

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: December 24, 2012

S. Farrell
Trinity College Dublin
D. Kutscher
NEC
C. Dannewitz
University of Paderborn
B. Ohlman
A. Keranen
Ericsson
P. Hallam-Baker
Comodo Group Inc.
June 22, 2012

Naming Things with Hashes
draft-farrell-decade-ni-08

Abstract

This document defines a set of ways to identify a thing (a digital object in this case) using the output from a hash function, specifying a new URI scheme for this, a way to map those to http URLs, and binary and human "speakable" formats for these names. The various formats are designed to support, but not require, a strong link to the referenced object such that the referenced object may be authenticated to the same degree as the reference to it.

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 December 24, 2012.

Copyright Notice

Copyright (c) 2012 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.

Table of Contents

1.	Introduction	3
2.	Hashes are what Count	4
3.	Named Information (ni) URI Format	5
3.1.	Content Type Query String Attribute	8
4.	.well-known URL Format	9
5.	URL Segment Format	10
6.	Binary Format	10
7.	Human-speakable Format	11
8.	Examples	12
9.	IANA Considerations	14
9.1.	Assignment of Named Information (ni) URI Scheme	14
9.2.	Assignment of Named Information for Humans (nih) URI Scheme	15
9.3.	Well Known 'ni' URI	15
9.4.	ni Hash Algorithm Registry	16
9.5.	Creation of ni parameter registry	17
10.	Security Considerations	17
11.	Acknowledgements	19
12.	References	20
12.1.	Normative References	20
12.2.	Informative References	21
	Authors' Addresses	22

1. Introduction

Names or identifiers are used in various protocols for identifying resources. In many scenarios those names or identifiers contain values that are hash function outputs. However, different deployments have chosen various different ways to include hash function outputs in such names or identifiers. This document specifies standard ways to do that to aid interoperability. We begin with a few example uses for the various ways to include a hash in a name as they are defined later in this document, for example Figure 1 shows an example of the "Named Information" (ni) URI scheme defined here.

```
ni:///sha-256;UyaQV-Ev4rdLoHyJJWCi110HfrYv9E1aGQA1M02X_-Q
```

Figure 1: Example ni URI

Hash function outputs can be used to ensure uniqueness in terms of mapping URIs [[RFC3986](#)] to a specific resource, or to make URIs hard to guess for security reasons. Since there is no standard way to interpret those strings, today in general only the creator of the URI knows how to use the hash function output. Other protocols, such as application layer protocols for accessing "smart objects" in constrained environments also require more compact (e.g., binary) forms of such identifiers, while in other situations people may have to input such values or talk about them, e.g., in a voice call.

As another example, protocols for accessing in-network storage servers need a way to identify stored resources uniquely and in a location-independent way so that replicas on different servers can be accessed by the same name. Also, such applications may require verifying that a resource representation that has been obtained actually corresponds to the name that was used to request the resource, i.e., verifying the integrity of the name-data binding.

Similarly, in the context of information-centric networking [[ref.netinf-design](#)] [[ref.ccn](#)] and elsewhere there is value in being able to compare a presented resource against the URI that was dereferenced in order to access that resource. If a cryptographically-strong comparison function can be used then this allows for many forms of in-network storage, without requiring as much trust in the infrastructure used to present the resource. The outputs of hash functions can be used in this manner, if presented in a standard way.

Additional applications might include creating references from web pages delivered over HTTP/TLS; DNS resource records signed using

DNSSEC or data values embedded in certificates, Certificate Revocation Lists (CRLs), or other signed data objects.

The "ni" URI scheme defined here is very similar to the "magnet link" informally defined in various other protocols. [[magnet](#)]

Media content-type, alternative locations for retrieval and other additional information about a resource named using this scheme can be provided using a query string. A companion specification [[I-D.hallambaker-decade-ni-params](#)] describes specific values that can be used in such query strings for these various purposes and other extensions to this basic format specification.

In addition, we also define a ".well-known" URL equivalent, and a way to include a hash as a segment of an HTTP URL, as well as a binary format for use in protocols that require more compact names and a human-speakable text form that could be used, e.g. for reading out (parts of) the name over a voice connection.

Not all uses of these names require use of the full hash output - truncated hashes can be safely used in some environments. For this reason, we define a new IANA registry for hash functions to be used with this specification so as not to mix strong and weak (truncated) hash algorithms in other protocol registries.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

Syntax definitions in this memo are specified according to ABNF [[RFC5234](#)].

2. Hashes are what Count

This section contains basic considerations related to how we use hash function outputs that are common to all formats.

When verifying whether two names refer to same object, an implementation MUST only consider the digest algorithm and the digest value, i.e., it MUST NOT consider other fields defined below (such as an authority field from a URI or any parameters). Implementations MUST consider two hashes identical, regardless of encoding, if the decoded hashes are based on the same algorithm and have the same length and the same binary value. In that case, the two names can be treated as referring to the same thing.

The sha-256 algorithm as specified in [[SHA-256](#)] is mandatory to

implement, that is, implementations MUST be able to generate/send and to accept/process names based on a sha-256 hash. However implementations MAY support additional hash algorithms and MAY use those for specific names, for example in a constrained environment where sha-256 is non-optimal or where truncated names are needed to fit into corresponding protocols (when a higher collision probability can be tolerated).

Truncated hashes MAY be supported if needed. When a hash value is truncated the name MUST indicate this. Therefore we use different hash algorithm strings for these, such as sha-256-32 for a 32-bit truncation of a sha-256 output. (A 32-bit truncated hash is essentially useless for security but might be useful for naming.)

When a hash value is truncated to N bits the left-most N bits, that is, the most significant N bits in network byte order, from the binary representation of the hash value MUST be used as the truncated value. An example of a 128-bit hash output truncated to 32 bits is shown in Figure 2.

128-bit hash: 0x265357902fe1b7e2a04b897c6025d7a2
32-bit truncated hash: 0x26535790

Figure 2: Example of Truncated Hash

When the input to the hash algorithm is a public key value, as may be used by various security protocols, the hash SHOULD be calculated over the public key in an X.509 SubjectPublicKeyInfo structure ([Section 4.1 of \[RFC5280\]](#)). This input has been chosen primarily for compatibility with DANE [[I-D.ietf-dane-protocol](#)], but also includes any relevant public key parameters in the hash input, which is sometimes necessary for security reasons. This does not force use of X.509 or full compliance with [[RFC5280](#)] since formatting any public key as a SubjectPublicKeyInfo is relatively straightforward and well supported by libraries.

Any of the formats defined below can be used to represent the resulting name for a public key.

Other than in the above special case where public keys are used, we do not specify the hash function input here. Other specifications are expected to define this.

3. Named Information (ni) URI Format

A Named Information (ni) URI consists of the following nine

components:

Scheme Name [Required] The scheme name is 'ni'.

Colon and Slashes [Required] The literal "://"

Authority [Optional] The optional authority component may assist applications in accessing the object named by an ni URI. There is no default value for the authority field. (See [[RFC3986](#)] [Section 3.2.2](#) for details.) While ni names with and without an authority differ syntactically from ni names with different authorities, all three refer to the same object if and only if the digest algorithm, length, and value are the same.

One slash [Required] The literal "/"

Digest Algorithm [Required] The name of the digest algorithm, as specified in the IANA registry defined in [Section 9.4](#) below.

Separator [Required] The literal ";"

Digest Value [Required] The digest value MUST be encoded using the base64url [[RFC4648](#)] encoding.

Query Parameter separator [Optional] '?' The query parameter separator acts as a separator between the digest value and the query parameters (if specified). For compatibility with IRIs, non-ASCII characters in the query part MUST be encoded as UTF-8, and the resulting octets MUST be %-encoded (see [[RFC3986](#)] [Section 2.1](#)).

Query Parameters [Optional] A tag=value list of optional query parameters as are used with HTTP URLs [[RFC2616](#)] with a separator character '&' between each. For example, "foo=bar&baz=bat"

It is OPTIONAL for implementations to check the integrity of the URI/resource mapping when sending, receiving or processing "ni" URIs.

Escaping of characters follows the rules in [RFC 3986](#). This means that %-encoding is used to distinguish between reserved and unreserved functions of the same character in the same URI component. As an example, an ampersand ('&') is used in the query part to separate attribute-value pairs; an ampersand in a value therefore has to be escaped as '%26'. Note that the set of reserved characters differs for each component, as an example, a slash ('/') does not have any reserved function in a query part and therefore does not have to be escaped. However, it can still appear escaped as '%2f' or '%2F', and implementations have to be able to understand such escaped

forms. Also note that any characters outside those allowed in the respective URI component have to be escaped.

The Named Information URI adapts the URI definition from the URI Generic Syntax [[RFC3986](#)]. We start with the base URI production:

```
URI = scheme ":" hier-part [ "?" query ] [ "#" fragment ]  
    ; from RFC 3986
```

Figure 3: URI syntax

Adapting that for the Named Information URI:

```
NI-URI      = ni-scheme ":" ni-hier-part [ "?" query ]  
              ; adapted from "URI" in RFC 3986  
              ; query is from RFC 3986, Section 3.4  
ni-scheme    = "ni"  
ni-hier-part = "://" [ authority ] "/" alg-val  
              ; authority is from RFC 3986, Section 3.2  
alg-val      = alg ";" val  
              ; adapted from "hier-part" in RFC 3986  
alg          = 1*unreserved  
val          = 1*unreserved  
              ; unreserved is from RFC3986, Section 2.3
```

Figure 4: ni Name syntax

The "val" field MUST contain the output of base64url encoding the result of applying the hash function ("alg") to its defined input, which defaults to the object bytes that are expected to be returned when the URI is dereferenced.

Relative ni URIs can occur. In such cases, the algorithm in [[RFC3986](#)] [Section 5](#) applies. As an example, in Figure 5, the absolute URI for "this third document" is "ni://example.com/sha-256-128;...".


```
<html>
<head>
  <title>ni: relative URI test</title>
  <base href="ni://example.com">
</head>

<body>
  <p>Please check <a href="sha-256;f40xZX...">this document</a>.
    and <a href="sha-256;UyaQV...">this other document</a>.
    and <a href="sha-256-128;...">this third document</a>.
  </p>
</body>
</html>
```

Figure 5: Example HTML with relative ni URI

The authority field in an ni URI is not quite the same as that from an HTTP URL, even though the same values (e.g. DNS names) may be usefully used in both. For an ni URI, the authority does not control nearly as much of the structure of the "right hand side" of the URI. With ni URIs we also define standard query string attributes and of course have a strictly defined way to include the hash value.

Internationalisation of strings within ni names is handled exactly as for http URIs - see [[I-D.ietf-httpbis-p1-messaging](#)] [Section 2.7](#).

[3.1](#). Content Type Query String Attribute

The semantics of a digest being used to establish a secure reference from an authenticated source to an external source may be a function of associated meta data such as the content type. The Content Type "ct" parameter specifies the MIME Content Type of the associated data as defined in [[RFC4288](#)]. See [Section 9.5](#) for the associated IANA registry for ni parameter names. as shown in Figure 6. Implementations of this specification MUST support parsing the ct= query string attribute name.

```
ni:///sha-256-32;f40xZQ?ct=text/plain
```

Figure 6: Example ni URI with Content Type

Protocols making use of ni URIs will need to specify how to verify name-data integrity for the MIME Content Types that they need to process and will need to take into account possible Content-Transfer-Encodings and other aspects of MIME encoding.

Implementations of this specification SHOULD support name-data

integrity validation for at least the application/octet-stream Content Type with no Content-Transfer-Encoding (which is equivalent to binary or 8bit). Additional Content Types and Content-Transfer-Encodings can of course also be supported, but are OPTIONAL.

If a) the user agent is sensitive to the Content Type and b) the ni name used has a ct= query string attribute and c) the object is retrieved (from a server) using a protocol that specifies a Content Type, then, if the two Content Types match, all is well. If, in this situation, the Content Types do not match, then the client SHOULD ignore the Content Type received from the server and use that from the ni name, or handle that situation as a potential security error. Content Type matching rules are defined in [\[RFC2045\] Section 5.1](#).

4. .well-known URL Format

We define a mapping between URIs following the ni URI scheme and HTTP [\[RFC2616\]](#) or HTTPS [\[RFC2617\]](#) URLs that makes use of the .well-known URI [\[RFC5785\]](#) by defining an "ni" suffix (see [Section 9](#)).

The HTTP(S) mapping MAY be used in any context where clients with support for ni URIs are not available.

Since the .well-known name-space is not intended for general information retrieval, if an application de-references a .well-known/ni URL via HTTP(S), then it will often receive a 3xx HTTP redirection response. A server responding to a request for a .well-known/ni URL will often therefore return a 3xx response and a client sending such a request MUST be able to handle that, as should any fully compliant HTTP [\[RFC2616\]](#) client.

For an ni name of the form "ni://n-authority/alg;val?query-string" the corresponding HTTP(S) URL produced by this mapping is "http://h-authority/.well-known/ni/alg/val?query-string", where "h-authority" is derived as follows: If the ni name has a specified authority (i.e., the n-authority is non-empty) then the h-authority MUST have the same value. If the ni name has no authority specified (i.e. the n-authority string is empty), a h-authority value MAY be derived from the application context. For example, if the mapping is being done in the context of a web page then the origin [\[RFC6454\]](#) for that web site can be used. Of course, there are in general no guarantees that the object named by the ni URI will be available via the corresponding HTTP(S) URL. But in the case that any data is returned, the retriever can determine whether or not it is content that matches the ni URI.

If an application is presented with a HTTP(S) URL with "/.well-

known/ni/" as the start of its pathname component, then the reverse mapping to an ni URI either including or excluding the authority might produce an ni URI that is meaningful, but there is no guarantee that this will be the case.

When mapping from an ni URI to a .well-known URL, an implementation will have to decide between choosing an "http" or "https" URL. If the object referenced does in fact match the hash in the URL, then there is arguably no need for additional data integrity, if the ni URI or .well-known URL was received "securely." However TLS also provides confidentiality, so there can still be reasons to use the "https" URL scheme even in this case. Additionally, web server policy such as [[I-D.ietf-websec-strict-transport-sec](#)] may dictate that data might only be available over "https". In general however, whether to use "http" or "https" is something that needs to be decided by the application.

5. URL Segment Format

Some applications may benefit from using hashes in existing HTTP URLs or other URLs. To do this one simply uses the "alg-val" production from the ni name scheme ABNF which may be included for example in the pathname, query string or even fragment components of HTTP URLs [[RFC2616](#)]. In such cases there is nothing present in the URL that ensures that a client can depend on compliance with this specification, so clients MUST NOT assume that any URL with a pathname component that matches the "alg-val" production was in fact produced as a result of this specification. That URL might or might not be related to this specification, only the context will tell.

6. Binary Format

If a more space-efficient version of the name is needed, the following binary format can be used. The binary format name consists of two fields: a header and the hash value. The header field defines how the identifier has been created and the hash value contains a (possibly truncated) result of a one-way hash over whatever is being identified by the hash value. The binary format of a name is shown in Figure 7.

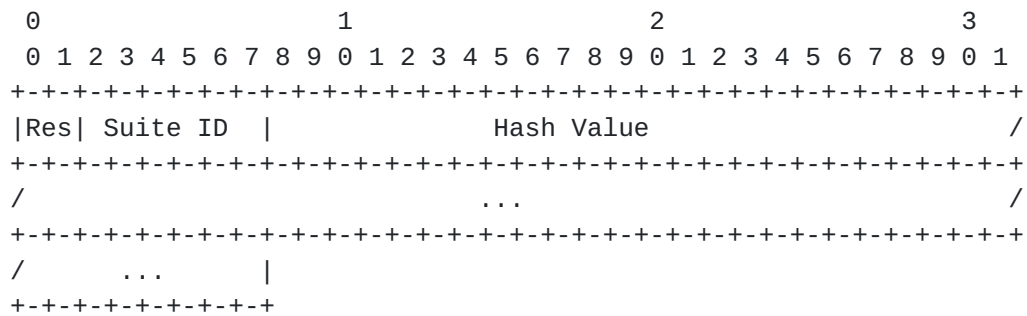


Figure 7: Binary Name Format

The Res field is a reserved 2-bit field for future use and MUST be set to zero for this specification.

The hash algorithm and truncation length are specified by the Suite ID. For maintaining efficient encoding for the binary format, only a few hash algorithms and truncation lengths are supported. See [Section 9.4](#) for details.

A hash value that is truncated to 120 bits will result in the overall name being a 128-bit value which may be useful for protocols that can easily use 128-bit identifiers.

7. Human-speakable Format

Sometimes a name may need to be referred to in a format that is easy or unambiguous for humans to read out, for example, over the phone. The ni URI format is not very suited for this, as, for example, base64url uses both upper and lower case which can easily cause confusion. For this particular purpose, ("speaking" the value of an ni URI) the following more verbose but less ambiguous (when spoken) URI format is defined with scheme name "nih", standing for "Named Information for Humans." (Or possibly "Not Invented Here," which is clearly false, and therefore worth including :-). Note that nih URIs are not intended for general uses, but only for this specific use-case.

The justification for using a URI scheme for this is that that might help a user agent for the speaker to better display the value, or perhaps if there was some use-case for a machine to speak the value. We do not include the query string since there is no way to ensure that its value might be spoken unambiguously, and similarly for the authority, where e.g. internationalised forms of domain name could not be usefully spoken. This leaves the hash value as the only part of the ni URI that we feel can be usefully included. But since speakers or listeners (or speech recognition) may err, we also

include a check-digit to catch common errors.

Fields in nih URIs are separated by a semi-colon (;) character. The first field is a hash algorithm string, as in the ni URI format. The hash value is represented using lower-case ASCII hex characters, for example an octet with the decimal value 58 (0x3A) is encoded as '3a'. This is the same as base16 encoding as defined in [RFC 4648](#) [[RFC4648](#)] except using lower-case letters.

The hash value is OPTIONALLY followed by a semi-colon ';' then a checkdigit. The checkdigit MUST be calculated using Luhn's mod N algorithm (with N=16) as defined in [[ISOIEC7812](#)], (see also http://en.wikipedia.org/wiki/Luhn_mod_N_algorithm). The input to the calculation is the ASCII-HEX encoded hash value (i.e. "val" in the ABNF production below). This maps the ASCII-HEX so that '0'=0,...'9'=9,'a'=10,...'f'=15. None of the other fields are input when calculating the checkdigit.

```
humanname = "nih:" alg-val [ ";" checkdigit ]
alg        = 1*unreserved
val        = 1*unreserved
checkdigit = unreserved
```

Figure 8: Human-speakable syntax

For algorithms that have a Suite ID reserved (see Figure 11), the alg field MAY contain the ID value as a ASCII encoded decimal number instead of the hash name string (for example, "3" instead of "sha-256-120"). Implementations MUST be able to match the decimal ID values for the algorithms and hash lengths that they support even if they do not support the binary format.

There is no such thing as a relative nih URI.

8. Examples

The following ni URI is generated from the text "Hello World!" (without the quotes, being 12 characters), using the sha-256 algorithm shown with and without an authority field:

```
ni:///sha-256;f40xZX_x_F05LcGBSKHWXfwtSx-j1ncoSt3SABJtkGk
```

```
ni://example.com/sha-256;f40xZX_x_F05LcGBSKHWXfwtSx-j1ncoSt3SABJtkGk
```

The following HTTP URL represents a mapping from the previous ni name based on the algorithm outlined above.

http://example.com/.well-known/ni/sha-256/
f40xZX_x_F05LcGBSKHWXfwtSx-j1ncoSt3SABJtkGk

Given the DER-encoded SubjectPublicKeyInfo in Figure 9 we derive the names shown in Figure 10 for this value.

```

00000000 30 82 01 22 30 0d 06 09 2a 86 48 86 f7 0d 01 01
00000020 01 05 00 03 82 01 0f 00 30 82 01 0a 02 82 01 01
00000040 00 a2 5f 83 da 9b d9 f1 7a 3a 36 67 ba fd 5a 94
00000060 0e cf 16 d5 5a 55 3a 5e d4 03 b1 65 8e 6d cf a3
00000100 b7 db a4 e7 cc 0f 52 c6 7d 35 1d c4 68 c2 bd 7b
00000120 9d db e4 0a d7 10 cd f9 53 20 ee 0d d7 56 6e 5b
00000140 7a ae 2c 5f 83 0a 19 3c 72 58 96 d6 86 e8 0e e6
00000160 94 eb 5c f2 90 3e f3 a8 8a 88 56 b6 cd 36 38 76
00000200 22 97 b1 6b 3c 9c 07 f3 4f 97 08 a1 bc 29 38 9b
00000220 81 06 2b 74 60 38 7a 93 2f 39 be 12 34 09 6e 0b
00000240 57 10 b7 a3 7b f2 c6 ee d6 c1 e5 ec ae c5 9c 83
00000260 14 f4 6b 58 e2 de f2 ff c9 77 07 e3 f3 4c 97 cf
00000300 1a 28 9e 38 a1 b3 93 41 75 a1 a4 76 3f 4d 78 d7
00000320 44 d6 1a e3 ce e2 5d c5 78 4c b5 31 22 2e c7 4b
00000340 8c 6f 56 78 5c a1 c4 c0 1d ca e5 b9 44 d7 e9 90
00000360 9c bc ee b0 a2 b1 dc da 6d a0 0f f6 ad 1e 2c 12
00000400 a2 a7 66 60 3e 36 d4 91 41 c2 f2 e7 69 39 2c 9d
00000420 d2 df b5 a3 44 95 48 7c 87 64 89 dd bf 05 01 ee
00000440 dd 02 03 01 00 01

00000000 53 26 90 57 e1 2f e2 b7 4b a0 7c 89 25 60 a2 d7
00000020 53 87 7e b6 2f f4 4d 5a 19 00 25 30 ed 97 ff e4

```

Figure 9: A SubjectPublicKeyInfo used in examples and its sha-256 hash


```

+-----+
| URI:                                     |
| ni:///sha-256;UyaQV-Ev4rdLoHyJJWci110HfrYv9E1aGQAlM02X_-Q |
+-----+
| .well-known URL (split over 2 lines): |
| http://example.com/.well-known/ni/sha256/ |
| UyaQV-Ev4rdLoHyJJWci110HfrYv9E1aGQAlM02X_-Q |
+-----+
| URL Segment:                           |
| sha-256;UyaQV-Ev4rdLoHyJJWci110HfrYv9E1aGQAlM02X_-Q |
+-----+
| Binary name (ASCII hex encoded) with 120-bit truncated hash value |
| which is Suite ID 0x03: |
| 0353 2690 57e1 2fe2 b74b a07c 8925 60a2 |
+-----+
| Human-speakable form of a name for this key (truncated to 120 bits |
| in length) with checkdigit: |
| nih:sha-256-120;53269057e12fe2b74ba07c892560a2;f |
+-----+
| Human-speakable form of a name for this key (truncated to 32 bits |
| in length) with checkdigit: |
| nih:sha-256-32;53269057;b |
+-----+
| Human-speakable form using decimal presentation of the |
| algorithm ID (sha-256-120) with checkdigit: |
| nih:3;53269057e12fe2b74ba07c892560a2;f |
+-----+

```

Figure 10: Example Names

9. IANA Considerations

9.1. Assignment of Named Information (ni) URI Scheme

The procedures for registration of a URI scheme are specified in [RFC 4395](#) [[RFC4395](#)]. The following is the proposed assignment template.

URI scheme name: ni

Status: Permanent

URI scheme syntax. See [Section 3](#)

URI scheme semantics. See [Section 3](#)

Encoding considerations. See [Section 3](#)

Applications/protocols that use this URI scheme name: General applicability.

Interoperability considerations: Defined here.

Security considerations: See [Section 10](#)

Contact: Stephen Farrell, stephen.farrell@cs.tcd.ie

Author/Change controller: IETF

References: As specified in this document

[9.2.](#) Assignment of Named Information for Humans (nih) URI Scheme

The procedures for registration of a URI scheme are specified in [RFC 4395](#) [[RFC4395](#)]. The following is the proposed assignment template.

URI scheme name: nih

Status: Permanent

URI scheme syntax. See [Section 7](#)

URI scheme semantics. See [Section 7](#)

Encoding considerations. See [Section 7](#)

Applications/protocols that use this URI scheme name: General applicability.

Interoperability considerations: Defined here.

Security considerations: See [Section 10](#)

Contact: Stephen Farrell, stephen.farrell@cs.tcd.ie

Author/Change controller: IETF

References: As specified in this document

[9.3.](#) Well Known 'ni' URI

The procedures for registration of a Well Known URI entry are specified in [RFC 5785](#) [[RFC5785](#)]. The following is the proposed assignment template.

URI suffix: ni

Change controller: IETF

Specification document(s): This document

Related information: None

9.4. ni Hash Algorithm Registry

IANA is requested to create a new registry for hash algorithms as used in the name formats specified here and called the "ni Hash Algorithm Registry". Future assignments are to be made through Expert Review [[RFC5226](#)]. This registry has five fields, the suite ID, the hash algorithm name string, the truncation length, the underlying algorithm reference and a status field that indicates if algorithm is current or deprecated and should no longer be used. The status field can have the value "current" or "deprecated". Other values are reserved for possible future definition.

If the status is "current", then that does not necessarily mean that the algorithm is "good" for any particular purpose, since the cryptographic strength requirements will be set by other applications or protocols.

A request to mark an entry as "deprecated" can be done by sending a mail to the Designated Expert. Before approving the request, the community **MUST** be consulted via a "call for comments" of at least two weeks by sending a mail to the IETF discussion list.

Initial values are specified below. The Designated Expert **SHOULD** generally approve additions that reference hash algorithms that are widely used in other IETF protocols. In addition, the Designated Expert **SHOULD NOT** accept additions where the underlying hash function (with no truncation) is considered weak for collisions. Part of the reasoning behind this last point is that inclusion of code for weak hash functions, e.g. the MD5 algorithm, can trigger costly false-positives if code is audited for inclusion of obsolete ciphers.

The suite ID field ("ID") can be empty, or can have values between 0 and 63, inclusive. Because there are only 64 possible values, this field is **OPTIONAL** (leaving it empty if omitted). Where the binary format is not expected to be used for a given hash algorithm, this field **SHOULD** be omitted. If an entry is registered without a suite ID, the Designated Expert **MAY** allow for later allocation of a suite ID, if that appears warranted. The Designated Expert **MAY** consult the community via a "call for comments" by sending a mail to the IETF discussion list before allocating a suite ID.

ID	Hash name string	Value length	Reference	Status
0	Reserved			
1	sha-256	256 bits	[SHA-256]	current
2	sha-256-128	128 bits	[SHA-256]	current
3	sha-256-120	120 bits	[SHA-256]	current
4	sha-256-96	96 bits	[SHA-256]	current
5	sha-256-64	64 bits	[SHA-256]	current
6	sha-256-32	32 bits	[SHA-256]	current
32	Reserved			

Figure 11: Suite Identifiers

The Suite ID value 32 is reserved for compatibility with ORCHIDs [[RFC4843](#)].

The referenced hash algorithm matching to the Suite ID, truncated to the length indicated, according to the description given in [Section 2](#), is used for generating the hash. The Designated Expert is responsible for ensuring that the document referenced for the hash algorithm meets the "specification required" rule."

9.5. Creation of ni parameter registry

IANA is requested to create a new registry entitled "Named Information URI Parameter Definitions".

The policy for future assignments to the registry is Expert Review, and as for the ni Hash Algorithm Registry above, the Designated Expert is responsible for ensuring that the document referenced for the parameter definition meets the "specification required" rule."

The fields in this registry are the parameter name, a description and a reference. The parameter name **MUST** be such that it is suitable for use as a query string parameter name in an ni URI.

The initial contents of the registry are:

Parameter	Meaning	Reference
-----	-----	-----
ct	Content Type	[RFC-THIS]

10. Security Considerations

No secret information is required to generate or verify a name of the form described here. Therefore a name like this can only provide evidence for the integrity for the referenced object and the proof of

integrity provided is only as good as the proof of integrity for the name from which we started. In other words, the hash value can provide a name-data integrity binding between the name and the bytes returned when the name is de-referenced using some protocol.

Disclosure of a name value does not necessarily entail disclosure of the referenced object but may enable an attacker to determine the contents of the referenced object by reference to a search engine or other data repository or, for a highly formatted object with little variation, by simply guessing the value and checking if the digest value matches. So the fact that these names contain hashes does not protect the confidentiality of the object that was input to the hash.

The integrity of the referenced content would be compromised if a weak hash function were used. SHA-256 is currently our preferred hash algorithm which is why we've only added SHA-256 based suites to the initial IANA registry.

If a truncated hash value is used, certain security properties will be affected. In general a hash algorithm is designed to produce sufficient bits to prevent a 'birthday attack' collision occurring. To ensure that the difficulty of discovering two pieces of content that result in the same digest with a work factor $O(2^x)$ by brute force requires a digest length of $2x$. Many security applications only require protection against a 2nd pre-image attack which only requires a digest length of x to achieve the same work factor. Basically, the shorter the hash value used, the less security benefit you can possibly get.

An important thing to keep in mind is not to make the mistake of thinking two names are the same when they aren't. For example, a name with a 32 bit truncated sha-256 hash is not the same as a name with the full 256 bits of hash output, even if the hash value for one is a prefix of that for the other.

The reason for this is that if an application treats those as the same name then that might open up a number of attacks. For example, if I publish an object with the full hash, then I probably (in general) don't want some other application to treat a name with just the first 32 bits of that as referring to the same thing, since the 32 bit name will have lots of colliding objects. If `ni` or `nih` URIs become widely used, there will be many cases where names will occur more than once in application protocols, and it'll be unpredictable which instance of the name would be used for name-data integrity checking, leading to threats. For this reason, we require that the algorithm, length and value all match before we consider two names to be the same.

The fact that an ni URI includes a domain name in the authority field by itself implies nothing about the relationship between the owner of the domain name and any content referenced by that URI. While a name-data integrity service can be provided using ni URIs, that does not in any sense validate the authority part of the name. For example, there is nothing to stop anyone creating an ni URI containing a hash of someone else's content. Application developers MUST NOT assume any relationship between the registrant of the domain name that is part of an ni URI and some matching content just because the ni URI matches that content.

If name-data integrity is successfully validated, and the hash is strong and long enough, then the "web origin" [[RFC6454](#)] for the bytes of the named object is really going to be the place from which you got the ni name and not the place from which you got the bytes of the object. This appears to offer a potential benefit if using ni names for, for example, scripts included from a HTML page accessed via server-authenticated https. If name-data integrity is not validated (and it is optional), or fails, then the web origin is, as usual, the place from which the object bytes were received. Applications making use of ni names SHOULD take this into account in their trust models.

Some implementations might mis-handle ni URIs that include non-base64 characters, whitespace or other non-conforming strings and that could lead to erroneously considering names to be the same when they are not. An ni URI that is malformed in such ways MUST NOT be treated as matching any other ni URI. Implementers need to check the behaviour of libraries for such parsing problems.

[11.](#) Acknowledgements

This work has been supported by the EU FP7 project SAIL. The authors would like to thank SAIL participants to our naming discussions, especially Jean-Francois Peltier, for their input.

The authors would also like to thank Carsten Bormann, Martin Durst, Tobias Heer, Bjoern Hoehrmann, Tero Kivinen, Barry Leiba, Larry Masinter, David McGrew, Alexey Melnikov, Bob Moskowitz, Jonathan Rees, Eric Rescorla, Zach Shelby, Martin Thomas, for their comments and input to the document. Thanks, in particular, to James Manger for correcting the examples.

[12.](#) References

12.1. Normative References

- [I-D.ietf-httpbis-p1-messaging]
Fielding, R., Lafon, Y., and J. Reschke, "HTTP/1.1, part 1: URIs, Connections, and Message Parsing", [draft-ietf-httpbis-p1-messaging-19](#) (work in progress), March 2012.
- [ISOIEC7812]
ISO, ""ISO/IEC 7812-1:2006 Identification cards -- Identification of issuers -- Part 1: Numbering system",", October 2006, <http://www.iso.org/iso/iso_catalogue/catalogue_tc/catalogue_detail.htm?csnumber=39698>.
- [RFC2045] Freed, N. and N. Borenstein, "Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies", [RFC 2045](#), November 1996.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", [RFC 2616](#), June 1999.
- [RFC2617] Franks, J., Hallam-Baker, P., Hostetler, J., Lawrence, S., Leach, P., Luotonen, A., and L. Stewart, "HTTP Authentication: Basic and Digest Access Authentication", [RFC 2617](#), June 1999.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, [RFC 3986](#), January 2005.
- [RFC4288] Freed, N. and J. Klensin, "Media Type Specifications and Registration Procedures", [BCP 13](#), [RFC 4288](#), December 2005.
- [RFC4395] Hansen, T., Hardie, T., and L. Masinter, "Guidelines and Registration Procedures for New URI Schemes", [BCP 35](#), [RFC 4395](#), February 2006.
- [RFC4648] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", [RFC 4648](#), October 2006.
- [RFC5234] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, [RFC 5234](#), January 2008.
- [RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S.,

Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", [RFC 5280](#), May 2008.

[RFC5785] Nottingham, M. and E. Hammer-Lahav, "Defining Well-Known Uniform Resource Identifiers (URIs)", [RFC 5785](#), April 2010.

[SHA-256] NIST, "United States National Institute of Standards and Technology (NIST), FIPS 180-2: Secure Hash Standard", August 2002, <<http://csrc.nist.gov/publications/fips/fips180-2/fips180-2.pdf>>.

12.2. Informative References

[I-D.hallambaker-decade-ni-params]
Hallam-Baker, P., Stradling, R., Farrell, S., Kutscher, D., and B. Ohlman, "The Named Information (ni) URI Scheme: Optional Features", [draft-hallambaker-decade-ni-params-03](#) (work in progress), June 2012.

[I-D.ietf-dane-protocol]
Hoffman, P. and J. Schlyter, "The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA", [draft-ietf-dane-protocol-23](#) (work in progress), June 2012.

[I-D.ietf-websec-strict-transport-sec]
Hodges, J., Jackson, C., and A. Barth, "HTTP Strict Transport Security (HSTS)", [draft-ietf-websec-strict-transport-sec-09](#) (work in progress), June 2012.

[RFC4843] Nikander, P., Laganier, J., and F. Dupont, "An IPv6 Prefix for Overlay Routable Cryptographic Hash Identifiers (ORCHID)", [RFC 4843](#), April 2007.

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 5226](#), May 2008.

[RFC6454] Barth, A., "The Web Origin Concept", [RFC 6454](#), December 2011.

[magnet] Wikipedia article, "Magnet URI Scheme", April 2012, <http://en.wikipedia.org/wiki/Magnet_link>.

[ref.ccn] Jacobson et al., "Networking Named Content", CoNEXT 2009 ,

December 2009.

[ref.netinf-design]

Ahlgren, D'Ambrosio, Dannewitz, Marchisio, Marsh, Ohlman,
Pentikousis, Rembarz, Strandberg, and Vercellone, "Design
Considerations for a Network of Information", Re-Arch 2008
Workshop , December 2008.

Authors' Addresses

Stephen Farrell
Trinity College Dublin
Dublin, 2
Ireland

Phone: +353-1-896-2354
Email: stephen.farrell@cs.tcd.ie

Dirk Kutscher
NEC
Kurfuersten-Anlage 36
Heidelberg,
Germany

Phone:
Email: kutscher@neclab.eu

Christian Dannewitz
University of Paderborn
Paderborn
Germany

Email: cdannewitz@upb.de

Borje Ohlman
Ericsson
Stockholm S-16480
Sweden

Email: Borje.Ohlman@ericsson.com

Ari Keranen
Ericsson
Jorvas 02420
Finland

Email: ari.keranen@ericsson.com

Phillip Hallam-Baker
Comodo Group Inc.

Email: philliph@comodo.com