

Network Working Group  
Internet-Draft  
Expires: September 2, 2006

P. Nikander  
Ericsson Research Nomadic Lab  
J. Laganier  
DoCoMo Euro-Labs  
F. Dupont  
CELAR  
March 1, 2006

**An IPv6 Prefix for Overlay Routable Cryptographic Hash Identifiers  
(ORCHID)  
draft-laganier-ipv6-khi-01**

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 [Section 6 of BCP 79](#). This document may not be modified, and derivative works of it may not be created, except to publish it as an RFC and to translate it into languages other than English.

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  
<http://www.ietf.org/ietf/1id-abstracts.txt>.

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

This Internet-Draft will expire on September 2, 2006.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

This document introduces Overlay Routable Cryptographic Hash

Identifiers (ORCHID) as a new, experimental class of IPv6-address-like identifiers. These identifiers are intended to be used as end-point identifiers at applications and APIs and not as identifiers for network location at the IP layer, i.e., locators. They are designed to appear as application layer entities and at the existing IPv6 APIs, but they should not appear in actual IPv6 headers. To make them more like vanilla IPv6 addresses, they are expected to be routable at an overlay level. Consequently, while they are considered as non-routable addresses from the IPv6 layer point of view, all existing IPv6 applications are expected to be able to use them in a manner compatible with current IPv6 addresses.

This document requests IANA to allocate a temporary prefix out of the IPv6 addressing space for Overlay Routable Cryptographic Hash Identifiers.

## **1. Introduction**

This document introduces Overlay Routable Cryptographic Hash Identifiers (ORCHID), a new class of IP-address-like identifiers. These identifiers are intended to be globally unique in a statistical sense (see [Section 4](#)), non-routable at the IP layer, and routable at some overlay layer. The identifiers are securely bound, via a secure hash function, to the concatenation of an input bitstring and a context tag. Typically, but not necessarily, the input bitstring will include a suitably encoded public cryptographic key.

### **1.1. Rationale and intent**

These identifiers are expected to be used at the existing IPv6 APIs and application protocols between consenting hosts. They may be defined and used in different contexts, suitable for different overlay protocols. Examples of these include Host Identity Tags (HIT) in the Host Identity Protocol (HIP) [[I-D.ietf-hip-base](#)] and Temporary Mobile Identifiers (TMI) for Mobile IPv6 Privacy Extension [[I-D.dupont-mip6-privacyext](#)].

As these identifiers are expected to be used alongside with IPv6 addresses at both applications and APIs, co-ordination is desired to make sure that an ORCHID is not inappropriately taken for a vanilla IPv6 address and vice versa. In practice, allocation of a separate prefix for ORCHIDs seems to suffice, making them compatible with IPv6 addresses at the upper layers while simultaneously making it trivial to prevent their usage at the IP layer.

While being technically possible to use ORCHIDs between consenting hosts without any co-ordination with the IETF and the IANA, the



authors would consider such practice potentially dangerous. A specific danger would be realised if the IETF community later decided to use the ORCHID prefix for some different purpose. In that case, hosts using the ORCHID prefix would be, for practical purposes, unable to use the prefix for the other, new purpose. That would lead to partial balkanisation of the Internet, similar to what has happened as a result of historical hijackings of non-RFC1918 IPv4 addresses for private use.

The whole need for the proposed allocation grows from the desire to be able to use ORCHIDs with existing applications and APIs. This desire leads to the potential conflict, mentioned above. Resolving the conflict requires the proposed allocation.

One can argue that the desire to use these kinds of identifiers via existing APIs is architecturally wrong, and there is some truth in that argument. Indeed, it would be more desirable to introduce a new API and update all applications to use identifiers, rather than locators, via that new API. That is exactly what we expect to happen in the longer run.

However, given the current state of the Internet, we do not consider it viable to introduce any changes that, at once, require applications to be rewritten and host stacks to be updated. Rather than that, we believe in piece-wise architectural changes that require only one of the existing assets to be touched. ORCHIDs are designed to address this situation: to allow people to experiment with protocol stack extensions, such as secure overlay routing, HIP, or Mobile IP privacy extensions, without requiring them to update their applications. The goal is to facilitate large-scale experiments with minimum user effort.

For example, there already exists, at the time of this writing, HIP implementations that run fully in user space, using the operating system to divert a certain part of the IPv6 address space to a user level daemon for HIP processing. In practical terms, those implementations are already now using a certain IPv6 prefix for differentiating HIP identifiers from IPv6 addresses, allowing them both to be used by the existing applications via the existing APIs.

This document argues for no more than allocating an experimental prefix for such purposes, thereby paving the way for large-scale experiments with cryptographic identifiers without the dangers caused by address-space hijacking.

## **1.2. ORCHID properties**

ORCHIDs are designed to have the following properties:



- o Statistically uniqueness; see also [Section 4](#)
- o Secure binding to the input parameters used in their generation (i.e., the context identifier and a bitstring.)
- o Conformance with the IPv6 global unicast address format as defined in [Section 2.5.4 of \[RFC3513\]](#).
- o Aggregation under a single IPv6 prefix. Note that this is only needed due to the co-ordination need, as indicated above. Without such co-ordination need, the ORCHID name space could potentially be completely flat.
- o Non-routability at the IP layer, by design.
- o Routability at some overlay layer, making them, from an application point of view, semantically similar to IPv6 addresses.

As mentioned above, ORCHIDs are intended to be generated and used in different contexts, as suitable for different mechanisms and protocols. The context identifier is meant to be used to differentiate between the different contexts; see [Section 4](#) for a discussion of the related API and kernel level implementation issues, and [Section 5](#) for the design choices explaining why the context identifiers are used.

### **[1.3.](#) Expected use of ORCHIDs**

Examples of identifiers and protocols that are expected to adopt the ORCHID format include Host Identity Tags (HIT) in the Host Identity Protocol [[I-D.ietf-hip-base](#)] and the Temporary Mobile Identifiers (TMI) in the Simple Privacy Extension for Mobile IPv6 [[I-D.dupont-mip6-privacyext](#)]. The format is designed to be extensible to allow other experimental proposals to share the same name space.

### **[1.4.](#) Action plan**

This document requests IANA to allocate an experimental prefix out of the IPv6 addressing space for Overlay Routable Cryptographic Hash Identifiers.

## **[2.](#) Cryptographic Hash Identifier Construction**

An ORCHID is generated using the algorithm below. The algorithm takes a bitstring and a context identifier as input and produces an ORCHID as output.



Input        := any bitstring  
Hash Input := Context ID | Input  
Hash        := SHA1( Expand( Hash Input ) )  
ORCHID      := Prefix | Encode\_n( Hash )

where:

|            : Denotes concatenation of bitstrings

Input        : A bitstring unique or statistically unique within a given context. The bitstring is intended to be associated with the to-be-created ORCHID, in the given context.

Context ID : A randomly generated value defining the expected usage context for the particular ORCHID.

As a baseline (TO BE DISCUSSED), we propose sharing the name space introduced for CGA Type Tags; see <http://www.iana.org/assignments/cga-message-types> and [RFC 3972](#).

Expand( ) : An expansion function designed to overcome recent attacks on SHA-1.

As a baseline (TO BE DISCUSSED), we propose using the method defined in [[I-D.irtf-cfrg-sha1-ime](#)].

Alternatively, it would be possible to use some other hash function, such as SHA-256, instead of SHA-1.

Encode\_n( ) : An extraction function which output is obtained by extracting an <n>-bits-long bitstring from the argument bitstring.

As a baseline (TO BE DISCUSSED), we propose taking <n> middlemost bits from the SHA1 output.

Prefix      : A constant ( 128 - <n> bits long ) bitstring value, TBD, assigned by IANA.

To form an ORCHID, two pieces of input data are needed. The first piece can be any bitstring, but is typically expected to contain a public cryptographic key and some other data. The second piece is a context identifier, which is an 128-bits-long datum, allocated as specified in [Section 7](#). Each specific experiment (such as HIP HITs or MIP6 TMIs) is expected to allocate their own, specific context





identifier.

The input bitstring and context identifier are concatenated to form an input datum, which is then fed to a cryptographic hash function. The result of the hash function is processed by an encoding function, resulting in an n-bits-long value. This value is prepended with the ORCHID prefix. The result is the ORCHID, an 128-bits-long bitstring that can be used at the IPv6 APIs in hosts participating to the particular experiment.

### **3. Routing Considerations**

ORCHIDs are designed to serve as location independent end-point-identifiers rather than IP-layer locators. Therefore, routers MAY be configured not to forward any packets containing an ORCHID as a source or a destination address. If the destination address is a ORCHID but the source address is a valid unicast source address, routers MAY be configured to generate an ICMP Destination Unreachable, Administratively Prohibited message.

Due to the experimental nature of ORCHIDs, router software MUST NOT include any special handling code for ORCHIDs. In other words, the non-routability property of ORCHIDs, if implemented, MUST be implemented via configuration and NOT by hard-wired software code. At this time, it is RECOMMENDED that the default router configuration does not handle ORCHIDs in any special way. In other words, there is no need to touch existing or new routers due to this experiment. If such reason should later appear, for example, due to a faulty implementation leaking ORCHIDs to the IP layer, the prefix can be and should be blocked by a simple configuration rule.

#### **3.1. Overlay Routing**

As mentioned multiple times, ORCHIDs are designed to be non-routable at the IP layer. However, there are multiple ongoing research efforts for creating various overlay routing and resolution mechanisms for flat identifiers. For example, the Host Identity Indirection Infrastructure (Hi3) [[hi3](#)] proposal outlines a way for using a Distributed Hash Table to forward HIP packets based on the Host Identity Tag.

What is common to the various research proposals is that they create a new kind of resolution or routing infrastructure on the top of the existing Internet routing structure. In practical terms, they allow delivery of packets based on flat, non-routable identifiers, utilising information stored in a distributed data base. Usually the database used is based on Distributed Hash Tables. This effectively



creates a new routing network on the top of the existing IP-based routing network, capable of routing packets that are not addressed by IP addresses but some other kind of identifiers.

Typical benefits from overlay routing include location independence, more scalable multi-cast, any-cast, and multi-homing support than in IP, and better DoS resistance than in the vanilla Internet. The main drawback is typically an order of magnitude slower performance, caused by an easily largish number of extra look-up or forwarding steps needed. Consequently, in most practical cases the overlay routing system is used only during initial protocol state set-up (cf. TCP handshake), after which the communicating end-points exchange packets directly with IP, bypassing the overlay network.

The net result of the typical overlay routing approaches is a communication service whose basic functionality is comparable to that of provided by classical IP but that provides considerably better resilience than vanilla IP in dynamic networking environments. Some experiments also introduce additional functionality, such as enhanced security or ability to effectively route through several IP addressing domains.

The authors expect ORCHIDs to become fully routable, via one or more overlay systems, before the end of the experiment.

#### **4. Collision Considerations**

As noted above, the aim is that ORCHIDs are globally unique in a statistical sense. That is, given the ORCHID referring to a given entity, the probability of the same ORCHID being used to refer to another entity elsewhere in the Internet must be sufficiently low so that it can be ignored for most practical purposes. We believe that the presented design meets this goal; see [Section 5](#).

Consider next the very rare case that some ORCHID happens to refer to two different entities at the same time at two different locations in the Internet. Even in that case the probability of this fact becoming visible (and therefore a matter of consideration) at any single location in the Internet is negligible. For the vast majority of cases the two simultaneous uses of the ORCHID will never cross each other. However, while rare such collisions are still possible. This section gives reasonable guidelines on how to mitigate the consequences in the case such a collision happens.

As mentioned above, ORCHIDs are expected to be used at the legacy IPv6 APIs between consenting hosts. The context ID is intended to differentiate between the various experiments, or contexts, sharing



the ORCHID name space. However, the context ID is not present in the ORCHID itself, but only in front of the input bitstring as an input to the hash function. While this may lead to certain implementation-related complications, we believe that the trade-off of allowing the hash result part of an ORCHID being longer more than pays off the cost.

Now, because ORCHIDs are not routable at the IP layer, in order to send packets using ORCHIDs at the API level, the sending host must have additional overlay state within the stack in order to determine parameters (e.g. what locators) to use in the outgoing packet. An underlying assumption here, and a matter of fact in the proposals that the authors are aware of, is that there is an overlay protocol for setting up and maintaining this additional state. It is assumed that the state-set-up protocol carries the input bitstring, and that the resulting ORCHID-related state in the stack can be associated back with the appropriate context and state-set-up protocol.

Even though ORCHID collisions are expected to be extremely rare, two kinds of collisions may still happen. First, it is possible that two different input bitstrings within the same context may map to the same ORCHID. In that case, the state-set-up mechanism is expected to resolve the conflict, for example, by indicating to the peer that the ORCHID in question is already in use.

A second type of collision may happen if two input bitstrings, used in different usage contexts, map to the same ORCHID. In this case the main confusion is about which context to use. In order to prevent these types of collisions, it is RECOMMENDED that implementations that simultaneously support multiple different contexts maintain a node-wide unified database of known ORCHIDs, and indicate a conflict if any of the mechanisms attempt to register a ORCHID that is already in use. For example, if a given ORCHID is already being used as a HIT in HIP, it cannot simultaneously be used as a TMI in Mobile IP. Instead, if Mobile IP attempts to use the ORCHID, it will be notified (by the kernel) that the ORCHID in question is already in use.

## **5. Design Choices**

The design of this name space faces two competing forces:

- As many bits as possible should be preserved for the hash result.
- It should be possible to share the name space between multiple mechanisms.

The desire to have a long hash result requires the prefix to be as short as possible, and to use few (if any) bits for additional



encoding. The present design takes this desire to the maxim: all the bits beyond the prefix are used as hash output. This leaves no bits in the ORCHID itself available for identifying the context. Additionally, due to security considerations, the present design **REQUIRES** that the hash function used in constructing ORCHIDs be constant; see [Section 6](#).

The authors explicitly considered including a hash extension mechanism, similar to the one in CGA [[RFC3972](#)], but decided to leave it out. There were two reasons: desire for simplicity, and the somewhat unclear IPR situation around the hash extension mechanism. If there is a future revision of this document, we strongly advise the future authors to reconsider the decision.

The desire to allow multiple mechanism to share the name space has been resolved by including the context identifier in the hash function input. While this does not allow the mechanism to be directly inferred from a ORCHID, it allows one to verify that a given input bitstring and ORCHID belong to a given context, with high probability; but see also [Section 6](#).

## **6. Security Considerations**

ORCHIDs are designed to be securely bound to the context identifier and the bitstring used as the input parameters during their generation. To provide this property, the ORCHID generation algorithm relies on the second-preimage resistance (a.k.a. one-way) property of the hash function used in the generation [[RFC4270](#)]. To have this property, and to avoid collisions, it is important that the allocated prefix is as short as possible, leaving as many bits as possible for the hash output.

All mechanism using ORCHIDs **MUST** use exactly the same mechanism for generating a ORCHID from the input bitstring. Allowing different mechanisms, without explicitly encoding the mechanism in the ORCHID itself, would allow so called bidding down attacks. That is, if multiple different hash functions were allowed in constructing ORCHIDs in a given shared name space, and if one of the hash functions became insecure, that would allow attacks against even those ORCHIDs that had been constructed using the other, still secure hash functions.

Due to the desire to keep the hash output value as long as possible, the present design allows only one method for constructing ORCHIDs from input bitstrings. If other methods (perhaps using more secure hash functions) are later needed, they **MUST** use a different prefix. Consequently, the suggested method to react to the hash result





becoming too short, due to increased computational power or to the used hash function becoming insecure due to advances in cryptography, is to allocate a new prefix and cease to use the present one.

As of today, SHA-1 applied in conjunction with a proper expansion function of the hash input is considered as satisfying the second-preimage resistance requirement [[I-D.irtf-cfrg-sha1-ime](#)]. Hash output of at least 100 bits, but preferably up to 120 bits, is considered to have a low enough probability of collisions.

In order to preserve a low enough probability of collisions (see [Section 4](#)), each method MUST utilize a mechanism that makes sure that the distinct input bitstrings are either unique or statistically unique, within that context. There are several possible methods to ensure that; for example, one can include into the input bitstring a globally maintained counter value, a pseudo-random number of sufficient entropy (minimum 120 bits), or a randomly generated public cryptographic key. The Context ID makes sure that input bitstrings from different contexts never overlap. These together make sure that the probability of collisions is determined only by the probability of natural collisions in the hash space and is not increased by a possibility of colliding input bit strings.

## 7. IANA Considerations

IANA is requested to allocate a temporary non-routable prefix from the IPv6 address space. As per Sections [2.5.1](#) and [2.5.4](#) of [[RFC3513](#)], the prefix must be allocated from the 0000::/3 block, since ORCHIDs do not have a 64-bit interface identifier part. The allocation will require updating <http://www.iana.org/assignments/ipv6-address-space>

As a baseline (TO BE DISCUSSED), we propose an 8-bit prefix to be allocated from the 1000::/4 block. During the discussions related to this draft, it was suggested that other identifier spaces may be later allocated from this block. However, this document does not define such a policy or allocations.

The Context Identifier (or Context ID) is a randomly generated value defining the usage context of a ORCHID. This document defines no specific value.

As a baseline (TO BE DISCUSSED), we propose sharing the name space introduced for CGA Type Tags. Hence, defining new values would follow the rules of [Section 8 of \[RFC3972\]](#), i.e., on a First Come First Served basis. The policy will require updating the policy for <http://www.iana.org/assignments/cga-message-types>



## **8. Acknowledgments**

Julien Laganier is partly funded by Ambient Networks, a research project supported by the European Commission under its Sixth Framework Program.

Special thanks to Geoff Huston for his sharp but constructive critic during the development of this memo. Tom Henderson helped to clarify a number of issues.

## **9. Version history**

### **9.1. -00 to -01**

The name Keyed Hash Identifier (KHI) was replaced with Overlay Routable Cryptographic Hash Identifier (ORCHID). However, the draft name was not changed.

More text added to explain the rationale behind the proposed allocation.

Text changed to emphasise that while ORCHIDs are expected to be non-routable at the IP-layer, they are expected to become fully routable and/or resolvable at some upper, overlay layer, thereby making their basic semantics fully compatible with IPv6 addresses.

Removed the proposed expiration date. If such an expiration date is needed, it can be added later during the discussions.

## **10. References**

### **10.1. Normative references**

- [I-D.irtf-cfrg-sha1-ime] Blumenthal, U., Jutla, C., and A. Patthak, "SHA1-IME: A SHA-1 Variant with Provably Good Message Expansion Code", November 2005.
- [RFC3513] Hinden, R. and S. Deering, "Internet Protocol Version 6 (IPv6) Addressing Architecture", [RFC 3513](#), April 2003.
- [RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", [RFC 3972](#), March 2005.



## **10.2.   Informative references**

- [I-D.dupont-mip6-privacyext]  
Dupont, F., "A Simple Privacy Extension for Mobile IPv6",  
[draft-dupont-mip6-privacyext-03](#) (work in progress),  
January 2006.
  
- [I-D.ietf-hip-base]  
Moskowitz, R., "Host Identity Protocol",  
[draft-ietf-hip-base-00](#) (work in progress), June 2004.
  
- [RFC3587] Hinden, R., Deering, S., and E. Nordmark, "IPv6 Global  
Unicast Address Format", [RFC 3587](#), August 2003.
  
- [RFC4270] Hoffman, P. and B. Schneier, "Attacks on Cryptographic  
Hashes in Internet Protocols", [RFC 4270](#), November 2005.
  
- [hi3] Nikander, P., Arkko, J., and B. Ohlman, "Host Identity  
Indirection Infrastructure (Hi3)", Nov 2004.



Authors' Addresses

Pekka Nikander  
Ericsson Research Nomadic Lab  
JORVAS FI-02420  
FINLAND

Phone: +358 9 299 1  
Email: pekka.nikander@nomadiclab.com

Julien Laganier  
DoCoMo Communications Laboratories Europe GmbH  
Landsberger Strasse 312  
Munich 80687  
Germany

Phone: +49 89 56824 231  
Email: julien.ietf@laposte.net  
URI: <http://www.docomolab-euro.com/>

Francis Dupont  
CELAR

Email: Francis.Dupont@point6.net





## 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 [BCP 78](#) and [BCP 79](#).

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](mailto:ietf-ipr@ietf.org).

## Disclaimer of Validity

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 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.

## Copyright Statement

Copyright (C) The Internet Society (2006). This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

## Acknowledgment

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

