

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: August 31, 2015

L. Daigle
TCE
March 2015

Internet Application Identifier Architecture
draft-daigle-appidarch-00.txt

Abstract

This document outlines a general architecture for Internet applications, through the perspective of applications identifiers. It provides a survey of past approaches, drawing out common elements and highlighting common traps and roadblocks.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on August 31, 2015.

Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the [Trust Legal Provisions](#) and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
2.	Basic components of Application Identifier Architecture . . .	2
3.	Application Identifier Architectures in More Detail	2
3.1.	System	2

3.2.	Identifiers	3
3.3.	Identified	3
4.	Survey of existing work	5
4.1.	Domain Name System	5
4.2.	World Wide Web	6
4.3.	Classic URIs	8
4.4.	IP addresses	10
5.	Common design choices and challenges	10
5.1.	Identifiers	10
5.2.	Identified	10
6.	Issues in (mis)using identifiers	11
7.	IANA Considerations	11
8.	Security Considerations	11
9.	References	11
	Author's Address	11

[1.](#) Introduction

This document posits a high level architecture of Internet application identifier systems, as well as a survey of IETF efforts dealing in standardization of Internet applications and services using those identifier systems.

Status of this revision: this is a very drafty -00 document. The hope and expectation is that it is enough to stimulate some thought and discussion to flesh out future documents.

[2.](#) Basic components of Application Identifier Architecture

There are 3 basic components of an Application Identifier Architecture:

- o System
- o Identifiers
- o Identified

The System is the context in which the identifiers are created, used and in which they are intended to make sense. This is usually transparent, except when identifiers are used outside of this context, causing greater or lesser problems over time. This is discussed in more detail, below.

Identifiers are typically strings of bits or characters. They may have multiple representations.

The concept of what is being Identified is also dependent on the System -- whether it's Internet hosts, services, documents, parts of documents, people or other actors from the physical world, their

representatives in the Internet, etc.

3. Application Identifier Architectures in More Detail

3.1. System

Daigle

Expires August 31, 2015

[Page 2]

As noted above, the System is the context in which the identifiers make sense. In the Domain Name System, for example, the system initially consisted of the set of hosts attached to the Internet, and it has generalized to the set of addressable Internet services. These are organized into ?domains?, which are operated under the control of a single entity, and individual domains are completely independent of each other.

3.2. Identifiers

Identifiers may identify a thing that exists ? content, service, location ? or is posited to exist. Typical actions on identifiers include:

- o "Minting" -- creation of an identifier, usually including association with the identified thing
- o Transformation -- changing the bits (characters) of the identifier by some set of rules and/or to conform to some structure; relative or absolute
- o Comparison -- of identifiers. Are 2 different identifiers the same? Identifying the same thing? Expressing relationship between things?
- o Resolution -- using the identifier to access what it identifies
- o Validation -- confirmation that the identifier conforms to the system?s rules (syntax)
- o Status check -- has the identifier been minted? Is it active within the system?
- o Authentication -- confirmation that the identifier association is valid (as minted)
- o Lookup -- some systems support look up ? finding identifier entries based on partial fragments, typically leading characters (bits)
- o Search -- some systems support searches for identifiers based on fragments of the related data
- o Subscribe -- subscribing to an identifier allows you to get periodic updates as to state of the identifier/identified.

3.3. Identified

Identifiers can be associated with just about any level of concept,

construct, network or software element, or entity in the physical world. The range of possible identified things is generally scoped by the System.

From the perspective of application architectures, there are 4 levels of things that may be identified, and may have individual identifiers:

- o Entity/resource -- the thing itself. For example, the published work "Moby Dick"
- o Instance -- a specific copy of the thing, e.g., a copy of "Moby Dick".
- o Properties -- the characteristics of the thing. The set of properties discussed is generally constrained by the System.
- o Relationship -- an identifier may identify something as "part-of" a larger entity.

There are actions that may be performed on the things identified:

- o Assign properties -- associate values with properties of identified item
- o Get properties
- o Intrinsic -- E.g., format, number of words
- o Applied -- director's cut
- o Publish -- put copy somewhere
- o First instance
- o Cache/replica
- o Get (a copy)
- o Any copy
- o Specific service
- o Closest
- o Cheapest
- o Authenticate -- confirm (cryptographically) that the resource is genuinely the one expected / related to identifier
- o Comparison
- o Equivalence

- o Send -- a reflection of "get"?

- o Subscribe

Daigle

Expires August 31, 2015

[Page 4]

- o Search

4. Survey of existing work

4.1. Domain Name System

The Domain Name System (DNS) was designed to provide identifiers to allow storage and retrieval of (sets of) properties associated with Internet hosts and services ? real and virtual.

DNS identifiers are hierarchical, dot-separated labels, typically expressed as characters. Host names are a subset of domain name identifiers, with some restrictions on the permissible characters.

- o "Minting" -- the authority for a domain can create labels within that domain. So-called "synthetic" domain labels are created dynamically.
- o Transformation -- relative domain names can be understood within the context of a ?search domain?
- o Comparison -- domain names are matched on an octet-by-octet basis (IDNs are not considered here)
- o Resolution -- DNS resolution means "DNS lookup" -- using the DNS infrastructure to retrieve resource records associated with the domain name. Resolution can be tailored to retrieve particular types of resource records (e.g., A records, or AAAA records, or MX records)
- o Validation -- any string that conforms to DNS syntax may be considered valid.
- o Status check -- DNS does not distinguish between "inactive" and "not minted". That is, either every label in the hierarchical domain name is accessible in an authoritative domain server (in which case the domain name is "minted" and "active") or the DNS will return the value that it does not exist. (Not true in DNSSEC?)
- o Authentication -- domain names are not authenticated (see below for DNSSEC).
- o Lookup -- DNS resolution is lookup of domain names
- o Search -- DNS does not support search (wildcards?)
- o Subscribe -- N/A

DNS identifies resource records. The resource records are themselves descriptions of Internet hosts, services, and other

information stored in the DNS, but fundamentally the DNS identifier is for resource records.

- o Entity/resource -- a set of resource records
- o Instance -- copies of resource records may be stored in caching servers; there are no special identifiers to distinguish primary or cached results
- o Properties -- N/A
- o Relationship -- N/A

There are actions that may be performed on the things identified:

- o Assign properties -- resource records have time to live (TTL) and serial numbers included
- o Get properties -- parsed as part of the response from the server.
- o Publish -- publishing a DNS resource record amounts to updating a DNS zone file with the record.
- o Get (a copy) -- resolve the domain name identifier
- o Authenticate -- DNSSEC is used to authenticate that the resource records/response received for domain name resolution are as they were published.
- o Comparison -- of RRs?
- o Send -- N/A
- o Search -- N/A

[4.2.](#) World Wide Web

The World Wide Web (WWW) is largely defined by the HTTP protocol. "Pages" defined in HTML are the primary design target, though these days much content of varying formats is delivered via the HTTP protocol. For the sake of discussion, we'll say that WWW identifiers are HTTP scheme URIs.

- o "Minting" -- typically, HTTP URIs are not composed consciously, so much as assembled practically with components of the domain name hosting the web server and some path components based on how the website is laid out hierarchically (which may or may not relate to

Daigle

Expires August 31, 2015

[Page 6]

an underlying file structure)

- o Transformation -- HTTP URIs may be relative (to current page in the "hierarchy", current domain authority etc)
- o Comparison -- HTTP URIs are not inherently comparable except by characterwise comparison or determining relative relationships
- o Resolution -- HTTP URIs are resolved by parsing the authority component from the URI, connecting to the server, and requesting the resource associated with the path part of the URI.
- o Validation -- any string that conforms to HTTP URI syntax may be considered valid.
- o Status check -- HTTP does not distinguish between "inactive" and "not minted". That is, either the HTTP server is available and the resource is found on it (in which case the URI is "minted" and "active") or the server (or path) are not found.
- o Authentication -- HTTP URIs are not authenticated (see below for authentication of servers).
- o Lookup -- N/A
- o Search -- HTTP does not support search (within server?)
- o Subscribe -- N/A

HTTP URIs identify "web pages" or "resources".

- o Entity/resource -- web content (page)
- o Instance -- copies of pages may be stored in caching servers; there are no special identifiers to distinguish primary or cached results
- o Properties -- properties may be embedded within the HTML document, but there are no special identifiers to query/retrieve properties; as part of the HTTP protocol, capabilities may be negotiated
- o Relationship -- relative URIs (?)

There are actions that may be performed on the things identified:

- o Assign properties -- the web server may assign properties to a web page.
- o Get properties -- parsed as part of the response from the server.

- o Publish -- publishing a web page is done on the backend, out of band of the WWW system
- o Get (a copy) -- resolve the URI

Daigle

Expires August 31, 2015

[Page 7]

- o Authenticate -- certs are used, within HTTP, to authenticate the server is empowered to operate for a given domain name. Individual pages are not authenticated.
- o Comparison -- many web pages look alike -- but there is no inherent way to claim two web pages (documents) are "the same".
- o Send -- N/A
- o Search -- within the WWW there is no support for search (all search is achieved as an external system).

4.3. Classic URIs

The advent of the WWW heralded a burst of development of standards for applications and content on the Internet. Much work was done to elaborate a general system of identifiers, supporting a broad range of application needs -- Uniform Resource Identifiers in the large, encompassing Locators (identifiers of Internet "location"), Names (persistent, location-independent identifiers for resources), Characteristics (metadata about resources), Agents (for composing actions).

In the large, the classic perspective on URIs was that they would identify all resources (documents, services, media, components, classes, parameters etc) that would be referenced within Internet protocols.

- o "Minting" -- dependent on the URI scheme. The HTTP URI scheme is outlined above as a dynamic URI creation example. Some (namespaces of) URNs require more explicit minting of identifiers to be used in URNs (e.g., ISBN URNs).
- o Transformation -- dependent on the URI scheme. URNs were intended to be (authoritatively) transformed into URLs identifying the location of the desired resource at a given point in time.
- o Comparison -- Scheme-dependent -- there is no URI-wide support for comparing URIs (except byte-wise equality).
- o Resolution -- Scheme-dependent. Some URI schemes do not have Internet-based resolution capabilities.
- o Validation -- any string that conforms to URI syntax may be considered valid. Individual schemes may include validation services (e.g., out of band lookup services, built-in checksums, etc).
- o Status check -- Generally, URIs do not distinguish between

"inactive" and "not minted". That is, successful resolution

implies minted, unsuccessful resolution is ambiguous. Individual schemes may provide more refined methods of confirmation of status (SIP URIs?)

- o Authentication -- There is no URI-wide support for URI authentication.
- o Lookup -- N/A
- o Search -- URIs are not inherently searchable
- o Subscribe -- N/A

Classically, URIs identify Internet "resources". However, URIs have been found in contexts that are disjoint from the Internet (e.g., XML component identification). "Resource" is deliberately general -- could be documents, services, etc. Each URI scheme refines its scope of intended resources.

Entity/resource -- typically Internet content or service, but see comment above.

Instance -- Most URI schemes do not support identification of instances. However, URNs are intended to identify locations of multiple instances of a given resource.

Properties -- URIs may identify URCs (Uniform Resource Characteristics) an articulation of the properties of a given resource. (URCs were never finally standardized).

Relationship -- relative URIs (?)

There are actions that may be performed on the things identified:

- o Assign properties -- publish a URC
- o Get properties -- retrieve a URC associated with the URI.
- o Publish -- scheme-dependent
- o Get (a copy) -- resolve the URI
- o Authenticate -- certs are used, within some URI schemes, to authenticate the server is empowered to operate for a given domain name. Individual resources are not authenticated. (Unless you have a URC with a checksum?)
- o Comparison -- many resources look alike -- but there is no inherent way to claim two resources are "the same".

- o Send -- N/A
- o Search -- within the URI space there is no support for search (all search is achieved as an external system).

4.4. IP addresses

To be added... the interesting thing about IP addresses is considering the context in which they are defined, and then contrasting that with the places they turn up.

5. Common design choices and challenges

In creating a system, there are important common design choices that need to be made. Sometimes the answer is implicit within the overall design constraints of the system. Other times, considerable effort is required to refine specifications and make appropriate choices. As noted, these are common design questions. This is an area where understanding of previous systems? design discussions can be particularly helpful (in order not to repeat them needlessly).

5.1. Identifiers

In defining identifiers for a system, is the intention that they:

- o Name something -- the identifier will be associated with one entity (or instance of an entity), wherever that entity may be located.
- o Locate something -- identify the location of an entity (at some point in time).
- o Are Smart or dumb identifiers -- "smart" identifiers have structures that can be parsed to determine something about the thing identified (e.g., domain in which it is stored); "dumb" identifiers are opaque and must be resolved within the system.
- o Have uniqueness -- is the identifier/resource binding unique?
- o Have scope -- is the uniqueness (or other properties) only maintained within some limited scope, or is it global?
- o Permanent -- what is the expected level of permanence of the identifier's relevance (the ability to use it, the binding between the identifier and the identified resource). Or, are they transient identifiers?

5.2. Identified

The specifics of the identified resource need to be defined, as well.

- o Instances -- can there be multiple instances of a single resource? How can they be distinguished and/or how can two instances be equated. This is important if one needs to be able to cache

instances or otherwise validate "local" copies.

- o Scope of applicability -- what constitutes a "resource" in this system?

6. Issues in (mis)using identifiers

Things like using IP addressed out of context of the routing system ? assumptions about uniqueness and volatility may be improper.

7. IANA Considerations

This memo includes no request to IANA.

8. Security Considerations

This document is about considering applications systems. Security is important to applications, but is not specifically called out here.

9. References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997, <<http://xml.resource.org/public/rfc/html/rfc2119.html>>.

Author's Address

Leslie Daigle
ThinkingCat Enterprises
Leesburg, VA 20176
US

Email: ldaigle@thinkingcat.com

