Hardware Discovery and Maintenance Workflows in Shasta Systems

May 9, 2019

{spresser,bshields}@cray.com

CRAY

CONTENT

- Purpose
  - To introduce the hardware management concepts and workflows expected to be used in Shasta
- Background
- Services
- Workflow examples
- Summary
- Q&A


Background

HARDWARE

- Mountain
  - Highly integrated
  - Scale-optimized cabinets for density, cooling, and high network bandwidth
- River
  - COTS (Common/Commercial Off The Shelf) hardware
  - Flexible support for arbitrary nodes

SPECIFICATIONS USED

- Redfish
- SNMP (Simple Network Management Protocol)
- IPMI (Intelligent Platform Management Interface)
- SSDP (Simple Service Discovery Protocol)

TASKS

- Endpoint Discovery
  - Finding hardware on the network
- Initialization
  - Configuring hardware to work with the system
- Geolocation
  - Determining where hardware is physically located
- Inventory Discovery
  - Determining the components in the individual compute nodes


Services

SERVICES

- IDEALS (Ideal Design of Equipment, Architecture and Layout Service)
  - Describes the system “as designed”
- HSM (Hardware State Manager)
  - Manages system state – power, hardware present and missing, etc.
- MEDS (Mountain Endpoint Discovery Service)
  - Manages endpoint discovery and geolocation for Mountain hardware
- REDS (River Endpoint Discovery Service)
  - Manages endpoint discovery, initialization and geolocation for River hardware

SERVICE INTERACTIONS

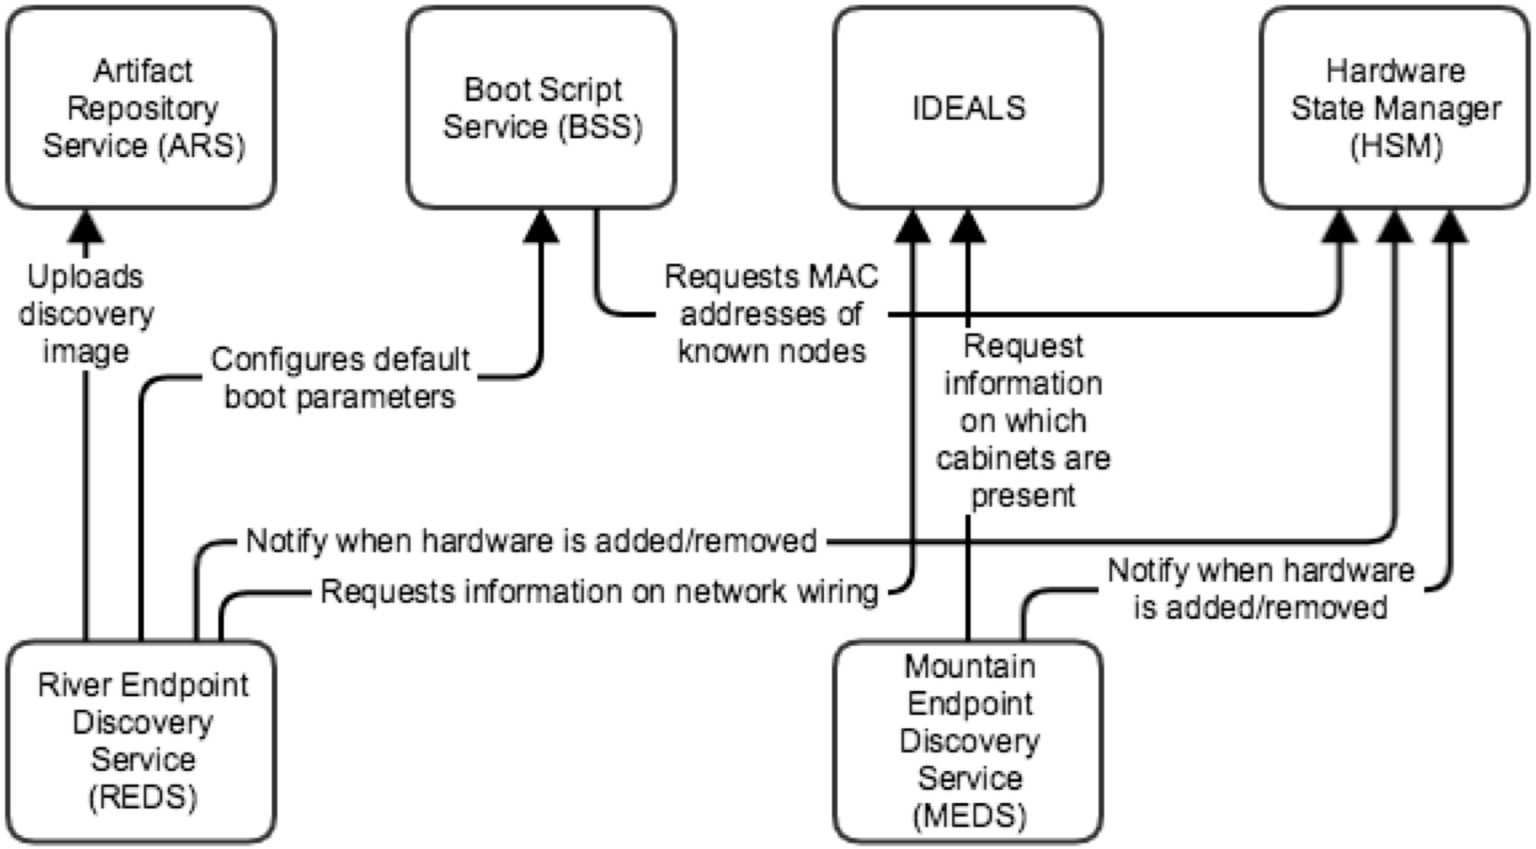

MEDS

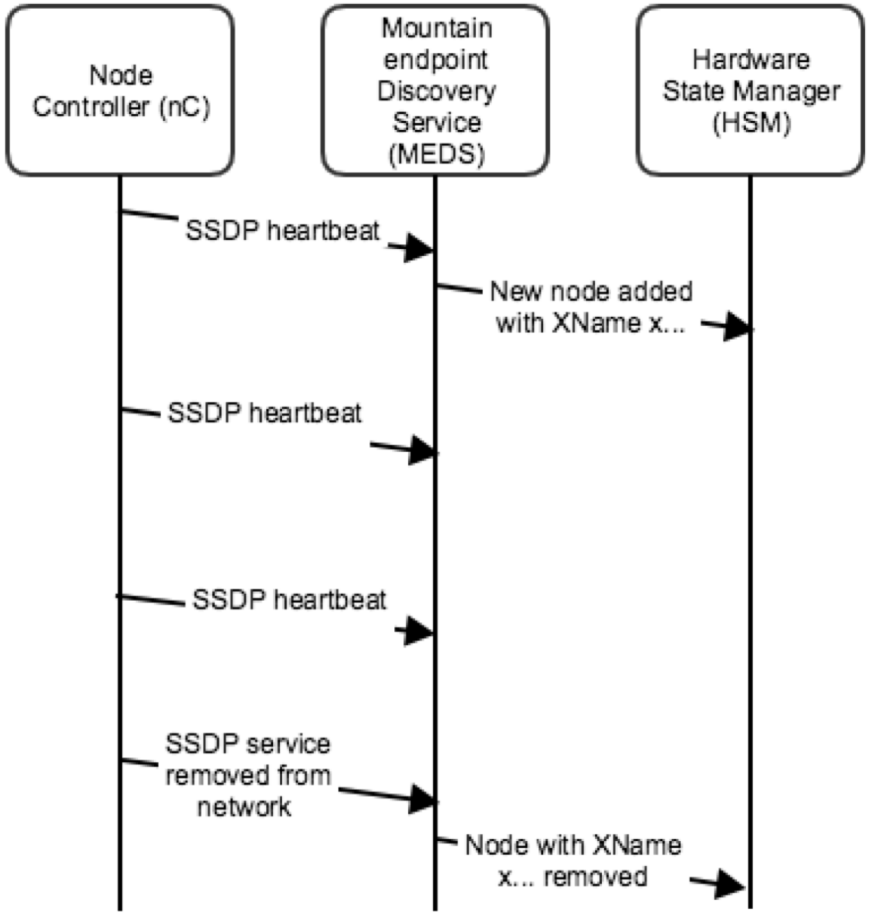

REDS

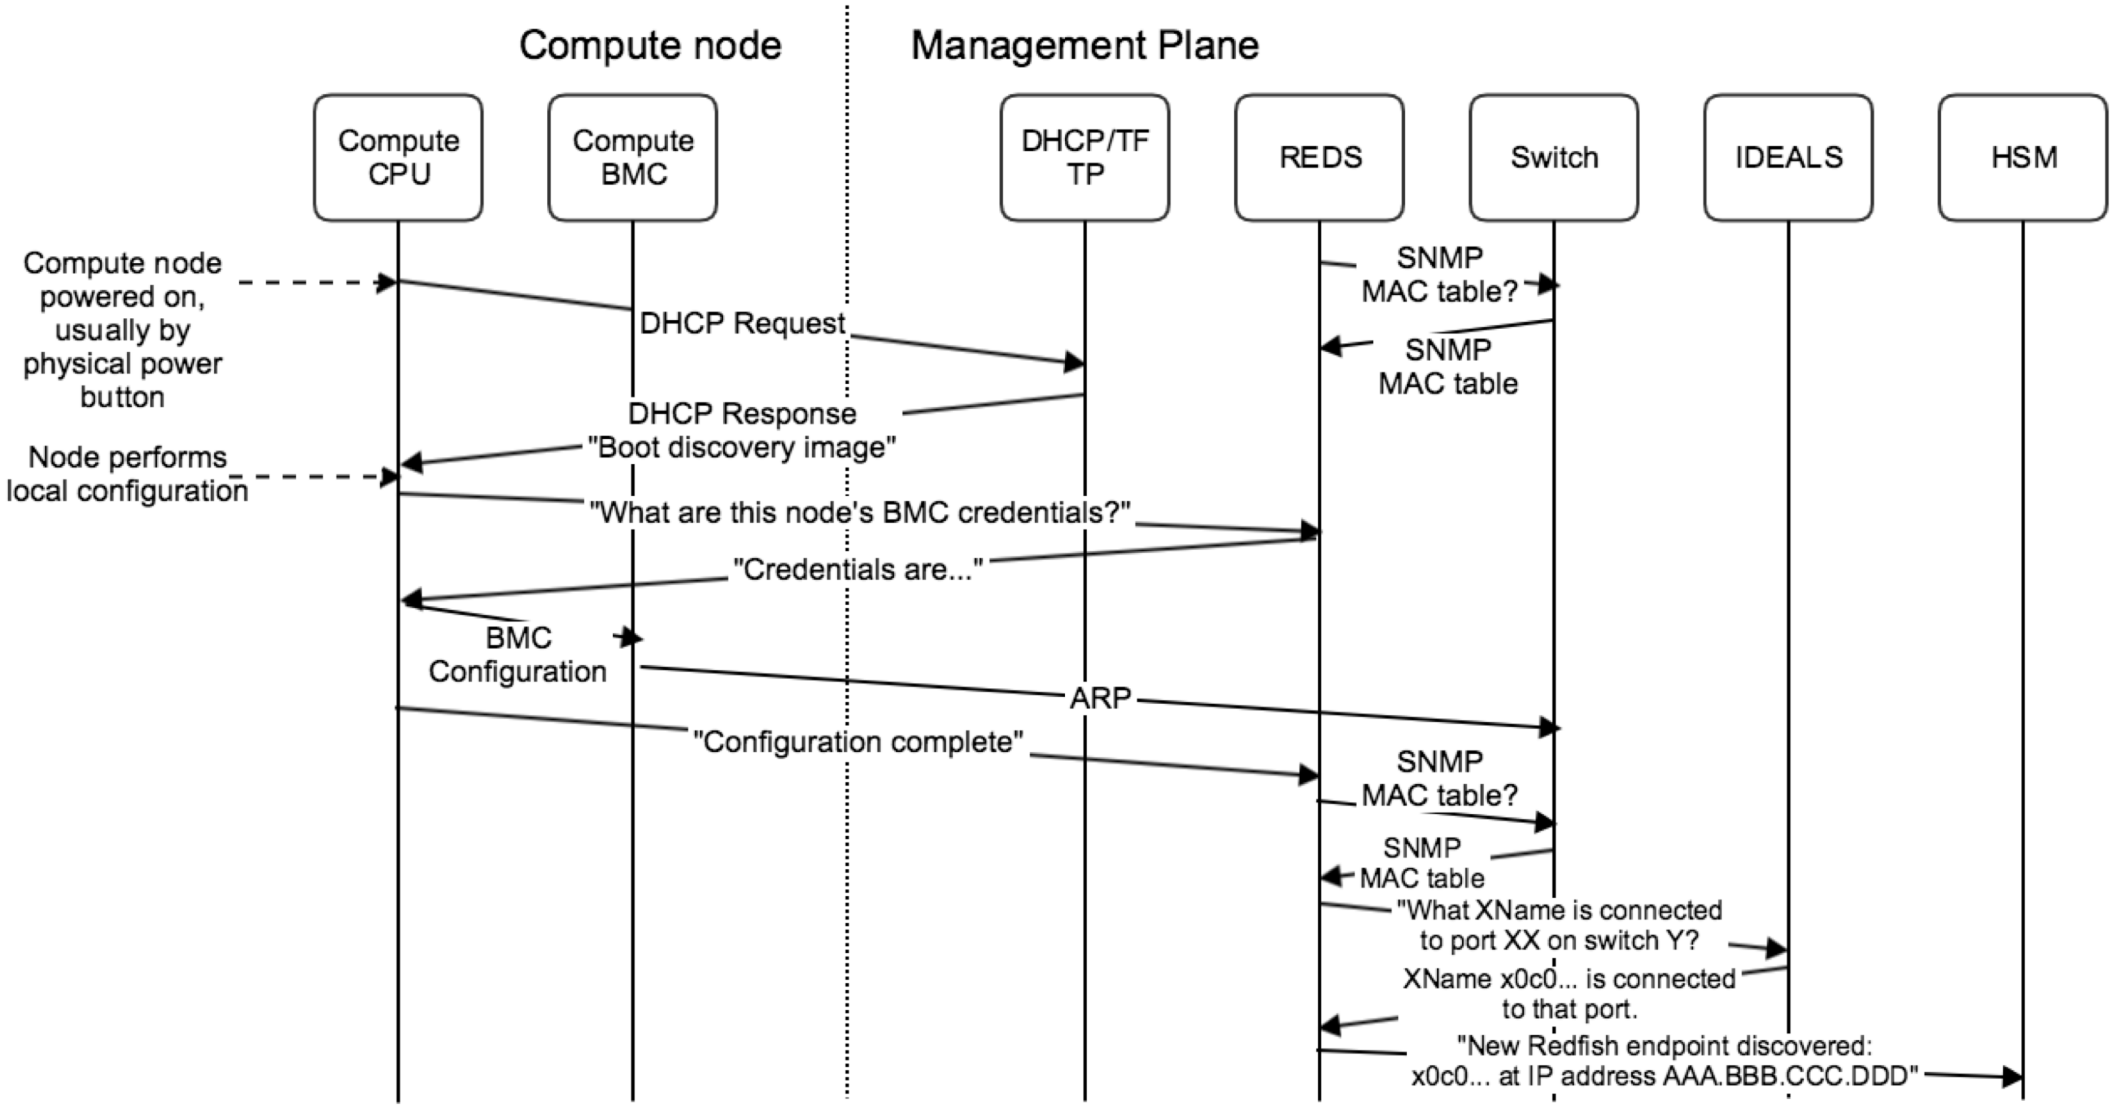


Workflows

NEW NODE

- Plug it in
- If it’s River, power it on, wait for it to power off

REMOVE A NODE

- Verify node drained in workload manager
- Power off, unplug and remove

REPLACE NODE

- Combine “Remove a node” with “New node”

INITIAL DISCOVERY

- Focus on first power-on of compute nodes in new system
- Computes need to run discovery when first powered on
- Power on hardware (may happen automatically; otherwise do manually)
- Hardware discovery should be automatic
- Note any hardware that fails discovery, but will not block use of the system

EXPAND/CONTRACT SYSTEM

- Update IDEALS
- Follow “New node” for new hardware
- Follow ”Remove a node” for removing hardware

COMPARISON TO XC SYSTEMS

- XC required running xtdiscover on hardware change
  - Two-hour runtime
  - All hardware had to ”bounce” correctly during discovery
  - Discovery couldn’t run while system in use
  - Difficult or impossible to incrementally build system
- We attempted to address all these concerns


Summary

WORKFLOWS

- Designed to be simple
- Discovery processes are on-going and as independent as possible
- No major discovery process when building system
- Tried to learn lessons from previous solutions, we think this is much easier

SAFE HARBOR STATEMENT

This presentation may contain forward-looking statements that are based on our current expectations. Forward looking statements may include statements about our financial guidance and expected operating results, our opportunities and future potential, our product development and new product introduction plans, our ability to expand and penetrate our addressable markets and other statements that are not historical facts.

These statements are only predictions and actual results may materially vary from those projected. Please refer to Cray's documents filed with the SEC from time to time concerning factors that could affect the Company and these forward-looking statements.

THANK YOU

QUESTIONS?

{spresser,bshields}@cray.com

cray.com

@cray_inc

linkedin.com/company/cray-inc-/