

ANIMA
Internet-Draft
Intended status: Informational
Expires: April 21, 2016

T. Eckert
M. Behringer
Cisco
October 19, 2015

Using Autonomic Control Plane for Stable Connectivity of Network OAM
draft-eckert-anima-stable-connectivity-02

Abstract

OAM (Operations, Administration and Management) processes for data networks are often subject to the problem of circular dependencies when relying on network connectivity of the network to be managed for the OAM operations itself. Provisioning during device/network bring up tends to be far less easy to automate than service provisioning later on, changes in core network functions impacting reachability can not be automated either because of ongoing connectivity requirements for the OAM equipment itself, and widely used OAM protocols are not secure enough to be carried across the network without security concerns.

This document describes how to integrate OAM processes with the autonomic control plane (ACP) in Autonomic Networks (AN). to provide stable and secure connectivity for those OAM processes.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on April 21, 2016.

Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

| | | |
|--------|---|----|
| 1. | Introduction | 2 |
| 1.1. | Self dependent OAM connectivity | 2 |
| 1.2. | Data Communication Networks (DCNs) | 3 |
| 1.3. | Leveraging the ACP | 3 |
| 2. | Solutions | 4 |
| 2.1. | Stable connectivity for centralized OAM operations . . . | 4 |
| 2.1.1. | Simple connectivity for non-autonomic NOC application devices | 4 |
| 2.1.2. | Limitations and enhancement overview | 5 |
| 2.1.3. | Simultaneous ACP and data plane connectivity | 6 |
| 2.1.4. | IPv4 only NOC application devices | 7 |
| 2.1.5. | Path selection policies | 8 |
| 2.1.6. | Autonomic NOC device/applications | 10 |
| 2.1.7. | Encryption of data-plane connections | 10 |
| 2.1.8. | Long term direction of the solution | 11 |
| 2.2. | Stable connectivity for distributed network/OAM functions | 12 |
| 3. | Security Considerations | 12 |
| 4. | No IPv4 for ACP | 14 |
| 5. | Further considerations | 14 |
| 6. | IANA Considerations | 14 |
| 7. | Acknowledgements | 14 |
| 8. | Change log [RFC Editor: Please remove] | 15 |
| 9. | References | 15 |
| | Authors' Addresses | 16 |

1. Introduction

1.1. Self dependent OAM connectivity

OAM (Operations, Administration and Management) processes for data networks are often subject to the problem of circular dependencies when relying on network connectivity of the network to be managed for the OAM operations itself:

The ability to perform OAM operations on a network device requires first the execution of OAM procedures necessary to create network

connectivity to that device in all intervening devices. This typically leads to sequential, 'expanding ring configuration' from a NOC. It also leads to tight dependencies between provisioning tools and security enrollment of devices. Any process that wants to enroll multiple devices along a newly deployed network topology needs to tightly interlock with the provisioning process that creates connectivity before the enrollment can move on to the next device.

When performing change operations on a network, it likewise is necessary to understand at any step of that process that there is no interruption of connectivity that could lead to removal of connectivity to remote devices. This includes especially change provisioning of routing, security and addressing policies in the network that often occur through mergers and acquisitions, the introduction of IPv6 or other mayor re-hauls in the infrastructure design.

All this circular dependencies make OAM processes complex and potentially fragile. When automation is being used, for example through provisioning systems or network controllers, this complexity extends into that automation software.

1.2. Data Communication Networks (DCNs)

In the late 1990'th and early 2000, IP networks became the method of choice to build separate OAM networks for the communications infrastructure in service providers. This concept was standardized in G.7712/Y.1703 and called "Data Communications Networks" (DCN). These where (and still are) physically separate IP(/MPLS) networks that provide access to OAM interfaces of all equipment that had to be managed, from PSTN switches over optical equipment to nowadays ethernet and IP/MPLS production network equipment.

Such DCN provide stable connectivity not subject to aforementioned problems because they are separate network entirely, so change configuration of the production IP network is done via the DCN but never affects the DCN configuration. Of course, this approach comes at a cost of buying and operating a separate network and this cost is not feasible for many networks, most notably smaller service providers, most enterprises and typical IoT networks.

1.3. Leveraging the ACP

One goal of the Autonomic Networks Autonomic Control plane (ACP) is to provide similar stable connectivity as a DCN, but without having to build a separate DCN. It is clear that such 'in-band' approach can never achieve fully the same level of separation, but the goal is to get as close to it as possible.

This solution approach has several aspects. One aspect is designing the implementation of the ACP in network devices to make it actually perform without interruption by changes in what we will call in this document the "data-plane", aka: the operator or controller configured services planes of the network equipment. This aspect is not currently covered in this document.

Another aspect is how to leverage the stable IPv6 connectivity provided by the ACP to build actual OAM solutions. This is the current scope of this document.

2. Solutions

2.1. Stable connectivity for centralized OAM operations

In the most common case, OAM operations will be performed by one or more applications running on a variety of centralized NOC systems that communicate with network devices. We describe differently advanced approaches to leverage the ACP for stable connectivity leveraging the ACP. The descriptions will show that there is a wide range of options, some of which are simple, some more complex.

Most easily we think there are three stages of interest:

- o There are simple options described first that we consider to be good starting points to operationalize the use of the ACP for stable connectivity.
- o There are more advanced intermediate options that try to establish backward compatibility with existing deployed approaches such as leveraging NAT. Selection and deployment of these approaches needs to be carefully vetted to ensure that they provide positive ROI. This very much depends on the operational processes of the network operator.
- o It seems clearly feasible to build towards a long-term configuration that provides all the desired operational, zero touch and security benefits of an autonomic network, but a range of details for this still have to be worked out.

2.1.1. Simple connectivity for non-autonomic NOC application devices

In the most simple deployment case, the ACP extends all the way into the NOC via a network device that is set up to provide access into the ACP natively to non-autonomic devices. It acts as the default-router to those hosts and provides them with only IPv6 connectivity into the ACP - but no IPv4 connectivity. NOC devices with this setup

need to support IPv6 but require no other modifications to leverage the ACP.

This setup is sufficient for troubleshooting OAM operations such as SSH into network devices, NMS that perform SNMP read operations for status checking, for software downloads into autonomic devices and so on. In conjunction with otherwise unmodified OAM operations via separate NOC devices/applications it can provide a good subset of the interesting stable connectivity goals from the ACP.

Because the ACP provides 'only' for IPv6 connectivity, and because the addressing provided by the ACP does not include any addressing structure that operations in a NOC often relies on to recognize where devices are on the network, it is likely highly desirable to set up DNS so that the ACP IPv6 addresses of autonomic devices are known via domain names with logical names. For example, if DNS in the network was set up with names for network devices as devicename.noc.example.com, then the ACP address of that device could be mapped to devicename-acp.noc.exmample.com.

2.1.2. Limitations and enhancement overview

This most simple type of attachment of NOC applications to the ACP suffers from a range of limitations:

1. NOC applications can not directly probe whether the desired so called 'data-plane' network connectivity works because they do not directly have access to it. This problem is not dissimilar to probing connectivity for other services (such as VPN services) that they do not have direct access to, so the NOC may already employ appropriate mechanisms to deal with this issue (probing proxies).
2. NOC applications need to support IPv6 which often is still not the case in many enterprise networks.
3. Performance of the ACP will be limited versus normal 'data-plane' connectivity. The setup of the ACP will often support only non-hardware accelerated forwarding. Running a large amount of traffic through the ACP, especially for tasks where it is not necessary will reduce its performance/effectiveness for those operations where it is necessary or highly desirable.
4. Security of the ACP is reduced by exposing the ACP natively (and unprotected) into a LAN In the NOC where the NOC devices are attached to it.

These four problems can be tackled independently of each other by solution improvements. Combining these solutions improvements together ultimately leads towards the the target long term solution.

2.1.3. Simultaneous ACP and data plane connectivity

Simultaneous connectivity to both ACP and data-plane can be achieved in a variety of ways. If the data-plane is only IPv4, then any method for dual-stack attachment of the NOC device/application will suffice: IPv6 connectivity from the NOC provides access via the ACP, IPv4 will provide access via the data-plane. If as explained above in the most simple case, an autonomic device supports native attachment to the ACP, and the existing NOC setup is IPv4 only, then it could be sufficient to simply attach the ACP device(s) as the IPv6 default-router to the NOC LANs and keep the existing IPv4 default router setup unchanged.

If the data-plane of the network is also supporting IPv6, then the NOC devices that need access to the ACP should have a dual-homing IPv6 setup. One option is to make the NOC devices multi-homed with one logical or physical IPv6 interface connecting to the data-plane, and another into the ACP. The LAN that provides access to the ACP should then be given an IPv6 prefix that shares a common prefix with the IPv6 ULA of the ACP so that the standard IPv6 interface selection rules on the NOC host would result in the desired automatic selection of the right interface: towards the ACP facing interface for connections to ACP addresses, and towards the data-plane interface for anything else. If this can not be achieved automatically, then it needs to be done via simple IPv6 static routes in the NOC host.

Providing two virtual (eg: dot1q subnet) connections into NOC hosts may be seen as undesired complexity. In that case the routing policy to provide access to both ACP and data-plane via IPv6 needs to happen in the NOC network itself: The NOC application device gets a single attachment interface but still with the same two IPv6 addresses as in before - one for use towards the ACP, one towards the data-plane. The first-hop router connecting to the NOC application device would then have separate interfaces: one towards the data-plane, one towards the ACP. Routing of traffic from NOC application hosts would then have to be based on the source IPv6 address of the host: Traffic from the address designated for ACP use would get routed towards the ACP, traffic from the designated data-plane address towards the data-plane.

In the most simple case, we get the following topology: Existing NOC application devices connect via an existing NOC LAN and existing first hop Rtr1 to the data-plane. Rtr1 is not made autonomic, but instead the edge router of the Autonomic network ANrtr is attached via a

separate interface to Rtr1 and ANrtr provides access to the ACP via ACPAccessLan. Rtr1 is configured with the above described IPv6 source routing policies and the NOC-app-devices are given the secondary IPv6 address for connectivity into the ACP.

```

--... (data-plane)
NOC-app-device(s) -- NOClan -- Rtr1
--- ACPAccessLan -- ANrtr ... (ACP)

```

Figure 1

If Rtr1 was to be upgraded to also implement Autonomic Networking and the ACP, the picture would change as follows:

```

----- ... (data-plane)
NOC-app-device(s) ----- NOClan --- ANrtr1
      . . ----- ... (ACP)
      \-/
      (ACP to data-plane loopback)

```

Figure 2

In this case, ANrtr1 would have to implement some more advanced routing such as cross-VRF routing because the data-plane and ACP are most likely run via separate VRFs. A simple short-term workaround could be a physical external loopback cable into two ports of ANrtr1 to connect the data-plane and ACP VRF as shown in the picture.

2.1.4. IPv4 only NOC application devices

With the ACP being intentionally IPv6 only, attachment of IPv4 only NOC application devices to the ACP requires the use of IPv4 to IPv6 NAT. This NAT setup could for example be done in Rtr1 in above picture to also support IPv4 only NOC application devices connected to NOClan.

To support connections initiated from IPv4 only NOC applications towards the ACP of network devices, it is necessary to create a static mapping of ACP IPv6 addresses into an unused IPv4 address space and dynamic or static mapping of the IPv4 NOC application device address (prefix) into IPv6 routed in the ACP. The main issue in this setup is the mapping of all ACP IPv6 addresses to IPv4. Without further network intelligence, this needs to be a 1:1 address mapping because the prefix used for ACP IPv6 addresses is too long to be mapped directly into IPv4 on a prefix basis.

One could implement in router software dynamic mappings by leveraging DNS, but it seems highly undesirable to implement such complex technologies for something that ultimately is a temporary problem (IPv4 only NOC application devices). With today's operational directions it is likely more preferable to automate the setup of 1:1 NAT mappings in that NAT router as part of the automation process of network device enrollment into the ACP.

The ACP can also be used for connections initiated by the network device into the NOC application devices. For example syslog from autonomic devices. In this case, static mappings of the NOC application devices IPv4 addresses are required. This can easily be done with a static prefix mapping into IPv6.

Overall, the use of NAT is especially subject to the RoI considerations, but the methods described here may not be too different from the same problems encountered totally independent of AN/ACP when some parts of the network are to introduce IPv6 but NOC application devices are not (yet) upgradeable.

2.1.1.5. Path selection policies

As mentioned above, the ACP is not expected to have high performance because its primary goal is connectivity and security, and for existing network device platforms this often means that it is a lot more effort to implement that additional connectivity with hardware acceleration than without - especially because of the desire to support full encryption across the ACP to achieve the desired security.

Some of these issues may go away in the future with further adoption of the ACP and network device designs that better tender to the needs of a separate OAM plane, but it is wise to plan for even long-term designs of the solution that does NOT depend on high-performance of the ACP. This is opposite to the expectation that future NOC application devices will have IPv6, so that any considerations for IPv4/NAT in this solution are temporary.

To solve the expected performance limitations of the ACP, we do expect to have the above describe dual-connectivity via both ACP and data-plane between NOC application devices and AN devices with ACP. The ACP connectivity is expected to always be there (as soon as a device is enrolled), but the data-plane connectivity is only present under normal operations but will not be present during eg: early stages of device bootstrap, failures, provisioning mistakes or during network configuration changes.

The desired policy is therefore as follows: In the absence of further security considerations (see below), traffic between NOC application and AN devices should prefer data-plane connectivity and resort only to using the ACP when necessary, unless it is an operation known to be so much tied to the cases where the ACP is necessary that it makes no sense to try using the data plane. An example here is of course the SSH connection from the NOC into a network device to troubleshoot network connectivity. This could easily always rely on the ACP. Likewise, if a NOC application is known to transmit large amounts of data, and it uses the ACP, then its performance need to be controlled so that it will not overload the ACP performance. Typical examples of this are software downloads.

There is a wide range of methods to build up these policies. We describe a few:

DNS can be used to set up names for the same network devices but with different addresses assigned: One name (name.noc.example.com) with only the data-plane address(es) (IPv4 and/or IPv6) to be used for probing connectivity or performing routine software downloads that may stall/fail when there are connectivity issues. One name (name-acp.noc.example.com) with only the ACP reachable address of the device for troubleshooting and probing/discovery that is desired to always only use the ACP. One name with data plane and ACP addresses (name-both.noc.example.com).

Traffic policing and/or shaping of at the ACP edge in the NOC can be used to throttle applications such as software download into the ACP.

MP-TCP is a very attractive candidate to automate the use of both data-plane and ACP and minimize or fully avoid the need for the above mentioned logical names to pre-set the desired connectivity (data-plane-only, ACP only, both). For example, a set-up for non MP-TCP aware applications would be as follows:

DNS naming is set up to provide the ACP IPv6 address of network devices. Unbeknownst to the application, MP-TCP is used. MP-TCP mutually discovers between the NOC and network device the data-plane address and carries all traffic across it when that MP-TCP sub-flow across the data-plane can be built.

In the Autonomic network devices where data-plane and ACP are in separate VRFs, it is clear that this type of MP-TCP sub-flow creation across different VRFs is new/added functionality. Likewise the policies of preferring a particular address (NOC-device) or VRF (AN device) for the traffic is potentially also a policy not provided as a standard.

2.1.6. Autonomic NOC device/applications

Setting up connectivity between the NOC and autonomic devices when the NOC device itself is non-autonomic is as mentioned in the beginning a security issue. It also results as shown in the previous paragraphs in a range of connectivity considerations, some of which may be quite undesirable or complex to operationalize.

Making NOC application devices autonomic and having them participate in the ACP is therefore not only a highly desirable solution to the security issues, but can also provide a likely easier operationalization of the ACP because it minimizes NOC-special edge considerations - the ACP is simply built all the way automatically, even inside the NOC and only authorized and authenticate NOC devices/applications will have access to it.

Supporting the ACP all the way into an application device requires implementing the following aspects in it: AN bootstrap/enrollment mechanisms, the secure channel for the ACP and at least the host side of IPv6 routing setup for the ACP. Minimally this could all be implemented as an application and be made available to the host OS via eg: a tap driver to make the ACP show up as another IPv6 enabled interface.

Having said this: If the structure of NOC applications is transformed through virtualization anyhow, then it may be considered equally secure and appropriate to construct a (physical) NOC application system by combining a virtual AN/ACP enabled router with non-AN/ACP enabled NOC-application VMs via a hypervisor, leveraging the configuration options described in the previous sections but just virtualizing them.

2.1.7. Encryption of data-plane connections

When combining ACP and data-plane connectivity for availability and performance reasons, this too has an impact on security: When using the ACP, the traffic will be mostly encryption protected, especially when considering the above described use of AN application devices. If instead the data-plane is used, then this is not the case anymore unless it is done by the application.

The most simple solution for this problem exists when using AN NOC application devices, because in that case the communicating AN NOC application and the AN network device have certificates through the AN enrollment process that they can mutually trust (same AN domain). In result, data-plane connectivity that does support this can simply leverage TLS/dTLS with mutual AN-domain certificate authentication - and does not incur new key management.

If this automatic security benefit is seen as most important, but a "full" ACP stack into the NOC application device is unfeasible, then it would still be possible to design a stripped down version of AN functionality for such NOC hosts that only provides enrollment of the NOC host into the AN domain to the extent that the host receives an AN domain certificate, but without directly participating in the ACP afterwards. Instead, the host would just leverage TLS/dTLS using its AN certificate via the data-plane with AN network devices as well as indirectly via the ACP with the above mentioned in-NOC network edge connectivity into the ACP.

When using the ACP itself, TLS/dTLS for the transport layer between NOC application and network device is somewhat of a double price to pay (ACP also encrypts) and could potentially be optimized away, but given the assumed lower performance of the ACP, it seems that this is an unnecessary optimization.

2.1.8. Long term direction of the solution

If we consider what potentially could be the most lightweight and autonomic long term solution based on the technologies described above, we see the following direction:

1. NOC applications should at least support IPv6. IPv4/IPv6 NAT in the network to enable use of ACP is long term undesirable. Having IPv4 only applications automatically leverage IPv6 connectivity via host-stack options is likely non-feasible (NOTE: this has still to be vetted more).
2. Build the ACP as a lightweight application for NOC application devices so ACP extends all the way into the actual NOC application devices.
3. Leverage and as necessary enhance MP-TCP with automatic dual-connectivity: If the MP-TCP unaware application is using ACP connectivity, the policies used should add sub-flow(s) via the data-plane and prefer them.
4. Consider how to best map NOC application desires to underlying transport mechanisms: With the above mentioned 3 points, not all options are covered. Depending on the OAM operation, one may still want only ACP, only data-plane, or automatically prefer one over the other and/or use the ACP with low performance or high-performance (for emergency OAM actions such as countering DDoS). It is as of today not clear what the simplest set of tools is to enable explicitly the choice of desired behavior of each OAM operations. The use of the above mentioned DNS and MP-TCP

mechanisms is a start, but this will require additional thoughts. This is likely a specific case of the more generic scope of TAPS.

2.2. Stable connectivity for distributed network/OAM functions

The ACP can provide common direct-neighbor discovery and capability negotiation and stable and secure connectivity for functions running distributed in network devices. It can therefore eliminate the need to re-implement similar functions in each distributed function in the network. Today, every distributed protocol does this with functional elements usually called "Hello" mechanisms and with often protocol specific security mechanisms.

KARP has tried to start provide common directions and therefore reduce the re-invention of at least some of the security aspects, but it only covers routing-protocols and it is unclear how well it applicable to a potentially wider range of network distributed agents such as those performing distributed OAM functions. The ACP can help in these cases.

This section is TBD for further iterations of this draft.

3. Security Considerations

We discuss only security considerations not covered in the appropriate sub-sections of the solutions described.

Even though ACPs are meant to be isolated, explicit operator misconfiguration to connect to insecure OAM equipment and/or bugs in ACP devices may cause leakage into places where it is not expected. Mergers/Aquisitions and other complex network reconfigurations affecting the NOC are typical examples.

ULA addressing as proposed in this document is preferred over globally reachable addresses because it is not routed in the global Internet and will therefore be subject to more filtering even in places where specific ULA addresses are being used.

Randomn ULA addressing provides more than sufficient protection against address collision even though there is no central assignment authority. This is helped by the expectation, that ACPs are never expected to connect all together, but only few ACPs may ever need to connect together, eg: when mergers and aquisitions occur.

If packets with unexpected ULA addresses are seen and one expects them to be from another networks ACP from which they leaked, then some form of ULA prefix registrastion (not allocation) can be beneficial. Some voluntary registries exist, for example

<https://www.sixxs.net/tools/grh/ula/>, although none of them is preferable because of being operated by some recognized authority. If an operator would want to make its ULA prefix known, it might need to register it with multiple existing registries.

ULA Centrally assigned ULA addresses (ULA-C) was an attempt to introduce centralized registration of randomly assigned addresses and potentially even carve out a different ULA prefix for such addresses. This proposal is currently not proceeding, and it is questionable whether the stable connectivity use case provides sufficient motivation to revive this effort.

Using current registration options implies that there will not be reverse DNS mapping for ACP addresses. For that one will have to rely on looking up the unknown/unexpected network prefix in the registry registry to determine the owner of these addresses.

Reverse DNS resolution may be beneficial for specific already deployed insecure legacy protocols on NOC OAM systems that intend to communicate via the ACP (eg: TFTP) and leverages reverse-DNS for authentication. Given how the ACP provides path security except potentially for the last-hop in the NOC, the ACP does make it easier to extend the lifespan of such protocols in a secure fashion as far as just the transport is concerned. The ACP does not make reverse DNS lookup a secure authentication method though. Any current and future protocols must rely on secure end-to-end communications (TLDP, dTLS) and identification and authentication via the certificates assigned to both ends. This is enabled by the certificate mechanisms of the ACP.

If DNS and especially reverse DNS are set up, then it should be set up in an automated fashion, linked to the autonomic registrar backend so that the DNS and reverse DNS records are actually derived from the subject name elements of the ACP device certificates in the same way as the autonomic devices themselves will derive their ULA addresses from their certificates to ensure correct and consistent DNS entries.

If an operator feels that reverse DNS records are beneficial to its own operations but that they should not be made available publically for "security" by concealment reasons, then the case of ACP DNS entries is probably one of the least problematic use cases for split-DNS: The ACP DNS names are only needed for the NOC applications intending to use the ACP - but not network wide across the enterprise.

4. No IPv4 for ACP

The ACP is targeted to be IPv6 only, and the prior explanations in this document show that this can lead to some complexity when having to connect IPv4 only NOC solutions, and that it will be impossible to leverage the ACP when the OAM agents on an ACP network device do not support IPv6. Therefore, the question was raised whether the ACP should optionally also support IPv4.

The decision not to include IPv4 for ACP as something that is considered in the use cases in this document is because of the following reasons:

In SP networks that have started to support IPv6, often the next planned step is to consider moving out IPv4 from a native transport as just a service on the edge. There is no benefit/need for multiple parallel transport families within the network, and standardizing on one reduces OPEX and improves reliability. This evolution in the data plane makes it highly unlikely that investing development cycles into IPv4 support for ACP will have a longer term benefit or enough critical short-term use-cases. Support for only IPv4 for ACP is purely a strategic choice to focus on the known important long term goals.

In other type of networks as well, we think that efforts to support autonomic networking is better spent in ensuring that one address family will be support so all use cases will long-term work with it, instead of duplicating effort into IPv4. Especially because auto-addressing for the ACP with IPv4 would be more ecomplex than in IPv6 due to the the IPv4 addressing space.

5. Further considerations

6. IANA Considerations

This document requests no action by IANA.

7. Acknowledgements

This work originated from an Autonomic Networking project at cisco Systems, which started in early 2010 including customers involved in the design and early testing. Many people contributed to the aspects described in this document, including in alphabetical order: BL Balaji, Steinthor Bjarnason, Yves Herthoghs, Sebastian Meissner, Ravi Kumar Vadapalli. The author would also like to thank Michael Richardson, James Woodyatt and Brian Carpenter for their review and comments.

8. Change log [RFC Editor: Please remove]

02: Updated references.

02: Modified ULA text to not suggest ULA-C as much better anymore, but still mention it.

02: Added explanation why no IPv4 for ACP.

01: Added security section discussing the role of address prefix selection and DNS for ACP. Title change to emphasize focus on OAM. Expanded abstract.

00: Initial version.

9. References

[I-D.behringer-anima-reference-model]

Behringer, M., Carpenter, B., Eckert, T., Ciavaglia, L., Liu, B., Jeff, J., and J. Strassner, "A Reference Model for Autonomic Networking", draft-behringer-anima-reference-model-04 (work in progress), October 2015.

[I-D.ietf-anima-autonomic-control-plane]

Behringer, M., Bjarnason, S., BL, B., and T. Eckert, "An Autonomic Control Plane", draft-ietf-anima-autonomic-control-plane-01 (work in progress), October 2015.

[I-D.ietf-anima-bootstrapping-keyinfra]

Pritikin, M., Richardson, M., Behringer, M., and S. Bjarnason, "Bootstrapping Key Infrastructures", draft-ietf-anima-bootstrapping-keyinfra-01 (work in progress), October 2015.

[I-D.irtf-nmrg-an-gap-analysis]

Jiang, S., Carpenter, B., and M. Behringer, "General Gap Analysis for Autonomic Networking", draft-irtf-nmrg-an-gap-analysis-06 (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.

[RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, DOI 10.17487/RFC4193, October 2005, <<http://www.rfc-editor.org/info/rfc4193>>.

Authors' Addresses

Toerless Eckert
Cisco

Email: eckert@cisco.com

Michael H. Behringer
Cisco

Email: mbehring@cisco.com