

A URN Namespace for Network Resources

Status of This Document

Grid Final Draft (GFD), Informational.

Copyright Notice

Copyright © Open Grid Forum (2011-2012). Some Rights Reserved. Distribution is unlimited.

Abstract

This document specifies the procedure to create Uniform Resource Names (URNs) in the `urn:ogf:network` namespace. URNs in this namespace can be used to define logical network resources, such a devices, (logical) ports, (logical) links, and topologies.

Contents

Abstract	1
Contents	1
1 Introduction	3
2 Registration	3
2.1 Namespace Identifier	3
2.2 Document Version	3
3 Syntax	3
3.1 Syntactic Structure	3
3.2 Encoding	5
3.3 Rules for Lexical Equivalence	5
3.4 Assignment	5
3.5 Validation	5
3.6 URN Rewriting	5

4	Namespace Considerations	6
4.1	Scope	6
4.2	Resource Type Described	6
4.3	Identifier uniqueness considerations	7
5	Community Considerations	7
5.1	Process of Organisation Identifier Assignment	7
5.2	Process of Network Resource Identifier Assignment	7
5.3	Identifier persistence considerations	8
6	Examples	9
7	Security Considerations	9
8	Prior Usage	10
8.1	GLIF Community	10
8.2	PerfSONAR Community	11
8.3	AutoGOLE Community	12
8.4	Backwards Compatibility	12
9	Contributors	13
	Intellectual Property Statement	13
	Disclaimer	14
	Full Copyright Notice	14
	References	14

1 Introduction

Uniform Resource Names (URNs) are persistent, globally unique identifiers [RFC 2141].

Topology exchange between network operators requires globally unique identifiers for network resources. The `urn:ogf:network` namespace provides globally unique identifiers for naming network resources without central registration.

This document defines and registers the `urn:ogf:network` namespace in accordance with [GFD.191]. It defines a procedure how organisation can create a globally unique organisation identifier, which can be used to prefix locally unique resource identifiers to form globally unique resource identifiers.

By re-using the domain name system with a date, no additional central registry or procedural overhead is required to create a globally unique organisation identifier.

Notational Conventions

The keywords “MUST”, “MUST NOT”, “REQUIRED”, “SHALL”, “SHALL NOT”, “SHOULD”, “SHOULD NOT”, “RECOMMENDED”, “MAY”, and “OPTIONAL” are to be interpreted as described in [RFC 2119].

2 Registration

2.1 Namespace Identifier

“`urn:ogf:network:`” The OGF is the Namespace Organisation for the `urn:ogf:network` namespace.

2.2 Document Version

Registration version number: 1

Registration date: 2012-09-01

3 Syntax

3.1 Syntactic Structure

A network resource URN (NURN) is defined by the following rules. These rules follow Augmented BFR [RFC 5234] format.

```

NURN = "urn:ogf:network:" ORGID ":" OPAQUE-PART *1QUERY *1FRAGMENT
ORGID = FQDN ":" DATE ; ID of assigning organisation
FQDN = 1*(ALPHA / DIGIT / "-" / ".") ; Domain name
DATE = YEAR *1(MONTH *1DAY) ; Date of creation of ORGID
YEAR = 4DIGIT
MONTH = 2DIGIT
DAY = 2DIGIT
OPAQUE-PART = *(ALPHA / DIGIT / OTHER)
OTHER = ALLOWED / EXTENSION
ALLOWED = "+" / "," / "-" / "." / ":" / ";" / "=" / "_"
EXTENSION = "!" / "$" / "(" / ")" / "*" / "@" / "~" / "&"
QUERY = "?" *QFCHAR
FRAGMENT = "#" *QFCHAR
QFCHAR = ALPHA / DIGIT / OTHER

```

ALPHA, DIGIT and HEXDIG are defined by [RFC 5234], OTHER is almost equal to <other> as defined by [RFC 2141], but lacks the single quote (') character and includes the ampersand (&) and tilde characters (~). QFCHAR is a subset of pchar as defined by [RFC 3986], only lacking the single quote character (') and percentage encoding ("% HEXDIG HEXDIG).

ALLOWED characters MAY be used for the assignment of network resource URNs. EXTENSION characters SHOULD NOT be used to assign network resource URNs. To allow for future extensions, parsers SHOULD accept network resource URNs with EXTENSION characters.

The QUERY and FRAGMENT parts MUST NOT be present in any **assigned** URN. This specification reserves their use for future standardization.

A network resource URN MUST NOT contain percentage-encoded characters ("% HEXDIG HEXDIG). It should also be noted that the following characters (which are either allowed by the URI or URN specification) MUST NOT be used in the OPAQUE-PART of a network resource URN: "%", "/", "?", "#", and "'".

DOMAIN is a fully qualified domain name (FQDN) of the URN assigning organisation in LDR format [RFC 5890]. Valid examples are `example.net` and `example.xn--jxalpdlp`. DATE is a date (either year, year+month or year+month+day). The combination of DOMAIN and DATE forms the *organisation identifier*, ORGID.

The full length of a NURN MUST NOT exceed 255 characters.

OPAQUE-PART is opaque, and MUST NOT be parsed or interpreted by any organisation except for the organisation that assigned the URN.

3.2 Encoding

A network resource URN uses a subset of 7-bit ASCII characters. No percentage-encoded characters are allowed.

3.3 Rules for Lexical Equivalence

Network resource URNs are lexical equivalent if and only if they are byte-equivalent after case normalisation.

Consider the following URNs:

- 1- urn:ogf:network:example.net:2012:path-glif-0418
- 2- UrN:oGf:NeTwOrK:eXaMpLe.NeT:2012:pAtH-gLiF-0418
- 3- URN:OGF:NETWORK:EXAMPLE.NET:2012:PATH-GLIF-0418

URNs 1, 2, and 3 are lexically equivalent to each other.

3.4 Assignment

The ORGID part must belong to the assignment organisation, as described in section 5.1.

Assigned network resource URNs **MUST NOT** contain a fragment or query part.

The characters defined in **EXTENSION SHOULD NOT** be used in assignment of network resources URNs, and are reserved for future use. Only characters in **ALPHA / DIGIT / ALLOWED SHOULD** be used in the **OPAQUE-PART**.

The length of the URN **MUST NOT** exceed 255 bytes.

3.5 Validation

A network resource URN that does not follow the specified syntax **SHOULD** be rejected.

No specific validation service or resolution service is defined in this document.

A recipient should either use the credibility of the sender or some other mechanism to judge the correctness of a given URN.

3.6 URN Rewriting

A recipient **MUST NOT** rewrite the URN if the rewriting results in a URN which is not lexically equivalent to the received URN. In particular, percentage-decoding of the URN as

described in section 6.2.2.2. of [RFC 3986] MUST NOT take place.

If two lexical equivalent URNs with different capitalisation have been received, the recipient MAY pick one of the two capitalisations, and use that in all communications, effectively rewriting the URNs.

With the above exception, URNs SHOULD retain the same capitalisation in a message exchange.

4 Namespace Considerations

4.1 Scope

The `urn:ogf:network` namespace is created to allow network operators to uniquely define resources in their network and facilitate unambiguous exchange of topology data with other network operators.

The only requirement for naming network resources is administrative ownership of the domain name used for `DOMAIN` on the `DATE` of the identifier assignment (see section 5.1). No other central registration is required.

The intended use of the `urn:ogf:network` namespace is to describe logical network resources roughly on OSI layers 1 and OSI layer 2. “Logical network resources” intends to mean elements in a functional topology description, rather than physical resources. It is expected that a peering network is only interested in the functional description of the network, not of its (physical) implementation. Nevertheless, this document does not forbid the description of other resources, such a physical network resources for inventory management.

4.2 Resource Type Described

The exact type of resource described by a URN can not and MUST NOT be determined from the syntax of the URN. This information MUST be provided by the context or through other means by the data exchange protocol.

Network resources URNs SHOULD identify *manifestations* of a network resources — they should refer to a functional component in a network that remains in place for a prolonged period of time. New version of the resource SHOULD NOT receive a new identifier. The change of attributes over time SHOULD be dealt with by a protocol, not by a change of the URN.

4.3 Identifier uniqueness considerations

URN identifiers **MUST** be assigned uniquely – they are assigned to at most one resource, and **MUST NOT** be re-assigned.

URN assigning organisations **MUST** follow these requirements before assigning URNs to network resources.

A single network resource **MAY** be identified by multiple URNs.

5 Community Considerations

5.1 Process of Organisation Identifier Assignment

An organisation that wishes to become an assigning organisation, must pick a globally unique organisation identifier.

An organisation identifier consists of two components, a fully qualified domain name and a date, which must both be chosen by the assigning organisation.

The assigning organisation **MUST** be the administrative contact of the chosen domain [RFC 5890] for at least the duration of the date.

It is **RECOMMENDED** that the date is a year. Organisation that expect their DNS registration to be more volatile **SHOULD** pick a more fine grained date specification (year+month or year+month+day).

There is no need for the assigning organisation to register themselves at the Open Grid Forum (the Namespace Organisation for the `urn:ogf:network` namespace).

An organisation **MAY** use multiple organisation identifiers. For example, an organisation may pick a new organisation identifier in order to create a new syntax for their **OPAQUE-PART** syntax.

5.2 Process of Network Resource Identifier Assignment

An assigning organisation assigns **OPAQUE-PARTs** to its network resources. The following requirements apply to the **OPAQUE-PART**:

- The **OPAQUE-PART** **MUST** uniquely define at most one network resource;
- The **OPAQUE-PART** **MUST NOT** be re-assigned;
- The **OPAQUE-PART** **SHOULD NOT** specify any properties of the network resource;

- The OPAQUE-PART MAY contain some structure according to some policy internal to the assigning organisation.
- The OPAQUE-PART MUST have a valid syntax (use only allowed characters, does not exceed maximum length).

The reason that the OPAQUE-PART SHOULD NOT contain any properties is because a URN MUST be persistent: it MUST NOT change, even after the properties of the described resource change. Naming these properties in the URN gives a false sense of meaning to the URN. Peer may inadvertently assume the identifier describes certain properties, and act upon that, even if the properties have long changed.

Good examples of URNs:

```
urn:ogf:network:example.net:2012:9ad7ef-mcasip-139284
urn:ogf:network:example.net:2012:link:2013:175
urn:ogf:network:example.net:2012:port:eth:98a-3470
```

Not so good examples of URNs:

```
urn:ogf:network:example.net:2012:link:24x7-protected:925-175
urn:ogf:network:example.net:2012:sw3.rtr.example.net:port3-1:vlan118
link:eth:US_CHI-NL_AMS-3937
```

A useful syntax for OPAQUE-PART is <type>:<year of creation>:<sequence number>, e.g. `port:2013:129`.

While `port:2013:129` contains attributes (type and year of creation), these may be acceptable as they will never change. `link:24x7-protected:925-175` contains attributes about the type of link, which may change in the future, and is therefore not a good URN. `link:eth:US_CHI-NL_AMS-3937` is also not a good URN. It contains the end points of a path, which are unlikely to change. However, if the path is actually an Ethernet LAN, it is possible to add another end-point, changing these properties. The network domains along the path may use this identifier for monitoring and do not accept a change in identifier. For that reason, it is best never to add attributes to a URN identifier.

5.3 Identifier persistence considerations

[RFC 3406] requires that URNs MUST NOT be re-assigned. Ever. In practice, it is impossible to control what identifiers will be assigned in a few decades from now.

The requirement of the date in the organisation is sufficient guarantee. If a domain name is transferred, or an organisation decides to start over with the assignment of local identifiers, it is easy enough to create a new organisational identifier.

Any organisation that wishes to assign names in the `urn:ogf:network` namespace must do so after due diligence, and make sure that no re-assignment occurs within the namespace(s) of the organisation and that the assigned name does not contain attributes which can change during the lifetime of the resource.

6 Examples

Syntactically valid network resource URNs which MAY be assigned include:

```
urn:ogf:network:example.net:2012:9ad7ef-mcasip-139284
urn:ogf:network:example.net:2012:link:925-175
urn:ogf:network:example.net:2012:port:2011-07-129
urn:ogf:network:example.net:20120916:4A6173706572
urn:ogf:network:example.net:2012:
urn:ogf:network:example.net:2012:l=214.56:x=a5y
```

The following URNs contain characters in the extension range. While they SHOULD NOT be assigned to network resources, a recipient SHOULD accept these examples:

```
urn:ogf:network:example.net:20120916:4A6173706572(AMS-GEN)
urn:ogf:network:example.net:20120916:l=*:x=a5y
```

The following example is a syntactically valid URN, which contains a query part and hence MUST NOT be assigned to a network resource, but MAY be used to query for a network resource, provided that subsequent standards define the syntax of the query part:

```
urn:ogf:network:example.net:2012:portgroup:ams3?vlan=18
```

The following URNs is *invalid*, and SHOULD be rejected by a recipient, because a slash is not allowed in a URN (this is a limitation of all URNs, not just this specification):

```
urn:ogf:network:example.net:2012:port:sw3:eth4/1:vlan18
```

7 Security Considerations

While this specification goes to great length to avoid accidental naming collisions, malicious software can easily craft a NURN to collide with an existing NURN. Recipients of a NURN MUST take such risks in consideration.

Recipients of a NURN MUST NOT assume that a NURN was crafted by the domain specified in the DOMAIN part of the NURN, without a proper validation check.

The allowed syntax is so limited that it is not expected that similar-looking malicious NURNs will be an issue. Users and applications should be able to detect the differences between `urn:ogf:network:example.c0m:4638I27` and `urn:ogf:network:example.com:4638127`.

Software that takes input from a user **MUST** ensure that the NURN is syntactically correct before transmitting it. For example, it **SHOULD** remove any trailing spaces from the user input.

Information in the OPAQUE-PART **MUST NOT** be interpreted to have any meaning whatsoever. While the originating domain may have included meaningful attributes in the NURN, these attributes may be out-of-date.

8 Prior Usage

URN identifiers in the `urn:ogf:network` namespace have been in use in three communities, GLIF, PerfSONAR, and AutoGOLE, with mutually conflicting syntaxes.

8.1 GLIF Community

The Global Lambda Integrated Facility (GLIF) is a community of research and education networks. Operators in this community agreed to use unique identifier for *lightpaths*, dedicated inter-domain circuits for researchers.

These identifiers take the form:

```
GLOBAL-ID = "urn:ogf:network:" DOMAIN ":" LOCAL-PART
DOMAIN = 1*(ALPHA / DIGIT / "-" / ".") ; Domain name
LOCAL-PART = 1*(ALPHA / DIGIT / "-" / ".")
```

For example:

```
urn:ogf:network:canarie.ca:kisti-uninett-glif-001
urn:ogf:network:es.net:4005
urn:ogf:network:dcn.internet2.edu:6811
```

The syntax is described in [GLIF-ID].

Identifiers described by the GLIF community generally do not contain a date, although it is possible to construct a URN which is both a valid NURN and GLOBAL-ID.

8.2 PerfSONAR Community

PerfSONAR is a distributed system for network performance monitoring on paths crossing several networks. Much of the perfSONAR protocols are standardised by the OGF in the Network Measurement (NM) and Network Measurement and Control (NMC) working groups. URNs in the `urn:ogf:network` namespace are used for topology description.

These identifiers take the form:

```
PS-URN = "urn:ogf:network" 1DOMAIN-PART *1NODE-PART *1PORT-PART
        *1LINK-PART *1PATH-PART *1SERVICE-PART *1WILDCARD
DOMAIN-PART = ":domain=" 1*DOMAIN
NODE-PART = ":node=" 1*PART-CHAR
PORT-PART = ":port=" 1*PART-CHAR
LINK-PART = ":link=" 1*PART-CHAR
PATH-PART = ":path=" 1*PART-CHAR
SERVICE-PART = ":service=" 1*PART-CHAR
DOMAIN = 1*(ALPHA / DIGIT / "-" / ".") ; Domain name
PART-CHAR = (ALPHA / DIGIT / "-" / "." / "/" / "_")
WILDCARD = ":*" ; Used for queries.
```

For example:

```
urn:ogf:network:domain=example.net
urn:ogf:network:domain=example.net:node=packrat
urn:ogf:network:domain=example.net:link=WASH_to_ATLA
urn:ogf:network:domain=example.net:node=packrat:port=eth0
urn:ogf:network:domain=example.net:port=Interface_To_Geant
urn:ogf:network:domain=example.net:node=packrat:service=Optical_Converter
urn:ogf:network:domain=example.net:node=packrat:port=eth0:link=WASH_to_ATLA
urn:ogf:network:domain=example.net:node=AMS:port=3/1:link=AMS-GEN
urn:ogf:network:domain=example.net:path=IN2P3_Circuit
urn:ogf:network:domain=example.net:node=packrat:*
```

The syntax is described in [perfSONAR-URN].

Identifiers described by the perfSONAR topology service are **not** valid network resource URNs. Note that the PS-URN syntax allows a slash in a URN, even though this is not allowed by [RFC 2141].

The meaning of the perfSONAR URNs is fundamentally different from network resource URNs: whereas perfSONAR URNs should specifically be parsed to find properties of the resource, this is not allowed for network resource URNs.

This document does not define a specific migration strategy for perfSONAR URNs.

8.3 AutoGOLE Community

Historically the AutoGOLE community used an invalid variant of the network resource URN. AutoGOLE is a proof-of-concept architecture where over ten organisations provide a persistent testbed to show their ability to perform automatic network provisioning across network domains. The resources used in that demo follow the following syntax.

```
AUTOGOLE-URN = "urn:ogf:network:" TYPE ":" NETWORK *1LOCAL-PART
TYPE = ("stp" / "nsa" / "nsnetwork")
NETWORK = 1*(ALPHA / DIGIT / "-" / ".") ; Human readable string
LOCAL-PART = ":" 1*(ALPHA / DIGIT / "-" / ".")
```

No formal publication has been made to describe this syntax.

The historic AutoGOLE Identifiers are **not** valid network resource URNs. A drawback of these AutoGOLE identifiers is that they use a custom name to identify networks, and subsequently the organisational identifiers of the assigning organisations. Deploying this syntax on a large scale would require the set up of a namespace registry.

The AutoGOLE community is currently in the process of adopting the valid network resource URNs.

8.4 Backwards Compatibility

Applications that wish to be backward compatible with the GLIF-based, PerfSONAR-based and AutoGOLE-based URNS, are recommended to accept:

```
BC-NURN = "urn:ogf:network:" 1*(ALPHA / DIGIT / OTHER / "/")
```

Applications that decide to be liberal in the URN that they accept **MUST** anticipate that other clients may do a more thorough syntax check and reject these URNs. In particular, the slash is formally not allowed in URNs.

Applications that merely accept URNs according to the BC-NURN syntax can still be compatible with this specification. However, as soon as a possibility exists that the application sends out URNs that do not comply to NURN syntax, then the application is no longer compatible with the specification described in this document.

9 Contributors

Freek Dijkstra

SARA

Science Park 121

1098 XG Amsterdam

The Netherlands

Email: Freek.Dijkstra@sara.nl

Jeroen van der Ham

Faculty of Science, Informatics Institute, University of Amsterdam

Science Park 904, 1098 XH Amsterdam

The Netherlands

Email: vdham@uva.nl

Acknowledgments

The authors would like to thank the following people (in arbitrary order).

Jason Zurawski for also initiating this work in the network markup language (NML) working group.

Jens Jensen, Richard Hughes-Jones, Greg Newby, Joel Replegle, and Alan Sill for their help in establishing the `urn:ogf` namespace.

Aaron Brown and others at Internet2 for defining the `urn:ogf:network` namespace specification in the perfSONAR community and the network measurement (NM) working group.

Lars Fischer, Tom Lehman, Ronald van der Pol, and Thomas Tam for defining the `urn:ogf:network` namespace specification in the GLIF community.

Björn Höhrmann and Kadir Karaca Koçer of the urn and urn-nid mailing lists at the IETF for useful advice on long term requirements for URNs.

Intellectual Property Statement

The OGF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any

such rights. Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the OGF Secretariat.

The OGF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this recommendation. Please address the information to the OGF Executive Director.

Disclaimer

This document and the information contained herein is provided on an “As Is” basis and the OGF disclaims all warranties, express or implied, including but not limited to any warranty that the use of the information herein will not infringe any rights or any implied warranties of merchantability or fitness for a particular purpose.

Full Copyright Notice

Copyright © Open Grid Forum (2011-2012). Some Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included as references to the derived portions on all such copies and derivative works. The published OGF document from which such works are derived, however, may not be modified in any way, such as by removing the copyright notice or references to the OGF or other organizations, except as needed for the purpose of developing new or updated OGF documents in conformance with the procedures defined in the OGF Document Process, or as required to translate it into languages other than English. OGF, with the approval of its board, may remove this restriction for inclusion of OGF document content for the purpose of producing standards in cooperation with other international standards bodies.

The limited permissions granted above are perpetual and will not be revoked by the OGF or its successors or assignees.

References

- [GFD.191] Freek Dijkstra, Richard Hughes-Jones, Gregory B. Newby and Joel Replogle. Procedure for Registration of Subnamespace Identifiers in the URN:OGF Hierarchy. December 2011. URL <http://www.ogf.org/documents/GFD.191.pdf>.
- [GLIF-ID] Lars Fischer, Tom Lehman, Ronald van der Pol and Thomas Tam. Global Light-path Identifiers Naming Scheme. Version 2.2. April 2009. URL <http://www.glif.is/working-groups/tech/global-identifiers-v2.2.pdf>.
- [perfSONAR-URN] Aaron Brown (editor). A short description of the URN scheme. Retrieved February 2010. URL <http://code.google.com/p/perfsonar-ps/wiki/URNs>.
- [RFC 2119] Scott Bradner. Key words for use in RFCs to Indicate Requirement Levels. RFC 2119 (Best Current Practice), March 1997. URL <http://tools.ietf.org/html/rfc2119>.
- [RFC 2141] Ryan Moats. URN Syntax. RFC 2141 (Standards Track), May 1997. URL <http://tools.ietf.org/html/rfc2141>.
- [RFC 3406] Leslie L. Daigle, Dirk-Willem van Gulik, Renato Iannella and Patrik Fältström. Uniform Resource Names (URN) Namespace Definition Mechanisms. RFC 3406 (Best Current Practice), October 2002. URL <http://tools.ietf.org/html/rfc3406>.
- [RFC 3986] Tim Berners-Lee, Roy T. Fielding, and Larry Masinter. Uniform Resource Identifier (URI): Generic Syntax. RFC 3986 (Standards Track), January 2005. URL <http://tools.ietf.org/html/rfc3986>.
- [RFC 5234] Dave Crocker and Paul Overell. Augmented BNF for Syntax Specifications: ABNF. RFC 5234 (Standards Track), January 2008. URL <http://tools.ietf.org/html/rfc5234>.
- [RFC 5890] John C. Klensin. Internationalized Domain Names in Applications (IDNA): Definitions and Document Framework. RFC 5890 (Standards Track), August 2010. URL <http://tools.ietf.org/html/rfc5890>.