

Unique Local IPv6 Unicast Addresses

A review from an RIR perspective

Geoff Huston

August 2003

Background

- IPv6 address architecture includes the requirement for local-use addresses that are:
 - Useable in a local (non-connected context)
 - Span more than a link
 - Are not components of a provider aggregate address block
 - Not intended to be globally routed
 - Unique (no NATS!)
 - Unicast addresses

Local Use Addresses

- The IETF IPv6 Working Group is considering alternatives to Site-Local Addresses

(This presentation is not intended to be a repeat of the Site-Local debate!)

- One proposal is to use a block of the Global Unicast Address space for “local” use
 - Where “*local*” implies “*not directly anticipated to be globally routed*”
 - See [draft-hinden-ipv6-global-local-addr-02.txt](#) for the complete text of the proposal

Questions Raised by the Proposal

- See `draft-huston-ipv6-local-use-comments-00.txt`
- What are the desirable characteristics of Local Use addresses?
- What distribution mechanisms are called for?
- Is there a role for the RIRs here?
- If so what issues would this raise for the RIRs to consider?

Characteristics of Local Use Addresses

1. Exclusive use of a common prefix drawn from the global unicast address space for all local use addresses (FC00::/7)
2. Unique assignment of a fixed size local use address block (/48) from within the pool of addresses defined by this prefix, using a Global ID as the block prefix.
3. There is no internal structure within the global ID, and these global IDs cannot be aggregated in a routing context.
4. The assignment information must be recorded and stored in a reliable manner.
5. Local Use Addresses are not intended to be passed within the global routing environment

The Proposal

- Use /48 blocks drawn from FC00::/7

An End user may either:

use a random number pick to draw a /48 block
from FC00::/8

or:

obtain a unique /48 block from a registry that
manages FD00::/8

A Local Use Registry System

A Local Use Registry system should be:

- Readily accessible for anybody
- Highly automated
- No justification required
- Limited identity requirement
- Rapid turnaround
- Inexpensive
- Allocate randomly from the block
- Transparency of charges
- Allow for once-and-forever allocation services
- Allow for agency structures
- Reliable and enduring records of unique allocations
- Limited publication of allocations

RIR Considerations

- Service model choice (renewable, non-renewable)
- Transaction model rather than membership-based
- Service fees to be cost-based
- Record management
- High-volume low-value transaction model
- Preserve Local Use address characteristics (non-aggregatable, no public per-allocation records, stable allocations, non-hoardable)
- Regulatory issues (competition, fee setting, equal access)

RIR Considerations

- This can be seen as a distinct service activity, not a seamless adjunct to existing activities:
 - Transactions, not membership
 - High volume, low value
 - Automated applications without evaluation
 - Limited publication of allocations
- Considerations:
 - Local agency activities?
 - Wholesaling?
 - Transfers?
 - ?

Broader Considerations

- What is the distinction between global local use and global unicast addresses?
 - And who gets to decide which is which?
- What is the benefit of having two classes of addresses?
 - Leased, qualified, provider aggregated
 - Enduring, unqualified, provider independent
- Can these local use addresses really be isolated from the routing system?
 - Is this a repeat of the V4 swamp construction?
 - And if so is this necessarily a bad thing?