



Document Number: DSP0266

Date: 2015-04-21

Version: 0.96.0

Redfish Scalable Platforms Management API Specification

Information for Work-in-Progress version:

IMPORTANT: This document is not a standard. It does not necessarily reflect the views of the DMTF or all of its members. Because this document is a Work in Progress, it may still change, perhaps profoundly. This document is available for public review and comment until the stated expiration date.

It expires on: 2015-08-24

Provide any comments through the DMTF Feedback Portal: <http://www.dmtf.org/standards/feedback>

Document Type: Specification

Document Status: Work in Progress

Document Language: en-US

Copyright Notice

Copyright © 2014-2015 Distributed Management Task Force, Inc. (DMTF). All rights reserved.

DMTF is a not-for-profit association of industry members dedicated to promoting enterprise and systems management and interoperability. Members and non-members may reproduce DMTF specifications and documents, provided that correct attribution is given. As DMTF specifications may be revised from time to time, the particular version and release date should always be noted.

Implementation of certain elements of this standard or proposed standard may be subject to third party patent rights, including provisional patent rights (herein "patent rights"). DMTF makes no representations to users of the standard as to the existence of such rights, and is not responsible to recognize, disclose, or identify any or all such third party patent right, owners or claimants, nor for any incomplete or inaccurate identification or disclosure of such rights, owners or claimants. DMTF shall have no liability to any party, in any manner or circumstance, under any legal theory whatsoever, for failure to recognize, disclose, or identify any such third party patent rights, or for such party's reliance on the standard or incorporation thereof in its product, protocols or testing procedures. DMTF shall have no liability to any party implementing such standard, whether such implementation is foreseeable or not, nor to any patent owner or claimant, and shall have no liability or responsibility for costs or losses incurred if a standard is withdrawn or modified after publication, and shall be indemnified and held harmless by any party implementing the standard from any and all claims of infringement by a patent owner for such implementations.

For information about patents held by third-parties which have notified the DMTF that, in their opinion, such patent may relate to or impact implementations of DMTF standards, visit <http://www.dmtf.org/about/policies/disclosures.php>.

CONTENTS

- 1. Scope
- 2. Normative References
- 3. Terms and Definitions
- 4. Symbols and Abbreviated Terms
- 5. Overview
 - 5.1. Principal Goals & Scope
 - 5.1.1. REST based
 - 5.1.2. Follow OData Conventions
 - 5.1.3. Model Oriented
 - 5.1.4. Separation of Protocol from Data Model
 - 5.1.5. Hypermedia API Service Endpoint
 - 5.1.6. Scope
 - 5.1.7. Limitations
 - 5.2. Service Elements
 - 5.2.1. Synchronous and Asynchronous Operation Support
 - 5.2.2. Eventing Mechanism
 - 5.2.3. Actions
 - 5.2.4. Service Entry Point Discovery
 - 5.2.5. Remote Access Support
 - 5.3. Security
- 6. Protocol Details
 - 6.1. Use of HTTP
 - 6.1.1. URIs
 - 6.1.2. HTTP Methods
 - 6.1.3. HTTP Redirect
 - 6.1.4. Media Types
 - 6.1.5. ETags
 - 6.2. Protocol Version
 - 6.3. Requests
 - 6.3.1. Request Headers
 - 6.3.2. Read Requests (GET)
 - 6.3.2.1. Service Root Request
 - 6.3.2.2. Metadata Document Request
 - 6.3.2.3. OData Service Document Request
 - 6.3.2.4. Resource Retrieval Requests
 - 6.3.2.4.1. Query Parameters
 - 6.3.2.4.2. Retrieving Collections
 - 6.3.3. HEAD
 - 6.3.4. Data Modification Requests
 - 6.3.4.1. Update (PATCH)
 - 6.3.4.2. Replace (PUT)
 - 6.3.4.3. Create (POST)
 - 6.3.4.4. Delete (DELETE)
 - 6.3.4.5. Actions (POST)
 - 6.4. Responses
 - 6.4.1. Response Headers
 - 6.4.2. Status Codes
 - 6.4.3. Metadata Responses
 - 6.4.3.1. Service Metadata
 - 6.4.3.1.1. Referencing OEM Extensions
 - 6.4.3.1.2. Annotations
 - 6.4.3.2. OData Service Document
 - 6.4.4. Resource Responses
 - 6.4.4.1. Context Property
 - 6.4.4.1.1. Select List
 - 6.4.4.2. Resource Identifier Property
 - 6.4.4.3. Type Property

- 6.4.4.4. ETag Property
 - 6.4.4.5. Primitive Properties
 - 6.4.4.5.1. DateTime Values
 - 6.4.4.6. Structured Properties
 - 6.4.4.7. Collection Properties
 - 6.4.4.8. Actions Property
 - 6.4.4.8.1. Action Representation
 - 6.4.4.8.2. Allowable Values
 - 6.4.4.9. Links Property
 - 6.4.4.9.1. Reference to a Single Related Resource
 - 6.4.4.9.2. Array of References to Related Resources
 - 6.4.4.10. OEM Property
 - 6.4.4.11. Extended Information
 - 6.4.4.11.1. Extended Object Information
 - 6.4.4.11.2. Extended Property Information
 - 6.4.4.12. Additional Annotations
 - 6.4.5. Resource Collections
 - 6.4.5.1. Context Property
 - 6.4.5.2. Resource Count Property
 - 6.4.5.3. Resource Members Property
 - 6.4.5.4. Partial Results
 - 6.4.5.5. Additional Annotations
 - 6.4.6. Error Responses
 - 6.4.6.1. Extended Information Object
7. Data Model & Schema
- 7.1. Type Identifiers
 - 7.1.1. Type Identifiers in JSON
 - 7.2. Common Naming Conventions
 - 7.3. Localization Considerations
 - 7.4. Schema Definition
 - 7.4.1. Common Annotations
 - 7.4.1.1. Description
 - 7.4.1.2. Long Description
 - 7.4.2. Schema Documents
 - 7.4.2.1. Referencing other Schemas
 - 7.4.2.2. Namespace Definitions
 - 7.4.3. Resource Type Definitions
 - 7.4.4. Resource Properties
 - 7.4.4.1. Property Types
 - 7.4.4.1.1. Primitive Types
 - 7.4.4.1.2. Structured Types
 - 7.4.4.1.3. Enums
 - 7.4.4.1.4. Collections
 - 7.4.4.2. Non-Nullable properties
 - 7.4.4.3. Read-only properties
 - 7.4.4.4. Required Properties
 - 7.4.4.5. Units of Measure
 - 7.4.4.6. Language-Dependent Property Values
 - 7.4.5. Reference Properties
 - 7.4.5.1. Expanded References
 - 7.4.5.2. Expanded Resources
 - 7.4.6. Resource Actions
 - 7.4.7. Resource Extensibility
 - 7.4.7.1. Property Extensions
 - 7.4.7.2. Custom Actions
 - 7.4.7.3. Custom Annotations
 - 7.5. Common Redfish Resource Properties
 - 7.5.1. Id

- 7.5.2. Name
- 7.5.3. Description
- 7.5.4. Modified
- 7.5.5. Status
- 7.5.6. Links
- 7.5.7. Actions
- 7.5.8. OEM
- 7.5.9. Settings
- 7.5.10. SettingsResult
- 7.6. Redfish Resources
 - 7.6.1. Current Configuration
 - 7.6.2. Settings
 - 7.6.3. Services
 - 7.6.4. Registry
- 7.7. Special Resource Situations
 - 7.7.1. Absent Resources
 - 7.7.2. Schema Variations
- 8. Service Details
 - 8.1. Eventing
 - 8.1.1. Event Message Subscription
 - 8.1.2. Event Message Objects
 - 8.1.3. Subscription Cleanup
 - 8.2. Asynchronous Operations
 - 8.3. Resource Tree Stability
 - 8.4. Discovery
 - 8.4.1. UPnP Compatibility
 - 8.4.2. USN Format
 - 8.4.3. M-SEARCH Response
 - 8.4.4. Notify, Alive, and Shutdown messages
- 9. Security
 - 9.1. Goals
 - 9.2. Protocols
 - 9.2.1. Discovery
 - 9.2.2. Transport
 - 9.2.2.1. TLS
 - 9.2.2.2. Cipher suites
 - 9.2.2.3. Certificates
 - 9.3. Sensitive Data
 - 9.4. Authentication
 - 9.4.1. HTTP Header Security
 - 9.4.1.1. HTTP Redirect
 - 9.4.2. Extended Error Handling
 - 9.4.3. HTTP Header Authentication
 - 9.4.3.1. BASIC authentication
 - 9.4.3.2. Digest authentication
 - 9.4.3.3. Negotiate
 - 9.4.3.4. Request / Message Level Authentication
 - 9.4.3.5. Certificate based authentication?
 - 9.4.4. Session Management
 - 9.4.4.1. Session Lifecycle Management
 - 9.4.4.2. Login Sessions
 - 9.4.4.3. Login
 - 9.4.4.4. Logout
 - 9.4.4.5. X-Auth-Token HTTP Header
 - 9.4.5. AccountService
 - 9.4.6. Async Tasks
 - 9.4.7. Event Subscriptions
 - 9.4.8. Privilege Model / Authorization

- 9.4.8.1. Profiles and Roles
- 9.4.9. Role Based Privilege
- 9.5. Data Model Validation
 - 9.5.1. Schema
- 9.6. Logging
 - 9.6.1. Required data for security log entries
 - 9.6.2. Completeness of Logging
 - 9.6.3. Content of Audit Logs
- 10. ANNEX A (informative)
 - 10.1. Change Log

Foreword

The Redfish Scalable Platform Management API ("Redfish") was prepared by the Scalable Platforms Management Forum of the DMTF.

DMTF is a not-for-profit association of industry members dedicated to promoting enterprise and systems management and interoperability. For information about the DMTF, see <http://www.dmtf.org>.

Acknowledgments

The DMTF acknowledges the following individuals for their contributions to this document:

- Jeff Autor - Hewlett-Packard Company
- David Brockhaus - Emerson Network Power
- Richard Brunner - VMware Inc.
- P Chandrasekhar - Dell Inc
- Chris Davenport - Hewlett-Packard Company
- Gamma Dean - Emerson Network Power
- Wassim Fayed - Microsoft Corporation
- Mike Garrett - Hewlett-Packard Company
- Steve Geffin - Emerson Network Power
- Jon Hass - Dell Inc
- Jeff Hilland - Hewlett-Packard Company
- Chris Hoffman - Emerson Network Power
- John Leung - Intel Corporation
- Milena Natanov - Microsoft Corporation
- Michael Pizzo - Microsoft Corporation
- Irina Salvan - Microsoft Corporation
- Hemal Shah - Broadcom Corporation
- Jim Shelton - Emerson Network Power
- Tom Slaughter - Intel Corporation
- Donnie Sturgeon - Emerson Network Power
- Pawel Szymanski - Intel Corporation
- Paul Vancil - Dell Inc

1. Scope

The Redfish Scalable Platforms Management API ("Redfish") is a new interface that uses RESTful interface semantics to access data defined in model format to perform out of band systems management. It is suitable for a wide range of servers, from stand-alone servers to rack mount and bladed environments but scales equally well for large scale cloud environments.

There are several out of band systems management standards (defacto and de jour) available in the industry. They all either vary widely in implementation, were developed for single server embedded environments or have their roots in antiquated software modeling constructs. There is no single industry standard that is simple to use, based on emerging programming standards, embedded friendly and capable of meeting large scale data center & cloud needs.

2. Normative References

The following referenced documents are indispensable for the application of this document. For dated or versioned references, only the edition cited (including any corrigenda or DMTF update versions) applies. For references without a date or version, the latest published edition of the referenced document (including any corrigenda or DMTF update versions) applies.

- IETF RFC 2616, R. Fielding et al., HTTP/1.1, <http://www.ietf.org/rfc/rfc2616.txt>
- IETF RFC 2617 J. Franks et al., HTTP Authentication: Basic and Digest Access Authentication, <http://www.ietf.org/rfc/rfc2617.txt>
- IETF RFC 3986 T. Berners-Lee et al, Uniform Resource Identifier (URI): Generic Syntax, <http://www.ietf.org/rfc/rfc3986.txt>
- IETF RFC 4627, D. Crockford, The application/json Media Type for JavaScript Object Notation (JSON), <http://www.ietf.org/rfc/rfc4627.txt>
- IETF RFC 4627, L. Dusseault et al, PATCH method for HTTP, <http://www.ietf.org/rfc/rfc5789.txt>
- IETF RFC 5988, M. Nottingham, Web linking, <http://www.ietf.org/rfc/rfc5988.txt>
- IETF RFC 6901, P. Bryan, Ed. et al, JavaScript Object Notation (JSON) Pointer, <http://www.ietf.org/rfc/rfc6901.txt>
- IETF RFC 6906, E. Wilde, The 'profile' Link Relation Type, <http://www.ietf.org/rfc/rfc6906.txt>
- ISO/IEC Directives, Part 2, Rules for the structure and drafting of International Standards, <http://isotc.iso.org/livelink/livelink.exe?func=ll&objId=4230456&objAction=browse&sort=subtypeH>
- JSON Schema, Core Definitions and Terminology, Draft 4 <http://tools.ietf.org/html/draft-zyp-json-schema-04.txt>
- JSON Schema, Interactive and Non-Interactive Validation, Draft 4 <http://tools.ietf.org/html/draft-fge-json-schema-validation-00.txt>
- OData Version 4.0 Part 1: Protocol. 24 February 2014. <http://docs.oasis-open.org/odata/odata/v4.0/os/odata-v4.0-part1-protocol.html>
- OData Version 4.0 Part 2: URL Conventions. 24 February 2014. <http://docs.oasis-open.org/odata/odata/v4.0/os/odata-v4.0-part2-url-conventions.html>
- OData Version 4.0 Part 3: Common Schema Definition Language (CSDL). 24 February 2014. <http://docs.oasis-open.org/odata/odata/v4.0/os/odata-v4.0-part3-csdl.html>
- OData Version 4.0: Core Vocabulary. 24 February 2014. <http://docs.oasis-open.org/odata/odata/v4.0/os/vocabularies/Org.OData.Core.V1.xml>
- OData Version 4.0 JSON Format. 24 February 2014. <http://docs.oasis-open.org/odata/odata-json-format/v4.0/os/odata-json-format-v4.0-os.html>
- OData Version 4.0: Units of Measure Vocabulary. 24 February 2014. <http://docs.oasis-open.org/odata/odata/v4.0/os/vocabularies/Org.OData.Measures.V1.xml>

3. Terms and Definitions

In this document, some terms have a specific meaning beyond the normal English meaning. Those terms are defined in this clause.

The terms "shall" ("required"), "shall not", "should" ("recommended"), "should not" ("not recommended"), "may", "need not" ("not required"), "can" and "cannot" in this document are to be interpreted as described in ISO/IEC Directives, Part 2, Annex H. The terms in parenthesis are alternatives for the preceding term, for use in exceptional cases when the preceding term cannot be used for linguistic reasons. Note that ISO/IEC Directives, Part 2, Annex H specifies additional alternatives. Occurrences of such additional alternatives shall be interpreted in their normal English meaning.

The terms "clause", "subclause", "paragraph", and "annex" in this document are to be interpreted as described in ISO/IEC Directives, Part 2, Clause 5.

The terms "normative" and "informative" in this document are to be interpreted as described in ISO/IEC Directives, Part 2, Clause 3. In this document, clauses, subclauses, or annexes labeled "(informative)" do not contain normative content. Notes and examples are always informative elements.

The following additional terms are used in this document.

Term	Definition
Collection	A Collection is a resource that acts as a container of other Resources. The members of a collection usually have similar characteristics. The container processes messages sent to the container. The members of the container process messages sent only to that member without affecting other members of the container.

Term	Definition
CRUD	Basic intrinsic operations used by any interface: Create, Read, Update and Delete.
Event	A record that corresponds to an individual alert.
Managed System	In the context of this specification, a managed system is a system that provides information or status, or is controllable, via a Redfish-defined interface.
Message	A complete request or response, formatted in HTTP/HTTPS. The protocol, based on REST, is a request/response protocol where every Request should result in a Response.
Operation	The HTTP request methods which map generic CRUD operations. These are POST, GET, PUT/PATCH, HEAD and DELETE.
OData	The Open Data Protocol, as defined in OData-Protocol .
OData Service Document	The name for a resource that provides information about the Service Root. The Service Document document provides a standard format for enumerating the resources exposed by the service that enables generic hypermedia-driven OData clients to navigate to the resources of the Redfish Service.
Redfish Alert Receiver	The name for the functionality that receives alerts from a Redfish Service. This functionality is typically software running on a remote system that is separate from the managed system.
Redfish Client	Name for the functionality that communicates with a Redfish Service and accesses one or more resources or functions of the Service.
Redfish Schema	The Schema definitions for Redfish resources. It is defined according to OData Schema notation that can be directly translated to a JSON Schema representation.
Redfish Service	Also referred to as the "Service". The collection of functionality that implements the protocols, resources, and functions that deliver an instantiation of a Redfish Schema and its associated behaviors for one or more managed systems.
Redfish Service Entry Point	Also referred to as "Service Entry Point". The interface through which a particular instance of a Redfish Service is accessed. A Redfish Service may have more than one Service Entry Point.
Request	A message from a Client to a Server. It consists of a request line (which includes the Operation), request headers, an empty line and an optional message body.
Resource	A Resource is addressable by a URI and is able to receive and process messages. A Resource can be either an individual entity, or a collection that acts as a container for several other entities.
Resource Tree	A Resource Tree is a tree structure of JSON encoded resources accessible via a well-known starting URI. A client may discover the resources available on a Redfish Service by following the resource links from the base of the tree. NOTE for Redfish client implementation: Although the resources are a tree, the references between resources may result in graph instead of a tree. Clients traversing the resource tree must contain logic to avoid infinite loops.
Response	A message from a Server to a Client in response to a request message. It consists of a status line, response headers, an empty line and an optional message body.
Service Root	The term Service Root is used to refer to a particular resource that is directly accessed via the service entry point. This resource serves as the starting point for locating and accessing the other resources and associated metadata that together make up an instance of a Redfish Service.
Subscription	A configuration setting in the event service that specifies where to send events.

4. Symbols and Abbreviated Terms

The following additional abbreviations are used in this document.

Term	Definition
CRUD	Create, Replace, Update and Delete
CSRF	Cross-Site Request Forgery

Term	Definition
HTTP	Hypertext Transfer Protocol
HTTPS	Hypertext Transfer Protocol over TLS
IP	Internet Protocol
JSON	JavaScript Object Notation
KVM-IP	Keyboard, Video, Mouse redirection over IP
NIC	Network Interface Card
PCI	Peripheral Component Interconnect
PCIe	PCI Express
TCP	Transmission Control Protocol
XSS	Cross-Site Scripting

5. Overview

The Redfish Scalable Platform Management API ("Redfish") is a management standard using a data model representation inside of a hypermedia RESTful interface. Because it is based on REST, Redfish is easier to use and implement than many other solutions. Since it is model oriented, it is capable of expressing the relationships between components in modern systems as well as the semantics of the services and components within them. It is also easily extensible. By using a hypermedia approach to REST, Redfish can express a large variety of systems from multiple vendors. By requiring JSON representation, a wide variety of resources can be created in a denormalized fashion not only to improve scalability, but the payload can be easily interpreted by most programming environments as well as being relatively intuitive for a human examining the data. The model is exposed in terms of an interoperable OData Schema, with the payload of the messages being expressed in JSON following OData JSON conventions. The schema (available in both XML and JSON formats) includes annotations to facilitate automatic translation of the schema to JSON Schema. The ability to externally host the Schema definition of the resources in a machine-readable format allows the meta data to be associated with the data without encumbering Redfish services with the meta data, thus enabling more advanced client scenarios as found in many data center and cloud environments.

5.1. Principal Goals & Scope

There are many principles and goals of Redfish as an architecture, a protocol and a data representation. It is intended that this architecture support a wide variety of systems found in service today - from stand alone machines to racks of equipment found in cloud service environment. Extensibility is a key goal, as is forward compatibility and deterministic behavior. Leveraging the protocols and standards widely accepted and used in environments today is a key strategy to achieve these goals. Simplicity to the extent possible is another goal, achieved by making as few operations and as few instances as possible in the model. Matching the programming environments that are being widely adopted today is another goal.

The following design tenets govern the design of the Redfish Scalable Platform Management API, with the key goals elaborated below:

- RESTful interface using a JSON payload and entity data model
- Separation of protocol from data model, allowing them to be revised independently
- Specified versioning rules for protocol and schema
- Leverage strength of internet protocol standards where it meets architectural requirements, such as JSON, HTTP, OData, and the RFCs referenced by this document.
- Focused on scalable environments but capable of managing current server environments
- Focus on out-of-band access -- implementable on existing BMC and firmware products
- Present value-add features alongside standardized items
- Functionality must be usable by non-computer-science professionals
- Data definitions as obvious in context as possible
- Opaque view of implementation architecture

5.1.1. REST based

This document defines a RESTful interface. Many service applications are exposed RESTful interfaces.

There are several reasons to define a RESTful interface:

- It enables a light weight implementation, where economy is a necessity (smaller data transmitted than SOAP, fewer layers to the protocol than WS-Man).
- It is on a trajectory to become a prevalent access method in the industry.
- It is easy to learn and easy to document.
- There are a number of toolkits & development environments that can be used for REST.
- It supports data model semantics and maps easily to the common CRUD operations.
- It fits with our design principle of simplicity.
- It is equally applicable to software application space as it is for embedded environments thus enabling convergence and sharing of code of components within the management ecosystem.
- It is schema agnostic so adapts well to any modeling language.
- By using it, Redfish can leverage existing security & discovery mechanisms in the industry.

5.1.2. Follow OData Conventions

With the popularity of RESTful APIs, there are nearly as many RESTful interfaces as there are applications. While following REST patterns helps promote good practices, due to design differences between the many RESTful APIs there is no interoperability between them.

OData defines a set of common RESTful conventions and markup which, if adopted, provides for interoperability between APIs.

Adopting OData conventions for describing schema, url conventions, and naming and structure of common properties in a JSON payload, not only encapsulate best practices for RESTful APIs but further enables Redfish services to be consumed by a growing ecosystem of generic client libraries, applications, and tools.

5.1.3. Model Oriented

Bitwise prior solutions (like IPMI) have difficulty showing relationships. Model orientation solves this. But current models that have evolved over time have become extremely complex, requiring many IOs to gather information. For that reason, as well as the complexity of their protocols and operations, they have received little implementation attention. Some have their roots in modeling multiple domains (printers, switches, software, etc). Additionally the expression of meta data in those models have only been adopted by niche markets.

The Redfish model is built for managing systems. All resources are defined in OData Schema and represented in JSON format. OData is an industry standard that encapsulates best practices for RESTful services and provides interoperability across services of different types. JSON is being widely adopted in multiple disciplines and has a large number of tools and programming languages that accelerate development when adopting these approaches.

5.1.4. Separation of Protocol from Data Model

The protocol operations are specified independently of the data model. The protocol is also versioned independently of the data model. The expectation is that the protocol version changes extremely infrequently, while the data model version is allowed to change as needed. This implies that innovation should happen primarily in the data model, not the protocol. It allows the data model to be extended and changed as needed without requiring the protocol or API version to change.

5.1.5. Hypermedia API Service Endpoint

Like other hypermedia APIs, Redfish has a single service endpoint URI and all other resources are accessible via opaque URIs referenced from the root. Any resource discovered through links found by accessing the root service or any service or resource referenced using references from the root service will conform to the same version of the protocol supported by the root service.

Note that the ServiceRoot Schema places requirements on the last segment of the path for the URIs discoverable through the service root.

5.1.6. Scope

The scope of this specification is to define the next generation systems management interface. This includes defining both the protocol and data model, as well as other architectural components needed for systems management environments.

Specifically, this document is intended to enable an open, industry-standard solution as proprietary or single-vendor efforts are not acceptable for target audience. The focus is on out-of-band access for large scale environments, though this architecture is capable of being the architectural successor to many of the current management standards.

5.1.7. Limitations

Redfish does not guarantee that client software will never need to be updated. Examples that may require updates include accommodation of new types of systems or their components, data model updates, and so on. System optimization for an application will always require architectural oversight. However, Redfish does attempt to minimize instances of forced upgrades to clients using Schemas, strict versioning and forward compatibility rules and through separation of the protocol from the data model.

Redfish does not enable a client to read a Resource Tree and write it to another Redfish Service. This is not possible as it is a hypermedia API. Only the root object has a well known URI. The resource topology reflects the topology of the system and devices it represents. Consequently, different server or device types will result in differently shaped resource trees, potentially even for identical systems from the same manufacturer.

Additionally, not all Redfish resources are simple read/write resources. Implementations may follow other interaction patterns discussed later. As an example, user credentials or certificates cannot simply be read from one service and transplanted to another. Another example is the use of Setting Data instead of writing to the same resource that was read from.

There is no raw/pass-thru interface as part of the standard.

5.2. Service Elements

5.2.1. Synchronous and Asynchronous Operation Support

While the majority of operations in this architecture are synchronous in nature, some operations can take a long time to execute, more time than a client typically wants to wait. For this reason, some operations can be asynchronous at the discretion of the service. The request portion of an asynchronous operation is no different from the request portion of a synchronous operation.

The use of HTTP Response codes enable a client to determine if the operation was completed synchronously or asynchronously. For more information see the section on [Tasks](#).

5.2.2. Eventing Mechanism

In some situations it is useful for a service to provide messages to clients that fall outside the normal request/response paradigm. These messages, called events, are used by the service to asynchronously notify the client of some significant state change or error condition, usually of a time critical nature.

Only one style of eventing is currently defined by this specification - push style eventing. In push style eventing, when the server detects the need to send an event, it uses an HTTP POST to push the event message to the client. Clients can subscribe to the eventing service to enable reception of events by creating a ListenerDestination subscription entry in the Event Service, or an administrator can create subscriptions as part of the Redfish service configuration. All subscriptions are persistent configuration settings.

Events originate from a specific resource. Not all resources are able to generate events. Those resources capable of generating events might not generate any events unless a subscription has been created to listen for the event. An administrator or client creates a subscription by sending a "subscribe" message to the Event Service. A subscribe message is sent using HTTP POST to the Event Subscriptions collection.

The Section on [Eventing](#) further in this specification discusses the details of the eventing mechanism.

5.2.3. Actions

Operations can be divided into two sets: intrinsic and extrinsic. Intrinsic operations, often referred to as CRUD, are mapped to [HTTP methods](#). The protocol also has the ability to support extrinsic operations -- those operations that do not map easily to CRUD. Examples of extrinsic would be items that collectively would be better performed if done as a set (for scalability, ease of interface, server side semantic preservation or similar reasons) or operations that have no natural mapping to CRUD operations. One examples is system reset. It is possible to combine multiple operations into a single action. A system reset could be modeled as an update to state, but semantically the client is actually requesting a state change and not simply changing the value in the state.

In Redfish, these extrinsic operations are called **actions** and are discussed in detail in different parts of this specification.

The Redfish Schema defines certain standard actions associated with [common Redfish resources](#). For these standard actions, the Redfish Schema contains the normative language on the behavior of the action. OEM extensions are also allowed to the [schema](#), including defining [actions](#) for existing resources.

5.2.4. Service Entry Point Discovery

While the service itself is at a well-known URI, the service host must be discovered. Redfish, like UPnP, uses SSDP for discovery. SSDP is supported in a wide variety of devices, such as printers. It is simple, lightweight, IPv6 capable and suitable for implementation in embedded environments. Redfish is investigating additional service entry point discovery (e.g. DHCP-based) approaches.

For more information, see the section on [Discovery](#)

5.2.5. Remote Access Support

A wide variety of remote access and redirection services are supported in this architecture. Critical to out-of-band environments are mechanisms to support Serial Console

access, Keyboard Video and Mouse re-direction (KVM-IP), Command Shell (i.e. Command Line interface) and remote Virtual Media. Support for Serial Console, Command Shell, KVM-IP and Virtual Media are all encompassed in this standard and are expressed in the Redfish Schema. This standard does not define the protocols or access mechanisms for accessing those devices and services. The Redfish Schema provides for the representation and configuration of those services, establishment of connections to enable those services and the operational status of those services. However, the specification of the protocols themselves are outside the scope of this specification.

5.3. Security

The challenge with security in a remote interface that is programmatic is to ensure both the interfaces used to interact with Redfish and the data being exchanged are secured. This means designing the proper security control mechanisms around the interfaces and securing the channels used to exchange the data. As part of this, specific behaviors are to be put in place including defining and using minimum levels of encryption for communication channels etc.

6. Protocol Details

The Redfish Scalable Platform Management API is based on REST and follows OData conventions for interoperability, as defined in [OData-Protocol](#), JSON payloads, as defined in [OData-JSON](#), and a machine-readable representation of schema, as defined in [OData-Schema](#). The schemas include annotations to enable direct translation to JSON Schema for validation and consumption by tools supporting JSON Schema. Following these common standards and conventions increases interoperability and enables leveraging of existing tool chains.

Redfish follows the OData minimal conformance level for clients consuming minimal metadata.

Throughout this document, we refer to Redfish as having a protocol mapped to a data model. More accurately, HTTP is the application protocol that will be used to transport the messages and TCP/IP is the transport protocol. The RESTful interface is a mapping to the message protocol. For simplicity though, we will refer to the RESTful mapping to HTTP, TCP/IP and other protocol, transport and messaging layer aspects as the Redfish protocol.

The Redfish protocol is designed around a web service based interface model, and designed for network and interaction efficiency for both user interface (UI) and automation usage. The interface is specifically designed around the REST pattern semantics.

[HTTP methods](#) are used by the Redfish protocol for common CRUD operations and to retrieve header information.

[Actions](#) are used for expanding operations beyond CRUD type operations, but should be limited in use.

[Media types](#) are used to negotiate the type of data that is being sent in the body of a message.

[HTTP status codes](#) are used to indicate the server's attempt at processing the request. [Extended error handling](#) is used to return more information than the HTTP error code provides.

The ability to send secure messages is important; the [Security](#) section of this document describes specific TLS requirements.

Some operations may take longer than required for synchronous return semantics. Consequently, deterministic [asynchronous semantic](#) are included in the architecture.

6.1. Use of HTTP

HTTP is ideally suited to a RESTful interface. This section describes how HTTP is used in the Redfish interface and what constraints are added on top of HTTP to assure interoperability of Redfish compliant implementations.

- A Redfish interface shall be exposed through a web service endpoint implemented using Hypertext Transfer Protocols, version 1.1 ([RFC2616](#)).

6.1.1. URIs

A URI is used to identify a resource, including the base service and all Redfish resources.

- A URI shall be a unique identifier to a resource.
- A URI shall be treated by the client as opaque, and thus should not be attempted to be understood or deconstructed by the client

To begin operations, a client must know the URI for a resource.

- Performing a GET operation yields a representation of the resource containing properties and links to associated resources.

The base resource URI is well known and is based on the protocol version. Discovering the URIs to additional resources is done through observing the associated resource links returned in previous responses. This type of API that is consumed by navigating URIs returned by the service is known as a Hypermedia API.

The URI is the primary unique identifier of resources. Redfish considers 3 parts of the URI as described in [RFC3986](#).

The first part includes the scheme and authority portions of the URI. The second part includes the root service and version. The third part is a unique resource identifier.

For example, in the following URL:

```
Example: https://mgmt.vendor.com/redfish/v1/Systems/1
```

- The first part is the scheme and authority portion (<https://mgmr.vendor.com>).
- The second part is the root service and version (/redfish/v1).
- The third part is the unique resource path (Systems/1).

The scheme and authority part of the URI shall not be considered part of the unique *identifier* of the resource. This is due to redirection capabilities and local operations which may result in the variability of the connection portion. The remainder of the URI (the service and resource paths) is what *uniquely identifies* the resource, and this is what is returned in all Redfish payloads.

- The unique identifier part of a URI shall be unique within the implementation.

For example, a POST may return the following URI in the Location header of the response (indicating the new resource created by the POST):

```
Example: /redfish/v1/Systems/2
```

Assuming the client is connecting through an appliance named "mgmt.vendor.com", the full URI needed to access this new resource is <https://mgmt.vendor.com/redfish/v1/Systems/2>.

URIs, as described in [RFC3986](#), may also contain a query (?query) and a frag (#frag) components. Queries are addressed in the section [Query Parameters](#). Fragments (frag) shall be ignored by the server when used as the URI for submitting an operation.

6.1.2. HTTP Methods

An attractive feature of the RESTful interface is the very limited number of operations which are supported. The following table describes the general mapping of operations to HTTP methods. If the value in the column entitled "required" has the value "yes" then the HTTP method shall be supported by a Redfish interface.

HTTP Method	Interface Semantic	Required
POST	Object create, Object action, Eventing	Yes
GET	Object or Collection retrieval	Yes
PUT	Object replace	No
PATCH	Object update	Yes
DELETE	Object delete	Yes
HEAD	Object or Collection header retrieval	No

Other HTTP methods are not allowed and shall receive a [405](#) response.

6.1.3. HTTP Redirect

HTTP redirect allows a service to redirect a request to another URL. Among other things, this enables Redfish resources to alias areas of the data model.

- All Redfish Clients shall correctly handle HTTP redirect.

NOTE: Refer to the [Security](#) section for security implications of HTTP Redirect

6.1.4. Media Types

Some resources may be available in more than one type of representation. The type of representation is indicated by the media type.

In HTTP messages the media type is specified in the Content-Type header. A client can tell a service that it wants the response to be sent using certain media types by setting the HTTP Accept header to a list of the acceptable media types.

- All resources shall be made available using the JSON media type "application/json".
- Redfish services shall make every resource available in a representation based on JSON, as specified in [RFC4627](#). Receivers shall not reject a message because it is encoded in JSON, and shall offer at least one response representation based on JSON. An implementation may offer additional representations using non-JSON media types.

Clients may request compression by specifying an [Accept-Encoding header](#) in the request.

- Responses to GET requests shall only be compressed if requested by the client.
- Services should support gzip compression when requested by the client.

6.1.5. ETags

In order to reduce the cases of unnecessary RESTful accesses to resources, the Redfish Service should support associating a separate ETag with each resource.

- Implementations should support returning [ETag properties](#) for each resource.
- Implementations should support returning ETag headers for each response that represents a single resource. Implementations shall support returning ETag headers for certain requests and responses as listed in the [Security](#) section.

The ETag is generated and provided as part of the resource payload because the service is in the best position to know if the new version of the object is different enough to be considered substantial. There are two types of ETags: weak and strong.

- Weak model -- only "important" portions of the object are included in formulation of the ETag. For instance, meta-data such as a last modified time should not be included in the ETag generation. The "important" properties that determine ETag change include writable settings and changeable attributes such as UUID, FRU data, serial numbers, etc.
- Strong model -- all portions of the object are included in the formulation of the ETag.

This specification does not mandate a particular algorithm for creating the ETag, but ETags should be highly collision-free. An ETag could be a hash, a generation ID, a time stamp or some other value that changes when the underlying object changes.

If a client PUTs or PATCHes a resource, it should include an ETag in the HTTP If-Match/If-None-Match header from a previous GET.

In addition to returning the ETag property on each resource,

- A Redfish Service should return the ETag header on client PUT/POST/PATCH
- A Redfish Service should return the ETag header on a GET of an individual resource

The format of the ETag header is:

```
ETag W/"<string>"
```

6.2. Protocol Version

The protocol version is separate from the version of the resources or the version of the schema supported by them.

Each version of the Redfish protocol is strongly typed. This is accomplished using the URI of the Redfish service in combination with the resource obtained at that URI, called the ServiceRoot.

The root URI for this version of the Redfish protocol shall be "/redfish/v1".

While the major version of the protocol is represented in the URI, the major version, minor version and errata version of the protocol are represented in the Version property of the ServiceRoot resource, as defined in the Schema for that resource. The protocol version is a string of the form:

MajorVersion.MinorVersion.Errata

where:

- *MajorVersion* = integer: something in the class changed in a backward incompatible way.
- *MinorVersion* = integer: a minor update. New functionality may have been added but nothing removed. Compatibility will be preserved with previous minorversions.

- *Errata* = integer: something in the prior version was broken and needed to be fixed.

Any resource discovered through links found by accessing the root service or any service or resource referenced using references from the root service shall conform to the same version of the protocol supported by the root service.

A GET on the resource "/redfish" shall return the following body:

```
{
  "v1": "/redfish/v1"
}
```

6.3. Requests

This section describes the requests that can be sent to Redfish services.

6.3.1. Request Headers

HTTP defines headers that can be used in request messages. The following table defines those headers and their requirements for Redfish services.

- Redfish services shall understand and be able to process the headers in the following table as defined by the HTTP 1.1 specification if the value in the Required column is set to "Yes".
- Redfish services shall understand and be able to process the headers in the following table as defined by the HTTP 1.1 specification if the value in the Required column is set to "Conditional" under the conditions noted in the description.
- Redfish services should understand and be able to process the headers in the following tables as defined by the HTTP 1.1 specification if the value in the Required column is set to "No".

Header	Required	Supported Values	Description
Accept	Yes	RFC 2616, Section 14.1	Indicates to the server what media type(s) this client is prepared to accept. <code>application/json</code> shall be supported for requesting resources and <code>application/xml</code> shall be supported for requesting metadata.
Accept-Encoding	Yes	RFC 2616, Section 14.4	Indicates if <code>gzip</code> encoding can be handled by the client
Accept-Language	No	RFC 2616, Section 14.4	This header is used to indicate the language(s) requested in the response. If this header is not specified, the appliance default locale will be used.
Content-Type	Conditional	RFC 2616, Section 14.17	Describes the type of representation used in the message body. <code>charset=utf-8</code> shall be supported for requests that have a body. Shall be required if there is a request body.
Content-Length	No	RFC 2616, Section 14.3	Describes the size of the message body. An optional means of indicating size of the body uses Transfer-Encoding: <code>chunked</code> , which does not use the Content-Length header. If a service does not support Transfer-Encoding and needs Content-Length instead, the service will respond with status code 411 .
Max-OData-Version	No	4.0	Indicates the maximum version of OData that an odata-aware client understands
OData-Version	Yes	4.0	If provided, services shall reject requests which specify an unsupported OData version.
Authorization	Conditional	RFC 2617, Section 2	Required for Basic Authorization
User-Agent	Yes	RFC 2616, Section 14.43	Required for tracing product tokens and their version. Multiple product tokens may be listed.
Host	Yes	RFC 2616, Section 14.23	Required to allow support of multiple origin hosts at a single IP address.

Header	Required	Supported Values	Description
Origin	Yes	W3C CORS, Section 5.7	Used to allow web applications to consume Redfish service while preventing CSRF attacks.
Via	No	RFC 2616, Section 14.45	Indicates network hierarchy and recognizes message loops. Each pass inserts its own VIA.
Max-Forwards	No	RFC 2616, Section 14.31	Limits gateway and proxy hops. Prevents messages from remaining in the network indefinitely.
If-Match	Conditional	RFC 2616, Section 14.31	If-Match shall be supported for Atomic requests on AccountService objects. If-Match shall be supported on requests for resources for which the service returns ETags.
If-None-Match	No	RFC 2616, Section 14.31	If this HTTP header is present, the service will only return the requested resource if the current ETag of that resource does not match the ETag sent in this header. If the ETag specified in this header matches the resource's current ETag, the status code returned from the GET will be 304.

- Redfish services shall understand and be able to process the headers in the following table as defined by this specification if the value in the Required column is set to "yes" .

Header	Required	Supported Values	Description
X-Auth-Token	Yes	Opaque encoded octet strings	Used for bearer authentication of user sessions. The token value shall be indistinguishable from random.

6.3.2. Read Requests (GET)

The GET method is used to retrieve a representation of a resource. That representation can either be a single resource or a collection. The service will return the representation using one of the media types specified in the Accept header, subject to requirements in the Media Types section [Media Types](#). If the Accept header is not present, the service will return the resources representations as application/json.

- The HTTP GET method shall be used to retrieve a resource without causing any side effects.
- The service shall ignore the content of the body on a GET.
- The GET operation shall be idempotent in the absence of outside changes to the resource.

6.3.2.1. Service Root Request

The root URL for Redfish version 1 services shall be "/redfish/v1".

Additionally, the latest supported Redfish service shall be aliased at "/redfish". In this case the endpoint at "/redfish" may be an HTTP redirect to "/redfish/v1".

The root URL for the service returns a RootService resource as defined by this specification.

6.3.2.2. Metadata Document Request

Redfish services shall expose a [metadata document](#) describing the service at the "/redfish/v1/\$metadata" resource. This metadata document describes the resources and collections available at the root, and references additional metadata documents describing the full set of resource types exposed by the service.

6.3.2.3. OData Service Document Request

Redfish services shall expose an [OData Service Document](#), at the "/redfish/v1/odata" resource. This service document provides a standard format for enumerating the resources exposed by the service, enabling generic hypermedia-driven OData clients to navigate to the resources of the service.

6.3.2.4. Resource Retrieval Requests

6.3.2.4.1. Query Parameters

When the resource addressed is a collection, the client can use the following paging query options to specify that a subset of the members be returned.

Attribute	Description	Example
\$skip	Integer indicating the number of resources in the collection to skip before retrieving the first resource.	http://collection?\$skip=5
\$top	Integer indicating the number of collection members to include in the response. The minimum value for this parameter is 1. The default behavior is to return all members.	http://collection?\$top=30

- Services should support the \$top and \$skip query parameters.
- Implementation shall return the 501, Not Implemented, status code for any query parameters starting with "\$" that are not supported, and should return an [extended error](#) indicating the requested query parameter(s) not supported for this resource.
- Implementations shall ignore unknown or unsupported query parameters that do not begin with "\$".

6.3.2.4.2. Retrieving Collections

Retrieving a collection is done by sending the HTTP GET method to the URI for the collection. The response will be a [resource collection representation](#) that includes the collection's attributes as well as the list of the members of the collection. A subset of the members can be returned using [client paging query parameters](#).

No requirements are placed on implementations to return a consistent set of members when a series of requests using paging query parameters are made over time to obtain the entire set of members. It is possible that this could result in missed or duplicate elements being retrieved if multiple GETs are used to retrieve a collection using paging.

- Clients shall not make assumptions about the URIs for the resource members of a collection.
- Retrieved collections shall always include the [count](#) property to specify the total number of members in the collection.
- If only a portion of the collection is returned due to client-specified paging query parameters or services returning [partial results](#), then the total number of resources across all pages shall be returned in the count property.

6.3.3. HEAD

The HEAD method differs from the GET method in that it **MUST NOT** return message body information. However, all of the same meta information and status codes in the HTTP headers will be returned as though a GET method were processed, including authorization checks.

- Services may support the HEAD method in order to return meta information in the the form of HTTP response headers.
- Services may support the HEAD method in order to verify link validity.
- Services may support the HEAD method in order to verify resource accessibility
- Services shall not support any other use of the HEAD method.
- The HEAD method shall be idempotent in the absence of outside changes to the resource.

6.3.4. Data Modification Requests

6.3.4.1. Update (PATCH)

The PATCH method is the preferred method used to perform updates on pre-existing resources. Changes to the resource are sent in the request body. Properties not specified in the request body are not directly changed by the PATCH request. The response is either empty or a representation of the resource after the update was done. The implementation may reject the update operation on certain fields based on its own policies and, if so, shall not apply any of the update requested. Updates to resources are idempotent.

- Services shall support the PATCH method to update a resource. If the resource can never be updated, status code [405](#) shall be returned.
- Services may return a representation of the resource after any server-side transformations in the body of the response.
- If a property in the request can never be updated, such as when a property is read only, a status code of [200](#) shall be returned along with a representation of the resource containing an [annotation](#) specifying the non-updatable property. In this success case, other properties may be updated in the resource.
- Services should return status code [405](#) if the client specifies a PATCH request against a collection.
- The PATCH operation should be idempotent in the absence of outside changes to the resource provided it is used with ETags to prevent subsequent PATCH attempts. Note that the ETAG value should change as the result of this operation.

OData markup ([resource identifiers](#), [type](#), [etag](#) and [links](#)) are ignored on Update.

6.3.4.2. Replace (PUT)

The PUT method is used to completely replace a resource. Properties omitted from the request body are reset to their default value.

- Services may support the PUT method to replace a resource in whole. If a service does not implement this method, status code [405](#) shall be returned.
- Services may return a representation of the resource after any server-side transformations in the body of the response.

- Services should return status code [405](#) if the client specifies a PUT request against a collection.
- The PUT operation should be idempotent in the absence of outside changes to the resource, with the possible exception that ETAG values may change as the result of this operation.

6.3.4.3. Create (POST)

The POST method is used to create new resources. The POST request is submitted to the resource collection in which the new resource is to belong.

- Services shall support the POST method for creating resources. If the resource does not offer anything to be created, a status code [405](#) shall be returned.
- The POST operation shall not be idempotent.

The body of the create request contains a representation of the object to be created. The service can ignore any service controlled attributes (e.g. id), forcing those attributes to be overridden by the service. The service shall set the Location header to the URI of the newly created resource. The response to a successful create request should be 201 (Created) and may include a response body containing the representation of the newly created resource.

6.3.4.4. Delete (DELETE)

The DELETE method is used to remove a resource.

- Services shall support the DELETE method for resources that can be deleted. If the resource can never be deleted, status code [405](#) shall be returned.
- Services may return a representation of the just deleted resource in the response body.
- Services should return status code [405](#) if the client specifies a DELETE request against a collection.
- The DELETE operation shall not be not idempotent.

6.3.4.5. Actions (POST)

The POST method is used to initiate operations on the object (such as Actions).

- Services shall support the POST method for sending actions.
- The POST operation may not be idempotent.

Custom actions are requested on a resource by sending the HTTP POST method to the URI of the action. If the [actions property](#) within a resource does not specify a target property, then the URI of an action shall be of the form:

ResourceUri/Actions/QualifiedActionName

where

- *ResourceUri* is the URL of the resource which supports invoking the action.
- "Actions" is the name of the property containing the actions for a resource, as defined by this specification.
- *QualifiedActionName* is the name of the action qualified by the namespace alias.

The first parameter of a bound function is the resource on which the action is being invoked. The remaining parameters are represented as name/value pairs in the body of the request.

Clients can query a resource directly to determine the [actions](#) that are available as well as [valid parameter values](#) for those actions. Some parameter information may require the client to examine the schema corresponding to the resource.

For instance, if a schema document <http://schemas.dmtf.org/redfish/v1/ComputerSystem.0.96.0> defines a Reset action, in the `ComputerSystem.0.96.0` namespace, with the alias "ComputerSystem", bound to the `ComputerSystem.Actions` type, such as this example:

```
<Schema Name="ComputerSystem.0.96.0" Alias="ComputerSystem">
...
  <Action Name="Reset" Isbound="true">
    <Parameter Name="Resource" Type="ComputerSystem.Actions" />
    <Parameter Name="ResetType" Type="ComputerSystem.ResetType" />
  </Action>
...
</Schema>
```

And a computer system resource contains an [Actions](#) property such as this:

```

"Actions": {
  "#ComputerSystem.Reset": {
    "target": "/redfish/v1/Systems/1/Actions/ComputerSystem.Reset",
    "ResetType@DMTF.AllowableValues": [
      "On",
      "ForceOff",
      "GracefulRestart",
      "ForceRestart",
      "Nmi",
      "ForceOn",
      "PushPowerButton"
    ]
  }
}

```

Then the following would represent a possible request for the Action:

```

POST /redfish/v1/Systems/1/Actions/ComputerSystem.Reset
{
  "ResetType": "On"
}

```

6.4. Responses

Redfish defines four types of responses:

- [Metadata Responses](#) - Describe the resources and types exposed by the service to generic clients.
- [Resource Responses](#) - JSON representation of an individual resource.
- [Resource Collection Responses](#) - JSON representation of a collections of resources.
- [Error Responses](#) - Top level JSON response providing additional information in the case of an HTTP error.

6.4.1. Response Headers

HTTP defines headers that can be used in response messages. The following table defines those headers and their requirements for Redfish services.

- Redfish services shall be able to return the headers in the following table as defined by the HTTP 1.1 specification if the value in the Required column is set to "yes" .
- Redfish services should be able to return the headers in the following tables as defined by the HTTP 1.1 specification if the value in the Required column is set to "no".
- Redfish clients shall be able to understand and be able to process all of the headers in the following table as defined by the HTTP 1.1. specification.

Header	Required	Supported Values	Description
OData-Version	Yes	4.0	Describes the OData version of the payload that the response conforms to.
Content-Type	Yes	RFC 2616, Section 14.17	Describes the type of representation used in the message body. <code>application/json</code> shall be supported. <code>charset=utf-8</code> shall be supported.
Content-Encoding	No	RFC 2616, Section 14.17	Describes the encoding that has been performed on the media type
Content-Length	No	RFC 2616, Section 14.3	Describes the size of the message body. An optional means of indicating size of the body uses Transfer-Encoding: chunked, which does not use the Content-Length header. If a service does not support Transfer-Encoding and needs Content-Length instead, the service will respond with status code 411.
ETag	Conditional	RFC 2616, Section 14.19	An identifier for a specific version of a resource, often a message digest. Etags shall be included on Account objects.

Header	Required	Supported Values	Description
Server	Yes	RFC 2616, Section 14.38	Required to describe a product token and its version. Multiple product tokens may be listed.
Location	Conditional	RFC 2616, Section 14.30	Indicates a URI that can be used to request a representation of the resource. Shall be returned if a new resource was created. Location and X-Auth-Token shall be included on responses which create user sessions.
Cache-Control	Yes	RFC 2616, Section 14.9	This header shall be supported and is meant to indicate whether a response can be cached or not.
Via	No	RFC 2616, Section 14.45	Indicates network hierarchy and recognizes message loops. Each pass inserts its own VIA.
Max-Forwards	No	RFC 2616, Section 14.31	Limits gateway and proxy hops. Prevents messages from remaining in the network indefinitely.
Link	No	RFC 5988, Section 5	Exposes additional metadata about response object. Shall only be returned in response to a HEAD operation.
Access-Control-Allow-Origin	Yes	W3C CORS, Section 5.1	Prevents or allows requests based on originating domain. Used to prevent CSRF attacks.
Allow	Yes	POST, PUT, PATCH, DELETE	Returned on GET or HEAD operation to indicate the other allowable operations for this resource. Shall be returned with a 405 (Method Not Allowed) response to indicate the valid methods for the specified Request URI.
WWW-Authenticate	Yes	RFC 2617	Required for Basic and other optional authentication mechanisms. See the [Security][#Security] section for details.

- Redfish services shall understand and be able to process the headers in the following table as defined by this specification if the value in the Required column is set to "yes".

Header	Required	Supported Values	Description
X-Auth-Token	Yes	Opaque encoded octet strings	Used for bearer authentication of user sessions. The token value shall be indistinguishable from random.

6.4.2. Status Codes

HTTP defines status codes that can be returned in response messages.

Where the HTTP status code indicates a failure, the response body contains an [extended error resource](#) to provide the client more meaningful and deterministic error semantics.

- Services shall return the extended error resource as described in this specification in the response body when a status code of 400 or 500 is returned.
- Services should return the extended error resource as described in this specification in the response body when a status code 400 or greater is returned.
- Extended error messages MUST NOT provide privileged info when authentication failures occur

NOTE: Refer to the [Security](#) section for security implications of extended errors

The following table lists some of the common HTTP status codes. Other codes may be returned by the service as appropriate. See the Description column for a description of the status code and additional requirements imposed by this specification.

- Clients shall understand and be able to process the status codes in the following table as defined by the HTTP 1.1 specification and constrained by additional requirements defined by this specification.
- Services shall respond with these status codes as appropriate.
- Exceptions from operations shall be mapped to HTTP status codes.
- Redfish services should not return the status code 100. Using the HTTP protocol for a multi-pass data transfer should be avoided, except upload of extremely large data.

HTTP Status Code	Description
200 OK	The request was successfully completed and includes a representation in its body.
201 Created	A request that created a new resource completed successfully. The Location header shall be set to the canonical URI for the newly created resource. A representation of the newly created resource may be included in the response body.
202 Accepted	The request has been accepted for processing, but the processing has not been completed. The Location header shall be set to the URI of a Task resource that can later be queried to determine the status of the operation. A representation of the Task resource may be included in the response body.
204 No Content	The request succeeded, but no content is being returned in the body of the response.
301 Moved Permanently	The requested resource resides under a different URI
302 Found	The requested resource resides temporarily under a different URI.
304 Not Modified	The service has performed a conditional GET request where access is allowed, but the resource content has not changed. Conditional requests are initiated using the headers If-Modified-Since and/or If-None-Match (see HTTP 1.1, sections 14.25 and 14.26) to save network bandwidth if there is no change.
400 Bad Request	The request could not be processed because it contains missing or invalid information (such as validation error on an input field, a missing required value, and so on). An extended error shall be returned in the response body, as defined in section Extended Error Handling .
401 Unauthorized	The authentication credentials included with this request are missing or invalid.
403 Forbidden	The server recognized the credentials in the request, but those credentials do not possess authorization to perform this request.
404 Not Found	The request specified a URI of a resource that does not exist.
405 Method Not Allowed	The HTTP verb specified in the request (e.g. DELETE, GET, HEAD, POST, PUT, PATCH) is not supported for this request URI. The response shall include an Allow header which provides a list of methods that are supported by the resource identified by the Request-URI.
406 Not Acceptable	The Accept header was specified in the request and the resource identified by this request is not capable of generating a representation corresponding to one of the media types in the Accept header.
409 Conflict	A creation or update request could not be completed, because it would cause a conflict in the current state of the resources supported by the platform (for example, an attempt to set multiple attributes that work in a linked manner using incompatible values).
410 Gone	The requested resource is no longer available at the server and no forwarding address is known. This condition is expected to be considered permanent. Clients with link editing capabilities SHOULD delete references to the Request-URI after user approval. If the server does not know, or has no facility to determine, whether or not the condition is permanent, the status code 404 (Not Found) SHOULD be used instead. This response is cacheable unless indicated otherwise.
411 Length Required	The request did not specify the length of its content using the Content-Length header (perhaps Transfer-Encoding: chunked was used instead). The addressed resource requires the Content-Length header.
412 Precondition Failed	Precondition (If Match or If Not Modified) check failed.
415 Unsupported Media Type	The request specifies a Content-Type for the body that is not supported.
500 Internal Server Error	The server encountered an unexpected condition that prevented it from fulfilling the request. An extended error shall be returned in the response body, as defined in section Extended Error Handling .
501 Not Implemented	The server does not (currently) support the functionality required to fulfill the request. This is the appropriate response when the server does not recognize the request method and is not capable of supporting the method for any resource.
503 Service Unavailable	The server is currently unable to handle the request due to temporary overloading or maintenance of the server.

6.4.3. Metadata Responses

Metadata describes resources, collections, capabilities and service-dependent behavior to generic consumers, including OData client tools and applications with no specific understanding of this specification. Clients are not required to request metadata if they already have sufficient understanding of the target service; for example, to request and interpret a JSON representation of a resource defined in this specification.

6.4.3.1. Service Metadata

The service metadata describes top-level resources and resource types of the service according to [OData Schema](#). The Redfish Service Metadata is represented as an XML document with a root element named "Edmx", defined in the <http://docs.oasis-open.org/odata/ns/edmx> namespace, and with an OData Version attribute equal to "4.0".

```
<edmx:Edmx xmlns:edmx="http://docs.oasis-open.org/odata/ns/edmx" Version="4.0">
  <!-- edmx:Reference and edmx:Schema elements go here -->
</edmx:Edmx>
```

The service metadata shall include the namespaces for each of the DMTF resource types, along with the "DMTFExtensions.0.96.0" namespace. These references may use the standard Uri for the hosted schema definitions (i.e., on <http://dmtf.org/schema>) or a Url to a local version of the schema that shall be identical to the hosted version. The namespace shall be aliased with the version-independent namespace name.

```
<edmx:Reference Uri="http://schemas.dmtf.org/redfish/v1/ServiceRoot.0.96.0">
  <edmx:Include Namespace="ServiceRoot.0.96.0" Alias="ServiceRoot"/>
</edmx:Reference>
<edmx:Reference Uri="http://schemas.dmtf.org/redfish/v1/AccountService.0.96.0">
  <edmx:Include Namespace="AccountService.0.96.0" Alias="AccountService"/>
</edmx:Reference>
...
<edmx:Reference Uri="http://schemas.dmtf.org/redfish/v1/VirtualMedia.0.96.0">
  <edmx:Include Namespace="VirtualMedia.0.96.0" Alias="VirtualMedia"/>
</edmx:Reference>
<edmx:Reference Uri="http://schemas.dmtf.org/redfish/v1/DMTFExtensions.0.96.0">
  <edmx:Include Namespace="DMTFExtensions.0.96.0" Alias="DMTF"/>
</edmx:Reference>
```

The service metadata shall include an entity container that defines the top level resource and collections. This entity container shall extend the ServiceContainer defined in the ServiceRoot.0.96.0 schema and may include additional resources or collections.

```
<edmx:DataServices>
  <Schema xmlns="http://docs.oasis-open.org/odata/ns/edm" Namespace="Service">
    <EntityContainer Name="Service" Extends="ServiceRoot.ServiceContainer"/>
  </Schema>
</edmx:DataServices>
```

6.4.3.1.1. Referencing OEM Extensions

The metadata document may reference additional schema documents describing OEM-specific extensions used by the ServiceRoot, for example custom types for additional collections.

```
<edmx:Reference Uri="http://contoso.org/Schema/CustomTypes">
  <edmx:Include Namespace="CustomTypes" Alias="Contoso"/>
</edmx:Reference>
```

6.4.3.1.2. Annotations

The service can annotate sets, types, actions and parameters with Redfish-defined or custom annotation terms. These annotations are typically in a separate Annotations file referenced from the service metadata document using the IncludeAnnotations directive. The alias of the namespace containing system annotations shall be "Annotations".

```
<edmx:Reference Uri="http://service/metadata/Service.Annotations">
  <edmx:IncludeAnnotations TermNamespace="Annotations.0.96.0" Alias="Annotations"/>
</edmx:Reference>
```

The annotation file itself specifies the Target schema element being annotated, the Term being applied, and the value of the term:

```
<Annotations Target="ComputerSystem.Reset/ResetType">
  <Annotation Term="Annotations.AdditionalValues">
    <Collection>
      <String>Update and Restart</String>
      <String>Update and PowerOff</String>
    </Collection>
  </Annotation>
</Annotations>
```

6.4.3.2. OData Service Document

The OData Service Document serves as a top-level entry point for generic OData clients.

```
{
  "@odata.context": "/redfish/v1/$metadata",
  "value": [
    {
      "name": "Service",
      "kind": "Singleton",
      "url": "/redfish/v1"
    },
    {
      "name": "Systems",
      "kind": "Singleton",
      "url": "/redfish/v1/Systems"
    },
    {
      "name": "Chassis",
      "kind": "Singleton",
      "url": "/redfish/v1/Chassis"
    },
    {
      "name": "Managers",
      "kind": "Singleton",
      "url": "/redfish/v1/Managers"
    },
    ...
  ]
}
```

The OData Service Document shall be returned as a JSON object, using the MIME type `application/json`.

The JSON object shall contain a context property named `"@odata.context"` with a value of `"/redfish/v1/$metadata"`. This context tells a generic OData client how to find the [service metadata](#) describing the types exposed by the service.

The JSON object shall include a property named `"value"` whose value is a JSON array containing an entry for the [service root](#) and each resource that is a direct child of the service root.

Each entry shall be represented as a JSON object and shall include a `"name"` property whose value is a user-friendly name of the resource, a `"kind"` property whose value is `"Singleton"` for individual resources (including collection resources) or `"EntitySet"` for top-level resource collections, and a `"url"` property whose value is the relative URL for the top-level resource.

6.4.4. Resource Responses

Resources are returned as JSON payloads, using the MIME type `application/json`.

6.4.4.1. Context Property

Responses that represent a single resource shall contain a context property named `"@odata.context"` describing the source of the payload. The value of the context property shall be the context URL that describes the resource according to [OData-Protocol](#).

The context URL for a resource that exists within a collection is of the form:

MetadataUrl.#Collection[(Selectlist)]/\$entity

Where:

- *MetadataUrl* = the metadata url of the service (`/redfish/v1/$metadata`)
- *Collection* = the collection resource. For contained resources this includes the path from the root collection or singleton resource to the containment property.
- *Selectlist* = comma-separated [list of properties](#) included in the response if the response includes a subset of properties defined for the represented resources.

The context URL for a resource that is a top-level singleton resource is of the form:

MetadataUrl.#SingletonName[(Selectlist)]

Where:

- *MetadataUrl* = the metadata url of the service (`/redfish/v1/$metadata`)
- *SingletonName* = the name of the top-level singleton resource
- *Selectlist* = comma-separated [list of properties](#) included in the response if the response includes a subset of properties defined for the represented resources.

6.4.4.1.1. Select List

If a response contains a subset of the properties defined in the schema for a type, then the context URL shall specify the subset of properties included. An asterisk (*) can be used to specify "all structural properties" for a given resource.

Expanded navigation properties shall be included in the select list if the result includes a subset of the properties defined for the expanded resource.

For example, the following context URL specifies that the result contains a single resource from the Members collection nested under the Links property of the Systems resource:

```
"@odata.context": "/redfish/v1/$metadata#Systems/Links/Members/$entity",
```

6.4.4.2. Resource Identifier Property

Resources in a response shall include a unique identifier property named `"@odata.id"`. The value of the identifier property shall be the [unique identifier](#) for the resource.

Resource Identifiers shall be represented in JSON payloads as uri paths relative to the schema portion of the uri. That is, they shall always start with `"/redfish/"`.

The resource identifier is the canonical URL for the resource and can be used to retrieve or edit the resource, as appropriate.

6.4.4.3. Type Property

All resources in a response shall include a type property named `"@odata.type"`. The value of the type property shall be an absolute URL that specifies the type of the resource and shall be of the form:

#Namespace.TypeName

Where:

- *Namespace* = The full namespace name of the schema in which the type is defined. For Redfish resources this will be the versioned namespace name.
- *TypeName* = The name of the type of the resource.

The client may issue a GET request to this URL using a content type of `application/xml` in order to retrieve a document containing the [definition of the resource](#).

6.4.4.4. ETag Property

ETags provide the ability to conditionally retrieve or update a resource. Resources should include an ETag property named "@odata.etag". The value of the ETag property is the [Etag](#) for a resource.

6.4.4.5. Primitive Properties

Primitive properties are returned as JSON values.

6.4.4.5.1. DateTime Values

DateTime values shall be returned as JSON strings according to the ISO 8601 "extended" format:

```
YYYY-MM-DD T hh:mm:ss[.SSS] (Z | (+ | -) hh:mm)
```

6.4.4.6. Structured Properties

Structured properties, defined as [complex types](#) or [expanded resource types](#), are returned as JSON objects. The type of the JSON object is specified in schema definition of the property containing the structured value.

6.4.4.7. Collection Properties

Collection-valued properties are returned as JSON arrays, where each element of the array is a JSON object whose type is specified in the schema document describing the containing type.

Collection-valued properties may contain a subset of the members of the full collection. In this case, the collection-valued property shall be annotated with a next link property. The property representing the next link shall be a peer of the collection-valued property, with the name of the collection-valued property suffixed with "@odata.nextLink". The value of the next link property shall be an opaque URL that the client can use to retrieve the next set of collection members. The next link property shall only be present if the number of resources requested is greater than the number of resources returned.

Collection-valued properties shall be annotated with a count. The property representing the count is a peer of the collection-valued property, with the name of the collection-valued property suffixed with "@odata.count". The value of the count is the total number of members available in the collection.

Collection-valued properties shall not be null. Empty collections shall be returned in JSON as an empty array.

6.4.4.8. Actions Property

Available actions for a resource are represented as individual properties nested under a single structured property on the resource named "Actions".

6.4.4.8.1. Action Representation

Actions are represented by a property nested under "Actions" whose name is the unique URI that identifies the action. This URI shall be of the form:

```
*#NamespaceAlias.ActionName*
```

Where:

- *NamespaceAlias* = The namespace alias used in the reference to the schema in which the action is defined. For Redfish resources this shall be the version-independent namespace alias.
- *ActionName* = The name of the action

The client may issue a GET request to this URL using a content type of `application/xml` in order to retrieve the schema document containing the [definition of the action](#).

The value of the property is a JSON object containing a property named "target" whose value is a relative or absolute URL used to invoke the action.

The property representing the available action may be annotated with the [AllowableValues](#) annotation in order to specify the list of allowable values for a particular parameter.

For example, the following property represents the Reset action, defined in the ComputerSystem.0.96.0 namespace (aliased with the version-independent "ComputerSystem"):

```
"#ComputerSystem.Reset": {
  "target": "/redfish/v1/Systems/1/Actions/ComputerSystem.Reset",
  "ResetType@DMTF.AllowableValues": [ "On", "Off", "Reset" ]
}
```

Given this, the client could invoke a POST request to `/redfish/v1/Systems/1/Actions/ComputerSystem.Reset` with the following body:

```
{
  "ResetType": "On"
}
```

6.4.4.8.2. Allowable Values

The property representing the action may be annotated with the "AllowableValues" annotation in order to specify the list of allowable values for a particular parameter.

The set of allowable values is specified by including a property whose name is the name of the parameter followed by "@DMTF.AllowableValues", and whose value is a JSON array of strings representing the allowable values for the parameter.

6.4.4.9. Links Property

[References](#) to other resources are represented by the links property on the resource.

The links property shall be named "Links" and shall contain a property for each navigation property defined in the schema for that type. For single-valued navigation properties, the value of the property shall be the [single related resource id](#). For collection-valued navigation properties, the value of the property shall be the [array of related resource ids](#).

The links property shall also include an [Oem property](#) for navigating vendor-specific links.

6.4.4.9.1. Reference to a Single Related Resource

A reference to a single resource is returned as a JSON object containing a single [resource-identifier-property](#) whose name is the name of the relationship and whose value is the uri of the referenced resource.

```
{
  "Links" : {
    "ManagedBy": {
      "@odata.id": "/redfish/v1/Chassis/Enc11"
    }
  }
}
```

6.4.4.9.2. Array of References to Related Resources

A reference to a collection of zero or more related resources is returned as an array of JSON objects whose name is the name of the relationship. Each member of the array is a JSON object containing a single [resource-identifier-property](#) whose value is the uri of the referenced resource.

```
{
  "Links" : {
    "Contains" : [
      {
        "@odata.id": "/redfish/v1/Chassis/1"
      },
      {
        "@odata.id": "/redfish/v1/Chassis/Enc11"
      }
    ]
  }
}
```

6.4.4.10. OEM Property

OEM-specific properties are nested under an [OEM property](#). The name of the OEM property shall be "Oem" and its value shall be a JSON object whose properties represent OEM extensions.

For example:

```

{
  "StandardProperty": "value",
  "Oem": {
    "Acme": {
      "@odata.type": "http://acme.com/schema/extensions.v.v.v#acme.acmetype",
      "AcmeSpecificProperty": "value"
    },
    "Acme:Type2": {
      "@odata.type": "http://acme.com/schema/extensions.v.v.v#acme.acmetype2",
      "AcmeSpecificProperty2": "value"
    },
    "EID:232": {
      "@odata.type": "http://eid.org/schemas/eid232.v.v.v",
      "EnterpriseSpecificProperty": "value"
    }
  }
}

```

Contents of the Oem object must be valid JSON and must have a [type property](#) per this specification. Any other requirements, validation or contents are beyond the scope of this specification.

6.4.4.11. Extended Information

Response objects may include extended information, for example properties that are not able to be updated. This information is represented as an annotation applied to a [specific property](#) of the JSON response or an [entire JSON object](#).

6.4.4.11.1. Extended Object Information

A JSON object can be annotated with "@DMTF.ExtendedInfo" in order to specify object-level status information.

```

{
  "@odata.context": "/redfish/v1/$metadata#SessionService/Links/Sessions/Links/Members/$entity",
  "@odata.id": "/redfish/v1/SessionService/Sessions/Administrator1",
  "@odata.type": "#Session.0.94.0.Session",
  "Id": "Administrator1",
  "Name": "User Session",
  "Description": "Manager User Session",
  "Modified": "2013-01-31T23:45:08+00:00",
  "UserName": "Administrator",
  "Oem": {},
  "@DMTF.ExtendedInfo" : {
    "code": "Base.0.96.ResourceCannotBeDeleted",
    "message": "The delete request failed because the resource requested cannot be deleted",
    "@message.severity": "Critical",
    "@message.resolution": "Do not attempt to delete a non-deletable resource."
  }
}

```

The value of the property is an [extended information object](#).

6.4.4.11.2. Extended Property Information

An individual property within a JSON object can be annotated with extended information using "@DMTF.ExtendedInfo", prepended with the name of the property.

```

{
  "@odata.context": "/redfish/v1/$metadata/Sessions/Links/Members/$entity",
  "@odata.id": "/redfish/v1/Sessions/Administrator1",
  "@odata.type": "#Session.0.96.0.Session",
  "Id": "Administrator1",

```

```

    "Name": "User Session",
    "Description": "Manager User Session",
    "Modified": "2013-01-31T23:45:08+00:00",
    "UserName": "Administrator",
    "UserName@DMTF.ExtendedInfo" : {
        "code": "Base.0.96.PropertyNotWriteable",
        "target": "UserName",
        "message": "The property %1 is a read only property and cannot be assigned a value",
        "@message.parameters": [
            {
                "UserName"
            }
        ],
        "@message.severity": "Warning",
        "@message.resolution": "Remove the property from the request body and resubmit the request if the operation failed"
    },
    "Oem": {}
}

```

The value of the property is an [extended information object](#).

6.4.4.12. Additional Annotations

A resource representation in JSON may include additional annotations represented as properties whose name is of the form:

[PropertyName]@Namespace.TermName

where

- *PropertyName* = the name of the property being annotated. If omitted, the annotation applies to the entire resource.
- *Namespace* = the name or alias of the namespace where the annotation term is defined. This namespace must be referenced by the [metadata document](#) specified in the [context url](#) of the request.
- *TermName* = the name of the annotation term being applied to the resource or property of the resource.

The client can get the definition of the annotation from the [service metadata](#), or may ignore the annotation entirely, but should not fail reading the resource due to unrecognized annotations, including new annotations defined within the DMTF namespace.

6.4.5. Resource Collections

Resource collections are returned as a JSON object. The JSON object shall include a [context](#), [resource count](#), and array of [values](#), and may include a [next link](#) for partial results.

6.4.5.1. Context Property

Responses that represent a collection of resources shall contain a context property named "@odata.context" describing the source of the payload. The value of the context property shall be the context URL that describes the resources according to [OData-Protocol](#).

The context URL for a resource collection is of the form:

MetadataUrl.#Collection[(SelectList)]

Where:

- *MetadataUrl* = the metadata url of the service (*/redfish/v1/\$metadata*)
- *Collection* = the collection resource. For contained resources this includes the path from the root collection or singleton resource to the containment property.
- *SelectList* = comma-separated [list of properties](#) included in the response if the response includes a subset of properties defined for the represented resources.

6.4.5.2. Resource Count Property

The total number of resources available in the collection is represented through the count property. The count property shall be named "@odata.count" and its value shall be an integer representing the total number of records in the result. This count is not affected by the \$top or \$skip [query parameters](#).

6.4.5.3. Resource Members Property

The members of the collection of resources are returned as a JSON array. The name of the property representing the members of the collection shall be "value".

6.4.5.4. Partial Results

Responses representing a single resource shall not be broken into multiple results.

Collections of resources, or resource ids, may be returned in multiple partial responses. For partial collections the service includes a next link property named "@odata.nextLink". The value of the next link property shall be an opaque URL that the client can use to retrieve the next set of resources. The next link shall only be returned if the number of resources requested is greater than the number of resources returned.

The value of the [count property](#) represents the total number of resources available if the client enumerates all pages of the collection.

6.4.5.5. Additional Annotations

A JSON object representing a collection of resources may include additional annotations represented as properties whose name is of the form:

@Namespace.TermName

where

- *Namespace* = the name or alias of the namespace where the annotation term is defined. This namespace shall be referenced by the [metadata document](#) specified in the [context url](#) of the request.
- *TermName* = the name of the annotation term being applied to the resource collection.

The client can get the definition of the annotation from the [service metadata](#), or may ignore the annotation entirely, but should not fail reading the response due to unrecognized annotations, including new annotations defined within the DMTF namespace.

6.4.6. Error Responses

HTTP response status codes alone often do not provide enough information to enable deterministic error semantics. For example, if a client does a PATCH and some of the properties do not match while others are not supported, simply returning an HTTP status code of 400 does not tell the client which values were in error. Error responses provide the client more meaningful and deterministic error semantics.

Error responses are defined by an extended error resource, represented as a single JSON object with a property named "error". The value of this property shall be an [extended information object](#).

6.4.6.1. Extended Information Object

Extended Information Objects provide additional information about an [object](#), [property](#), or [error response](#).

Extended information is represented as a JSON object with the following properties:

Property	Description
code	String indicating a specific error or message (not to be confused with the HTTP status code). This code can be used to access a detailed message from a message registry.
message	A human readable error message indicating the semantics associated with the error.
target	An optional string defining the target of the particular error.
details	An optional array of JSON objects with code, message, target, severity, and resolution properties, providing more detailed information about the error.
message.parameters	An optional array of strings representing the substitution parameter values for the message. The Severity attribute is an annotation specified in the DMTF namespace and shall be prefixed with the alias "message".
message.severity	An optional string representing the severity of the error. The Severity attribute is an annotation specified in the DMTF namespace and shall be prefixed with the alias "message".
message.resolution	An optional string describing recommended action(s) to take to resolve the error. The Resolution attribute is an annotation specified in the DMTF namespace and shall be prefixed with the alias "message".

```

{
  "error": {
    "code": "400",
    "message": "The update operation failed.",
    "details": [
      {
        "code": "Base.0.96.PropertyValueNotInList",
        "target": "IndicatorLED",
        "message": "The value %1 for the property %2 is not in the list of acceptable values",
        "@message.parameters": [
          "RED",
          "IndicatorLED"
        ],
        "@message.severity": "Warning",
        "@message.resolution": "Remove the property from the request body and resubmit the request if the operation failed"
      },
      {
        "code": "Base.0.96.PropertyNotWriteable",
        "target": "SKU",
        "message": "The property %1 is a read only property and cannot be assigned a value",
        "@message.parameters": [
          "SKU"
        ],
        "@message.severity": "Warning",
        "@message.resolution": "Remove the property from the request body and resubmit the request if the operation failed"
      }
    ]
  }
}

```

Error codes identify specific errors defined in a message registry.

The value of the code property shall be of the form

RegistryName.MajorVersion.MinorVersion.MessageKey

where

- *RegistryName* is the name of the registry. The registry name shall be Pascal-cased.
- *MajorVersion* is a positive integer representing the major version of the registry
- *MinorVersion* is a positive integer representing the minor version of the registry
- *MessageKey* is a human-readable key into the registry. The message key shall be Pascal-cased and shall not include spaces, periods or special chars.

The client can use the error code to search the message registry for the corresponding message.

The message registry approach has advantages for internationalization (since the registry can be translated easily) and light weight implementation (since large strings need not be included with the implementation).

7. Data Model & Schema

One of the key tenants of the Redfish interface is the separation of protocol and data model. This section describes common data model, resource, and Redfish Schema requirements.

- Each resource shall be strongly typed according to a [resource type definition](#). The type shall be defined in a [schema document](#) and identified by a unique [type identifier](#).

7.1. Type Identifiers

Types are identified by a *Type URI*. The full URI for a type is of the form:

#Namespace.TypeName

where:

- *Namespace* = the full name or alias of the namespace in which the type is defined
- *TypeName* = the name of the type

The full namespace for types defined by this specification is of the form:

ResourceTypeName.MajorVersion.MinorVersion.Errata

where

- *ResourceTypeName* = the name of the resource type. For [structured \(complex\) types](#), [enumerations](#), and [actions](#), this is generally the name of the containing resource type.
- *MajorVersion* = integer: something in the class changed in a backward incompatible way.
- *MinorVersion* = integer: a minor update. New properties may have been added but nothing removed. Compatibility will be preserved with previous minorversions.
- *Errata* = integer: something in the prior version was broken and needed to be fixed.

An example of a valid type namespace might be "System.0.96.0".

7.1.1. Type Identifiers in JSON

Types used within a JSON payload shall be defined in, or referenced, by the [service metadata](#).

Resource types defined by this specification shall be referenced in JSON documents using the full (versioned) namespace name.

Non-resource types (for example enumerations, complex types, and actions) shall be referenced in JSON documents using the version-independent namespace alias defined in the [service metadata](#).

NOTE: Refer to the [Security](#) section for security implications of Data Model & Schema

7.2. Common Naming Conventions

The Redfish interface is intended to be easily readable and intuitive. Thus, consistency helps the consumer who is unfamiliar with a newly discovered property understand its use. While this is no substitute for the normative information in the specification and schema, the following rules help with readability and client usage.

Resource Name, Property Names, and constants such as Enumerations shall be Pascal-cased

- The first letter of each word shall be upper case with spaces between words shall be removed (eg PowerState, SerialNumber.)
- No underscores are used.
- Both characters are capitalized for two-character acronyms (eg IPAddress, RemoteIP)
- Only the first character of acronyms with three or more characters is capitalized, except the first word of a Pascal-cased identifier (eg Wwn, VirtualWwn)

Exceptions are allowed for the following cases:

- Well-known technology names like "iSCSI"
- Product names like "iLO"
- Well-known abbreviations or acronyms

For attributes that have units, or other special meaning, the unit identifier should be appended to the name. The current list includes:

- Bandwidth (Mbps), (eg PortSpeedMbps)
- CPU speed (Mhz), (eg ProcessorSpeedMhz)
- Memory size (MegaBytes, MB), (eg MemoryMB)
- Counts of items (Count), (eg ProcessorCount, FanCount)
- The State of a resource (State) (eg PowerState.)
- State values where "work" is being done end in (ing) (eg Applying, Clearing)

7.3. Localization Considerations

The Redfish architecture supports localized strings but does not impose any specific requirement for localization upon services. However, current market forces seem to require localization as necessary (e.g. schemas).

Schema-supplied display strings may be localized as necessary, but a Schema file may only contain one language. Alternate language schemas may be published and available to Redfish clients, but need not be provided via the Redfish schema store.

Property names defined within a Redfish schema are never localized. User-supplied string-valued property values such as an asset tag may be localized. Localizable string valued properties should be annotated with the [IsLanguageDependent annotation term](#).

7.4. Schema Definition

Individual resources and their dependent types and actions are defined within a [schema document](#).

7.4.1. Common Annotations

All Redfish types and properties shall include [description](#) and [long description](#) annotations.

7.4.1.1. Description

The Description annotation can be applied to any type, property, action or parameter in order to provide a human-readable description of the schema element.

The `Description` annotation is defined in <http://docs.oasis-open.org/odata/odata/v4.0/os/vocabularies/Org.OData.Core.V1.xml>.

7.4.1.2. Long Description

The LongDescription annotation term can be applied to any type, property, action or parameter in order to provide a formal specification of the schema element.

The `LongDescription` annotation term is defined in <http://docs.oasis-open.org/odata/odata/v4.0/os/vocabularies/Org.OData.Core.V1.xml>.

7.4.2. Schema Documents

Individual resources are defined as entity types within a schema document according to [OData-Schema](#). The representation may include annotations to facilitate automatic generation of JSON Schema capable of validating JSON payloads.

The outer element of the schema document shall be the `Edmx` element, and shall have a `Version` attribute with a value of "4.0".

```
<edmx:Edmx xmlns:edmx="http://docs.oasis-open.org/odata/ns/edmx" Version="4.0">
  <!-- edmx:Reference and edmx:DataService elements go here -->
</edmx:Edmx>
```

7.4.2.1. Referencing other Schemas

Schemas may reference types defined in other schema documents by including a `Reference` element.

The reference element specifies the `Uri` of the schema document describing the referenced type and has one or more child `Include` elements that specify the `Namespace` attribute containing the types to be referenced, along with an optional `Alias` attribute for that namespace.

Type definitions generally reference the OData and DMTF namespaces for common type annotation terms, and resource type definitions reference the DMTF Resource.0.96.0 namespace for base types. Schemas that include measures such as temperature, speed, or dimensions generally include the [OData Measures namespace](#).


```

<edmx:Reference Uri="http://docs.oasis-open.org/odata/odata/v4.0/cs01/vocabularies/Org.OData.Core.V1.xml">
  <edmx:Include Namespace="Org.OData.Core.V1" Alias="OData"/>
</edmx:Reference>
<edmx:Reference
  Uri="http://docs.oasis-open.org/odata/odata/v4.0/os/vocabularies/Org.OData.Measures.V1.xml">
  <edmx:Include Namespace="Org.OData.Measures.V1" Alias="OData.Measures"/>
</edmx:Reference>
<edmx:Reference Uri="http://schemas.dmtf.org/redfish/v1/Extensions.0.96.0">
  <edmx:Include Namespace="DMTFExtensions.0.96.0" Alias="DMTF"/>
</edmx:Reference>
<edmx:Reference Uri="http://schemas.dmtf.org/redfish/v1/Resource.0.96.0">
  <edmx:Include Namespace="Resource.0.96.0" Alias="Resource"/>
</edmx:Reference>

```

7.4.2.2. Namespace Definitions

Resource types are defined within a namespace. The namespace is defined through a Schema element that contains attributes for declaring the Namespace and local Alias for the schema.

The Schema element is a child of the DataServices element, which is a child of the Edmx element.

```

<edmx:DataServices>
  <Schema xmlns="http://docs.oasis-open.org/odata/ns/edm" Namespace="MyTypes.0.96.0" Alias="MyTypes">

    <!-- Type definitions go here -->

  </Schema>
</edmx:DataServices>

```

7.4.3. Resource Type Definitions

Resource types are defined within a namespace using EntityType elements. The Name attribute specifies the name of the resource and the BaseType specifies the base type, if any.

Redfish resources derive from a common Resource base type named "Resource" in the Resource.0.96.0 namespace.

The EntityType contains the property and navigation property elements that define the resource, as well as annotations describing the resource.

```

<EntityType Name="TypeA" BaseType="Resource.Resource">
  <Annotation Term="Core.Description" String="This is the description of TypeA."/>
  <Annotation Term="Core.LongDescription" String="This is the specification of TypeA."/>

  <!-- Property and Navigation Property definitions go here -->

</EntityType>

```

All resources shall include Description and LongDescription annotations.

7.4.4. Resource Properties

Structural properties of the resource are defined using the Property element. The Name attribute specifies the name of the property, and the Type its type. Properties that must have a non-nullable value include the nullable attribute with a value of "false".

```

<Property Name="Property1" Type="Edm.String" Nullable="false">
  <Annotation Term="Core.Description" String="This is a property of TypeA."/>
  <Annotation Term="Core.LongDescription" String="This is the specification of Property1."/>

```

```
<Annotation Term="OData.Permissions" EnumMember="OData.Permissions/Read" />
<Annotation Term="DMTF.Required" />
<Annotation Term="OData.Measures.Unit" String="Watts" />
</Property>
```

All properties shall include [Description](#) and [LongDescription](#) annotations.

Properties that are read-only are annotated with the [Permissions](#) annotation with a value of `ODataPermissions/Read`.

Properties that are required to be implemented by all services are annotated with the [required](#) annotation.

Properties that have units associated with them can be annotated with the [units](#) annotation

7.4.4.1. Property Types

Type type of a property is specified by the `Type` attribute. The value of the type attribute may be a [primitive type](#), a [structured type](#), an [enumeration type](#) or a [collection](#) of primitive, structured or enumeration types.

7.4.4.1.1. Primitive Types

Primitive types are prefixed with the "Edm" namespace prefix.

Redfish services support the following primitive types:

Type	Meaning
Edm.Boolean	True or False
Edm.DateTimeOffset	Date and time with a time-zone
Edm.Decimal	Numeric values with fixed precision and scale
Edm.Double	IEEE 754 binary64 floating-point number (15-17 decimal digits)
Edm.Duration	Signed duration in days, hours, minutes, and (sub)seconds
Edm.Int64	Signed 64-bit integer
Edm.String	Sequence of UTF-8 characters

7.4.4.1.2. Structured Types

Structured types are defined within a [namespace](#) using `ComplexType` elements. The `Name` attribute of the complex type specifies the name of the structured type. Complex types can include a `BaseType` attribute to specifies the base type, if any.

Structured types may be reused across different properties of different resource types.

```
<ComplexType Name="PropertyTypeA">
  <Annotation Term="Core.Description" String="This is type used to describe a structured property." />
  <Annotation Term="Core.LongDescription" String="This is the specification of the type." />

  <!-- Property and Navigation Property definitions go here -->

</ComplexType>
```

Structured types can contain [properties](#), [navigation properties](#) and annotations.

Structured types shall include [Description](#) and [LongDescription](#) annotations.

7.4.4.1.3. Enums

Enumeration types are defined within a [namespace](#) using `EnumType` elements. The `Name` attribute of the enumeration type specifies the name of the enumeration type.

Enumeration types may be reused across different properties of different resource types.

EnumType elements contain Member elements that define the members of the enumeration. The Member elements contain a Name attribute that specifies the string value of the member name.

```
<EnumType Name="EnumTypeA">
  <Annotation Term="Core.Description" String="This is the EnumTypeA enumeration."/>
  <Annotation Term="Core.LongDescription" String="This is used to describe the EnumTypeA enumeration."/>
  <Member Name="MemberA">
    <Annotation Term="Core.Description" String="Description of MemberA"/>
  </Member>
  <Member Name="MemberB">
    <Annotation Term="Core.Description" String="Description of MemberB"/>
  </Member>
</EnumType>
```

Enumeration Types shall include [Description](#) and [LongDescription](#) annotations.

Enumeration Members shall include [Description](#) annotations.

7.4.4.1.4. Collections

The [type](#) attribute may specify a collection of [primitive](#), [structured](#) or [enumeration](#) types.

The value of the type attribute for a collection-valued property is of the form:

Collection(*NamespaceQualifiedTypeName*)

where *NamespaceQualifiedTypeName* is the namespace or alias qualified name of the primitive, structured, or enumeration type.

7.4.4.2. Non-Nullable properties

Properties may include the [Nullable](#) attribute with a value of false to specify that the property cannot contain null values. A property with a nullable attribute with a value of "true", or no nullable attribute, can accept null values.

```
<Property Name="Property1" Type="Edm.String" Nullable="false">
```

7.4.4.3. Read-only properties

The [Permissions](#) annotation term can be applied to a property with the value of [OData.Permissions/Read](#) in order to specify that it is read-only.

```
<Annotation Term="OData.Permissions" EnumMember="OData.Permissions/Read"/>
```

The [Permissions](#) annotation term is defined in <http://docs.oasis-open.org/odata/odata/v4.0/os/vocabularies/Org.OData.Core.V1.xml>.

7.4.4.4. Required Properties

The [Required](#) annotation term is used to specify that a property is required to be supported by services. Properties not annotated with the [Required](#) annotation, or annotated with a [Boolean](#) attribute with a value of "false", are optional.

If an implementation supports a property, it shall always provide a value for that property. If a value is unknown, then null is an acceptable values in most cases. Properties not returned from a GET operation shall indicate that the property is not currently supported by the implementation.

```
<Annotation Term="DMTF.Required"/>
```

The [Required](#) annotation term is defined in <http://schemas.dmtf.org/redfish/v1/DMTFExtensions.0.96.0>.

7.4.4.5. Units of Measure

In addition to following [naming conventions](#), properties representing units of measure shall be annotated with the Units annotation term in order to specify the units of measurement for the property.

```
<Annotation Term="OData.Measures.Unit" String="Watts"/>
```

The Unit annotation term is defined in <http://docs.oasis-open.org/odata/odata/v4.0/os/vocabularies/Org.OData.Measures.V1.xml>.

The string values for the Units of Measure annotation shall be taken from **TODO**.

7.4.4.6. Language-Dependent Property Values

Property names are never localized. User-supplied string-valued property values such as an asset tag may be localized. Localizable string valued properties should be annotated with the IsLanguageDependent annotation term.

Properties not annotated with the IsLanguageDependent annotation, or annotated with a Boolean attribute with a value of "false", do not have values that vary based on localization.

```
<Annotation Term="OData.IsLanguageDependent"/>
```

The IsLanguageDependent annotation term is defined in <http://docs.oasis-open.org/odata/odata/v4.0/os/vocabularies/Org.OData.Core.V1.xml>.

7.4.5. Reference Properties

Properties that reference other resources are represented as navigation properties. The NavigationProperty element specifies the Name and namespace (or alias) qualified Type of the related resource(s).

If the property references a single type, the value of the type attribute is the namespace (or alias) qualified name of the related resource type.

```
<NavigationProperty Name="RelatedType" Type="MyTypes.TypeB">
  <Annotation Term="Core.Description" String="This property references a related resource."/>
  <Annotation Term="Core.LongDescription" String="This is the specification of the related property."/>
  <Annotation Term="DMTF.ExpandReferences"/>
</NavigationProperty>
```

If the property references a collection of resources, the value of the type attribute is of the form:

Collection(NamespaceQualifiedTypeName)

where NamespaceQualifiedTypeName is the namespace (or alias) qualified name of the type of related resources.

```
<NavigationProperty Name="RelatedType" Type="Collection(MyTypes.TypeB)">
  <Annotation Term="Core.Description" String="This property represents a collection of related resources."/>
  <Annotation Term="Core.LongDescription" String="This is the specification of the related property."/>
  <Annotation Term="DMTF.ExpandReferences"/>
</NavigationProperty>
```

All navigation properties shall include [Description](#) and [LongDescription](#) annotations.

7.4.5.1. Expanded References

Navigation properties in a Redfish JSON payload are expanded to include the [related resource id](#) or [collection of related resource ids](#). This behavior is expressed using the ExpandReferences annotation.

```
<Annotation Term="DMTF.ExpandReferences" />
```

The `ExpandReferences` annotation term is defined in <http://schemas.dmtf.org/redfish/v1/DMTFExtensions.0.96.0>.

7.4.5.2. Expanded Resources

This term can be applied to a navigation property in order to specify that the default behavior for the service is to expand the related `resource` or `collection of resources` in responses.

```
<Annotation Term="DMTF.ExpandResources" />
```

The `ExpandResources` annotation term is defined in <http://schemas.dmtf.org/redfish/v1/DMTFExtensions.0.96.0>.

7.4.6. Resource Actions

Actions are grouped under a property named "Actions".

```
<Property Name="Actions" Type="MyType.Actions">
```

The type of the `Actions` property is a `structured type` with a single OEM property whose type is a structured type with no defined properties.

```
<ComplexType Name="Actions">
  <Property Name="OEM" Type="MyType.OEMActions" />
</ComplexType>

<ComplexType Name="OEMActions" />
```

Individual actions are defined within a `namespace` using `Action` elements. The `Name` attribute of the action specifies the name of the action. The `Isbound` attribute specifies that the action is bound to (appears as a member of) a resource or structured type.

The `Action` element contains one or more `Parameter` elements that specify the `Name` and `Type` of each parameter.

The first parameter is called the "binding parameter" and specifies the resource or `structural type` that the action appears as a member of (the type of the `Actions` property on the resource). The remaining `Parameter` elements describe additional parameters to be passed to the action.

```
<Action Name="MyAction" Isbound="true">
  <Parameter Name="Thing" Type="MyType.Actions" />
  <Parameter Name="Parameter1" Type="Edm.Boolean" />
</Action>
```

7.4.7. Resource Extensibility

Vendors can define additional `properties`, `links`, and `actions` for common Redfish resources using the OEM property on resources, links, and actions.

While the information and semantics of these extensions are outside of the standard, the schema representing the data, the resource itself, and the semantics around the protocol shall conform to the requirements in this specification.

7.4.7.1. Property Extensions

Resources contain a property called "Oem" whose value is an empty "OEM" `complex type` defined for the resource.

```
<ComplexType Name="MyType.OEMType"/>
```

Vendor-specific information is defined in OEM-specific complex types.

```
<Schema Name="acme.v.v.v" Alias="acme">
  <ComplexType Name="acmetype">
    <Property Name="AcmeSpecificProperty" Type="Edm.String"/>
  </ComplexType>
</Schema>
```

The OEM-specific types are exposed as dynamic properties of the root "OEM" type named using a unique OEM identifier that shall be either an ICANN-recognized domain name including the top-level domain suffix, or an IANA-assigned Enterprise ID prefaced with "EID:". Organizations using a '.com' top level domain may omit the suffix (e.g. ACME.com may use 'ACME', but ACME.org must use 'ACME.org' as their OEM property name). This property name may be followed by a colon and any string to allow further namespacing of vendor objects.

```
"Oem": {
  "Acme": {
    "@odata.type": "acme.v.v.v#acme.acmetype",
    "AcmeSpecificProperty": "value"
  }
}
```

7.4.7.2. Custom Actions

OEM-specific actions can be defined by defining actions bound to the OEM property of the [resource's Actions](#) property type.

```
<Action Name="Ping" Isbound="true">
  <Parameter Name="acmetype" Type="MyType.OEMActions"/>
</Action>
</Schema>
```

Such bound actions appear in the JSON payload as properties of the Oem type, nested under an [Actions](#) property.

```
"Actions": {
  "OEM": {
    "acme.v.v.v#acme.Ping": {
      "target": "/redfish/v1/Systems/1/Actions/OEM/acme.Ping"
    }
  }
}
```

7.4.7.3. Custom Annotations

This specification defines a set of common annotations for extending the definition of resource types used by Redfish. In addition, services may define custom annotations.

Services may apply annotations to resources in order to provide service-specific information about the type, such as whether the service supports modifications of particular properties.

Services can apply annotations to existing resources where those resources don't already define a value for the annotation. Services cannot change the value of an annotation applied as part of the resource definition.

Because [service annotations](#) may be applied to existing resource definitions, they are generally specified in a service-specific metadata document referenced by the [service metadata](#).

7.5. Common Redfish Resource Properties

This section contains a set of common properties across all Redfish resources. The property names in this section shall not be used for any other purpose, even if they are not implemented in a particular resource.

Common properties are defined in the base Resource.0.96.0.Resource schema.

7.5.1. Id

The Id property of a resource identifies the resource within a collection.

7.5.2. Name

The Name property is used to convey a human readable moniker for the resource. The type of the Name property shall be string.

7.5.3. Description

The Description property is used to convey a human readable description of the resource. The type of the Description property shall be string.

7.5.4. Modified

The Modified property contains the time stamp equal to the last time the Redfish service modified this resource. The format of this property shall be the [Standard Timestamp Format](#).

7.5.5. Status

The Status property represents the status of a resource.

The value of the status property is a common status object type as defined by this specification. By having a common representation of status, clients can depend on consistent semantics. The Status object is capable of indicating the current intended state, the state the resource has been requested to change to, the current actual state and any problem affecting the current state of the resource.

7.5.6. Links

The [Links property](#) represents the links associated with the resource, as defined by that resources schema definition. All navigation properties defined for a resource shall be nested under the links property.

7.5.7. Actions

The [Actions property](#) contains the actions supported by a resource.

7.5.8. OEM

The [OEM property](#) is used for OEM extensions as defined in [Schema Extensibility](#).

7.5.9. Settings

The Settings property contains a URI that the client can use to PUT properties or PATCH property changes for resources that are not directly [updatable](#).

If the resource itself is read only but has a partner Setting resource, this is used to make changes at some point in the future to the resource.

7.5.10. SettingsResult

The `SettingsResult` property represents the result of the last Setting Data.

The value of the `SettingsResult` property is a JSON object containing the results of applying a Setting resource to this resource. This `SettingsResult` object contains the following information about the last Setting Data apply result, which includes:

- Time of the attempted application
- ETag of the Setting Data object that was applied
- Redfish [Extended Information](#) containing status information

7.6. Redfish Resources

Collectively known as the Redfish Schema, the set of resource descriptions contains normative requirements on implementations conforming to this specification.

Redfish Resources are one of several general kinds:

- Root Service Resource
 - Contains the mapping of a particular service instance to applicable subtending resources.
 - Contains the UUID of a service instance. This UUID would be the same UUID returned via SSDP discovery.
- Current Configuration Resources, contain a mixture of:
 - Inventory (static and read-only)
 - Health Telemetry (dynamic and read-only)
 - Current Configuration Settings (dynamic and read/write)
 - Current Metric values
- Setting Resources
 - Dynamic, Read/Write Pending Configuration Settings
- Services
 - Common services like Eventing, Tasks, Sessions
- Registry Resources
 - Static, Read-Only JSON encoded information for Event and Message Registries

7.6.1. Current Configuration

Current Configuration resources represent the service's knowledge of the current state and configuration of the resource. This may be directly updatable with a PATCH or it may be read-only by the client and the client must PATCH to a separate Setting resource.

7.6.2. Settings

Setting resources represent the future state and configuration of the resource. This property is always paired with a Current Configuration object. Where the Current Configuration object represents the current state, the settings resource represents the future intended state. The state of the resource is changed either directly, such as with a POST of an action or PUT of a reset or indirectly, such as when a user reboots a machine outside of the Redfish service.

7.6.3. Services

Service resources represent components of the Redfish Service itself as well as dependent resources. While the complete list is discoverable only by traversing the Redfish Service tree, the list includes services like the Eventing service, Task management and Session management.

7.6.4. Registry

Registry resources are those resources that assist the client in interpreting Redfish resources beyond the Schema definitions. Examples of registries include Message Registries, Event Registries and enumeration registries, such as those used for BIOS. In registries, a identifier is used to retrieve more information about a given resource, event, message or other item. This can include other properties, property restrictions and the like. Registries are themselves resources.

7.7. Special Resource Situations

There are some situations that arise with certain kinds of resources that need to exhibit common semantic behavior.

7.7.1. Absent Resources

Resources may be either absent or their state unknown at the time a client requests information about that resource. For removed resources where the URI is expected to

remain constant (such as when a fan is removed), the Resource should represent the State property of the Status object as "Absent". In this circumstance, any required or supported properties for which there is no known value shall be represented as null.

7.7.2. Schema Variations

There are cases when deviations from the published schema are necessary. An example is BIOS where different servers may have minor variations in available configuration settings. A provider may build a single schema that is a superset of the individual implementations. In order to support these variations, Redfish supports omitting parameters defined in the class schema in the current configuration object. The following rules apply:

- All Redfish services must support attempts to set unsupported configuration elements in the Setting Data by marking them as exceptions in the Setting Data Apply status structure, but not failing the entire configuration operation.
- The support of a specific property in a resource is signaled by the presence of that property in the Current Configuration object. If the element is missing from Current Configuration, the client may assume the element is not supported on that resource.
- For ENUM configuration items that may have variation in allowable values, a special read-only capabilities element will be added to Current Configuration which specifies limits to the element. This is an override for the schema only to be used when necessary.

Providers may split the schema resources into separate files such as Schema + String Registry, each with a separate URI and different Content-Encoding.

- Resources may communicate omissions from the published schema via the Current Configuration object if applicable.

8. Service Details

8.1. Eventing

This section covers the REST-based mechanism for subscribing to and receiving event messages.

The Redfish service requires a client or administrator to create subscriptions to receive events. A subscription is created when an administrator sends an HTTP POST message to the URI of the subscription resource. This request includes the URI where an event-receiver client expects events to be sent, as well as the type of events to be sent. The Redfish service will then, when an event is triggered within the service, send an event to that URI.

- Services shall support "push" style eventing for all resources capable of sending events.
- Services shall not "push" events (using HTTP POST) unless an event subscription has been created. Either the client or the service can terminate the event stream at any time by deleting the subscription. The service may delete a subscription if the number of delivery errors exceeds pre-configured thresholds.
- Services shall respond to a successful subscription with HTTP status 201 and set the HTTP Location header to the address of a new subscription resource. Subscriptions are persistent and will remain across event service restarts.
- Clients shall terminate a subscription by sending an HTTP DELETE message to the URI of the subscription resource.
- Services may terminate a subscription by sending a special "subscription terminated" event as the last message. Future requests to the associated subscription resource will respond with HTTP status 404.

There are two types of events generated in a Redfish service - life cycle and alert.

Life cycle events happen when resources are created, modified or destroyed. Not every modification of a resource will result in an event - this is similar to when ETags are changed and implementations may not send an event for every resource change. For instance, if an event was sent for every Ethernet packet received or every time a sensor changed 1 degree, this could result in more events than fits a scalable interface. This event usually indicates the resource that changed as well as, optionally, any attributes that changed.

Alert events happen when a resource needs to indicate an event of some significance. This may be either directly or indirectly pertaining to the resource. This style of event usually adopts a message registry approach similar to extended error handling in that a MessageID will be included. Examples of this kind of event are when a chassis is opened, button is pushed, cable is unplugged or threshold exceeded. These events usually do not correspond well to life cycle type events hence they have their own category.

NOTE: Refer to the [Security](#) section for security implications of Eventing

8.1.1. Event Message Subscription

The client locates the eventing service through traversing the Redfish service interface. When the eventing service has been discovered, clients subscribe to messages by sending a HTTP POST to the URL of the collection for subscriptions in the Eventing Service for which they are requesting events. This should be found off of the root service as described in the schema for that service.

The specific syntax of the subscription body is found in the Redfish Schema.

On success, the "subscribe" action shall return with HTTP status 201 (CREATED) and the Location header in the response shall contain a URI giving the location of the newly created "subscription" resource. The body of the response, if any, shall contain a representation of the subscription resource. Sending an HTTP GET to the subscription resource shall return the configuration of the subscription.

8.1.2. Event Message Objects

Event message objects POSTed to the specified client endpoint shall contain the properties as described in the Redfish Event Schema.

This event message structure supports a message registry. In a message registry approach there is a message registry that has a list or array of MessageIDs in a well known format. These MessageIDs are terse in nature and thus they are much smaller than actual messages, making them suitable for embedded environments. In the registry, there is also a message. The message itself can have arguments as well as default values for Severity and RecommendedActions.

The MessageID property contents shall be of the form

RegistryName.MajorVersion.MinorVersion.MessageKey

where

- *RegistryName* is the name of the registry. The registry name shall be Pascal-cased.
- *MajorVersion* is a positive integer representing the major version of the registry
- *MinorVersion* is a positive integer representing the minor version of the registry
- *MessageKey* is a human-readable key into the registry. The message key shall be Pascal-cased and shall not include spaces, periods or special chars.

8.1.3. Subscription Cleanup

To unsubscribe from the messages associated with this subscription, the client or administrator simply sends an HTTP DELETE request to the subscription resource URI.

In order to avoid "orphan" subscriptions (subscriptions not cleaned up by the client, e.g., in the case the client has died or simply forgets to delete a subscription), the event service can be configured to automatically delete subscriptions under the following circumstances:

- The service received an error POSTing to the configured event destination URL (client-url field of the subscription) some service-defined number of consecutive times within a configurable time window. See the properties defined in the EventService schema for details.

8.2. Asynchronous Operations

Services that support asynchronous operations will implement the Task service & Task resource.

The Task service is used to describe the service that handles tasks. It contains a collection of zero or more task resources. The Task resource is used to describe a long running operation that is spawned when a request will take longer than a few seconds, such as when a service is instantiated. Clients will poll the URI of the task resource to determine when the operation has completed and if it was successful.

The Task structure in the Redfish Schema contains the exact structure of a Task. The type of information it contains are start time, end time, task state, task status, response (completion codes) as well as potential links to sub-tasks that were spawned.

Each task has a number of possible states. The exact states and their semantics are defined in the Task resource of the Redfish Schema.

When a client issues a request for a long-running operation, the service returns a status of 202 (Accepted).

Any response with a status code of 202 (Accepted) shall include a location header containing the URL of a monitor for the task and may include a wait header to specify the amount of time the client should wait before querying status of the operation.

The response body of a 202 (Accepted) should contain an instance of the Task resource describing the state of the task.

As long as the operation is in process, the service shall continue to return a status code of 202 (Accepted) when querying the status monitor returned in the location header.

Once the operation has completed, the status monitor shall return a status code of OK (200) and include in the body of the response a mime message containing the results of the initial operation, as if it had completed synchronously.

The client can continue to get information about the status by directly querying the Task resource using the [resource identifier](#) returned in the body of the 202 (Accepted) response.

- Services that support asynchronous operations shall implement the Task resource
- The response to an asynchronous operation shall return a status code of 202 (Accepted) and set the HTTP response header "Location" to the URI of a status monitor

associated with the activity. The response may also include a wait header specifying the amount of time the client should wait before polling for status. The response body should contain a representation of the Task resource in the same major media type (e.g. JSON, XML) that would have been used to return a synchronous response.

- GET requests to either the Task monitor or the Task Resource shall return the current status of the operation without blocking.
- Operations using HTTP GET, PUT, PATCH should always be synchronous.
- Clients shall be prepared to handle both synchronous and asynchronous responses for requests using HTTP DELETE and HTTP POST methods.

8.3. Resource Tree Stability

The Resource Tree must be consistent on a single service across device reboot and A/C power cycle, and must withstand a reasonable amount of configuration change (e.g. adding an adapter to a server). The resource Tree on one service may not be consistent across instances of devices. The client must walk the data model and discover resources to interact with them. It is possible that some resources will remain very stable from system to system (e.g. BMC network settings) -- but it is not an architectural guarantee.

- A Resource Tree should remain stable across Service restarts and minor device configuration changes
- A Resource Tree shall not be expected by the client to be consistent between instances of services

8.4. Discovery

Automatic discovery of managed devices supporting the Redfish Scalable Platform Management API is accomplished using the Simple Service Discovery Protocol (SSDP). This protocol allows for network-efficient discovery without resorting to ping-sweeps, router table searches, or restrictive DNS naming schemes. Use of SSDP is optional, and if implemented, shall allow the user to disable the protocol through the 'Manager Network Service' resource.

As the objective of discovery is for client software to locate Redfish-compliant managed devices, the primary SSDP functionality incorporated is the M-SEARCH query. Redfish also follows the SSDP extensions and naming used by UPnP where applicable, such that Redfish-compliant systems can also implement UPnP without conflict.

8.4.1. UPnP Compatibility

For compatibility with general purpose SSDP client software, primarily UPnP, TCP port 1900 will be used for all SSDP traffic. It is recommended that devices also respond to M-SEARCH queries for UPnP Root Devices (with NT:upnp:rootdevice), with appropriate descriptors and XML documents.

8.4.2. USN Format

The UUID supplied in the USN field shall equal the UUID returned for the Manager implementing the Redfish service. If there are multiple / redundant managers, the UUID shall remain static regardless of redundancy failover. The Unique ID shall be in the canonical UUID format, followed by '!:dmf-org'

8.4.3. M-SEARCH Response

The managed device must respond to M-SEARCH queries searching for Search Target (ST) of the Redfish Service from clients with the AL pointing to the Redfish service root URI. Redfish device shall also respond to M-SEARCH queries for Search Target type of "ssdp:all".

Redfish Service root Search Target (ST): URN:dmtf-org:service:redfish-rest:1

The URN in the reply shall use a service name of 'redfish-rest:' followed by the major version of the Redfish specification. If the minor version of the Redfish Specification to which the service conforms is a non-zero value, and that version is backwards-compatible with previous minor revisions, then that minor version shall be appended, preceeded with a colon. For example, a service conforming to a Redfish specification version "1.4" would reply with a service of "redfish-rest:1:4".

An example response to an M-SEARCH multicast or unicast query shall follow the format shown below. Fields in brackets are placeholders for device-specific values.

```
HTTP/1.1 200 OK
CACHE-CONTROL:<seconds, at least 1800>
ST:urn:dmtf-org:service:redfish-rest:1
USN:uuid:<UUID of Manager>::urn:dmtf-org:service:redfish-rest:1
AL:<URL of Redfish service root>
EXT:
```

8.4.4. Notify, Alive, and Shutdown messages

Redfish devices may implement the additional SSDP messages defined by UPnP to announce their availability to software. This capability, if implemented, must allow the end

user to disable the traffic separately from the M-SEARCH response functionality. This allows users to utilize the discovery functionality with minimal amounts of network traffic generated.

9. Security

9.1. Goals

- Privilege Model to Monitor and Manage:
 - System Settings
 - BIOS Configuration
 - System Power States
 - Sensor Information (power/thermal/health)
 - Network Settings
 - Storage Settings
 - Logs
 - Redfish Service Configuration
 - Account Management
 - Network Settings
 - Logs
 - Firmware versions
 - OEM vendor-specific features and functionality
- Permission/ authorization model shall be consistent between instances of Redfish compliant devices
 - Define a minimum baseline for the permission/ authorization model
- Infrastructure Authentication
- CURL compatibility
- Automated clients
- Embedded Service Processors

9.2. Protocols

9.2.1. Discovery

9.2.2. Transport

TODO: communication requirements for Redfish peers - clients or other servers.

9.2.2.1. TLS

Implementations shall support TLS v1.1 or later

Implementations shall only use compliant TLS connections to transport [Sensitive Data](#). Including to any third party authentication services or clients.

9.2.2.2. Cipher suites

Implementations should support AES-256 based ciphers from the TLS suites.

Redfish implementations should consider supporting ciphers similar to below which enable authentication and identification without use of trusted certificates.

```
TLS_PSK_WITH_AES_256_GCM_SHA384
TLS_DHE_PSK_WITH_AES_256_GCM_SHA384
TLS_RSA_PSK_WITH_AES_256_GCM_SHA384
```

Additional advantage with using above recommended ciphers is -

"AES-GCM is not only efficient and secure, but hardware implementations can achieve high speeds with low cost and low latency, because the mode can be pipelined."

References to RFCs -

```
http://tools.ietf.org/html/rfc5487
http://tools.ietf.org/html/rfc5288
```

9.2.2.3. Certificates

Implementations shall support replacement of the default certificate if one is provided, with a certificate having at least a 4096 bit RSA key and sha512-rsa signature.

9.3. Sensitive Data

- Sensitive data shall minimally include
 - User credentials (usernames, passwords)
 - Private Keys
 - Persistent Session Keys
 - Password Complexity Requirements
 - Critical Security Parameters (CSPs) as defined below

Critical Security Parameters (CSP): Security-related information (e.g., cryptographic keys, authentication data such as passwords and PINs) appearing in plaintext or otherwise unprotected form and whose disclosure or modification can compromise the security of a cryptographic module or the security of the information protected by the module.

REF: <http://csrc.nist.gov/publications/fips/fips140-1/fips1401.pdf> - This is used only as a source for the definition of CSPs. No additional requirements for compatibility with the FIPS standard should be inferred.

9.4. Authentication

- Default Credentials

Services should NOT implement default credentials for any account installed on the spec compliant device, with a well known password.

- Password Complexity

A Redfish device shall support a set of configurable password complexity rules, should include length, character set

When an authentication failure occurs a Redfish device shall NOT provide password complexity requirements to the user

- Account Lockout

A Redfish implementations should revoke login privilege after a configurable number of authentication failures

- Authentication Failure Delays

A Redfish device shall implement progressive authentication attempt

- Authentication Failure Policy

A Redfish device shall implement configurable options to manage login behavior when user authentication failures occur

```
- Number of failures
- Time between failures
- Increments between progressive delays
- Specific actions which happen automatically when failure events occur
- Time since last failed authentication attempt
```

9.4.1. HTTP Header Security

- All write activities shall be authenticated, i.e. POST, PUT/PATCH, and DELETE, except for
 - The POST operation to the Sessions service/object needed for authentication

- Extended error messages shall NOT provide privileged info when authentication failures occur
- REST objects shall not be available unauthenticated, except for
 - The root object which is needed to identify the device and service locations
- Unauthenticated REST operation results shall not contain [Sensitive Data](#).
- External services linked via extref references are not part of this spec, and may have other security requirements.
- CORS headers are not recommended. Services may choose to implement them.

9.4.1.1. HTTP Redirect

- When there is a HTTP Redirect the privilege requirements for the target resource shall be enforced

9.4.2. Extended Error Handling

- Authentication shall occur when [Sensitive Data](#) is present in any part of the REST object.
- Extended error messages shall NOT provide privileged info when authentication failures occur

9.4.3. HTTP Header Authentication

- HTTP Headers for authentication shall be processed before other headers that may affect the response, i.e.: etag, If-Modified, etc.
- HTTP Cookies shall NOT be used to authenticate any activity i.e.: GET, POST, PUT/PATCH, and DELETE.

9.4.3.1. BASIC authentication

HTTP BASIC authentication as defined by [RFC2617](#) shall be supported, and shall only use compliant TLS connections to transport the data between any third party authentication service and clients.

9.4.3.2. Digest authentication

Implementations may support HTTP Digest authentication mechanism

9.4.3.3. Negotiate

Implementations may support the HTTP Negotiate authentication mechanism

9.4.3.4. Request / Message Level Authentication

Every request that establishes a secure channel shall be accompanied by an authentication header.

9.4.3.5. Certificate based authentication?

Implementations should support certificate based authentication.

9.4.4. Session Management

9.4.4.1. Session Lifecycle Management

Session management is left to the implementation of the Redfish Service. This includes orphaned session timeout and number of simultaneous open sessions.

- **A Redfish Service shall provide login sessions compliant with this specification.**

9.4.4.2. Login Sessions

For functionality requiring multiple Redfish operations, a standard Login session is specified. The URI used for session management is specified in the /redfish resource with the property "SessionService".

```
{
  ...
  "SessionService": {
    "@odata.id": "/redfish/v1/SessionService"
  },
}
```

```

    ...
  }

```

9.4.4.3. Login

A session is created by an HTTP POST to the SessionService/Sessions resource, including the following POST body:

```

{
  "UserName": "<username>"
  "Password": "<password>"
}

```

The Origin header should be saved in reference to this session creation and compared to subsequent requests using this session to verify the request has been initiated from an authorized client domain.

The return includes an X-Auth-Token header with a session token and Location header.

The return JSON body includes a representation of the newly created session object:

```

<operation> <uri> HTTP/1.1
<header>
<header>
Location: "/redfish/v1/SessionService/Sessions/Administrator1"
X-Auth-Token: <token string>
<header>

{
  "@odata.context": "/redfish/v1/$metadata#SessionService/Links/Sessions/$entity",
  "@odata.id": "/redfish/v1/SessionService/Sessions/Administrator1",
  "UserName": "<username>"
}

```

9.4.4.4. Logout

Logout is accomplished by performing a DELETE to the Session resource provided by the Login operation including the X-Auth-Token header. Optionally, the service may also support logout upon DELETE to the Sessions resource without specifying the individual session URI.

The ability to DELETE to a Session resource allows the service to support logging out of one or more sessions from a different session if the user has sufficient privilege to do so.

9.4.4.5. X-Auth-Token HTTP Header

Implementations shall only use compliant TLS connections to transport the data between any third party authentication service and clients.

- Limited Lifetime
 - Implementations shall support session timeout, session idle time is defined as time from the last accepted transaction.
 - Session timeout shall default to a finite limit.
 - Implementations should NOT support infinite session times.

9.4.5. AccountService

- User passwords should be stored with one-way encryption techniques.
- Implementations may support exporting user accounts with passwords, but shall do so using encryption methods to protect them.
- The root REST object should be available unauthenticated, but shall not contain any [Sensitive Data](#).
 - Authentication shall occur when [Sensitive Data](#) is present in the root REST object.
- User accounts shall support ETags and shall support atomic operations
 - Implementations may reject requests which do not include an ETag
- User Management activity is atomic

- Extended error messages shall NOT provide privileged info when authentication failures occur

9.4.6. Async Tasks

- Irrespective of which users/ privileged context was used to start an async task the information in the status object shall be used to enforce the privilege(s) required to access that object.

9.4.7. Event Subscriptions

- The Redfish device shall check the privilege of the subscriber before pushing event data object to the destination
- The Redfish device shall encrypt event data when there is [Sensitive Data](#) in the event data object before pushing it
- The Redfish device may verify the destination for identity purposes before pushing event data object to the Destination

9.4.8. Privilege Model / Authorization

- The Authorization subsystem has the following components to it
 - Profiles - A Profile is a collection of Roles from the Authorization perspective (NOTE: Profiles per Redfish constitute more than Roles. REF: See section ... for more info.)
 - Roles - A Role is a collection of Privileges
 - Privileges - A Privileges is a permission to perform a specific action/ activity
 - NOTE: Specific sets of privileges can be assigned to a user without using/ modifying/ creating/ leveraging pre-canned roles
- When a Profile like a System Profile is implemented all the required Roles per Redfish which constitute that Profile shall be implemented
 - System Administrator and System Operator are 2 Roles are which are required to be implemented in the System Profile
 - System Administrator - The user with this role monitors and configures the system and / or the Redfish device
- When a Role like a System Operator is implemented all Privileges required per Redfish which constitute that Role are required to be implemented
 - The Power Control privilege is required to be implemented per Redfish
- User Management
 - Assigning privileges to users, either local to the Redfish device or users / user groups from the directory services infrastructure like AD/ LDAP
 - Assign privileges individually (Users created by cherry-picking privileges) OR
 - Assign privileges via pre-defined roles (Roles required by Redfish spec or OEM implementation) OR
 - Assign privileges via custom roles (End user admin defined roles which can be reused)
- Implementations shall enforce the same privilege model for ETag related activity as is enforced for the data being represented by the ETag. For example, when activity requiring privileged access to read data item represented by ETag requires the same privileged access to read the ETag.
- Privileges
 - shall implement a set of pre-defined privileges
 - Login
 - Configure Manager
 - Configure Users
 - NOTE: The Login privilege is automatically assigned to all users
 - shall implement a set of pre-defined privileges as required per implemented profiles
 - System Profile has the Power Control privilege
 - Other Profiles will have appropriate required privileges
 - may implement a set of OEM privileges
 - Remote Console (Remote Keyboard, Video, Mouse)
 - Remote Media (Remote mounting of media like USB storage, file shares etc)
 - Diagnostic capability
 - Clear Logs
 - OEM Privileges shall follow the requirements below

- Privileges that grant permission to affect/ modify the object and/ or object extensions in a Profile, themselves belong to that Profile
- Privileges that make use of methods / derived methods / OEM methods in a profile, themselves belong to that Profile

9.4.8.1. Profiles and Roles

- shall implement a set of pre-defined roles based on profiles implemented by the Redfish device
- System Profile 1) System Administrator - Monitor and configures the system and the Redfish device
 - Configure BMC
 - Configure Users
 - Power Control
 - All OEM privileges relevant to the System Profile 2) System Operator - Performs system management tasks like power control but not configuring the power subsystem
 - Power Control
 - Will not include any OEM privileges 3) User (Read-Only with very low privilege reads)
- Network Profile 1) Network Administrator - Monitor and configures the network component
 - TBD 2) Network Operator - Performs network management tasks like network operations including interface up/ down but not any persistent network configuration changes
 - TBD 3) User (Read-Only with very low privilege reads)
- Storage Profile 1) Storage Administrator - Monitor and configures the storage component
 - TBD 2) Storage Operator - Performs storage management tasks like storage operations including acknowledging bad drives to trigger rebuilding a logical volume but not any persistent storage configuration changes
 - TBD 3) User (Read-Only with very low privilege reads)
- NOTE: All OEM privileges that modify objects and/ or object extensions in a Profile, shall be assigned to the Administrator Role in that Profile.
- Pre-defined roles shall NOT be modifiable with respect to privileges assigned
- shall allow Redfish user to define custom roles
- should allow Redfish user to name custom roles to enable reuse in user context
- may implement a set of OEM roles

9.4.9. Role Based Privilege

A Redfish service may contain resources that require separated privileges. A Redfish service may limit REST operations against specific resources based upon user privileges.

- A Redfish service may limit individual resource access by Redfish clients based upon role-based privileges
- A Redfish service may limit individual resource access by Redfish clients based upon specific user account information (future -- not defined yet)

9.5. Data Model Validation

9.5.1. Schema

Server and Client implementations should check supplied data against schema and perform data validation checks to prevent vulnerabilities caused by later processing errors.

When there is a disagreement between a Server and Client on schema validation, the server may enforce its version and reject the request.

Clients shall NOT perform data interpolation unless the schema permits that.

Privileges should NOT be modified without a strong security related requirement. Schema validation shall include privilege checks when privilege requirements have been modified.

NOTE: Privilege changes as part of schema updates/ changes shall be captured in the schema change log.

Idempotent actions shall be rejected when there is a security reason to do so.

Resource definitions shall include required privileges to perform read/ RW actions on that resource.

Resource tree stability - Permissions on resources should be stable as well.

Custom Actions - Privilege model shall be applied consistently to both the body and the response. Where applicable the privilege model defined for the URI should be inherited for custom actions.

9.6. Logging

9.6.1. Required data for security log entries

Implementations shall log authentication requests including failures. Authentication login/logout log entries shall contain a user identifier that can be used to uniquely identify the client and a time stamp.

Logs shall include detailed privileged info, but shall NOT include the [Sensitive Data](#) outside a privileged user or a privileged security context.

9.6.2. Completeness of Logging

- Every entity from the originator of the RESTful service call, through every intermediary, to the very last entity in the call chain, log an entry in their audit log for the call activity triggered/ taken/ ... This means same as any RESTful service call, the audit log entry will 'be complete' for the activity performed within said entity.
 - shall - All write activities i.e. POST, PUT/ PATCH and DELETE
 - NOTE: When a new log entry is created logging the occurrence of that event is not required.
 - shall - Have the ability to log the privileged reads i.e. GETs
 - This ability may be turned on by default.
- Rejection of idempotent actions due to security reasons shall be logged

9.6.3. Content of Audit Logs

Details : Need to generate events for the following

1. logon, log-off, modification of user accounts
2. successful and rejected login attempts,
3. successful and rejected connections to nodes and other resource access attempts
4. details about the modification of user accounts
5. all changes to the system configuration,
6. information about the use of built-in utilities running in Redfish compliant-devices(e.g. low-level diagnostic tools),
7. information about accessing the system interfaces of the Redfish compliant-devices
8. network addresses and protocols (e.g. workstation IP address and protocol used for access)
9. activation and de-activation of protection measures

The file where the events are written, one or more messages per event should at least have the following information :

- User ID
- Date, time
- Event type
- Event description

10. ANNEX A (informative)

10.1. Change Log

Version	Date	Description
0.94.0	2015-13-1	Initial merge of 0.91 and 0.92 versions
0.96.0	2015-3-3	Near-final Chassis and ComputerSystem schemas. Introduction of referenceable (array) members and use for power metrics, thermal metrics. Introduction of SessionService. Added JSONSchemaFile to OData metadata, mockups. Miscellaneous clean-up.