Network Working Group Internet-Draft Expires: July 27, 2008 Intended Status: Informational

Mechanisms for use in pointing to overlay networks, nodes, or resources <u>draft-hardie-p2poverlay-pointers-00.txt</u>

## Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with <u>Section 6 of BCP 79</u>.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <a href="http://www.ietf.org/ietf/lid-abstracts.txt">http://www.ietf.org/ietf/lid-abstracts.txt</a>.

The list of Internet-Draft Shadow Directories can be accessed at <a href="http://www.ietf.org/shadow.html">http://www.ietf.org/shadow.html</a>.

This Internet-Draft will expire on July 27, 2008.

## Copyright Notice

Copyright (C) The IETF Trust (2008).

## Abstract

Discovering overlay networks and the resources found within in them presents a number of bootstrapping problems. While those hard problems are under discussion, this draft proposes a small set of mechanisms which are intended to be generically useful for providing pointers to peer-to-peer overlay networks in web pages, email messages, and other textual media. While the mechanisms described below each meet similar needs, they are not mutually exclusive; it is expected that each will find some useful deployment during the early days of peer-to-peer overlay deployment.

#### 1. Requirements notation

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 [<u>RFC 2119</u>].

## 2. Introduction

This document proposes mechanisms for providing pointers to peer-to-peer overlay networks and resources. These mechanisms are intended to be useful in textual media (web pages, email messages, etc.). While it is commonly true that peer-to-peer networks avoid external dependencies on the DNS or other infrastructure, these mechanisms are intended to be useful in contexts where that infrastructure is present but no connection to a specific overlay has yet been made. These mechanisms are intended to be useable, in other words, as external pointers to specific overlays, their nodes, and their resources.

This is not meant to imply that infrastructure like the global DNS would always be required. IP addresses from a local address realm or resolution services from within another overlay might, for example, be used instead of the global DNS. They may also be used after a connection has been made to a specific overlay to point to particular resources or nodes.

## **<u>3</u>**. Overlay description media type.

For email, the web, and other textual media which might carry pointers to overlay networks, one basic mechanism for providing pointers is to use existing protocols and URI schemes to dereference a resource which contains sufficient information to contact an overlay. For this to be maximally effective, a media type which has an organized method for presenting this data is needed. With such a media type, the pointer can be provided using existing URI schemes, e.g. http://example.org/overlay-pointer.odd. A client can use that URI to retrieve the resource. Once it has the resource, the client can use the structured information it contains to access the overlay. A multipart MIME type could also contain this MIME type, so that it could be carried by applications for which multipart MIME is common practice. Below is a registration of an xml-based MIME type intended for this, application/overlay-pointer+xml, along with a namespace registration, and a schema registration in RelaxNG.

The schema as set forth below contains only a single top-level container, called "availableOverlayDetails", with the following elements: ianaName (which may optionally note an ianaAlgorithm); enrollmentServerContact (which may contain one or more dnsName, ipv6Address, or ipv4Address elements); availableServices (which contains a list of serviceName elements); a joinPolicy (which contains a list of one or more restrictionsImposed); and a

Hardie Expires July 27, 2008 [Page 2]

pointerCreationDate. The "ianaName" and "ianaAlgorithm" are intended to be tokens from the registry set out in <u>Section 6.1</u>; most of the rest are either the standard data types for XML constructs of their type or simply text. Clearly, many other additions or choices are possible here, and the author expects robust discussion to inform the final schema.

## **<u>3.1</u>** Registration of media type application/overlay-pointer+xml.

Content-type registration for 'application/overlay-pointer+xml'

This specification requests the registration of a new MIME type according to the procedures of RFC 4288 [7] and guidelines in RFC 3023 [5].

MIME media type name: application

MIME subtype name: overlay-pointer+xml

Mandatory parameters: none

Optional parameters: charset

Indicates the character encoding of enclosed XML.

- Encoding considerations: Uses XML, which can employ 8-bit characters, depending on the character encoding used. See <u>RFC</u> <u>3023</u> [5], <u>Section 3.2</u>.
- Security considerations: This content type is designed to carry structured information about overlay networks. In cases where that information should be confidential or subject to access control, it should be protected with mechanisms appropriate to the using protocol. Those might include use of a transport which provides confidentiality, object encryption, or some combination.

Interoperability considerations: None

- Published specification: RFCXXXX [NOTE TO IANA/RFC-EDITOR: Please replace XXXX with the RFC number of this specification.]
- Applications which use this media type: Applications providing pointers to overlay networks.

Additional information:

Magic Number: None

File Extension: .odd

Macintosh file type code: 'TEXT'

Personal and email address for further information: Ted Hardie hardie@qualcomm.com

Intended usage: LIMITED USE

Author:

This specification is related to the work of the P2PSIP working group, but is an individual submission.

Change controller:

The IESG <iesg@ietf.org>

## 3.2 Schema namespace registration.

```
Overlay Pointer Registration
```

URI: urn:ietf:params:xml:ns:overlaypointer1

Registrant Contact: Ted Hardie.

XML:

## BEGIN

```
<?xml version="1.0"?>
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML Basic 1.0//EN"
  "http://www.w3.org/TR/xhtml-basic/xhtml-basic10.dtd">
<html xmlns="http://www.w3.org/1999/xhtml">
<head>
  <meta http-equiv="content-type"
        content="text/html;charset=iso-8859-1"/>
 <title>Overlay Pointer Namespace</title>
</head>
<body>
  <h1>Namespace for OverlayPointer</h1>
 <h2>urn:ietf:params:xml:ns:overlaypointer1</h2>
See <a href="[URL of published RFC]">RFCXXXX
    [NOTE TO IANA/RFC-EDITOR:
    Please replace XXXX with the RFC number of this
    specification.]</a>.
</body>
</html>
```

Hardie Expires July 27, 2008 [Page 4]

```
URI: urn:ietf:params:xml:ns:overlaypointer1
 Registrant Contact: Ted Hardie
    (hardie@qualcomm.com).
 Relax NG Schema:
<?xml version="1.0" encoding="UTF-8"?> <grammar
 ns="urn:ietf:params:xml:ns:overlaypointer1"
 xmlns="http://relaxng.org/ns/structure/1.0"
 xmlns:a="http://relaxng.org/ns/compatibility/annotations/1.0"
 datatypeLibrary="http://www.w3.org/2001/XMLSchema-datatypes">
  <div>
 <a:documentation> Overlay Pointer Description
     </a:documentation>
     <define name="availableOverlayDetails">
   element name="availableOverlayDetails"{
      element typeName {
              attribute ianaName {text},
              attribute ianaAlgorithm {text}?
              }
      element enrollmentServerContact {
              attribute dnsName {hostname} |
              attribute ipv6Address {text} |
              attribute ipv4Address {IP}
              }*
      element availableServices {
              attribute serviceName {text}*
              }
      element joinPolicy {
              attribute restrictionsImposed {text}*
      element pointerCreationDate {date}
       }
  </div>
</grammar>
 [EDITOR's Note: the "text" type of ipv6Address is a hack, because
  IP seemed to be limited to IPv4. A better data type should
  clearly be used here. There are also two different forms for
  RelaxNG schemas, and input on which one to use as the normative
```

Hardie Expires July 27, 2008 [Page 5]

version is requested.]

#### 4. Overlay pointer URI.

The use of an existing protocol to dereference a descriptive document has the advantage that it may use deployed protocols and URI schemes. The obvious disadvantage is that it introduces a layer of indirection and may introduce delay in the beginning of overlay-specific protocol processing or may frustrate the protocol processing where the dereferencing fails. A URI scheme specifically for overlay pointers is proposed (as a provisional URI registration) to provide for a direct indicator. While this may require configuration (associating the URI scheme with a handler that invokes the overlay processing), that configuration avoids the layer of indirection mentioned above.

#### **<u>4.1</u>** Registration template for an overlay URI.

URI scheme semantics.

The authority section of the URI contains a hostname or IP address associated with an enrollment server or introducer for the overlay. It may also contain a port on which enrollment services are running. The otype, or overlay type, indicates the registered type for the overlay (e.g., Pastry); these are tokens registered by IANA, in the "Peer to Peer Overlay Network types" registry. The service parameters (if present), indicate the services that an overlay offers. In the following example:

```
overlay://enrollment.example.org/;otype=Pastry;service=mass-storage
```

the URI provides a pointer to an enrollment server for an overlay running the Pastry DHT and providing a service of "mass-storage". While it is assumed that some uses of this URI might create agreements for the meaning of specific services (e.g. p2psip might create an "ICE" service), the current registration treats these as free text so that other users can mint new ones as needed. If discussion during the provisional phase indicates the chance of confusion among these, a parameter registry would be created as part of the transition from provisional to permanent registration.

Hardie Expires July 27, 2008 [Page 6]

Encoding considerations. No special considerations

Applications/protocols that use this URI scheme name.

Applications which need to provide a protocol pointer to an overlay network's enrollment servers or introducers.

Interoperability considerations. None known.

Security considerations.

As currently constructed, this URI scheme's authority section is expected to contain a hostname or IPv6 address . It would be possible to have SRV records or a DDDS application choose entry points based on the authority's DNS name instead. That would provide a better chance that a DoS attack against a specific introducer or enrollment server could not eliminate the ability of new nodes to join an overlay. It does, however, create another layer of indirection and make the use of an IP address in the authority section problematic; in this instance, the ability of someone controlling the zone to add or update the records associated with a server instance was judged a better fit for the problem space.

Contact.

Ted Hardie <hardie@qualcomm.com>

Author/Change controller.

Ted Hardie <hardie@qualcomm.com>

References.

draft-hardie-p2poverlay-pointers-00.txt
<u>RFC 3986</u>

## 5. Overlay node pointer URI.

Among the bootstrapping problems presented by peer to peer overlays is the lack of a generic way to represent nodes within an overlay, resources stored at that node or resources stored within the overlay. A URI scheme focused on the node and its resources is proposed here, along with context-dependent ways to use it that allow for it to represent resources stored within an overlay. This URI scheme registration is intended to be provisional. It contains a significant limitation that deserves to be highlighted: although the typical "host" portion of an authority section for a URI allows DNS names, IP addresses, or a registered name, this proposal limits the authority section to registered names which are current node identifiers for a specific overlay.

A second point to note is that the absence of an authority section indicates that the resource noted in the query section is somewhere within the overlay, but the URI does not establish at which node-id. Where the URI contains a pointer to the overlay context, this provides a mechanism to give an external reference to a resource within a specific overlay without referencing a node-id. Within the context of a specific overlay, this would allow the overlay client to invoke overlay-specific search mechanisms to establish one or more appropriate node-ids offering the service or hosting the resource (and thus to construct "full" pointer URIs).

A basic example of the node-id use is:

overlay-node://22301203/?resource=example.iso

The authority points to a specific node-id; the query section points to a resource stored at that node. It is, however, valid only within a context that already understands what overlay that node occurs in. In order to create a context that establishes that, this registration re-uses the methods discussed above to set the context.

The following two examples are presented without percent-encoding to highlight the relation to the sections above, but percent-encoding of the context-setting URIs would be required. See <u>Section 5.1</u> for the actual syntax.

Note that the following lines use  $\$  to indicate that the full URI is carried across two lines. In this example, the context is set with reference to a specific overlay-pointer+xml resource:

?resource=service-instance

In this example, the context is set with a reference to a specific "overlay" URI:

overlay-node://22301203/;context="overlay://enrollment.example.org/\
;otype=pastry"\
?resource=example.iso

In this example, the context is set with reference to a specific "overlay" URI, but no node ID is given. This is how you would construct a URI for "ICE services, offered within a specific overlay":

overlay-node:///;context="overlay://introducer.example.org/;otype=pastry"\
?resource=ICE-services

Hardie Expires July 27, 2008

5.1 Registration template for an overlay node pointer URI.

```
URI scheme name.
   "overlay-node" is requested
Status.
  provisional
URI scheme syntax.
    overlay-node-URI = overlay-node-scheme "://" [overlay-authority]
    "/" [";"" context=" overlay-context-uri] ["?" query] ["#"fragment]
              ; query and fragment are as defined in rfc 3986
overlay-scheme = "overlay-node"
overlay-authority= [ userinfo "@" ] reg-name [ ":" port ]
          ;reg-name is defined in <u>RFC 3986</u> *note limitation from host*
overlay-context-uri = HTTP-URI | overlay-URI
    ;HTTP-URI is defined in RFC 2616
    ;overlay-URI in <u>draft-hardie-p2poverlay-pointers-00.txt</u>
otype = "otype=" token
               ;valid tokens are in the "Peer to Peer
               ;Overlay Network types" IANA registry
service="service=" *1ALPHANUM
URI scheme semantics.
   See section 5 of draft-hardie-p2poverlay-pointers-00.txt.
Encoding considerations.
   No special considerations
Applications/protocols that use this URI scheme name.
   Applications which need to provide a protocol pointer to an
   overlay network node, its resources, or resources stored within
   a specific overlay network.
Interoperability considerations.
   None known.
Security considerations.
   The authority section contains a node-id valid within a specific
   overlay. Global context is not required; if present, it is
   given using the "context" parameter. Where it is not present
   and the context is incorrect, it is possible that the effort to
   retrieve a resource will fail or will return incorrect results.
   Careful naming of resources within an overlay may mitigate this
   problem, but any security system must be aware that overlay-node
   URIs without a context parameter have different characteristics
```

from URIs that fit in within a known global context like the

Contact. Ted Hardie <hardie@qualcomm.com>

DNS.

Author/Change controller. Ted Hardie <hardie@qualcomm.com>

References. <u>draft-hardie-p2poverlay-pointers-00.txt</u> RFC 3986

# <u>6</u>. IANA Considerations

This document registers a new media type, an XML schema, two provisional URI schemes, and creates a registry for peer to peer overlay types. The media type registration is in <u>section 3.1</u>. The XML schema is in <u>section 3.2</u>. The first URI scheme registration is in <u>section 4.1</u>; the second is in <u>section 5.1</u>. The registry for peer-to-peer overlay network types is in <u>section 6.1</u>, below.

Hardie Expires July 27, 2008 [Page 9]

6.1 Peer-to-peer overlay network type registry.

IANA is requested to create a registry titled "Peer-to-Peer Overlay Network Types". The Registry shall have the following fields: Type Name, Algorithm Type, Record Creator, Record Creator contact, Documentation URI. An example is given below:

Type Name: Pastry Algorithm Type: Distributed Hash Table Record Creator: IESG Record Creator contact: iesg@ietf.org Docmentation URI: <u>http://freepastry.org/</u>

The registry is to permit new entries using the "Expert Review" guidelines as described in <u>RFC 2434[RFC2434]</u>. The Expert Reviewer is requested to pay particular attention to the Algorithm Type field and to limit the creation of new values for that field where the algorithms are variants of a fundamental type. Since the main purpose of the registry is to enable clients to determine their interoperability with a specific mechanism, however, the Expert Reviewer should not limit the creation of new Type Names, except where the documentation provided either clearly indicates full interoperability with an existing Type Name of is of insufficient completeness to make any determination on interoperability. The Record Creator contact field SHOULD contain at least an email address and MAY contain any other contact data desired by the Record Creator.

#### 7. Security Considerations

Security considerations for each of the three methods given above is documented within each method. The general security issue here is that providing a pointer signals a point at which the overlay may be touched or resource retrieved; disclosure of that indicates either a point of attack, where a specific resource resides, or both.

For overlay networks concerned with chosen location attacks, disclosure of a service or resources at a particular node-id may be problematic, as it might assist the attacker in choosing a location to attack. For an attacker with access to multiple clients or the ability to mint new identities, this is not a large advantage, as the attacker could join the overlay, collect the same information, and then re-join.

## 8. Acknowledgments.

Vidya Narayanan, Lakshminath Dondeti, and Spencer Dawkins commented on early versions of this document.

[Page 10]

## 9. References

## 9.1 Normative References

[KeyWords] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", <u>RFC 2119</u>, <u>BCP 14</u>, March 1997.

[URI] Berners-Lee T. et al, "URI Generic Syntax", <u>RFC 3986</u>, STD 66, January 2005.

[HTTP] Fielding, R. et al, "Hypertext Transfer Protocol --HTTP/1.1", <u>RFC 2616</u> June 1999.

[URI-REG] Hansen, T. et al. "Guidelines and Registration Procedures for New URI Schemes", <u>RFC 4395</u>, <u>BCP 115</u>, February 2006.

[ABNF] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", <u>RFC 4234</u>, October 2005.

## 9.2 Informative References

Author's Addresses

Ted Hardie Qualcomm, Inc. Email: hardie@qualcomm.com

#### Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights 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; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in <u>BCP 78</u> and <u>BCP 79</u>.

Copies of IPR disclosures made to the IETF Secretariat 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 IETF on-line IPR repository at http://www.ietf.org/ipr.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in  $\frac{BCP}{78}$ , and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM 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.

## Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.