

On Using JSON-LD to Create Evolvable RESTful Services

Markus Lanthaler^{1,2}

¹ Institute for Information Systems and Computer Media
Graz University of Technology
Graz, Austria

mail@markus-lanthaler.com

Christian Gütl^{1,2}

² School of Information Systems
Curtin University of Technology
Perth, Australia

christian.guetl@iicm.tugraz.at

ABSTRACT

As the amount of data and devices on the Web experiences exponential growth issues on how to integrate such hugely heterogeneous components into a scalable system become increasingly important. REST has proven to be a viable solution for such large-scale information systems. It provides a set of architectural constraints that, when applied as a whole, result in benefits in terms of loose coupling, maintainability, evolvability, and scalability. Unfortunately, some of REST's constraints such as the ones that demand self-descriptive messages or require the use of hypermedia as the engine of application state are rarely implemented correctly. This results in tightly coupled and thus brittle systems. To solve these and other issues, we present JSON-LD, a community effort to standardize a media type targeted to machine-to-machine communication with inherent hypermedia support and rich semantics. Since JSON-LD is 100% compatible with traditional JSON, developers can continue to use their existing tools and libraries. As we show in the paper, JSON-LD can be used to build truly RESTful services that, at the same time, integrate the exposed data into the Semantic Web. The required additional design costs are significantly outweighed by the achievable benefits in terms of loose coupling, evolvability, scalability, self-descriptiveness, and maintainability.

Categories and Subject Descriptors

H.3.4 [Information Storage and Retrieval]: Systems and Software – *Semantic Web, Web 2.0, World Wide Web (WWW)*.
H.4.3 [Information Systems Applications]: Communications Applications – *Internet*. D.2.11 [Software]: Software Architectures – *Service-oriented architecture (SOA)*

General Terms

Design, Standardization

Keywords

Web services; REST; Semantic Web; Linked Data; JSON-LD; Web of Things

1. INTRODUCTION

The Internet has experienced exponential growth, yet, it is expected that in the near future the amount of data generated by machines (e.g. sensors) will exceed that created by humans by several orders of magnitude. While the lower-level technical problems of connecting such a large number of machines are

being solved, issues on how to integrate these hugely heterogeneous datasets into a scalable system become increasingly important. Reusing the Web's underlying architectural style, i.e., REST [1], has proven to be a viable solution to transform islands of data into an integrated Web of Data. It provides a set of architectural constraints that, when applied as a whole, result in a concrete system architecture that “emphasizes scalability of component interactions, generality of interfaces, independent deployment of components, and intermediary components to reduce interaction latency, enforce security, and encapsulate legacy systems” [1].

While some of REST's constraints such as stateless interaction, uniform interface, identification of resources, or manipulation of resources through representations are well understood, others are rarely implemented correctly; regardless of a service claiming to be RESTful or not. Primarily the constraints that demand self-descriptive messages and require the use of hypermedia as the engine of application state are often ignored. Instead of creating specific media types, often, general media types with poor semantics are used which results in reduced visibility and requires out-of-band knowledge to process a message. Similarly, instead of including valid state transitions in the form of hyperlinks in responses sent from the server, such knowledge is often documented out-of-band and consequently hardcoded into the client. To solve these issues, there are basically three options.

The first one, and often advocated as the cleanest solution, is to create a new media type which specifies the application's semantics and supports the required hypermedia controls to fulfill REST's hypermedia constraint. Unfortunately, this approach is not as straightforward as it might seem at the first sight. On one hand, it is not trivial to design a media type that is general enough for a broad range of applications, yet useful. On other hand it is difficult to find broad acceptance for a media type that is just usable in a very small application domain. Obviously, if the media type introduces a new serialization format, no existing client libraries can be used to parse representations. This then forces all clients to implement parsers specifically designed for this new media type. While such an approach might provide the best possible efficiency, it does not scale when the number of services or even just the number of entities using different media types in a single service increases. This is often criticized as media type explosion. In principle the same applies to media types that build on top of existing media types. A common pattern is to add, e.g., a `+xml` suffix to the media type identifier to describe that it is based on XML's syntax. Even though this practice has been standardized, at least for XML for more than a decade, some client libraries still do not understand it. Furthermore, this pattern cannot be used in content type negotiation. It can thus just be seen as a

hint to the developer to describe how to process such a representation. The practice of defining specialized media types for entities might also result in tighter coupled systems at the model layer as it makes it convenient for developers to reuse them as an application-level data model which is then shared among all system components (see [2] for an excellent discussion about the different coupling facets).

The second option is to extend an existing media type's application semantics with custom semantics. However, even though the Web is more than two decades old, surprisingly few media types exist that provide hypermedia support which goes beyond pure "GET-links" that do not have any support for semantic annotations. The reason for this is that it is hard to get consensus, and consequently support, for new processing mechanisms as provided by a media type. The Atom protocol suite ([3], [4]) defines media types which are among the most widely used for such an approach. Their main application domain is the syndication and manipulation of feeds of articles or blog posts, but since Atom's model has been specified general and extensible enough to support a broad range of applications, it is often used in other contexts as well; Microsoft's OData [5], Google's GData [6], and our previous work SAPS [7] are just a few examples. Such use beyond the original scenario is enabled by allowing the semantic description of links in feeds as well as in entries. Technically this works by setting the `rel` attribute of the `atom:link` element to a custom value instead of using one of the specified values. If an IRI (Internationalized Resource Identifier) is used to denote such a link's semantics, a developer can avoid name collisions and allow users to look its description up by simply dereferencing that IRI.

Finally, the third option is to use a completely application-agnostic media type that focuses on the presentation of raw data and thus allowing the serialization syntax to be separated from the application semantics. XML and JSON are the most popular media types of that class, but, since both have no inherent support for hyperlinks, it is impossible to use them to build truly RESTful services without documenting out-of-band how hyperlinks are represented. Traditionally, this was also true for RDF (in all its serialization formats) as the used IRIs were not meant to be dereferenced—similar to namespace declarations in XML documents. This shortcoming, that prevented networking effects to arise, was addressed by the introduction of the Linked Data principles [15] which demand dereferenceable HTTP IRIs that return useful information. Unsurprisingly, the Semantic Web, or Web of Data to avoid unnecessary misconceptions stemming from the historically AI-heavy term "Semantic Web", gained huge traction from this initiative.

Since the Linked Data principles align well with the REST architectural style (see [17] for an extensive analysis) it would just seem natural to combine their strengths. Nevertheless, the two remain largely separated in practice. Instead of providing Linked Data via RESTful Web services, current efforts deploy centralistic SPARQL endpoints or upload static dumps of the data. This rarely reflects the nature of the data, i.e., descriptions of interlinked resources. In our opinion this stems from the fact that average Web developers fear to use Semantic Web technologies; a phenomenon we denoted as *Semaphobia* [18]. Developers are often overwhelmed by the (perceived) complexity or intimidated by the AI-heavy history of the Semantic Web. The prevalent terminology, suffused with words such as *Ontology*, just seems to fuel

their misconceptions, while others are waiting for a killer application making it a classical chicken-and-egg problem. Another common perception is that the Semantic Web is a disruptive technology making it a show-stopper for enterprises needing to evolve their systems and build upon existing infrastructure investments. This relies on the fact that RDF is traditionally triple-centric whereas most developers program their systems in an entity centric (think object oriented) manner. Obviously, some developers are also just reluctant to use new technologies. Therefore, to mitigate these problems, these "new" technologies have to be introduced incrementally.

The XML serialization format of RDF, RDF/XML [19], is a great example of this. While it has been around for over a decade with very little uptake, RDFa [20] had gained a lot of momentum recently. It uses the same underlying data model as RDF, but instead of creating a new serialization format, it is used as a semantic layer on top of the ubiquitous (X)HTML. This gives Web developers a way to easily add semantic annotations to HTML documents. Since this makes it much easier for search engines to extract structured data from Web pages, they started to use it to improve their search algorithms and to present the results in a visually more appealing way which is a clear incentive for developers to annotate their Web pages. The recent introduction of *schema.org* can therefore be seen a major step forward for the Semantic Web as it allows a broad range of data, ranging from events and recipes to products and people, to be annotated with a shared vocabulary which is understood by all major search engines. Unfortunately, a similar approach for machines talking to each other via Web services is still missing. This is the gap that JSON-LD, the approach we are presenting in this paper, is trying to fill. It uses, similar to RDFa, an already successful syntax, i.e., JSON, and adds a semantic layer on top of it.

The remainder of this paper is organized as follows. In section 2 we give an overview of related work. Then, in section 3, we present JSON-LD and its data model. Section 4 shows how JSON-LD can be used to create evolvable RESTful services and finally, section 5 concludes the paper and gives an overview of future work.

2. RELATED WORK

According to ProgrammableWeb's statistics [8], three out of four APIs are RESTful and roughly half of them use JSON as the data format. It is interesting to observe that some of the most used APIs such as, e.g., Facebook's Graph API, Twitter's Streaming API, or Foursquare's API are now JSON-only. One of the reasons why JSON overtook XML as the primary data format in Web APIs might be the inherent impedance mismatch between XML and object oriented programming constructs (the so called O/X impedance mismatch) which often results in severe interoperability problems. The fundamental problem is that the XML Schema language has a number of type system constructs which simply do not exist in commonly used object oriented programming languages such as, e.g., Java. This leads to interoperability problems because each program stack has its own way of mapping the various XSD type system constructs to objects in the target platform's programming language and vice versa. Recent extensions for common languages such as C_o or LINQ (Language Integrated Query) for C# or E4X (ECMAScript for XML) for JavaScript ease handling of XML enormously but are not always available. In fact, XML was not even intended to be a generic

data-interchange format but designed as a lightweight subset of SGML to simplify electronic publishing in multiple media.

In contrast to XML, JSON, the JavaScript Object Notation, was specifically designed as a lightweight, language-independent data-interchange format that is easy to parse and generate. At the same time it is much less complex than XML. But, this simplicity comes at a price. JSON has neither native hypermedia support, nor does it support namespaces or semantic annotations. There have been various proposals to solve these shortcomings and all of them have in common that they specify a set of reserved keywords to express certain aspects such as, e.g., hyperlinks.

The most prominent examples of trying to add hypermedia support to JSON are probably JSON Schema [10] and its trimmed down counterpart JSON Reference [11]. Both define a special keyword `$ref` to denote a hyperlink. While, as the name suggests, JSON Schema puts that type information in a schema describing the document, JSON Reference uses the `$ref` keyword directly within the document. It can thus be seen as a static serialization of the same type but it lacks support for semantic annotation to describe its relation to the current document (this is possible with JSON Schema). Two related solutions for this issue are HAL and Collection+JSON, but in contrast to the previously mentioned approaches which augment JSON, they represent a new media type on their own. HAL [12] uses the `_links` keyword instead of `$ref` but instead of setting its value directly to the link's target, an object whose keys are the link relations and whose values are the link targets is used. It has also support to embed external resources within a representation. Often this is important as it allows applications to greatly decrease the number of required HTTP requests. Collection+JSON [13] is basically a JSON version of the Atom protocol suite to manage simple lists of entities. This media type not only specifies how links (which can be templated) are represented, but also how HTTP can be used to manipulate the various representations.

Similar to the above described proposals, but with a different goal in mind, various approaches have been presented to add semantic annotations or namespace support to JSON. These two aspects can be considered to be roughly the same as the idea of semantic annotations is to define the semantics of a concept in a special namespace to avoid collisions when the same terms are reused in different documents. The different proposals can be classified into two groups based on whether namespaces are supposed to be dereferenceable or not. In the first group, where namespaces are just used to avoid collisions and are thus not expected to be dereferenceable, often DNS-style names (`com.example.projects.namespacesInJSON`) are used [14]; the syntactic differences of the proposals are negligible. The second group of approaches assumes namespaces to be dereferenceable to be able to retrieve further information about them. As such they are mostly trying to create a JSON serialization format for RDF and thus offer much more functionality such as typing or internationalization support. As part of their effort to standardize a JSON serialization format for RDF, the RDF Working group already compared most of the existing approaches [16]; therefore we would like to refer the interested reader to that document for a detailed comparison. Summarized, it can be said that most of the approaches create a new media type with specific processing mechanisms and that the main differences between them are whether they are triple- or entity-centric and the degree by which they rely on microsyntaxes. This determines how familiar a repre-

sentation will look to a JSON developer; an important aspect for the acceptance of such a format. Unfortunately, most of the approaches fall short in this respect.

To overcome these and other shortcomings we introduced SEREDASj [18] in previous work. It is a description language for *SEmantic REstful DAta Services* and focuses on the description of JSON resource representations and their interconnections. It also allows these representations to be transformed to RDF. From working with different developers, we found that the separation of the description layer and the data is suboptimal as it effectively creates a second layer of interconnected resources on top of the data in JSON documents which increases the cognitive load on the developer. It was difficult for developers to understand documents without at the same time looking at the SEREDASj description document. Also, the syntax, which followed JSON Schema's approach, was often considered to be too verbose.

After having worked for quite some time on SEREDASj and having built several promising prototypes, we discovered the JSON-LD project. As it had almost the same goals we were trying to achieve and followed a very similar approach, we joined the, back then, still small community. After understanding its goals and mindset we decided to discontinue our work on SEREDASj in favor of JSON-LD as we believed we could achieve more in less time by joining forces, instead of working on similar, yet different approaches competing to solve similar problems.

In spite of being a comparatively young project, JSON-LD has already had a turbulent history. According to Manu Sporny [21], the work was started internally at Digital Bazaar in March 2010. This was shortly before at the W3C RDF Next Steps Workshop the desire of the community for a JSON-based RDF format was found [22]. Consequently the RDF Working Group at W3C started working on a JSON-based RDF serialization on two fronts. It decided to quickly standardize Talis' triple-centric RDF/JSON [9] for RDF experts needing a JSON-based serialization and to incubate on JSON-LD for average Web developers without RDF background. Unsurprisingly, this strategy soon ended in a general confusion as to the exact target group it is attempting to address and what the outcome should be [21]; the group did not share a common vision. Finally, in August 2011 Thomas Steiner, the appointed co-editor, pulled the "emergency brake" [23] and the work in the RDF Working Group was stopped. Despite these happenings, we continued to work as part of the JSON-LD community to improve the syntax and created a W3C community group [24] instead of waiting for the RDF Working Group to decide on how to proceed.

3. JSON-LD

JSON-LD is an attempt to create a simple method to not only express Linked Data in JSON but also to add semantics to existing JSON documents. It has been designed to be as simple as possible, very terse, and human readable. Furthermore, it was a goal to require as little effort as possible from developers to transform their plain old JSON to semantically rich JSON-LD. Consequently, an entity-centric approach was followed whereas traditional Semantic Web technologies are usually triple-centric. While the initial versions of JSON-LD [25] looked like a more or less direct translation of Turtle to JSON, the syntax was dramatically changed in the latest versions and allows now data to be serialized in a way that is often indistinguishable from normal JSON [26]. This is remarkable since JSON is used to serialize a directed

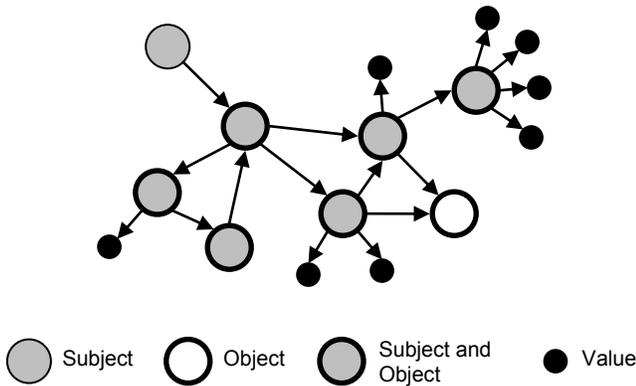


Figure 1. A Linked Data graph

graph that potentially contains cycles while its native data model is a tree.

Figure 1 shows JSON-LD's data model, a Linked Data graph. Nodes in the graph are called *subjects* or *objects* and edges are called *properties* (predicates in RDF). A subject is a node with at least one outgoing edge whereas an object is a node with at least one incoming edge. This implies that a node can be a subject and an object at the same time. To be unambiguously identifiable and referenceable, a subject should be labeled with an IRI. This is not a strict requirement though; JSON-LD also supports unlabeled nodes. Even though such nodes do not fulfill the requirements of Linked Data, they are supported as they allow certain use cases which require just locally referenceable data. The same applies to properties (edges): if they are labeled with an IRI they are referenceable from other documents and thus Linked Data; otherwise they are just traditional JSON properties that only have a meaning in the specific document they are used. The situation is slightly different for objects. If an object is labeled by an IRI, it is called an object; if it is labeled by something that is not an IRI, e.g. a number, it is denoted as a *value*, i.e., a literal in RDF.

Given the reluctance of average Web developers to use semantic technologies, huge efforts have been put into JSON-LD so that developers do not have to be knowledgeable about other semantic Web technologies. All a developer needs to know is JSON and two keywords (@context and @id) to use JSON-LD's basic functionality. Since JSON-LD is 100% compatible with plain old JSON, developers can continue to use their existing tools and libraries. This is especially important for enterprises as it allows them to add meaning to their JSON documents in a way that is not disruptive to their operations and is transparent to their current customers. At the same time JSON-LD is expressive enough to support all major RDF concepts.

The basic idea of JSON-LD is to create a description of the data in the form of a so called context. It links, similarly to SEREDASj description documents, objects and their properties in a JSON document to concepts in an ontology. Furthermore, it allows values to be type-coerced and language tagged. A context can either be directly embedded in a JSON-LD document or put into a separate file and referenced from different documents. This, and the fact that plain old JSON documents can reference a context

```

1  {
2    "@context": {
3      "foaf": "http://xmlns.com/foaf/0.1/",
4      "title": "foaf:title",
5      "name": "foaf:name",
6      "homepage": {
7        "@id": "foaf:workplaceHomepage",
8        "@type": "@id"
9      }
10   },
11   "@id": "http://me.markus-lanthaler.com",
12   "@type": "foaf:Person",
13   "title": [
14     {"@value": "Dipl.Ing.", "@language": "de"},
15     {"@value": "MSc", "@language": "en"}
16   ],
17   "name": "Markus Lanthaler",
18   "homepage": "http://www.tugraz.at/"
19 }

```

Listing 1. An exemplary JSON-LD document

via an HTTP link header, provides a smooth upgrade path for existing infrastructure as it allows most of the functionality without having to change the contents of an existing document at all.

Listing 1 contains a simple JSON-LD document that describes a person based on the FOAF vocabulary [27]. First, a prefix for the FOAF vocabulary is defined (line 3) in the embedded context (lines 2-10) to abbreviate the long concept IRIs. Then, in lines 4-7, the three JSON properties `title`, `name`, and `homepage` are mapped to concepts in the FOAF vocabulary. Additionally, the value of the `homepage` property is specified to be of the type `@id`, i.e., it is specified to be an IRI (line 8). Finally, in line 11, the person described in the document is unambiguously identified by an IRI to make it possible to reference this person in other documents. The same mechanism allows JSON-LD documents containing more information to be transcluded which enables clients to discover new data by simply following those links; this principle is known as *Follow Your Nose* [28]. By having all data semantically annotated as in the example, a machine client can be programmed to “understand” that the document is about a person (line 12) and to figure out which properties specify the person's title (and in which language it is; lines 13-16), name (line 17) and the homepage of the organization it works for (line 18). A JSON-LD publisher is free to choose between using terms that are mapped to concept IRIs in a vocabulary via a context as in the example and using these IRIs directly in the document. Since this flexibility results in variability that makes it more difficult to

```

1  {
2    "@id": "http://me.markus-lanthaler.com",
3    "@type": "http://xmlns.com/foaf/0.1/Person",
4    "http://xmlns.com/foaf/0.1/title": [
5      {"@value": "Dipl.Ing.", "@language": "de"},
6      {"@value": "MSc", "@language": "en"}
7    ],
8    "http://xmlns.com/foaf/0.1/name":
9      "Markus Lanthaler",
10   "http://xmlns.com/foaf/0.1/
11     workplaceHomepage":
12     { "@id": "http://www.tugraz.at/" }

```

Listing 2. The expanded form of the document in Listing 1

process the data, JSON-LD specifies two special document forms: expanded and normalized.

The expanded form (Listing 2) is a JSON-LD document where all terms and prefixes have been expanded into full IRIs and all type and language coercions are defined inline so that the context can be eliminated from the document without losing any information. It can thus be seen as an explicit version of the document. Normalization goes a step further by also labeling all blank nodes in order to create an output graph with a deterministic structure. This is sometimes necessary since a directed graph, the underlying data model of JSON-LD, can express the same information in a variety of different ways; digitally signing a graph, e.g., would otherwise be impossible. Since there are no unlabeled nodes, the normalized form of the document in Listing 1 would look almost the same as the expanded version in Listing 2. The only difference being that the properties in the normalized form would be in a different order as they are sorted according to a deterministic algorithm. For greater flexibility, JSON-LD also defines a *framing* API which allows a developer to transform a document into a form that is convenient to process for a specific application. The developer defines a frame, i.e., a template, which is then used to restructure the data contained in an arbitrary JSON-LD document into the desired form. This allows the developer to subsequently work with the framed document just as with any other JSON document which means that usually all existing JSON tools and workflows can be retained.

Conversion of a JSON-LD document, especially one in the expanded or normalized form, to RDF triples is straightforward. A subject, which could also be used as an object in another triple, is defined by `@id`. All other JSON-LD properties are converted to predicates. Finally, literal values are either taken directly from a property's value, or created by taking the value of `@value` and adding language (`@language`) and/or type information (`@type`). The example in Listing 1 could thus be converted to, e.g., a Turtle document as shown in Listing 3.

4. EVOLVABLE RESTFUL SERVICES WITH JSON-LD

As mentioned in the introduction, the Hypermedia as the Engine of Application State (HATEOAS) constraint is one of the least understood constraints, and thus seldom implemented correctly. Annoyed by the fact that a lot of services claim to be RESTful regardless of violating the hypermedia constraint, Fielding [29] made it very clear that hypermedia is a fundamental requirement but since the term REST is so widely misused, there are efforts in the community to look for an alternative term, such as *Hypermedia API*, to denote truly RESTful services.

A lot of systems, regardless of claiming to be RESTful, rely heavily on implicit state control-flow which is characteristic of the RPC-style. The allowed messages and how they have to be interpreted depends on previously exchanged messages and thus in which implicit state the system is in. Third parties or intermediaries trying to interpret the conversation need the full state transition table and the initial state to understand the communication which is often not available or not practical. This also makes it difficult or virtually impossible to recover from partial failures in such distributed systems.

To solve these issues and assure evolvability, the use of hypermedia is a core tenet of the REST architectural style. It refers to

```
1 @prefix rdf: <http://www.w3.org/1999/02/22-  
   rdf-syntax-ns#> .  
2 @prefix foaf: <http://xmlns.com/foaf/0.1/> .  
3  
4 <http://me.markus-lanthaler.com>  
5   rdf:type foaf:Person ;  
6   foaf:title  
7     "Dipl.Ing."@de ,  
8     "MSc"@en ;  
9   foaf:name "Markus Lanthaler" ;  
10  foaf:workplaceHomepage  
11    <http://www.tugraz.at/> .
```

Listing 3. Triples extracted from Listing 1 converted to Turtle

the use of hypermedia controls in resource representations as a way of navigating the state machine of an application. "A REST API should be entered with no prior knowledge beyond the initial URI (bookmark) and set of standardized media types. [...] From that point on, all application state transitions must be driven by client selection of server-provided choices that are present in the received representations or implied by the user's manipulation of those representations." [29] While the human Web is unquestionably based on this type of interaction and state control-flow where very little is known a priori, machine-to-machine communication is often based on static contracts and out-of-band knowledge resulting in tight coupling. Such approaches might work in the short term but are condemned to break in the long term since assumptions about server resources will break eventually as resources evolve over time. Parastatidis et al. [30] define the set of legal interactions necessary to achieve a specific, application-dependent goal as the *domain application protocol* of a service. The protocol defines the interactions rules between the different participants. Consequently, the application state is a snapshot of the system at an instant in time. This coincides with Fielding's definition [1] of application state which defines it as the "pending requests, the topology of connected components (some of which may be filtering buffered data), the active requests on those connectors, the data flow of representations in response to those requests, and the processing of those representations as they are received by the user agent." Accordingly, the overall system state consists of the application state and the server state. By using the notion of a domain application protocol the phrase "hypermedia as the engine of application state" can now be explained as the use of hypermedia controls to advertise valid state transitions at runtime instead of agreeing on static contracts at design time. Changes in the domain application protocol can thus be dynamically communicated to clients. This brings some of the human Web's adaptivity to the Web of machines and allows the building of loosely coupled and evolvable systems. Rather than requiring an understanding of a specific IRI structure, clients only need to understand the semantics or business context in which a link appears [30].

The creation of a truly RESTful service can be dramatically simplified by the use of a well-defined, generic media type with inherent hypermedia support. Developers can then fully concentrate on defining the domain application protocol instead of having to design new media types which often are specializations of existing syntaxes. This can be compared to the omnipresence of HTML and browsers in the human Web where such standardization is already well advanced. JSON-LD was specifically designed for such purposes. It has built-in support for domain

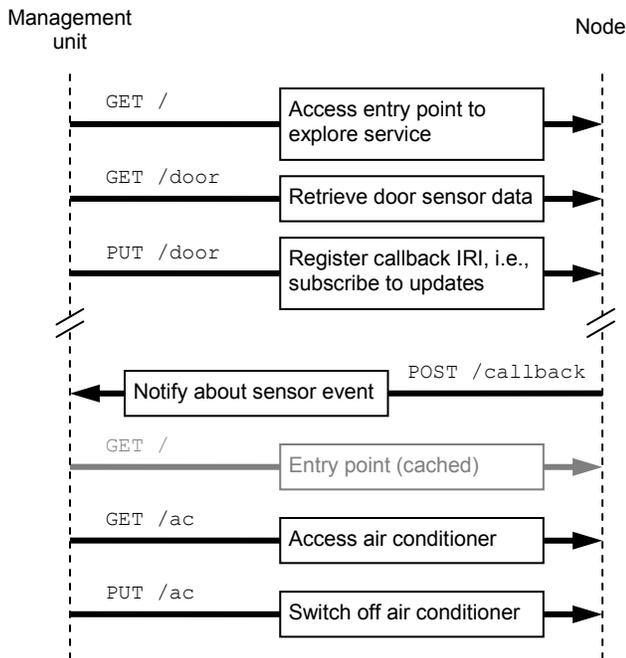


Figure 2. Interaction between management unit and node

semantics yet its syntax is completely independent thereof. A developer can thus use JSON-LD as a generic media type for a broad range of applications. If Semantic Web technologies, i.e., ontologies, are used to express the domain application protocol, it is possible to leverage the rich underlying model and existing tools for tasks such as data validation or even to make implicit knowledge explicit. This standardization allows the development process to be further streamlined as it enables the creation of reusable client libraries.

The core process to create a RESTful service would then be to define and describe the application semantics. Every concept gets assigned an IRI where the definition of that concept can be looked

```

1  {
2    "@context": "http://api.com/dap.jsonld",
3    "@id": "/",
4    "location": "http://api.com/room/48",
5    "sensors": [
6      {
7        "@id": "/temperature",
8        "@type": "dap:TemperatureSensor"
9      },
10     {
11       "@id": "/door",
12       "@type": "dap:DoorSensor"
13     }
14   ],
15   "actuators": [
16     {
17       "@id": "/ac",
18       "@type": "dap:AirConditioner"
19     }
20   ]
21 }

```

Listing 4. The node's "homepage"

up. Please note that this process is almost the same as the one chosen by, e.g., Atom. The only difference is that instead of requiring out-of-band knowledge to look up the link relation's semantics, i.e., the knowledge about the IANA registry [31], the semantics can be directly accessed by simply dereferencing the IRI. A developer could then go a step further and describe the semantics in a machine processable way. It would, e.g., be possible to describe the allowed value range for a property by using OWL [32], a standardized ontology. This would also make it possible to automatically check the consistency of a domain application protocol and to generate human readable documentation from it. Clearly, this goes far beyond what is achievable with the traditional definition of media types as the descriptions can be reasoned with by computer programs.

As a real-world proof of concept for the principles described in this paper, we have designed and implemented a Web of Things consisting of a central management unit and nodes equipped with sensors and actuators. It clearly illustrates the way in which such a distributed system can be modeled and implemented by using JSON-LD and Semantic Web technologies. For example, assume we have a node in each room with a door sensor and an air conditioner connected to it. An application running on the management unit could then be programmed to turn the air conditioner off when the door is open.

Figure 2 illustrates how the management unit and the nodes would interact. First, the management unit would establish a connection to the node's entry point. The response could contain a representation as shown in Listing 4 (context defined externally for the sake of brevity). In the next step, the application on the management unit would choose which link to follow. Instead of having to rely on type information (line 12), it would also be possible to create a more specific property than just `sensors`; this is a decision the domain application protocol designer has to take. Given that the application is interested in the door sensor, it would dereference the link in line 11 and the node would return a representation similar to the one shown in Listing 5. Finally, the management unit would add its own callback IRI to the list of subscribers (lines 7-10). As soon as the reading of the sensor changes, the node would notify all subscribers, which at this point includes the management unit. This would then trigger a process similar to the one just described, but this time, the management unit would look up the air conditioner and update its representation to turn it off.

While this seems overly chatty, it is worth noting that responses can be cached and that the developer is free to optimize representations to reduce the number of required roundtrips. For example, Listing 4 can be cached for long periods as it is not expected to

```

1  {
2    "@context": "http://api.com/dap.jsonld",
3    "@id": "/door",
4    "@type": "dap:DoorSensor",
5    "node": "/",
6    "reading": "dap:closed",
7    "subscribers": [
8      "http://log.example.com/",
9      "http://accesscontrol.example.com/"
10   ]
11 }

```

Listing 5. The door sensor data

```

1  {
2    "@context": "http://api.com/dap.jsonld",
3    "@id": "/",
4    "location": "http://api.com/room/48",
5    "sensors": [
6      {
7        "@id": "/temperature",
8        "@type": "dap:TemperatureSensor",
9        "reading": 32.5,
10       "subscribers": []
11      },
12      {
13        "@id": "/door",
14        "@type": "dap:DoorSensor",
15        "reading": "dap:closed",
16        "subscribers": [
17          "http://log.example.com/",
18          "http://accesscontrol.example.com/"
19        ]
20      }
21    ],
22    "actuators": [
23      {
24        "@id": "/ac",
25        "@type": "dap:AirConditioner",
26        "status": "dap:off"
27      }
28    ]
29  }

```

Listing 6. Optimization of Listing 4 with embedded resources aiming to reduce the number of required HTTP requests

change often. Alternatively, it could be heavily optimized to a form similar to the one shown in Listing 6 where all the required data is directly embedded into the “homepage” instead of being transcluded; obviously this renders caching almost useless as the representation would change continuously. Such an optimization can be compared to the use of image sprites or data URIs in HTML pages to reduce the number of required HTTP requests. Whilst it is completely acceptable for a service to make such structural changes, semantic changes to the domain application protocol have to be carefully evaluated as they might change the contract and break existing consumers relying on them.

If OWL’s expressivity is used to describe the relationships between different resources, much smarter clients can be built. It is, e.g., possible to create a generic crawler which indexes all available sensors; by leveraging OWL’s class hierarchy this would also work for sensor types that are still unknown at design time. Similarly, a logging service can be created that stores reported sensor readings. It is enough if such a service understands which property contains the reading; it is not necessary that it understands the reading itself. Such flexibility greatly contributes to the extensibility and evolvability of a system.

5. CONCLUSIONS AND FUTURE WORK

In this paper we presented JSON-LD, a community effort to standardize Linked Data in JSON. In contrast to previous approaches, great efforts have been made to keep the approach as simple as possible and to create a syntax that results in serializations that follow the structure of how JSON is typically used by Web developers. The syntax is 100% compatible with JSON and flexible enough to provide a smooth upgrade path for existing infrastructure. This means that developers have to change neither their

workflows nor their existing tools and programming libraries. We hope this lowers the entry barrier to publish Linked Data in the form of RESTful services and results in a broader adoption of the underlying ideas and principles. The additional design costs are significantly outweighed by the achievable benefits in terms of loose coupling, evolvability, scalability, self-descriptiveness, and maintainability. Our experiments in the context of a large-scale Web of Things project have delivered proofs that the presented approach is practical. In fact, JSON-LD could be a first step toward standardizing semantic RESTful Web services and could form the basis for various efforts that previously could not seem to find any common ground. By working with a community with different backgrounds, we hope to be able to create a balanced solution that builds on the efforts put into RDF in the last decade while having the potential to positively affect the Web as a whole. Given that JSON-LD is still a work in progress, we would like to take the chance to invite the interested reader to join the community at <http://www.json-ld.org>.

As we have shown, JSON-LD itself is not a complete technology stack – it needs ontologies to express domain semantics. In future work we would like to investigate how a lightweight ontology to support a wide range of application domains could be modeled. Furthermore, we would like to explore various ideas to create smarter service clients.

6. REFERENCES

- [1] R.T. Fielding, “Architectural styles and the design of network-based software architectures,” PhD dissertation, Department of Information and Computer Science, University of California, Irvine, 2000.
- [2] C. Pautasso and E. Wilde, “Why is the Web Loosely Coupled? A Multi-Faceted Metric for Service Design,” in *Proceedings of the 18th International Conference on World Wide Web (WWW)*, 2009, pp. 911-920.
- [3] The Atom Syndication Format. [Online]. Available: <http://tools.ietf.org/html/rfc4287>.
- [4] The Atom Publishing Protocol. [Online]. Available: <http://tools.ietf.org/html/rfc5023>.
- [5] Open Data Protocol. [Online]. Available: <http://www.odata.org/>.
- [6] Google Data Protocol. [Online]. Available: <http://code.google.com/apis/gdata/>.
- [7] M. Lanthaler and C. Gütl, “SAPS: Semantic AtomPub-based Services,” in *Proceedings of the 11th IEEE/IPSJ International Symposium on Applications and the Internet (SAINT)*, 2011, pp. 382-387.
- [8] T. Vitvar and J. Musser, “ProgrammableWeb.com: Statistics, trends, and best practices,” *Keynote of the Web APIs and Service Mashups Workshop at the European Conference on Web Services*, 2010.
- [9] RDF/JSON, Talis Systems Ltd., 2011. [Online]. Available: <http://docs.api.talis.com/platform-api/output-types/rdf-json>. [Accessed: 15-Jan-2012].
- [10] K. Zyp and G. Court, “JSON Schema”, 2010. [Online]. Available: <http://tools.ietf.org/html/draft-zyp-json-schema-03>. [Accessed: 18-Jan-2011].

- [11] P. Bryan and K. Zyp, "JSON Reference", 2011. [Online]. Available: <http://tools.ietf.org/html/draft-pbryan-zyp-json-ref-01>. [Accessed: 20-Dec-2011].
- [12] M. Kelly, "HAL - Hypertext Application Language", 2011. [Online]. Available: http://stateless.co/hal_specification.html. [Accessed: 20-Dec-2011].
- [13] M. Amundsen, "Collection+JSON - Document Format", 2011. [Online]. Available: <http://amundsen.com/media-types/collection/format/>. [Accessed: 24-Oct-2011].
- [14] Y. Goland, "Adding Namespaces to JSON", 2006. [Online]. Available: <http://www.goland.org/jsonnamespace/>. [Accessed: 05-Jan-2012].
- [15] T. Berners-Lee, "Linked Data," *Design Issues for the World Wide Web*, 2006. [Online]. Available: <http://www.w3.org/DesignIssues/LinkedData.html>. [Accessed: 06-Jun-2010].
- [16] RDF Working Group, "JSON Serializations by Example", 2011. [Online]. Available: <http://www.w3.org/2011/rdfwg/wiki/JSON-Serialization-Examples>. [Accessed: 28-Jul-2011].
- [17] K. R. Page, D. C. De Roure, and K. Martinez, "REST and Linked Data: a match made for domain driven development?" in *Proceedings of the 2nd International Workshop on RESTful Design (WS-REST)*, 2011, pp. 22-25.
- [18] M. Lanthaler and C. Gütl, "A Semantic Description Language for RESTful Data Services to Combat Semaphobia," in *Proceedings of the 2011 5th IEEE International Conference on Digital Ecosystems and Technologies (DEST)*, 2011, pp. 47-53.
- [19] RDF/XML Syntax Specification (Revised). [Online]. Available: <http://www.w3.org/TR/2004/REC-rdf-syntax-grammar-20040210/>.
- [20] RDFa in XHTML: Syntax and Processing. [Online]. Available: <http://www.w3.org/TR/2008/REC-rdfa-syntax-20081014>.
- [21] M. Sporny, "Linked JSON: RDF for the Masses," *The Beautiful, Tormented Machine*, 2011. [Online]. Available: <http://manu.sporny.org/2011/linked-json/>. [Accessed: 28-Apr-2011].
- [22] I. Herman, "W3C Workshop — RDF Next Steps Workshop Report," 2010. [Online]. Available: <http://www.w3.org/2009/12/rdf-ws/Report.html>. [Accessed: 05-Aug-2010].
- [23] T. Steiner, "JSON Emergency Brake," *RDF Working Group mailing list*, 2011. [Online]. Available: <http://lists.w3.org/Archives/Public/public-rdf-wg/2011Aug/0131.html>. [Accessed: 23-Aug-2011].
- [24] JSON for Linking Data Community Group, *W3C Community and Business Groups*. [Online]. Available: <http://www.w3.org/community/json-ld/>.
- [25] JSON-LD - Linked Data Expression in JSON, Unofficial Draft 30 May 2010. [Online]. Available: <http://json-ld.org/spec/ED/json-ld-syntax/20100529/>.
- [26] JSON-LD Syntax 1.0, Unofficial Draft 18 March 2012. [Online]. Available: <http://json-ld.org/spec/ED/json-ld-syntax/20120318/>.
- [27] D. Brickley and L. Miller, FOAF Vocabulary Specification 0.98. 2010. [Online]. Available: <http://xmlns.com/foaf/spec/>. [Accessed: 17-Jan-2011].
- [28] L. Dodds and I. Davis, *Linked Data Patterns - A pattern catalogue for modelling, publishing, and consuming Linked Data*, 2011, pp. 44-55. [Online]. Available: <http://patterns.dataincubator.org/book/linked-data-patterns.pdf>. [Accessed: 07-Oct-2011].
- [29] R. T. Fielding, "REST APIs must be hypertext-driven," *Untangled musings of Roy T. Fielding*, 2008. [Online]. Available: <http://roy.gbiv.com/untangled/2008/rest-apis-must-be-hypertext-driven>. [Accessed: 02-Jun-2010].
- [30] S. Parastatidis, J. Webber, G. Silveira, and I. S. Robinson, "The Role of Hypermedia in Distributed System Development," in *Proceedings of the 1st International Workshop on RESTful Design (WS-REST)*, 2010, pp. 16-22.
- [31] Link Relations, IANA. [Online]. Available: <http://www.iana.org/assignments/link-relations/link-relations.xml>.
- [32] OWL 2 Web Ontology Language. [Online]. Available: <http://www.w3.org/TR/2009/REC-owl2-primer-20091027/>.