any time.
8. Remove the second ShMM-500 carrier.
9. After removing the ShMM-500 carrier, insert a second ShMM-700 carrier into its slot. It should come up as a backup Shelf Manager. Do a switchover. The transition procedure is now complete.

Note: As implied above, a switchover from a ShMM-700 back to a ShMM-500 is NOT supported. Once a ShMM-700 becomes active, the only supported action is to remove the second ShMM-500 carrier. Also, a ShMM-500 carrier should never be inserted into a shelf with a running ShMM-700.

Appendix B Customer support

If you are having problems with the Pigeon Point Shelf Manager or ShMM product, please contact your supplier for the Pigeon Point products with questions and problem reports.

If you have any questions about direct purchase of Pigeon Point products (one of the *Pigeon Point Board Management Reference* designs, for instance), please contact nVent; for guidance on such contacts, visit www.pigeonpoint.com/contact.html.

Appendix C Revision history

This section lists the changes that have been made in each revision of this document, starting with the 2.1.0 release.

C.1 Release 2.1.0

Section 7.3: corrects the boundaries of the Flash partitions maintained on the ShMM-500 by FOSL for 16MB Flash devices.

- Section 7.3: covers the boundaries of the Flash partitions maintained on the ShMM-500 by FOSL for 32MB Flash and 64MB Flash devices.

C.2 Release 2.2.0

- Section 3.4.5: elaborates on the algorithm for computation of IP address for endpoints of USB network interface.
- Section 3.3: introduces new configuration parameters:
ALLOW_ALL_COMMANDS_FROM_IPMB, ALLOW_CHANGE_EVENT_RECEIVER, ALLOW_RESET_STANDALONE, DEFAULT_RMCP_NETMASK, IPMB_LINK_ISOLATION_TIMEOUT, MAX_INCOMING_IPMB_REQUESTS, MAX_OEM_FILTERS, SENSOR_POLL_INTERVAL, SHELF_FRU_IPMB_SOURCE1, SHELF_FRU_IPMB_SOURCE2, SWITCHOVER_ON_HANDLE_OPEN.
- Section 3.5.1: adds an option to limit the search for potential sources of Shelf FRU Information on IPMB-0 to the two well-known IPMB-0 locations defined by the configuration variables *SHELF_FRU_IPMB_SOURCE1, SHELF_FRU_IPMB_SOURCE2.*
- Section 3.8: covers a new option for local sensors to be configured when the Shelf Manager is started. The Sensor Device Records (SDRs) defining these sensors are read from the file */var/nvdata/user_sdr.*
- Section 4.3: provides the details of command-line invocations of the Shelf Manager.
- Section 4.5: describes Shelf Manager operation on radial shelves.

C.3 Release 2.3.0

- Section 3.2.1: corrects references to the U-Boot environment variable *gatewayip* that was previously incorrectly referenced as *gateway.*
- Section 3.4.2.1: explains the configuration parameter *INITIAL_SLOW_LINK_DELAY.*
- Introduces the parallel usage of two network interfaces. In parallel mode, instead of having a single RMCP network address that is switched between the two network interfaces, the Shelf Manager supports RMCP on both interfaces with different IP addresses, as two separate IPMI channels (channels 1 and 2).
- Section 3.4.9: adds an option for IP addresses for the Shelf Manager to be assigned by a DHCP server.

- Section 3.3: introduces the new configuration parameters:
`COOLING_FAN_DECREASE_TIMEOUT`, `COOLING_FAN_INCREASE_TIMEOUT`,
`CPLD_ACTIVE_WORKAROUND`, `DEFAULT_GATEWAY_IP_ADDRESS2`,
`DEFAULT_RMCP_IP_ADDRESS2`, `INITIAL_SLOW_LINK_DELAY`,
`FAN_LEVEL_STEP_DOWN`, `FAN_LEVEL_STEP_UP`, `IPMB_RETRY_TIMEOUT_MSEC`,
`MAX_NODE_BUSY_RETRANSMISSIONS`, `NORMAL_STABLE_TIME`,
`PREFERRED_DHCP_SERVER`, `SYSTEM_MANAGER_TRUNCATES_SEL`, `USE_DHCP`,
`USE_SECOND_CHANNEL`.
- Section 3.9: adds coverage for the Auxiliary Firmware Revision being set when the Shelf Manager is started. It is read from the file `/var/nvdata/aux-fw-revision`.
- Section 3.10.1: adds an option for the Shelf Manager to obtain date and time via the Network Time Protocol.
- Section 4.6: introduces the option for SEL truncation to be done automatically under control of the System Manager.

C.4 Release 2.4.0

- Section 3.4.9: adds a description of the Request Identifier (Request ID) byte of the DHCP Client Identifier, along with coverage of the default inclusion of the Shelf Address in the DHCP Client Identifier.
- Section 3.3: introduces the new configuration parameter `RMCP_WITHOUT_SHELF_FRU`.
- Section 4.4.1: provides a description of the new Redundancy and CPLD sensor that is provided by both active and backup Shelf Managers.
- Section 7.3: deletes the table describing the Flash partitions for a ShMM-500 with 32MB Flash devices; this configuration is not offered.

C.5 Release 2.4.1

- Section 3.2.1: changes the default value of the U-Boot environment variable `hostname` to `shmm300` or `shmm500`, based on the ShMM type.
- Section 3.2.4: a new section that describes how to establish the secondary RC script.
- Section 3.3: contains the content of the former section 3.4 “Setting up Shelf Manager Configuration File”.
- Section 3.3: uses the new configuration parameter name `DHCP_SERVER_ADDRESS` instead of the original name `PREFERRED_DHCP_SERVER`.
- Section 3.5.2: adds caveats about using the `eeepromw` utility.
- Section 3.5.3: a new section that describes how to set up the Shelf FRU Information using CLI commands.
- Section 3.11: adds a description of how ShMM POST results are reported using the IPMI command “Get Self Test Results”.

C.6 Release 2.4.2

- Section 3.3.1: adds (in a new subsection) clarifications for the role of a carrier-specific Shelf Manager configuration file in determining the effective values of Shelf Manager configuration variables.

C.7 Release 2.4.4

- Section 3.3.4: adds (in a new subsection) a description of the new configuration parameter *VERBOSITY_CONSOLE*.

C.8 Release 2.5.0

- Section 3.3: adds new configuration parameters *DETECT_DEADLOCKS*, *EXIT_IF_HEALTHY_LOST_IN_STANDALONE_MODE*, *EXTERNAL_EVENT_HANDLER*, *INNER_SEQUENCE_NUMBER_IN_SEND_MSG_RESPONSE*, *IPMC_PRESERVE_ON_REVISION_CHANGE*.
- New section 3.12: describes the new external event handling facility, including how to configure and use it.
- New section 4.8: describes the deadlock detection facility.
- Entire document: removes coverage of first generation ShMM-300 mezzanine. Starting with release 2.5.0, the ShMM-300 is not supported on new releases of the Shelf Manager.

C.9 Release 2.5.2

- Section 3.3: adds new configuration parameters: *COOLING_MANAGEMENT*, *HPDL*, *HPDL_ON_SUBSIDIARY_FRUS*, *ISOLATE_MUX_ON_GPIO8*, *PET_FORMAT*, *SWAPPED_CROSS_CONNECTS*, *ENABLE_DIRECT_SHELF_FRU_WRITE*. Modifies the description of the configuration parameter *SWITCHOVER_ON_HANDLE_OPEN* to reflect that programmatic deactivation also triggers a switchover.
- Section 3.4.2.2: describes the second instance of the Shelf Manager IP Connection record in the Shelf FRU Information as the source of the IP address, netmask and default gateway for the second RMCP network interface.
- New section 3.4.2.3: describes support for site-dependent ShMC cross-connects.
- New section 3.6: describes configuring Carrier and Shelf Attributes using HPDL (the Pigeon Point Hardware Platform Description Language).
- New section 3.7: describes configuring the cooling management strategy on HPDL-based platforms.
- Section 3.8: specifies that the section applies only to systems where HPDL is not used.
- New section 3.13: describes how to configure the Platform Event Trap format.
- Section 4.4.1: describes the event message format for the Redundancy and CPLD state sensor.
- Section 4.4.3: augments the list of reboot causes that are reported by the reboot reason sensor.

- New section 4.7: describes the cooling state sensors.
- Section 4.8: adds information about the direct deadlock detection mechanism.
- Section 7: changes the heading structure to eliminate an unnecessary heading level.

C.10 Release 2.5.3

- Section 3.3: corrects the description of the configuration parameter *COOLING_MANAGEMENT*.
- Section 3.6.3: amends the description of the procedure to locate HPDL data, taking into account the new environment variables that can provide the location of that data.
- New section 4.9: describes the provisions for isolating faulty subsidiary buses behind a multiplexer on the ShMM carrier, for carriers that have GPIO control of that multiplexer.

C.11 Release 2.6.0

- Entire document: adds coverage of ShMM-1500R.
- Sections 3.1, 3.10.1: changes the default time zone name from “UTC0” to “UTC”.
- Section 3.3: adds new configuration parameters:
ACTIVATE_LOCAL_WITHOUT_SHELF_FRU, ATCA_TESTER_COMPATIBILITY, DEFAULT_VLAN_ID, DEFAULT_VLAN_ID2, DHCP_FOR_RMCP_ONLY, TURBO_MODE_MIN_FAN_FAILURES.
- New section 3.4.8: describes assigning VLAN IDs in accordance with version 2.0 of the IPMI specification.
- Section 3.12.1: in the PEF configuration example, corrects “control 7” to “control 1”.
- New section 4.5.1: describes operation in ShMM-500-based Shelves with radial IPMB-0.
- New section 4.5.2: describes operation in ShMM-1500-based Shelves with radial IPMB-0.
- New section 5: introduces the use of the IPMI analysis tools.
- Section 7.3: adds tables describing the Flash partitions for a ShMM-1500 with 32 MB and 64MB Flash devices.
- Section 7.6: adds descriptions of new options and specifiers for the reliable upgrade utility: *-a, -R, -C.*
- Section 7.8.2: adds an example reliable upgrade scenario for the ShMM-1500.

C.12 Release 2.6.1

- Section 3.2.1: adds descriptions of new U-Boot variables *log_max* and *log_remote*.
- Section 3.3: adds new configuration parameters: *ENABLE_INTEGRALHPI, INTEGRALHPI_DOMAIN_ID, TACHOMETER_THRESHOLD_UPDATE_DELAY.*
- Section 3.3: corrects the description of the configuration parameter *REDUNDANCY_NET_ADAPTER2.*
- New section 3.13.1: adds a parsed example of an SNMP Trap in default (IPMI PET v1.0) format.
- New section 3.14: describes configuration of the IntegralHPI subsystem.

C.13 Release 2.6.4

- Section 2.4.4: lists the full set of ShMM variants, including those that have reached an end of life status. Describes new ShMM-1500 models with encryption code removed.
- Section 3.2.1: adds descriptions of the HPDL carrier-specific option *CARRIER_FRU_LOCATION*.
- Section 3.3: adds new configuration parameters:
BOARD_LAN_PARAMETERS_CHANNEL_LIST, INITIALIZATION_SCRIPT, PEF_USE_KEYED_ALARMS, SEL_FILE_COMPRESSION_ENABLED, SEL_FILE_JOURNALING_ENABLED, SEL_FILE_WRITE_DELAY, SHELF_MANAGER_CONFIGURATION_IN_SHELF_FRU_INFO, TIMEPROTO, TIMESERVER.
- New section 3.3.2: describes a new mechanism for retrieving configuration variables from Shelf FRU Information.
- New section 3.4.2.4: describes bonded usage of the two network interfaces.
- New section 3.6.4.1: describes how to specify location of the Carrier FRU Information.
- Section 3.14.2: adds a note about the special version of the HPI client library.
- Section 4.3: adds a description of the option *-i* for Shelf Manager command-line invocations; adds a note about Shelf Manager configurations with encryption code present or removed.
- New section 4.10: describes how to use a new facility in which the Shelf Manager stores and assigns LAN configuration parameters to boards or modules in a shelf that implement LAN channels in the local management controller.
- Section 7.6: adds a note about reliable upgrade from an encryption-present RFS to an encryption-absent RFS.

C.14 Release 2.6.4.2

- Section 3.4.9: adds a description of the DHCP client restart and status commands; describes provisions for storing, in a flash file, a TFTP server name and bootfile name received from the DHCP server.

C.15 Release 2.6.4.4

- Section 3.3: adds coverage of a new configuration parameter:
SWITCHOVER_ON_BROKEN_LINK_BACKUP_DELAY.
- Section 3.4.2.1: provides user guidance for a special situation when a hub board is inserted into a shelf with cross-connect Shelf Manager links where no hub boards were present before.

C.16 Release 2.7.0

- Section 3.3: adds new configuration parameters:
BOARD_LAN_PARAMETERS_CHANNEL_SYNCHRONOUS, BOARD_LAN_PARAMETERS_USE_DHCP,

COOLING_KEEP_POWERED_OFF_FRUS_IN_M1. The increased length of the configuration parameter string *BOARD_LAN_PARAMETERS_CHANNEL_LIST* is noted.

- Section 3.4.2.1: adds a description of active-standby management of the network interface on the backup Shelf Manager, when RMCP address propagation is in effect.
- Section 3.7.1: adds a description of the effects of the new configuration variable *COOLING_KEEP_POWERED_OFF_FRUS_IN_M1* and a description of the handling cold-sensitive FRUs in the default cooling algorithm.
- Section 3.14.2: corrects the path to the OpenHPI client library.
- New section 4.4.1: describes in detail the initialization of redundant Shelf Managers.
- Section 4.10.2: adds a description of how the Shelf Manager retrieves LAN parameters from a DHCP Server.
- New section 4.10.2.3: describes the configuration of a Linux DHCP Server to enable retrieving LAN Configuration information from it.
- New section 4.10.4: describes the synchronous assignment of LAN configuration parameters to a board.

C.17 Release 2.7.1

- Section 3.3: adds new configuration parameters:
ALLOW_POWER_UNRELATED_FRU_IN_CRITICAL_STATE,
REAPPLY_POWER_MAX_COOLING_STATE.
- Section 3.7.1: adds a description of the effects of the new configuration variables *ALLOW_POWER_UNRELATED_FRU_IN_CRITICAL_STATE* and *REAPPLY_POWER_MAX_COOLING_STATE*.
- Section 3.10.1: clarifies that ShMM timestamps correspond to the local time zone.
- Section 4.4.3: adds descriptions of two more states of the Reboot Reason sensor.
- New section 4.4.1: describes the HPI System Event sensor.
- Section 5.2, new section 5.4.3: describes the IPMB board trace mode.
- Section 7.6: adds a description of a new parameter *conf* for the upgrade script *step4hshm*.

C.18 Release 2.7.2

- Sections 7.6, 7.7: adds a description of a new parameter *flip* for the upgrade script *step4hshm*. Notes that the option *-u* is no longer supported.
- New section 3.15: describes configuration of local services.

C.19 Release 2.7.3

- Section 3.3: adds new configuration parameter:
COOLING_NO_POWER_DOWN_IN_CRITICAL_ALERT.
- Section 3.15: adds more information about configuration of the SNMP service.

C.20 Release 2.7.4

- Section 3.3: adds new configuration parameter:
REDUNDANCY_COMPRESSION_THRESHOLD.

C.21 Release 2.8.0

- Section 3.3: adds new configuration parameters: *NO_M0_M1_EVENT_AT_STARTUP*, *PET_OEM_WITH_SEVERITY_STRING*, *RUN_HPI_SNMP_SUBAGENT*, *SPECIAL_HPI_ENTITY_FOR_AMC_CARRIER*.
- Section 3.3: adds a description of the *corrupted_images* variable.
- Section 3.3: adds coverage of the */etc/localtime* file.
- Section 3.3: adds a description of double-active state detection.
- Section 4.12: adds descriptions of the ShMM tests accessible via Diagnostic Initiators.
- Section 9: adds a description of HPI-based upgrades.

C.22 Release 2.8.1

- Section 3.3: adds a new configuration parameter:
HPDL_NETWORK_ELEMENT_ID_EEPROM_OFFSET.
- Section 3.3: fixes a typo in the configuration parameter name
SHELF_MANAGER_CONFIGURATION_IN_SHELF_FRU_INFO.

C.23 Release 2.8.2

- Section 3.3: adds a new configuration parameter: *DHCP_SPECIAL_CLIENT_ID_FORMAT.*

C.24 Release 3.0.0

- Section 2: adds a description of the ShMM-700.
- Section 2.2: updates Figure 1.
- Section 2.4.1: adds HPDL coverage.
- Section 2.4.2: updates Figure 2.
- Section 2.4.4: adds text and a new column for the ShMM-700 in Table 3.
- Section 3: adds ShMM-700 coverage.
- Section 3.3: clarifies the description of the configuration parameter:
ATCA_TESTER_COMPATIBILITY.
- Section 3.3: describes new configuration variables: *AXIE_TIMING_ASYMMETRIC_MATCH*, *ENABLE_LOCKS_LOGGING*, *ENABLE_RTC_TRICKLE_CHARGER*, *ISOLATE_MUX_BUS*.
- Section 3.3: describes new IntegralHPI-related configuration variables:
INTEGRALHPI_SHELF_ENTITY_LOCATION,
INTEGRALHPI_DEFAULT_AUTO_INSERT_TIMEOUT,
INTEGRALHPI_DEFAULT_AUTO_INSERT_TIMEOUT_MSEC,

INTEGRALHPI_DEFAULT_AUTO_EXTRACT_TIMEOUT, and
INTEGRALHPI_DEFAULT_AUTO_EXTRACT_TIMEOUT_MSEC.

- Section 3.4.7: describes the new utility *hwri* that replaces *cpld* utility on the ShMM-700.
- Section 3.5.2: adds ShMM-700 coverage.
- Section 3.10.1: clarifies the behavior of the system clock before the time is successfully obtained from the time server.
- Section 4.3: corrects the example of the Shelf Manager startup log.
- Section 4.4: adds ShMM-700 coverage.
- Section 4.8: adds ShMM-700 coverage.
- Section 4.9: adds ShMM-700 coverage.
- Section 4.9: adds coverage of the new configuration parameter *ISOLATE_MUX_BUS*.
- Section 4.12: clarifies current support on the ShMM-500/1500.
- Section 5: notes current absence of IPMB trace facility on ShMM-700.
- Section 6: adds ShMM-700 coverage.
- Section 7: notes ShMM-500/1500 focus for reliable upgrade coverage.
- New section 8: covers ShMM-700-specific reliable upgrades.
- New section 9: covers HPI-based firmware upgrades for all three ShMM types.
- Renumbered section 8 - Appendix C: updated section numbers reflect new section 8 and section 9.
- New Appendix A: describes how to convert a shelf from ShMM-500-based to ShMM-700-based Shelf Managers.
- Appendix B and Appendix C: retitled content from previous sections 10 and 11.

C.25 Release 3.1.1

- Section 5: adds coverage of HPM.2 IPMI Trace Payloads in board trace mode.
- Section 5: removes the limitation for ShMM-700 of not supporting controlled mode.
- Section 8.7: adds a note about encryption of SmartFusion cSoC upgrade images.

C.26 Release 3.2.0

- Section 3.3: describes new configuration variables: *ALLOWED_CIPHER_SUITES*, *AXIE_SEQUENCING_ENUM_READY_TIMEOUT*, *AXIE_SEQUENCING_M4_TIMEOUT*, *AXIE_SEQUENCING_WAIT_DELAY*, *AXIE_SYSTEM_MODULE_IPMB_ADDRESS*, *AXIE_NOT_READY_STARTUP_DELAY*, *TURBO_MODE_MIN_MISSING_FAN_TRAYS*, *WATCHDOG_TIMEOUT*.
- New section 3.16: describes how to configure the maximum speed of I²C buses on the ShMM-700.

- Section 4.5.3: introduces the *ripmb_prog* utility for programming of ShMM-700 routing element SmartFusion cSoC(s).
- Section 4.9: adds a description of the HPDL-based method of I²C bus fault isolation.

C.27 Release 3.3.0

- Section 3.3: describes new configuration variables:
HPM2_ENABLE_INACTIVE_STATE_MANAGEMENT,
HPM2_ENABLE_SHMM_INACTIVE_STATE, *HPM3_ENABLE*,
IPMB_LOCAL_FAILURE_RETRIES, *RMCP_PLUS_ONLY*.
- Section 3.4.9: describes the HPM.3 Client Identifier format as default format for DHCP IP address assignments.
- Section 4.4.3: adds a new cause of the last reboot.

C.28 Release 3.4.0

- Section 3.3: describes new configuration variable *SAVE_CRASH_INFORMATION*. The *COOLING_IGNORE_LOCAL_CONTROL* variable was removed, since it was not used in the Shelf Manager. The default value of the *HPDL_ON_SUBSIDIARY_FRUS* variable was corrected to *TRUE*. Clarification was added for the *IPMB_RETRY_TIMEOUT_MSEC* variable. The value *CRITICAL_ALERT* of the variable *REAPPLY_POWER_MAX_COOLING_STATE* is no longer valid.
- Section 3.3: describes new configuration variables *COOLING_EXHAUST_SENSOR_LIST*, *COOLING_INTAKE_SENSOR_LIST*, *DEFAULT_HPM2_CLIENT_IP_ADDRESS*, *DEFAULT_HPM2_GATEWAY_IP_ADDRESS*, *DEFAULT_HPM2_NETMASK*, *HPM2_ENABLE_CLIENT_SESSIONS*, *HPM2_CLIENT_MIRROR_CHANNEL*, *HPM2_NET_ADAPTER*.
- New section 3.4.3: covers configuration considerations for the HPM.2 channel.
- Section 3.15: explains how to establish IPv6 access to FTP, TELNET and SSH, which is available only on the ShMM-700.
- Section 3.16: adds explanation on bus mapping for ShMM-700.
- Section 4.4.3: adds coverage of the */var/nvdata/reboot_reason.log* file, which contains information about recent reboots.
- Section 4.5.3: explains how the redundant Shelf Managers communicate with each other over a cross link.
- Section 4.10.2: corrects the name of the LAN configuration parameters file on the ShMM.
- Appendix A: adds coverage of particular aspects of converting a live shelf from ShMM-500-based management to ShMM-700-based management.
- Appendix B: updates contact information.

C.29 Release 3.5.0

- Section 3.3: describes a new IntegralHPI-related configuration variable and new Shelf Manager configuration variable controlling the SNMP trap version format:
INTEGRALHPI_SESSION_EVENTS_CAPACITY,
PET_SNMP_VERSION,
DHCP_REQUEST_OPTION_LIST.
- New section 3.4.10: covers configuration of IPv6-based services of the Shelf Manager.
- New section 8.4: adds a ShMM-700-specific description of the *user_partition* and *flash_reset* variables in U-Boot and details for 'user' partition format changes.
- Sections 3.2.1, 3.4.7, 4.2, 6.3 and 6.4: changes instances of the term “factory defaults” to alternate terms appropriate to each context.

C.30 Release 3.5.0.1

- Section 3.3: describes new configuration variable
COOLING_NO_ALL_FRUS_POWER_DOWN_IN_CRITICAL_ALERT.
- Section 4.4.2: describes the previously undocumented state “Active with Backup unavailable” of the CPLD State/HWRI State sensor on the physical Shelf Manager.

C.31 Release 3.5.1

- Section 3.4: adds notes on the risks of doing manual changes in the ShMM network configuration separately from the Shelf Manager.
- Section 3.4.10: updated with information about support in this release of PICMG 3.0 R3.0 IPv6 amendments in ECN 3.0-3.0-001.
- Section 3.15: adds information about configuring IPv6 support for various network services.

C.32 Release 3.6.0

- Section 3.3: describes new configuration variables *DHCPV6_REQUESTED_OPTIONS*, *DHCPV6_OPTIONS_SCRIPT*, *DHCPV6_SERVER_ADDRESSES*, *ENABLE_NEW_DHCPV4_CLIENT*, *USE_DHCPV6*.
- Section 3.4.9: adds information about the new DHCPv4 client.
- New section 3.4.11: covers the DHCPv6 client for the Shelf Manager.
- Section 4.9.2: adds information on how to cause a switchover to be initiated if the number of failed I²C multiplexer reset attempts exceeds a certain threshold.

C.33 Release 3.6.1

- Overall: implements branding changes to reflect the acquisition of Pigeon Point Systems by Pentair.
- Section 3.2.1: adds descriptions of new U-Boot variable *log_remote2*.
- Section 8.6: adds the *--clone* option for the *rupgrade* utility.

- Appendix A: adds more detail on converting a live shelf from ShMM-500-based management to ShMM-700-based management.

C.34 Release 3.6.1.2

- Section 3.3: describes configuration variables `AXIE_ROOT_PREFERENCE_CHANNELS`, `AXIE_SEQUENCING_SYSTEM_MANAGER_TIMEOUT`, `PREFER_CHASSIS_HPDL_SDRS_FROM_FILE`. Fixes the default value of `MAX_PENDING_IPMB_REQUESTS` configuration variable.
- Section 8.6: describes the `--install-only` option of the `rupgrade` utility and the `--activate` option of the `backend` utility.
- New section 8.9: describes the `imls` utility.
- New section 8.10.4: adds an example of a reliable upgrade scenario using the `--install-only` option.

C.35 Release 3.6.1.5

- Section 3.3: describes the new configuration variable `DISABLE_SESSION_AUDIT_EVENTS`.
- Section 5.2: adds a description of `-r` option to enable radial IPMB link monitoring in the `ipmb_traced` daemon.

C.36 Release 3.6.1.6

- Section 3.3: corrects the description of the configuration variable `DISABLE_SESSION_AUDIT_EVENTS`.
- Section 5.2: adds a description of a new parameter of `-r` option in the `ipmb_traced` daemon.
- Section 8.6: describes the `-copy-home` option of the `rupgrade` utility.

C.37 Release 3.7.0

- Section 3.3: adds the description of new configuration variables `USE_IPMB_A_ONLY`, `PET_PORT` and `BOARD_LAN_PARAMETERS_ASSIGN_FROM_SHM`.
- Section 4.5.3.1: adds coverage for the SmartFusion A2F200 cSoC device. Shelf Manager Release 3.7.0 supports ShMM-700R modules that use a SmartFusion A2F200 or A2F060 device and Carrier boards that use SmartFusion A2F200 or A2F060 devices to implement a Radial IPMB-0 topology.

C.38 Release 3.7.1

- Overall: implements branding changes to reflect the split of nVent from Pentair.

C.39 Release 3.7.1.3

- Section 2.3: the description of the BMR variants is made more specific.

C.40 Release 3.7.1.6

- Section 3.12.2: an example of the external event handler script is added.

C.41 Release 3.7.1.8

- Section 2.4.4: updates the list of ShMM models. Updates the statement on ECCN for the product.

C.42 Release 3.7.2

- New section 3.4.4: Gigabit Ethernet support description for ShMM-700. Also, corresponding variables have been added to Table 5.