Network Working Group B. Liu
Internet-Draft S. Jiang
Intended status: Standards Track Huawei Technologies
Expires: December 14, 2015 June 12, 2015

Intent Distribution for Autonomic Networking
draft-liu-anima-intent-distribution-00

Abstract

This document describes the requirements of distributing intent information in an autonomic network. Ideally, the intent may be generated/injected at an arbitrary autonomic node and be distributed among the whole autonomic domain. Then this document resolves the distribution requirements into protocol design requirements. Specifically, this document introduces a solution which is some extension based on the Anima signalling protocol (GDNP, Generic Discovery and Negotiation Protocol).

Table of Contents

1. Introduction

This document describes the requirements of distributing intent information in an autonomic network. Ideally, the intent may be generated/injected at an arbitrary autonomic node and be distributed among the whole autonomic domain. Then this document resolves the distribution requirements into protocol design requirements. Specifically, this document introduces a solution which is some extension based on the Anima signalling protocol (GDNP, Generic Discovery and Negotiation Protocol).

2. Intent Distribution Requirements

2.1. Distributed to the Whole Domain

When the intent is injected at an arbitrary autonomic node, the node MUST be able to distribute it to the whole nodes in the domain. This requirement does not necessarily mean the node need to send the intent to all nodes through unicast or multicast all by itself; there might be distribution function infrastructure that could be used/ triggered by the node.

2.1.1. Autonomic Domain Boundary

The domain boundary devices are supposed to know themselves as boundary.When the messages come to the devices, they won't distribute them anymore so that the messages are only distributed with the domain.

[Editor's Notes] It is a practical issue that how an autonomic node knows itself is the domain boundary. It is not in the scope of this document.

2.2. De-coupling of Intent Content and Bearing Protocol

The content of intent SHOULD NOT be coupled with the bearing protocol.

2.3. Avoiding Signaling Storm

If flooding mechanism is used, then there should be a mechanism to prevent the packets which carrying the intent to travel around the domain again and again.

2.4. Arbitary Intent Injecting Point (Optional)

The intent SHOULD be injected at any autonomic node, rather than a pre-specified Node.

Discuss: may be only within a group of autonomic nodes, it supports input at "any" node.

2.5. Conflict Handling (Optional)

So long as it supports arbitrary point where to inject an intent, there is possibility that two nodes advertise the same intent but with different contents. Hence, there should be a mechanism to handle the conflicted intent.

3. Protocol Design

3.1. Protocol Requirements

3.1.1. Multicast and Unicast Communication

The following communication modes need to be supported to distribute intents.

- o On Link Multicast

  This is a basic distribution behavior. The message is multicasted to all the other nodes on the same link.

- o Off Link Multicast

  Normally, an autonomic domain is not only limited within a link. Thus, off link multicast is needed to reach the other nodes out of the initiator's link.

  When there is off-link multicast, there needs to be flooding control mechanisms as described in Section 3.2.3.

- o Point to Point

  Besides multicast, the intent might be distributed only between two nodes. Thus, point to point unicast communication is also needed.

3.1.2. Messages Interaction

The following message interaction modes need to be supported.

- o Unsolicited advertisement

  This is the most typical use case of intent distribution. The intent is advertised by one of the autonomic node and flooded to all the others in the same autonomic domain. The process is stateless ,which means there is no need to pre-establish connections between autonomic nodes for intent exchange, hence the autonomic nodes are always monitoring the intent coming.

  [Editor's Note] This document doesn't achieve unsolicit advertisement for intent as a new explicit message type, instead, it is by the ASA interpreting the Intent Option (defined below) and resolving it into GDNP synchronization behavior at each hop. However, Unsolicit Advertisement might be a generic function that reused by various ASA. So, if the

ANIMA signalling is going to provide such function at message level in the future, this document would update accordingly.

o Request-Response

This is mostly for point-to-point intent exchange. For example, when a new device gets online, it request intent from it’s neighbor, then the neighbor will distribute the common intent that shared among all the nodes within the autonomic domain to the new device.

3.1.3. Fragmentation Considerations

Since the intent distribution runs over GDNP, it does not provide any explicit fragmentation/reassembly support.

[Editor’s Note] However, there might be concern that when an intent packet needs to be split, it might need to be split into fragments each of which could be interpreted individually, thus there is no need to wait for the assembling of all fragments. However, this is only a hypothetical use case.

3.2. Intent Distribution over Anima Signaling Protocol

Since there is a signalling protocol under development in Anima working group, it is reasonable to leverage the current protocol to do intent distribution.

This section makes some extension to the signalling protocol to fulfil the requirements described above. Specifically, the extension is based on the 03 version of the GDNP protocol [I-D.carpenter-anima-gdn-protocol] .

3.2.1. Intent Option

The content of intent is encapsulated as a dedicated option, so that it could be carried by various type of messages if needed.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         OPTION_INTENT         |          option-len           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|F|  Flood TTL  |    Reserved   |        Sequence Number        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Domain ID of the Source Node of the Message           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              TBD                              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                         Intent Content                        |
|          (variable length, Supporting arbitrary format)       |
.                                                               .
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

OPTION_INTENT: Identifies the Intent Option type. 16-bit.

option-len: Length of the whole intent. 16-bit.

F bit: Flooding flag bit. When the flag is set, it means the intent needs to be flooded.

Flood TTL: Limits the hops that an Intent message could travel. 8-bit.

Reserved: Set to zero, ignored on receipt. 8-bit.

Sequence Number: A sequence number to identify an intent option. 16-bit. Each time one node sends an Intent Option, the sequence number MUST be increased.

Node ID: Identifies the source of the Intent Option. 32-bit. This documents assumes that each autonomic nodes has a Node ID available after the bootstrapping process described in [I-D.pritikin-anima-bootstrapping-keyinfra] . The Node ID may generated based on the domain certificate issued to the node during bootstrapping.

Intent Content: The intent content, such as the intent specified in [I-D.du-anima-an-intent].

3.2.2. Node Behavior

o Initiating Node

a) Assuming there is an ASA in charge of the intent distribution.

b) The ASA generates the Intent Option and calls the GDNP module to send the Intent Option in a Request Message (as defined in Section 3.6.4 of [I-D.carpenter-anima-gdn-protocol]).

c) If this is a flooding intent, the ASA sets the F flag and calls the GDNP module to multicast the message to all it's neighbors in a Discovery Message (as defined in Section 3.6.2 of [I-D.carpenter-anima-gdn-protocol]).

o Receiving Node

a) Assuming there is an ASA in charge of the intent distribution.

b) The GDNP module extracts the Intent Option and handle it up to the ASA.

c) If the F flag is not set, the node calls the GDNP module to response a Negotiation-Ending message with a Accept Option; if set, then no need to response. [Open Question] Does nodes need to response for the flooding intent?

d) If it is a flooding intent, the node multicast the option again to all it's neighbors.

[Editor's Note] The behavior as described above could also be achieved through Usolicit Synchronization message/function which was briefly discussed in the previous version of GDNP. If the Unsolicited Synchronization is added back to the GDNP, this document should also consider the relevant solution accordingly.

3.2.3. Flooding Control

o Loop Avoidance

When messages are flooded off link, it is highly possible that the message would be flooded back to the initiator again, thus there would be a large amount of duplicated messages circling around the network. So, there needs to be relevant mechanism to avoid/limit the packets loop.

To achieve this goal, the nodes need to do the following actions:

a) The node maintains a flooding state table which stores each interface's record that whether a specific intent option had been received or sent from it. The option identification could be the combination of the Sequence Number and Node ID in the Intent Option.

b) The node MUST NOT send a flooding Intent Option message to the interfaces that had received or sent the same Intent Option.

o Flooding TTL

In case an Intent is occasionally looped around, the Flooding TTL is to guarantee the packet would not travel in a infinite loop in the network.

[Open Question] Is Flooding TTL a redundant field?

4. Security Considerations

Intents could significantly influence the network nodes' behavior, so authentication is strongly required.

However, the authentication could be done at multiple layers:

- Outer layer authentication: the ASAs' communication is within a protected channel such as ACP (Autonomic Control Plane, [I-D.behringer-anima-autonomic-control-plane]).

- Inner layer authentication: the ASAs' communication might not be within a protected channel, then there should be embedded protection in intent distribution itself.

[Open Question] As described in section 7.3 of [I-D.irtf-nmrg-autonomic-network-definitions], intent distribution is considered as a function of the ACP. Do we consider to remove this limitation? ACP is a good secure channel for distributing intent, but maybe not a mandatory channel.

5. IANA Considerations

TBD.

6. Acknowledgements

This document is inherited from [I-D.carpenter-anima-gdn-protocol] and [I-D.behringer-anima-reference-model]. So thanks all the contributors of the two work items.

This document was produced using the xml2rfc tool [RFC2629].

7. References

7.1. Normative References

[I-D.carpenter-anima-gdn-protocol]
Carpenter, B. and B. Liu, "A Generic Discovery and Negotiation Protocol for Autonomic Networking", draft-carpenter-anima-gdn-protocol-03 (work in progress), April 2015.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2629] Rose, M., "Writing I-Ds and RFCs using XML", RFC 2629, June 1999.

7.2. Informative References

[I-D.behringer-anima-autonomic-control-plane]
Behringer, M., Bjarnason, S., BL, B., and T. Eckert, "An Autonomic Control Plane", draft-behringer-anima-autonomic-control-plane-02 (work in progress), March 2015.

[I-D.behringer-anima-reference-model]
Behringer, M., Carpenter, B., Eckert, T., Ciavaglia, L., and B. Liu, "A Reference Model for Autonomic Networking", draft-behringer-anima-reference-model-01 (work in progress), April 2015.

[I-D.irtf-nmrg-autonomic-network-definitions]
Behringer, M., Pritikin, M., Bjarnason, S., Clemm, A., Carpenter, B., Jiang, S., and L. Ciavaglia, "Autonomic Networking - Definitions and Design Goals", draft-irtf-nmrg-autonomic-network-definitions-07 (work in progress), March 2015.

[I-D.pritikin-anima-bootstrapping-keyinfra]
Pritikin, M., Behringer, M., and S. Bjarnason, "Bootstrapping Key Infrastructures", draft-pritikin-anima-bootstrapping-keyinfra-01 (work in progress), February 2015.

Authors' Addresses

Bing Liu
Huawei Technologies
Q14, Huawei Campus
No.156 Beiqing Road
Hai-Dian District, Beijing 100095
P.R. China

Email: leo.liubing@huawei.com

Sheng Jiang
Huawei Technologies
Q14, Huawei Campus
No.156 Beiqing Road
Hai-Dian District, Beijing 100095
P.R. China

Email: jiangsheng@huawei.com