

CRAY

**Redfish APIs on Next Generation Cray
Hardware
CUG 2018**

Steven J. Martin, Cray Inc.



Modernizing Cray Systems Management Use of Redfish APIs on Next Generation Cray Hardware

Steven Martin, David Rush,
Kevin Hughes, Matt Kelly - Cray Inc.
{stevem, rushd, khughes, mpkelly}@cray.com

Please hold questions for the end of the presentation!



COMPUTE

STORE

ANALYZE

Modernizing Cray Systems Management



- **This presentation will cover the following**
 - Motivation to change Cray system management
 - Quick intro to DMTF and Redfish
 - Endpoint discovery
 - Redfish control
 - Redfish telemetry
 - Redfish and network boot
 - Cray and 3rd party tool integration



Motivation to Change Cray System Management

- **Requirements and expectations have evolved**
 - Proprietary management systems seen as limiting going forward
 - Customers are asking for open interfaces
 - Common interfaces for custom and COTS (commercial off-the-shelf) hardware
- **Industry standards are now more capable**
 - Redfish is replacing IPMI as the low level interface
 - New RESTful API enable discoverable capabilities
 - Open and documented APIs
 - JSON data formats human and machine friendly

Introduction to DMTF and Redfish



- **Distributed Management Task Force (DMTF)**

- The DMTF creates open manageability standards spanning diverse emerging and traditional IT infrastructures including cloud, virtualization, network, servers and storage.

- **Redfish:**

- DMTF's Redfish® is a standard API designed to deliver simple and secure management for converged, hybrid IT and the Software Defined Data Center (SDDC).
- Now also [ISO/IEC 30115:2018](#)





Endpoint Discovery

- **Controller IP addressing**

- Support customer assigned addresses
- Supports IPv4/IPv6
- 'explicit' endpoint discovery
- Expected to integrate with customer networks

- **Name Resolution**

- Redfish is HTTPS, requires X.509 Certificates
- Clients ***MUST*** validate server (BMC) certificate FQDN
- Prevent sending HTTP Basic credentials imposter



Endpoint Discovery (continued)


- **COTS rackmount equipment**

- Barcode scan of MAC addresses
- Query MAC addresses from Ethernet Switch ports
- SSDP, only answers "I'm somewhere over here"
 - COTS BMC is not location aware

- **Custom compute equipment**

- Embedded Ethernet switches, can't map 1:1 BMC to switch ports
- SSDP, answers "I am exactly here..."
 - BMC is location aware

Redfish Control



- **Commercial off-the-shelf (COTS) support Redfish**

- COTS systems "#ComputerSystem.Reset"
 - Both support:
 - On, ForceOff, GracefulRestart, Nmi, PushPowerButton
 - The "COTS-A" system adds:
 - GracefulShutdown, ForceRestart, ForceOn

- **Cray custom (optimized) systems supporting Redfish**

- Same Redfish APIs implemented in the Cray firmware
- Common higher level management tools

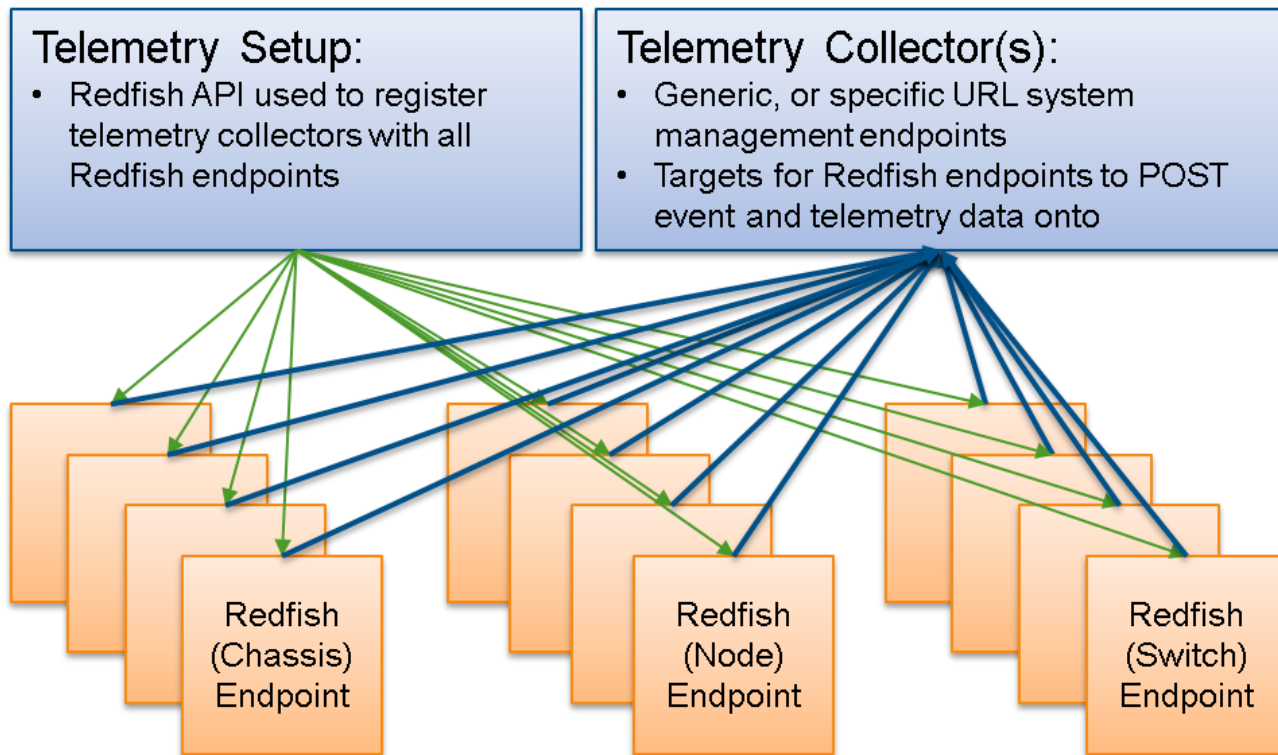


Redfish Control: Chassis Enclosure

```
sms:~user> curl -u root:XXXXXXXXX https://cmm7/redfish/v1/Chassis/Enclosure
```

```
{ ...  
  "Actions": {  
    "#Chassis.Reset": {  
      "ResetType@Redfish.AllowableValues": [ "On", "Off", "ForceOff" ],  
      "target": "/redfish/v1/Chassis/Enclosure/Actions/Chassis.Reset"  
    }  
  },  
  "ChassisType": "Enclosure", "Id": "Enclosure",  
  "Links": {  
    "Contains": [  
      {"@odata.id": "/redfish/v1/Chassis/B0"}, {"@odata.id": "/redfish/v1/Chassis/M2"},  
      {"@odata.id": "/redfish/v1/Chassis/B1"}, {"@odata.id": "/redfish/v1/Chassis/M1"},  
      {"@odata.id": "/redfish/v1/Chassis/M0"}, {"@odata.id": "/redfish/v1/Chassis/B2"},  
      ...  
    ], "ManagedBy": [{ "@odata.id": "/redfish/v1/Managers/BMC" }]  
  }, "Manufacturer": "Cray Inc", "Name": "Enclosure",  
  "PowerState": "Off", "Status": { "State": "Enabled" }  
}
```

Redfish Telemetry



COMPUTE

STORE

ANALYZE

Redfish Telemetry: EventService Subscribe



```
sms~user> curl -u root:XXXXXXXX -X \  
      POST https://COTS-B/redfish/v1/EventService/Subscriptions/  
{  
  "Destination": "http://SMS-Telemetry/TelemetryEndpoint",  
  "Context": "power",  
  "EventTypes": [  
    "Alert",  
    "Status Change",  
    "Value"  
  ],  
  "Protocol": "Redfish"  
}
```

Redfish and Network Boot



- **Cray custom compute equipment**
 - Contains node management Ethernet
 - UEFI IPv4/IPv6 Network Boot
 - Supports commonly used network bootstrap programs PXE and iPXE
- **Flexibility to boot over HSN or NMN**
 - High speed network (HSN)
 - Node management network (NMN)
- **Redfish controls to select boot device/network**

Cray and 3rd Party Tool Integration



- **Cray tools will simplify administration tasks**
 - Orchestrate parallel Redfish operations for scaling
 - Manage telemetry collection
- **Flexible environment supporting**
 - Cray tools
 - 3rd party tools
 - Open source tools (DMTF, and other)
 - Customer developed tools

Cray and 3rd Party Tool Integration



- **Redfish Tools Released on GitHub**

- “DMTF invites review and open source contributions from the industry in the collaborative community environment of its public GitHub repo”.
- <https://www.dmtf.org/content/new-redfish-tools-released-github>

- <https://github.com/DMTF/Redfishtool>

- Handy CLI (developed by the DMTF) Requires Python 3
- Helps check compliance with the standard



Wrap Up

- **Cray is making fundamental changes in systems management**
 - Moving to Redfish as the hardware interface
 - Open well documented REST APIs
 - Support for open interfaces and tools
- **We are working on this now, and interested in your feedback!**

Q&A

Steven J. Martin
stevem@cray.com

Legal Disclaimer



Information in this document is provided in connection with Cray Inc. products. No license, express or implied, to any intellectual property rights is granted by this document.

Cray Inc. may make changes to specifications and product descriptions at any time, without notice.

All products, dates and figures specified are preliminary based on current expectations, and are subject to change without notice.

Cray hardware and software products may contain design defects or errors known as errata, which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Cray uses codenames internally to identify products that are in development and not yet publicly announced for release. Customers and other third parties are not authorized by Cray Inc. to use codenames in advertising, promotion or marketing and any use of Cray Inc. internal codenames is at the sole risk of the user.

Performance tests and ratings are measured using specific systems and/or components and reflect the approximate performance of Cray Inc. products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance.

The following are trademarks of Cray Inc. and are registered in the United States and other countries: CRAY and design, SONEXION, URIKA and YARCDATA. The following are trademarks of Cray Inc.: CHAPEL, CLUSTER CONNECT, CLUSTERSTOR, CRAYDOC, CRAYPAT, CRAYPORT, DATAWARP, ECOPHLEX, LIBSCI, NODEKARE, REVEAL. The following system family marks, and associated model number marks, are trademarks of Cray Inc.: CS, CX, XC, XE, XK, XMT and XT. The registered trademark LINUX is used pursuant to a sublicense from LMI, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis. Other trademarks used on this website are the property of their respective owners.