IETF 121 – 6lo

Path-Aware Semantic Addressing (PASA) for Low power and Lossy Networks

~~draft-ietf-6lo-path-aware-semantic-addressing-07~~

draft-ietf-6lo-path-aware-semantic-addressing-08

L. Iannone, G. Li, D. Lou, P. Liu, R. Long, K. Makhijani, P. Thubert

IETF 121 – Dublin

Since IETF 120

Path-Aware Semantic Addressing (PASA) for Low power and Lossy Networks

draft-ietf-6lo-path-aware-semantic-addressing-08

Status | IESG evaluation record | IESG writeups | Email expansions | History

Versions:

00 01 02 03 04 05 06 07 08

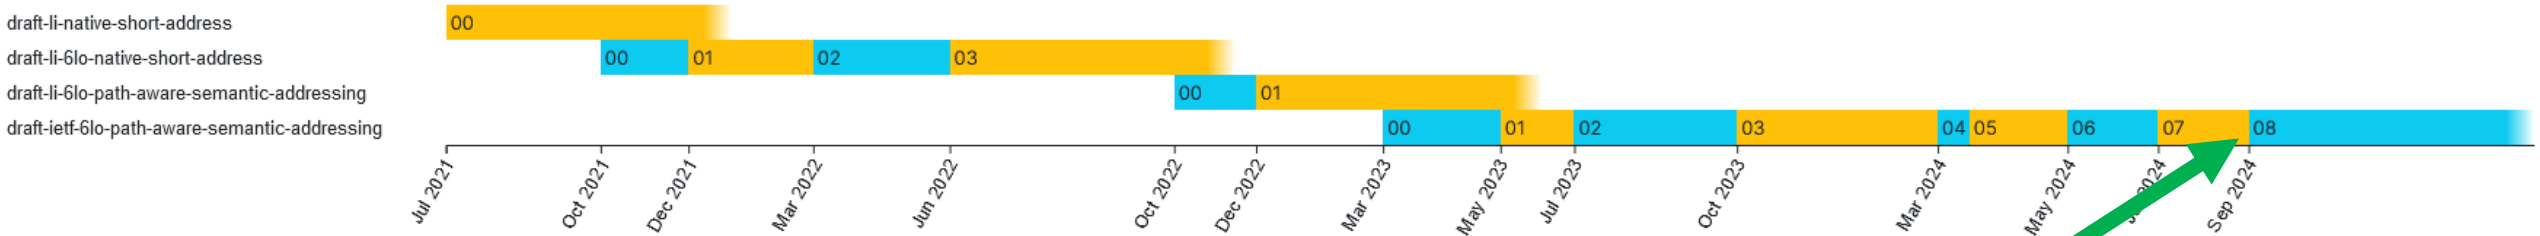

draft-ietf-6lo-path-aware-semantic-address-08.txt
September 2024
Main changes: Address residual issues of RTGDIR review by Joel Halpern

Main Content Changes (I)

No changes in the document structure

- Still done some word-smithing here and there and typo fixes.

Main Content Changes (II)

12. Reliability Considerations

Because PASA uses algorithmically generated addresses based on the network topology, nodes do not generate and store forwarding table entries in the normal case. They are limited to have a default gateway and the ND table. One of the potential issues is the risk of renumbering of addresses in case of topology changes. Because of the applicability domain of PASA, the common case of topology change is known in advance and can be planned, so to reduce disruption due to renumbering (see Section 4).

Another case is temporary link failures or node temporary failures, where the network is still able to provide connectivity through alternative links, which is strictly related to the underlying technology, the network topology, the deployed redundancy, and the expected reliability. Failures may raise the issue of topology changes and re-numbering. Such a issues can be avoided or at least mitigated following the procedures in [RFC8505] and [I-D.iannone-6lo-nd-gaao] keeping state in non-volatile memory.

More complex reliability scenarios, including the case for multiple root nodes and alternative solutions, are beyond the scope of this document, which is focused only on the address allocation framework and stateless forwarding. A more in-depth discussion about reliability, including the case of multiple roots, can be found in [I-D.li-6lo-pasa-reliability]. Furthermore, specific reliability solutions depend as well on the specific Address Assignment Function used (different from the one presented in this document).

12. Reliability Considerations

Because PASA uses algorithmically generated addresses, based on the network topology, nodes do not generate and store forwarding table entries in the normal case. They are limited to have a default gateway and the ND table. One of the potential issues is the risk of renumbering of addresses in case of topology changes. Because of the applicability domain of PASA, the common case of topology change is known in advance and can be planned, so to reduce disruption due to renumbering (see Section 4).

Another case is temporary link failures or node temporary failures, where the network is still able to provide connectivity through alternative links, which is strictly related to the underlying technology, the network topology, the deployed redundancy, and the expected reliability. Failures may raise the issue of topology changes and re-numbering. Such issues can be avoided, or at least mitigated, following the procedures in [RFC8505] and [I-D.iannone-6lo-nd-gaao] keeping state in non-volatile memory.

Reliability of external connectivity, with more than one node functioning as gateway, can be achieved in several ways. One simple solution is to use a multi topology approach, where each gateway acts as a root for a logically independent topology, identified via a different prefix. The multiple topologies can either be used at the same time or with a primary/backup policies. This solution is particularly suitable in case the PASA domain is multihomed.

An alternative solution is to separate root and gateway roles, setting up the topology so that some of the children of the root will also function as gateways, offering external connectivity. In this way traffic destined outside the local PASA domain will still still forwarded using a simple default route toward the root, and then sent outside when they reach one of the root's children or the root itself. This second solution allows to accommodate load balancing external connectivity through the selection of the nodes that offer gateway service.

A third solution consist in creating a PASA Root backup with the same address using the Virtual Router Redundancy Protocol (VRRP [RFC9568]). However, in order to offer full resiliancy, the address allocation state in the primary PASA Root has to be duplicated in the secondary PASA Root.

One last resort, to ensure reliability, is to use a routing protocol, however, such a solution, would annihilate the advantages of the PASA addressing scheme, namely the stateless forwarding.

A more in-depth discussion about reliability, including the case of multiple roots, can be found in [I-D.li-6lo-pasa-reliability]. Furthermore, specific reliability solutions depend as well on the specific Address Assignment Function used (different from the one presented in this document).

Revised Section "Reliability Considerations" to address concern about multiple external connectivity

Possible solutions:

- Multi-topology
  - Suitable if PASA domain is multi-homed
- Root/Gateway roles separation with some of the children of the root are also gateways
  - Suitable when performing load balance of external connections
- Uses VRRP (Virtual Router Redundancy Protocol)
  - Address allocation state to be mirrored on both routers
- Last resort: routing protocol but annihilating PASA benefits (stateless forwarding)

Next Steps

- RTGDIR & GENART Reviews Addressed
  - Getting closer to WGLC

THANKS!

IETF 121 – 6lo

Reliability Considerations of Path-Aware Semantic Addressing

~~draft-li-6lo-pasa-reliability-03~~

draft-li-6lo-pasa-reliability-04

IETF 121 – Dublin

Since IETF 119*

Reliability Considerations of Path-Aware Semantic Addressing

draft-li-6lo-pasa-reliability-04

Status | Email expansions | History

Versions:

00 | 01 | 02 | 03 | 04

This document is an Internet-Draft (I-D). Anyone may submit an I-D to the IETF. This I-D is **not endorsed by the IETF** and has **no formal standing** in the IETF standards process.

draft-li-6lo-pasa-reliability-04.txt
September 2024
Main changes: Address residual issues of RTGDIR review by Joel Halpern

***This document was not presented in IETF 120**

Main Changes (I)

No changes in the document structure

Table of Contents

Table of Contents

- Still done quite some word-smithing here and there and typo fixes.

Main Changes (II)

5. Links/Nodes Failure Detection and Recovery

Previous sections describe actions and possible solutions to failure events, but did not discuss how failures are detected. This memo assumes that depending on the specific technology in use, and the level of desired reliability, the most suitable failure detection mechanism is used to trigger the above-described actions. It is considered not desirable to define one single failure detection technique to be used in the context of PASA, neither to define new ones.
The link failure could be detected by leveraging layer 2 feedbacks, like for instance the lack of acknowledgement upon packet transmission. It can also be detected using existing network layer solutions, like for instance the Bidirectional Forwarding Detection (BFD) [RFC7130] or IPv6 specific mechanisms [RFC5534].

Another aspect of the general failure management is to recover from failures, going back to the original state. In the context of PASA there are a couple of possible approaches that can be used, e.g. by using PASA addresses lifetime. Addresses can be assigned associated with a lifetime. When such lifetime expires, node have to undergo the same initial procedure for address allocation. This is also a good moment to check whether a certain link or node is back to normal functioning. If it is not the case, the algorithmic procedure will anyway create topologies that do not consider failed links/nodes. A faster alternative approach could be based, like in the case of failure detection, on periodic checks that may leverage on layer 2 features or on some neighbor discovery messages. The former method is more effective, while the latter introduces communication overhead.

5. Links/Nodes Failure Detection and Recovery

Previous sections describe actions and possible solutions to failure events, but did not discuss how failures are detected. This memo assumes that depending on the specific technology in use, and the level of desired reliability, the most suitable failure detection mechanism is used to trigger the above-described actions. It is considered not desirable to define one single failure detection technique to be used in the context of PASA, neither to define new ones. The link failure could be detected by leveraging layer 2 feedbacks, like for instance the lack of acknowledgement upon packet transmission. It can also be detected using existing network layer solutions, like for instance the Bidirectional Forwarding Detection (BFD) [RFC7130] or IPv6 specific mechanisms [RFC5534].

Another aspect of the general failure management is to recover from failures, going back to the original state. In particular, since, according to [RFC8505] and [I-D.ietf-6lo-path-aware-semantic-addressing], nodes are supposed to keep addressing state in non-volatile memory, upon failure recovery nodes can just re-register addresses rather than restart the addressing process. This allows as well to signal node recovery, since the recovered node will re-register the previously obtained address, signaling that it is back. For link recovery detection, in the context of PASA, there are a couple of possible approaches that can be used, e.g. by using PASA addresses lifetime. Addresses can be assigned associated with a lifetime. When such lifetime expires, node have to undergo the same initial procedure to re-obtain the same address allocation. This is also a good moment to check whether a certain link or node is back to normal functioning. If it is not the case, the algorithmic procedure will anyway create topologies that do not consider failed links/nodes. However, this will cause renumbering, which, depending on the topology and the location of the failure, may not be the best solution. A faster alternative approach could be based, like in the case of failure detection, on periodic checks that may leverage on layer 2 features or on some neighbor discovery messages. The former method is more effective, while the latter introduces communication overhead.

- Added text about use of non-volatile memory
  - In line with text in the main PASA document and GAAO document

Main Changes (III)

```
6.  Resiliency

   Real resiliency provided by the different approaches depends on the
   specific topology.

   The single-address solution may introduce more state.  Indeed, the
   root has the overview of the PASA network.  It knows all nodes'
   addresses, the alternative links and the broken links.  It is able to
   compute a usable path towards a destination.  This comes with the
   benefit of potentially being able to find a higher number of
   alternative paths, hence, in the end providing a stronger protection
   against multiple failures.  The PASA Router and PASA Host are rather
   dummy, performing PASA stateless forwarding.  They only are aware of
   the link state toward their direct neighbors, and act accordingly.
   However, the use of source routing may create MTU issues if the path
   is too long.

   The multi-address approach leverages more on the stateless forwarding
   of PASA.  The root is in general unaware of nodes' addresses, and the
   network topology.  In case of failure, a redirection rule is set on
   the root, hence the amount of states is proportional to the number of
   failures.  This means less state overall, but may be less robust to
   multiple failures.  Differently from the single address solution, a
   small state is also required on PASA Routers, because if a link fails
   a redirect rule has to be used.

   The above-mentioned pros and cons need to be pondered when choosing a
   reliability solution to be deployed in an PASA domain.
```

```
6.  Resiliency

   Real resiliency provided by the different approaches presented in
   this document depends on the specific topology.

   The single-address solution may introduce more state.  Indeed, the
   root has full knowledge of the PASA network.  It knows all nodes'
   addresses, the alternative links and the broken links.  It is able to
   compute a usable path towards a destination.  This comes with the
   benefit of potentially being able to find a higher number of
   alternative paths, hence, in the end, providing a stronger protection
   against multiple failures.  The PASA Router and PASA Host are rather
   dummy, performing PASA stateless forwarding.  They only are aware of
   the link state toward their direct neighbors, and act accordingly.
   However, the use of source routing may create MTU issues if the path
   is too long.

   The multi-address approach leverages more on the stateless forwarding
   of PASA.  The root is in general unaware of nodes' addresses, and the
   network topology.  In case of failure, a redirection rule is set on
   the root, hence the amount of states is proportional to the number of
   failures.  This means less state overall, but may be less robust to
   multiple failures.  Differently from the single address solution, a
   small state is also required on PASA Routers, because if a link fails
   a redirect rule has to be used.

   The above-mentioned pros and cons need to be pondered when choosing a
   reliability solution to be deployed in an PASA domain.

   Both approaches, presented in previous sections, rely on the presence
   of more than one PASA Root, which provides resilient connectivity
   toward other networks (or the Internet).  In case of one PASA Root
   failure, the procedures described in the present document apply, and
   external connectivity is provided by the other PASA Roots.  The same
   applies if one of the link between the PASA Roots and their children
   fail; procedures described in this document provide resilient
   connectivity.  Resiliency in the external connectivity depend on the
   specific deployment, provisioning, and technology used, which are out
   of the scope of this document.
```

- Added text about the use of multiple PASA Roots

Next Steps

THANKS!