


The SOA Magazine

Feature Article



Standards Supporting RESTful Services

by Raj Balasubramanian, Benjamin Carlyle and Cesare Pautasso

Published: January 14, 2010 (SOA Magazine Issue XXXV: January 2010)

[Download PDF](#)

[Digg This](#) • [De.licio.us](#) • [Slashdot](#) • [Technorati](#) • [StumbleUpon](#) • [Google Bookmark](#)

Abstract: The REST and SOA styles are both derived from a long history of developing architectural expertise. REST is derived from the needs of the Web to enable services owned by often competing organizations to interoperate, to rapidly evolve, to operate at very large scale, and to do so in a way that supports long-term backwards-compatibility concerns. SOA is rooted in enterprise architectures that are each aimed at producing maximum value for their respective business stakeholders. This article, comprised of an excerpt from the upcoming book "SOA with REST" [REF-1], focuses on the Web standards most relevant to the design and development of RESTful services.

Introduction

Although there are some standards bodies, there is no single governing organization that controls the Web. For the most part the success of the Web is attributed to the largely decentralized nature of its governance, established by the choice of organizations and the small set of technology standards governing the Web. In large part, the Web consists of several implementations of software products, protocols, services and applications from different vendors, individuals and organizations built in compliance with the same specification and standards.

HTTP and HTML are the key standards that drive the Web, and these are both centrally controlled albeit with some interplay between the IETF and the W3C. The decoupling of methods versus the media type has allowed a great deal of diversity to develop, and the Web has evolved freely in terms of image and video media types for example. This decoupling of standards bodies from each other, and an ability to let the market decide on the relative success of different media types has as you say been important in the flexibility of the Web and for its ability to deal with change. However, individual specifications both for method and for media type are strongly centralized and very importantly are not specific to service or its implementations.

In this article we will look at key Web standards and the organizations that produce and maintain them and discuss how these standards can be applied to RESTful services both in the Internet and enterprise spaces.

Standards Governing a RESTful Service

A RESTful service will employ a set of methods inheriting from the uniform interface, support certain media types for the content to be exposed and consumed as well as the supporting infrastructure to host these services. Looking at the Web, it is clear that not only are there standard methods and media types, but importantly that these are decoupled from each other. Methods and media types are also decoupled from the resource identifier through the Uniform Interface constraint. As we are going to see, each standard is defined and enforced within each of these architectural layers (Figure 1).

The layers help to logically cluster the landscape of the REST-related standards. Although the network and subsequent layers are important, they are not relevant in this discussion around RESTful services, hence we

start with the Transport layer which defines network connection and flow control between two end points. The Application/Protocol layer defines the application layer of the OSI stack. The Identifier/Linking layer addresses uniquely identifying resources. The Representation/Content layer addresses mechanisms used to handle resource representation.

These 4 layers form the basis on which the consumers and providers are built. The Component used to consume/provide the services and the Programmability options for the provider and consumer leverage the underlying layers for defining RESTful Services.

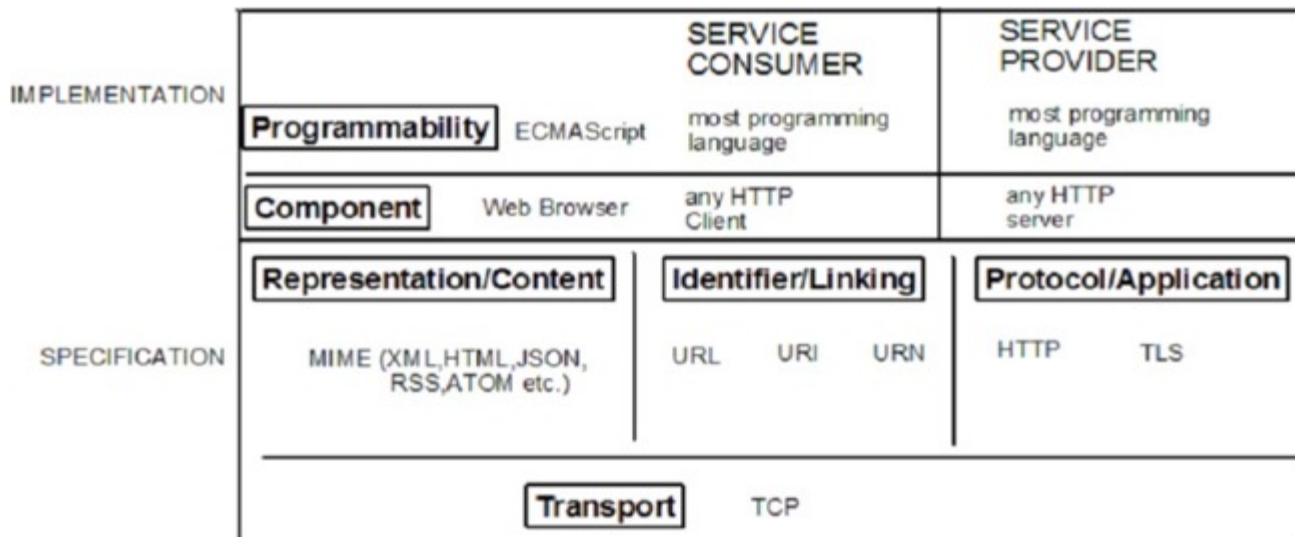


Figure 1

The primary standards around RESTful service definition of the Protocol/Application layer is HTTP. This HyperText Transfer Protocol is the foundational standard the Web is built on and also RESTful services will be built as well. The HTTP standard implements the Uniform Contract pattern by defining key methods and semantics around the methods.

The main standard defining the unique identification of resources which will be used by the ensuing layers to globally address and link to resources stem from the Uniform Resource Identifier (URI) standard and the Uniform Resource Name (URN) and Uniform Resource Link (URL) subset of this standard. RFC 3986 is the authoritative source for standards around URI. The concept of URL is again used by the Web in how Web sites and Web resources are addressed. Any Web resource that needs to be identified, such as an image, PDF document, HTML document or Javascript file all have a unique identifier in the URL used to access that resource.

A typical URL has 5 parts: `{scheme}:[//{authority}]{path}[?{query}][#{fragment}]`

The Representation/Content layer is another important layer for a RESTful service and contains the Multi-purpose Internet Mail Extensions (MIME) standard which defines the content/media types and associated metadata. The content types or media types define the representation of the resource in the message payload. To do so, the corresponding media type identifier is exchanged in an HTTP Header field called Content-Type. For example the payload of an XML document is indicated in the HTTP header as follows

```
Content-type: application/xml
```

Similarly the payload of an HTML document is indicated in the HTTP header as follows

```
Content-type: text/html
```

There are several standard content types defined as part of the MIME standard and an enterprise can leverage any or all of them. The standard also gives a well defined extensibility mechanism, where the set of existing content-types can be extended if so desired. One can use the `application/vnd.*` namespace for vendor specific namespace extensions. For custom mime-type extensions the enterprises still have to go through registration steps through the standards body. This could pose challenges for wider adoption of RESTful services from an enterprise standpoint if they chose to employ custom mime-types.

As we move to service specific layers be it for consumers or providers, the components and programmability (primarily supported by ECMAScript) become specialized and employ software implementations of the underlying standards and specifications. Specifically the components in both the service consumer and provider implement HTTP and the supporting standards as part of standard libraries - either as reference implementation of the specification or as Open Source projects in a variety of programming languages.

The obvious consumer from the Web perspective has been the Web browsers (like Lynx, Netscape Navigator, Microsoft IE, Mozilla Firefox, Opera, Chrome, etc.), but the consumers could be generic applications or other services built leveraging the HTTP client libraries (e.g., libcurl) or specialized tools (e.g., curl, wget). Similarly the well known service providers for the Web are Web servers (like Apache HTTPD Server, Microsoft IIS, IBM HTTP Server, Jetty, Tomcat, etc.), but the providers could be implemented using standard HTTP libraries on any of the available platforms. For example, Restlet is an Open Source implementation of the HTTP Library with REST specific nuances in Java that can be hosted on any J2EE container or implemented as a standalone service.

We will see more example of software implementations for developing RESTful services on variety of platforms both from software vendors as well as from several Open Source projects in the section on Realization of REST in Software.

Standards Bodies

So far, we've described several standards that can support a RESTful service. These standards are all governed by various standards organizations that are independent bodies or consortia of vendors and individuals (and organizations) or government sponsored entities. The standard bodies maintain and govern these standards and specifications. To an enterprise starting to adopt RESTful services in achieving their SOA, one or more of these standards organization might provide a good model to start your own service governance model. For practical purposes, most of the relevant RESTful standards are governed by Internet Engineering Task Force (IETF) and World Wide Web Consortium (W3C). The primary organization for the various standards of interest from the previous section is highlighted in Figure 2.

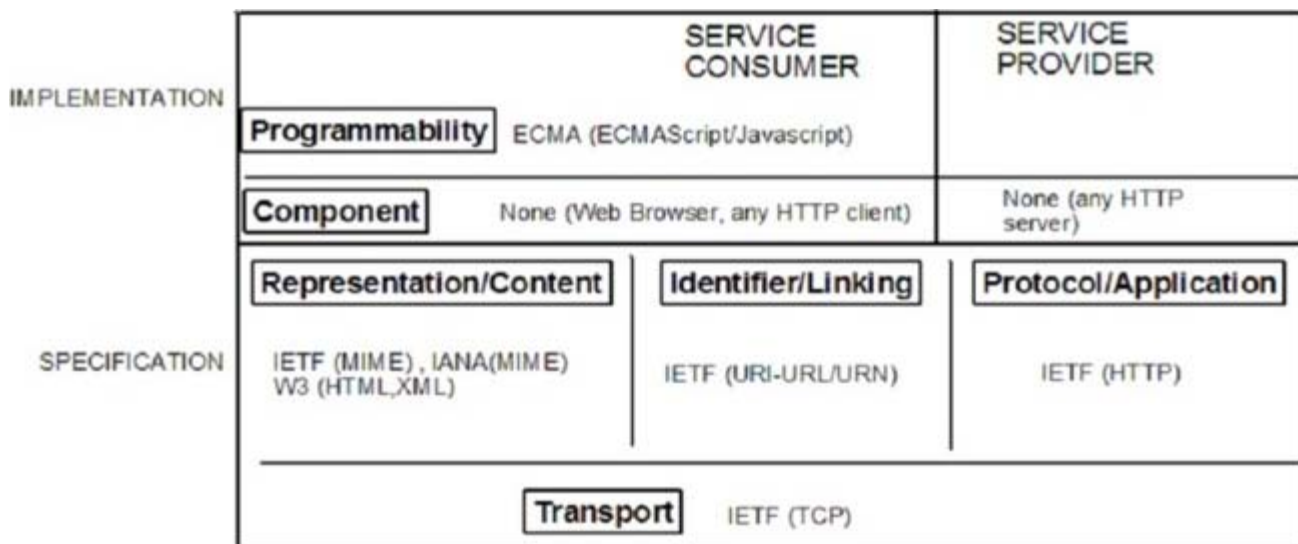


Figure 2

The IETF Governance Model

The Internet Engineering Task Force (IETF) is the primary body behind the Internet. It is responsible for specifications governing application protocols such as HTTP, transports, routing, security, real-time communications, operational management, and a range of other aspects of the modern Internet that we tend today to take for granted.

The IETF standards process is a combination of bottom-up development and top-down control mechanisms.

Anyone may pen an internet draft, on any topic they choose. These drafts expire within six months, and must generally be adopted by a working group to achieve an RFC or "request for comments" status. When a working group has chosen to promote a specific draft, it is put through review by the Internet Engineering Steering Group (IESG) who may seek to further steer the specification before it is adopted or rejected.

Each working group is limited to a specific charter. A working group's charter must be approved by the Internet Architecture Board (IAB) for architectural consistency and integrity.

A group known as the Internet Assigned Numbers Authority (IANA) controls allocation of numbers and identifiers to fit within small registries. For example, the set of identifiers for media types on the Internet is governed by the IANA. Allocation of well-known TCP port numbers is also a responsibility of this group. The IANA provides registries of important identifiers that can act as a quick reference into applicable standards and avoid identifier clashes from occurring. A separate body known as the RFC editor is responsible for assigning identifiers to standards themselves as the move out of draft phase and become RFCs.

Much of the work of the IETF occurs through mailing lists, although regular meetings are held to achieve face-to-face contact. These meetings allow working groups to accelerate their standards activities. These conferences are also the forum for the Birds of a Feather (BOF) meetings that are usually required to form and establish the charter of a new working group.

An Internet Draft may go through several stages along the standards track. Firstly it is submitted to the IESG, and a "last call" is published for comments before becoming an RFC. The last call lasts for two to four weeks. If it makes it to this stage, the RFC has become a Proposed Standard and will remain so for at least a further six months.

The next step along the standards track is to become a Draft Standard. This requires a continuing consensus around the RFC, as well as the existence of two independent and interoperable implementations of the standard. If the document makes it to Draft Standard status, it may after a few more years become a full standard known as an Internet Standard.

The "rough consensus and working code" approach to standards development of the IETF reflects a lean organization, while the long and cautious standardization process reflects the long-term nature of Internet architecture and the need to get things right. Importantly, the IETF only covers governance of standards. It does not enforce compliance to those standards. New development activity can occur outside of the IETF, and the results can be standardized only when they are already known to work well. The standardization process itself almost depends on this behavior, in an approach that takes the best that is out there and encourages rather than demands consistency with the standard. The network effects of a consistent standard are usually sufficient to ensure compliance to the standard, as interoperable technology has a greater reach than proprietary technology in a large heterogeneous network such as the Internet.

There are supporting organizations like IANA (Internet Assigned Numbers Authority) and ISOC (Internet Society). IANA distributes and maintains unique numbering of ports, IP address, DNS Routes and Internet related resources. IANA also maintains the list of Media-Types and associated parameters.

The W3C Governance Model

The World Wide Web Consortium (W3C) began as an organization designed to feed specifications into the IETF for formal standardization. The W3C has a particular focus on the Web. Unlike the IETF which is a body consisting of benevolent individuals, the W3C is a body made up of often competing organizations. The W3C governs standards for HTML, XML, the Semantic Web, SOAP-based Web Services, and a range of other activities.

Each member organization has a seat on the Advisory Committee, can access member-only information, can make member submissions, and may make use of W3C logos and promotional materials. Members may participate in working, interest, and coordination groups, in workshops and symposia, or on the W3C (management) team as W3C fellows.

The W3C team provides technical leadership and organizes and manages W3C activities. The Advisory Board provides ongoing guidance to the team on issues of strategy, management, legal matters, process, and conflict resolution. The Technical Architecture Group (TAG) documents and builds consensus around principles of Web architecture, interpreting and clarifying these principles where necessary. It resolves

issues involving general Web architecture, and helps coordinate cross-technology architecture developments inside and outside the W3C.

The general public may participate in mailing lists and workshops, comment on upcoming standards, or may be invited as experts in working groups.

Other Governing Bodies

A number of other standards bodies exist with varying degrees of influence on the Web and in Web-related domains. Bodies such as OASIS focus on SOAP-based Web Services. The microformats community standardizes rich metadata markup for inclusion within HTML documents. Some bodies such as railML.org are focused on the development of media types that address particular industry concerns. Many standards bodies are interrelated, sharing members and technology foundations. There are specific vendors, organization, individuals and Open Source projects that govern specific programming languages and hence the implementation of the HTTP specific libraries. Some of these vendors have started to extend these traditional HTTP libraries and implement libraries that are useful in building RESTful services.

Conclusion

To summarize, let's revisit the primary standards and governing bodies most relevant to RESTful service development.

Standard	Organization	Location
TCP	IETF	http://tools.ietf.org/html/rfc793
HTTP	IETF	http://www.ietf.org/rfc/rfc2616.txt
URI/URL/URN	IETF	http://www.ietf.org/rfc/rfc3986.txt
MIME	IETF	http://www.ietf.org/rfc/rfc2045.txt
HTML	W3C	http://www.w3.org/TR/REC-html40/
XML	W3C	http://www.w3.org/XML/
XSD	W3C	http://www.w3.org/XML/Schema
RSS	Harvard Law	http://cyber.law.harvard.edu/rss/rss.html
ATOM	IETF	http://www.ietf.org/rfc/rfc4287.txt
MIME type	IANA	http://www.iana.org/assignments/media-types/
JSON	JSON.ORG	http://www.json.org/
Javascript	ECMA	http://www.ecma-international.org/memento/TC39.htm

Access to these standards specifications is also available via www.soaspecs.com.

References

[REF-1] "SOA with REST" by Raj Balasubramanian, Benjamin Carlyle, Thomas Erl, Cesare Pautasso, Prentice Hall, Copyright 2010

THE PRENTICE HALL SERVICE-ORIENTED COMPUTING SERIES FROM THOMAS ERL



[Home](#) [Past Issues](#) [Contributors](#) [What is SOA?](#) [SOA Glossary](#) [SOA School](#) [SOA Books](#) [About/RSS](#)
[Legal](#)

Copyright © 2006-2010
SOA Systems Inc.