

OPEN

Compute Summit

January 28-29, 2014

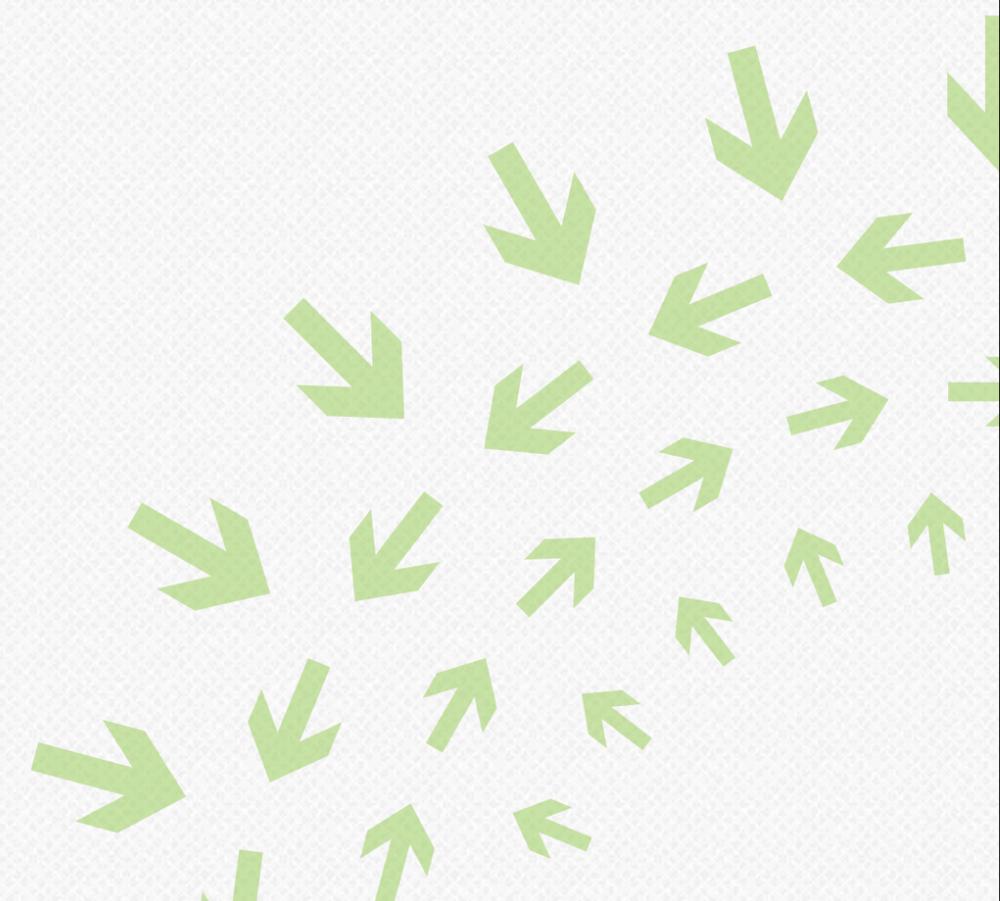
San Jose





HPM.1 Firmware Update

Hank Bruning
JBlade
hank@jblade.com



Status

1. The OCP Hardware Management subcommittee is collecting requirements for firmware update.
2. This presentation is a subset of the what was presented at the OCP Hardware Management workshop on October 1, 2013
3. HPM.1 has not been chosen as the firmware update method



Agenda

1. Goal

2. HPM.1 Overview

3. Patents





Goals

Goal

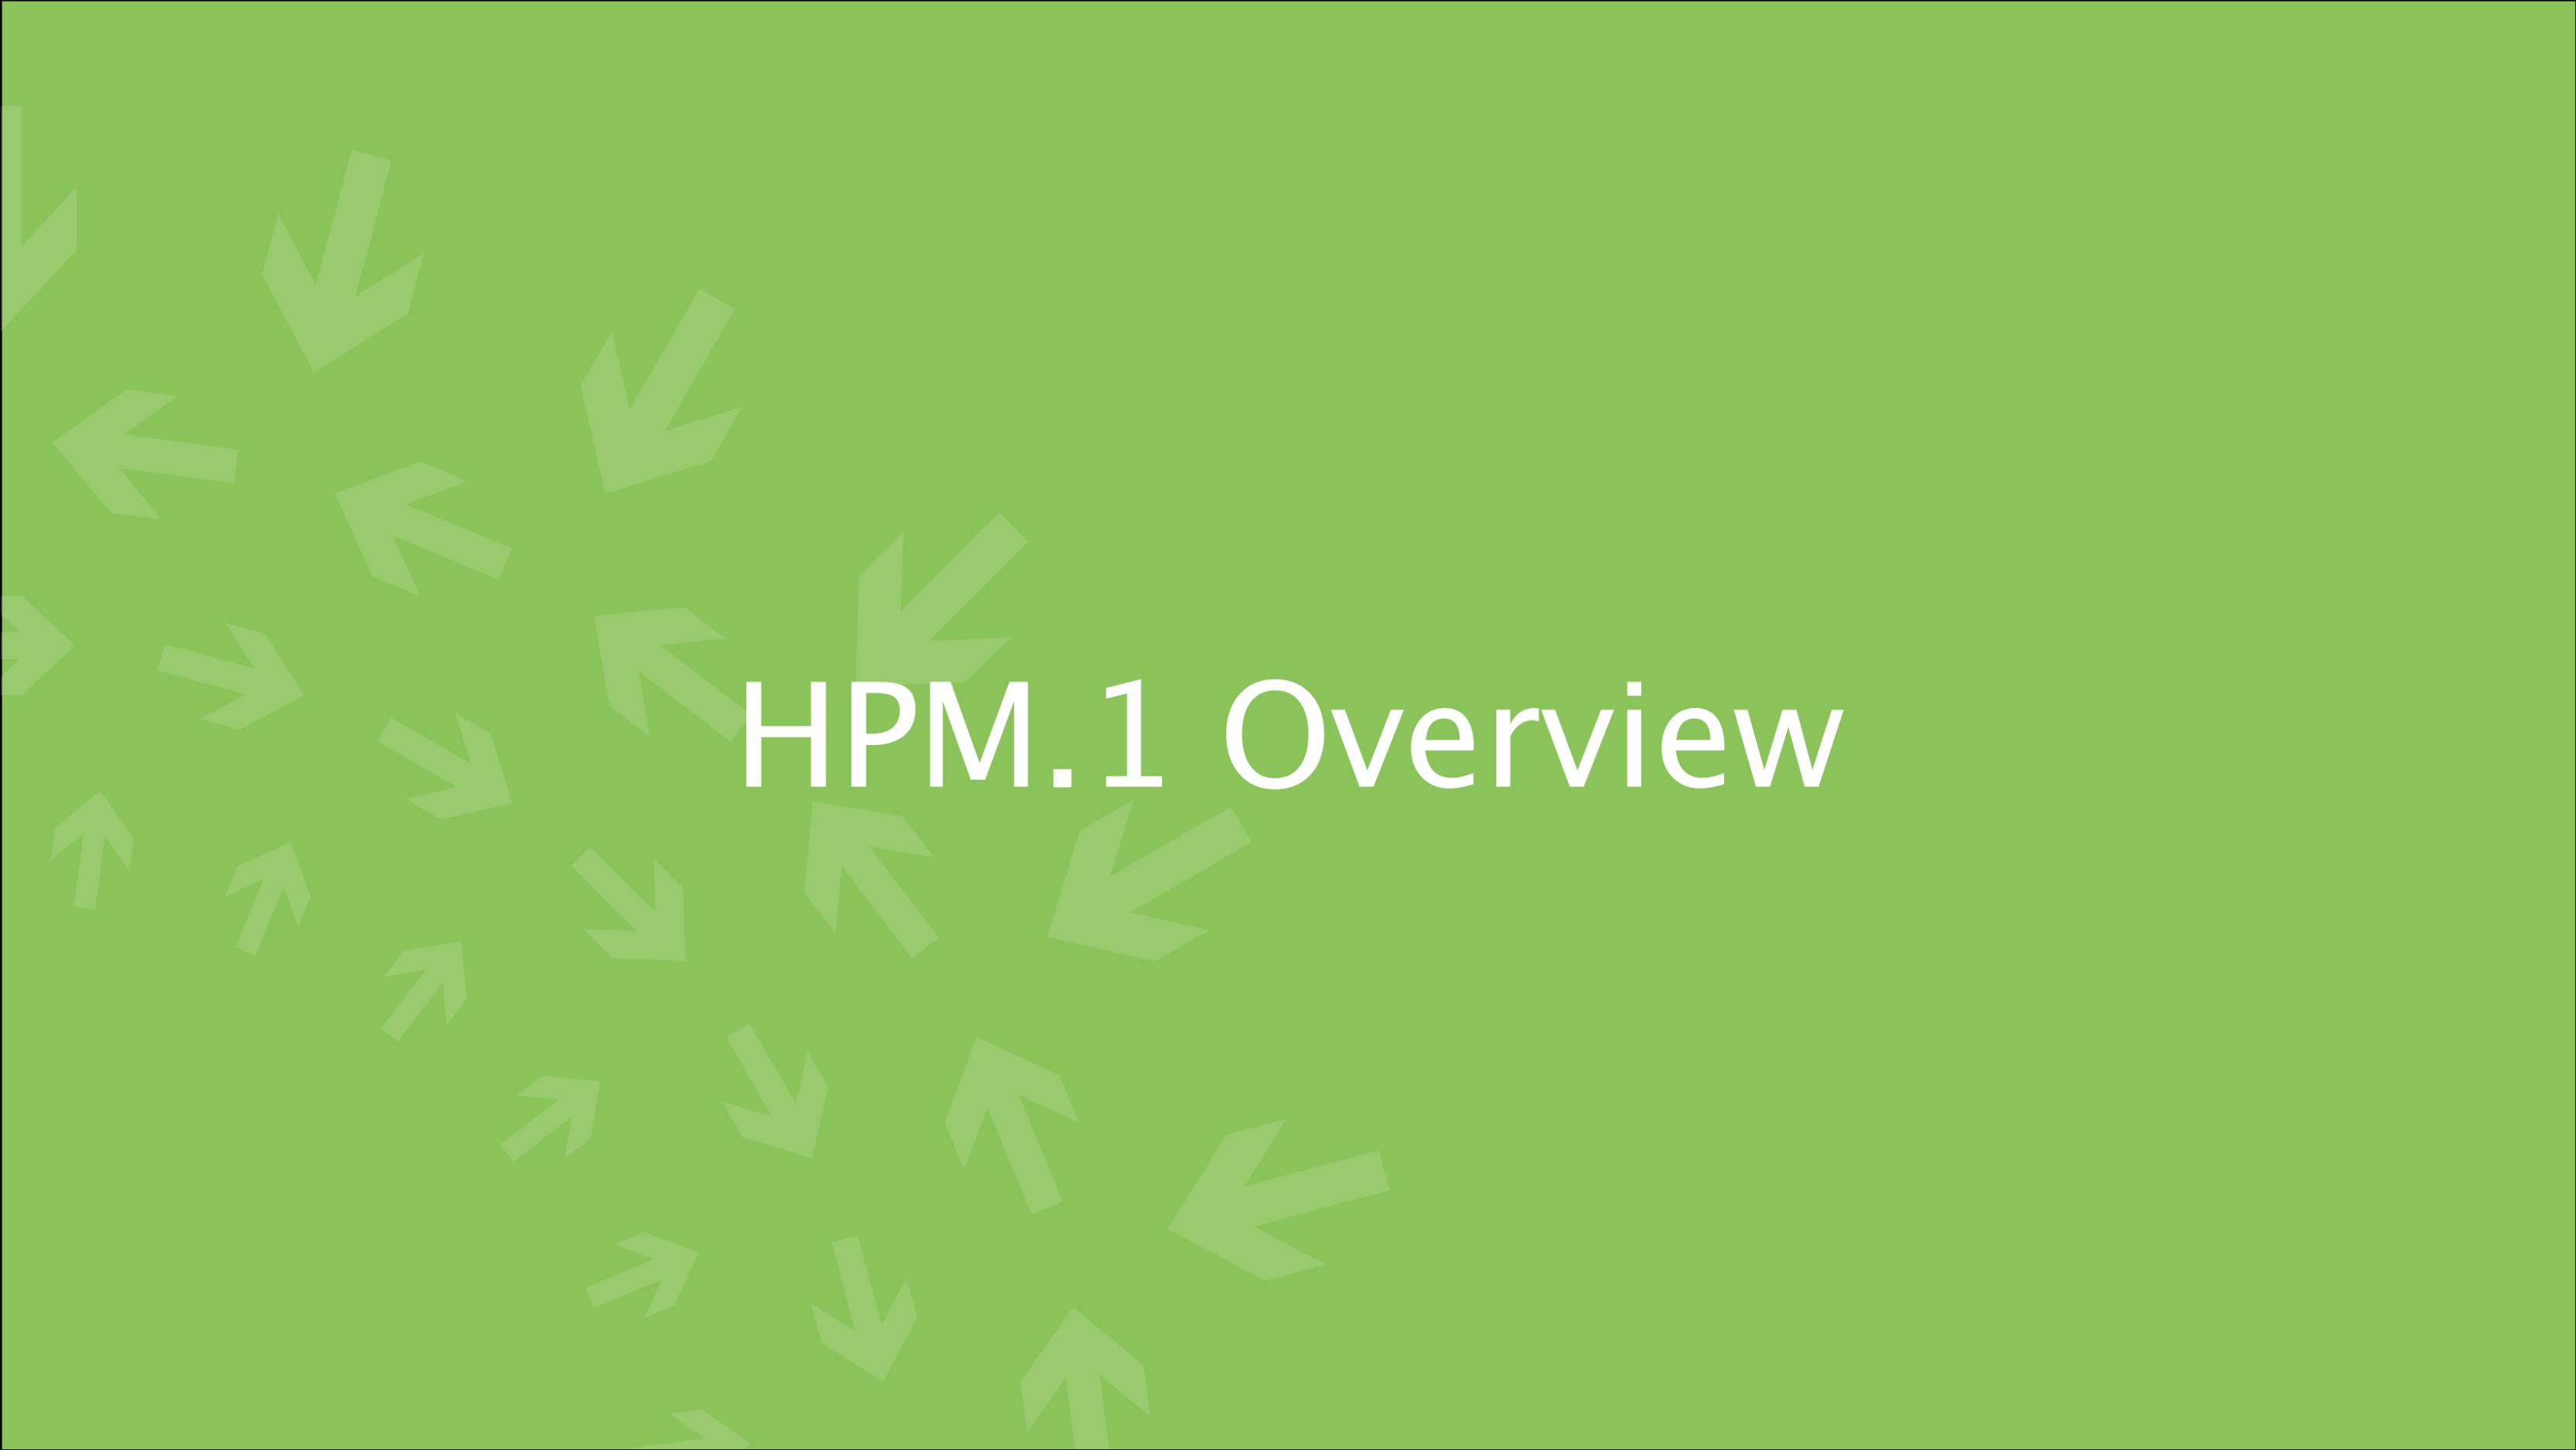
- Provide a uniform Firmware update process which updates Flash, FPGAs, Crosspoint switches, etc for all internet enabled hardware in an OCP rack.
- Requirements
 - Patent Free
 - Open Source implementation
 - Operating System Independent
 - Does not require physical access to hardware



Goal

- The target hardware for the firmware update ranges from the simplest device to the most complex.
 - Power Distribution units
 - Servers (this OCP working group)
 - RAIDs
 - Ethernet, Fiber Channel and Infiniband switches





HPM.1 Overview

HPM.1 Overview

- Define by PICMG in 2007 for Telcos
 - HPM.1 is one of 45+ PICMG specs developed in last 20 years
 - Available for purchase, \$100, at www.picmg.org
- 300+ companies supported the spec
 - Alcatel, AMD, Cisco, Fujitsu, Huawei, IBM, Lucent, Nokia, Sun....
- Of the 4 IPMI Promoters (they control the IPMI spec)
 - Intel, HP and NEC supported HPM.1
 - Dell did not vote, not a PICMG member



HPM.1 Open Source

- Open source C++ implementation of the RMCP client was published in 2007 by Kontron and present in ipmitool.
- No open source implementation of what the BMC would execute.



HPM.1 Vendors

- Telco Hardware

 - AdvanTech, Kontron, Emerson Network Power (old Motorola Computer Group), Radisys

- Data Center Hardware

 - None as of January 2014

- BMC Software Vendors

 - Pigeon Point Systems

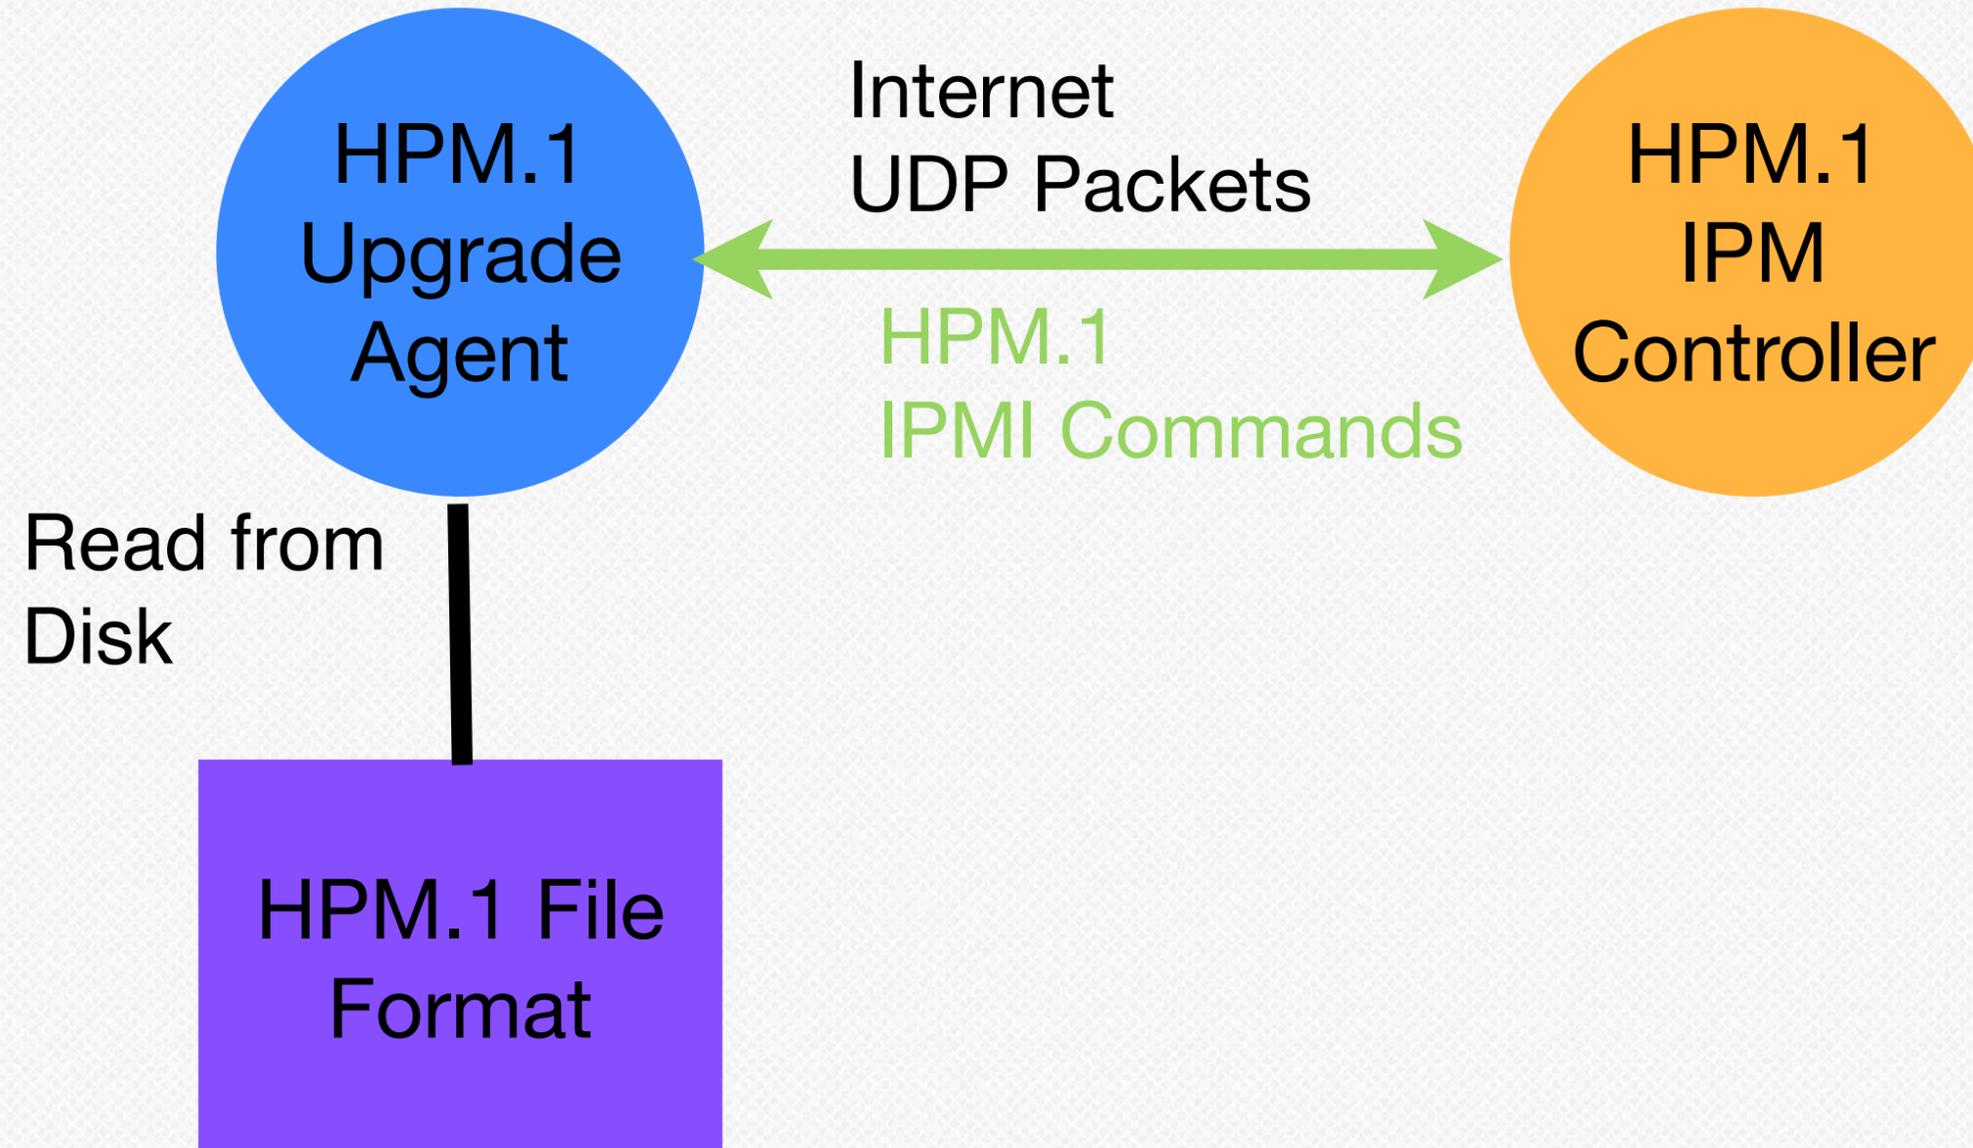


HPM.1 Specification

- HPM.1 is small
 - 66 pages long
 - 184 Numbered requirements
- Defines 4 functional areas
 - Firmware image file format
 - IPMI Commands
 - IPM Controller requirements (software in the server)
 - Upgrade agent (software outside the server. ipmitool)



HPM.1 Functional Areas



HPM.1 File Header

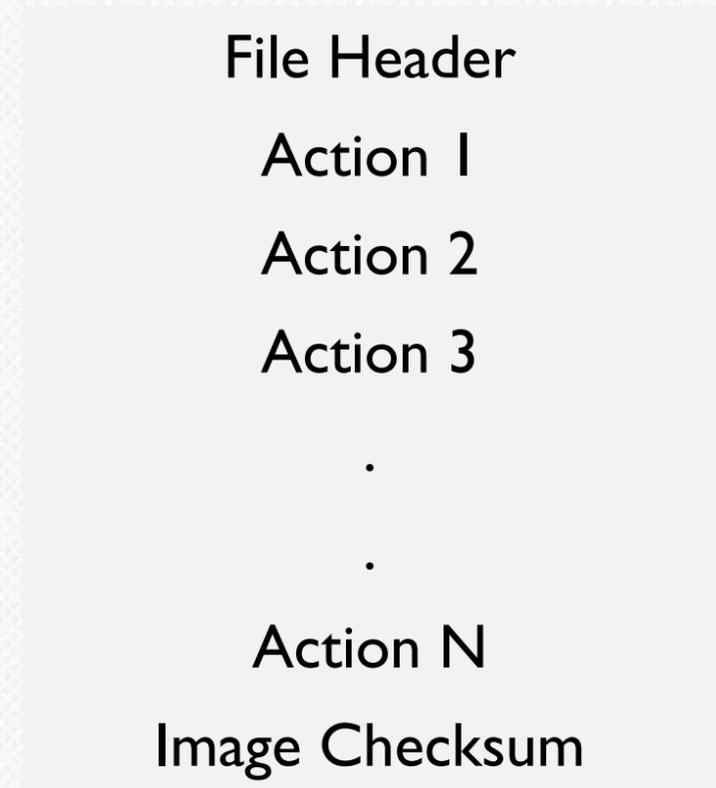
- Upgrade agent identifies file using Manufacture, Device and ProductID from IPMI Command “Get Device ID”
- Revision number
 - Mandatory two bytes. Format is X.Y
 - Optional trailing 3 bytes. Any format.
 - Each action containing firmware may have it's own free form version number

Image Signature
Device ID
Manufacturer ID
Product ID
Image create time
Image Capabilities
Components
Self test timeout
Roll back timeout
Inaccessible timeout
Earliest compatible
Revision of file
OEM Info
Header checksum



HPM.1 File Format

- File is not sent as a single image
- Each action is sent by itself to IPM Controller
 - A single image can contain a mix of x86, ARM, FPGA, etc firmware each as an action.
 - Actions move the sequencing of update process from DC operators to HW vendors.
 - An action may take seconds or many minutes
- An action operates on one of eight devices
- Action can be one of Backup/Prepare/Upload firmware
- A device can be operated on by multiple actions
- Actions may contain different firmware architectures
- Actions can contain a self test ← Big Win



HPM.1 IPMI Commands

Command Name	Mandatory/Optional
Get target upgrade capabilities	M
Get component properties	M
Abort firmware upgrade	O
Initiate upgrade action	O/M
Upload firmware block	M
Finish firmware upload	M
Get upgrade status	O/M
Activate firmware	M
Query self-test results	O/M
Query rollback status	O/M
Initiate manual rollback	O/M

O/M- the command may be mandatory depending of fields in “Get Component Properties” or IPM implementation. See the spec for details.



HPM.1 IPMI Commands

- There are only five IPMI commands necessary for minimal HPM.1 support.
- I estimate that 70% of existing monolithic servers, RAIDs, Switches using IPMI RMCP can be firmware upgraded to support HPM.1. No hardware changes necessary



HPM.1 Upgrade Agent

- Runs outside the server/Raid/Switch and communicates with an IPM Controller using IPMI RMCP packets(UDP protocol)
- Uses same security, authentication and integrity as all other IPMI commands
- Responsible for finding HPM.1 image and sending to to IPM Controller
- An Upgrade agent may dynamically create the upgrade image.

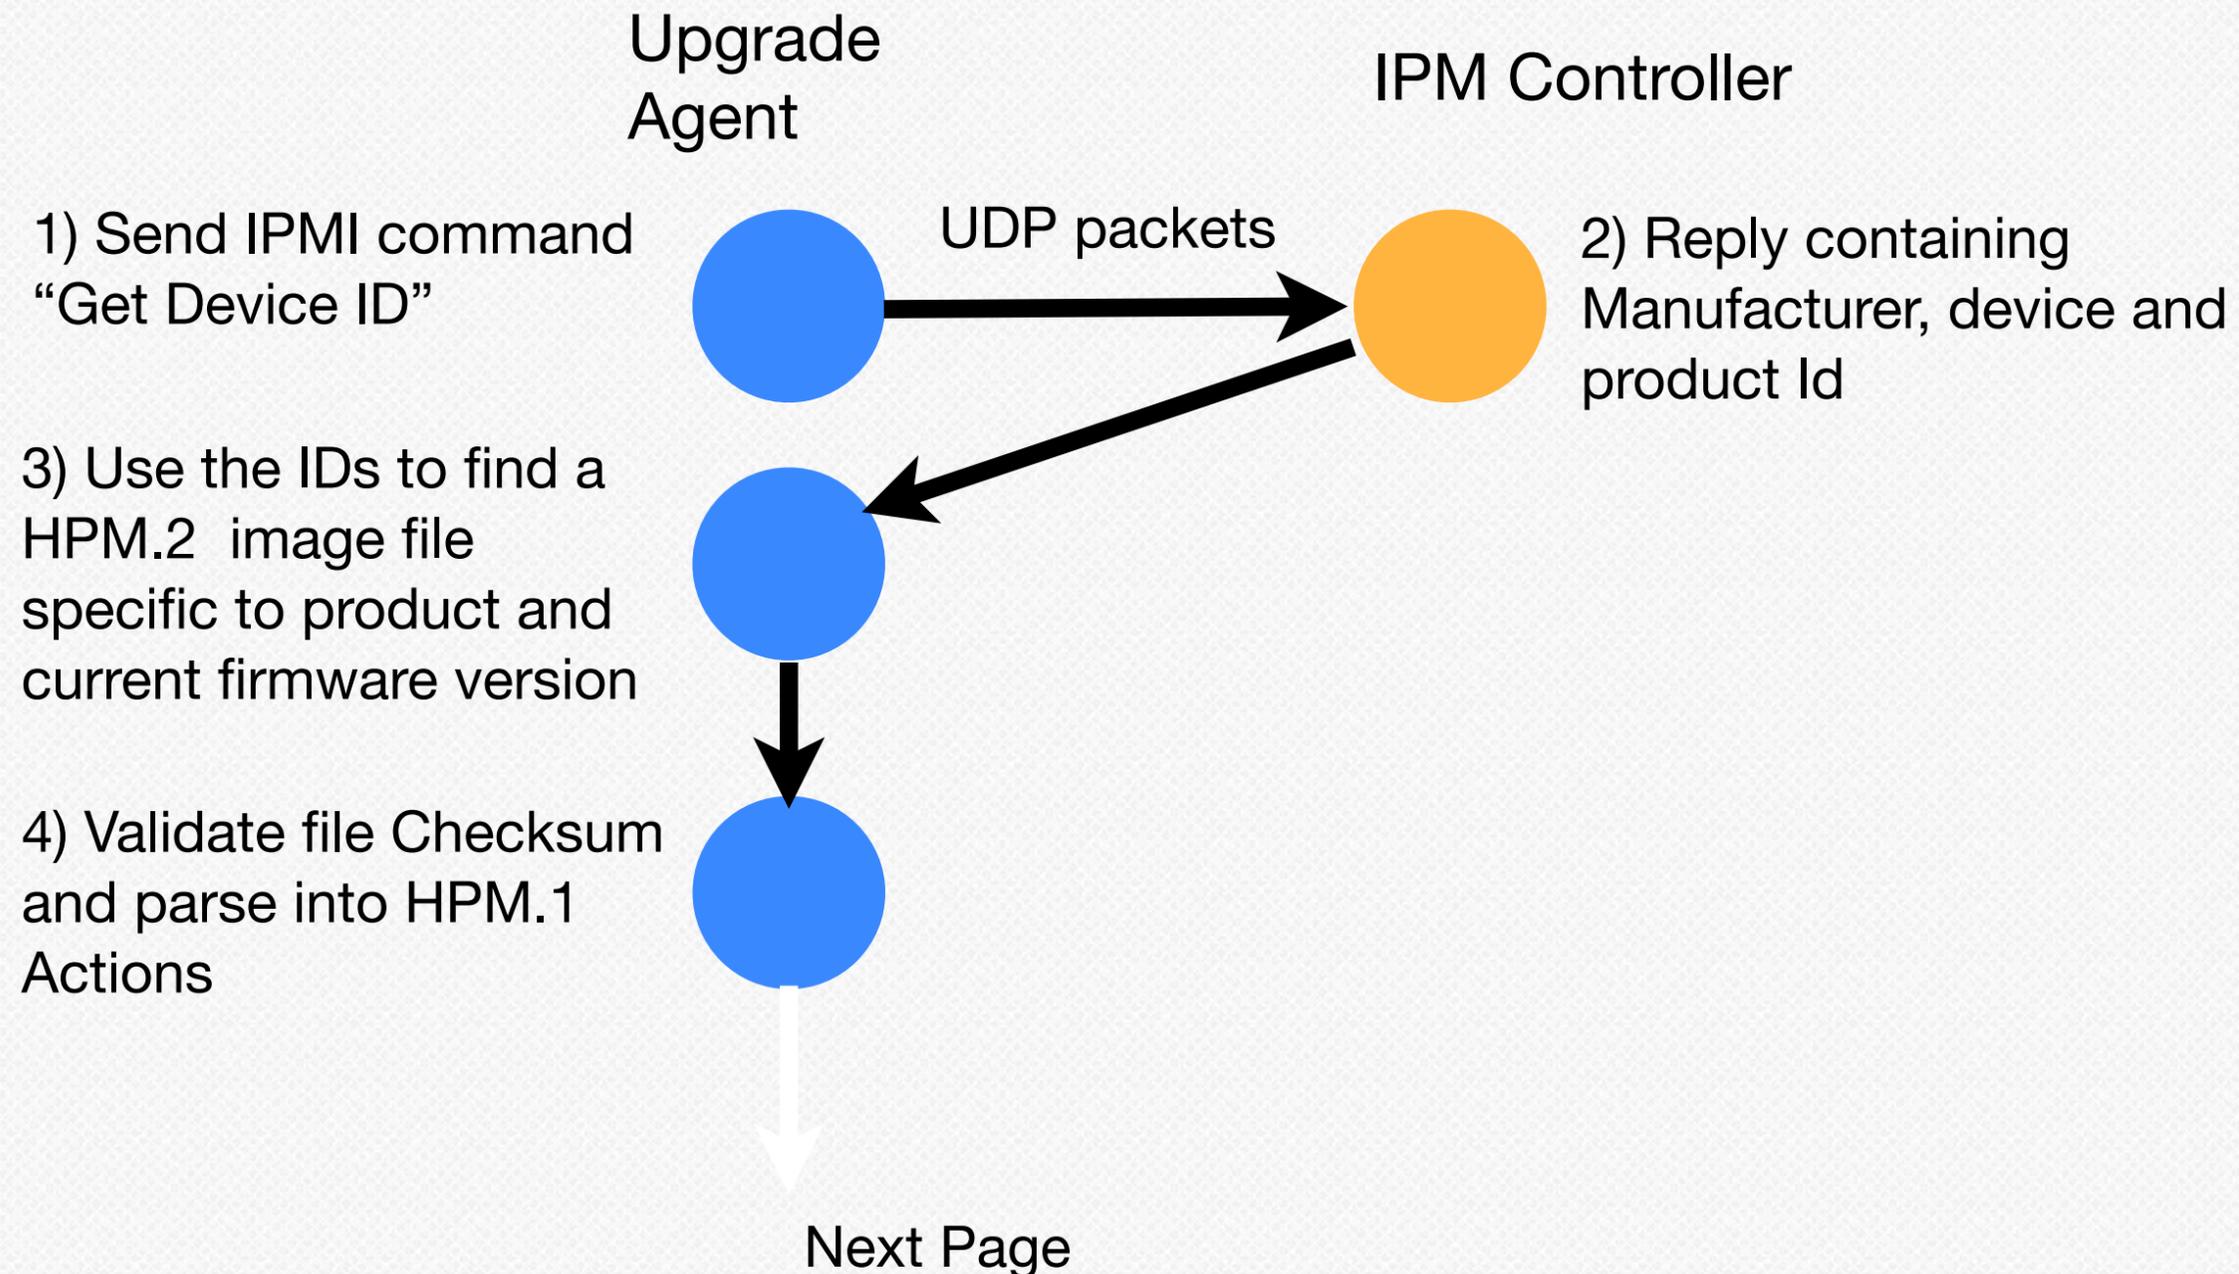


HPM.1 IPMI Controller

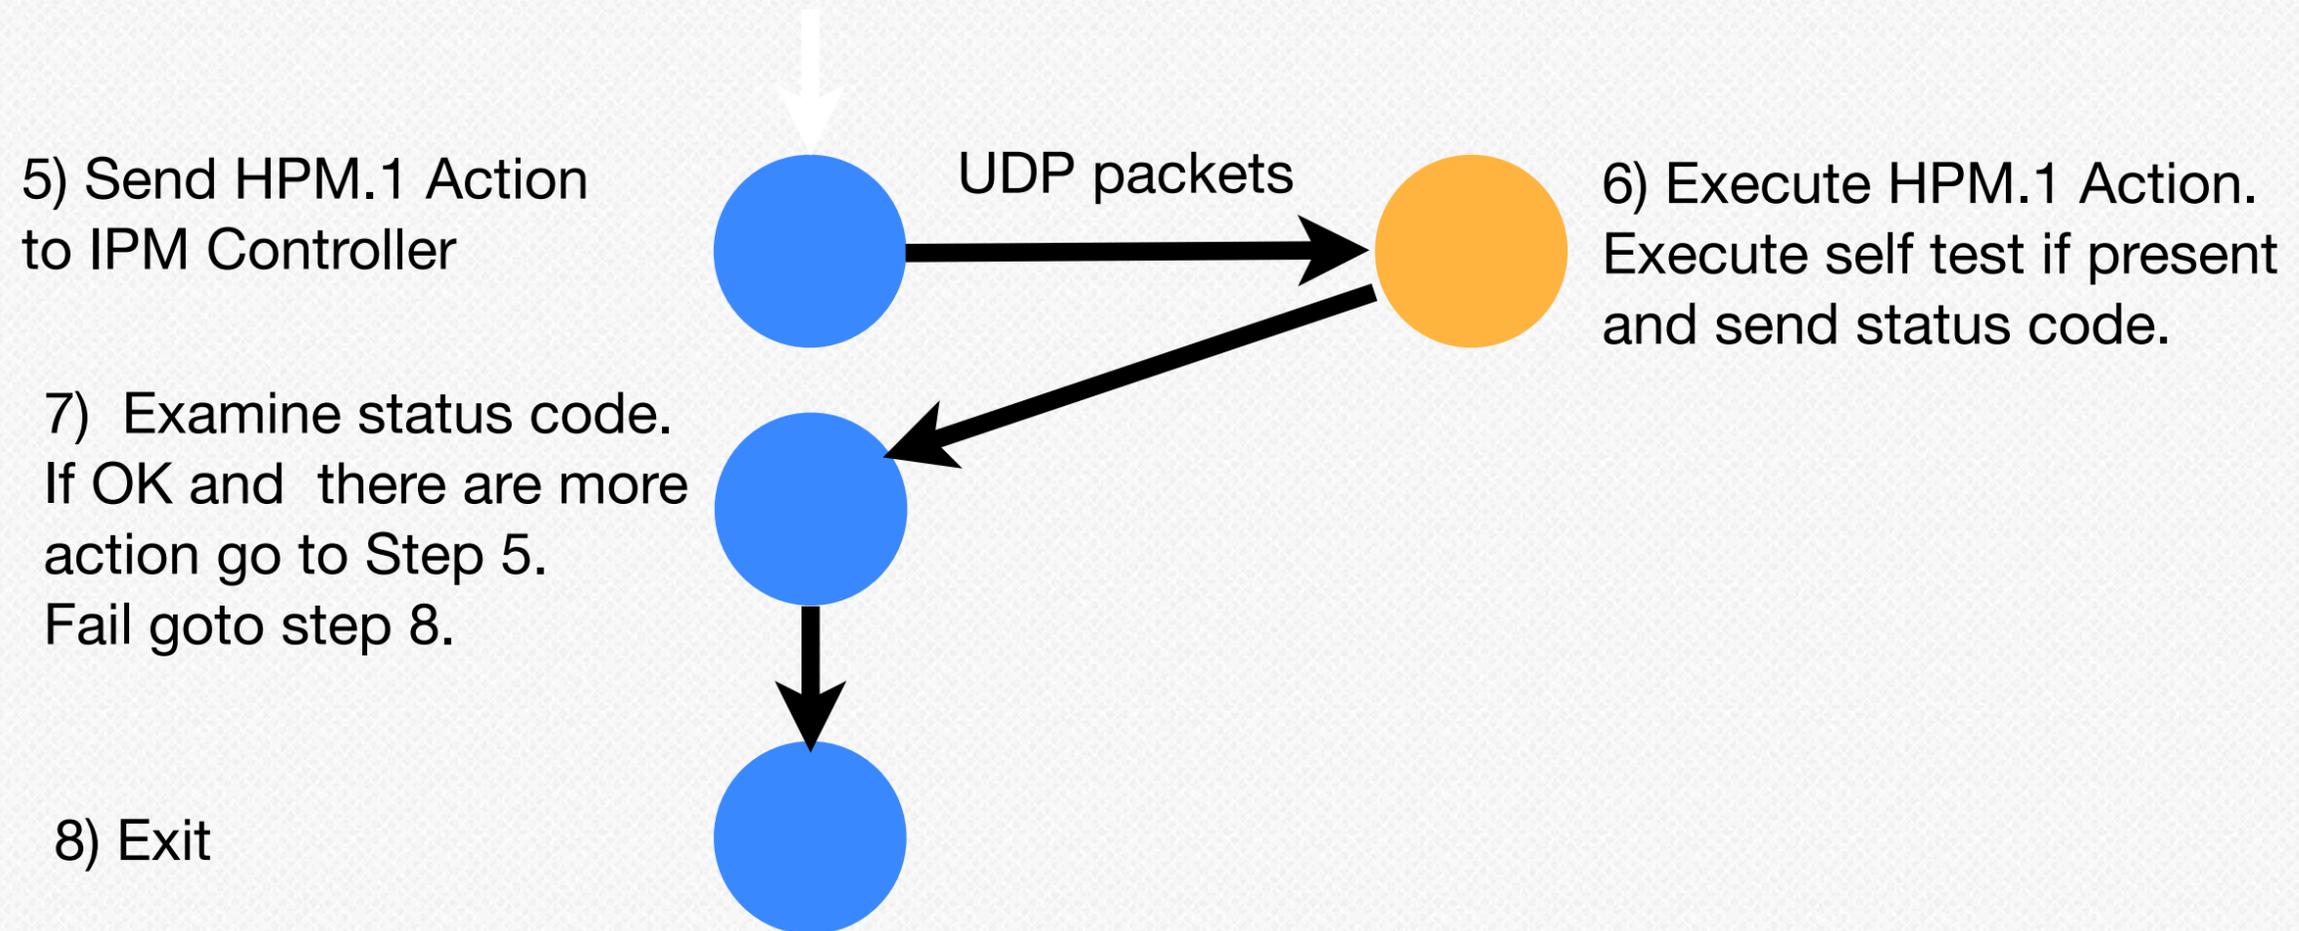
- Contained in the server/RAID/Switch
- Received HPM.1 Actions which prepare the device for upgrade, write firmware to storage or test the firmware.
- Executes HPM.1 rollback functions to recover from when the firmware test fails

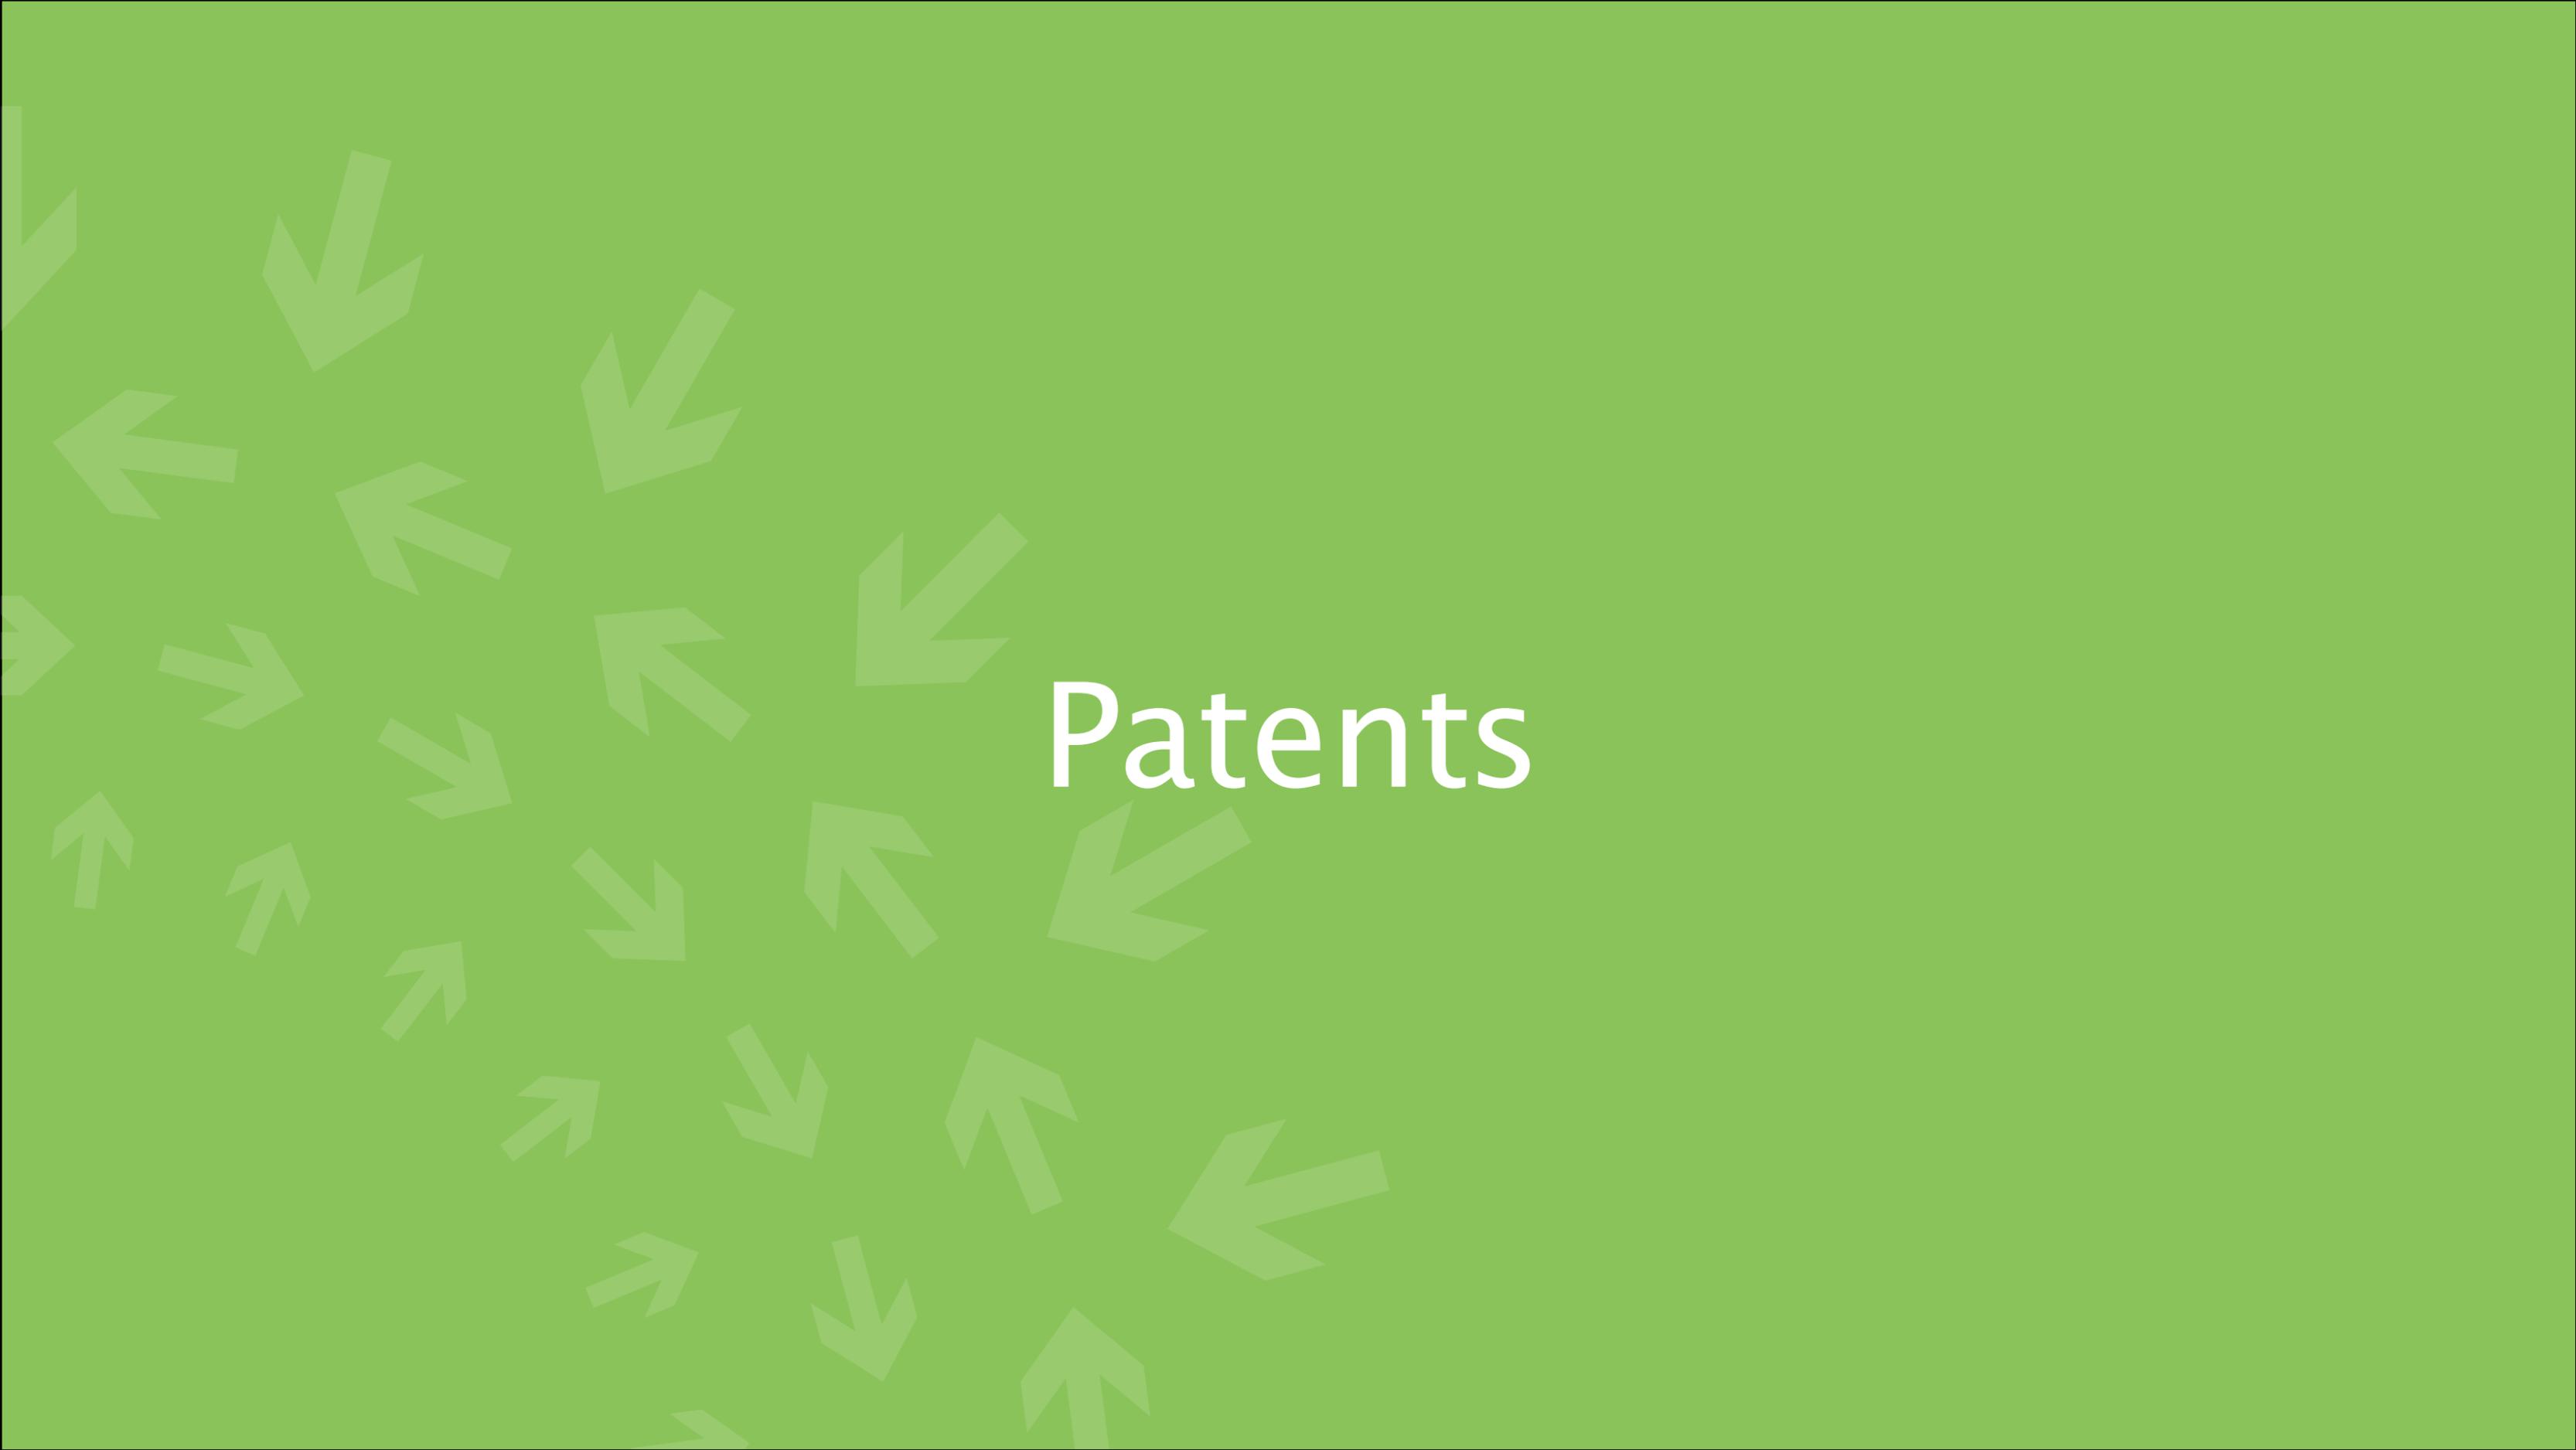


Upgrade Process



Upgrade Process





Patents

HPM.1 Patent Status

- To my knowledge HPM.1 does not infringe on any patents in any country
- Since the introduction of HPM.1 in 2007 there have been no allegations that it infringes on any patent



NetApp Patent

- The patent 7,484,084 January 27, 2009 by Network Appliances requires:
 - A BMC
 - A NIC with DMC(to send the firmware image to memory)
 - A host operating system
- Firmware update procedures must be designed so they do not infringe
- HPM.1 does not need a Host OS. No provision in '084 patent for upgrading firmware without a Host OS



Avoiding '084

- Monolithic servers using a Disposable OS and a NIC doing DMA of the firmware may infringe
- Monolithic servers updating firmware with a Disposable OS have both BMC and an Host OS fulfilling 2 of the '084 requirements
- Only thing missing is a NIC that does DMA
- It may be possible to avoid '084 when a DMA enabled NIC is present but it is hard

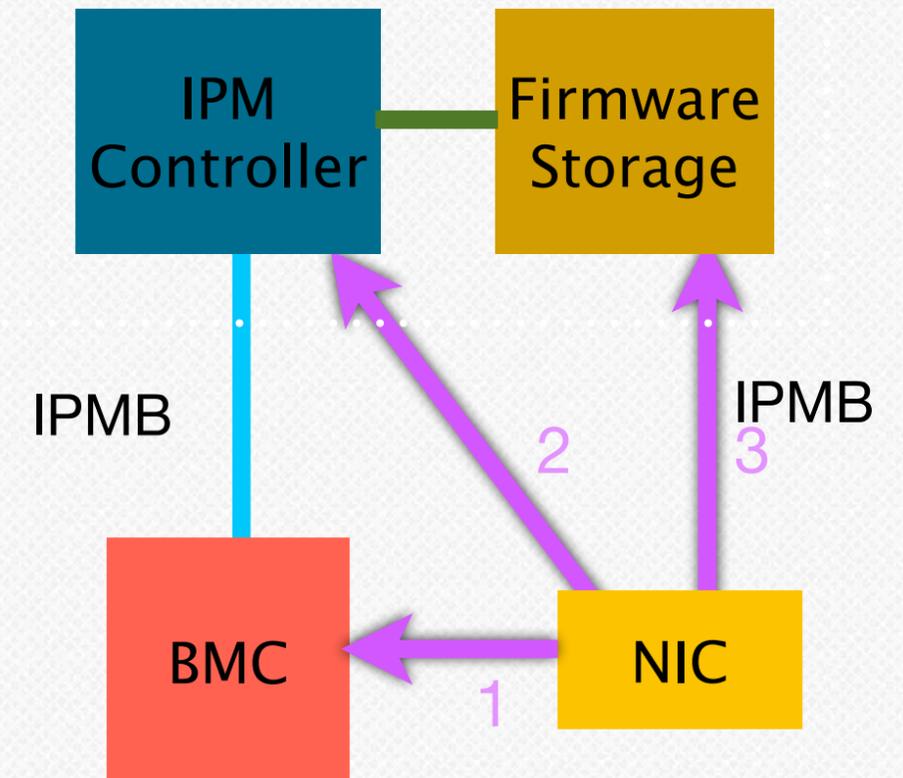


Firmware update without HPM.1

- You can design around '084 with or without HPM.1
- '084 mandates a BMC
- Options
 - Each node has a IPM Controller not a BMC
 - Don't implement one of the required features from IPMI 2.0 Table 3-1
 - Don't have an Operating System present

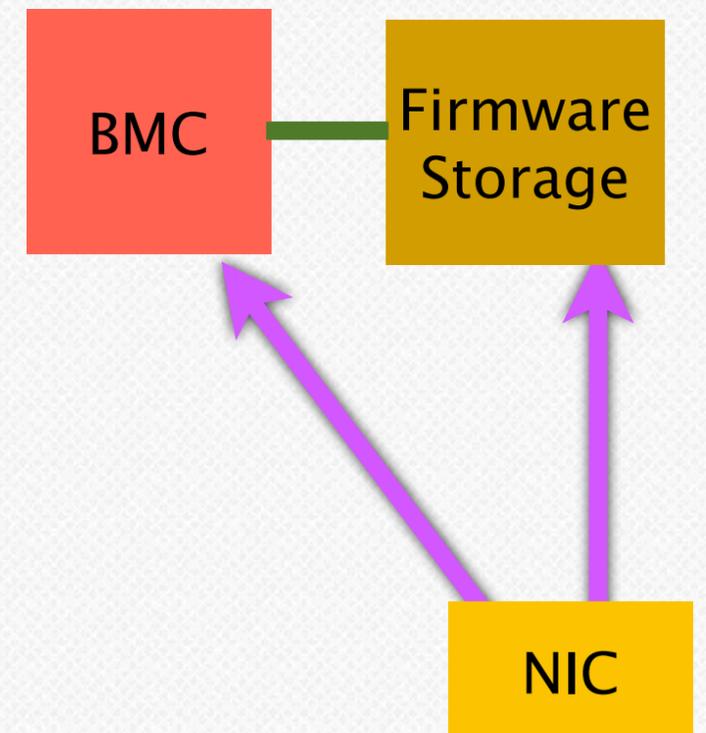
HPM.1 is not necessary to avoid '084

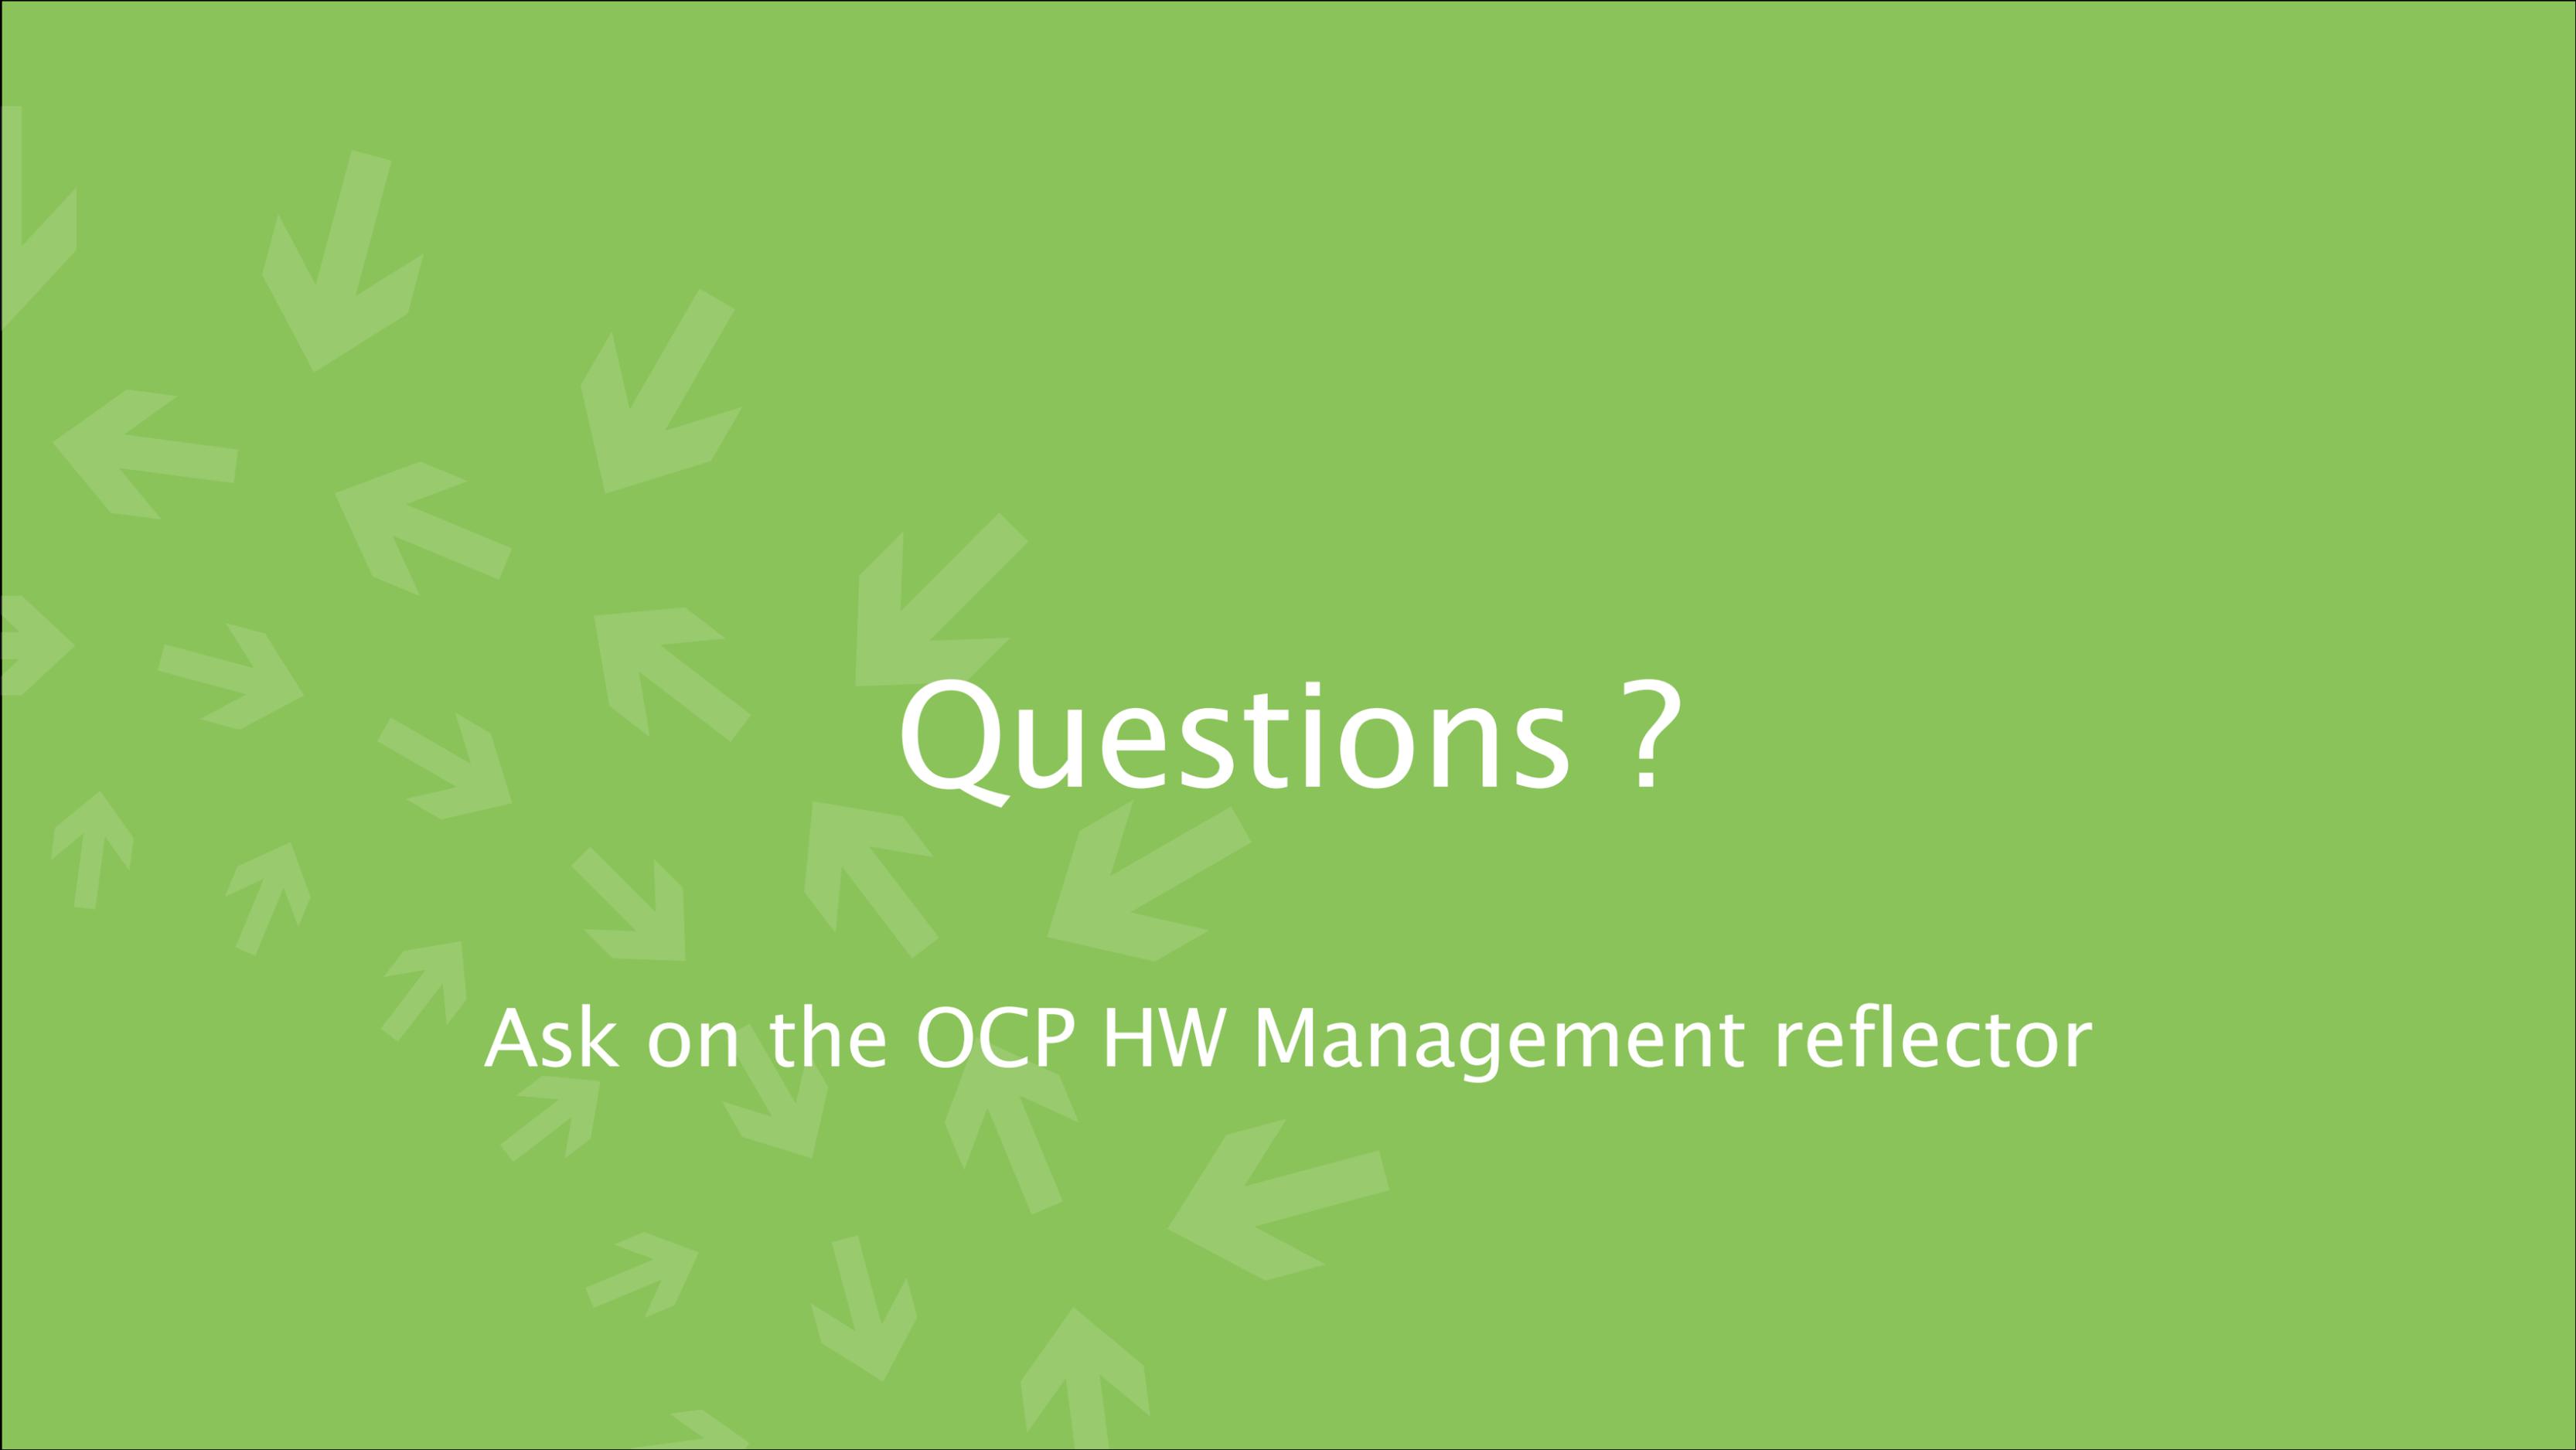
No BMC on Node. Uses IPM Controller



Firmware update with HPM.1

- Design around '084 with or without HPM.1
- Each node has a IPM Controller not a BMC
- Nodes share a physical or logical NIC that can DMA firmware into storage, RAM or the BMC





Questions ?

Ask on the OCP HW Management reflector