

ANIMA  
Internet Draft  
Intended status: Standards Track  
Expires: January 13, 2019

T.S.Choi  
T.S.Jeong  
ETRI  
J.K.Choi  
J.S.Han  
KAIST  
October 14, 2018

## Trust networking and procedures for Autonomic Networking

draft-choi-anima-trust-networking-01

### Abstract

This document describes trust networking as an application of autonomic networking. The objective of trustworthy autonomic networking is providing trust networking environment where all autonomic nodes can communicate without any security concern. It defines a trust networking domain and describes how to configure and maintain the trust networking domain. While communication within the trust networking domain is done with trust, the communication with external nodes should be done via a specific autonomic service agent (ASA) called "trust gateway". The trust gateway ASA performs trust evaluation of the external nodes and enforces domain specific policies to keep the domain trustworthy.

### 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."



Copyright Notice

Copyright (c) 2018 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.

1.	Introduction .....	4
2.	Background .....	4
2.1.	Security Model and its Limitations .....	5
2.2.	Trust Model and Trust Relations .....	6
2.3.	Comparisons of Security and Trust Model .....	7
3.	Trust Networking Framework .....	8
3.1.	Defining Trust Networking Domain .....	9
3.2.	Protecting Trust Networking Domain .....	9
3.3.	Expanding Trust Networking Domain .....	10
3.4.	Communicating with External Entities .....	11
4.	Differences between trust networking and ANIMA security framework .....	1
4.1.	Domain as a Whole .....	12
4.2.	Individual Nodes (Domain members) .....	13
4.3.	Domain Boundary.....	13
4.4.	Topology .....	14
4.5.	Technology .....	15
4.6.	Connection to the Internet .....	15
4.7.	Security, Trust and Privacy Model .....	16
4.8.	Operations .....	16
5.	Trust networking domain as an application of autonomic networking .....	1
5.1.	Definition of a Trust networking domain .....	18
5.2.	Configuration of Trust networking domain .....	19
5.3.	Communication between Trusted Autonomic Nodes within a trust networking domain .....	20
5.4.	Communication between trusted autonomic nodes and external nodes .....	20
6.	Trust Networking in the Autonomic Networking Infrastructure ..	21
6.1.	Identification of Trust networking domain and Trusted Autonomic Node .....	22
6.2.	Discovery of Trust networking domain .....	23
6.3.	Signaling Between Trusted Autonomic Nodes .....	23
6.4.	Trust Evaluation .....	24
7.	Procedures for trust networking .....	25
7.1.	Building a trust networking domain .....	25
7.1.1.	Domain initialization .....	25
7.1.2.	Node registration .....	26
7.2.	Evicting existing node from trust networking domain .....	



7.3. Terminating trust networking domain ..... 28

7.4. Communication among trust networking domains ..... 28

    7.4.1. Trustworthy networking within a single trust networking domain ..... 28

    7.4.2. Trustworthy networking between trust networking domains ..... 28

..... 28

8. Security Considerations ..... 30

..... 30

9. IANA Considerations ..... 31

10. Acknowledgements ..... 31

11. Contributors ..... 31

12. References ..... 31

    12.1. Normative References ..... 31

..... 31

    12.2. Informative References ..... 31

..... 31

## 1. Introduction

The document describes the concept of trust networking as an application of Autonomic Networking Architecture. It defines a trust networking domain in compliance with reference model of autonomic networking. By definition of autonomic domain [rfc7575 Autonomic Networking Definitions and Design Goals] the trust networking domain is defined as a collection of autonomic nodes which trust other nodes in the same trust networking domain. That means, communications within the trust networking domain with sufficient trust level can be done without any further security concerns. For example, assume that a subnet properly protected from external threats and all nodes in the subnet are verified through trust evaluation procedures, then the communications within the subnet can be done with confidence that nodes do no harm to each other.

This document first defines a trust networking domain and then describes how to configure the trust networking domain and keep the domain trustworthy. This document also describes a trust networking framework that consists of interconnected trust networking domains. The framework guides how to define the trust networking domain, how to manage members of the domain, how to protect the domain from hostile external world, how to expand the domain, and how to handle communications with external entities. Finally this documents shows how to apply the trust networking framework to the existing IP based network with minor modifications

## 2. Background

One of the biggest problems in the current Internet is protecting information assets against divergent attacks. In the beginning of





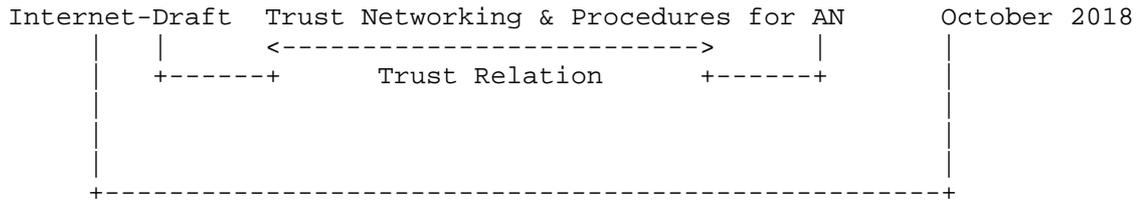


Figure 2. Trust Model

The trust relation used in the trust model is assumed to be reflexive, symmetric, and transitive. Reflexive means that entity A trust itself, denoting as AA. Symmetric relation assumes that two entities A and B satisfy AB and BA at the same time, denoting as AB. Transitive means that for three entities A, B, and C, if AB and BC then A C. If all entities in a given group satisfy all three characteristics, the group is declared as a trust equivalent class. We can easily guess the role of the trust model as formation of a trust equivalent class for the set of entities trusting each other.

The trust model should provide safe and reliable communication environment to entities without requiring additional security features on the entities. Thanks to the transitive trust relation, if an external entity is trusted by one member of the domain as a trust equivalence class, other members in that domain also can trust the external entity. By restricting the domain to trusted entities, the environment can be kept safe and reliable.

### 2.3. Comparisons of Security and Trust Model

The trust model is opposite in almost every aspect as shown in Table 1. First of all, the trust model is based on confidence that entities in a trust networking domain never do harm, while the security model is based on suspicion that adversaries attacks anytime. The relationship in trust model is binary in the sense that an entity trust another specific entity, but relationship in the security model is unary because the entity itself must protect regardless of other entities. With respect of rules, trust model keeps trusted IDs as a white list but security model keeps threatening entities as a black list. Thus, behavior of entities in the trust model is proactive while the security model acts in reactive manner. That leads the policy of the trust model is to prevent risk by communicating only with trusted entities, but policy of the security model monitors all communications to detect and remove threatening actions. The trust model provides mechanisms for

Internet-Draft Trust Networking & Procedures for AN October 2018  
 accepting entities or domains after verifying their trust, while the security model provides mechanisms for watching the traffic and blocking the threatening traffics. As the result, the network space of the trust model starts with a restricted space and incrementally grows as new entities or domains are accepted, while the network space of security model starts as an unrestricted and open space, but the space may be diminished by excluding misbehaving entities.

Table 1 Comparison of Trust and Security Model

	Trust Model	Security Model
based on	confidence	suspicion
relationship	binary	unary
rules	white list	black list
behavior	proactive	reactive
policy	prevention	detect and remove
mechanism	verify and accept	watch and block
network space	unrestricted and diminishing	restricted and expanding

### 3. Trust Networking Framework

The purpose of the trustworthy communication framework is to provide safe and reliable environment to entities without requiring additional security features. For keeping the environment trustworthy, the domain accepts only eligible entities. However, this restriction seems contradict to global scalability that requires the domain being open to everyone. Our solution is the incremental strategy, where a domain starts from a small and restricted network space and gradually expands to a global scale

Internet-Draft Trust Networking & Procedures for AN October 2018  
network space by accepting external entities or collaborating with  
other domains. This section discusses technical issues on the  
trustworthy communication framework.

### 3.1. Defining Trust Networking Domain

A primitive domain can be defined as the network space that is autonomous, isolated, and well protected from external attacks. For example, isolated home or enterprise network can be defined as a domain. If all hosts in the domain are disinfected and communication links are not exposed, the domain can be declared as a trust networking domain. The trust networking domain is not always a physical network space but sometime it can be formed by a logical group of users with mutual trust. In any case, the entities in the domain forms a trust equivalence class and communication with other entities in the domain is allowed without any protection.

To keep to domain trustworthy only qualified entities can be accepted as a member of the domain, and misbehaving entities have to be removed from the domain. For maintenance of a domain, the behavior of entities in the domain may be monitored, and if suspicious activities are discovered, the corresponding entity must be removed.

### 3.2. Protecting Trust Networking Domain

The domain representing an autonomous network space can take role of security unit as well as packet processing unit. The isolated domain from external world does not allow communication with external entities. For opening the domain to untrusty external world, well-defined interfaces are required to protect the domain. Let's call this protected domain an "insulated trust networking domain". As an example of insulated trust networking domain, we can imagine the local area network with firewalls on all links to the external Internet. The local area network is not isolated but is insulated from attacks injected through the external links.

The proposed framework assumes that each domain has at least one gateway that performs security functions for the domain. The gateway identifies external entities, evaluate trust level, accepts or rejects the packets according to the trust levels of external entities. And also the gateway will forward only authorized and sterilized packets to peer domain for keeping its reputation or trust level. In the sense that gateways performs security functions on the behalf of the entities inside of the domain, the security of

Internet-Draft Trust Networking & Procedures for AN October 2018  
 entities is said to be delegated to gateways. This delegated security has great benefit in applying complex security functions to devices with a limited or no processing power.

### 3.3. Expanding Trust Networking Domain

If all communications are limited within a trust networking domain, the serious scalability arises with respect to global communication. Now, we have to consider expansion of trust networking domain, starting from a small trust networking domain to a global scale network. First, consider the situation that an entity outside of domain tries to communicate with an entity inside of the domain. For trustworthy communication across border of domain, the entity must be a member of the domain. The domain gateway performs well-defined procedure for checking identity and evaluating the trust level of the external entity, and then only qualified entities are allowed to communicate with entities in the domain. Also the link connecting the domain with external entities should be secure enough for the trust level. This is one way to expand a domain.

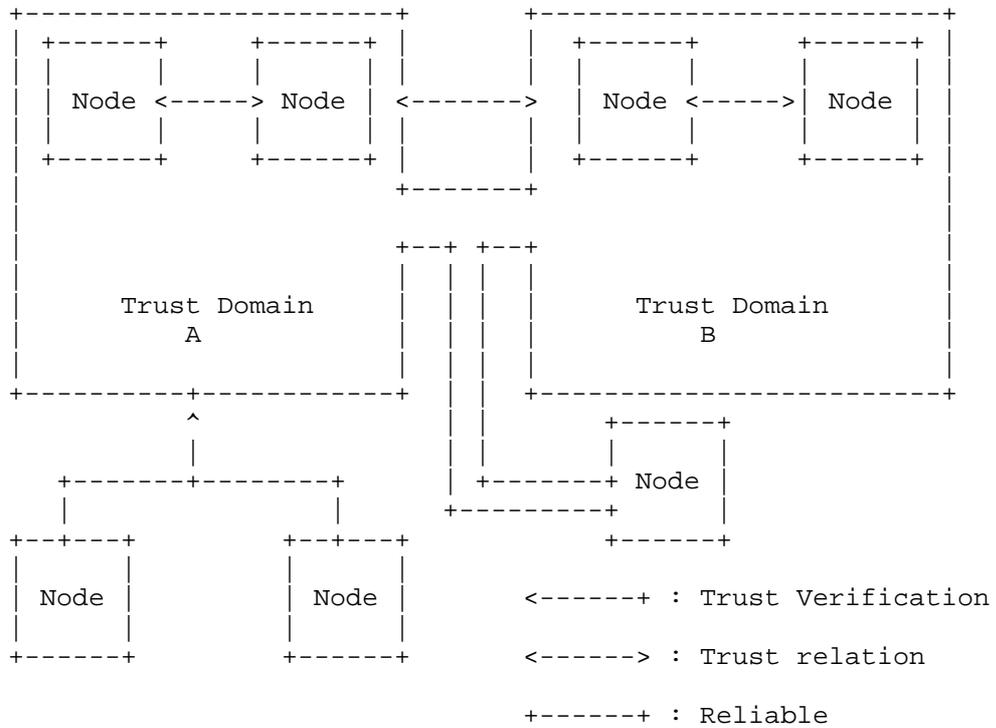


Figure 3. Expansion of Trust Model

Expanding a domain by accepting new entities has limitation when reaching the maximum number of entities being managed by a single domain. The other solution is collaboration of domains. Suppose two domains trust each other and those are connected by reliable links, then entities within one domain can trust entities within another domain.

Figure 3 shows a trust networking domain with trusted entities and 3 ways how to expand the domain. First, new entities can join to the domain after passing trust verification. Second, a remote entity can join to the domain via reliable channel. And third, when two domains may have trust agreement and connected by reliable channel, all entities in one domain can exchange packets in the pre-agreed trust level.

#### 3.4. Communicating with External Entities

As already seen, the communication inside of a domain requires no further security. However, communication with entities outside of the domain needs special care. Assume that all communication with external entities must take place at the special entity called a gateway, which enforce well-defined procedure communication for external entities. As explained in Section 3.3.2, an insulated trust networking domain has one or more gateways to perform trust verification for every packet injected to the domain. When a packet arrives at the gateway of a domain, the gateway first check whether the source ID of the packet is in the trusted ID list. If exists, the packet is accepted. Otherwise, the gateway lookups the trusted domain list to find sending domain of the packet. If the sending domain is in the list, the packet can also be accepted and ID of the packet is saved in the trust ID list. This mean that the gateway believes the trusted sending domain not to send harmful packets. If ID of the packet is not in the trusted ID list nor the sending domain is in the trust networking domain list, then verification procedure for individual ID has to be performed. The procedure is somewhat similar to accepting new entities in the domain. The overall procedure of a gateway is shown in Figure 4.

This section describes major differences between the proposed trust autonomic domain (TAD) and ANIMA security framework. The differences are explained based on a following set of criteria defined in the draft-carpenter-limited-domains-03: domain as a whole, domain members, domain boundary, topology, technology, connection to the Internet, security/trust/privacy model, and operation since our proposed domain and that of ANIMA are kinds of limited domains.

#### 4.1. Domain as a Whole

Networking is a very complex task and traditional way of handling the complexity is layering, where each layer takes a specific role and provides its services to the next higher later. This layering architecture decomposes the whole networking task functions vertically. However, the network in general spans physical or logical regions. Each region may have distinct features, such as different physical media, separate administration, and diverse networking requirement. The concept of domain in this document is defined as the networking region that shares common characteristics and also is distinguished from the rest of the network. Traditional layers cover its own regions implicitly; the physical layer spans the range covering electric signals. The data link covers the range connected by layer 2 bridges, and the network layer covers the whole devices connected by routers, and so on. Instead of implicit regions of the layers, a domain can be defined as any region of the network which is distinguishable from the rest of the network. It can be defined as a region covered by electric signal, a home network owned by a single user, a virtual private network overlaid on the Internet, a social network composed of members. Thus, it can be defined by any layer.

In the context of TAD, the domain can be defined by trust. That means all members within a TAD trust each other so that the members can communicate with others without any concern of security. For this, TAD needs to add an additional ASA which performs a role of domain administrator. Its main functionality is to manage trust policies including allocating trust level to domains and their members. Domain administrator can extend the functionality of ANIMA MASA or define a new ASA for the purpose of the domain administration. The details of domain administrator is specified in Section 5 below.

As defined in the previous section, the domain covers a specific region of the network, to where a set of nodes belongs. Since a domain shares common characteristics, any node within the domain must be able to communicate with other nodes in the domain. The node as a member of a domain can be host, networking devices, applications depending on the characteristics of the domain. For keeping the same characteristics, a node trying to be a new member of the domain must prove its functionalities to all or a designated member of the domain. Joining to a domain may be accomplished by simply plugging interfaces to the networking device or well-defined interactions enforced by domain administrator. The joining procedure may be implicit when a domain has fixed and permanent members, or explicit in case that a node can join or leave the domain.

In the sense of TAD, a node is assumed as a host that has communication functions required by the domain. Since a TAD is defined under the intent of trust, a node should have identifiable and authenticatable ID. TAD utilizes a concept of self-certifying ID. The self-certifying ID can be newly defined. However, in the context of TDA as an application use case of ANIMA, we can utilize IdevID as a self-certifiable ID and preferably extend IdevID with public key information as an option to ensure the global uniqueness.

#### 4.3. Domain Boundary

Since a domain is a set of nodes that shares common characteristics, only nodes within a domain can communicate. In other words, a node within a domain cannot communicate with nodes outside of the domain. However, we can assume special nodes that belongs multiple domains simultaneously. Let's call a node joining more than two domains a "gateway". A gateway node must be equipped with multiple functionalities, each for the joined domain. The role of gateway is conveying interactions of one domain to other domains. Of course, conveying interaction may include necessary functions such as interpretation, filtering, transformation etc. From outside of a domain, the internals of the domain is hidden and the boundary of the domain composed of gateways are only exposed. All interactions passing the boundary of a domain must performed by at least one of the gateways whose role is to enforce necessary gatewaying procedures.

In the context of TAD, all members of a TAD trust each other, but cannot trust nodes outside of the domain. The only way for an internal node to communicate with external nodes is passing through a gateway of the domain. Once the gateway receives communication request from a node outside of the domain, it authenticates the node and evaluates the trustworthiness of the node. If the external node is trustworthy and communication channel between gateway and the node is safe and reliable enough for the domain trust level, the gateway accepts communication and injects the communication possibly with transformation. Unlike ANIMA which assumes IP based communications by every domains, TAD may allow any networking technology besides IP. Therefore, a gateway is a mandatory component where the need for it is implicit in ANIMA due to the homogenous nature networking technology used in a domain. The details of domain gateway functionality is specified in Section 5 below.

#### 4.4. Topology

As defined in Section 4.1, a domain is a range of network where all members can communicate. The communication can be done in either specific layer protocols or any common functionalities. For example, if domain is defined by local area network, the domain may use local IP addresses, link-local or site-local. For domains defined by virtual network overlaid on global Internet may use global IP addresses with filtering functions.

As already explained in section 4.3, some special nodes may belong to multiple domains. In this case the range of the domains that involve the same nodes can be viewed as overlapped domains. The node belonging multiple domains should have multiple functionalities, one of each domain. Those functionalities should be separated. We can find similar situation in multi-homed IP host in the Internet, where the host has separate IP addresses, one for each IP address domain.

In the context of TAD, domains also have self-certifying ID as an ordinary node to become a member of another domain. The domain administrator must take a role of the required procedures of the parent domain such as trust evaluation, join and leave. Also the

#### 4.5. Technology

In the context of TAD, any technology is allowed for the domain since a domain has its own mechanisms hidden from outside. Apart from the existing Internet using global IP addresses, each domain may use its own routing or forwarding mechanisms, such as Ethernet, MPLS, or Upper-Layer IDs. Only requirement for inter-domain communication is that the gateway must aware of mechanisms for both domain and takes a role of translation. Note that each domain has a domain specific addressing scheme and identification of nodes/domains must be done by globally unique identifier. With global ID a node can join a domain or move from one domain to another. In this case a node acquires a domain specific address when joining the domain.

#### 4.6. Connection to the Internet

In the context of TAD, the existing Internet can be viewed as a huge domain with global coverage. Nodes or domains with IP capability can join the global Internet domain as members. Since the existing Internet has no notion of ID, let us assume the global Internet domain top-level domain where every domain can join. Each domain with its specific mechanism can join the global Internet domain permanently or intermittently. The communication from one domain to another domain through the global Internet domain is done by the normal IP communication. However, the gateway of each domain must translate its internal communication mechanism to that of the corresponding IP address communications. More specifically, Inter-domain communication is done by global ID and the ID is translated into domain-specific address when passing the domain boundary. This ID based communication may be encapsulated in IP packet when traversing the global Internet domain. To allow this translation, the ID to IP address mapping system must be provided, where IP address is the gateway address of the domain that involves the node with the ID.

One of implication of a domain is secure protection of the domain internals from the rest of the network. That is members of a domain should be identified, authenticated, and authorized. According to domain's policies, well-defined procedures must be enforced to a node to become a member of the domain.

In TAD all members of the domain must have the same or higher trust level than the domain requires. That means, whenever a new node tries to be a member of the domain or an external node tries to communicate with an internal node, the domain administrator must authenticate and evaluate the node. Only the node passing the evaluation procedure is allowed to communicate. In this case communication must be done via channels safe and reliable enough for the trust level. In some cases where the channel is not safe nor reliable, the communicating nodes must authenticate or encrypt the traffic. Note that whether the traffic is protected or not depends on the risk level of the channel and trust level of the domain. Unlike the VPN that protects all channels in the same security protocols, channels for a domain are additionally protected only when the risk level of a specific channel is higher than required.

#### 4.8. Operations

In addition to trust relation between nodes within a domain, the environment of the domain must be considered. Environment of a domain includes factors affecting domain operation such as communication channels among nodes, operation skills of domain administrator, reliability of devices, etc. To be protected from the rest of networks, a domain should be securely protected from external attacks.

Since communications within a TAD are carried out on the mutual-trust basis, the domain administrator should keep the domain trustworthy by accepting only trusted members, monitoring traffic to detect suspicious behavior, and periodic auditing the logs of domain members, and so on.



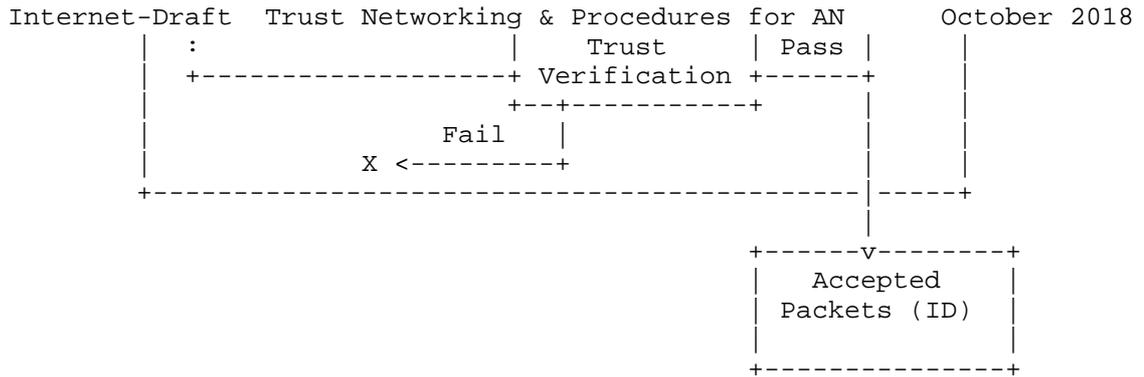


Figure 4. Packet Processing at the Gateway

5.1. Definition of a Trust networking domain

A trust networking domain is defined as a collection of autonomic nodes trusting each other. Since all nodes within a trust networking domain maintains certain trust level set by the domain, communications within the domain can be done without any further security concern. However, communications with external node require additional verification phase before the communications actually begin. The verification is performed at the border of the domain, where external nodes are checked if their trust level are sufficiently high for the domain. In the sense that the domain as a collection of node are protected from external world, it seems "zone defense" rather than "individual defense" of the traditional security scheme.

Figure 5 shows the high-level architectural view of trust networking domain. Autonomic nodes has the interface with management function. Trust management functions define the trusted autonomic nodes according to their trust level. They also define the trust networking domain by grouping or classifying autonomic nodes. At the same trust networking domain, an autonomic node directly communicates with each other. The control and management functions at the trust networking domain are defined at the interfaces between autonomic nodes and management plane. There are trust gateway for an autonomic node to communicate with different trust networking domains or non-autonomic nodes since there is no direct communication path. Trust gateway is used to communicate autonomic nodes with different trust networking domains

Internet-Draft Trust Networking & Procedures for AN October 2018  
 or the non-autonomous nodes. An autonomous node can communicate remote  
 autonomous nodes or non-autonomous nodes through trust gateway. In  
 these cases, the traditional trust evaluation and/or certificate  
 procedures can be applied at trust gateway. Trust evaluation  
 procedure is running by management plane of autonomous networking.

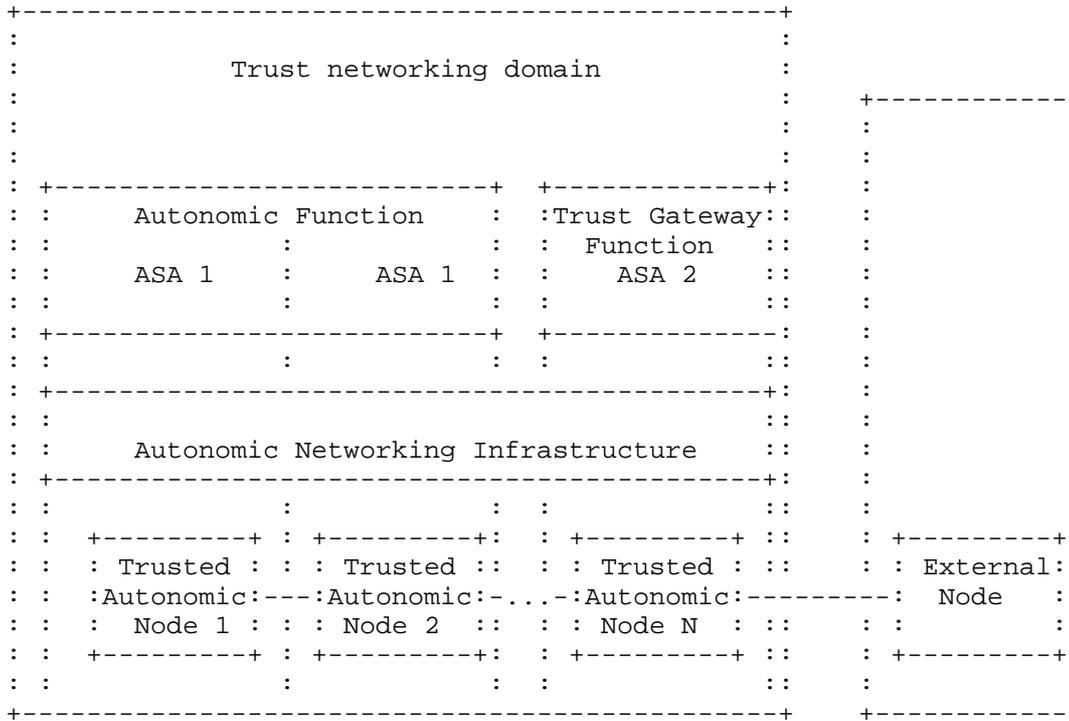


Figure 5. Trust networking domain at the Autonomous Networking

## 5.2. Configuration of Trust networking domain

A trust networking domain is consisted of a group of autonomous nodes. The network management plane communicates with a list of autonomous nodes to build the trust networking domain. The trust management information database which contains a list of autonomous nodes according to the trust level of each domain is built at the bootstrapping time or at the instance of request.

At the bootstrapping time, the management plane securely distributes the trust information of each domain to the corresponding autonomic nodes. The membership management is done by management plane when the autonomic nodes can be joined to or leaved from each trust networking domain.

At the instance that an autonomic node request to build a trust networking domain to the management plane, trust management function confirm to build a trust networking domain after completing the proper trust evaluation procedures.

If an autonomic node could not continue to be a member of the certain trust networking domain, it notify to management plane for leave. Similarly, if the trust management functions decide that an autonomic node is not relevant to stay in a certain trust networking domain, they notify the corresponding autonomic node for leave and update the trust management information database.

Within a trust networking domain, an autonomic node can communicate each other without any additional security and certificate procedure. In a case, an autonomic node may register multiple trust networking domains simultaneously.

### 5.3. Communication between Trusted Autonomic Nodes within a trust networking domain

At the same trust networking domain, autonomic nodes directly communicate with each other. Autonomic nodes can discover other nodes at the same trust networking domain. It requires control or management information between autonomic nodes and control/management plane. It can be pre-configured during bootstrapping. The control information between autonomic nodes can be used to identify the trust networking domain. The autonomic nodes can easily communicate with each other at the same trust networking domain by enabling self-managing capability of autonomic networking. The autonomic service agents can be implemented for trusted communication.

### 5.4. Communication between trusted autonomic nodes and external nodes

Autonomic nodes must communicate with autonomic nodes of the different trust networking domain. They also communicate with the non-autonomic nodes.

Trust gateway can help that an autonomic node communicate with the autonomic nodes with different trust networking domain or the non-autonomic nodes. Some autonomic service agents (ASA) may include the trust gateway functions for communicating autonomic nodes with different trust networking domain, which is in the reference model for Autonomic Networking [I-D.ietf-anima-reference-model].

## 6. Trust Networking in the Autonomic Networking Infrastructure

This section describes trust networking of autonomic network. Within a trust networking domain, an autonomic node is credited by their trust level from management plane.

The trust management plane maintains the trust information tables up to date. The trust management plane is tracking of trust status of each autonomic node as an application of autonomic networking. The trust information table contains the trust information of autonomic nodes based on the trust networking domain. All the interactions between autonomic nodes should be verified according to trust evaluation procedures of management plane.

The autonomic nodes within the same trust networking domain create and maintain network connectivity without additional complexity. Trust provisioning among autonomic nodes is to exempt any additional processing (like identification, addressing, routing, forwarding, and security, etc.) to maintain autonomic networking within the same trust networking domain.

The interactions between autonomic nodes are based on the trust evaluation of the trust networking domain. The trust information is used to leverage the direct interactions between autonomic nodes. Trust gateway can help to the interaction of autonomic nodes with different trust networking domains or with non-autonomic nodes.

The trust management plane is used to handle the trust level of each autonomic node with proper trust evaluation procedure.

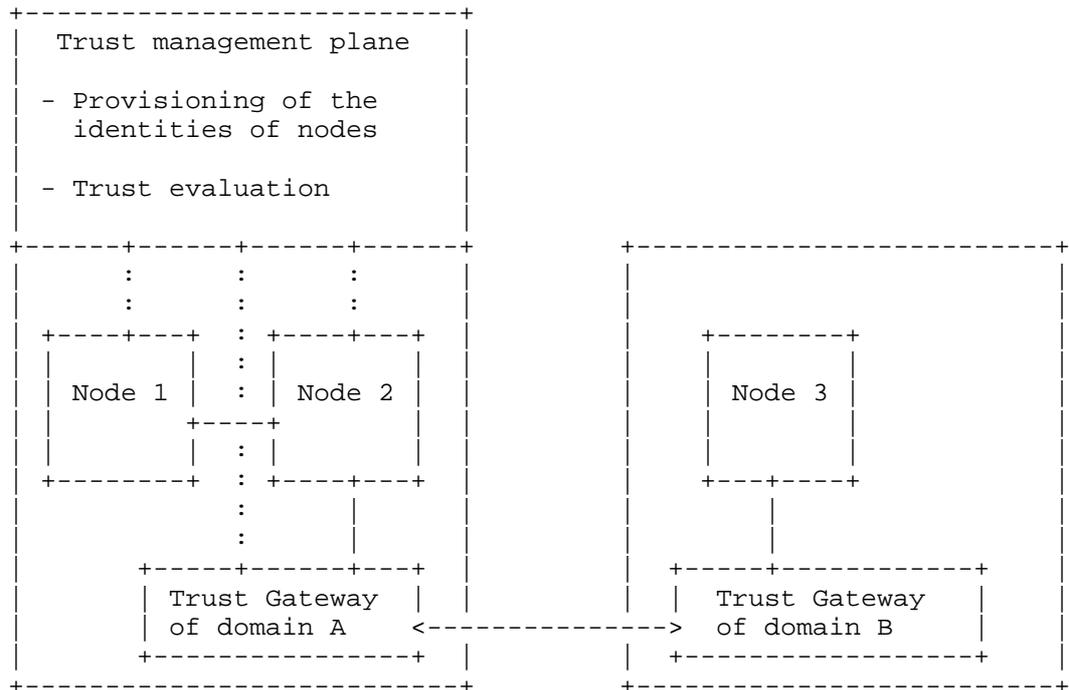


Figure 6. Trust provisioning at the Autonomic Networking

### 6.1. Identification of Trust networking domain and Trusted Autonomic Node

This section describes trust level. An autonomic node can initiate to create their own trust networking domain. The management plane provides that an autonomic node can build the relevant trust networking domain by identifying the corresponding autonomic nodes. Specific policies can be applied to build trust networking domain.

In a trust networking domain, each autonomic node should be identified by the relevant naming and addressing schemes, which are also compliant with the Reference Model for Autonomic Networking [I-D.ietf-anima-reference-model]. Before data exchange, the autonomic nodes obtains the identities (e.g., IP address and port number,

etc.) of destination nodes and the corresponding trust networking domain. In a case, the MAC address can be also used for identification.

The trust management information database is used for the discovery of autonomic nodes at the same trust networking domain. The autonomic nodes with the same trust networking domain may use the relevant identification schemes. In the trust management information database, a list of autonomic nodes are classified into the relevant identification code which indicates the same trust networking domain. The identification code for a trust networking domain may contain name/nickname and number as well as IP address and port number, etc.

## 6.2. Discovery of Trust networking domain

The trust management information database is used for the discovery of autonomic nodes at the same trust networking domain. Before data exchange, an autonomic node looks up the trust management information database to find the destination autonomic nodes. If the destination node belongs to the same trust networking domain with original autonomic node, it is possible to initiate data exchange.

## 6.3. Signaling Between Trusted Autonomic Nodes

At the same trust networking domain, an autonomic nodes communicate with each other. For data exchange, the autonomic node should discover each other by accessing the trust management information database of management plane.

After discovery of destination autonomic node, the signaling protocol like "A Generic Autonomic Signaling Protocol (GRASP)" [I-D.ietf-anima-grasp] are needed to initiate data exchange. Within the same trust networking domain, an autonomic node directly communicates with each other after completing signaling procedure, in which the connectivity among autonomic nodes are securely and automatically maintained. The pre-configuration between autonomic nodes can be done during bootstrapping. The autonomic control plane at the Reference Model for Autonomic Networking [I-D.ietf-anima-reference-model] can be either implemented to carry signaling protocol.

For data exchange with different trust networking domains or non-autonomic nodes, the trust gateway provides proper interworking

functions for data exchange and signaling since there is no direct communication paths between them. The trust gateway provides the relevant control and management information to extend data exchange with different trust networking domains or non-autonomic nodes. The authentication and certificate procedures equivalent with the trust networking domain can be applicable to provide external connectivity.

#### 6.4. Trust Evaluation

Trust evaluation of network is the way of calculating trust for networking services. It requires data collection from various sources. Physical data sources are collected from the capability of data processing, storage, and communication through network. In cyber world, logical data sources are software that work on computing algorithm, storage, and networking. In the social world, human produces various data through user interfaces.

In the physical network, trust can be measured by counting on their trustworthiness of network elements. In the cyber world, software can be accidentally or maliciously altered or destroyed during control, computing, and communicating instances. The unexpected behaviors of software is detected or monitored to evaluate and update their trust level. In the social world, human behaviors can be measured by considering its trustworthiness in terms of ability, honesty and benevolence. Social trust reflects individual human activity. Human interacts with others honestly and kindly so that their trust level is affected by some risks.

For trust evaluation, the collected data are categorized into two types of attributes and indicators namely, qualitative and quantitative. Trust index is used to calculate the certain trust level of each network entity. As the results of trust evaluation, trustor finally make a decision. The network management plane provides to calculate the trust level of the network elements from various data sources and store their values to trust management information database.

The trust management information contains the trust level of autonomic nodes. The interactions inside a trust networking domain are analyzed and accumulated to evaluate the trust level of each node. The trust level of autonomic node is contained at the trust

Internet-Draft Trust Networking & Procedures for AN October 2018  
management information database. All the interactions between  
autonomic nodes in a same trust networking domain is validated by  
the trust evaluation procedure.

The trust evaluation procedure is fed by the following inputs.

- o Pre-provisioned or manually configured by policy or management information
- o Analysis from interactions between autonomic nodes
- o The accumulated history information of trust verifications such as authentication of non-autonomic nodes and validity of application specific transactions.
- o other unaccepted or unexpected behaviors

While autonomic nodes communicate with each other, they choose the relevant trust management protocol whether they meet trust requirements in the same trust networking domain or not. Trust management protocol between autonomic nodes and trust management database is needed to check trust evaluation. Trust evaluation procedure between autonomic nodes at same trust networking domain are taken for trust identification.

If the prerequisite and pre-configuration procedures are already taken for trust management, simple and light-weight solution can be applicable for communication between autonomic nodes.

## 7. Procedures for trust networking

### 7.1. Building a trust networking domain

#### 7.1.1. Domain initialization

To build a new trust networking domain, the domain administrator needs to initiate the functionalities of trust networking domain as follows:

##### - Domain administration

To initialize a domain with respect to the trust, the domain administrator needs to configure policies of trust and membership. To manage the trust level, the domain administrator sets the



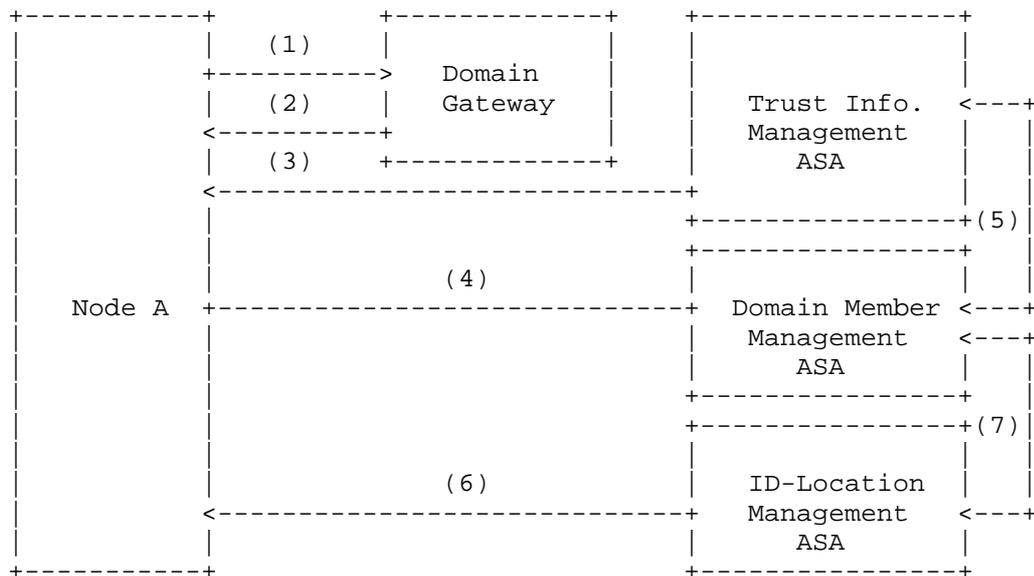


Figure 9. Procedures of node registration

- (1) Node A connects to the network of trust networking domain;
- (2) The domain assigns a private IP address to Node A. The domain gateway is assigned as the default gateway for IP network;
- (3) Trust information management ASA analyses the trust information of node A;
- (4) Node A request to join the domain;
- (5) Domain membership management ASA of the domain administrator receives the requests and decides to approve Node A, based on the domain policy and trust level of Node A;
- (6) ID-Location management ASA of the domain administrator issues a new identifier of Node A;
- (7) ID-Location management ASA archives Node A's identifier and private IP address.

## 7.2. Evicting existing node from trust networking domain

(Editors' note) This section describes how to evict existing node in trust networking domain including trust management procedures. Further details are for further study.

## 7.3. Terminating trust networking domain

(Editors' note) This section describes how to terminate trust networking domain including signalling procedures with child nodes (or domains) and parent domains. Further details are for further study.

## 7.4. Communication among trust networking domains

This section describes trustworthy communication between nodes within a single trust networking domain and between nodes separated into multiple trust networking domains.

### 7.4.1. Trustworthy networking within a single trust networking domain

In order for the two hosts to send and receive messages to each other, a networking path must first be established. If two hosts are located in the same domain, they already have trust relationship with each other which means no additional security procedures are needed.

### 7.4.2. Trustworthy networking between trust networking domains

Two hosts are in different domains. It means that they do not know each other's IP address directly. The domain administrator provides IP address of each hosts for trustworthy networking between two hosts in different domains. If a Host 2 wants to perform trustworthy networking with a Host 1 in other domain, it is possible to establish a networking path between two nodes through interactions between domain administration functions and access and delivery control functions. Figure 10 shows an overview of trustworthy networking between trust networking domains.



- (1) Host 2 requests IP address of Host 1 to the domain administration ASA 2 through the ID of the host 1;
- (2) The domain administration ASA 2 requests IP address of the Host 1 to the domain administration ASA 1;
- (3) The domain administration ASA 1 obtains IP address of the Host 1 and reply ID and IP address of the Host 1 to domain administration ASA 2, and it replies to Host 2;
- (4) Host 2 requests a trust level of Host 1 through the domain administration ASA 2;
- (5) The domain administration ASA 2 checks a trust level of Host 2 through the trust information management ASA and requests a trust level of Host 1 to domain administration ASA 1;
- (6) The domain administration function 1 obtains the trust level of Host 1 through the trust information management ASA and replies it to the domain administration ASA 2, and the result replies to Host 2;
- (7) The access and delivery control ASA 2 forms a routing path with the access and delivery control function 1 through the ID-based routing ASA;
- (8) The Host 2 and the Host 1 establish a reliable link through the domain gateway ASA of each trust networking domain;
- (9) Networking path established between Host 1 and Host 2.

#### 8. Security Considerations

Data exchange between autonomic nodes at the trust networking domain must be secured. The signaling or management protocols for trust identification and discovery of trust networking domain are secure. The control/management plane for trust management is self-protecting. The autonomic node in a trust networking domain should be certified by its identity. The pre-configuration information of autonomic nodes from trust management information database should be certified during bootstrapping time.

For data exchange with different trust networking domain or non-autonomic network, the trust gateway should be securely implemented. Trust gateway maintains the same trust level for cross-domain applications or interaction with non-autonomic network.

This document requests no action by IANA.

## 10. Acknowledgements

## 11. Contributors

## 12. References

### 12.1. Normative References

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

Eckert, T., Behringer, M., and S. Bjarnason, "An Autonomic Control Plane (ACP)", draft-ietf-anima-autonomic-control-plane-13 (work in progress), December 2017.

[I-D.ietf-anima-grasp]

Bormann, C., Carpenter, B., and B. Liu, "A Generic Autonomic Signaling Protocol (GRASP)", draft-ietf-anima-grasp-15 (work in progress), April 2018.

[ITU-T Y.3052] Overview of trust provisioning in information and communication technology infrastructures and service, March 2017

[ITU-T Y.3053] Framework of trustworthy networking with trust-centric network domains, January 2018

### 12.2. Informative References

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

Behringer, M., Carpenter, B., Eckert, T., Ciavaglia, L., Pierre, P., Liu, B., Nobre, J., and J. Strassner, "A Reference Model for Autonomic Networking", draft-ietf-anima-reference-model-08 (work in progress), February 2018.

Tae Sang Choi  
Electronics and Telecommunication Research Institute (ETRI)  
218 Gajeong-ro, Gajeong-dong, Yuseong-gu, Daejeon  
Korea

Email: choits@etri.re.kr

Jun Kyun Choi (editor)  
Korea Advanced Institute of Science and Technology (KAIST)  
193 Munji Ro, Yuseong-gu, Daejeon  
Korea

Email: jkchoi59@kaist.ac.kr

Tae Su Jeong  
Electronics and Telecommunication Research Institute (ETRI)  
218 Gajeong-ro, Gajeong-dong, Yuseong-gu, Daejeon  
Korea

Email: tsjeong@etri.re.kr

Nam Seok Ko  
Electronics and Telecommunication Research Institute (ETRI)  
218 Gajeong-ro, Gajeong-dong, Yuseong-gu, Daejeon  
Korea

Email: nsko@etri.re.kr

Jae Seob Han  
Korea Advanced Institute of Science and Technology (KAIST)  
193 Munji Ro, Yuseong-gu, Daejeon  
Korea

Email: j89449@kaist.ac.kr

ANIMA WG  
Internet-Draft  
Intended status: Standards Track  
Expires: January 2, 2019

T. Eckert  
Huawei  
Jul 1, 2018

DNS-SD compatible service discovery in GRASP  
draft-eckert-anima-grasp-dnssd-01

Abstract

DNS Service Discovery (DNS-SD) defines the common framework for applications to announce and discover services. This includes service names, service instance names, common parameters for selecting a service instance (weight, priority) as well as service specific parameters.

GRASP is intended to also be used for service discovery. Reinventing service discovery for GRASP with a similar set of fetures would result in duplication of work. Therefore, this document defines how to use GRASP to announce and discover services in a way that inherits DNS-SD features and also tries to be compatible in spirit as much as possibel while still maintaining the intended simplicity of GRASP.

The goal of this document is to permit defining service and their parameters once and then use that in GRASP, mDNS and (unicast) DNS. Future work can also define DNS-SD <-> GRASP gateway functions.

In support of service discovery, this document also defines name discovery and schemes for reuseable elements in GRASP objectives which are designed to be extensible so that future work that identifies elements required across multiple objectives do not need to define a scheme how to do this.

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 <https://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 January 2, 2019.

#### Copyright Notice

Copyright (c) 2018 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 (<https://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. Overview . . . . .	2
2. Specification (Normative) . . . . .	4
2.1. Service and Name Objectives . . . . .	4
2.2. Objective Value Reuseable Elements Structure . . . . .	4
2.3. Reuseable Elements . . . . .	6
2.3.1. Sender Loop Count . . . . .	6
2.3.2. Service Element . . . . .	6
2.3.3. Name Element . . . . .	9
3. Explanations (Informative) . . . . .	11
3.1. Using GRASP service announcements . . . . .	11
3.2. Further comparison with DNS-SD . . . . .	13
3.3. Open Issues . . . . .	13
4. Security Considerations . . . . .	14
5. IANA Considerations . . . . .	14
6. Acknowledgements . . . . .	15
7. Change log [RFC Editor: Please remove] . . . . .	15
7.1. 01 - . . . . .	15
7.2. 00 - Initial version . . . . .	15
8. References . . . . .	15
8.1. Normative References . . . . .	15
8.2. Informative References . . . . .	15
Author's Address . . . . .	15

#### 1. Overview

DNS Service Discovery (DNS-SD) defines the common framework for applications to announce and discover services. This includes service names, service instance names, common parameters for

selecting a service instance (weight, priority) as well as service specific parameters.

GRASP is intended to also be used for service discovery. Reinventing service discovery for GRASP with a similar set of fetures would result in duplication of work. Therefore, this document defines how to use GRASP to announce and discover services in a way that inherits DNS-SD features and also tries to be compatible in spirit as much as possibel while still maintaining the intended simplicity of GRASP.

The goal of this document is to permit defining service and their parameters once and then use that in GRASP, mDNS and (unicast) DNS. Future work can also define DNS-SD <-> GRASP gateway functions.

GRASP exists as so-called GRASP-Domains, which are networks across which GRASP is run. This document primarily defines how to perform service discovery across such a domain leveraging GRASPs options to perform unsolicited flooding of announcements or flooding of requests and finding the closest service instances. The initial use case of this document is to support what in DNS-SD is done via mDNS but in larger networks - GRASP-Domains. Beside the efficient flooding, GRASP provides reliability and security (depending on the so called substrate used by GRASP, such as the autonomic control plane - ACP). Providing compatibility with existing mDNS service announcer or clients is possible, but not described in this version of the document.

The encoding of information choosen in this document does not try to use GRASP solely as a transport layer, but to also leverage the CBOR structure of GRASP messages to natively encode the message elements required for services in a way that is most simple - instead of using GRASP only as e.g.: an encapsulation of otherwise unchanged DNS message encodings. This is done to minimize the amount of coding required (and not require any DNS code unless future gateway functions are requireed), to increase the simplicity, minimize the amount of data on the wire and allow easier extensibility. On the downside, the mechanisms provided here do not cover the whole slew of possible options of DNS/DNS-SD, but instead only those deemed to be required. Others can be added later.

In support of service discovery, this document also defines name discovery and schemes for reuseable elements in GRASP objectives which are designed to be extensible so that future work that identifies elements required across multiple objectives do not need to define a scheme how to do this.

## 2. Specification (Normative)

### 2.1. Service and Name Objectives

Unsolicited, flooded announcements (M\_FLOOD) in GRASP and solicited flooded discovery (M\_DISCOVERY) operate on the unit of GRASP objective-names. Therefore a scheme is required to indicate services via objective-names. Note: future work may want to reuse the encodings related to services (defined below in this document) inside other (multicast or unicast only) objective exchanges, in which case the service names are not impacted.

When an objective is meant to be solely about a service name as defined and registered according to RFC6335, the objective MUST use an objective-name of SRV.<service-name>. This naming scheme allows to avoid creating duplicate and potentially inconsistent registration of names for those objectives vs. registrations done for example for DNS-SD. The primary use case for this naming scheme are therefore service names that are intended to be used in both DNS-SD and GRASP.

When an objective is meant announcement and discovery of a DNS compatible <name> such as "www-internal" in "www-internal.example.com", the objective SHOULD use an objective-name of NAME.<name>. See Section 2.3.3 for more details.

See Section 5 for the detailed IANA asks relating to these definitions.

### 2.2. Objective Value Reuseable Elements Structure

Because service discovery, as explained in the prior section, needs to utilize different objectives, it requires cross-objective standardized encoding of the elements of services. GRASP did not define standardized message elements for the message body (called "objective-value") of GRASP messages. Therefore, this document introduces such a feature.

[RFC-editor: please remove all occurrences of XXXX in rfcXXXX with the RFC number assigned to this document and remove this edit note.]

```
objective-value /= { 1*elements }
elements        // = ( @rfcXXXX: { 1*relement } )

relement = ( relement-codepoint => relement-value )
relement-codepoint = uint
relement-value      = any
```

If an objective wants to use reusable elements, the objective-value MUST be a CBOR map and the reusable elements are found under the key "@rfcXXXX". Objectives that do not want reusable elements as defined here can use any objective-value format including a CBOR map, but they can not use the "@rfcXXXX" key if they use a map. This approach was chosen as the hopefully least intrusive mechanism given how by nature all of "objective-value" is meant to be defined by individual objective definitions.

The value of "@rfcXXXX" is a map of reusable elements. Each relement has an IANA registered element-name and codepoint (see Section 5). The element-name is for documentation purposes only, CBOR encodings only use the numeric codepoint for encoding efficiency to minimize the risk for this solution to not be applicable to low-bitrate networks such as in IoT.

Format and semantic of the relement-value is determined by the specification of the reusable element as is the fact whether more than one instances of the same reusable element are permitted.

Reusable elements SHOULD be defined to be extensible. The methods used depend on the complexity of the element and the likely need to extend/modify the element with backward or non-backward compatible information. The following is a set of initial options to choose from:

Element values that are a map MUST permit and reserve key value 0 (numerical) for private extensions of the element defined by the individual objective.

Element values that are a map MUST NOT use bareword key values starting with a "\_". These too are for private extensions defined by the individual objective.

Element values SHOULD be defined so that additional keys in maps and additional elements at the end of arrays can be ignored by prior versions of the definition. Whenever a newer definition is made for an element where this rule is violated, the element SHOULD be changed

in a way for older version recipients to recognize that it is not compatible with it.

One method to indicate compatibility is a traditional version "`<major>.<minor>`". Within the same `<major>` version number, increasing `<minor>` version numbers must be backward compatible. Different `<major>` version numbers are not expected to be compatible with each other. If they are, then this can be indicated by including multiple version numbers.

A compressed form of version compatibility information is the use of a simple bitmask element where each bit indicates a version that the represented data is compatible with.

## 2.3. Reuseable Elements

### 2.3.1. Sender Loop Count

```
relement-codepoint ::= ( &(sender-loop-count:1) => 1..255 )
```

Sender-loop-count is set by the sender of an objective message to the same value as the loop-count of the message. On receipt, distance = ( sender-loop-count - loop-count ) is the distance of the sender from the receiver in hops. This element can be used for informational purposes in M\_FLOOD and M\_DISCOVERY messages and may be required to be used in these messages by the specification of other elements (such as the service element described below). This element MUST occur at most once. If a receiver expects to use the distance but sender-loop-count was not announced, then distance SHOULD be assumed to be 255 by the receiver.

### 2.3.2. Service Element

The `srv-element` (service element) is a reuseable element to request or announce a service instance or to request and list service instance names.

```
relement-codepoint // = ( &(srv-element:2) => context-element )
```

```
context-element = {
    ?( &(private:0)           => any),
    ?( &(msg-type:1)         => msg-type),
    ?( &(service:2)          => tstr),
    *( &(instance:3)         => tstr),
    ?( &(domain:4)           => tstr),
    ?( &(priority:5)         => 0..65535 ),
    ?( &(weight:6)           => 0..65535 ),
    *( &(kvpairs:7)          => { *(tstr: any) },
    ?( &(range:8)            => 0..255 ),
    *( &(clocator:9)         => clocator),
}
```

```
clocator = [ context, locator-option ]
```

```
context = cstr
```

```
locator-option = ; from GRASP
```

```
msg-type = &( describe: 0, describe-request:1,
              enumerate:2, enumerate-request:3)
```

**Service:** A service name registered according to RFC6335. If it is not present, then objective-name MUST be SRV.<service-name> where <service-name> is the service-name.

**Instance:** The <Instance> of a DNS-SD Service Instance Name ( <Instance> . <Service> . <Domain> ). It is optional, see Section 3.2.

**Domain:** The equivalent of the <Domain> field of a DNS-SD Service Instance Name. If domain is not present, this is equivalent to ".local" in DNS (as introduced by mDNS) and implies the unnamed "local" domain, which is the GRASP domain across which the message is transmitted.

**Priority, Weight:** Service Instance selection criteria as defined in RFC2782. If either one is not present, its value defaults to 0.

**Kvpairs:** Map of key/value pairs that are service parameters in the same format as the key/value pairs in TXT field(s) of DNS-SD TXT records as defined in RFC6763, section 6.3.

**Range:** Allows to flexibly combine distance and priority/weight based service selection according to the definition of distance in Section 2.3.1.

If min-distance is the distance of the closest service announcer, and min-range the range announced by it, then the recipient MUST consider the priority/weight of all service announcers that are not further away than (min-distance + min-range). If not included, range defaults to 255.

If range is announced, the sender-loop-count element MUST also be announced.

**Clocator:** The "contextual locator" allows to indicate zero or more locators for the indicated service instance. The context element indicates in which context the locator-option is to be resolved. The reserved context value of "" (empty string) indicates the GRASP domain used, aka: the "local" context in which the service announcement is made. The reserved context value of "0" indicates the default routing context of the announcing node. This is often called "global table", "VRF 0" or "default VRF" on nodes using the "VRF" abstraction. Any other value is a string specifying a context such as another VRF.

The mechanism by which originator and recipient of the srv-element agree on common naming for contexts is outside the scope of this specification. The context therefore allows to indicate locators both for the context through which the GRASP message distributed the srv-element (GRASP domain) as well as that for other contexts. Assume the GRASP domain is the ACP, then clocators in ACP would have a context of "", clocators in the global routing table (part of the data-plane) a context of "0", and clocators on other VRFs (also part of data-plane) a clocator that is their string name.

If no locators are indicated, then the locator of the service(s) is the optional locator-option of the GRASP message in which the objective is contained meant to be used for the service(s) indicated and the clocator implied is "".

If locator(s) are indicated, the messages location-option must be ignored for the service (but may be necessary to be present for other purposes of the objective).

**Msg-type** Type (aka: intention) of the srv-element. If not present, it is assumed to be "describe".

**Describe:** Describes one service instance. At least one clocator is required for a positive response, all other fields are permitted, but optional. "Describe" is used in M\_FLOOD for unsolicited announcements of services (flooded), in M\_RESPONSE messages for solicited announcements of a service and in M\_NEGOTIATE for negotiated announcements (both unicasted). If clocator is not

included, then all fields except service and instance (and msg-type and private) must not be included and the srv-element provides a negative reply: No information about this service/service instance. This is only permitted in unicasted "describe" messages.

**Describe-request:** Request for a "describe" reply. It is used in M\_DISCOVERY (flooded) for solicited discovery of services or in M\_REQ\_SYN (unicasted) for negotiated discovery of service instance(s). In "describe-request", only service is mandatory (but can be provided via the objective-name field of the message), and domain is optional. "Instance" is optional. If provided, then the recipient is asked to provide information about the named instance only. All other fields of srv-element are to be ignored by the receiver in this specification, but a semantic for setting them may be introduced in followup work, specifically to filter replies by the indicated fields.

"Describe-request" without instance MAY be answered by "Enumerate" (see below) if the responder has so many instances that it thinks the initiator should rather first select one or fewer instances and ask for their description. The sender of the "Describe-request" MUST be prepared to accept that answer and as necessary follow up with "Describe-request" with the instance names of interest.

**Enumerate:** Used in the same GRASP messages as "describe", but instead of providing information about one service instance, it is listing service instance names. The purpose of enumerate is the same as browsing a service in DNS-SD. It would be followed by some human or automated selection of one or more instances and then a "describe" M\_REQ\_SYN request for those instances sent to the source of the "enumerate" to learn about the locators and other parameters of the service instances.

In this specification, all fields other than service, instance and domain (and msg-type and private) must be unset in "enumerate".

**Enumerate-request:** Requests an "enumerate" reply. It is used in the same way as "Describe-request" except that instance would usually not be set (because in that case it is more useful to send a "Describe-request").

### 2.3.3. Name Element

The NAME,<name> elements is meant to provide basic name resolution comparable to mDNS name resolution for GRASP domains where this is

desirable and no better name resolution exist - for example in the ACP where there is no requirement for DNS.

Because the GRASP service lookup (unlike) DNS does not mandate that nodes have names (not even service instance names), the use of names is primarily meant to support legacy software. New designs should instead look up only services and service instance names, and nodes should announce their names as service instance names for the services they offer:

For example consider a GRASP (ACP) domain of "example.com". The node providing some "www" service could have a name "www-internal" which means GRASP objective NAME.www-internal, that objective value would include primarily the nodes IP address(es) and the port number for the www service would have to be guessed (80). Better, the node would announce GRASP objective SRV.www and the objective value would include the service instance name www-internal and the (TCP) port information (80 or a non-default port).

```
relement-codepoint ::= ( &(name-element:3) => context-element )

context-element ::= {
    *( &name:10)          => tstr),
}

ipv6-address-option = [O_IPv4_ADDRESS, ipv6-address]
ipv4-address-option = [O_IPv6_ADDRESS, ipv6-address]
locator-option /= ipv4-address-option
locator-option /= ipv6-address-option
```

Name information is carried in the name-element relement. It is a context-element like the one used for srv-element except that it adds the name component and that it does not permit the service and instance components and that it allows only describe and describe-request values in the msg-type. Clocators MUST use the ipv6-address-option or ipv4-address-option in the locator-option component.

TBD: Unclear if/how we should best formalize the differences in the context element permitted information between services and names. The above is quite informal.

Priority, weight, kvpairs, range (and of course private) MAY be used in describe messages to support multiple instances of the same name, as used for name anycast/prioritycast.

Nodes may have multiple names. These can be listed in the name component. If a nodes names have the notion of a primary name and secondary names then the primary name should be the first in the list of names. In DNS-SD, the name pointed to by CNAME RRs can be considered to be the primary name. A describe-request for a non-primary name SHOULD return in the list of names the requested name and the primary name.

Note that there is no reverse lookup defined in this version of the document (no lookup from IP address to name).

### 3. Explanations (Informative)

#### 3.1. Using GRASP service announcements

TBD: This section contains a range of details that should become normative in later versions.

This section provides a step by step walk-through of how to use GRASP service announcements and compares it to DNS-SD.

The most simple method to use GRASP service discovery is to select (and if still necessary, register) a <service-name> and start one or more agents (e.g.: ASAs) announcing their service instance(s) via GRASP. At minimum, an agent should periodically (default 60 seconds) announce the service instance via GRASP M\_FLOOD messages as an objective SRV.<service-name> with a srv-element and a sender-loop-count element (default 255). The ttl of the GRASP message should be 3.5 times the announcement period, e.g.: 210000 msec.

Consumers of the service will use GRASP to learn of the service instances and select one. This approach is most similar to the use of DNS-SD with mDNS except that the scope of the announcement is a whole GRASP domain (such as the ACP) as opposed to a single IP subnet in mDNS and that mDNS primarily relies on request & reply but in its standard not on periodic unsolicited announcements. We describe here the unsolicited flooding option via M\_FLOOD first because it is recommended for services with a dense population of service consumers and it is most simple to describe.

On the service announcer, the parameters priority, weight and range of the service instance can be selected from intent or configuration - or left at default. The default range 255 will result in selection of a random target of the service like in DNS-SD. Setting priority/

weight allows to prioritize and weigh the selection as in DNS-SD. Setting range to 0 allows to select the closest target, priority/weight are only compared between targets of the same shortest distance. Distance based options are not available in DNS-SD because it does not expect that network distance is available to arbitrary DNS-SD client. It is available to GRASP clients though. Using  $0 < \text{range} < 255$  allows for a hybrid priority/weight and distance based service selection (e.g.: Select the highest priority instance within a range of 5 hops).

If the service is a non-GRASP service, then the result of the service discovery has to be a transport locator to which the client can open a connection and talk the protocol implied by the service. This transport locator(s) have to be put into the clocator parameter. The context of the clocator would normally be "", aka: the transport locator is in the IP reachability associated with the GRASP domain (e.g.: IPv6 of the ACP for ACP GRASP domain).

If an ACP service is announced via ACP GRASP, then the locator(s) can be `O_IPv6_LOCATOR` or `O_FQDN_LOCATOR`. The `O_IPv6_LOCATOR` is used if the service is defined to be available via some transport layer port (TCP, UDP or other). The determination of the actual transport connection to be used is the same as in DNS-SD: If the transport protocol is not TCP or UDP, it has to be implied by the specification of `<service-name>` or can be detailed in kvpairs which carries the same information as DNS-TXT TXT RRs of the service. Alternatively, the `transport-proto` field of the locator can contain any valid IP protocol directly (TBD), which is not possible in DNS-SD.

Like DNS-SD, service discovery via GRASP does not require allocation and use of well-known ports for services. Unlike DNS-SD, there is no need in GRASP to define service instance names or target names. In DNS SD, PTR RRs resolve from a service name to a set of service instance named. SRV and TXT RRs resolve from service instance names to service instance parameters including the target. A target is the DNS host name of the service instance. It gets resolved via A/AAAA RRs to IPv4/IPv6 addresses of the targ. In GRASP service discovery, host names are not used. Service instance names are optional too. Service instance names are useful for human diagnostics and human selection of service instances. In fully automated environments, they can be are less important. For diagnostic purposes, it is recommended to give service instances service instance names in GRASP service announcements.

A locator with `O_URI_LOCATOR` type can be used in GRASP to indicate a URI for the transport method for a service instance. If the URI includes a host part, care must be taken to use only IP addresses in the host part if the context of the GRASP domain does not support

host name resolution - such as the ACP - or to use the GRASP name resolution mechanisms described elsewhere in this document. And that the addresses indicated are also reachable in the GRASP domain. For example, in service announcements across a DULL GRASP domain, only the IPv6 link-local addresses on that subnet must be used (this applies equally when using the O\_IPv6\_LOCATOR).

Instead of using M\_FLOOD to periodically announce service instances, M\_DISCOVERY can be used to actively query for service instances. The msg-type type must then be "describe-request". Because no periodic flooding is necessary, this solution is more lightweight for the network when the number of requesting clients is small. Note though that the M\_DISCOVERY will terminate as soon as a provider of the objective is found, so the service instances found will be based on distance and therefore selection of instance by priority and weight will not work equally well as with M\_FLOOD. Consider for example a central service instance in the NOC that should always be used (for example for centralized operational diagnostics) unless the WAN connection is broken, in which case distributed backup service instances should be used. With the current logic of M\_DISCOVERY this is not possible.

### 3.2. Further comparison with DNS-SD

Neither the GRASP SRV.\* objective-name, the service name nor any other parameter explicitly indicate the second label "\_tcp" or "\_udp" of DNS-SD entries. DNS-SD, RFC6763 explains how this is an unnecessary, historic artefact.

This version of the document does not define an equivalent to "\_sub" structuring of service enumeration.

This version of the document does not define mechanisms for reverse resolution of arbitrary services: An inquirer may unicast M\_SYNC\_REC to a node with a series of objectives with specific service names of interest and describe-request, but there is no indication of "ANY" service.

### 3.3. Open Issues

TBD: Examine limitations mentioned in "in this version of the text/document".

TBD: The GRASP specification does currently only permit TCP and UDP for the transport-proto element. This draft should expand the GRASP definitions to permit any valid IP protocol. We just need to decide whether this should only apply to the locator in the srv element or

also retroactive to the locator-option in GRASP messages (maybe not there ?).

TBD: A fitting CBOR representation for a kvpair key without value needs to be specified so that it can be distinguished from an empty value as outlined in RFC6763 section 6.4.

TBD: In this version, every service/service-instance is an element by itself. Future versions of this document may add more encoding options to allow more compact encoding of recurring fields.

TBD: Is there a way in CDDL to formally define the string names of the relement-codepoint's ?

#### 4. Security Considerations

TBD.

#### 5. IANA Considerations

This document requests a new "GRASP Objective Value Standard Elements" table in the GRASP Parameter Registrar. The values in this table are names and a unique numerical value assigned to each name. Future values MUST be assigned using the RFC Required policy defined by [RFC8126]. The numerical value is simply to be assigned sequentially. The following initial values are assigned by this document:

sender-loop-count 1 [defined in rfcXXXX]

srv-element 2 [defined in rfcXXXX]

name-element 3 [defined in rfcXXXX]

This document updates the handling of the "GRASP Objective Names" Table introduced in the GRASP IANA considerations as follows:

Assignments for objective-names of the form "SRV.<text>" and "NAME.<text>" are special.

Assignment of "SRV.<text>" can only be requested if <text> is also a registered service-name according to RFC6335. The specification required for registration of a "GRASP Objective Name" MUST declare that the intended use of the objective name in GRASP is intended to be compatible with the intended use of the registered service name.

Registration of "SRV.<text>" in the "GRASP Objective Name" table is optional, but recommended for all new service-names that are meant to

be used with GRASP. Non-registration can for example happen with DNS-SD <-> GRASP gateways that inject pre-existing service-names into GRASP. Note that according to the GRASP RFC, registration is mandatory, so this exemption for "SRV.<text>" is also an update to that specification.

There MUST NOT be any assignment for objective names of the form "NAME.<text>". These names are simply used by GRASP nodes without registration (just like names in mDNS).

## 6. Acknowledgements

## 7. Change log [RFC Editor: Please remove]

### 7.1. 01 -

Only refreshing, no changes since -00.

### 7.2. 00 - Initial version

## 8. References

### 8.1. Normative References

[I-D.ietf-anima-grasp]

Bormann, C., Carpenter, B., and B. Liu, "A Generic Autonomic Signaling Protocol (GRASP)", draft-ietf-anima-grasp-15 (work in progress), July 2017.

[RFC6763]

Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", RFC 6763, DOI 10.17487/RFC6763, February 2013, <<https://www.rfc-editor.org/info/rfc6763>>.

### 8.2. Informative References

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

Eckert, T., Behringer, M., and S. Bjarnason, "An Autonomic Control Plane (ACP)", draft-ietf-anima-autonomic-control-plane-16 (work in progress), June 2018.

Author's Address

Toerless Eckert  
Futurewei Technologies Inc.  
2330 Central Expy  
Santa Clara 95050  
USA

Email: tte+ietf@cs.fau.de

ANIMA  
Internet-Draft  
Intended status: Standards Track  
Expires: January 3, 2019

T. Eckert, Ed.  
Huawei  
Jul 2, 2018

Autoconfiguration of NOC services in ACP networks via GRASP  
draft-eckert-anima-noc-autoconfig-00

Abstract

This document defines standards for the autoconfiguration of crucial NOC services on ACP nodes via GRASP. It enables secure remote access to zero-touch bootstrapped ANI devices via SSH/Netconf with Radius/Diameter authentication and authorization and provides lifecycle autoconfiguration for other crucial services such as syslog, NTP (clock synchronization) and DNS for operational purposes.

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 <https://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 January 3, 2019.

Copyright Notice

Copyright (c) 2018 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 (<https://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.	Overview	2
1.2.	ACP nodes supporting NOC autoconfiguration	3
1.3.	Use of ACP GRASP for autoconfiguration	3
1.4.	GRASP parameters	4
2.	Services	7
2.1.	Syslog	7
2.2.	NTP	10
2.3.	DNS for operations	12
2.4.	Radius	13
2.5.	Diameter	14
2.6.	SSH server	14
3.	Security Considerations	14
4.	Change log [RFC Editor: Please remove]	15
5.	References	15
5.1.	Normative References	15
5.2.	Informative References	16
	Author's Address	16

## 1. Introduction

### 1.1. Overview

This document defines standards for the autoconfiguration of NOC services on ACP nodes via GRASP.

One key purpose of this autoconfiguration is the seamless step from zero-touch bootstrap in ANI devices over to a securely remotely configurable ACP node. For this purpose, this document defines the autoconfiguration of the SSH/Netconf server on the ACP node with autoconfigured authentication across the ACP with Radius or Diameter.

Both for this initial bringup as well as for ongoing operations of ACP nodes, this document describes autoconfiguration of syslog services as well as autoconfiguration of DNS for operational purposes. Syslog is also specified to make tracking of device configuration state easier and allow to fully automate the bringup of the ACP node after zero-touch bootstrap.

## 1.2. ACP nodes supporting NOC autoconfiguration

This document introduces the term ACPna nodes to indicate nodes supporting ACP that also support the requirements described in this document: ACP (n)oc (a)utoconfigurable.

If an ACPna node supports zero-touch bootstrap of the ACP where no configuration is possible before the ACP is enabled, then the NOC autoconfiguration features described in this document SHOULD be enabled by default on such an ACPna node after this zero-touch bootstrap, because the autoconfiguration of these NOC services can be the only method for the ACPna node to become operationally accessible from the NOC so that it can further be configured. ANI nodes are nodes supporting ACP and BRSKI ([I-D.ietf-anima-bootstrapping-keyinfra]). BRSKI bootstrap is an instance of such a zero-touch bootstrap requiring auto-enablement of NOC autoconfiguration after zero-toch bootstrap.

If an ACPna node was not zero-touch bootstrapped, then NOC autoconfiguration SHOULD be enabled whenever ACP is enabled but may be separately configurable.

## 1.3. Use of ACP GRASP for autoconfiguration

Autoconfiguration of ACNna services utilizes the ACP instance of GRASP, ([I-D.ietf-anima-grasp] as defined in [I-D.ietf-anima-autonomic-control-plane]). It leverages and extends the GRASP objective definitions of [I-D.eckert-anima-grasp-dnssd]. Those objective elements allow to create DNS-SD compatible service announcements with flexible priority/weight and distance based selection across multiple service instances and per-service parameters.

Nodes in a NOC supporting a particular service announce it via the appropriate GRASP objective into ACP GRASP. The NOC nodes therefore need to have access to the ACP, either directly because they are ACP nodes or because they use the ACP connect function (see [I-D.ietf-anima-autonomic-control-plane]). ACPna nodes receive these announcements and auto-configure the services tied to them. In most instances, the service announcement from the NOC is for a server instance that a client instance on the ACPna node connects to, for example a syslog server in the NOC. In another instance, the NOC service is that of an authentication service and the ACPna nodes will enable a server instance that leverages the authentication service in the NOC.

Note: Currently, this document does not define the option of an mDNS/DNS-SD -> ACP GRASP gateway function to enable NOC nodes without

GRASP implementations to utilize mDNS/DNS-SD to announce their services and then expect an appropriate translation function to convert these announcements into GRASP objectives. This document does define all the GRASP objectives so that that it would be possible to define such a gateway function, but some loss of functionality would exist. For once, GRASP does support network distance based service selection (e.g., select a server from the closest NOC), whereas no such mechanism exists in DNS-SD. In addition, this document believes that support of GRASP software to announce services from NOC systems is very easy to accomplish.

#### 1.4. GRASP parameters

Unless otherwise described, all GRASP objective announcements described in this document SHOULD default to the following GRASP parameters. These parameters MAY all be configurable on the NOC nodes.

- o M\_FLOOD GRASP message, periodically sent once every 60 second. Random phase vs. full minutes (so different service announcements are distributed over time in the network).
- o ttl of 210000 msec (3.5 times 60 seconds).
- o locator-option is the ACP address of the announcing node unless the announcement is done from a third-party, for example if the announcing server does not support GRASP but GRASP is run on another NOC node.
- o objective-name is 'SRV.<name>', where <name> is an [RFC6335] registered service name for the service in question.
- o objective-flags is sync-only, loop-count is 255.
- o objective-value MUST comply with the requirements of [I-D.eckert-anima-grasp-dnssd].

```

[M_FLOOD, 12340815, h'fd89b714f3db0000200000064000001', 210000,
  ["SRV.syslog", 4, 255,
    { rfcXXXX: {
      &(sender-loop-count:1) => 255,
      &(srv-element:2) => {
        &(msg-type:1) => &(describe: 0),
        &(service:2) => "syslog",
        &(instance:3) => "east-coast-primary",
        &(priority:5) => 0,
        &(weight:6) => 65535,
        &(kvpairs:7) => { "replicate" => 2 },
        &(range:8) => 2,
      }
    }
  ],
  [O_IPv6_LOCATOR,
    h'fd89b714f3db0000200000064000001', TLS12, 514]
]

```

Figure 1: SRV.syslog example

The above example shows the default values for a "syslog" service announcement using the objective-value elements defined in [I-D.eckert-anima-grasp-dnssd]. SRV.syslog is the standard objective name for the "syslog" service, as is SRV.<any> for the <any> service. The announcer of this objective also provides the syslog service as it is announcing its own address in the locator option. It provides syslog on the standard syslog TCP port 514 using TLS12.

The DNS-SD equivalent service attributes are carried in the srv-element. The msg-type indicates that this objective is a service announcement. The instance of "" indicates that this service announcement for the ACP itself, and not for e.g. the data-plane. It is shown here just for illustration purposes and can be left out in encoding because it is the default. Likewise, the service element is redundant and shown only for illustrative purposes. Priority and weight have the same semantic as in DNS-SD SRV records. In this case, the service announcement indicates the highest priority (0) and the highest weight (65535). Kvpairs includes service specific options.

Going beyond the capabilities, the range parameter indicates that the client of this service should select this announced service not only by priority/weight but primarily by the distance in terms of network hop-count between this service announcer and the client: The client is expected to select the best service announcement by priority and weight only between alternatives that are not more than two network

hops apart in distance to the client. Otherwise the client should pick the closer one.

To allow the client to know the distance to a service announcement, the sender-loop-count parameter is included in the announcement. It MUST be set by the sender to the same value (255 in this example) as the loop-count in the GRASP header. The loop-count in the header is hop-by-hop reduced. When the GRASP message arrives at the client, the difference between sender-loop-count and loop-count is the distance to the service announcer in hops.

```

;
; Following GRASP header definitions from GRASP
;
flood-message = [M_FLOOD, session-id, initiator, ttl,
                 +[objective, (locator-option / [])]]
objective = ["SRV.<rfc6335-name>", objective-flags, loop-count,
            objective-value]

objective-flags = sync-only ; as in GRASP spec
sync-only       = 4         ; M_FLOOD only requires synchronization
loop-count      = 255      ; recommended
;
; Following GRASP objective-value definitions from GRASP DNS-SD
;
objective-value = { 1*elements }
elements        = ( @rfcXXXX: { 1*relement } )
relement       //= ( &(sender-loop-count:1) => 1..255 )
relement       //= ( &(srv-element:2) => context-element )
context-element = {
    ?( &(private:0)      => any),
    ?( &(msg-type:1)     => msg-type),
    ?( &(service:2)      => tstr),
    *( &(instance:3)    => tstr),
    ?( &(domain:4)       => tstr),
    ?( &(priority:5)    => 0..65535 ),
    ?( &(weight:6)       => 0..65535 ),
    *( &(kvpairs:7)     => { *(tstr: any) },
    ?( &(range:8)        => 0..255 ),
    *( &(clocator:9)    => clocator),
}
;
TLS12 = 257

```

Figure 2: GRASP service definition

The above picture shows the complete CDDL definition of a GRASP M\_FLOOD message indicating a service together with the objective-value encoding. Some of the context-element options are not used in this document (TBD - remove before going RFC).

The value 257 is defined to indicate TLS12 ([RFC5246]) to be used in the protocol field of GRASP locators to indicate that a TCP port is intended to be used with TLS version 1.2. Values 1..255 are reserved for IP protocol numbers.

## 2. Services

### 2.1. Syslog

ACPna nodes SHOULD support autoconfiguring of syslog via the SRV.syslog objective.

When an ACPna node discovers one or more SRV.syslog announcements across the ACP, it SHOULD perform syslog operations to the best available discovered server.

Local configuration of syslog on the ACPna node SHOULD have no impact on the autoconfigured syslog operations, or else, misconfiguration could cause to failure of the autoconfigured syslog operations. Instead, configured syslog operations should just operate as ships-in-the-night to the GRASP learned autoconfigured syslog operations.

Severity of syslog messages SHOULD be 5 (Notice) (see [RFC5424]), and all messages that are necessary to support normal remote operations of the node should be assigned severities higher (numerically lower) or equal to 5/Notice.

Syslog service announcements SHOULD include the instance option, indicating the unique name of the service instance described by the GRASP objective. This serves diagnostics and avoids having to identify service instances by the address(es) in the locator-options. In the example Figure 1, the instance name is "east-coast-primary".

The syslog facility value is a choice of the ACPna node, the autoconfigured syslog server must be able to deal with any syslog facility code received. If an ACPna node has no pre-established standard for the facility-code, then the value of local7 (23) MAY be used.

For resilience, it may be appropriate to receive syslog messages on more than one server. A server can indicate this via the "replicate" keyword in the GRASP objective-value kvpair element. The value of the "replicate" keyword indicates the maximum number of syslog

servers that the client SHOULD autoconfigure syslog to. After selecting the best service announcement, the client looks up the value N of the "replicate" keyword of that best servers announcement and selects the best N-1 service announcements and ultimately logs to all N. ACPna nodes SHOULD support autoconfigured syslog to up to 3 servers simultaneously.

Autoconfigured syslog SHOULD support TLS1.2, TCP and UDP. Because ACP provides encryption, use of just TCP instead of TLS should be sufficient and may achieve higher performance. Use of UDP should be avoided because of the potential to loose packets and not supporting congestion control.

If a syslog server supports more than one transport option (TLS1.2, TCP, UDP), it SHOULD announce them via a single GRASP objective and list them via clocator options of the srv-element because the locator-option in the GRASP header (as shown in example Figure 1) allows only one locator-option. The order of the clocator options in the indicates the preference of the server. From this list, the client supports the first option supported also by the client and ignore the others. The context of the clocator would normally be "", indicating that the locator-option address is reachable via the ACP.

Instead of (or in addition to) using multiple clocator options, a server can also announce multiple SRV.syslog objectives, but in that case each of them would be considered to be a different service instance considered by the the client when selecting the (set of) best service instances. If a service announcement indicates via the "replicate" keyword that the client should log to three service instances, and announce three separate SRV.syslog objectives, each one with a different locator-option, then the client might select to log to all three of them - instead of - which is more likely the desired option - for the client to log to actually three different servers. Hence the use of multiple clocator options that are examined by clients only after server selection is done.

When a client uses TLS, it MUST use its ACP domain certificate for authentication. Likewise, the syslog server MUSTS use its ACP domain certificate.

Logging by default uses the ACP, in the clocator option, this is indicated via a context value of "". Servers may also indicate support for logging across the data-plane, which may provide higher performance but may fail if reachability in the data-plane does not exist, so care must be taken when announcing this option. For example, in managed MPLS/VPN networks where the ACP extends across P/PE and CE devices, the global routing table on a CE device is often not the same as that on P/PE devices, and therefore CE devices may

not be able to log to "0". In this case, the syslog server should instead announce a deployment choosen name for the context, such as "VRF0". Clients would only take such a clocator into account if there is a local configuration that maps the context name to a routing table. In this example, only P/PE nodes would have this configuration, therefore allowing the CE nodes to ignore this clocator; And if this clocator was the only locator-option in the GRASP objective, then the whole objective MUST be ignored by the client when selecting the best possible service instance. Note that for contexts other than the ACP (""), both IPv4 and IPv6 are possible, depending on what version(s) of IP are deployed in the data-plane.

Failure to connect to a choosen service instance SHOULD be taken into accounts by clients when selecting service instances to log to. For UDP locator-options, ICMP/ICMPv6 error indications are such connection failures. For TCP/TLS connections, connection failure includes TCP and TLS failures as well as keepalive failures. When failures occur, clients should attempt to re-connect with exponential timeouts, starting with 5 seconds and staying at 320 seconds or until the GRASP service announcement expires and is not refreshed.

When connecting to a server fails, the ACPna client SHOULD connect to the next best available server in the meantime. ACPna client SHOULD support connecting to up to four service instances if any connections fail. If for example the client is logging to two service instances because 2 is indicated in the "replicate" option of the service announcements, and one fails, the client will attempt to re-connect to it while in parallel establishing syslog connection to a third-best service-instance.

When establishing connection to a new syslog service instance, ACPna clients SHOULD log with severity 5 an indication of this event, indicating its own ACP address, the ACP address and if existing instance name of the new syslog service instance and the reason. Like any other autoconfigured syslog message, this would go to all syslog connections and therefore show up on the redundant syslog servers, allowing to recognize failure of connectivity to another syslog server - and tracing of client logs across syslog servers if the client changes them.

Examples:

```
ACP: fd89:b714:f3db::0200:0000:6400:0042 start logging to:  
fd89:b714:f3db::0200:0000:6400:0001/east-coast-primary,TLS reason:  
starting up
```

```
ACP: fd89:b714:f3db::0200:0000:6400:0042 start logging to:
fd89:b714:f3db::0200:0000:6400:0001/east-coast-primary,TLS reason:
new better service instance
```

```
ACP: fd89:b714:f3db::0200:0000:6400:0042 stop logging to:
fd89:b714:f3db::0200:0000:6400:0001/east-coast-secondary,TLS reason:
connection failure
```

When failures to deliver syslog messages to ANY syslog servers happen, clients SHOULD track the this and indicate loss of messages via the next working syslog connection. Note that due to the possibility of ICMP/ICMPv6 errors, only the successful delivery of messages via TLS or TCP should be tracked. TBD: need to check if this can reasonably be recommended, pending on availability of e.g. TAPS API spec to know whetherer a TCP write was sent and acknowledged by the receiver (given how there are no reply messages in syslog).

## 2.2. NTP

Time synchronization is one of the most fundamental functionality for network devices for a variety of functions to work and also for diagnostics to be comparable across the network. If problems propagate fast across the network, the client generated timestamp of events in syslog messages (or other diagnostics function) allows to trace event propagation and decude causality. This may require network clock synchronization in the order of milliseconds, something which is easily achievable in todays network devices via NTP.

ACPna nodes SHOULD support autoconfiguration of clock synchronization through NTP ([RFC5905]) with the following autoconfiguration semantics.

The GRASP objective for NTP is SRV.ntp. This does not distinguish between NTPv4 and NTPv3 because NTPv4 is fully backward compatible with NTPv3, so server and client will negotiate between these two versions.

The kvpair key "stratum" has a numeric value and indicates the stratum or level of a server in a synchronization tree. The value of 1 indicates the root of the distribution tree. Servers that synchronize from the master have a stratum of 2, and so on.

The kvpair key "minpoll" indicates the lowest periodic polling that the client will perform against the server. Announcing a large numeric value allows for a server to reduce the amount of NTP messages from clients, but slows down convergence time of clientsnumber of service instances that simultaneously bootstrap.

The kvpair key "key" indicates the NTPv3 authentication mechanism. When present, clients MUST use the value as the key to perform NTPv3 (MD5) hash authentication of message with this service instance. Note that the encryption of the ACP serves as protection of distributing such a cleartext symmetric key via GRASP to clients.

TBD: Understand NTPv4 autokey and define appropriate kvpair to enable auto-configuring it, especially when the service instance announcement indicates the use of the data-plane.

The autoconfiguration described in the following paragraphs is for leafs of the clock distribution graph, e.g., nodes that do only aim to obtain synchronized time from a server. Configuration of the server hierarchy is left to explicit configuration.

Clients SHOULD select service instance(s) with the worst (highest) stratum value. In the face of multiple equal options, clients have to pick the best ones based on the standard selection criteria priority/weight and range, allowing for distributed NTP server deployment by e.e., setting range to 1, or via centralized deployment with multiple servers, setting range to 255 and priority/weight accordingly. Making the stratum the primary selection criteria allows in the future to also introduce autoconfiguration of servers in the NTP clock distribution tree without incurring the problem that a large number of clients would then select higher stratum servers (and overload them).

Like most other autoconfigured services, the autoconfigured NTP time synchronization SHOULD take precedence over explicit configured NTP options to ensure that time synchronization is not subject to misconfiguration of individual nodes (but only subject to misconfiguration of servers).

The kvpair "TZ" option allows to signal the time zone of the ACP network to clients. Its value is a string indicating the time zone of all nodes in the ACP network. Care must be taken not to use this option in networks extending across multiple time zones. Because time zone distribution does not work automatically across larger networks with multiple time zones, overriding the signalled time zone SHOULD be possible through local configuration.

TBD: references for time zone spec standards and also for DST rule indications.

### 2.3. DNS for operations

Availability of DNS names for network operations/troubleshooting is today mostly an convenience in network operations, but with IPv6 evolving the need to use DNS names even in CLI based network diagnostics is raising - because IPv6 addresses often are more difficult to memorize by operators. More and more network features also support configuration that instead of addresses include domain names or URLs, and ultimately, any non-fully autoconfigured functions should rather rely on domain-names and URLs instead of just addresses for greater flexibility and reliability in the face of address changes.

In the face of this, ACPna nodes SHOULD support autoconfiguration of DNS for operational purposes. "For operation purposes" implies that the use of the autoconfigured DNS servers SHOULD NOT be used for DNS services offered to users of the data plane, such as DNS proxy services. This would cause the ACP to effectively carry user traffic, whenever a client DNS request to an ACPna node with a DNS proxy would be forwarded to an autoconfigured server via the ACP.

The GRASP objective name for such OAM use of DNS is OAM-DNS. It is explicitly not SRV.dns to highlight that this instance of DNS is copied for operational purposes only to isolate it from user issues (performance across the ACP and attacks). Utilizing different DNS contexts also allows to set up split-horizon DNS where all the operationally relevant DNS names are only made available via the DNS servers or zones available across the ACP.

The value of the "search-list" kvpair option is a ";" (semicolon) separated list of domain name prefixes that should be searched by the client for non-FQDN that they need to resolve. "local-arpa" is the prefix to use for reverse IPv4/IPv6 address lookups. If for example "local-arpa" is set to "arpa.example.com", then the clients should first look up IPv4/IPv6 addresses in "ipv6.arpa.example.com"/"in-addr.arpa.example.com." before resorting to lookup in the Internet global "ipv6.arpa."/ "in-addr.arpa.". For RFC1918/ULA addresses, no fallback to the global reverse lookup prefixes should be done.

ACPna nodes SHOULD look up their name via a reverse lookup of their ACP address, and then auto-configure this name.

There are no service specifics for the selection of DNS servers. A ACPna node simply uses the standard priority/weight/range options to select a DNS server. It MAY prefer a server with TCP locator-option simply because that allows in most cases faster discovery of connectivity problems than a UDP connection.

TBD: Note that it is fairly easy to re-use the autoconfiguration scheme described here to provide auto-configuration of DNS for user DNS services with the help of the ACP. The objective name would have to be changed and the clocators would have to indicate a data-plane context, so that user requests are carried across the data-plane from DNS proxies to DNS servers. It is unclear if this service should be described in this document though.

#### 2.4. Radius

Radius [RFC2865] is a protocol used for AAA service - Authentication, Authorization and Accounting. Autodiscovery of Radius servers across the ACP for ACPna nodes serves the purpose to enable authentication and authorization of other ACPna autoconfigured services such as below described Section 2.6.

ACPna nodes MUST support Radius and/or Diameter autoconfiguration if they support any of the autoconfigured services depending on such an authentication service.

The GRASP objective name is SRV.radius. The UDP or TCP port of the locator-option in the GRASP header or the clocator option indicate the UDP or TCP port of the Radius servers authentication connection. The context of a clocator MUST be "" to indicate the ACP - because the Radius connections MUST pass across the ACP to be protected against eavesdropping - and the radius security methods described here are not sufficiently secure to allow passing them across the data-plane.

The kvpair "secret\_key" string value indicates the secret key to use on the connection to the Radius server. The optional "acct\_port" numeric value indicate the UDP/TCP port of any accounting connection supported by the radius server. The protocol (UDP vs. TCP) is the same as the one in the chosen locator-option/clocator.

There are no service specific selection rules. TCP is preferred for faster recognition of a failed server and reselection of an alternative server.

The specific data/authentication/authorization configuration required on the Radius server depends on the OAM service authenticated/authorized and is described in its section in this document.

TBD: Should we define AVpair or different objective names to distinguish what services can be authenticated? Would be easier if we found another service than SSH/Netconf.

## 2.5. Diameter

TBD. Alternative to Radius. Author would welcome suggesting what parameters are relevant for a diameter authentication service.

## 2.6. SSH server

ACPna nodes supporting SSH server functionality for remote management access via CLI, NETCONF or other methods SHOULD auto-enable SSH server functionality across the ACP whenever they are aware from ACP GRASP of RADIUS (Section 2.4) or DIAMETER (Section 2.5) authentication servers. ACPna nodes that support ACPna SSH server functionality MUST support authentication via either RADIUS and/or Diameter.

If both protocols are supported by the ACPna node, the ACPna node SHOULD select the authentication server based on the service priority parameters across both protocols. E.g., if a RADIUS server has a higher priority in GRASP than the DIAMETER server, the ACPna node should authenticate against the RADIUS server.

When valid authentication server(s) are discovered, the SSH server is autoconfigured. It SHOULD only listen to the standard SSH port with the ACP address of the node but not be reachable from the data-plane. It MUST NOT be modifyable by configuration (only by auto-configuration). If autoconfiguration of an SSH server on the standard SSH port conflicts with explicitly configured SSH server for the data-plane due to software limitations or complexity, the autoconfigured SSH server MAY be started on a node-type specific and not dynamically selected port number. This port number must be well-known to OAM operations as there is no method provided to signal it to the SSH client side.

Note that this document does not define any standards for the exact message options for authentication or authorization. Especially authorization, such as privilege level that permits to change configuration is likely using vendor specific methods, and Radius/Diameter servers must be capable to recognize the type of client as they had to without this autoconfiguration.

## 3. Security Considerations

There is no protection against "unauthorized" ACP nodes to generate service announcements, because there is no authorization scheme in GRASP. Discovery of unauthorized announcers is easy though because the service announcements are flooded across the ACP and are therefore easily visible on nodes that may specifically observe announcements to discover unauthorized ones.

A possible framework to define authorization could rely on defining roles for ACP nodes either through additional parameters in their ACP domain certificate or following initial provisioning, and then lock down the ability for later configuration to enable services (and their GRASP announcements) to only those included in the role assigned to the node. This is outside the scope of this document.

#### 4. Change log [RFC Editor: Please remove]

00: Initial version.

#### 5. References

##### 5.1. Normative References

[I-D.eckert-anima-grasp-dnssd]

Eckert, T., "DNS-SD compatible service discovery in GRASP", draft-eckert-anima-grasp-dnssd-01 (work in progress), July 2018.

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

Eckert, T., Behringer, M., and S. Bjarnason, "An Autonomic Control Plane (ACP)", draft-ietf-anima-autonomic-control-plane-16 (work in progress), June 2018.

[I-D.ietf-anima-grasp]

Bormann, C., Carpenter, B., and B. Liu, "A Generic Autonomic Signaling Protocol (GRASP)", draft-ietf-anima-grasp-15 (work in progress), July 2017.

[RFC2865] Rigney, C., Willens, S., Rubens, A., and W. Simpson, "Remote Authentication Dial In User Service (RADIUS)", RFC 2865, DOI 10.17487/RFC2865, June 2000, <<https://www.rfc-editor.org/info/rfc2865>>.

[RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, DOI 10.17487/RFC5246, August 2008, <<https://www.rfc-editor.org/info/rfc5246>>.

[RFC5424] Gerhards, R., "The Syslog Protocol", RFC 5424, DOI 10.17487/RFC5424, March 2009, <<https://www.rfc-editor.org/info/rfc5424>>.

[RFC5905] Mills, D., Martin, J., Ed., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, DOI 10.17487/RFC5905, June 2010, <<https://www.rfc-editor.org/info/rfc5905>>.

[RFC6335] Cotton, M., Eggert, L., Touch, J., Westerlund, M., and S. Cheshire, "Internet Assigned Numbers Authority (IANA) Procedures for the Management of the Service Name and Transport Protocol Port Number Registry", BCP 165, RFC 6335, DOI 10.17487/RFC6335, August 2011, <<https://www.rfc-editor.org/info/rfc6335>>.

## 5.2. Informative References

[I-D.ietf-anima-bootstrapping-keyinfra] Pritikin, M., Richardson, M., Behringer, M., Bjarnason, S., and K. Watsen, "Bootstrapping Remote Secure Key Infrastructures (BRSKI)", draft-ietf-anima-bootstrapping-keyinfra-16 (work in progress), June 2018.

## Author's Address

Toerless Eckert (editor)  
Futurewei Technologies Inc.  
2330 Central Expy  
Santa Clara 95050  
USA

Email: [tte+ietf@cs.fau.de](mailto:tte+ietf@cs.fau.de)

ANIMA WG  
Internet-Draft  
Intended status: Standards Track  
Expires: 3 May 2021

T. Eckert, Ed.  
Futurewei USA  
M. Behringer, Ed.

S. Bjarnason  
Arbor Networks  
30 October 2020

An Autonomic Control Plane (ACP)  
draft-ietf-anima-autonomic-control-plane-30

Abstract

Autonomic functions need a control plane to communicate, which depends on some addressing and routing. This Autonomic Control Plane should ideally be self-managing, and as independent as possible of configuration. This document defines such a plane and calls it the "Autonomic Control Plane", with the primary use as a control plane for autonomic functions. It also serves as a "virtual out-of-band channel" for Operations, Administration and Management (OAM) communications over a network that provides automatically configured hop-by-hop authenticated and encrypted communications via automatically configured IPv6 even when the network is not configured, or misconfigured.

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 <https://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 3 May 2021.

Copyright Notice

Copyright (c) 2020 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 (<https://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 (Informative) . . . . .	6
1.1. Applicability and Scope . . . . .	9
2. Acronyms and Terminology (Informative) . . . . .	11
3. Use Cases for an Autonomic Control Plane (Informative) . . . . .	16
3.1. An Infrastructure for Autonomic Functions . . . . .	17
3.2. Secure Bootstrap over a not configured Network . . . . .	17
3.3. Data-Plane Independent Permanent Reachability . . . . .	17
4. Requirements (Informative) . . . . .	19
5. Overview (Informative) . . . . .	20
6. Self-Creation of an Autonomic Control Plane (ACP) (Normative) . . . . .	21
6.1. Requirements for use of Transport Layer Security (TLS) . . . . .	22
6.2. ACP Domain, Certificate and Network . . . . .	23
6.2.1. ACP Certificates . . . . .	24
6.2.2. ACP Certificate AcpNodeName . . . . .	26
6.2.2.1. AcpNodeName ASN.1 Module . . . . .	29
6.2.3. ACP domain membership check . . . . .	30
6.2.3.1. Realtime clock and Time Validation . . . . .	33
6.2.4. Trust Anchors (TA) . . . . .	33
6.2.5. Certificate and Trust Anchor Maintenance . . . . .	34
6.2.5.1. GRASP objective for EST server . . . . .	35
6.2.5.2. Renewal . . . . .	37
6.2.5.3. Certificate Revocation Lists (CRLs) . . . . .	37
6.2.5.4. Lifetimes . . . . .	38
6.2.5.5. Re-enrollment . . . . .	38
6.2.5.6. Failing Certificates . . . . .	40
6.3. ACP Adjacency Table . . . . .	41
6.4. Neighbor Discovery with DULL GRASP . . . . .	41
6.5. Candidate ACP Neighbor Selection . . . . .	45
6.6. Channel Selection . . . . .	45
6.7. Candidate ACP Neighbor verification . . . . .	49
6.8. Security Association (Secure Channel) protocols . . . . .	49
6.8.1. General considerations . . . . .	50
6.8.2. Common requirements . . . . .	51
6.8.3. ACP via IPsec . . . . .	52
6.8.3.1. Native IPsec . . . . .	52
6.8.3.1.1. RFC8221 (IPsec/ESP) . . . . .	53

6.8.3.1.2. RFC8247 (IKEv2) . . . . .	54
6.8.3.2. IPsec with GRE encapsulation . . . . .	55
6.8.4. ACP via DTLS . . . . .	56
6.8.5. ACP Secure Channel Profiles . . . . .	58
6.9. GRASP in the ACP . . . . .	59
6.9.1. GRASP as a core service of the ACP . . . . .	59
6.9.2. ACP as the Security and Transport substrate for GRASP . . . . .	59
6.9.2.1. Discussion . . . . .	62
6.10. Context Separation . . . . .	63
6.11. Addressing inside the ACP . . . . .	63
6.11.1. Fundamental Concepts of Autonomic Addressing . . . . .	63
6.11.2. The ACP Addressing Base Scheme . . . . .	65
6.11.3. ACP Zone Addressing Sub-Scheme (ACP-Zone) . . . . .	67
6.11.4. ACP Manual Addressing Sub-Scheme (ACP-Manual) . . . . .	68
6.11.5. ACP Vlong Addressing Sub-Scheme (ACP-VLong-8/ ACP-VLong-16 . . . . .	69
6.11.6. Other ACP Addressing Sub-Schemes . . . . .	70
6.11.7. ACP Registrars . . . . .	71
6.11.7.1. Use of BRSKI or other Mechanism/Protocols . . . . .	71
6.11.7.2. Unique Address/Prefix allocation . . . . .	72
6.11.7.3. Addressing Sub-Scheme Policies . . . . .	72
6.11.7.4. Address/Prefix Persistence . . . . .	74
6.11.7.5. Further Details . . . . .	74
6.12. Routing in the ACP . . . . .	74
6.12.1. ACP RPL Profile . . . . .	75
6.12.1.1. Overview . . . . .	75
6.12.1.1.1. Single Instance . . . . .	75
6.12.1.1.2. Reconvergence . . . . .	76
6.12.1.2. RPL Instances . . . . .	77
6.12.1.3. Storing vs. Non-Storing Mode . . . . .	77
6.12.1.4. DAO Policy . . . . .	77
6.12.1.5. Path Metric . . . . .	77
6.12.1.6. Objective Function . . . . .	77
6.12.1.7. DODAG Repair . . . . .	77
6.12.1.8. Multicast . . . . .	78
6.12.1.9. Security . . . . .	78
6.12.1.10. P2P communications . . . . .	78
6.12.1.11. IPv6 address configuration . . . . .	78
6.12.1.12. Administrative parameters . . . . .	79
6.12.1.13. RPL Packet Information . . . . .	79
6.12.1.14. Unknown Destinations . . . . .	79
6.13. General ACP Considerations . . . . .	80
6.13.1. Performance . . . . .	80
6.13.2. Addressing of Secure Channels . . . . .	80
6.13.3. MTU . . . . .	81
6.13.4. Multiple links between nodes . . . . .	81
6.13.5. ACP interfaces . . . . .	82

6.13.5.1.	ACP loopback interfaces . . . . .	82
6.13.5.2.	ACP virtual interfaces . . . . .	84
6.13.5.2.1.	ACP point-to-point virtual interfaces . . . . .	84
6.13.5.2.2.	ACP multi-access virtual interfaces . . . . .	84
7.	ACP support on L2 switches/ports (Normative) . . . . .	87
7.1.	Why (Benefits of ACP on L2 switches) . . . . .	87
7.2.	How (per L2 port DULL GRASP) . . . . .	88
8.	Support for Non-ACP Components (Normative) . . . . .	89
8.1.	ACP Connect . . . . .	89
8.1.1.	Non-ACP Controller / NMS system . . . . .	90
8.1.2.	Software Components . . . . .	92
8.1.3.	Auto Configuration . . . . .	93
8.1.4.	Combined ACP/Data-Plane Interface (VRF Select) . . . . .	94
8.1.5.	Use of GRASP . . . . .	96
8.2.	Connecting ACP islands over Non-ACP L3 networks (Remote ACP neighbors) . . . . .	97
8.2.1.	Configured Remote ACP neighbor . . . . .	97
8.2.2.	Tunneled Remote ACP Neighbor . . . . .	98
8.2.3.	Summary . . . . .	98
9.	ACP Operations (Informative) . . . . .	99
9.1.	ACP (and BRSKI) Diagnostics . . . . .	99
9.1.1.	Secure Channel Peer diagnostics . . . . .	103
9.2.	ACP Registrars . . . . .	104
9.2.1.	Registrar interactions . . . . .	104
9.2.2.	Registrar Parameter . . . . .	105
9.2.3.	Certificate renewal and limitations . . . . .	106
9.2.4.	ACP Registrars with sub-CA . . . . .	107
9.2.5.	Centralized Policy Control . . . . .	107
9.3.	Enabling and disabling ACP/ANI . . . . .	108
9.3.1.	Filtering for non-ACP/ANI packets . . . . .	108
9.3.2.	Admin Down State . . . . .	109
9.3.2.1.	Security . . . . .	110
9.3.2.2.	Fast state propagation and Diagnostics . . . . .	110
9.3.2.3.	Low Level Link Diagnostics . . . . .	111
9.3.2.4.	Power Consumption Issues . . . . .	112
9.3.3.	Interface level ACP/ANI enable . . . . .	112
9.3.4.	Which interfaces to auto-enable? . . . . .	112
9.3.5.	Node Level ACP/ANI enable . . . . .	114
9.3.5.1.	Brownfield nodes . . . . .	114
9.3.5.2.	Greenfield nodes . . . . .	115
9.3.6.	Undoing ANI/ACP enable . . . . .	116
9.3.7.	Summary . . . . .	117
9.4.	Partial or Incremental adoption . . . . .	117
9.5.	Configuration and the ACP (summary) . . . . .	118
10.	Summary: Benefits (Informative) . . . . .	119
10.1.	Self-Healing Properties . . . . .	119
10.2.	Self-Protection Properties . . . . .	121
10.2.1.	From the outside . . . . .	121

10.2.2.	From the inside	122
10.3.	The Administrator View	123
11.	Security Considerations	124
12.	IANA Considerations	129
13.	Acknowledgements	130
14.	Contributors	130
15.	Change log [RFC-Editor: Please remove]	131
15.1.	Summary of changes since entering IESG review	131
15.1.1.	Reviews (while in IESG review status) / status	131
15.1.2.	BRSKI / ACP registrar related enhancements	132
15.1.3.	Normative enhancements since start of IESG review	132
15.1.4.	Explanatory enhancements since start of IESG review	133
15.2.	draft-ietf-anima-autonomic-control-plane-30	134
15.3.	draft-ietf-anima-autonomic-control-plane-29	136
15.4.	draft-ietf-anima-autonomic-control-plane-28	138
15.5.	draft-ietf-anima-autonomic-control-plane-27	140
15.6.	draft-ietf-anima-autonomic-control-plane-26	140
15.7.	draft-ietf-anima-autonomic-control-plane-25	141
15.8.	draft-ietf-anima-autonomic-control-plane-24	144
15.9.	draft-ietf-anima-autonomic-control-plane-23	145
15.10.	draft-ietf-anima-autonomic-control-plane-22	146
16.	Normative References	148
17.	Informative References	151
Appendix A. Background and Futures (Informative)		
A.1.	ACP Address Space Schemes	160
A.2.	BRSKI Bootstrap (ANI)	161
A.3.	ACP Neighbor discovery protocol selection	162
A.3.1.	LLDP	162
A.3.2.	mDNS and L2 support	163
A.3.3.	Why DULL GRASP	163
A.4.	Choice of routing protocol (RPL)	163
A.5.	ACP Information Distribution and multicast	165
A.6.	CAs, domains and routing subdomains	166
A.7.	Intent for the ACP	167
A.8.	Adopting ACP concepts for other environments	168
A.9.	Further (future) options	170
A.9.1.	Auto-aggregation of routes	170
A.9.2.	More options for avoiding IPv6 Data-Plane dependencies	170
A.9.3.	ACP APIs and operational models (YANG)	171
A.9.4.	RPL enhancements	171
A.9.5.	Role assignments	172
A.9.6.	Autonomic L3 transit	172
A.9.7.	Diagnostics	172
A.9.8.	Avoiding and dealing with compromised ACP nodes	173
A.9.9.	Detecting ACP secure channel downgrade attacks	174

Appendix B. Unfinished considerations (To Be Removed From RFC)	175
B.1. Considerations for improving secure channel negotiation	175
B.2. ACP address verification	176
B.3. Public CA considerations	178
B.4. Hardening DULL GRASP considerations	179
Authors' Addresses	179

## 1. Introduction (Informative)

Autonomic Networking is a concept of self-management: Autonomic functions self-configure, and negotiate parameters and settings across the network. [RFC7575] defines the fundamental ideas and design goals of Autonomic Networking. A gap analysis of Autonomic Networking is given in [RFC7576]. The reference architecture for Autonomic Networking in the IETF is specified in the document [I-D.ietf-anima-reference-model].

Autonomic functions need an autonomically built communications infrastructure. This infrastructure needs to be secure, resilient and re-usable by all autonomic functions. Section 5 of [RFC7575] introduces that infrastructure and calls it the Autonomic Control Plane (ACP). More descriptively it would be the "Autonomic communications infrastructure for OAM and Control". For naming consistency with that prior document, this document continues to use the name ACP though.

Today, the OAM and control plane of IP networks is what is typically called in-band management/signaling: Its management and control protocol traffic depends on the routing and forwarding tables, security, policy, QoS and potentially other configuration that first has to be established through the very same management and control protocols. Misconfigurations including unexpected side effects or mutual dependences can disrupt OAM and control operations and especially disrupt remote management access to the affected node itself and potentially a much larger number of nodes for whom the affected node is on the network path.

For an example of inband management failing in the face of operator induced misconfiguration, see [FCC], for example III.B.15 on page 8: "...engineers almost immediately recognized that they had misdiagnosed the problem. However, they were unable to resolve the issue by restoring the link because the network management tools required to do so remotely relied on the same paths they had just disabled".

Traditionally, physically separate, so-called out-of-band (management) networks have been used to avoid these problems or at least to allow recovery from such problems. Worst case, personnel are sent on site to access devices through out-of-band management ports (also called craft ports, serial console, management ethernet port). However, both options are expensive.

In increasingly automated networks either centralized management systems or distributed autonomic service agents in the network require a control plane which is independent of the configuration of the network they manage, to avoid impacting their own operations through the configuration actions they take.

This document describes a modular design for a self-forming, self-managing and self-protecting ACP, which is a virtual out-of-band network designed to be as independent as possible of configuration, addressing and routing to avoid the self-dependency problems of current IP networks while still operating in-band on the same physical network that it is controlling and managing. The ACP design is therefore intended to combine as well as possible the resilience of out-of-band management networks with the low-cost of traditional IP in-band network management. The details how this is achieved are described in Section 6.

In a fully autonomic network node without legacy control or management functions/protocols, the Data-Plane would be for example just a forwarding plane for "Data" IPv6 packets, aka: packets other than the control and management plane packets that are forwarded by the ACP itself. In such networks/nodes, there would be no non-autonomous control or non-autonomous management plane.

Routing protocols for example would be built inside the ACP as so-called autonomous functions via autonomous service agents, leveraging the ACP's functions instead of implementing them separately for each protocol: discovery, automatically established authenticated and encrypted local and distant peer connectivity for control and management traffic, and common control/management protocol session and presentation functions.

When ACP functionality is added to nodes that have non-autonomous management plane and/or control plane functions (henceforth called non-autonomous nodes), the ACP instead is best abstracted as a special Virtual Routing and Forwarding (VRF) instance (or virtual router) and the complete pre-existing non-autonomous management and/or control plane is considered to be part of the Data-Plane to avoid introduction of more complex, new terminology only for this case.

Like the forwarding plane for "Data" packets, the non-autonomous control and management plane functions can then be managed/used via the ACP. This terminology is consistent with pre-existing documents such as [RFC8368].

In both instances (autonomous and non-autonomous nodes), the ACP is built such that it is operating in the absence of the Data-Plane, and in the case of existing non-autonomous (management, control) components in the Data-Plane also in the presence of any (mis-)configuration thereof.

The Autonomic Control Plane serves several purposes at the same time:

1. Autonomic functions communicate over the ACP. The ACP therefore directly supports Autonomic Networking functions, as described in [I-D.ietf-anima-reference-model]. For example, Generic Autonomic Signaling Protocol (GRASP - [I-D.ietf-anima-grasp]) runs securely inside the ACP and depends on the ACP as its "security and transport substrate".
2. A controller or network management system can use it to securely bootstrap network devices in remote locations, even if the (Data-Plane) network in between is not yet configured; no Data-Plane dependent bootstrap configuration is required. An example of such a secure bootstrap process is described in [I-D.ietf-anima-bootstrapping-keyinfra].
3. An operator can use it to access remote devices using protocols such as Secure SHell (SSH) or Network Configuration Protocol (NETCONF) running across the ACP, even if the network is misconfigured or not configured.

This document describes these purposes as use cases for the ACP in Section 3, it defines the requirements in Section 4. Section 5 gives an overview of how the ACP is constructed.

The normative part of this document starts with Section 6, where the ACP is specified. Section 7 explains how to support ACP on L2 switches (normative). Section 8 explains how non-ACP nodes and networks can be integrated (normative).

The remaining sections are non-normative: Section 10 reviews benefits of the ACP (after all the details have been defined), Section 9 provides operational recommendations, Appendix A provides additional explanations and describes additional details or future standard or proprietary extensions that were considered not to be appropriate for standardization in this document but were considered important to document. There are no dependencies against Appendix A to build a complete working and interoperable ACP according to this document.

The ACP provides secure IPv6 connectivity, therefore it can be used not only as the secure connectivity for self-management as required for the ACP in [RFC7575], but it can also be used as the secure connectivity for traditional (centralized) management. The ACP can be implemented and operated without any other components of autonomic networks, except for the GRASP protocol. ACP relies on per-link DULL GRASP (see Section 6.4) to autodiscover ACP neighbors, and includes the ACP GRASP instance to provide service discovery for clients of the ACP (see Section 6.9) including for its own maintenance of ACP certificates.

The document "Using Autonomic Control Plane for Stable Connectivity of Network OAM" [RFC8368] describes how the ACP alone can be used to provide secure and stable connectivity for autonomic and non-autonomic OAM applications, specifically for the case of current non-autonomic networks/nodes. That document also explains how existing management solutions can leverage the ACP in parallel with traditional management models, when to use the ACP and how to integrate with potentially IPv4 only OAM backends.

Combining ACP with Bootstrapping Remote Secure Key Infrastructures (BRSKI), see [I-D.ietf-anima-bootstrapping-keyinfra]) results in the "Autonomic Network Infrastructure" (ANI) as defined in [I-D.ietf-anima-reference-model], which provides autonomic connectivity (from ACP) with secure zero-touch (automated) bootstrap from BRSKI. The ANI itself does not constitute an Autonomic Network, but it allows the building of more or less autonomic networks on top of it - using either centralized, Software Defined Networking- (SDN-)style (see [RFC7426]) automation or distributed automation via Autonomic Service Agents (ASA) / Autonomic Functions (AF) - or a mixture of both. See [I-D.ietf-anima-reference-model] for more information.

### 1.1. Applicability and Scope

Please see the following Terminology section (Section 2) for explanations of terms used in this section.

The design of the ACP as defined in this document is considered to be applicable to all types of "professionally managed" networks: Service Provider, Local Area Network (LAN), Metro(politan networks), Wide Area Network (WAN), Enterprise Information Technology (IT) and ->"Operational Technology" (OT) networks. The ACP can operate equally on layer 3 equipment and on layer 2 equipment such as bridges (see Section 7). The hop-by-hop authentication, integrity-protection and confidentiality mechanism used by the ACP is defined to be negotiable, therefore it can be extended to environments with different protocol preferences. The minimum implementation

requirements in this document attempt to achieve maximum interoperability by requiring support for multiple options depending on the type of device: IPsec, see [RFC4301], and Datagram Transport Layer Security (DTLS, see Section 6.8.4).

The implementation footprint of the ACP consists of Public Key Infrastructure (PKI) code for the ACP certificate including "Enrollment over Secure Transport (EST, see [RFC7030]), the GRASP protocol, UDP, TCP and Transport Layer Security (TLS, see Section 6.1), for security and reliability of GRASP and for EST, the ACP secure channel protocol used (such as IPsec or DTLS), and an instance of IPv6 packet forwarding and routing via the Routing Protocol for Low-power and Lossy Networks (RPL), see [RFC6550], that is separate from routing and forwarding for the Data-Plane (user traffic).

The ACP uses only IPv6 to avoid complexity of dual-stack ACP operations (IPv6/IPv4). Nevertheless, it can without any changes be integrated into even otherwise IPv4-only network devices. The Data-Plane itself would not need to change and it could continue to be IPv4 only. For such IPv4-only devices, the IPv6 protocol itself would be additional implementation footprint that is only required for the ACP.

The protocol choices of the ACP are primarily based on wide use and support in networks and devices, well understood security properties and required scalability. The ACP design is an attempt to produce the lowest risk combination of existing technologies and protocols to build a widely applicable operational network management solution.

RPL was chosen because it requires a smaller routing table footprint in large networks compared to other routing protocols with an autonomically configured single area. The deployment experience of large scale Internet of Things (IoT) networks serves as the basis for wide deployment experience with RPL. The profile chosen for RPL in the ACP does not leverage any RPL specific forwarding plane features (IPv6 extension headers), making its implementation a pure control plane software requirement.

GRASP is the only completely novel protocol used in the ACP, and this choice was necessary because there is no existing suitable protocol to provide the necessary functions to the ACP, so GRASP was developed to fill that gap.

The ACP design can be applicable to devices constrained with respect to cpu and memory, and to networks constrained with respect to bitrate and reliability, but this document does not attempt to define the most constrained type of devices or networks to which the ACP is

applicable. RPL and DTLS for ACP secure channels are two protocol choices already making ACP more applicable to constrained environments. Support for constrained devices in this specification is opportunistic, but not complete, because the reliable transport for GRASP (see Section 6.9.2) only specifies TCP/TLS. See Appendix A.8 for discussions about how future standards or proprietary extensions/variations of the ACP could better meet different expectations from those on which the current design is based including supporting constrained devices better.

## 2. Acronyms and Terminology (Informative)

[RFC-Editor: Please add ACP, BRSKI, GRASP, MASA to <https://www.rfc-editor.org/materials/abbrev.expansion.txt>.]

[RFC-Editor: What is the recommended way to reference a hanging text, e.g. to a definition in the list of definitions? Up to -28, this document was using XMLv2 and the only option I could find for RFC/XML to point to a hanging text was `format="title"`, which leads to references such as `'->"ACP certificate" ()'`, aka: redundant empty parenthesis. Many reviewers were concerned about this. I created a ticket to ask for an xml2rfc enhancement to avoid this in the future: <https://trac.tools.ietf.org/tools/xml2rfc/trac/ticket/347>. When I changed to XMLv3 in version -29, I could get rid of the unnecessary `()` by using `format="none"`, but that format is declared to be deprecated in XMLv3. So I am not aware of any working AND "non-deprecated" option.]

[RFC-Editor: Question: Is it possible to change the first occurrences of [RFCxxxx] references to "rfcxxx title" [RFCxxxx]? the XML2RFC format does not seem to offer such a format, but I did not want to duplicate 50 first references - one reference for title mentioning and one for RFC number.]

This document serves both as a normative specification for how ACP nodes have to behave as well as describing requirements, benefits, architecture and operational aspects to explain the context. Normative sections are labelled "(Normative)" and use BCP 14 keywords. Other sections are labelled "(Informative)" and do not use those normative keywords.

In the rest of the document we will refer to systems using the ACP as "nodes". Typically, such a node is a physical (network equipment) device, but it can equally be some virtualized system. Therefore, we do not refer to them as devices unless the context specifically calls for a physical system.

This document introduces or uses the following terms (sorted alphabetically). Terms introduced are explained on first use, so this list is for reference only.

- ACP: "Autonomic Control Plane". The Autonomic Function as defined in this document. It provides secure zero-touch (automated) transitive (network wide) IPv6 connectivity for all nodes in the same ACP domain as well as a GRASP instance running across this ACP IPv6 connectivity. The ACP is primarily meant to be used as a component of the ANI to enable Autonomic Networks but it can equally be used in simple ANI networks (with no other Autonomic Functions) or completely by itself.
- ACP address: An IPv6 address assigned to the ACP node. It is stored in the `acp-node-name` of the `->"ACP certificate"`.
- ACP address range/set: The ACP address may imply a range or set of addresses that the node can assign for different purposes. This address range/set is derived by the node from the format of the ACP address called the "addressing sub-scheme".
- ACP connect interface: An interface on an ACP node providing access to the ACP for non ACP capable nodes without using an ACP secure channel. See Section 8.1.1.
- ACP domain: The ACP domain is the set of nodes with `->"ACP certificates"` that allow them to authenticate each other as members of the ACP domain. See also Section 6.2.3.
- ACP (ANI/AN) certificate: A [RFC5280] certificate (LDevID) carrying the `acp-node-name` which is used by the ACP to learn its address in the ACP and to derive and cryptographically assert its membership in the ACP domain.
- ACP `acp-node-name` field: An information field in the ACP certificate in which the ACP relevant information is encoded: the ACP domain name, the ACP IPv6 address of the node and optional additional role attributes about the node.
- ACP Loopback interface: The Loopback interface in the ACP Virtual Routing and Forwarding (VRF) that has the ACP address assigned to it. See Section 6.13.5.1.
- ACP network: The ACP network constitutes all the nodes that have access to the ACP. It is the set of active and transitively connected nodes of an ACP domain plus all nodes that get access to the ACP of that domain via ACP edge nodes.
- ACP (ULA) prefix(es): The /48 IPv6 address prefixes used across the ACP. In the normal/simple case, the ACP has one ULA prefix, see Section 6.11. The ACP routing table may include multiple ULA prefixes if the "rsub" option is used to create addresses from more than one ULA prefix. See Section 6.2.2. The ACP may also include non-ULA prefixes if those are configured on ACP connect interfaces. See Section 8.1.1.
- ACP secure channel: A channel authenticated via `->"ACP certificates"`

providing integrity protection and confidentiality through encryption. These are established between (normally) adjacent ACP nodes to carry traffic of the ACP VRF securely and isolated from Data-Plane traffic in-band over the same link/path as the Data-Plane.

- ACP secure channel protocol: The protocol used to build an ACP secure channel, e.g., Internet Key Exchange Protocol version 2 (IKEv2) with IPsec or Datagram Transport Layer Security (DTLS).
- ACP virtual interface: An interface in the ACP VRF mapped to one or more ACP secure channels. See Section 6.13.5.
- AN "Autonomic Network": A network according to [I-D.ietf-anima-reference-model]. Its main components are ANI, Autonomic Functions and Intent.
- (AN) Domain Name: An FQDN (Fully Qualified Domain Name) in the acp-node-name of the Domain Certificate. See Section 6.2.2.
- ANI (nodes/network): "Autonomic Network Infrastructure". The ANI is the infrastructure to enable Autonomic Networks. It includes ACP, BRSKI and GRASP. Every Autonomic Network includes the ANI, but not every ANI network needs to include autonomic functions beyond the ANI (nor Intent). An ANI network without further autonomic functions can for example support secure zero-touch (automated) bootstrap and stable connectivity for SDN networks - see [RFC8368].
- ANIMA: "Autonomic Networking Integrated Model and Approach". ACP, BRSKI and GRASP are specifications of the IETF ANIMA working group.
- ASA: "Autonomic Service Agent". Autonomic software modules running on an ANI device. The components making up the ANI (BRSKI, ACP, GRASP) are also described as ASAs.
- Autonomic Function: A function/service in an Autonomic Network (AN) composed of one or more ASA across one or more ANI nodes.
- BRSKI: "Bootstrapping Remote Secure Key Infrastructures" ([I-D.ietf-anima-bootstrapping-keyinfra]. A protocol extending EST to enable secure zero-touch bootstrap in conjunction with ACP. ANI nodes use ACP, BRSKI and GRASP.
- CA: "Certification Authority". An entity that issues digital certificates. A CA uses its private key to sign the certificates it issues. Relying parties use the public key in the CA certificate to validate the signature.
- CRL: "Certificate Revocation List". A list of revoked certificates. Required to revoke certificates before their lifetime expires.
- Data-Plane: The counterpoint to the ACP VRF in an ACP node: forwarding of user traffic and in non-autonomous nodes/networks also any non-autonomous control and/or management plane functions. In a fully Autonomic Network node, the Data-Plane is managed autonomically via Autonomic Functions and Intent. See Section 1 for more detailed explanations.
- device: A physical system, or physical node.

**Enrollment:** The process through which a node authenticates itself to a network with an initial identity, which is often called IDevID certificate, and acquires from the network a network specific identity, which is often called LDevID certificate, and certificates of one or more Trust Anchor(s). In the ACP, the LDevID certificate is called the ACP certificate.

**EST:** "Enrollment over Secure Transport" ([RFC7030]). IETF standard-track protocol for enrollment of a node with an LDevID certificate. BRSKI is based on EST.

**GRASP:** "Generic Autonomic Signaling Protocol". An extensible signaling protocol required by the ACP for ACP neighbor discovery. The ACP also provides the "security and transport substrate" for the "ACP instance of GRASP". This instance of GRASP runs across the ACP secure channels to support BRSKI and other NOC/OAM or Autonomic Functions. See [I-D.ietf-anima-grasp].

**IDevID:** An "Initial Device IDentity" X.509 certificate installed by the vendor on new equipment. Contains information that establishes the identity of the node in the context of its vendor/manufacturer such as device model/type and serial number. See [AR8021]. The IDevID certificate cannot be used as a node identifier for the ACP because they are not provisioned by the owner of the network, so they can not directly indicate an ACP domain they belong to.

**in-band (management/signaling):** In-band management traffic and/or control plane signaling uses the same network resources such as routers/switches and network links that it manages/controls. In-band is the standard management and signaling mechanism in IP networks. Compared to ->"out-of-band" it requires no additional physical resources, but introduces potentially circular dependencies for its correct operations. See ->"introduction".

**Intent:** Policy language of an autonomic network according to [I-D.ietf-anima-reference-model].

**Loopback interface:** See ->"ACP Loopback interface".

**LDevID:** A "Local Device IDentity" is an X.509 certificate installed during "enrollment". The Domain Certificate used by the ACP is an LDevID certificate. See [AR8021].

**Management:** Used in this document as another word for ->"OAM".

**MASA (service):** "Manufacturer Authorized Signing Authority". A vendor/manufacturer or delegated cloud service on the Internet used as part of the BRSKI protocol.

**MIC:** "Manufacturer Installed Certificate". This is another word to describe an IDevID in referenced materials. This term is not used in this document.

**native interface:** Interfaces existing on a node without configuration of the already running node. On physical nodes these are usually physical interfaces; on virtual nodes their equivalent.

**NOC:** Network Operations Center.

- node: A system supporting the ACP according to this document. Can be virtual or physical. Physical nodes are called devices.
- Node-ID: The identifier of an ACP node inside that ACP. It is the last 64 (see Section 6.11.3) or 78-bits (see Section 6.11.5) of the ACP address.
- OAM: Operations, Administration and Management. Includes Network Monitoring.
- Operational Technology (OT): [https://en.wikipedia.org/wiki/Operational\\_Technology](https://en.wikipedia.org/wiki/Operational_Technology): "The hardware and software dedicated to detecting or causing changes in physical processes through direct monitoring and/or control of physical devices such as valves, pumps, etc.". OT networks are today in most cases well separated from Information Technology (IT) networks.
- out-of-band (management) network: An out-of-band network is a secondary network used to manage a primary network. The equipment of the primary network is connected to the out-of-band network via dedicated management ports on the primary network equipment. Serial (console) management ports were historically most common, higher end network equipment now also has ethernet ports dedicated only for management. An out-of-band network provides management access to the primary network independent of the configuration state of the primary network. See ->"Introduction"
- (virtual) out-of-band network: The ACP can be called a virtual out-of-band network for management and control because it attempts to provide the benefits of a (physical) ->"out-of-band network" even though it is physically carried ->"in-band". See ->"introduction".
- root CA: "root Certification Authority". A ->"CA" for which the root CA Key update procedures of [RFC7030], Section 4.4 can be applied.
- RPL: "IPv6 Routing Protocol for Low-Power and Lossy Networks". The routing protocol used in the ACP. See [RFC6550].
- (ACP/ANI/BRSKI) Registrar: An ACP registrar is an entity (software and/or person) that is orchestrating the enrollment of ACP nodes with the ACP certificate. ANI nodes use BRSKI, so ANI registrars are also called BRSKI registrars. For non-ANI ACP nodes, the registrar mechanisms are undefined by this document. See Section 6.11.7. Renewal and other maintenance (such as revocation) of ACP certificates may be performed by other entities than registrars. EST must be supported for ACP certificate renewal (see Section 6.2.5). BRSKI is an extension of EST, so ANI/BRSKI registrars can easily support ACP domain certificate renewal in addition to initial enrollment.
- RPI: "RPL Packet Information". Network extension headers for use with the ->"RPL" routing protocols. Not used with RPL in the ACP. See Section 6.12.1.13.
- RPL: "Routing Protocol for Low-Power and Lossy Networks". The routing protocol used in the ACP. See Section 6.12.

sUDI: "secured Unique Device Identifier". This is another word to describe an IDevID in referenced material. This term is not used in this document.

TA: "Trust Anchor". A Trust Anchor is an entity that is trusted for the purpose of certificate validation. Trust Anchor Information such as self-signed certificate(s) of the Trust Anchor is configured into the ACP node as part of Enrollment. See [RFC5280], Section 6.1.1.

UDI: "Unique Device Identifier". In the context of this document unsecured identity information of a node typically consisting of at least device model/type and serial number, often in a vendor specific format. See sUDI and LDevID.

ULA: (Global ID prefix) A "Unique Local Address" (ULA) is an IPv6 address in the block fc00::/7, defined in [RFC4193]. ULA is the IPv6 successor of the IPv4 private address space ([RFC1918]). ULAs have important differences over IPv4 private addresses that are beneficial for and exploited by the ACP, such as the Locally Assigned Global ID prefix, which are the first 48-bits of a ULA address [RFC4193], section 3.2.1. In this document this prefix is abbreviated as "ULA prefix".

(ACP) VRF: The ACP is modeled in this document as a "Virtual Routing and Forwarding" instance (VRF). This means that it is based on a "virtual router" consisting of a separate IPv6 forwarding table to which the ACP virtual interfaces are attached and an associated IPv6 routing table separate from the Data-Plane. Unlike the VRFs on MPLS/VPN-PE ([RFC4364]) or LISP XTR ([RFC6830]), the ACP VRF does not have any special "core facing" functionality or routing/mapping protocols shared across multiple VRFs. In vendor products a VRF such as the ACP-VRF may also be referred to as a so called VRF-lite.

(ACP) Zone: An ACP zone is a set of ACP nodes using the same zone field value in their ACP address according to Section 6.11.3. Zones are a mechanism to support structured addressing of ACP addresses within the same /48-bit ULA prefix.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119],[RFC8174] when, and only when, they appear in all capitals, as shown here.

### 3. Use Cases for an Autonomic Control Plane (Informative)

This section summarizes the use cases that are intended to be supported by an ACP. To understand how these are derived from and relate to the larger set of use cases for autonomic networks, please refer to [RFC8316].

### 3.1. An Infrastructure for Autonomic Functions

Autonomic Functions need a stable infrastructure to run on, and all autonomic functions should use the same infrastructure to minimize the complexity of the network. In this way, there is only need for a single discovery mechanism, a single security mechanism, and single instances of other processes that distributed functions require.

### 3.2. Secure Bootstrap over a not configured Network

Today, bootstrapping a new node typically requires all nodes between a controlling node such as an SDN controller ("Software Defined Networking", see [RFC7426]) and the new node to be completely and correctly addressed, configured and secured. Bootstrapping and configuration of a network happens in rings around the controller - configuring each ring of devices before the next one can be bootstrapped. Without console access (for example through an out-of-band network) it is not possible today to make devices securely reachable before having configured the entire network leading up to them.

With the ACP, secure bootstrap of new devices and whole new networks can happen without requiring any configuration of unconfigured devices along the path: As long as all devices along the path support ACP and a zero-touch bootstrap mechanism such as BRSKI, the ACP across a whole network of unconfigured devices can be brought up without operator/provisioning intervention. The ACP also provides additional security for any bootstrap mechanism, because it can provide encrypted discovery (via ACP GRASP) of registrars or other bootstrap servers by bootstrap proxies connecting to nodes that are to be bootstrapped and the ACP encryption hides the identities of the communicating entities (pledge and registrar), making it more difficult to learn which network node might be attackable. The ACP certificate can also be used to end-to-end encrypt the bootstrap communication between such proxies and server. Note that bootstrapping here includes not only the first step that can be provided by BRSKI (secure keys), but also later stages where configuration is bootstrapped.

### 3.3. Data-Plane Independent Permanent Reachability

Today, most critical control plane protocols and OAM protocols are using the Data-Plane of the network. This leads to often undesirable dependencies between control and OAM plane on one side and the Data-Plane on the other: Only if the forwarding and control plane of the Data-Plane are configured correctly, will the Data-Plane and the OAM/control plane work as expected.

Data-Plane connectivity can be affected by errors and faults, for example misconfigurations that make AAA (Authentication, Authorization and Accounting) servers unreachable or can lock an administrator out of a device; routing or addressing issues can make a device unreachable; shutting down interfaces over which a current management session is running can lock an admin irreversibly out of the device. Traditionally only out-of-band access can help recover from such issues (such as serial console or ethernet management port).

Data-Plane dependencies also affect applications in a Network Operations Center (NOC) such as SDN controller applications: Certain network changes are today hard to implement, because the change itself may affect reachability of the devices. Examples are address or mask changes, routing changes, or security policies. Today such changes require precise hop-by-hop planning.

Note that specific control plane functions for the Data-Plane often want to depend on forwarding of their packets via the Data-Plane: Aliveness and routing protocol signaling packets across the Data-Plane to verify reachability across the Data-Plane, using IPv4 signaling packets for IPv4 routing vs. IPv6 signaling packets for IPv6 routing.

Assuming appropriate implementation (see Section 6.13.2 for more details), the ACP provides reachability that is independent of the Data-Plane. This allows the control plane and OAM plane to operate more robustly:

- \* For management plane protocols, the ACP provides the functionality of a Virtual out-of-band (VooB) channel, by providing connectivity to all nodes regardless of their Data-Plane configuration, routing and forwarding tables.
- \* For control plane protocols, the ACP allows their operation even when the Data-Plane is temporarily faulty, or during transitional events, such as routing changes, which may affect the control plane at least temporarily. This is specifically important for autonomic service agents, which could affect Data-Plane connectivity.

The document "Using Autonomic Control Plane for Stable Connectivity of Network OAM" [RFC8368] explains this use case for the ACP in significantly more detail and explains how the ACP can be used in practical network operations.

#### 4. Requirements (Informative)

The following requirements were identified for the design of the ACP based on the above use-cases (Section 3). These requirements are informative. The ACP as specified in the normative parts of this document is meeting or exceeding these use-case requirements:

- ACP1: The ACP should provide robust connectivity: As far as possible, it should be independent of configured addressing, configuration and routing. Requirements 2 and 3 build on this requirement, but also have value on their own.
- ACP2: The ACP must have a separate address space from the Data-Plane. Reason: traceability, debug-ability, separation from Data-Plane, infrastructure security (filtering based on known address space).
- ACP3: The ACP must use autonomically managed address space. Reason: easy bootstrap and setup ("autonomic"); robustness (admin cannot break network easily). This document uses Unique Local Addresses (ULA) for this purpose, see [RFC4193].
- ACP4: The ACP must be generic, that is it must be usable by all the functions and protocols of the ANI. Clients of the ACP must not be tied to a particular application or transport protocol.
- ACP5: The ACP must provide security: Messages coming through the ACP must be authenticated to be from a trusted node, and it is very strongly > recommended that they be encrypted.

Explanation for ACP4: In a fully autonomic network (AN), newly written ASAs could potentially all communicate exclusively via GRASP with each other, and if that was assumed to be the only requirement against the ACP, it would not need to provide IPv6 layer connectivity between nodes, but only GRASP connectivity. Nevertheless, because ACP also intends to support non-AN networks, it is crucial to support IPv6 layer connectivity across the ACP to support any transport and application layer protocols.

The ACP operates hop-by-hop, because this interaction can be built on IPv6 link local addressing, which is autonomic, and has no dependency on configuration (requirement 1). It may be necessary to have ACP connectivity across non-ACP nodes, for example to link ACP nodes over the general Internet. This is possible, but introduces a dependency against stable/resilient routing over the non-ACP hops (see Section 8.2).

## 5. Overview (Informative)

When a node has an ACP certificate (see Section 6.2.1) and is enabled to bring up the ACP (see Section 9.3.5), it will create its ACP without any configuration as follows. For details, see Section 6 and further sections:

1. The node creates a VRF instance, or a similar virtual context for the ACP.
2. The node assigns its ULA IPv6 address (prefix) (see Section 6.11 which is learned from the `acp-node-name` (see Section 6.2.2) of its ACP certificate (see Section 6.2.1) to an ACP loopback interface (see Section 6.11) and connects this interface into the ACP VRF.
3. The node establishes a list of candidate peer adjacencies and candidate channel types to try for the adjacency. This is automatic for all candidate link-local adjacencies, see Section 6.4 across all native interfaces (see Section 9.3.4). If a candidate peer is discovered via multiple interfaces, this will result in one adjacency per interface. If the ACP node has multiple interfaces connecting to the same subnet across which it is also operating as an L2 switch in the Data-Plane, it employs methods for ACP with L2 switching, see Section 7.
4. For each entry in the candidate adjacency list, the node negotiates a secure tunnel using the candidate channel types. See Section 6.6.
5. The node authenticates the peer node during secure channel setup and authorizes it to become part of the ACP according to Section 6.2.3.
6. Unsuccessful authentication of a candidate peer results in throttled connection retries for as long as the candidate peer is discoverable. See Section 6.7.
7. Each successfully established secure channel is mapped into an ACP virtual interface, which is placed into the ACP VRF. See Section 6.13.5.2.
8. Each node runs a lightweight routing protocol, see Section 6.12, to announce reachability of the ACP loopback address (or prefix) across the ACP.
9. This completes the creation of the ACP with hop-by-hop secure tunnels, auto-addressing and auto-routing. The node is now an ACP node with a running ACP.

### Note:

- \* None of the above operations (except the following explicit configured ones) are reflected in the configuration of the node.
- \* Non-ACP NMS ("Network Management Systems") or SDN controllers have to be explicitly configured for connection into the ACP.

\* Additional candidate peer adjacencies for ACP connections across non-ACP Layer-3 clouds requires explicit configuration. See Section 8.2.

The following figure illustrates the ACP.

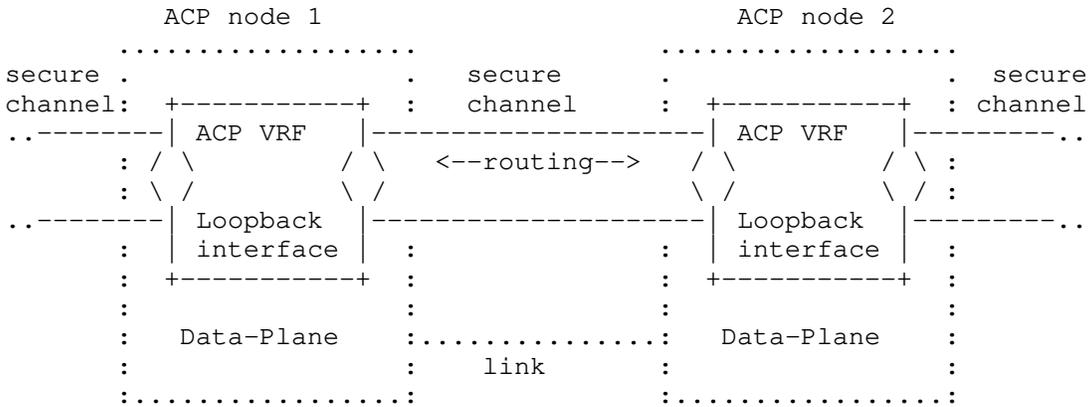


Figure 1: ACP VRF and secure channels

The resulting overlay network is normally based exclusively on hop-by-hop tunnels. This is because addressing used on links is IPv6 link local addressing, which does not require any prior set-up. In this way the ACP can be built even if there is no configuration on the node, or if the Data-Plane has issues such as addressing or routing problems.

6. Self-Creation of an Autonomic Control Plane (ACP) (Normative)

This section specifies the components and steps to set up an ACP. The ACP is automatically "self-creating", which makes it "indestructible" against most changes to the Data-Plane, including misconfigurations of routing, addressing, NAT, firewall or any other traffic policy filters that inadvertently or otherwise unavoidably would also impact the management plane traffic, such as the actual operator CLI session or controller NETCONF session through which the configuration changes to the Data-Plane are executed.

Physical misconfiguration of wiring between ACP nodes will also not break the ACP: As long as there is a transitive physical path between ACP nodes, the ACP should be able to recover given that it automatically operates across all interfaces of the ACP nodes and automatically determines paths between them.

Attacks against the network via incorrect routing or addressing information for the Data-Plane will not impact the ACP. Even impaired ACP nodes will have a significantly reduced attack surface against malicious misconfiguration because only very limited ACP or interface up/down configuration can affect the ACP, and pending on their specific designs these type of attacks could also be eliminated. See more in Section 9.3 and Section 11.

An ACP node can be a router, switch, controller, NMS host, or any other IPv6 capable node. Initially, it MUST have its ACP certificate, as well as an (empty) ACP Adjacency Table (described in Section 6.3). It then can start to discover ACP neighbors and build the ACP. This is described step by step in the following sections:

#### 6.1. Requirements for use of Transport Layer Security (TLS)

The following requirements apply to TLS required or used by ACP components. Applicable ACP components include ACP certificate maintenance via EST, see Section 6.2.5, TLS connections for Certificate Revocation List (CRL) Distribution Point (CRLDP) or Online Certificate Status Protocol (OCSP) responder (if used, see Section 6.2.3) and ACP GRASP (see Section 6.9.2). On ANI nodes these requirements also apply to BRSKI.

TLS MUST comply with [RFC7525] except that TLS 1.2 ([RFC5246]) is REQUIRED and that older versions of TLS MUST NOT be used. TLS 1.3 ([RFC8446]) SHOULD be supported. The choice for TLS 1.2 as the lowest common denominator for the ACP is based on current expected most likely availability across the wide range of candidate ACP node types, potentially with non-agile operating system TCP/IP stacks.

TLS MUST offer TLS\_ECDHE\_RSA\_WITH\_AES\_256\_GCM\_SHA384 and TLS\_ECDHE\_ECDSA\_WITH\_AES\_256\_GCM\_SHA384 and MUST NOT offer options with less than 256-bit symmetric key strength or hash strength of less than 384 bits. When TLS 1.3 is supported, TLS\_AES\_256\_GCM\_SHA384 MUST be offered and TLS\_CHACHA20\_POLY1305\_SHA256 MAY be offered.

TLS MUST also include the "Supported Elliptic Curves" extension, it MUST support the NIST P-256 (secp256r1(22)) and P-384 (secp384r1(24)) curves [RFC8422]. In addition, TLS 1.2 clients SHOULD send an ec\_point\_format extension with a single element, "uncompressed".

## 6.2. ACP Domain, Certificate and Network

The ACP relies on group security. An ACP domain is a group of nodes that trust each other to participate in ACP operations such as creating ACP secure channels in an autonomous peer-to-peer fashion between ACP domain members via protocols such as IPsec. To authenticate and authorize another ACP member node with access to the ACP Domain, each ACP member requires keying material: An ACP node MUST have a Local Device IDentity (LDevID) certificate, henceforth called the ACP certificate and information about one or more Trust Anchor (TA) as required for the ACP domain membership check (Section 6.2.3).

Manual keying via shared secrets is not usable for an ACP domain because it would require a single shared secret across all current and future ACP domain members to meet the expectation of autonomous, peer-to-peer establishment of ACP secure channels between any ACP domain members. Such a single shared secret would be an unacceptable security weakness. Asymmetric keying material (public keys) without certificates does not provide the mechanisms to authenticate ACP domain membership in an autonomous, peer-to-peer fashion for current and future ACP domain members.

The LDevID certificate is called the ACP certificate. The TA is the Certification Authority (CA) root certificate of the ACP domain.

The ACP does not mandate specific mechanisms by which this keying material is provisioned into the ACP node. It only requires the certificate to comply with Section 6.2.1, specifically to have the `acp-node-name` as specified in Section 6.2.2 in its domain certificate as well as those of candidate ACP peers. See Appendix A.2 for more information about enrollment or provisioning options.

This document uses the term ACP in many places where the Autonomic Networking reference documents [RFC7575] and [I-D.ietf-anima-reference-model] use the word `autonomic`. This is done because those reference documents consider (only) fully autonomic networks and nodes, but support of ACP does not require support for other components of autonomic networks except for relying on GRASP and providing security and transport for GRASP. Therefore, the word `autonomic` might be misleading to operators interested in only the ACP.

[RFC7575] defines the term "Autonomic Domain" as a collection of autonomic nodes. ACP nodes do not need to be fully autonomic, but when they are, then the ACP domain is an autonomic domain. Likewise, [I-D.ietf-anima-reference-model] defines the term "Domain Certificate" as the certificate used in an autonomic domain. The ACP

certificate is that domain certificate when ACP nodes are (fully) autonomic nodes. Finally, this document uses the term ACP network to refer to the network created by active ACP nodes in an ACP domain. The ACP network itself can extend beyond ACP nodes through the mechanisms described in Section 8.1.

#### 6.2.1. ACP Certificates

ACP certificates MUST be [RFC5280] compliant X.509 v3 ([X.509]) certificates.

ACP nodes MUST support handling ACP certificates, TA certificates and certificate chain certificates (henceforth just called certificates in this section) with RSA public keys and certificates with Elliptic Curve (ECC) public keys.

ACP nodes MUST NOT support certificates with RSA public keys of less than 2048-bit modulus or curves with group order of less than 256-bit. They MUST support certificates with RSA public keys with 2048-bit modulus and MAY support longer RSA keys. They MUST support certificates with ECC public keys using NIST P-256 curves and SHOULD support P-384 and P-521 curves.

ACP nodes MUST NOT support certificates with RSA public keys whose modulus is less than 2048 bits, or certificates whose ECC public keys are in groups whose order is less than 256-bits. RSA signing certificates with 2048-bit public keys MUST be supported, and such certificates with longer public keys MAY be supported. ECDSA certificates using the NIST P-256 curve MUST be supported, and such certificates using the P-384 and P-521 curves SHOULD be supported.

ACP nodes MUST support RSA certificates that are signed by RSA signatures over the SHA-256 digest of the contents, and SHOULD additionally support SHA-384 and SHA-512 digests in such signatures. The same requirements for digest usage in certificate signatures apply to ECDSA certificates, and additionally, ACP nodes MUST support ECDSA signatures on ECDSA certificates.

The ACP certificate SHOULD use an RSA key and an RSA signature when the ACP certificate is intended to be used not only for ACP authentication but also for other purposes. The ACP certificate MAY use an ECC key and an ECDSA signature if the ACP certificate is only used for ACP and ANI authentication and authorization.

Any secure channel protocols used for the ACP as specified in this document or extensions of this document MUST therefore support authentication (e.g. signing) starting with these type of certificates. See [RFC8422] for more information.

The reason for these choices are as follows: As of 2020, RSA is still more widely used than ECC, therefore the MUST for RSA. ECC offers equivalent security at (logarithmically) shorter key lengths (see [RFC8422]). This can be beneficial especially in the presence of constrained bandwidth or constrained nodes in an ACP/ANI network. Some ACP functions such as GRASP peer-2-peer across the ACP require end-to-end/any-to-any authentication/authorization, therefore ECC can only reliably be used in the ACP when it MUST be supported on all ACP nodes. RSA signatures are mandatory to be supported also for ECC certificates because CAs themselves may not support ECC yet.

The ACP certificate SHOULD be used for any authentication between nodes with ACP domain certificates (ACP nodes and NOC nodes) where a required authorization condition is ACP domain membership, such as ACP node to NOC/OAM end-to-end security and ASA to ASA end-to-end security. Section 6.2.3 defines this "ACP domain membership check". The uses of this check that are standardized in this document are for the establishment of hop-by-hop ACP secure channels (Section 6.7) and for ACP GRASP (Section 6.9.2) end-to-end via TLS.

The ACP domain membership check requires a minimum amount of elements in a certificate as described in Section 6.2.3. The identity of a node in the ACP is carried via the `acp-node-name` as defined in Section 6.2.2.

To support ECDH directly with the key in the ACP certificate, ACP certificates with ECC keys need to indicate to be Elliptic Curve Diffie-Hellman capable (ECDH): If the X.509v3 `keyUsage` extension is present, the `keyAgreement` bit must then be set. Note that this option is not required for any of the required ciphersuites in this document and may not be supported by all CA.

Any other fields of the ACP certificate are to be populated as required by [RFC5280]: As long as they are compliant with [RFC5280], any other field of an ACP certificate can be set as desired by the operator of the ACP domain through appropriate ACP registrar/ACP CA procedures. For example, other fields may be required for other purposes that the ACP certificate is intended to be used for (such as elements of a `SubjectName`).

For further certificate details, ACP certificates may follow the recommendations from [CABFORUM].

For diagnostic and other operational purposes, it is beneficial to copy the device identifying fields of the node's `IDevID` certificate into the ACP certificate, such as the [X.520], section 6.2.9 "serialNumber" attribute in the subject field distinguished name encoding. Note that this is not the certificate serial number. See

also [I-D.ietf-anima-bootstrapping-keyinfra] section 2.3.1. This can be done for example if it would be acceptable for the device's "serialNumber" to be signaled via the Link Layer Discovery Protocol (LLDP, [LLDP]) because like LLDP signaled information, the ACP certificate information can be retrieved by neighboring nodes without further authentication and be used either for beneficial diagnostics or for malicious attacks. Retrieval of the ACP certificate is possible via a (failing) attempt to set up an ACP secure channel, and the "serialNumber" usually contains device type information that may help to faster determine working exploits/attacks against the device.

Note that there is no intention to constrain authorization within the ACP or autonomic networks using the ACP to just the ACP domain membership check as defined in this document. It can be extended or modified with additional requirements. Such future authorizations can use and require additional elements in certificates or policies or even additional certificates. See the additional check against the id-kp-cmcRA [RFC6402] extended key usage attribute (Section 6.2.5) and for possible future extensions, see Appendix A.9.5.

#### 6.2.2. ACP Certificate AcpNodeName

```

acp-node-name = local-part "@" acp-domain-name
local-part = [ acp-address ] [ "+" rsub extensions ]
acp-address = 32HEXDIG | "0" ; HEXDIG as of RFC5234 section B.1
rsub = [ <subdomain> ] ; <subdomain> as of RFC1034, section 3.5
acp-domain-name = ; <domain> ; as of RFC 1034, section 3.5
extensions = *( "+" extension )
extension = 1*etext ; future standard definition.
etext      = ALPHA / DIGIT / ; Printable US-ASCII
            "!" / "#" / "$" / "%" / "&" / "'" /
            "*" / "-" / "/" / "=" / "?" / "^" /
            "_" / "`" / "{" / "|" / "}" / "~"

routing-subdomain = [ rsub "." ] acp-domain-name

```

#### Example:

```

given an ACP address   of fd89:b714:f3db:0:200:0:6400:0000
and an ACP domain-name of acp.example.com
and an rsub extension of area51.research

```

then this results in:

```

acp-node-name          = fd89b714f3db00000200000064000000
                       +area51.research@acp.example.com
acp-domain-name        = acp.example.com
routing-subdomain      = area51.research.acp.example.com

```

Figure 2: ACP Node Name ABNF

acp-node-name in above Figure 2 is the ABNF ([RFC5234]) definition of the ACP Node Name. An ACP certificate MUST carry this information. It MUST be encoded as a subjectAltName / otherName / AcpNodeName as described in Section 6.2.2.1.

Nodes complying with this specification MUST be able to receive their ACP address through the domain certificate, in which case their own ACP certificate MUST have a 32HEXDIG acp-address field. Acp-address is case insensitive because ABNF HEXDIG is. It is recommended to encode acp-address with lower case letters. Nodes complying with this specification MUST also be able to authenticate nodes as ACP domain members or ACP secure channel peers when they have a 0-value acp-address field and as ACP domain members (but not as ACP secure channel peers) when the acp-address field is omitted from their AcpNodeName. See Section 6.2.3.

acp-domain-name is used to indicate the ACP Domain across which ACP nodes authenticate and authorize each other, for example to build ACP secure channels to each other, see Section 6.2.3. acp-domain-name SHOULD be the FQDN of an Internet domain owned by the network administration of the ACP and ideally reserved to only be used for the ACP. In this specification it serves to be a name for the ACP that ideally is globally unique. When acp-domain-name is a globally unique name, collision of ACP addresses across different ACP domains can only happen due to ULA hash collisions (see Section 6.11.2). Using different acp-domain-names, operators can distinguish multiple ACP even when using the same TA.

To keep the encoding simple, there is no consideration for internationalized acp-domain-names. The acp-node-name is not intended for end user consumption. There is no protection against an operator to pick any domain name for an ACP whether or not the operator can claim to own the domain name. Instead, the domain name only serves as a hash seed for the ULA and for diagnostics to the operator. Therefore, any operator owning only an internationalized domain name should be able to pick an equivalently unique 7-bit ASCII acp-domain-name string representing the internationalized domain name.

"routing-subdomain" is a string that can be constructed from the acp-node-name, and it is used in the hash-creation of the ULA (see below). The presence of the "rsub" component allows a single ACP domain to employ multiple /48 ULA prefixes. See Appendix A.6 for example use-cases.

The optional "extensions" field is used for future standardized extensions to this specification. It MUST be ignored if present and not understood.

The following points explain and justify the encoding choices described:

1. Formatting notes:
  - 1.1 "rsub" needs to be in the "local-part": If the format just had routing-subdomain as the domain part of the acp-node-name, rsub and acp-domain-name could not be separated from each other to determine in the ACP domain membership check which part is the acp-domain-name and which is solely for creating a different ULA prefix.
  - 1.2 If both "acp-address" and "rsub" are omitted from AcpNodeName, the "local-part" will have the format "++extension(s)". The two plus characters are necessary so the node can unambiguously parse that both "acp-address" and "rsub" are omitted.
2. The encoding of the ACP domain name and ACP address as described in this section is used for the following reasons:
  - 2.1 The acp-node-name is the identifier of a node's ACP. It includes the necessary components to identify a node's ACP both from within the ACP as well as from the outside of the ACP.
  - 2.2 For manual and/or automated diagnostics and backend management of devices and ACPs, it is necessary to have an easily human readable and software parsed standard, single string representation of the information in the acp-node-name. For example, inventory or other backend systems can always identify an entity by one unique string field but not by a combination of multiple fields, which would be necessary if there was no single string representation.
  - 2.3 If the encoding was not that of such a string, it would be necessary to define a second standard encoding to provide this format (standard string encoding) for operator consumption.
  - 2.4 Addresses of the form <local>@<domain> have become the preferred format for identifiers of entities in many systems, including the majority of user identification in web or mobile applications such as multi-domain single-sign-on systems.
3. Compatibilities:
  - 3.1 It should be possible to use the ACP certificate as an LDevID certificate on the system for other uses beside the ACP. Therefore, the information element required for the ACP should be encoded so that it minimizes the possibility of creating incompatibilities with such other uses. The

attributes of the subject field for example are often used in non-ACP applications and should therefore not be occupied by new ACP values.

- 3.2 The element should not require additional ASN.1 en/decoding, because libraries to access certificate information especially for embedded devices may not support extended ASN.1 decoding beyond predefined, mandatory fields. subjectAltName / otherName is already used with a single string parameter for several otherNames (see [RFC3920], [RFC7585], [RFC4985], [RFC8398]).
- 3.3 The element required for the ACP should minimize the risk of being misinterpreted by other uses of the LDevID certificate. It also must not be misinterpreted to actually be an email address, hence the use of the otherName / rfc822Name option in the certificate would be inappropriate.

See section 4.2.1.6 of [RFC5280] for details on the subjectAltName field.

#### 6.2.2.1. AcpNodeName ASN.1 Module

The following ASN.1 module normatively specifies the AcpNodeName structure. This specification uses the ASN.1 definitions from [RFC5912] with the 2002 ASN.1 notation used in that document. [RFC5912] updates normative documents using older ASN.1 notation.

```

ANIMA-ACP-2020
  { iso(1) identified-organization(3) dod(6)
    internet(1) security(5) mechanisms(5) pkix(7) id-mod(0)
    id-mod-anima-acpnode-name-2020(IANA1) }

DEFINITIONS IMPLICIT TAGS ::=
BEGIN

IMPORTS
  OTHER-NAME
  FROM PKIX1Implicit-2009
    { iso(1) identified-organization(3) dod(6) internet(1)
      security(5) mechanisms(5) pkix(7) id-mod(0)
      id-mod-pkix1-implicit-02(59) }

  id-pkix
  FROM PKIX1Explicit-2009
    { iso(1) identified-organization(3) dod(6) internet(1)
      security(5) mechanisms(5) pkix(7) id-mod(0)
      id-mod-pkix1-explicit-02(51) } ;

  id-on OBJECT IDENTIFIER ::= { id-pkix 8 }

  AcpNodeNameOtherNames OTHER-NAME ::= { on-AcpNodeName, ... }

  on-AcpNodeName OTHER-NAME ::= {
    AcpNodeName IDENTIFIED BY id-on-AcpNodeName
  }

  id-on-AcpNodeName OBJECT IDENTIFIER ::= { id-on IANA2 }

  AcpNodeName ::= IA5String (SIZE (1..MAX))
    -- AcpNodeName as specified in this document carries the
    -- acp-node-name as specified in the ABNF in Section 6.1.2

END

```

Figure 3

### 6.2.3. ACP domain membership check

The following points constitute the ACP domain membership check of a candidate peer via its certificate:

- 1: The peer has proved ownership of the private key associated with the certificate's public key. This check is performed by the security association protocol used, for example [RFC7296], section 2.15.

- 2: The peer's certificate passes certificate path validation as defined in [RFC5280], section 6 against one of the TA associated with the ACP node's ACP certificate (see Section 6.2.4 below). This includes verification of the validity (lifetime) of the certificates in the path.
- 3: If the peer's certificate indicates a Certificate Revocation List (CRL) Distribution Point (CRLDP) ([RFC5280], section 4.2.1.13) or Online Certificate Status Protocol (OCSP) responder ([RFC5280], section 4.2.2.1), then the peer's certificate MUST be valid according to those mechanisms when they are available: An OCSP check for the peer's certificate across the ACP must succeed or the peer certificate must not be listed in the CRL retrieved from the CRLDP. These mechanisms are not available when the ACP node has no ACP or non-ACP connectivity to retrieve a current CRL or access an OCSP responder and the security association protocol itself has also no way to communicate CRL or OCSP check. Retries to learn revocation via OCSP/CRL SHOULD be made using the same backoff as described in Section 6.7. If and when the ACP node then learns that an ACP peer's certificate is invalid for which rule 3 had to be skipped during ACP secure channel establishment, then the ACP secure channel to that peer MUST be closed even if this peer is the only connectivity to access CRL/OCSP. This applies (of course) to all ACP secure channels to this peer if there are multiple. The ACP secure channel connection MUST be retried periodically to support the case that the neighbor acquires a new, valid certificate.
- 4: The peer's certificate has a syntactically valid acp-node-name field and the acp-domain-name in that peer's acp-node-name is the same as in this ACP node's certificate (lowercase normalized).

When checking a candidate peer's certificate for the purpose of establishing an ACP secure channel, one additional check is performed:

- 5: The acp-address field of the candidate peer certificate's AcpNodeName is not omitted but either 32HEXDIG or 0, according to Figure 2.

Technically, ACP secure channels can only be built with nodes that have an acp-address. Rule 5 ensures that this is taken into account during ACP domain membership check.

Nodes with an omitted acp-address field can only use their ACP domain certificate for non-ACP-secure channel authentication purposes. This includes for example NMS type nodes permitted to communicate into the ACP via ACP connect (Section 8.1)

The special value 0 in an ACP certificates acp-address field is used for nodes that can and should determine their ACP address through other mechanisms than learning it through the acp-address field in their ACP certificate. These ACP nodes are permitted to establish ACP secure channels. Mechanisms for those nodes to determine their ACP address are outside the scope of this specification, but this option is defined here so that any ACP nodes can build ACP secure channels to them according to Rule 5.

The optional rsub field of the AcpNodeName is not relevant to the ACP domain membership check because it only serves to structure routing and addressing within an ACP but not to segment mutual authentication/authorization (hence the name "routing subdomain").

In summary:

- \* Steps 1...4 constitute standard certificate validity verification and private key authentication as defined by [RFC5280] and security association protocols (such as Internet Key Exchange Protocol version 2 IKEv2 [RFC7296] when leveraging certificates.
- \* Steps 1...4 do not include verification of any pre-existing form of non-public-key-only based identity elements of a certificate such as a web servers domain name prefix often encoded in certificate common name. Step 5 is an equivalent step for the AcpNodeName.
- \* Step 4 constitutes standard CRL/OCSP checks refined for the case of missing connectivity and limited functionality security association protocols.
- \* Steps 1...4 authorize to build any secure connection between members of the same ACP domain except for ACP secure channels.
- \* Step 5 is the additional verification of the presence of an ACP address as necessary for ACP secure channels.
- \* Steps 1...5 therefore authorize to build an ACP secure channel.

For brevity, the remainder of this document refers to this process only as authentication instead of as authentication and authorization.

[RFC-Editor: Please remove the following paragraph].

Note that the ACP domain membership check does not verify the network layer address of the security association. See [ACPDRAFT], Appendix B.2 for explanations.

#### 6.2.3.1. Realtime clock and Time Validation

An ACP node with a realtime clock in which it has confidence, MUST check the time stamps when performing ACP domain membership check such as the certificate validity period in step 1. and the respective times in step 4 for revocation information (e.g., signingTimes in CMS signatures).

An ACP node without such a realtime clock MAY ignore those time stamp validation steps if it does not know the current time. Such an ACP node SHOULD obtain the current time in a secured fashion, such as via a Network Time Protocol signaled through the ACP. It then ignores time stamp validation only until the current time is known. In the absence of implementing a secured mechanism, such an ACP node MAY use a current time learned in an insecure fashion in the ACP domain membership check.

Current time MAY for example be learned unsecured via NTP ([RFC5905]) over the same link-local IPv6 addresses used for the ACP from neighboring ACP nodes. ACP nodes that do provide NTP insecure over their link-local addresses SHOULD primarily run NTP across the ACP and provide NTP time across the ACP only when they have a trusted time source. Details for such NTP procedures are beyond the scope of this specification.

Beside ACP domain membership check, the ACP itself has no dependency against knowledge of the current time, but protocols and services using the ACP will likely have the need to know the current time. For example, event logging.

#### 6.2.4. Trust Anchors (TA)

ACP nodes need TA information according to [RFC5280], section 6.1.1 (d), typically in the form of one or more certificate of the TA to perform certificate path validation as required by Section 6.2.3, rule 2. TA information MUST be provisioned to an ACP node (together with its ACP domain certificate) by an ACP Registrar during initial enrollment of a candidate ACP node. ACP nodes MUST also support renewal of TA information via EST as described below in Section 6.2.5.

The required information about a TA can consist of not only a single, but multiple certificates as required for dealing with CA certificate renewals as explained in Section 4.4 of CMP ([RFC4210]).

A certificate path is a chain of certificates starting at the ACP certificate (leaf/end-entity) followed by zero or more intermediate CA certificates and ending with the TA information, which are

typically one or two the self-signed certificates of the TA. The CA that signs the ACP certificate is called the assigning CA. If there are no intermediate CA, then the assigning CA is the TA. Certificate path validation authenticates that the ACP certificate is permitted by a TA associated with the ACP, directly or indirectly via one or more intermediate CA.

Note that different ACP nodes may have different intermediate CA in their certificate path and even different TA. The set of TA for an ACP domain must be consistent across all ACP members so that any ACP node can authenticate any other ACP node. The protocols through which ACP domain membership check rules 1-3 are performed need to support the exchange not only of the ACP nodes certificates, but also exchange of the intermedia TA.

ACP nodes MUST support for the ACP domain membership check the certificate path validation with 0 or 1 intermediate CA. They SHOULD support 2 intermediate CA and two TA (to permit migration to from one TA to another TA).

Certificates for an ACP MUST only be given to nodes that are allowed to be members of that ACP. When the signing CA relies on an ACP Registrar, the CA MUST only sign certificates with acp-node-name through trusted ACP Registrars. In this setup, any existing CA, unaware of the formatting of acp-node-name, can be used.

These requirements can be achieved by using a TA private to the owner of the ACP domain or potentially through appropriate contractual agreements between the involved parties (Registrar and CA). Using public CA is out of scope of this document. [RFC-Editor: please remove the following sentence]. See [ACPDRAFT], Appendix B.3 for further considerations.

A single owner can operate multiple independent ACP domains from the same set of TA. Registrars must then know which ACP a node needs to be enrolled into.

#### 6.2.5. Certificate and Trust Anchor Maintenance

ACP nodes MUST support renewal of their Certificate and TA information via EST and MAY support other mechanisms. See Section 6.1 for TLS requirements. An ACP network MUST have at least one ACP node supporting EST server functionality across the ACP so that EST renewal is useable.

ACP nodes SHOULD be able to remember the IPv6 locator parameters of the O\_IPv6\_LOCATOR in GRASP of the EST server from which they last renewed their ACP certificate. They SHOULD provide the ability for

these EST server parameters to also be set by the ACP Registrar (see Section 6.11.7) that initially enrolled the ACP device with its ACP certificate. When BRSKI (see [I-D.ietf-anima-bootstrapping-keyinfra]) is used, the IPv6 locator of the BRSKI registrar from the BRSKI TLS connection SHOULD be remembered and used for the next renewal via EST if that registrar also announces itself as an EST server via GRASP (see next section) on its ACP address.

The EST server MUST present a certificate that is passing ACP domain membership check in its TLS connection setup (Section 6.2.3, rules 1...4, not rule 5 as this is not for an ACP secure channel setup). The EST server certificate MUST also contain the id-kp-cmcRA [RFC6402] extended key usage attribute and the EST client MUST check its presence.

The additional check against the id-kp-cmcRA extended key usage extension field ensures that clients do not fall prey to an illicit EST server. While such illicit EST servers should not be able to support certificate signing requests (as they are not able to elicit a signing response from a valid CA), such an illicit EST server would be able to provide faked CA certificates to EST clients that need to renew their CA certificates when they expire.

Note that EST servers supporting multiple ACP domains will need to have for each of these ACP domains a separate certificate and respond on a different transport address (IPv6 address and/or TCP port), but this is easily automated on the EST server as long as the CA does not restrict registrars to request certificates with the id-kp-cmcRA extended usage extension for themselves.

#### 6.2.5.1. GRASP objective for EST server

ACP nodes that are EST servers MUST announce their service via GRASP in the ACP through M\_FLOOD messages. See [I-D.ietf-anima-grasp], section 2.8.11 for the definition of this message type:

Example:

```
[M_FLOOD, 12340815, h'fd89b714f3db0000200000064000001', 210000,
  [{"SRV.est", 4, 255 }],
  [O_IPv6_LOCATOR,
    h'fd89b714f3db0000200000064000001', IPPROTO_TCP, 443]]
]
```

Figure 4: GRASP SRV.est example

The formal definition of the objective in Concise data definition language (CDDL) (see [RFC8610]) is as follows:

```
flood-message = [M_FLOOD, session-id, initiator, ttl,
                 +[objective, (locator-option / [])]]
                 ; see example above and explanation
                 ; below for initiator and ttl

objective = ["SRV.est", objective-flags, loop-count,
            objective-value]

objective-flags = sync-only ; as in GRASP spec
sync-only       = 4        ; M_FLOOD only requires synchronization
loop-count      = 255      ; recommended as there is no mechanism
                   ; to discover network diameter.
objective-value = any      ; reserved for future extensions
```

Figure 5: GRASP SRV.est definition

The objective name "SRV.est" indicates that the objective is an [RFC7030] compliant EST server because "est" is an [RFC6335] registered service name for [RFC7030]. Objective-value MUST be ignored if present. Backward compatible extensions to [RFC7030] MAY be indicated through objective-value. Non [RFC7030] compatible certificate renewal options MUST use a different objective-name. Non-recognized objective-values (or parts thereof if it is a structure partially understood) MUST be ignored.

The M\_FLOOD message MUST be sent periodically. The default SHOULD be 60 seconds; the value SHOULD be operator configurable but SHOULD be not smaller than 60 seconds. The frequency of sending MUST be such that the aggregate amount of periodic M\_FLOODs from all flooding sources cause only negligible traffic across the ACP. The time-to-live (ttl) parameter SHOULD be 3.5 times the period so that up to three consecutive messages can be dropped before considering an announcement expired. In the example above, the ttl is 210000 msec, 3.5 times 60 seconds. When a service announcer using these parameters unexpectedly dies immediately after sending the M\_FLOOD, receivers would consider it expired 210 seconds later. When a receiver tries to connect to this dead service before this timeout, it will experience a failing connection and use that as an indication that the service instance is dead and select another instance of the same service instead (from another GRASP announcement).

The "SRV.est" objective(s) SHOULD only be announced when the ACP node knows that it can successfully communicate with a CA to perform the EST renewal/rekeying operations for the ACP domain. See also Section 11.

#### 6.2.5.2. Renewal

When performing renewal, the node SHOULD attempt to connect to the remembered EST server. If that fails, it SHOULD attempt to connect to an EST server learned via GRASP. The server with which certificate renewal succeeds SHOULD be remembered for the next renewal.

Remembering the last renewal server and preferring it provides stickiness which can help diagnostics. It also provides some protection against off-path compromised ACP members announcing bogus information into GRASP.

Renewal of certificates SHOULD start after less than 50% of the domain certificate lifetime so that network operations has ample time to investigate and resolve any problems that causes a node to not renew its domain certificate in time - and to allow prolonged periods of running parts of a network disconnected from any CA.

#### 6.2.5.3. Certificate Revocation Lists (CRLs)

The ACP node SHOULD support revocation through CRL(s) via HTTP from one or more CRL Distribution Points (CRLDP). The CRLDP(s) MUST be indicated in the Domain Certificate when used. If the CRLDP URL uses an IPv6 address (ULA address when using the addressing rules specified in this document), the ACP node will connect to the CRLDP via the ACP. If the CRLDP uses a domain name, the ACP node will connect to the CRLDP via the Data-Plane.

It is common to use domain names for CRLDP(s), but there is no requirement for the ACP to support DNS. Any DNS lookup in the Data-Plane is not only a possible security issue, but it would also not indicate whether the resolved address is meant to be reachable across the ACP. Therefore, the use of an IPv6 address versus the use of a DNS name doubles as an indicator whether or not to reach the CRLDP via the ACP.

A CRLDP can be reachable across the ACP either by running it on a node with ACP or by connecting its node via an ACP connect interface (see Section 8.1).

When using a private PKI for ACP certificates, the CRL may be need-to-know, for example to prohibit insight into the operational practices of the domain by tracking the growth of the CRL. In this case, HTTPS may be chosen to provide confidentiality, especially when making the CRL available via the Data-Plane. Authentication and authorization SHOULD use ACP certificates and ACP domain membership check. The CRLDP MAY omit the CRL verification during authentication of the peer to permit retrieval of the CRL by an ACP node with revoked ACP certificate. This can allow for that (ex) ACP node to quickly discover its ACP certificate revocation. This may violate the desired need-to-know requirement though. ACP nodes MAY support CRLDP operations via HTTPS.

#### 6.2.5.4. Lifetimes

Certificate lifetime may be set to shorter lifetimes than customary (1 year) because certificate renewal is fully automated via ACP and EST. The primary limiting factor for shorter certificate lifetimes is load on the EST server(s) and CA. It is therefore recommended that ACP certificates are managed via a CA chain where the assigning CA has enough performance to manage short lived certificates. See also Section 9.2.4 for discussion about an example setup achieving this. See also [I-D.ietf-acme-star].

When certificate lifetimes are sufficiently short, such as few hours, certificate revocation may not be necessary, allowing to simplify the overall certificate maintenance infrastructure.

See Appendix A.2 for further optimizations of certificate maintenance when BRSKI can be used ("Bootstrapping Remote Secure Key Infrastructures", see [I-D.ietf-anima-bootstrapping-keyinfra]).

#### 6.2.5.5. Re-enrollment

An ACP node may determine that its ACP certificate has expired, for example because the ACP node was powered down or disconnected longer than its certificate lifetime. In this case, the ACP node SHOULD convert to a role of a re-enrolling candidate ACP node.

In this role, the node does maintain the TA and certificate chain associated with its ACP certificate exclusively for the purpose of re-enrollment, and attempts (or waits) to get re-enrolled with a new ACP certificate. The details depend on the mechanisms/protocols used by the ACP Registrars.

Please refer to Section 6.11.7 and [I-D.ietf-anima-bootstrapping-keyinfra] for explanations about ACP Registrars and vouchers as used in the following text. When ACP is intended to be used without BRSKI, the details about BRSKI and vouchers in the following text can be skipped.

When BRSKI is used (i.e.: on ACP nodes that are ANI nodes), the re-enrolling candidate ACP node would attempt to enroll like a candidate ACP node (BRSKI pledge), but instead of using the ACP nodes IDevID certificate, it SHOULD first attempt to use its ACP domain certificate in the BRSKI TLS authentication. The BRSKI registrar MAY honor this certificate beyond its expiration date purely for the purpose of re-enrollment. Using the ACP node's domain certificate allows the BRSKI registrar to learn that node's acp-node-name, so that the BRSKI registrar can re-assign the same ACP address information to the ACP node in the new ACP certificate.

If the BRSKI registrar denies the use of the old ACP certificate, the re-enrolling candidate ACP node MUST re-attempt re-enrollment using its IDevID certificate as defined in BRSKI during the TLS connection setup.

Both when the BRSKI connection is attempted with the old ACP certificate or the IDevID certificate, the re-enrolling candidate ACP node SHOULD authenticate the BRSKI registrar during TLS connection setup based on its existing TA certificate chain information associated with its old ACP certificate. The re-enrolling candidate ACP node SHOULD only fall back to requesting a voucher from the BRSKI registrar when this authentication fails during TLS connection setup. As a countermeasure against attacks that attempt to force the ACP node to forget its prior (expired) certificate and TA, the ACP node should alternate between attempting to re-enroll using its old keying material and attempting to re-enroll with its IDevID and requesting a voucher.

When other mechanisms than BRSKI are used for ACP certificate enrollment, the principles of the re-enrolling candidate ACP node are the same. The re-enrolling candidate ACP node attempts to authenticate any ACP Registrar peers during re-enrollment protocol/mechanisms via its existing certificate chain/TA information and provides its existing ACP certificate and other identification (such as the IDevID certificate) as necessary to the registrar.

Maintaining existing TA information is especially important when enrollment mechanisms are used that unlike BRSKI do not leverage a mechanism (such as the voucher in BRSKI) to authenticate the ACP registrar and where therefore the injection of certificate failures could otherwise make the ACP node easily attackable remotely by

returning the ACP node to a "duckling" state in which it accepts to be enrolled by any network it connects to. The (expired) ACP certificate and ACP TA SHOULD therefore be maintained and attempted to be used as one possible credential for re-enrollment until new keying material is acquired.

When using BRSKI or other protocol/mechanisms supporting vouchers, maintaining existing TA information allows for re-enrollment of expired ACP certificates to be more lightweight, especially in environments where repeated acquisition of vouchers during the lifetime of ACP nodes may be operationally expensive or otherwise undesirable.

#### 6.2.5.6. Failing Certificates

An ACP certificate is called failing in this document, if/when the ACP node to which the certificate was issued can determine that it was revoked (or explicitly not renewed), or in the absence of such explicit local diagnostics, when the ACP node fails to connect to other ACP nodes in the same ACP domain using its ACP certificate. For connection failures to determine the ACP certificate as the culprit, the peer should pass the domain membership check (Section 6.2.3) and other reasons for the connection failure can be excluded because of the connection error diagnostics.

This type of failure can happen during setup/refresh of a secure ACP channel connections or any other use of the ACP certificate, such as for the TLS connection to an EST server for the renewal of the ACP domain certificate.

Example reasons for failing certificates that the ACP node can only discover through connection failure are that the domain certificate or any of its signing certificates could have been revoked or may have expired, but the ACP node cannot self-diagnose this condition directly. Revocation information or clock synchronization may only be available across the ACP, but the ACP node cannot build ACP secure channels because ACP peers reject the ACP node's domain certificate.

ACP nodes SHOULD support the option to determine whether its ACP certificate is failing, and when it does, put itself into the role of a re-enrolling candidate ACP node as explained above (Section 6.2.5.5).

### 6.3. ACP Adjacency Table

To know to which nodes to establish an ACP channel, every ACP node maintains an adjacency table. The adjacency table contains information about adjacent ACP nodes, at a minimum: Node-ID (identifier of the node inside the ACP, see Section 6.11.3 and Section 6.11.5), interface on which neighbor was discovered (by GRASP as explained below), link-local IPv6 address of neighbor on that interface, certificate (including acp-node-name). An ACP node MUST maintain this adjacency table. This table is used to determine to which neighbor an ACP connection is established.

Where the next ACP node is not directly adjacent (i.e., not on a link connected to this node), the information in the adjacency table can be supplemented by configuration. For example, the Node-ID and IP address could be configured. See Section 8.2.

The adjacency table MAY contain information about the validity and trust of the adjacent ACP node's certificate. However, subsequent steps MUST always start with the ACP domain membership check against the peer (see Section 6.2.3).

The adjacency table contains information about adjacent ACP nodes in general, independently of their domain and trust status. The next step determines to which of those ACP nodes an ACP connection should be established.

### 6.4. Neighbor Discovery with DULL GRASP

[RFC-Editor: GRASP draft is in RFC editor queue, waiting for dependencies, including ACP. Please ensure that references to I-D.ietf-anima-grasp that include section number references (throughout this document) will be updated in case any last-minute changes in GRASP would make those section references change.

Discovery Unsolicited Link-Local (DULL) GRASP is a limited subset of GRASP intended to operate across an insecure link-local scope. See section 2.5.2 of [I-D.ietf-anima-grasp] for its formal definition. The ACP uses one instance of DULL GRASP for every L2 interface of the ACP node to discover link level adjacent candidate ACP neighbors. Unless modified by policy as noted earlier (Section 5 bullet point 2.), native interfaces (e.g., physical interfaces on physical nodes) SHOULD be initialized automatically to a state in which ACP discovery can be performed and any native interfaces with ACP neighbors can then be brought into the ACP even if the interface is otherwise not configured. Reception of packets on such otherwise not configured interfaces MUST be limited so that at first only IPv6 Stateless Address Auto Configuration (SLAAC - [RFC4862]) and DULL GRASP work

and then only the following ACP secure channel setup packets - but not any other unnecessary traffic (e.g., no other link-local IPv6 transport stack responders for example).

Note that the use of the IPv6 link-local multicast address (ALL\_GRASP\_NEIGHBORS) implies the need to use Multicast Listener Discovery Version 2 (MLDv2, see [RFC3810]) to announce the desire to receive packets for that address. Otherwise DULL GRASP could fail to operate correctly in the presence of MLD snooping ([RFC4541]) switches that are not ACP supporting/enabled - because those switches would stop forwarding DULL GRASP packets. Switches not supporting MLD snooping simply need to operate as pure L2 bridges for IPv6 multicast packets for DULL GRASP to work.

ACP discovery SHOULD NOT be enabled by default on non-native interfaces. In particular, ACP discovery MUST NOT run inside the ACP across ACP virtual interfaces. See Section 9.3 for further, non-normative suggestions on how to enable/disable ACP at node and interface level. See Section 8.2.2 for more details about tunnels (typical non-native interfaces). See Section 7 for how ACP should be extended on devices operating (also) as L2 bridges.

Note: If an ACP node also implements BRSKI to enroll its ACP certificate (see Appendix A.2 for a summary), then the above considerations also apply to GRASP discovery for BRSKI. Each DULL instance of GRASP set up for ACP is then also used for the discovery of a bootstrap proxy via BRSKI when the node does not have a domain certificate. Discovery of ACP neighbors happens only when the node does have the certificate. The node therefore never needs to discover both a bootstrap proxy and ACP neighbor at the same time.

An ACP node announces itself to potential ACP peers by use of the "AN\_ACP" objective. This is a synchronization objective intended to be flooded on a single link using the GRASP Flood Synchronization (M\_FLOOD) message. In accordance with the design of the Flood message, a locator consisting of a specific link-local IP address, IP protocol number and port number will be distributed with the flooded objective. An example of the message is informally:

```
[M_FLOOD, 12340815, h'fe80000000000000c0011001feef0000', 210000,
  [{"AN_ACP", 4, 1, "IKEv2" }],
  [O_IPv6_LOCATOR,
    h'fe80000000000000c0011001feef0000', IPPROTO_UDP, 15000]]
  [{"AN_ACP", 4, 1, "DTLS" }],
  [O_IPv6_LOCATOR,
    h'fe80000000000000c0011001feef0000', IPPROTO_UDP, 17000]]
]
```

Figure 6: GRASP AN\_ACP example

The formal CDDL definition is:

```
flood-message = [M_FLOOD, session-id, initiator, ttl,
                 +[objective, (locator-option / [])]]

objective = ["AN_ACP", objective-flags, loop-count,
            objective-value]

objective-flags = sync-only ; as in the GRASP specification
sync-only = 4 ; M_FLOOD only requires synchronization
loop-count = 1 ; limit to link-local operation

objective-value = method-name / [ method, *extension ]
method = method-name / [ method-name, *method-param ]
method-name = "IKEv2" / "DTLS" / id
extension = any
method-param = any
id = text .regexp "[A-Za-z@_$(-)]*[A-Za-z0-9@_$(-)]*"

```

Figure 7: GRASP AN\_ACP definition

The objective-flags field is set to indicate synchronization.

The loop-count is fixed at 1 since this is a link-local operation.

In the above example the RECOMMENDED period of sending of the objective is 60 seconds. The indicated ttl of 210000 msec means that the objective would be cached by ACP nodes even when two out of three messages are dropped in transit.

The session-id is a random number used for loop prevention (distinguishing a message from a prior instance of the same message). In DULL this field is irrelevant but has to be set according to the GRASP specification.

The originator MUST be the IPv6 link local address of the originating ACP node on the sending interface.

The method-name in the 'objective-value' parameter is a string indicating the protocol available at the specified or implied locator. It is a protocol supported by the node to negotiate a secure channel. IKEv2 as shown above is the protocol used to negotiate an IPsec secure channel.

Method-params allows to carry method specific parameters. This specification does not define any method-param(s) for "IKEv2" or "DTLS". Method-params for these two methods that are not understood by an ACP node MUST be ignored by it.

extension(s) allows to define method independent parameters. This specification does not define any extensions. Extensions not understood by an ACP node MUST be ignored by it.

The locator-option is optional and only required when the secure channel protocol is not offered at a well-defined port number, or if there is no well-defined port number.

IKEv2 is the actual protocol used to negotiate an Internet Protocol security architecture (IPsec) connection. GRASP therefore indicates "IKEv2" and not "IPsec". If "IPsec" was used, this too could mean use of the obsolete older version IKE (v1) ([RFC2409]). IKEv2 has an IANA assigned port number 500, but in the above example, the candidate ACP neighbor is offering ACP secure channel negotiation via IKEv2 on port 15000 (purely to show through the example that GRASP allows to indicate the port number and it does not have to be the IANA assigned one).

There is no default UDP port for DTLS, it is always locally assigned by the node. For further details about the "DTLS" secure channel protocol, see Section 6.8.4.

If a locator is included, it MUST be an O\_IPv6\_LOCATOR, and the IPv6 address MUST be the same as the initiator address (these are DULL requirements to minimize third party DoS attacks).

The secure channel methods defined in this document use the objective-values of "IKEv2" and "DTLS". There is no distinction between IKEv2 native and GRE-IKEv2 because this is purely negotiated via IKEv2.

A node that supports more than one secure channel protocol method needs to flood multiple versions of the "AN\_ACP" objective so that each method can be accompanied by its own locator-option. This can use a single GRASP M\_FLOOD message as shown in Figure 6.

The use of DULL GRASP primarily serves to discover the link-local IPv6 address of candidate ACP peers on subnets. The signaling of the supported secure channel option is primarily for diagnostic purposes, but it is also necessary for discovery when the protocol has no well-known transport address, such as in the case of DTLS. [RFC-Editor: Please remove the following sentence]. See [ACPDRAFT], Appendix B.4.

Note that a node serving both as an ACP node and BRSKI Join Proxy may choose to distribute the "AN\_ACP" objective and the respective BRSKI in the same M\_FLOOD message, since GRASP allows multiple objectives in one message. This may be impractical though if ACP and BRSKI operations are implemented via separate software modules / ASAs.

The result of the discovery is the IPv6 link-local address of the neighbor as well as its supported secure channel protocols (and non-standard port they are running on). It is stored in the ACP Adjacency Table (see Section 6.3), which then drives the further building of the ACP to that neighbor.

Note that the DULL GRASP objective described intentionally does not include the ACP node's ACP certificate even though this would be useful for diagnostics and to simplify the security exchange in ACP secure channel security association protocols (see Section 6.8). The reason is that DULL GRASP messages are periodically multicasted across IPv6 subnets and full certificates could easily lead to fragmented IPv6 DULL GRASP multicast packets due to the size of a certificate. This would be highly undesirable.

#### 6.5. Candidate ACP Neighbor Selection

An ACP node determines to which other ACP nodes in the adjacency table it should attempt to build an ACP connection. This is based on the information in the ACP Adjacency table.

The ACP is established exclusively between nodes in the same domain. This includes all routing subdomains. Appendix A.6 explains how ACP connections across multiple routing subdomains are special.

The result of the candidate ACP neighbor selection process is a list of adjacent or configured autonomic neighbors to which an ACP channel should be established. The next step begins that channel establishment.

#### 6.6. Channel Selection

To avoid attacks, initial discovery of candidate ACP peers cannot include any non-protected negotiation. To avoid re-inventing and validating security association mechanisms, the next step after discovering the address of a candidate neighbor can only be to try first to establish a security association with that neighbor using a well-known security association method.

From the use-cases it seems clear that not all type of ACP nodes can or need to connect directly to each other or are able to support or prefer all possible mechanisms. For example, code space limited IoT

devices may only support DTLS because that code exists already on them for end-to-end security, but low-end in-ceiling L2 switches may only want to support Media Access Control Security (MacSec, see 802.1AE ([MACSEC]) because that is also supported in their chips. Only a flexible gateway device may need to support both of these mechanisms and potentially more. Note that MacSec is not required by any profiles of the ACP in this specification. Instead, MacSec is mentioned as a likely next interesting secure channel protocol. Note also that the security model allows and requires for any-to-any authentication and authorization between all ACP nodes because there is also end-to-end and not only hop-by-hop authentication for secure channels.

To support extensible secure channel protocol selection without a single common mandatory to implement (MTI) protocol, ACP nodes MUST try all the ACP secure channel protocols it supports and that are feasible because the candidate ACP neighbor also announced them via its AN\_ACP GRASP parameters (these are called the "feasible" ACP secure channel protocols).

To ensure that the selection of the secure channel protocols always succeeds in a predictable fashion without blocking, the following rules apply:

- \* An ACP node may choose to attempt to initiate the different feasible ACP secure channel protocols it supports according to its local policies sequentially or in parallel, but it MUST support acting as a responder to all of them in parallel.
- \* Once the first ACP secure channel protocol connection to a specific peer IPv6 address passes peer authentication, the two peers know each other's certificate because those ACP certificates are used by all secure channel protocols for mutual authentication. The peer with the higher Node-ID in the AcpNodeName of its ACP certificate takes on the role of the Decider towards the peer. The other peer takes on the role of the Follower. The Decider selects which secure channel protocol to ultimately use.
- \* The Follower becomes passive: it does not attempt to further initiate ACP secure channel protocol connections with the Decider and does not consider it to be an error when the Decider closes secure channels. The Decider becomes the active party, continues to attempt setting up secure channel protocols with the Follower. This process terminates when the Decider arrives at the "best" ACP secure channel connection option that also works with the Follower ("best" from the Deciders point of view).
- \* A peer with a "0" acp-address in its AcpNodeName takes on the role of Follower when peering with a node that has a non-"0" acp-address (note that this specification does not fully define the

behavior of ACP secure channel negotiation for nodes with a "0" ACP address field, it only defines interoperability with such ACP nodes).

In a simple example, ACP peer Node 1 attempts to initiate an IPsec via IKEv2 connection to peer Node 2. The IKEv2 authentication succeeds. Node 1 has the lower ACP address and becomes the Follower. Node 2 becomes the Decider. IKEv2 might not be the preferred ACP secure channel protocol for the Decider Node 2. Node 2 would therefore proceed to attempt secure channel setups with (in its view) more preferred protocol options (e.g., DTLS/UDP). If any such preferred ACP secure channel connection of the Decider succeeds, it would close the IPsec connection. If Node 2 has no preferred protocol option over IPsec, or no such connection attempt from Node 2 to Node 1 succeeds, Node 2 would keep the IPsec connection and use it.

The Decider SHOULD NOT send actual payload packets across a secure channel until it has decided to use it. The Follower MAY delay linking the ACP secure channel into the ACP virtual interface until it sees the first payload packet from the Decider up to a maximum of 5 seconds to avoid unnecessarily linking a secure channel that will be terminated as undesired by the Decider shortly afterwards.

The following sequence of steps show this example in more detail. Each step is tagged with [<step#>{:<connection>}]. The connection is included to more easily distinguish which of the two competing connections the step belongs to, one initiated by Node 1, one initiated by Node 2.

- [1] Node 1 sends GRASP AN\_ACP message to announce itself
- [2] Node 2 sends GRASP AN\_ACP message to announce itself
- [3] Node 2 receives [1] from Node 1
- [4:C1] Because of [3], Node 2 starts as initiator on its preferred secure channel protocol towards Node 1. Connection C1.
- [5] Node 1 receives [2] from Node 2
- [6:C2] Because of [5], Node 1 starts as initiator on its preferred secure channel protocol towards Node 2. Connection C2.
- [7:C1] Node1 and Node2 have authenticated each others certificate on connection C1 as valid ACP peers.
- [8:C1] Node 1 certificate has lower ACP Node-ID than Node2, therefore Node 1 considers itself the Follower and Node 2 the Decider on connection C1. Connection setup C1 is completed.
- [9] Node 1 refrains from attempting any further secure channel connections to Node 2 (the Decider) as learned from [2] because it knows from [8:C1] that it is the Follower relative to Node 1.
- [10:C2] Node1 and Node2 have authenticated each others certificate on connection C2 (like [7:C1]).
- [11:C2] Node 1 certificate has lower ACP Node-ID than Node2, therefore Node 1 considers itself the Follower and Node 2 the Decider on connection C2, but they also identify that C2 is to the same mutual peer as their C1, so this has no further impact: the roles Decider and Follower where already assigned between these two peers by [8:C1].
- [12:C2] Node 2 (the Decider) closes C1. Node 1 is fine with this, because of its role as the Follower (from [8:C1]).
- [13] Node 2 (the Decider) and Node 1 (the Follower) start data transfer across C2, which makes it become a secure channel for the ACP.

Figure 8: Secure Channel sequence of steps

All this negotiation is in the context of an "L2 interface". The Decider and Follower will build ACP connections to each other on every "L2 interface" that they both connect to. An autonomic node MUST NOT assume that neighbors with the same L2 or link-local IPv6 addresses on different L2 interfaces are the same node. This can only be determined after examining the certificate after a successful security association attempt.

The Decider SHOULD NOT suppress attempting a particular ACP secure channel protocol connection on one L2 interface because this type of ACP secure channel connection has failed to the peer with the same ACP certificate on another L2 interface: Not only the supported ACP secure channel protocol options may be different on the same ACP peer across different L2 interfaces, but also error conditions may cause inconsistent failures across different L2 interfaces. Avoiding such connection attempt optimizations can therefore help to increase robustness in the case of errors.

#### 6.7. Candidate ACP Neighbor verification

Independent of the security association protocol chosen, candidate ACP neighbors need to be authenticated based on their domain certificate. This implies that any secure channel protocol MUST support certificate based authentication that can support the ACP domain membership check as defined in Section 6.2.3. If it fails, the connection attempt is aborted and an error logged. Attempts to reconnect MUST be throttled. The RECOMMENDED default is exponential base 2 backoff with an initial retransmission time (IRT) of 10 seconds and a maximum retransmission time (MRT) of 640 seconds.

Failure to authenticate an ACP neighbor when acting in the role of a responder of the security authentication protocol MUST NOT impact the attempts of the ACP node to attempt establishing a connection as an initiator. Only failed connection attempts as an initiator must cause throttling. This rule is meant to increase resilience of secure channel creation. Section 6.6 shows how simultaneous mutual secure channel setup collisions are resolved.

#### 6.8. Security Association (Secure Channel) protocols

This section describes how ACP nodes establish secured data connections to automatically discovered or configured peers in the ACP. Section 6.4 above described how IPv6 subnet adjacent peers are discovered automatically. Section 8.2 describes how non IPv6 subnet adjacent peers can be configured.

Section 6.13.5.2 describes how secure channels are mapped to virtual IPv6 subnet interfaces in the ACP. The simple case is to map every ACP secure channel into a separate ACP point-to-point virtual interface Section 6.13.5.2.1. When a single subnet has multiple ACP peers this results in multiple ACP point-to-point virtual interfaces across that underlying multi-party IPv6 subnet. This can be optimized with ACP multi-access virtual interfaces (Section 6.13.5.2.2) but the benefits of that optimization may not justify the complexity of that option.

#### 6.8.1. General considerations

Due to Channel Selection (Section 6.6), ACP can support an evolving set of security association protocols and does not require support for a single network wide MTI. ACP nodes only need to implement those protocols required to interoperate with their candidate peers, not with potentially any node in the ACP domain. See Section 6.8.5 for an example of this.

The degree of security required on every hop of an ACP network needs to be consistent across the network so that there is no designated "weakest link" because it is that "weakest link" that would otherwise become the designated point of attack. When the secure channel protection on one link is compromised, it can be used to send/receive packets across the whole ACP network. Therefore, even though the security association protocols can be different, their minimum degree of security should be comparable.

Secure channel protocols do not need to always support arbitrary L3 connectivity between peers, but can leverage the fact that the standard use case for ACP secure channels is an L2 adjacency. Hence, L2 dependent mechanisms could be adopted for use as secure channel association protocols:

L2 mechanisms such as strong encrypted radio technologies or [MACSEC] may offer equivalent encryption and the ACP security association protocol may only be required to authenticate ACP domain membership of a peer and/or derive a key for the L2 mechanism. Mechanisms to auto-discover and associate ACP peers leveraging such underlying L2 security are possible and desirable to avoid duplication of encryption, but none are specified in this document.

Strong physical security of a link may stand in where cryptographic security is infeasible. As there is no secure mechanism to automatically discover strong physical security solely between two peers, it can only be used with explicit configuration and that configuration too could become an attack vector. This document therefore only specifies with ACP connect (Section 8.1) one

explicitly configured mechanism without any secure channel association protocol – for the case where both the link and the nodes attached to it have strong physical security.

#### 6.8.2. Common requirements

The authentication of peers in any security association protocol MUST use the ACP certificate according to Section 6.2.3. Because auto-discovery of candidate ACP neighbors via GRASP (see Section 6.4) as specified in this document does not communicate the neighbors ACP certificate, and ACP nodes may not (yet) have any other network connectivity to retrieve certificates, any security association protocol MUST use a mechanism to communicate the certificate directly instead of relying on a referential mechanism such as communicating only a hash and/or URL for the certificate.

A security association protocol MUST use Forward Secrecy (whether inherently or as part of a profile of the security association protocol).

Because the ACP payload of legacy protocol payloads inside the ACP and hop-by-hop ACP flooded GRASP information is unencrypted, the ACP secure channel protocol requires confidentiality. Symmetric encryption for the transmission of secure channel data MUST use encryption schemes considered to be security wise equal to or better than 256-bit key strength, such as AES256. There MUST NOT be support for NULL encryption.

Security association protocols typically only signal the End Entity certificate (e.g. the ACP certificate) and any possible intermediate CA certificates for successful mutual authentication. The TA has to be mutually known and trusted and therefore its certificate does not need to be signaled for successful mutual authentication. Nevertheless, for use with ACP secure channel setup, there SHOULD be the option to include the TA certificate in the signaling to aid troubleshooting, see Section 9.1.1.

Signaling of TA certificates may not be appropriate when the deployment is relying on a security model where the TA certificate content is considered confidential and only its hash is appropriate for signaling. ACP nodes SHOULD have a mechanism to select whether the TA certificate is signaled or not. Assuming that both options are possible with a specific secure channel protocol.

An ACP secure channel MUST immediately be terminated when the lifetime of any certificate in the chain used to authenticate the neighbor expires or becomes revoked. This may not be standard behavior in secure channel protocols because the certificate

authentication may only influence the setup of the secure channel in these protocols, but may not be re-validated during the lifetime of the secure connection in the absence of this requirement.

When specifying an additional security association protocol for ACP secure channels beyond those covered in this document, protocol options SHOULD be eliminated that are not necessary to support devices that are expected to be able to support the ACP to minimize implementation complexity. For example, definitions for security protocols often include old/inferior security options required only to interoperate with existing devices that will not be able to update to the currently preferred security options. Such old/inferior security options do not need to be supported when a security association protocol is first specified for the ACP, strengthening the "weakest link" and simplifying ACP implementation overhead.

### 6.8.3. ACP via IPsec

An ACP node announces its ability to support IPsec, negotiated via IKEv2, as the ACP secure channel protocol using the "IKEv2" objective-value in the "AN\_ACP" GRASP objective.

The ACP usage of IPsec and IKEv2 mandates a profile with a narrow set of options of the current standards-track usage guidance for IPsec [RFC8221] and IKEv2 [RFC8247]. These options result in stringent security properties and can exclude deprecated/legacy algorithms because there is no need for interoperability with legacy equipment for ACP secure channels. Any such backward compatibility would lead only to increased attack surface and implementation complexity, for no benefit.

#### 6.8.3.1. Native IPsec

An ACP node that is supporting native IPsec MUST use IPsec in tunnel mode, negotiated via IKEv2, and with IPv6 payload (e.g., ESP Next Header of 41). It MUST use local and peer link-local IPv6 addresses for encapsulation. Manual keying MUST NOT be used, see Section 6.2. Traffic Selectors are:

TSi = (0, 0-65535, :: - FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF)

TSr = (0, 0-65535, :: - FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF)

IPsec tunnel mode is required because the ACP will route/forward packets received from any other ACP node across the ACP secure channels, and not only its own generated ACP packets. With IPsec transport mode (and no additional encapsulation header in the ESP payload), it would only be possible to send packets originated by the ACP node itself because the IPv6 addresses of the ESP must be the same as that of the outer IPv6 header.

#### 6.8.3.1.1. RFC8221 (IPsec/ESP)

ACP IPsec implementations MUST comply with [RFC8221] (and its updates). The requirements from above and this section amend and superseded its requirements.

The IP Authentication Header (AH) MUST NOT be used (because it does not provide confidentiality).

For the required ESP encryption algorithms in section 5 of [RFC8221] the following guidance applies:

- \* ENCR\_NULL AH MUST NOT be used (because it does not provide confidentiality).
- \* ENCR\_AES\_GCM\_16 is the only MTI ESP encryption algorithm for ACP via IPsec/ESP (it is already listed as MUST in [RFC8221]).
- \* ENCR\_AES\_CBC with AUTH\_HMAC\_SHA2\_256\_128 (as the ESP authentication algorithm) and ENCR\_AES\_CCM\_8 MAY be supported. If either provides higher performance than ENCR\_AES\_GCM\_16 it SHOULD be supported.
- \* ENCR\_CHACHA20\_POLY1305 SHOULD be supported at equal or higher performance than ENCR\_AES\_GCM\_16. If that performance is not feasible, it MAY be supported.

IKEv2 indicates an order for the offered algorithms. The algorithms SHOULD be ordered by performance. The first algorithm supported by both sides is generally chosen.

Explanations:

- \* There is no requirement to interoperate with legacy equipment in ACP secure channels, so a single MTI encryption algorithm for IPsec in ACP secure channels is sufficient for interoperability and allows for the most lightweight implementations.
- \* ENCR\_AES\_GCM\_16 is an authenticated encryption with associated data (AEAD) cipher mode, so no additional ESP authentication algorithm is needed, simplifying the MTI requirements of IPsec for the ACP.

- \* There is no MTI requirement for the support of ENCR\_AES\_CBC because ENCR\_AES\_GCM\_16 is assumed to be feasible with less cost/higher performance in modern devices hardware accelerated implementations compared to ENCR-AES\_CBC.
- \* ENCR\_CHACHA20\_POLY1305 is mandatory in [RFC8221] because of its target use as a fallback algorithm in case weaknesses in AES are uncovered. Unfortunately, there is currently no way to automatically propagate across an ACP a policy to disallow use of AES based algorithms, so this target benefit of ENCR\_CHACHA20\_POLY1305 cannot fully be adopted yet for the ACP. Therefore, this algorithm is only recommended. Changing from AES to this algorithm at potentially big drop in performance could also render the ACP inoperable. Therefore, the performance requirement against this algorithm so that it could become an effective security backup to AES for the ACP once a policy to switch over to it or prefer it is available in an ACP framework.

[RFC8221] allows for 128-bit or 256-bit AES keys. This document mandates that only 256-bit AES keys MUST be supported.

When [RFC8221] is updated, ACP implementations will need to consider legacy interoperability, and the IPsec WG has generally done a very good job of taking that into account in its recommendations.

#### 6.8.3.1.2. RFC8247 (IKEv2)

[RFC8247] provides a baseline recommendation for mandatory to implement ciphers, integrity checks, pseudo-random-functions and Diffie-Hellman mechanisms. Those recommendations, and the recommendations of subsequent documents apply well to the ACP. Because IKEv2 for ACP secure channels is sufficient to be implemented in control plane software, rather than in ASIC hardware, and ACP nodes supporting IKEv2 are not assumed to be code-space constrained, and because existing IKEv2 implementations are expected to support [RFC8247] recommendations, this documents makes no attempt to simplify its recommendations for use with the ACP.

See [IKEV2IANA] for IANA IKEv2 parameter names used in this text.

ACP Nodes supporting IKEv2 MUST comply with [RFC8247] amended by the following requirements which constitute a policy statement as permitted by [RFC8247].

To signal the ACP certificate chain (including TA) as required by Section 6.8.2, "X.509 Certificate - Signature" payload in IKEv2 can be used. It is mandatory according to [RFC7296] section 3.6.

ACP nodes SHOULD set up IKEv2 to only use the ACP certificate and TA when acting as an IKEv2 responder on the IPv6 link local address and port number indicated in the AN\_ACP DULL GRASP announcements (see Section 6.4).

When CERTREQ is received from a peer, and does not indicate any of this ACP nodes TA certificates, the ACP node SHOULD ignore the CERTREQ and continue sending its certificate chain including its TA as subject to the requirements and explanations in Section 6.8.2. This will not result in successful mutual authentication but assists diagnostics.

Note that with IKEv2, failing authentication will only result in the responder receiving the certificate chain from the initiator, but not vice versa. Because ACP secure channel setup is symmetric (see Section 6.7), every non-malicious ACP neighbor will attempt to connect as an initiator though, allowing to obtain the diagnostic information about the neighbors certificate.

In IKEv2, ACP nodes are identified by their ACP address. The ID\_IPv6\_ADDR IKEv2 identification payload MUST be used and MUST convey the ACP address. If the peer's ACP certificate includes a 32HEXDIG ACP address in the acp-node-name (not "0" or omitted), the address in the IKEv2 identification payload MUST match it. See Section 6.2.3 for more information about "0" or omitted ACP address fields in the acp-node-name.

IKEv2 authentication MUST use authentication method 14 ("Digital Signature") for ACP certificates; this authentication method can be used with both RSA and ECDSA certificates, indicated by an ASN.1 object AlgorithmIdentifier.

The Digital Signature hash SHA2-512 MUST be supported (in addition to SHA2-256).

The IKEv2 Diffie-Hellman key exchange group 19 (256-bit random ECP), MUST be supported. Reason: ECC provides a similar security level to finite-field (MODP) key exchange with a shorter key length, so is generally preferred absent other considerations.

#### 6.8.3.2. IPsec with GRE encapsulation

In network devices it is often more common to implement high performance virtual interfaces on top of GRE encapsulation than on top of a "native" IPsec association (without any other encapsulation than those defined by IPsec). On those devices it may be beneficial to run the ACP secure channel on top of GRE protected by the IPsec association.

The requirements for ESP/IPsec/IKEv2 with GRE are the same as for native IPsec (see Section 6.8.3.1) except that IPsec transport mode and next protocol GRE (47) are to be negotiated. Tunnel mode is not required because of GRE. Traffic Selectors are:

TSi = (47, 0-65535, Initiator-IPv6-LL-addr ... Initiator-IPv6-LL-addr)

TSr = (47, 0-65535, Responder-IPv6-LL-addr ... Responder-IPv6-LL-addr)

If IKEv2 initiator and responder support IPsec over GRE, it will be preferred over native IPsec because of the way how IKEv2 negotiates transport mode (as used by this IPsec over GRE profile) versus tunnel mode as used by native IPsec (see [RFC7296], section 1.3.1). The ACP IPv6 traffic has to be carried across GRE according to [RFC7676].

#### 6.8.4. ACP via DTLS

This document defines the use of ACP via DTLS, on the assumption that it is likely the first transport encryption supported in some classes of constrained devices: DTLS is commonly used in constrained devices when IPsec is not. Code-space on those devices may be also be too limited to support more than the minimum number of required protocols.

An ACP node announces its ability to support DTLS version 1.2 ([RFC6347]) compliant with the requirements defined in this document as an ACP secure channel protocol in GRASP through the "DTLS" objective-value in the "AN\_ACP" objective (see Section 6.4).

To run ACP via UDP and DTLS, a locally assigned UDP port is used that is announced as a parameter in the GRASP AN\_ACP objective to candidate neighbors. This port can also be any newer version of DTLS as long as that version can negotiate a DTLS v1.2 connection in the presence of an DTLS v1.2 only peer.

All ACP nodes supporting DTLS as a secure channel protocol MUST adhere to the DTLS implementation recommendations and security considerations of BCP 195, BCP 195 [RFC7525] except with respect to the DTLS version. ACP nodes supporting DTLS MUST support DTLS 1.2. They MUST NOT support older versions of DTLS.

Unlike for IPsec, no attempts are made to simplify the requirements of the BCP 195 recommendations because the expectation is that DTLS would be using software-only implementations where the ability to reuse of widely adopted implementations is more important than minimizing the complexity of a hardware accelerated implementation which is known to be important for IPsec.

DTLS v1.3 ([I-D.ietf-tls-dtls13]) is "backward compatible" with DTLS v1.2 (see section 1. of DTLS v1.3). A DTLS implementation supporting both DTLS v1.2 and DTLS v1.3 does comply with the above requirements of negotiating to DTLS v1.2 in the presence of a DTLS v1.2 only peer, but using DTLS v1.3 when both peers support it.

Version v1.2 is the MTI version of DTLS in this specification because

- \* There is more experience with DTLS v1.2 across the spectrum of target ACP nodes.
- \* Firmware of lower end, embedded ACP nodes may not support a newer version for a long time.
- \* There are significant changes of DTLS v1.3, such as a different record layer requiring time to gain implementation and deployment experience especially on lower end, code space limited devices.
- \* The existing BCP [RFC7525] for DTLS v1.2 may equally take longer time to be updated with experience from a newer DTLS version.
- \* There are no significant use-case relevant benefits of DTLS v1.3 over DTLS v1.2 in the context of the ACP options for DTLS. For example, signaling performance improvements for session setup in DTLS v1.3 is not important for the ACP given the long-lived nature of ACP secure channel connections and the fact that DTLS connections are mostly link-local (short RTT).

Nevertheless, newer versions of DTLS, such as DTLS v1.3 have stricter security requirements and use of the latest standard protocol version is for IETF security standards in general recommended. Therefore, ACP implementations are advised to support all the newer versions of DTLS that can still negotiate down to DTLS v1.2.

[RFC-editor: if by the time of AUTH48, DTLS 1.3 would have evolved to be an RFC, then not only would the references to the DTLS v1.3 draft be changed to the RFC number, but that RFC is then going to be put into the normative list of references and the above paragraph is going to be amended to say: Implementations SHOULD support [DTLSv1.3-RFC]. This is not done right now, because there is no benefit in potentially waiting in RFC-editor queue for that RFC given how the text already lays out a non-normative desire to support DTLSv1.3.]

There is no additional session setup or other security association besides this simple DTLS setup. As soon as the DTLS session is functional, the ACP peers will exchange ACP IPv6 packets as the payload of the DTLS transport connection. Any DTLS defined security association mechanisms such as re-keying are used as they would be for any transport application relying solely on DTLS.

#### 6.8.5. ACP Secure Channel Profiles

As explained in the beginning of Section 6.6, there is no single secure channel mechanism mandated for all ACP nodes. Instead, this section defines two ACP profiles (baseline and constrained) for ACP nodes that do introduce such requirements.

An ACP node supporting the "baseline" profile MUST support IPsec natively and MAY support IPsec via GRE. An ACP node supporting the "constrained" profile node that cannot support IPsec MUST support DTLS. An ACP node connecting an area of constrained ACP nodes with an area of baseline ACP nodes needs to support IPsec and DTLS and supports therefore the baseline and constrained profile.

Explanation: Not all type of ACP nodes can or need to connect directly to each other or are able to support or prefer all possible secure channel mechanisms. For example, code space limited IoT devices may only support DTLS because that code exists already on them for end-to-end security, but high-end core routers may not want to support DTLS because they can perform IPsec in accelerated hardware but would need to support DTLS in an underpowered CPU forwarding path shared with critical control plane operations. This is not a deployment issue for a single ACP across these type of nodes as long as there are also appropriate gateway ACP nodes that support sufficiently many secure channel mechanisms to allow interconnecting areas of ACP nodes with a more constrained set of secure channel protocols. On the edge between IoT areas and high-end core networks, general-purpose routers that act as those gateways and that can support a variety of secure channel protocols is the norm already.

IPsec natively with tunnel mode provides the shortest encapsulation overhead. GRE may be preferred by legacy implementations because the virtual interfaces required by ACP design in conjunction with secure channels have in the past more often been implemented for GRE than purely for native IPsec.

ACP nodes need to specify in documentation the set of secure ACP mechanisms they support and should declare which profile they support according to above requirements.

## 6.9. GRASP in the ACP

### 6.9.1. GRASP as a core service of the ACP

The ACP MUST run an instance of GRASP inside of it. It is a key part of the ACP services. The function in GRASP that makes it fundamental as a service of the ACP is the ability to provide ACP wide service discovery (using objectives in GRASP).

ACP provides IP unicast routing via the RPL routing protocol (see Section 6.12).

The ACP does not use IP multicast routing nor does it provide generic IP multicast services (the handling of GRASP link-local multicast messages is explained in Section 6.9.2). Instead, the ACP provides service discovery via the objective discovery/announcement and negotiation mechanisms of the ACP GRASP instance (services are a form of objectives). These mechanisms use hop-by-hop reliable flooding of GRASP messages for both service discovery (GRASP M\_DISCOVERY messages) and service announcement (GRASP M\_FLOOD messages).

See Appendix A.5 for discussion about this design choice of the ACP.

### 6.9.2. ACP as the Security and Transport substrate for GRASP

In the terminology of GRASP ([I-D.ietf-anima-grasp]), the ACP is the security and transport substrate for the GRASP instance run inside the ACP ("ACP GRASP").

This means that the ACP is responsible for ensuring that this instance of GRASP is only sending messages across the ACP GRASP virtual interfaces. Whenever the ACP adds or deletes such an interface because of new ACP secure channels or loss thereof, the ACP needs to indicate this to the ACP instance of GRASP. The ACP exists also in the absence of any active ACP neighbors. It is created when the node has a domain certificate, and continues to exist even if all of its neighbors cease operation.

In this case ASAs using GRASP running on the same node would still need to be able to discover each other's objectives. When the ACP does not exist, ASAs leveraging the ACP instance of GRASP via APIs MUST still be able to operate, and MUST be able to understand that there is no ACP and that therefore the ACP instance of GRASP cannot operate.

The following explanation how ACP acts as the security and transport substrate for GRASP is visualized in Figure 9 below.

GRASP unicast messages inside the ACP always use the ACP address. Link-local addresses from the ACP VRF MUST NOT be used inside objectives. GRASP unicast messages inside the ACP are transported via TLS. See Section 6.1 for TLS requirements. TLS mutual authentication MUST use the ACP domain membership check defined in (Section 6.2.3).

GRASP link-local multicast messages are targeted for a specific ACP virtual interface (as defined Section 6.13.5) but are sent by the ACP into an ACP GRASP virtual interface that is constructed from the TCP connection(s) to the IPv6 link-local neighbor address(es) on the underlying ACP virtual interface. If the ACP GRASP virtual interface has two or more neighbors, the GRASP link-local multicast messages are replicated to all neighbor TCP connections.

TCP and TLS connections for GRASP in the ACP use the IANA assigned TCP port for GRASP (7107). Effectively the transport stack is expected to be TLS for connections from/to the ACP address (e.g., global scope address(es)) and TCP for connections from/to link-local addresses on the ACP virtual interfaces. The latter ones are only used for flooding of GRASP messages.

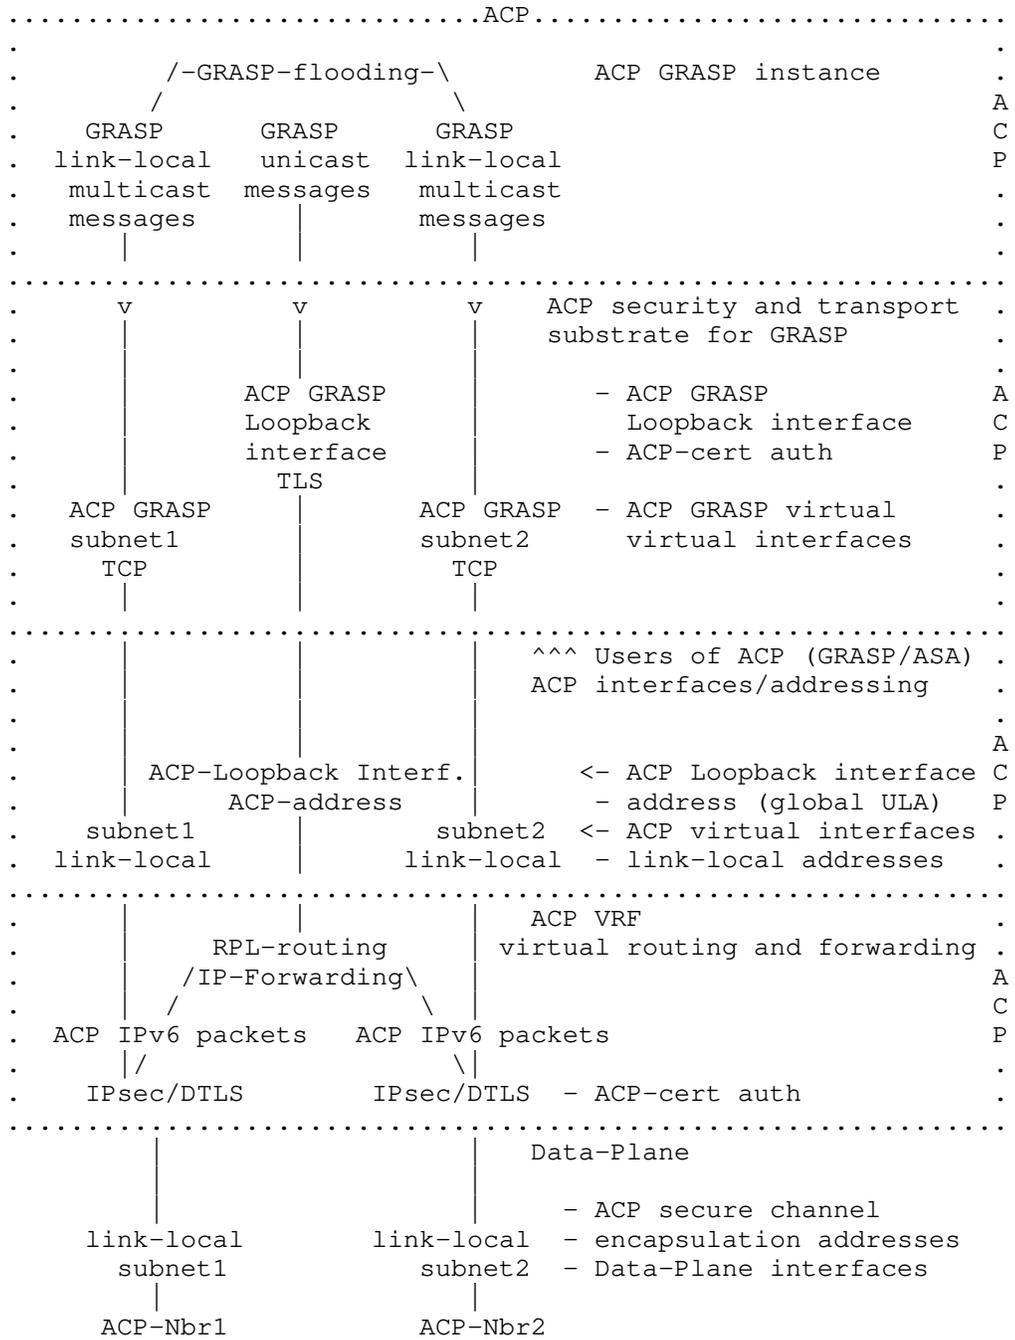


Figure 9: ACP as security and transport substrate for GRASP

### 6.9.2.1. Discussion

TCP encapsulation for GRASP M\_DISCOVERY and M\_FLOOD link local messages is used because these messages are flooded across potentially many hops to all ACP nodes and a single link with even temporary packet loss issues (e.g., WiFi/Powerline link) can reduce the probability for loss free transmission so much that applications would want to increase the frequency with which they send these messages. Such shorter periodic retransmission of datagrams would result in more traffic and processing overhead in the ACP than the hop-by-hop reliable retransmission mechanism by TCP and duplicate elimination by GRASP.

TLS is mandated for GRASP non-link-local unicast because the ACP secure channel mandatory authentication and encryption protects only against attacks from the outside but not against attacks from the inside: Compromised ACP members that have (not yet) been detected and removed (e.g., via domain certificate revocation / expiry).

If GRASP peer connections were to use just TCP, compromised ACP members could simply eavesdrop passively on GRASP peer connections for whom they are on-path ("man in the middle" - MITM) or intercept and modify them. With TLS, it is not possible to completely eliminate problems with compromised ACP members, but attacks are a lot more complex:

Eavesdropping/spoofing by a compromised ACP node is still possible because in the model of the ACP and GRASP, the provider and consumer of an objective have initially no unique information (such as an identity) about the other side which would allow them to distinguish a benevolent from a compromised peer. The compromised ACP node would simply announce the objective as well, potentially filter the original objective in GRASP when it is a MITM and act as an application level proxy. This of course requires that the compromised ACP node understand the semantics of the GRASP negotiation to an extent that allows it to proxy it without being detected, but in an ACP environment this is quite likely public knowledge or even standardized.

The GRASP TLS connections are run the same as any other ACP traffic through the ACP secure channels. This leads to double authentication/encryption, which has the following benefits:

- \* Secure channel methods such as IPsec may provide protection against additional attacks, for example reset-attacks.
- \* The secure channel method may leverage hardware acceleration and there may be little or no gain in eliminating it.

- \* There is no different security model for ACP GRASP from other ACP traffic. Instead, there is just another layer of protection against certain attacks from the inside which is important due to the role of GRASP in the ACP.

#### 6.10. Context Separation

The ACP is in a separate context from the normal Data-Plane of the node. This context includes the ACP channels' IPv6 forwarding and routing as well as any required higher layer ACP functions.

In classical network system, a dedicated VRF is one logical implementation option for the ACP. If possible by the systems software architecture, separation options that minimize shared components are preferred, such as a logical container or virtual machine instance. The context for the ACP needs to be established automatically during bootstrap of a node. As much as possible it should be protected from being modified unintentionally by ("Data-Plane") configuration.

Context separation improves security, because the ACP is not reachable from the Data-Plane routing or forwarding table(s). Also, configuration errors from the Data-Plane setup do not affect the ACP.

#### 6.11. Addressing inside the ACP

The channels explained above typically only establish communication between two adjacent nodes. In order for communication to happen across multiple hops, the autonomic control plane requires ACP network wide valid addresses and routing. Each ACP node creates a Loopback interface with an ACP network wide unique address (prefix) inside the ACP context (as explained in in Section 6.10). This address may be used also in other virtual contexts.

With the algorithm introduced here, all ACP nodes in the same routing subdomain have the same /48 ULA prefix. Conversely, ULA global IDs from different domains are unlikely to clash, such that two ACP networks can be merged, as long as the policy allows that merge. See also Section 10.1 for a discussion on merging domains.

Links inside the ACP only use link-local IPv6 addressing, such that each node's ACP only requires one routable address prefix.

##### 6.11.1. Fundamental Concepts of Autonomic Addressing

- \* Usage: Autonomic addresses are exclusively used for self-management functions inside a trusted domain. They are not used for user traffic. Communications with entities outside the

trusted domain use another address space, for example normally managed routable address space (called "Data-Plane" in this document).

- \* Separation: Autonomic address space is used separately from user address space and other address realms. This supports the robustness requirement.
- \* Loopback-only: Only ACP Loopback interfaces (and potentially those configured for "ACP connect", see Section 8.1) carry routable address(es); all other interfaces (called ACP virtual interfaces) only use IPv6 link local addresses. The usage of IPv6 link local addressing is discussed in [RFC7404].
- \* Use-ULA: For Loopback interfaces of ACP nodes, we use ULA with L=1 (as defined in section 3.1 of [RFC4193]). Note that the random hash for ACP Loopback addresses uses the definition in Section 6.11.2 and not the one of [RFC4193] section 3.2.2.
- \* No external connectivity: They do not provide access to the Internet. If a node requires further reaching connectivity, it should use another, traditionally managed address scheme in parallel.
- \* Addresses in the ACP are permanent, and do not support temporary addresses as defined in [RFC4941].
- \* Addresses in the ACP are not considered sensitive on privacy grounds because ACP nodes are not expected to be end-user hosts and ACP addresses do therefore not represent end-users or groups of end-users. All ACP nodes are in one (potentially federated) administrative domain. They are assumed to be to be candidate hosts of ACP traffic amongst each other or transit thereof. There are no transit nodes less privileged to know about the identity of other hosts in the ACP. Therefore, ACP addresses do not need to be pseudo-random as discussed in [RFC7721]. Because they are not propagated to untrusted (non ACP) nodes and stay within a domain (of trust), we also consider them not to be subject to scanning attacks.

The ACP is based exclusively on IPv6 addressing, for a variety of reasons:

- \* Simplicity, reliability and scale: If other network layer protocols were supported, each would have to have its own set of security associations, routing table and process, etc.
- \* Autonomic functions do not require IPv4: Autonomic functions and autonomic service agents are new concepts. They can be exclusively built on IPv6 from day one. There is no need for backward compatibility.
- \* OAM protocols do not require IPv4: The ACP may carry OAM protocols. All relevant protocols (SNMP, TFTP, SSH, SCP, RADIUS, Diameter, NETCONF ...) are available in IPv6. See also [RFC8368] for how ACP could be made to interoperate with IPv4 only OAM.

Further explanation about the addressing and routing related reasons for the choice of the autonomous ACP addressing can be found in Section 6.13.5.1.

#### 6.11.2. The ACP Addressing Base Scheme

The Base ULA addressing scheme for ACP nodes has the following format:

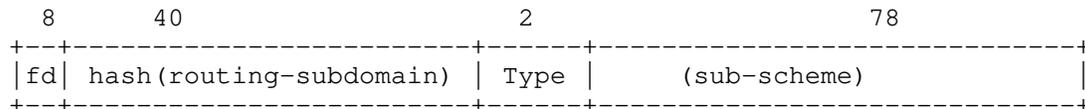


Figure 10: ACP Addressing Base Scheme

The first 48-bits follow the ULA scheme, as defined in [RFC4193], to which a type field is added:

- \* "fd" identifies a locally defined ULA address.
- \* The 40-bits ULA "global ID" (term from [RFC4193]) for ACP addresses carried in the acp-node-name in the ACP certificates are the first 40-bits of the SHA256 hash of the routing subdomain from the same acp-node-name. In the example of Section 6.2.2, the routing subdomain is "area51.research.acp.example.com" and the 40-bits ULA "global ID" 89b714f3db.
- \* When creating a new routing-subdomain for an existing autonomic network, it MUST be ensured, that rsub is selected so the resulting hash of the routing-subdomain does not collide with the hash of any pre-existing routing-subdomains of the autonomic network. This ensures that ACP addresses created by registrars for different routing subdomains do not collide with each other.
- \* To allow for extensibility, the fact that the ULA "global ID" is a hash of the routing subdomain SHOULD NOT be assumed by any ACP node during normal operations. The hash function is only executed during the creation of the certificate. If BRSKI is used, then the BRSKI registrar will create the acp-node-name in response to the EST Certificate Signing Request (CSR) Attribute Request message by the pledge.

- \* Establishing connectivity between different ACP (different acp-domain-name) is outside the scope of this specification. If it is being done through future extensions, then the rsub of all routing-subdomains across those autonomic networks need to be selected so the resulting routing-subdomain hashes do not collide. For example, a large cooperation with its own private TA may want to create different autonomic networks that initially should not be able to connect but where the option to do so should be kept open. When taking this future possibility into account, it is easy to always select rsub so that no collisions happen.
- \* Type: This field allows different address sub-schemes. This addresses the "upgradability" requirement. Assignment of types for this field will be maintained by IANA.

The sub-scheme may imply a range or set of addresses assigned to the node, this is called the ACP address range/set and explained in each sub-scheme.

Please refer to Section 6.11.7 and Appendix A.1 for further explanations why the following Sub-Addressing schemes are used and why multiple are necessary.

The following summarizes the addressing Sub-Schemes:

Type	Name	F-bit	Z	V-bits	Prefix
0x00	ACP-Zone	N/A	0	1 bit	/127
0x00	ACP-Manual	N/A	1	N/A	/64
0x01	ACP-VLong-8	0	N/A	8 bits	/120
0x01	ACP-VLong-16	1	N/A	16 bits	/112
0x10	Reserved / For future definition/allocation				
0x11	Reserved / For future definition/allocation				

Figure 11: Addressing Sub-Schemes

F-Bit and Z are two encoding fields explained below for the Sub-Schemes that introduce/use them. V-bits is the number of bits of addresses allocated to the ACP node. Prefix is the prefix the ACP node is announcing into the RPL routing protocol.

## 6.11.3. ACP Zone Addressing Sub-Scheme (ACP-Zone)

This sub-scheme is used when the Type field of the base scheme is 0x00 and the Z bit is 0x0.

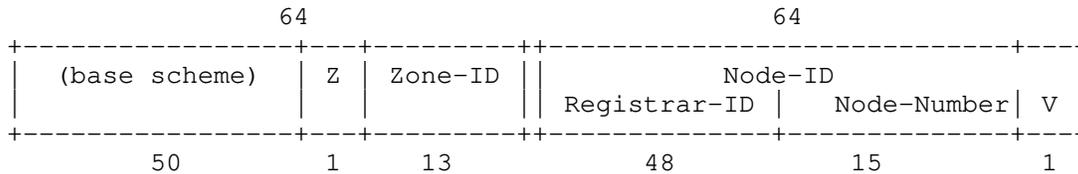


Figure 12: ACP Zone Addressing Sub-Scheme

The fields are defined as follows:

- \* Type: MUST be 0x0.
- \* Z: MUST be 0x0.
- \* Zone-ID: A value for a network zone.
- \* Node-ID: A unique value for each node.

The 64-bit Node-ID must be unique across the ACP domain for each node. It is derived and composed as follows:

- \* Registrar-ID (48-bit): A number unique inside the domain that identifies the ACP registrar which assigned the Node-ID to the node. One or more domain-wide unique identifiers of the ACP registrar can be used for this purpose. See Section 6.11.7.2.
- \* Node-Number: Number to make the Node-ID unique. This can be sequentially assigned by the ACP Registrar owning the Registrar-ID.
- \* V (1-bit): Virtualization bit: 0: Indicates the ACP itself ("ACP node base system"); 1: Indicates the optional "host" context on the ACP node (see below).

In the ACP Zone Addressing Sub-Scheme, the ACP address in the certificate has V field as all zero bits.

The ACP address set of the node includes addresses with any Zone-ID value and any V value. No two nodes in the same ACP can have the same Node-ID, but different Zone-IDs.

The Virtual bit in this sub-scheme allows the easy addition of the ACP as a component to existing systems without causing problems in the port number space between the services in the ACP and the existing system. V:0 is the ACP router (autonomic node base system), V:1 is the host with pre-existing transport endpoints on it that

could collide with the transport endpoints used by the ACP router. The ACP host could for example have a p2p virtual interface with the V:0 address as its router into the ACP. Depending on the software design of ASAs, which is outside the scope of this specification, they may use the V:0 or V:1 address.

The location of the V bit(s) at the end of the address allows the announcement of a single prefix for each ACP node. For example, in a network with 20,000 ACP nodes, this avoid 20,000 additional routes in the routing table.

It is RECOMMENDED that only Zone-ID 0 is used unless it is meant to be used in conjunction with operational practices for partial/incremental adoption of the ACP as described in Section 9.4.

Note: Zones and Zone-ID as defined here are not related to [RFC4007] zones or zone\_id. ACP zone addresses are not scoped (reachable only from within an RFC4007 zone) but reachable across the whole ACP. An RFC4007 zone\_id is a zone index that has only local significance on a node, whereas an ACP Zone-ID is an identifier for an ACP zone that is unique across that ACP.

6.11.4. ACP Manual Addressing Sub-Scheme (ACP-Manual)

This sub-scheme is used when the Type field of the base scheme is 0x00 and the Z bit is 0x1.

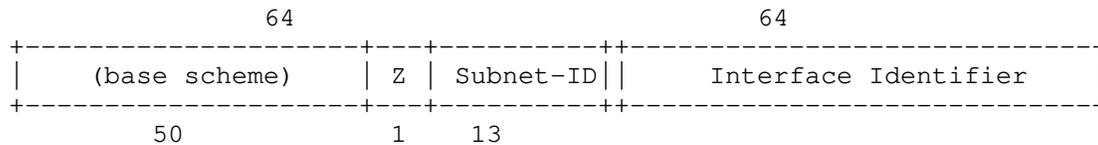


Figure 13: ACP Manual Addressing Sub-Scheme

The fields are defined as follows:

- \* Type: MUST be 0x0.
- \* Z: MUST be 0x1.
- \* Subnet-ID: Configured subnet identifier.
- \* Interface Identifier.

This sub-scheme is meant for "manual" allocation to subnets where the other addressing schemes cannot be used. The primary use case is for assignment to ACP connect subnets (see Section 8.1.1).

"Manual" means that allocations of the Subnet-ID need to be done today with pre-existing, non-autonomic mechanisms. Every subnet that uses this addressing sub-scheme needs to use a unique Subnet-ID (unless some anycast setup is done).

The Z bit field was added to distinguish Zone addressing and manual addressing sub-schemes without requiring one more bit in the base scheme and therefore allowing for the Vlong scheme (described below) to have one more bit available.

Manual addressing sub-scheme addresses SHOULD NOT be used in ACP certificates. Any node capable to build ACP secure channels and permitted by Registrar policy to participate in building ACP secure channels SHOULD receive an ACP address (prefix) from one of the other ACP addressing sub-schemes. Nodes not capable (or permitted) to participate in ACP secure channels can connect to the ACP via ACP connect interfaces of ACP edge nodes (see Section 8.1), without setting up an ACP secure channel. Their ACP certificate MUST omit the acp-address field to indicate that their ACP certificate is only usable for non- ACP secure channel authentication, such as end-to-end transport connections across the ACP or Data-Plane.

Address management of ACP connect subnets is done using traditional assignment methods and existing IPv6 protocols. See Section 8.1.3 for details. Therefore, the notion of V-bit many addresses assigned to the ACP nodes does not apply to this Sub-Scheme.

6.11.5. ACP Vlong Addressing Sub-Scheme (ACP-VLong-8/ACP-VLong-16)

This sub-scheme is used when the Type field of the base scheme is 0x01.

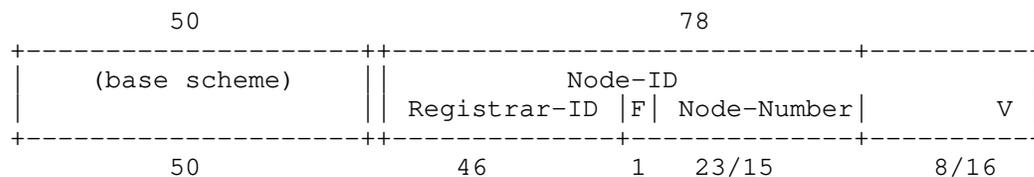


Figure 14: ACP Vlong Addressing Sub-Scheme

This addressing scheme foregoes the Zone-ID field to allow for larger, flatter routed networks (e.g., as in IoT) with 8421376 Node-Numbers (2<sup>23</sup>+2<sup>15</sup>). It also allows for up to 2<sup>16</sup> (i.e. 65536) different virtualized addresses within a node, which could be used to address individual software components in an ACP node.

The fields are the same as in the Zone-ID sub-scheme with the following refinements:

- \* F: format bit. This bit determines the format of the subsequent bits.
- \* V: Virtualization bit: this is a field that is either 8 or 16 bits. For F=0, it is 8 bits, for F=1 it is 16 bits. The V bits are assigned by the ACP node. In the ACP certificate's ACP address Section 6.2.2, the V-bits are always set to 0.
- \* Registrar-ID: To maximize Node-Number and V, the Registrar-ID is reduced to 46-bits. One or more domain-wide unique identifiers of the ACP registrar can be used for this purpose. See Section 6.11.7.2.
- \* The Node-Number is unique to each ACP node. There are two formats for the Node-Number. When F=0, the node-number is 23 bits, for F=1 it is 15 bits. Each format of node-number is considered to be in a unique number space.

The F=0 bit format addresses are intended to be used for "general purpose" ACP nodes that would potentially have a limited number (< 256) of clients (ASA/Autonomic Functions or legacy services) of the ACP that require separate V(irtual) addresses.

The F=1 bit Node-Numbers are intended for ACP nodes that are ACP edge nodes (see Section 8.1.1) or that have a large number of clients requiring separate V(irtual) addresses. For example, large SDN controllers with container modular software architecture (see Section 8.1.2).

In the Vlong addressing sub-scheme, the ACP address in the certificate has all V field bits as zero. The ACP address set for the node includes any V value.

#### 6.11.6. Other ACP Addressing Sub-Schemes

Before further addressing sub-schemes are defined, experience with the schemes defined here should be collected. The schemes defined in this document have been devised to allow hopefully sufficiently flexible setup of ACPs for a variety of situation. These reasons also lead to the fairly liberal use of address space: The Zone Addressing Sub-Scheme is intended to enable optimized routing in large networks by reserving bits for Zone-ID's. The Vlong addressing sub-scheme enables the allocation of 8/16-bit of addresses inside individual ACP nodes. Both address spaces allow distributed, uncoordinated allocation of node addresses by reserving bits for the registrar-ID field in the address.

### 6.11.7. ACP Registrars

ACP registrars are responsible to enroll candidate ACP nodes with ACP certificates and associated trust anchor(s). They are also responsible that an `acp-node-name` field is included in the ACP certificate carrying the ACP domain name and the ACP nodes ACP address prefix. This address prefix is intended to persist unchanged through the lifetime of the ACP node.

Because of the ACP addressing sub-schemes, an ACP domain can have multiple distributed ACP registrars that do not need to coordinate for address assignment. ACP registrars can also be sub-CAs, in which case they can also assign ACP certificates without dependencies against a (shared) TA (except during renewals of their own certificates).

ACP registrars are PKI registration authorities (RA) enhanced with the handling of the ACP certificate specific fields. They request certificates for ACP nodes from a Certification Authority through any appropriate mechanism (out of scope in this document, but required to be BRSKI for ANI registrars). Only nodes that are trusted to be compliant with the requirements against registrar described in this section can be given the necessary credentials to perform this RA function, such as credentials for the BRSKI connection to the CA for ANI registrars.

#### 6.11.7.1. Use of BRSKI or other Mechanism/Protocols

Any protocols or mechanisms may be used by ACP registrars, as long as the resulting ACP certificate and TA certificate(s) allow to perform the ACP domain membership described in Section 6.2.3 with other ACP domain members, and meet the ACP addressing requirements for its `acp-node-name` as described further below in this section.

An ACP registrar could be a person deciding whether to enroll a candidate ACP node and then orchestrating the enrollment of the ACP certificate and associated TA, using command line or web based commands on the candidate ACP node and TA to generate and sign the ACP certificate and configure certificate and TA onto the node.

The only currently defined protocol for ACP registrars is BRSKI ([I-D.ietf-anima-bootstrapping-keyinfra]). When BRSKI is used, the ACP nodes are called ANI nodes, and the ACP registrars are called BRSKI or ANI registrars. The BRSKI specification does not define the handling of the `acp-node-name` field because the rules do not depend on BRSKI but apply equally to any protocols/mechanisms an ACP registrar may use.

#### 6.11.7.2. Unique Address/Prefix allocation

ACP registrars MUST NOT allocate ACP address prefixes to ACP nodes via the `acp-node-name` that would collide with the ACP address prefixes of other ACP nodes in the same ACP domain. This includes both prefixes allocated by the same ACP registrar to different ACP nodes as well as prefixes allocated by other ACP registrars for the same ACP domain.

To support such unique address allocation, an ACP registrar MUST have one or more 46-bit identifiers unique across the ACP domain which is called the Registrar-ID. Allocation of Registrar-ID(s) to an ACP registrar can happen through OAM mechanisms in conjunction with some database / allocation orchestration.

ACP registrars running on physical devices with known globally unique EUI-48 MAC address(es) can use the lower 46 bits of those address(es) as unique Registrar-IDs without requiring any external signaling/configuration (the upper two bits, V and U are not uniquely assigned but functional). This approach is attractive for distributed, non-centrally administered, lightweight ACP registrar implementations. There is no mechanism to deduce from a MAC address itself whether it is actually uniquely assigned. Implementations need to consult additional offline information before making this assumption. For example by knowing that a particular physical product/MIC-chip is guaranteed to use globally unique assigned EUI-48 MAC address(es).

When the candidate ACP device (called Pledge in BRSKI) is to be enrolled into an ACP domain, the ACP registrar needs to allocate a unique ACP address to the node and ensure that the ACP certificate gets a `acp-node-name` field (Section 6.2.2) with the appropriate information - ACP domain-name, ACP-address, and so on. If the ACP registrar uses BRSKI, it signals the ACP `acp-node-name` field to the Pledge via the `EST /csrattrs` command (see [I-D.ietf-anima-bootstrapping-keyinfra], section 5.9.2 - "EST CSR Attributes").

[RFC-Editor: please update reference to section 5.9.2 accordingly with latest BRSKI draft at time of publishing, or RFC]

#### 6.11.7.3. Addressing Sub-Scheme Policies

The ACP registrar selects for the candidate ACP node a unique address prefix from an appropriate ACP addressing sub-scheme, either a zone addressing sub-scheme prefix (see Section 6.11.3), or a Vlong addressing sub-scheme prefix (see Section 6.11.5). The assigned ACP address prefix encoded in the `acp-node-name` field of the ACP certificate indicates to the ACP node its ACP address information.

The sub-addressing scheme indicates the prefix length: /127 for zone address sub-scheme, /120 or /112 for Vlong address sub-scheme. The first address of the prefix is the ACP address. All other addresses in the prefix are for other uses by the ACP node as described in the zone and Vlong addressing sub scheme sections. The ACP address prefix itself is then signaled by the ACP node into the ACP routing protocol (see Section 6.12) to establish IPv6 reachability across the ACP.

The choice of addressing sub-scheme and prefix-length in the Vlong address sub-scheme is subject to ACP registrar policy. It could be an ACP domain wide policy, or a per ACP node or per ACP node type policy. For example, in BRSKI, the ACP registrar is aware of the IDevID certificate of the candidate ACP node, which typically contains a "serialNumber" attribute in the subject field distinguished name encoding that is often indicating the node's vendor and device type and can be used to drive a policy selecting an appropriate addressing sub-scheme for the (class of) node(s).

ACP registrars SHOULD default to allocate ACP zone sub-address scheme addresses with Zone-ID 0.

ACP registrars that are aware of the IDevID certificate of a candidate ACP device SHOULD be able to choose the zone vs. Vlong sub-address scheme for ACP nodes based on the [X.520] "serialNumber" attribute in the subject field distinguished name encoding of the IDevID certificate, for example by the PID (Product Identifier) part which identifies the product type, or the complete "serialNumber". The PID for example could identify nodes that allow for specialized ASA requiring multiple addresses or non-autonomic VMs for services and those nodes could receive Vlong sub-address scheme ACP addresses.

In a simple allocation scheme, an ACP registrar remembers persistently across reboots its currently used Registrar-ID and for each addressing scheme (Zone with Zone-ID 0, Vlong with /112, Vlong with /120), the next Node-Number available for allocation and increases it during successful enrollment to an ACP node. In this simple allocation scheme, the ACP registrar would not recycle ACP address prefixes from no longer used ACP nodes.

If allocated addresses cannot be remembered by registrars, then it is necessary to either use a new value for the Register-ID field in the ACP addresses, or determine allocated ACP addresses from determining the addresses of reachable ACP nodes, which is not necessarily the set of all ACP nodes. Non-tracked ACP addresses can be reclaimed by revoking or not renewing their certificates and instead handing out new certificate with new addresses (for example with a new Registrar-ID value). Note that such strategies may require coordination amongst registrars.

#### 6.11.7.4. Address/Prefix Persistence

When an ACP certificate is renewed or rekeyed via EST or other mechanisms, the ACP address/prefix in the `acp-node-name` field MUST be maintained unless security issues or violations of the unique address assignment requirements exist or are suspected by the ACP registrar.

ACP address information SHOULD be maintained even when the renewing/rekeying ACP registrar is not the same as the one that enrolled the prior ACP certificate. See Section 9.2.4 for an example.

ACP address information SHOULD also be maintained even after an ACP certificate did expire or failed. See Section 6.2.5.5 and Section 6.2.5.6.

#### 6.11.7.5. Further Details

Section 9.2 discusses further informative details of ACP registrars: What interactions registrars need, what parameters they require, certificate renewal and limitations, use of sub-CAs on registrars and centralized policy control.

#### 6.12. Routing in the ACP

Once ULA address are set up all autonomic entities should run a routing protocol within the autonomic control plane context. This routing protocol distributes the ULA created in the previous section for reachability. The use of the autonomic control plane specific context eliminates the probable clash with Data-Plane routing tables and also secures the ACP from interference from the configuration mismatch or incorrect routing updates.

The establishment of the routing plane and its parameters are automatic and strictly within the confines of the autonomic control plane. Therefore, no explicit configuration is required.

All routing updates are automatically secured in transit as the channels of the ACP are encrypted, and this routing runs only inside the ACP.

The routing protocol inside the ACP is RPL ([RFC6550]). See Appendix A.4 for more details on the choice of RPL.

RPL adjacencies are set up across all ACP channels in the same domain including all its routing subdomains. See Appendix A.6 for more details.

#### 6.12.1. ACP RPL Profile

The following is a description of the RPL profile that ACP nodes need to support by default. The format of this section is derived from [I-D.ietf-roll-applicability-template].

##### 6.12.1.1. Overview

RPL Packet Information (RPI) defined in [RFC6550], section 11.2 defines the data packet artefacts required or beneficial in forwarding of packets routed by RPL. This profile does not use RPI for better compatibility with accelerated hardware forwarding planes which most often does not support the Hop-by-Hop headers used for RPI, but also to avoid the overhead of the RPI header on the wire and cost of adding/removing them.

##### 6.12.1.1.1. Single Instance

To avoid the need for RPI, the ACP RPL profile uses a simple destination prefix based routing/forwarding table. To achieve this, the profiles uses only one RPL instanceID. This single instanceID can contain only one Destination Oriented Directed Acyclic Graph (DODAG), and the routing/forwarding table can therefore only calculate a single class of service ("best effort towards the primary NOC/root") and cannot create optimized routing paths to accomplish latency or energy goals between any two nodes.

This choice is a compromise. Consider a network that has multiple NOCs in different locations. Only one NOC will become the DODAG root. Traffic to and from other NOCs has to be sent through the DODAG (shortest path tree) rooted in the primary NOC. Depending on topology, this can be an annoyance from a latency point of view or from minimizing network path resources, but this is deemed to be acceptable given how ACP traffic is "only" network management/control traffic. See Appendix A.9.4 for more details.

Using a single instanceID/DODAG does not introduce a single point of failure, as the DODAG will reconfigure itself when it detects Data-Plane forwarding failures including choosing a different root when the primary one fails.

The benefit of this profile, especially compared to other IGPs is that it does not calculate routes for node reachable through the same interface as the DODAG root. This RPL profile can therefore scale to much larger number of ACP nodes in the same amount of compute and memory than other routing protocols. Especially on nodes that are leafs of the topology or those close to those leafs.

#### 6.12.1.1.2. Reconvergence

In RPL profiles where RPL Packet Information (RPI, see Section 6.12.1.13) is present, it is also used to trigger reconvergence when misrouted, for example looping, packets are recognized because of their RPI data. This helps to minimize RPL signaling traffic especially in networks without stable topology and slow links.

The ACP RPL profile instead relies on quick reconverging the DODAG by recognizing link state change (down/up) and triggering reconvergence signaling as described in Section 6.12.1.7. Since links in the ACP are assumed to be mostly reliable (or have link layer protection against loss) and because there is no stretch according to Section 6.12.1.7, loops caused by loss of RPL routing protocol signaling packets should be exceedingly rare.

In addition, there are a variety of mechanisms possible in RPL to further avoid temporary loops RECOMMENDED to be used for the ACPL RPL profile: DODAG Information Objects (DIOs) SHOULD be sent 2 or 3 times to inform children when losing the last parent. The technique in [RFC6550] section 8.2.2.6. (Detaching) SHOULD be favored over that in section 8.2.2.5., (Poisoning) because it allows local connectivity. Nodes SHOULD select more than one parent, at least 3 if possible, and send Destination Advertisement Objects (DAO)s to all of them in parallel.

Additionally, failed ACP tunnels can be quickly discovered through the secure channel protocol mechanisms such as IKEv2 Dead Peer Detection. This can function as a replacement for a Low-power and Lossy Networks' (LLN's) Expected Transmission Count (ETX) feature that is not used in this profile. A failure of an ACP tunnel should immediately signal the RPL control plane to pick a different parent.

#### 6.12.1.2. RPL Instances

Single RPL instance. Default RPLInstanceID = 0.

#### 6.12.1.3. Storing vs. Non-Storing Mode

RPL Mode of Operations (MOP): MUST support mode 2 - "Storing Mode of Operations with no multicast support". Implementations MAY support mode 3 ("... with multicast support" as that is a superset of mode 2). Note: Root indicates mode in DIO flow.

#### 6.12.1.4. DAO Policy

Proactive, aggressive DAO state maintenance:

- \* Use K-flag in unsolicited DAO indicating change from previous information (to require DAO-ACK).
- \* Retry such DAO DAO-RETRIES(3) times with DAO- ACK\_TIME\_OUT(256ms) in between.

#### 6.12.1.5. Path Metric

Use Hopcount according to [RFC6551]. Note that this is solely for diagnostic purposes as it is not used by the objective function.

#### 6.12.1.6. Objective Function

Objective Function (OF): Use OF0 [RFC6552]. No use of metric containers.

rank\_factor: Derived from link speed: <= 100Mbps:  
LOW\_SPEED\_FACTOR(5), else HIGH\_SPEED\_FACTOR(1)

This is a simple rank differentiation between typical "low speed" or "IoT" links that commonly max out at 100 Mbps and typical infrastructure links with speeds of 1 Gbps or higher. Given how the path selection for the ACP focusses only on reachability but not on path cost optimization, no attempts at finer grained path optimization are made.

#### 6.12.1.7. DODAG Repair

Global Repair: we assume stable links and ranks (metrics), so there is no need to periodically rebuild the DODAG. The DODAG version is only incremented under catastrophic events (e.g., administrative action).

Local Repair: As soon as link breakage is detected, the ACP node send No-Path DAO for all the targets that were reachable only via this link. As soon as link repair is detected, the ACP node validates if this link provides a better parent. If so, a new rank is computed by the ACP node and it sends new DIO that advertise the new rank. Then it sends a DAO with a new path sequence about itself.

When using ACP multi-access virtual interfaces, local repair can be triggered directly by peer breakage, see Section 6.13.5.2.2.

stretch\_rank: none provided ("not stretched").

Data Path Validation: Not used.

Trickle: Not used.

#### 6.12.1.8. Multicast

Not used yet but possible because of the selected mode of operations.

#### 6.12.1.9. Security

[RFC6550] security not used, substituted by ACP security.

Because the ACP links already include provisions for confidentiality and integrity protection, their usage at the RPL layer would be redundant, and so RPL security is not used.

#### 6.12.1.10. P2P communications

Not used.

#### 6.12.1.11. IPv6 address configuration

Every ACP node (RPL node) announces an IPv6 prefix covering the addresses assigned to the ACP node via the AcpNodeName. The prefix length depends on the addressing sub-scheme of the acp-address, /127 for Zone Addressing Sub-Scheme and /112 or /120 for Vlong addressing sub-scheme. See Section 6.11 for more details.

Every ACP node MUST install a black hole (aka null) route if there are unused parts of the ACP address space assigned to the ACP node via its `AcpNextName`. This is superseded by longer prefixes assigned to interfaces for the address space actually used by the node. For example, when the node has an ACP-VLong-8 address space, it installs a /120 black hole route. If it then for example only uses the ACP address (first address from the space), it would assign that address via a /128 address prefix to the ACP loopback interface (see Section 6.13.5.1). None of those longer prefixes are announced into RPL.

For ACP-Manual address prefixes configured on an ACP node, for example for ACP connect subnets (see Section 8.1.1), the node announces the /64 subnet prefix.

#### 6.12.1.12. Administrative parameters

Administrative Preference ([RFC6550], 3.2.6 - to become root):  
Indicated in `DODAGPreference` field of DIO message.

- \* Explicit configured "root": 0b100
- \* ACP registrar (Default): 0b011
- \* ACP-connect (non-registrar): 0b010
- \* Default: 0b001.

#### 6.12.1.13. RPL Packet Information

RPI is not required in the ACP RPL profile for the following reasons.

One RPI option is the RPL Source Routing Header (SRH) [RFC6554] which is not necessary because the ACP RPL profile uses storing mode where each hop has the necessary next-hop forwarding information.

The simpler RPL Option header [RFC6553] is also not necessary in this profile, because it uses a single RPL instance and data path validation is also not used.

#### 6.12.1.14. Unknown Destinations

Because RPL minimizes the size of the routing and forwarding table, prefixes reachable through the same interface as the RPL root are not known on every ACP node. Therefore, traffic to unknown destination addresses can only be discovered at the RPL root. The RPL root SHOULD have attach safe mechanisms to operationally discover and log such packets.

As this requirement places additional constraints on the Data-Plane functionality of the RPL root, it does not apply to "normal" nodes that are not configured to have special functionality (i.e., the administrative parameter from Section 6.12.1.12 has value 0b001). If the ACP network is degraded to the point where there are no nodes that could be configured as root, registrar, or ACP-connect nodes, it is possible that the RPL root (and thus the ACP as a whole) would be unable to detect traffic to unknown destinations. However, in the absence of nodes with administrative preference other than 0b001, there is also unlikely to be a way to get diagnostic information out of the ACP, so detection of traffic to unknown destinations would not be actionable anyway.

### 6.13. General ACP Considerations

Since channels are by default established between adjacent neighbors, the resulting overlay network does hop-by-hop encryption. Each node decrypts incoming traffic from the ACP, and encrypts outgoing traffic to its neighbors in the ACP. Routing is discussed in Section 6.12.

#### 6.13.1. Performance

There are no performance requirements against ACP implementations defined in this document because the performance requirements depend on the intended use case. It is expected that full autonomic node with a wide range of ASA can require high forwarding plane performance in the ACP, for example for telemetry. Implementations of ACP to solely support traditional/SDN style use cases can benefit from ACP at lower performance, especially if the ACP is used only for critical operations, e.g., when the Data-Plane is not available. The design of the ACP as specified in this document is intended to support a wide range of performance options: It is intended to allow software-only implementations at potentially low performance, but can also support high performance options. See [RFC8368] for more details.

#### 6.13.2. Addressing of Secure Channels

In order to be independent of the Data-Plane routing and addressing, the GRASP discovered ACP secure channels use IPv6 link local addresses between adjacent neighbors. Note: Section 8.2 specifies extensions in which secure channels are configured tunnels operating over the Data-Plane, so those secure channels cannot be independent of the Data-Plane.

To avoid that Data-Plane configuration can impact the operations of the IPv6 (link-local) interface/address used for ACP channels, appropriate implementation considerations are required. If the IPv6

interface/link-local address is shared with the Data-Plane, it needs to be impossible to unconfigure/disable it through configuration. Instead of sharing the IPv6 interface/link-local address, a separate (virtual) interface with a separate IPv6 link-local address can be used. For example, the ACP interface could be run over a separate MAC address of an underlying L2 (Ethernet) interface. For more details and options, see Appendix A.9.2.

Note that other (non-ideal) implementation choices may introduce additional undesired dependencies against the Data-Plane. For example, shared code and configuration of the secure channel protocols (IPsec / DTLS).

#### 6.13.3. MTU

The MTU for ACP secure channels MUST be derived locally from the underlying link MTU minus the secure channel encapsulation overhead.

ACP secure Channel protocols do not need to perform MTU discovery because they are built across L2 adjacencies - the MTU on both sides connecting to the L2 connection are assumed to be consistent. Extensions to ACP where the ACP is for example tunneled need to consider how to guarantee MTU consistency. This is an issue of tunnels, not an issue of running the ACP across a tunnel. Transport stacks running across ACP can perform normal PMTUD (Path MTU Discovery). Because the ACP is meant to prioritize reliability over performance, they MAY opt to only expect IPv6 minimum MTU (1280) to avoid running into PMTUD implementation bugs or underlying link MTU mismatch problems.

#### 6.13.4. Multiple links between nodes

If two nodes are connected via several links, the ACP SHOULD be established across every link, but it is possible to establish the ACP only on a sub-set of links. Having an ACP channel on every link has a number of advantages, for example it allows for a faster failover in case of link failure, and it reflects the physical topology more closely. Using a subset of links (for example, a single link), reduces resource consumption on the node, because state needs to be kept per ACP channel. The negotiation scheme explained in Section 6.6 allows the Decider (the node with the higher ACP address) to drop all but the desired ACP channels to the Follower - and the Follower will not re-try to build these secure channels from its side unless the Decider shows up with a previously unknown GRASP announcement (e.g., on a different link or with a different address announced in GRASP).

### 6.13.5. ACP interfaces

The ACP VRF has conceptually two type of interfaces: The "ACP Loopback interface(s)" to which the ACP ULA address(es) are assigned and the "ACP virtual interfaces" that are mapped to the ACP secure channels.

#### 6.13.5.1. ACP loopback interfaces

For autonomous operations of the ACP, as described in Section 6 and Section 7, the ACP node uses the first address from the N bit ACP prefix ( $N = 128 - \text{number of Vbits of the ACP address}$ ) assigned to the node. This address is assigned with an address prefix of N or larger to a loopback interface.

Other addresses from the prefix can be used by the ACP of the node as desired. The autonomous operations of the ACP does not require additional global scope IPv6 addresses, they are instead intended for ASA or non-autonomous functions. Non fully autonomic components of the ACP such as ACP connect interfaces (see Figure 16) may also introduce additional global scope IPv6 addresses on other types of interfaces into the ACP.

[RFC-Editor: please remove this paragraph: Note to reviewers: Please do not complain again about an obsolete RFC number in the following paragraph. The text should make it clear that the reference was chosen to indicate a particular point in time, but not to recommend/use a particularly obsolete protocol spec.]

The use of loopback interfaces for global scope addresses is common operational configuration practice on routers, for example in IBGP connections since BGP4 (see [RFC1654]) or earlier. The ACP adopts and automates this operational practice.

A loopback interface for use with the ACP as described above is an interface behaving according to [RFC6724] Section 4., paragraph 2: Packets sent by the host of the node from the loopback interface behave as if they are looped back by the interface so that they look as if they originated from the loopback interface, are then received by the node and forwarded by it towards the destination.

The word loopback only indicates this behavior, but not the actual name of the interface type chosen in an actual implementation. A loopback interface for use with the ACP can be a virtual/software construct without any associated hardware, or it can be a hardware interface operating in loopback mode.

A loopback interface used for the ACP MUST NOT have connectivity to other nodes.

The following reviews the reasons for the choice of loopback addresses for ACP addresses is based on the IPv6 address architecture and common challenges:

1. IPv6 addresses are assigned to interfaces, not nodes. IPv6 continues the IPv4 model that a subnet prefix is associated with one link, see [RFC4291], Section 2.1.
2. IPv6 implementations commonly do not allow assignment of the same IPv6 global scope address in the same VRF to more than one interface.
3. Global scope addresses assigned to interfaces that are connecting to other nodes (external interfaces) may not be stable addresses for communications because any such interface could fail due to reasons external to the node. This could render the addresses assigned to that interface unusable.
4. If failure of the subnet does not result in bringing down the interface and making the addresses unusable, it could result in unreachability of the address because the shortest path to the node might go through one of the other nodes on the same subnet which could equally consider the subnet to be operational even though it is not.
5. Many OAM service implementations on routers cannot deal with more than one peer address, often because they do already expect that a single loopback address can be used, especially to provide a stable address under failure of external interfaces or links.
6. Even when an application supports multiple addresses to a peer, it can only use one address for a connection at a time with the most widely deployed transport protocols TCP and UDP. While [RFC6824] solves this problem, it is not widely adopted for router OAM services implementations.
7. To completely autonomously assign global scope addresses to subnets connecting to other nodes, it would be necessary for every node to have an amount of prefix address space in the order of the maximum number of subnets that the node could connect to and then the node would have to negotiate with adjacent nodes across those subnets whose address space to use for each subnet.
8. Using global scope addresses for subnets between nodes is unnecessary if those subnets only connect routers, such as ACP secure channels, because they can communicate to remote nodes via their global scope loopback addresses. Using global scope addresses for those external subnets is therefore wasteful for the address space and also unnecessarily increasing the size of routing and forwarding tables, which especially for the ACP is highly undesirable because it should attempt to minimize the per-node overhead of the ACP VRF.

9. For all these reasons, the ACP addressing schemes do not consider ACP addresses for subnets connecting ACP nodes.

Note that [RFC8402] introduces the term Node-SID to refer to IGP prefix segments that identify a specific router, for example on a loopback interface. An ACP loopback address prefix may similarly be called an ACP Node Identifier.

#### 6.13.5.2. ACP virtual interfaces

Any ACP secure channel to another ACP node is mapped to ACP virtual interfaces in one of the following ways. This is independent of the chosen secure channel protocol (IPsec, DTLS or other future protocol - standards or non-standards).

Note that all the considerations described here are assuming point-to-point secure channel associations. Mapping multi-party secure channel associations such as [RFC6407] is out of scope.

##### 6.13.5.2.1. ACP point-to-point virtual interfaces

In this option, each ACP secure channel is mapped into a separate point-to-point ACP virtual interface. If a physical subnet has more than two ACP capable nodes (in the same domain), this implementation approach will lead to a full mesh of ACP virtual interfaces between them.

When the secure channel protocol determines a peer to be dead, this SHOULD result in indicating link breakage to trigger RPL DODAG repair, see Section 6.12.1.7.

##### 6.13.5.2.2. ACP multi-access virtual interfaces

In a more advanced implementation approach, the ACP will construct a single multi-access ACP virtual interface for all ACP secure channels to ACP capable nodes reachable across the same underlying (physical) subnet. IPv6 link-local multicast packets sent into an ACP multi-access virtual interface are replicated to every ACP secure channel mapped into the ACP multicast-access virtual interface. IPv6 unicast packets sent into an ACP multi-access virtual interface are sent to the ACP secure channel that belongs to the ACP neighbor that is the next-hop in the ACP forwarding table entry used to reach the packets destination address.

When the secure channel protocol determines a peer to be dead for a secure channel mapped into an ACP multi-access virtual interface, this SHOULD result in signaling breakage of that peer to RPL, so it can trigger RPL DODAG repair, see Section 6.12.1.7.

There is no requirement for all ACP nodes on the same multi-access subnet to use the same type of ACP virtual interface. This is purely a node local decision.

ACP nodes MUST perform standard IPv6 operations across ACP virtual interfaces including SLAAC (Stateless Address Auto-Configuration) - [RFC4862]) to assign their IPv6 link local address on the ACP virtual interface and ND (Neighbor Discovery - [RFC4861]) to discover which IPv6 link-local neighbor address belongs to which ACP secure channel mapped to the ACP virtual interface. This is independent of whether the ACP virtual interface is point-to-point or multi-access.

"Optimistic Duplicate Address Detection (DAD)" according to [RFC4429] is RECOMMENDED because the likelihood for duplicates between ACP nodes is highly improbable as long as the address can be formed from a globally unique local assigned identifier (e.g., EUI-48/EUI-64, see below).

ACP nodes MAY reduce the amount of link-local IPv6 multicast packets from ND by learning the IPv6 link-local neighbor address to ACP secure channel mapping from other messages such as the source address of IPv6 link-local multicast RPL messages - and therefore forego the need to send Neighbor Solicitation messages.

The ACP virtual interface IPv6 link local address can be derived from any appropriate local mechanism such as node local EUI-48 or EUI-64 ("EUI" stands for "Extended Unique Identifier"). It MUST NOT depend on something that is attackable from the Data-Plane such as the IPv6 link-local address of the underlying physical interface, which can be attacked by SLAAC, or parameters of the secure channel encapsulation header that may not be protected by the secure channel mechanism.

The link-layer address of an ACP virtual interface is the address used for the underlying interface across which the secure tunnels are built, typically Ethernet addresses. Because unicast IPv6 packets sent to an ACP virtual interface are not sent to a link-layer destination address but rather an ACP secure channel, the link-layer address fields SHOULD be ignored on reception and instead the ACP secure channel from which the message was received should be remembered.

Multi-access ACP virtual interfaces are preferable implementations when the underlying interface is a (broadcast) multi-access subnet because they do reflect the presence of the underlying multi-access subnet into the virtual interfaces of the ACP. This makes it for example simpler to build services with topology awareness inside the ACP VRF in the same way as they could have been built running natively on the multi-access interfaces.

Consider also the impact of point-to-point vs. multi-access virtual interface on the efficiency of flooding via link local multicasted messages:

Assume a LAN with three ACP neighbors, Alice, Bob and Carol. Alice's ACP GRASP wants to send a link-local GRASP multicast message to Bob and Carol. If Alice's ACP emulates the LAN as per-peer, point-to-point virtual interfaces, one to Bob and one to Carol, Alice's ACP GRASP will send two copies of multicast GRASP messages: One to Bob and one to Carol. If Alice's ACP emulates a LAN via a multipoint virtual interface, Alice's ACP GRASP will send one packet to that interface and the ACP multipoint virtual interface will replicate the packet to each secure channel, one to Bob, one to Carol. The result is the same. The difference happens when Bob and Carol receive their packet. If they use ACP point-to-point virtual interfaces, their GRASP instance would forward the packet from Alice to each other as part of the GRASP flooding procedure. These packets are unnecessary and would be discarded by GRASP on receipt as duplicates (by use of the GRASP Session ID). If Bob and Carol's ACP would emulate a multi-access virtual interface, then this would not happen, because GRASP's flooding procedure does not replicate back packets to the interface that they were received from.

Note that link-local GRASP multicast messages are not sent directly as IPv6 link-local multicast UDP messages into ACP virtual interfaces, but instead into ACP GRASP virtual interfaces, that are layered on top of ACP virtual interfaces to add TCP reliability to link-local multicast GRASP messages. Nevertheless, these ACP GRASP virtual interfaces perform the same replication of message and, therefore, result in the same impact on flooding. See Section 6.9.2 for more details.

RPL does support operations and correct routing table construction across non-broadcast multi-access (NBMA) subnets. This is common when using many radio technologies. When such NBMA subnets are used, they MUST NOT be represented as ACP multi-access virtual interfaces because the replication of IPv6 link-local multicast messages will not reach all NBMA subnet neighbors. In result, GRASP message flooding would fail. Instead, each ACP secure channel across such an interface MUST be represented as a ACP point-to-point virtual interface. See also Appendix A.9.4.

Care needs to be taken when creating multi-access ACP virtual interfaces across ACP secure channels between ACP nodes in different domains or routing subdomains. If for example future inter-domain ACP policies are defined as "peer-to-peer" policies, it is easier to create ACP point-to-point virtual interfaces for these inter-domain secure channels.

## 7. ACP support on L2 switches/ports (Normative)

### 7.1. Why (Benefits of ACP on L2 switches)

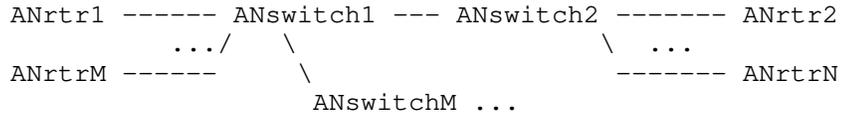


Figure 15: Topology with L2 ACP switches

Consider a large L2 LAN with ANrtr1...ANrtrN connected via some topology of L2 switches. Examples include large enterprise campus networks with an L2 core, IoT networks or broadband aggregation networks which often have even a multi-level L2 switched topology.

If the discovery protocol used for the ACP is operating at the subnet level, every ACP router will see all other ACP routers on the LAN as neighbors and a full mesh of ACP channels will be built. If some or all of the AN switches are autonomic with the same discovery protocol, then the full mesh would include those switches as well.

A full mesh of ACP connections can create fundamental scale challenges. The number of security associations of the secure channel protocols will likely not scale arbitrarily, especially when they leverage platform accelerated encryption/decryption. Likewise, any other ACP operations (such as routing) needs to scale to the number of direct ACP neighbors. An ACP router with just 4 physical interfaces might be deployed into a LAN with hundreds of neighbors connected via switches. Introducing such a new unpredictable scaling factor requirement makes it harder to support the ACP on arbitrary platforms and in arbitrary deployments.

Predictable scaling requirements for ACP neighbors can most easily be achieved if in topologies such as these, ACP capable L2 switches can ensure that discovery messages terminate on them so that neighboring ACP routers and switches will only find the physically connected ACP L2 switches as their candidate ACP neighbors. With such a discovery mechanism in place, the ACP and its security associations will only need to scale to the number of physical interfaces instead of a potentially much larger number of "LAN-connected" neighbors. And the ACP topology will follow directly the physical topology, something which can then also be leveraged in management operations or by ASAs.

In the example above, consider ANswitch1 and ANswitchM are ACP capable, and ANswitch2 is not ACP capable. The desired ACP topology is that ANrtr1 and ANrtrM only have an ACP connection to ANswitch1, and that ANswitch1, ANrtr2, ANrtrN have a full mesh of ACP connection

amongst each other. ANswitch1 also has an ACP connection with ANswitchM and ANswitchM has ACP connections to anything else behind it.

## 7.2. How (per L2 port DULL GRASP)

To support ACP on L2 switches or L2 switched ports of an L3 device, it is necessary to make those L2 ports look like L3 interfaces for the ACP implementation. This primarily involves the creation of a separate DULL GRASP instance/domain on every such L2 port. Because GRASP has a dedicated link-local IPv6 multicast address (ALL\_GRASP\_NEIGHBORS), it is sufficient that all packets for this address are being extracted at the port level and passed to that DULL GRASP instance. Likewise the IPv6 link-local multicast packets sent by that DULL GRASP instance need to be sent only towards the L2 port for this DULL GRASP instance (instead of being flooded across all ports of the VLAN to which the port belongs).

When Ports/Interfaces across which the ACP is expected to operate in an ACP-aware L2-switch or L2/L3-switch/router are L2-bridged, packets for the ALL\_GRASP\_NEIGHBORS multicast address MUST never be forward between these ports. If MLD snooping is used, it MUST be prohibited from bridging packets for the ALL\_GRASP\_NEIGHBORS IPv6 multicast address.

On hybrid L2/L3 switches, multiple L2 ports are assigned to a single L3 VLAN interface. With the aforementioned changes for DULL GRASP, ACP can simply operate on the L3 VLAN interfaces, so no further (hardware) forwarding changes are required to make ACP operate on L2 ports. This is possible because the ACP secure channel protocols only use link-local IPv6 unicast packets, and these packets will be sent to the correct L2 port towards the peer by the VLAN logic of the device.

This is sufficient when p2p ACP virtual interfaces are established to every ACP peer. When it is desired to create multi-access ACP virtual interfaces (see Section 6.13.5.2.2), it is REQUIRED not to coalesce all the ACP secure channels on the same L3 VLAN interface, but only all those on the same L2 port.

If VLAN tagging is used, then all the above described logic only applies to untagged GRASP packets. For the purpose of ACP neighbor discovery via GRASP, no VLAN tagged packets SHOULD be sent or received. In a hybrid L2/L3 switch, each VLAN would therefore only create ACP adjacencies across those ports where the VLAN is carried untagged.

In result, the simple logic is that ACP secure channels would operate over the same L3 interfaces that present a single flat bridged network across all routers, but because DULL GRASP is separated on a per-port basis, no full mesh of ACP secure channels is created, but only per-port ACP secure channels to per-port L2-adjacent ACP node neighbors.

For example, in the above picture, ANswitch1 would run separate DULL GRASP instances on its ports to ANrtr1, ANswitch2 and ANswitchI, even though all those three ports may be in the data plane in the same (V)LAN and perform L2 switching between these ports, ANswitch1 would perform ACP L3 routing between them.

The description in the previous paragraph was specifically meant to illustrate that on hybrid L3/L2 devices that are common in enterprise, IoT and broadband aggregation, there is only the GRASP packet extraction (by Ethernet address) and GRASP link-local multicast per L2-port packet injection that has to consider L2 ports at the hardware forwarding level. The remaining operations are purely ACP control plane and setup of secure channels across the L3 interface. This hopefully makes support for per-L2 port ACP on those hybrid devices easy.

In devices without such a mix of L2 port/interfaces and L3 interfaces (to terminate any transport layer connections), implementation details will differ. Logically most simply every L2 port is considered and used as a separate L3 subnet for all ACP operations. The fact that the ACP only requires IPv6 link-local unicast and multicast should make support for it on any type of L2 devices as simple as possible.

A generic issue with ACP in L2 switched networks is the interaction with the Spanning Tree Protocol. Without further L2 enhancements, the ACP would run only across the active STP topology and the ACP would be interrupted and re-converge with STP changes. Ideally, ACP peering SHOULD be built also across ports that are blocked in STP so that the ACP does not depend on STP and can continue to run unaffected across STP topology changes, where re-convergence can be quite slow. The above described simple implementation options are not sufficient to achieve this.

## 8. Support for Non-ACP Components (Normative)

### 8.1. ACP Connect



ACP connect has security consequences: All systems and processes connected via ACP connect have access to all ACP nodes on the entire ACP, without further authentication. Thus, the ACP connect interface and NOC systems connected to it needs to be physically controlled/secured. For this reason the mechanisms described here do explicitly not include options to allow for a non-ACP router to be connected across an ACP connect interface and addresses behind such a router routed inside the ACP.

Physical controlled/secured means that attackers can gain no access to the physical device hosting the ACP Edge Node, the physical interfaces and links providing the ACP connect link nor the physical devices hosting the NOC Device. In a simple case, ACP Edge node and NOC Device are co-located in an access controlled room, such as a NOC, to which attackers cannot gain physical access.

An ACP connect interface provides exclusively access to only the ACP. This is likely insufficient for many NMS hosts. Instead, they would require a second "Data-Plane" interface outside the ACP for connections between the NMS host and administrators, or Internet based services, or for direct access to the Data-Plane. The document "Using Autonomic Control Plane for Stable Connectivity of Network OAM" [RFC8368] explains in more detail how the ACP can be integrated in a mixed NOC environment.

An ACP connect interface SHOULD use an IPv6 address/prefix from the ACP Manual Addressing Sub-Scheme (Section 6.11.4), letting the operator configure for example only the Subnet-ID and having the node automatically assign the remaining part of the prefix/address. It SHOULD NOT use a prefix that is also routed outside the ACP so that the addresses clearly indicate whether it is used inside the ACP or not.

The prefix of ACP connect subnets MUST be distributed by the ACP edge node into the ACP routing protocol RPL. The NMS hosts MUST connect to prefixes in the ACP routing table via its ACP connect interface. In the simple case where the ACP uses only one ULA prefix and all ACP connect subnets have prefixes covered by that ULA prefix, NMS hosts can rely on [RFC6724] to determine longest match prefix routes towards its different interfaces, ACP and Data-Plane. With RFC6724, The NMS host will select the ACP connect interface for all addresses in the ACP because any ACP destination address is longest matched by the address on the ACP connect interface. If the NMS hosts ACP connect interface uses another prefix or if the ACP uses multiple ULA prefixes, then the NMS hosts require (static) routes towards the ACP interface for these prefixes.

When an ACP Edge node receives a packet from an ACP connect interface, the ACP Edge node MUST only forward the packet into the ACP if the packet has an IPv6 source address from that interface (this is sometimes called "RPF filtering"). This filtering rule MAY be changed through administrative measures. The more any such administrative action enable reachability of non ACP nodes to the ACP, the more this may cause security issues.

To limit the security impact of ACP connect, nodes supporting it SHOULD implement a security mechanism to allow configuration/use of ACP connect interfaces only on nodes explicitly targeted to be deployed with it (those in physically secure locations such as a NOC). For example, the registrar could disable the ability to enable ACP connect on devices during enrollment and that property could only be changed through re-enrollment. See also Appendix A.9.5.

ACP Edge nodes SHOULD have a configurable option to prohibit packets with RPI headers (see Section 6.12.1.13 across an ACP connect interface. These headers are outside the scope of the RPL profile in this specification but may be used in future extensions of this specification.

#### 8.1.2. Software Components

The previous section assumed that ACP Edge node and NOC devices are separate physical devices and the ACP connect interface is a physical network connection. This section discusses the implication when these components are instead software components running on a single physical device.

The ACP connect mechanism cannot only be used to connect physically external systems (NMS hosts) to the ACP but also other applications, containers or virtual machines. In fact, one possible way to eliminate the security issue of the external ACP connect interface is to collocate an ACP edge node and an NMS host by making one a virtual machine or container inside the other; and therefore converting the unprotected external ACP subnet into an internal virtual subnet in a single device. This would ultimately result in a fully ACP enabled NMS host with minimum impact to the NMS hosts software architecture. This approach is not limited to NMS hosts but could equally be applied to devices consisting of one or more VNF (virtual network functions): An internal virtual subnet connecting out-of-band management interfaces of the VNFs to an ACP edge router VNF.

The core requirement is that the software components need to have a network stack that permits access to the ACP and optionally also the Data-Plane. Like in the physical setup for NMS hosts this can be realized via two internal virtual subnets. One that is connecting to the ACP (which could be a container or virtual machine by itself), and one (or more) connecting into the Data-Plane.

This "internal" use of ACP connect approach should not be considered to be a "workaround" because in this case it is possible to build a correct security model: It is not necessary to rely on unprovable external physical security mechanisms as in the case of external NMS hosts. Instead, the orchestration of the ACP, the virtual subnets and the software components can be done by trusted software that could be considered to be part of the ANI (or even an extended ACP). This software component is responsible for ensuring that only trusted software components will get access to that virtual subnet and that only even more trusted software components will get access to both the ACP virtual subnet and the Data-Plane (because those ACP users could leak traffic between ACP and Data-Plane). This trust could be established for example through cryptographic means such as signed software packages.

#### 8.1.3. Auto Configuration

ACP edge nodes, NMS hosts and software components that as described in the previous section are meant to be composed via virtual interfaces SHOULD support on the ACP connect subnet StateLess Address Autoconfiguration (SLAAC - [RFC4862]) and route auto configuration according to [RFC4191].

The ACP edge node acts as the router towards the ACP on the ACP connect subnet, providing the (auto-)configured prefix for the ACP connect subnet and (auto-)configured routes into the ACP to NMS hosts and/or software components.

The ACP edge node uses the Route Information Option (RIO) of RFC4191 to announce aggregated prefixes for address prefixes used in the ACP (with normal RIO lifetimes. In addition, the ACP edge node also uses a RIO to announce the default route (::/0) with a lifetime of 0.

These RIOs allow to connect Type C hosts to the ACP via an ACP connect subnet on one interface and another network (Data Plane / NMS network) on the same or another interface of the Type C host, relying on other routers than the ACP edge node. The RIOs ensure that these hosts will only route the prefixes used in the ACP to the ACP edge node.

Type A/B host ignore the RIOs and will consider the ACP node to be their default router for all destination. This is sufficient when type A/B hosts only need to connect to the ACP but not to other networks. Attaching Type A/B hosts to both the ACP and other networks, requires either explicit ACP prefix route configuration on the Type A/B hosts or the combined ACP/Data-Plane interface on the ACP edge node, see Section 8.1.4.

Aggregated prefix means that the ACP edge node needs to only announce the /48 ULA prefixes used in the ACP but none of the actual /64 (Manual Addressing Sub-Scheme), /127 (ACP Zone Addressing Sub-Scheme), /112 or /120 (Vlong Addressing Sub-Scheme) routes of actual ACP nodes. If ACP interfaces are configured with non ULA prefixes, then those prefixes cannot be aggregated without further configured policy on the ACP edge node. This explains the above recommendation to use ACP ULA prefix covered prefixes for ACP connect interfaces: They allow for a shorter list of prefixes to be signaled via RFC4191 to NMS hosts and software components.

The ACP edge nodes that have a Vlong ACP address MAY allocate a subset of their /112 or /120 address prefix to ACP connect interface(s) to eliminate the need to non-autonomically configure/provision the address prefixes for such ACP connect interfaces.

8.1.4. Combined ACP/Data-Plane Interface (VRF Select)

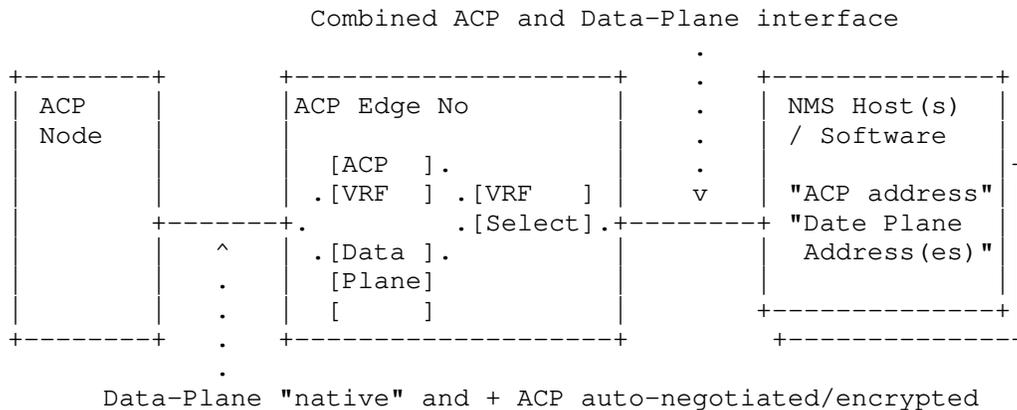


Figure 17: VRF select

Using two physical and/or virtual subnets (and therefore interfaces) into NMS Hosts (as per Section 8.1.1) or Software (as per Section 8.1.2) may be seen as additional complexity, for example with legacy NMS Hosts that support only one IP interface, or it may be insufficient to support [RFC4191] Type A or B host (see Section 8.1.3).

To provide a single subnet into both ACP and Data-Plane, the ACP Edge node needs to de-multiplex packets from NMS hosts into ACP VRF and Data-Plane. This is sometimes called "VRF select". If the ACP VRF has no overlapping IPv6 addresses with the Data-Plane (it should have no overlapping addresses), then this function can use the IPv6 Destination address. The problem is Source Address Selection on the NMS Host(s) according to RFC6724.

Consider the simple case: The ACP uses only one ULA prefix, the ACP IPv6 prefix for the Combined ACP and Data-Plane interface is covered by that ULA prefix. The ACP edge node announces both the ACP IPv6 prefix and one (or more) prefixes for the Data-Plane. Without further policy configurations on the NMS Host(s), it may select its ACP address as a source address for Data-Plane ULA destinations because of Rule 8 of RFC6724. The ACP edge node can pass on the packet to the Data-Plane, but the ACP source address should not be used for Data-Plane traffic, and return traffic may fail.

If the ACP carries multiple ULA prefixes or non-ULA ACP connect prefixes, then the correct source address selection becomes even more problematic.

With separate ACP connect and Data-Plane subnets and RFC4191 prefix announcements that are to be routed across the ACP connect interface, RFC6724 source address selection Rule 5 (use address of outgoing interface) will be used, so that above problems do not occur, even in more complex cases of multiple ULA and non-ULA prefixes in the ACP routing table.

To achieve the same behavior with a Combined ACP and Data-Plane interface, the ACP Edge Node needs to behave as two separate routers on the interface: One link-local IPv6 address/router for its ACP reachability, and one link-local IPv6 address/router for its Data-Plane reachability. The Router Advertisements for both are as described above (Section 8.1.3): For the ACP, the ACP prefix is announced together with RFC4191 option for the prefixes routed across the ACP and lifetime=0 to disqualify this next-hop as a default router. For the Data-Plane, the Data-Plane prefix(es) are announced together with whatever default router parameters are used for the Data-Plane.

In result, RFC6724 source address selection Rule 5.5 may result in the same correct source address selection behavior of NMS hosts without further configuration on it as the separate ACP connect and Data-Plane interfaces. As described in the text for Rule 5.5, this is only a MAY, because IPv6 hosts are not required to track next-hop information. If an NMS Host does not do this, then separate ACP connect and Data-Plane interfaces are the preferable method of attachment. Hosts implementing [RFC8028] should (instead of may) implement [RFC6724] Rule 5.5, so it is preferred for hosts to support [RFC8028].

ACP edge nodes MAY support the Combined ACP and Data-Plane interface.

#### 8.1.5. Use of GRASP

GRASP can and should be possible to use across ACP connect interfaces, especially in the architectural correct solution when it is used as a mechanism to connect Software (e.g., ASA or legacy NMS applications) to the ACP.

Given how the ACP is the security and transport substrate for GRASP, the requirements for devices connected via ACP connect is that those are equivalently (if not better) secured against attacks than ACP nodes that do not use ACP connect and run only software that is equally (if not better) protected, known (or trusted) not to be malicious and accordingly designed to isolate access to the ACP against external equipment.

The difference in security is that cryptographic security of the ACP secure channel is replaced by required physical security/control of the network connection between an ACP edge node and the NMS or other host reachable via the ACP connect interface. See Section 8.1.1.

When using "Combined ACP and Data-Plane Interfaces", care has to be taken that only GRASP messages intended for the ACP GRASP domain received from Software or NMS Hosts are forwarded by ACP edge nodes. Currently there is no definition for a GRASP security and transport substrate beside the ACP, so there is no definition how such Software/NMS Host could participate in two separate GRASP Domains across the same subnet (ACP and Data-Plane domains). At current it is assumed that all GRASP packets on a Combined ACP and Data-Plane interface belong to the GRASP ACP Domain. They SHOULD all use the ACP IPv6 addresses of the Software/NMS Hosts. The link-local IPv6 addresses of Software/NMS Hosts (used for GRASP M\_DISCOVERY and M\_FLOOD messages) are also assumed to belong to the ACP address space.

## 8.2. Connecting ACP islands over Non-ACP L3 networks (Remote ACP neighbors)

Not all nodes in a network may support the ACP. If non-ACP Layer-2 devices are between ACP nodes, the ACP will work across it since it is IP based. However, the autonomic discovery of ACP neighbors via DULL GRASP is only intended to work across L2 connections, so it is not sufficient to autonomically create ACP connections across non-ACP Layer-3 devices.

### 8.2.1. Configured Remote ACP neighbor

On the ACP node, remote ACP neighbors are configured explicitly. The parameters of such a "connection" are described in the following ABNF.

```

connection = [ method , local-addr, remote-addr, ?pmtu ]
method = [ "IKEv2", ?port ]
method =/ [ "DTLS", port ]
local-addr = [ address , ?vrf ]
remote-addr = [ address ]
address = ("any" | ipv4-address | ipv6-address )
vrf = tstr ; Name of a VRF on this node with local-address

```

Figure 18: Parameters for remote ACP neighbors

Explicit configuration of a remote-peer according to this ABNF provides all the information to build a secure channel without requiring a tunnel to that peer and running DULL GRASP inside of it.

The configuration includes the parameters otherwise signaled via DULL GRASP: local address, remote (peer) locator and method. The differences over DULL GRASP local neighbor discovery and secure channel creation are as follows:

- \* The local and remote address can be IPv4 or IPv6 and are typically global scope addresses.
- \* The VRF across which the connection is built (and in which local-addr exists) can to be specified. If vrf is not specified, it is the default VRF on the node. In DULL GRASP the VRF is implied by the interface across which DULL GRASP operates.
- \* If local address is "any", the local address used when initiating a secure channel connection is decided by source address selection ([RFC6724] for IPv6). As a responder, the connection listens on all addresses of the node in the selected VRF.
- \* Configuration of port is only required for methods where no defaults exist (e.g., "DTLS").

- \* If remote address is "any", the connection is only a responder. It is a "hub" that can be used by multiple remote peers to connect simultaneously - without having to know or configure their addresses. Example: Hub site for remote "spoke" sites reachable over the Internet.
- \* Pmtu should be configurable to overcome issues/limitations of Path MTU Discovery (PMTUD).
- \* IKEv2/IPsec to remote peers should support the optional NAT Traversal (NAT-T) procedures.

#### 8.2.2. Tunneled Remote ACP Neighbor

An IPinIP, GRE or other form of pre-existing tunnel is configured between two remote ACP peers and the virtual interfaces representing the tunnel are configured for "ACP enable". This will enable IPv6 link local addresses and DULL on this tunnel. In result, the tunnel is used for normal "L2 adjacent" candidate ACP neighbor discovery with DULL and secure channel setup procedures described in this document.

Tunneled Remote ACP Neighbor requires two encapsulations: the configured tunnel and the secure channel inside of that tunnel. This makes it in general less desirable than Configured Remote ACP Neighbor. Benefits of tunnels are that it may be easier to implement because there is no change to the ACP functionality - just running it over a virtual (tunnel) interface instead of only native interfaces. The tunnel itself may also provide PMTUD while the secure channel method may not. Or the tunnel mechanism is permitted/possible through some firewall while the secure channel method may not.

Tunneling using an insecure tunnel encapsulation increases on average the risk of a MITM downgrade attack somewhere along the underlay path that blocks all but the most easily attacked ACP secure channel option. ACP nodes supporting tunneled remote ACP Neighbors SHOULD support configuration on such tunnel interfaces to restrict or explicitly select the available ACP secure channel protocols (if the ACP node supports more than one ACP secure channel protocol in the first place).

#### 8.2.3. Summary

Configured/Tunneled Remote ACP neighbors are less "indestructible" than L2 adjacent ACP neighbors based on link local addressing, since they depend on more correct Data-Plane operations, such as routing and global addressing.

Nevertheless, these options may be crucial to incrementally deploy the ACP, especially if it is meant to connect islands across the Internet. Implementations SHOULD support at least Tunneled Remote ACP Neighbors via GRE tunnels - which is likely the most common router-to-router tunneling protocol in use today.

## 9. ACP Operations (Informative)

The following sections document important operational aspects of the ACP. They are not normative because they do not impact the interoperability between components of the ACP, but they include recommendations/requirements for the internal operational model beneficial or necessary to achieve the desired use-case benefits of the ACP (see Section 3).

- \* Section 9.1 describes recommended operator diagnostics capabilities of ACP nodes.
- \* Section 9.2 describes high level how an ACP registrar needs to work, what its configuration parameters are and specific issues impacting the choices of deployment design due to renewal and revocation issues. It describes a model where ACP Registrars have their own sub-CA to provide the most distributed deployment option for ACP Registrars, and it describes considerations for centralized policy control of ACP Registrar operations.
- \* Section 9.3 describes suggested ACP node behavior and operational interfaces (configuration options) to manage the ACP in so-called greenfield devices (previously unconfigured) and brownfield devices (preconfigured).

The recommendations and suggestions of this chapter were derived from operational experience gained with a commercially available pre-standard ACP implementation.

### 9.1. ACP (and BRSKI) Diagnostics

Even though ACP and ANI in general are taking out many manual configuration mistakes through their automation, it is important to provide good diagnostics for them.

Basic standardized diagnostics would require support for (yang) models representing the complete (auto-)configuration and operational state of all components: GRASP, ACP and the infrastructure used by them: TLS/DTLS, IPsec, certificates, TA, time, VRF and so on. While necessary, this is not sufficient:

Simply representing the state of components does not allow operators to quickly take action - unless they do understand how to interpret the data, and that can mean a requirement for deep understanding of all components and how they interact in the ACP/ANI.

Diagnostic supports should help to quickly answer the questions operators are expected to ask, such as "is the ACP working correctly?", or "why is there no ACP connection to a known neighboring node?"

In current network management approaches, the logic to answer these questions is most often built as centralized diagnostics software that leverages the above mentioned data models. While this approach is feasible for components utilizing the ANI, it is not sufficient to diagnose the ANI itself:

- \* Developing the logic to identify common issues requires operational experience with the components of the ANI. Letting each management system define its own analysis is inefficient.
- \* When the ANI is not operating correctly, it may not be possible to run diagnostics from remote because of missing connectivity. The ANI should therefore have diagnostic capabilities available locally on the nodes themselves.
- \* Certain operations are difficult or impossible to monitor in real-time, such as initial bootstrap issues in a network location where no capabilities exist to attach local diagnostics. Therefore, it is important to also define means of capturing (logging) diagnostics locally for later retrieval. Ideally, these captures are also non-volatile so that they can survive extended power-off conditions - for example when a device that fails to be brought up zero-touch is being sent back for diagnostics at a more appropriate location.

The simplest form of diagnostics answering questions such as the above is to represent the relevant information sequentially in dependency order, so that the first non-expected/non-operational item is the most likely root cause. Or just log/highlight that item. For example:

Q: Is ACP operational to accept neighbor connections:

- \* Check if any potentially necessary configuration to make ACP/ANI operational are correct (see Section 9.3 for a discussion of such commands).
- \* Does the system time look reasonable, or could it be the default system time after clock chip battery failure (certificate checks depend on reasonable notion of time)?.

- \* Does the node have keying material - domain certificate, TA certificates, ...>
- \* If no keying material and ANI is supported/enabled, check the state of BRSKI (not detailed in this example).
- \* Check the validity of the domain certificate:
  - Does the certificate validate against the TA?
  - Has it been revoked?
  - Was the last scheduled attempt to retrieve a CRL successful (e.g., do we know that our CRL information is up to date).
  - Is the certificate valid: validity start time in the past, expiration time in the future?
  - Does the certificate have a correctly formatted acp-node-name field?
- \* Was the ACP VRF successfully created?
- \* Is ACP enabled on one or more interfaces that are up and running?

If all this looks good, the ACP should be running locally "fine" - but we did not check any ACP neighbor relationships.

Question: why does the node not create a working ACP connection to a neighbor on an interface?

- \* Is the interface physically up? Does it have an IPv6 link-local address?
- \* Is it enabled for ACP?
- \* Do we successfully send DULL GRASP messages to the interface (link layer errors)?
- \* Do we receive DULL GRASP messages on the interface? If not, some intervening L2 equipment performing bad MLD snooping could have caused problems. Provide e.g., diagnostics of the MLD querier IPv6 and MAC address.
- \* Do we see the ACP objective in any DULL GRASP message from that interface? Diagnose the supported secure channel methods.
- \* Do we know the MAC address of the neighbor with the ACP objective? If not, diagnose SLAAC/ND state.
- \* When did we last attempt to build an ACP secure channel to the neighbor?
- \* If it failed, why:
  - Did the neighbor close the connection on us or did we close the connection on it because the domain certificate membership failed?
  - If the neighbor closed the connection on us, provide any error diagnostics from the secure channel protocol.
  - If we failed the attempt, display our local reason:
    - o There was no common secure channel protocol supported by the two neighbors (this could not happen on nodes supporting this specification because it mandates common support for IPsec).

- o The ACP certificate membership check (Section 6.2.3) fails:
  - + The neighbor's certificate is not signed directly or indirectly by one of the nodes TA. Provide diagnostics which TA it has (can identify whom the device belongs to).
  - + The neighbor's certificate does not have the same domain (or no domain at all). Diagnose domain-name and potentially other cert info.
  - + The neighbor's certificate has been revoked or could not be authenticated by OCSP.
  - + The neighbor's certificate has expired - or is not yet valid.
- Any other connection issues in e.g., IKEv2 / IPsec, DTLS?.

Question: Is the ACP operating correctly across its secure channels?

- \* Are there one or more active ACP neighbors with secure channels?
- \* Is the RPL routing protocol for the ACP running?
- \* Is there a default route to the root in the ACP routing table?
- \* Is there for each direct ACP neighbor not reachable over the ACP virtual interface to the root a route in the ACP routing table?
- \* Is ACP GRASP running?
- \* Is at least one SRV.est objective cached (to support certificate renewal)?
- \* Is there at least one BRSKI registrar objective cached (in case BRSKI is supported)
- \* Is BRSKI proxy operating normally on all interfaces where ACP is operating?
- \* ...

These lists are not necessarily complete, but illustrate the principle and show that there are variety of issues ranging from normal operational causes (a neighbor in another ACP domain) over problems in the credentials management (certificate lifetimes), explicit security actions (revocation) or unexpected connectivity issues (intervening L2 equipment).

The items so far are illustrating how the ANI operations can be diagnosed with passive observation of the operational state of its components including historic/cached/counted events. This is not necessary sufficient to provide good enough diagnostics overall:

The components of ACP and BRSKI are designed with security in mind but they do not attempt to provide diagnostics for building the network itself. Consider two examples:

1. BRSKI does not allow for a neighboring device to identify the pledges IDevID certificate. Only the selected BRSKI registrar can do this, but it may be difficult to disseminate information about undesired pledges from those BRSKI registrars to locations/nodes where information about those pledges is desired.
2. LLDP disseminates information about nodes to their immediate neighbors, such as node model/type/software and interface name/number of the connection. This information is often helpful or even necessary in network diagnostics. It can equally be considered to be too insecure to make this information available unprotected to all possible neighbors.

An "interested adjacent party" can always determine the IDevID certificate of a BRSKI pledge by behaving like a BRSKI proxy/registrar. Therefore, the IDevID certificate of a BRSKI pledge is not meant to be protected - it just has to be queried and is not signaled unsolicited (as it would be in LLDP) so that other observers on the same subnet can determine who is an "interested adjacent party".

#### 9.1.1. Secure Channel Peer diagnostics

When using mutual certificate authentication, the TA certificate is not required to be signaled explicitly because its hash is sufficient for certificate chain validation. In the case of ACP secure channel setup this leads to limited diagnostics when authentication fails because of TA mismatch. For this reason, Section 6.8.2 recommends to also include the TA certificate in the secure channel signaling. This should be possible to do without protocol modifications in the security association protocols used by the ACP. For example, while [RFC7296] does not mention this, it also does not prohibit it.

One common deployment use case where the diagnostic through the signaled TA of a candidate peer is very helpful are multi-tenant environments such as office buildings, where different tenants run their own networks and ACPs. Each tenant is given supposedly disjoint L2 connectivity through the building infrastructure. In these environments there are various common errors through which a device may receive L2 connectivity into the wrong tenants network.

While the ACP itself is not impacted by this, the Data-Plane to be built later may be impacted. Therefore, it is important to be able to diagnose such undesirable connectivity from the ACP so that any autonomic or non-autonomic mechanisms to configure the Data-Plane can accordingly treat such interfaces. The information in the TA of the peer can then ease troubleshooting of such issues.

Another example case is the intended or accidental re-activation of equipment whose TA certificate has long expired, such as redundant gear taken from storage after years.

A third example case is when in a mergers & acquisition case ACP nodes have not been correctly provisioned with the mutual TA of previously disjoint ACP. This is assuming that the ACP domain names were already aligned so that the ACP domain membership check is only failing on the TA.

A fourth example case is when multiple registrars were set up for the same ACP but without correctly setting up the same TA. For example, when registrars support to also be CA themselves but are misconfigured to become TA instead of intermediate CA.

## 9.2. ACP Registrars

As described in Section 6.11.7, the ACP addressing mechanism is designed to enable lightweight, distributed and uncoordinated ACP registrars that are providing ACP address prefixes to candidate ACP nodes by enrolling them with an ACP certificate into an ACP domain via any appropriate mechanism/protocol, automated or not.

This section discusses informatively more details and options for ACP registrars.

### 9.2.1. Registrar interactions

This section summarizes and discusses the interactions with other entities required by an ACP registrar.

In a simple instance of an ACP network, no central NOC component beside a TA is required. Typically, this is a root CA. One or more uncoordinated acting ACP registrar can be set up, performing the following interactions:

To orchestrate enrolling a candidate ACP node autonomically, the ACP registrar can rely on the ACP and use Proxies to reach the candidate ACP node, therefore allowing minimum pre-existing (auto-)configured network services on the candidate ACP node. BRSKI defines the BRSKI proxy, a design that can be adopted for various protocols that Pledges/candidate ACP nodes could want to use, for example BRSKI over CoAP (Constrained Application Protocol), or proxying of NETCONF.

To reach a TA that has no ACP connectivity, the ACP registrar would use the Data-Plane. ACP and Data-Plane in an ACP registrar could (and by default should be) completely isolated from each other at the network level. Only applications such as the ACP registrar would need the ability for their transport stacks to access both.

In non-autonomic enrollment options, the Data-Plane between a ACP registrar and the candidate ACP node needs to be configured first. This includes the ACP registrar and the candidate ACP node. Then any appropriate set of protocols can be used between ACP registrar and candidate ACP node to discover the other side, and then connect and enroll (configure) the candidate ACP node with an ACP certificate. NETCONF ZeroTouch ([RFC8572]) is an example protocol that could be used for this. BRSKI using optional discovery mechanisms is equally a possibility for candidate ACP nodes attempting to be enrolled across non-ACP networks, such as the Internet.

When candidate ACP nodes have secure bootstrap, such as BRSKI Pledges, they will not trust to be configured/enrolled across the network, unless being presented with a voucher (see [RFC8366]) authorizing the network to take possession of the node. An ACP registrar will then need a method to retrieve such a voucher, either offline, or online from a MASA (Manufacturer Authorized Signing Authority). BRSKI and NETCONF ZeroTouch are two protocols that include capabilities to present the voucher to the candidate ACP node.

An ACP registrar could operate EST for ACP certificate renewal and/or act as a CRL Distribution point. A node performing these services does not need to support performing (initial) enrollment, but it does require the same above described connectivity as an ACP registrar: via the ACP to ACP nodes and via the Data-Plane to the TA and other sources of CRL information.

#### 9.2.2. Registrar Parameter

The interactions of an ACP registrar outlined Section 6.11.7 and Section 9.2.1 above depend on the following parameters:

- \* A URL to the TA and credentials so that the ACP registrar can let the TA sign candidate ACP node certificates.
- \* The ACP domain-name.
- \* The Registrar-ID to use. This could default to a MAC address of the ACP registrar.

- \* For recovery, the next-useable Node-IDs for zone (Zone-ID=0) sub-addressing scheme, for Vlong /112 and for Vlong /120 sub-addressing scheme. These IDs would only need to be provisioned after recovering from a crash. Some other mechanism would be required to remember these IDs in a backup location or to recover them from the set of currently known ACP nodes.
- \* Policies if candidate ACP nodes should receive a domain certificate or not, for example based on the devices IDevID certificate as in BRSKI. The ACP registrar may have a whitelist or blacklist of devices [X.520] "serialNumbers" attribute in the subject field distinguished name encoding from their IDevID certificate.
- \* Policies what type of address prefix to assign to a candidate ACP devices, based on likely the same information.
- \* For BRSKI or other mechanisms using vouchers: Parameters to determine how to retrieve vouchers for specific type of secure bootstrap candidate ACP nodes (such as MASA URLs), unless this information is automatically learned such as from the IDevID certificate of candidate ACP nodes (as defined in BRSKI).

### 9.2.3. Certificate renewal and limitations

When an ACP node renews/rekeys its certificate, it may end up doing so via a different registrar (e.g., EST server) than the one it originally received its ACP certificate from, for example because that original ACP registrar is gone. The ACP registrar through which the renewal/rekeying is performed would by default trust the acp-node-name from the ACP nodes current ACP certificate and maintain this information so that the ACP node maintains its ACP address prefix. In EST renewal/rekeying, the ACP nodes current ACP certificate is signaled during the TLS handshake.

This simple scenario has two limitations:

1. The ACP registrars cannot directly assign certificates to nodes and therefore needs an "online" connection to the TA.
2. Recovery from a compromised ACP registrar is difficult. When an ACP registrar is compromised, it can insert for example a conflicting acp-node-name and create thereby an attack against other ACP nodes through the ACP routing protocol.

Even when such a malicious ACP registrar is detected, resolving the problem may be difficult because it would require identifying all the wrong ACP certificates assigned via the ACP registrar after it was compromised. And without additional centralized tracking of assigned certificates there is no way to do this.

#### 9.2.4. ACP Registrars with sub-CA

In situations, where either of the above two limitations are an issue, ACP registrars could also be sub-CAs. This removes the need for connectivity to a TA whenever an ACP node is enrolled, and reduces the need for connectivity of such an ACP registrar to a TA to only those times when it needs to renew its own certificate. The ACP registrar would also now use its own (sub-CA) certificate to enroll and sign the ACP nodes certificates, and therefore it is only necessary to revoke a compromised ACP registrars sub-CA certificate. Alternatively one can let it expire and not renew it, when the certificate of the sub-CA is appropriately short-lived.

As the ACP domain membership check verifies a peer ACP node's ACP certificate trust chain, it will also verify the signing certificate which is the compromised/revoked sub-CA certificate. Therefore, ACP domain membership for an ACP node enrolled from a compromised and discovered ACP registrar will fail.

ACP nodes enrolled by a compromised ACP registrar would automatically fail to establish ACP channels and ACP domain certificate renewal via EST and therefore revert to their role as a candidate ACP members and attempt to get a new ACP certificate from an ACP registrar - for example, via BRSKI. In result, ACP registrars that have an associated sub-CA makes isolating and resolving issues with compromised registrars easier.

Note that ACP registrars with sub-CA functionality also can control the lifetime of ACP certificates easier and therefore also be used as a tool to introduce short lived certificates and not rely on CRL, whereas the certificates for the sub-CAs themselves could be longer lived and subject to CRL.

#### 9.2.5. Centralized Policy Control

When using multiple, uncoordinated ACP registrars, several advanced operations are potentially more complex than with a single, resilient policy control backend, for example including but not limited to:

- \* Which candidate ACP node is permitted or not permitted into an ACP domain. This may not be a decision to be taken upfront, so that a policy per "serialNumber" attribute in the subject field distinguished name encoding can be loaded into every ACP registrar. Instead, it may better be decided in real-time including potentially a human decision in a NOC.
- \* Tracking of all enrolled ACP nodes and their certificate information. For example, in support of revoking individual ACP nodes certificates.

- \* More flexible policies what type of address prefix or even what specific address prefix to assign to a candidate ACP node.

These and other operations could be introduced more easily by introducing a centralized Policy Management System (PMS) and modifying ACP registrar behavior so that it queries the PMS for any policy decision occurring during the candidate ACP node enrollment process and/or the ACP node certificate renewal process. For example, which ACP address prefix to assign. Likewise the ACP registrar would report any relevant state change information to the PMS as well, for example when a certificate was successfully enrolled onto a candidate ACP node.

### 9.3. Enabling and disabling ACP/ANI

Both ACP and BRSKI require interfaces to be operational enough to support sending/receiving their packets. In node types where interfaces are by default (e.g., without operator configuration) enabled, such as most L2 switches, this would be less of a change in behavior than in most L3 devices (e.g. routers), where interfaces are by default disabled. In almost all network devices it is common though for configuration to change interfaces to a physically disabled state and that would break the ACP.

In this section, we discuss a suggested operational model to enable/disable interfaces and nodes for ACP/ANI in a way that minimizes the risk of operator action to break the ACP in this way, and that also minimizes operator surprise when ACP/ANI becomes supported in node software.

#### 9.3.1. Filtering for non-ACP/ANI packets

Whenever this document refers to enabling an interface for ACP (or BRSKI), it only requires to permit the interface to send/receive packets necessary to operate ACP (or BRSKI) - but not any other Data-Plane packets. Unless the Data-Plane is explicitly configured/enabled, all packets not required for ACP/BRSKI should be filtered on input and output:

Both BRSKI and ACP require link-local only IPv6 operations on interfaces and DULL GRASP. IPv6 link-local operations means the minimum signaling to auto-assign an IPv6 link-local address and talk to neighbors via their link-local address: SLAAC (Stateless Address Auto-Configuration - [RFC4862]) and ND (Neighbor Discovery - [RFC4861]). When the device is a BRSKI pledge, it may also require TCP/TLS connections to BRSKI proxies on the interface. When the device has keying material, and the ACP is running, it requires DULL GRASP packets and packets necessary for the secure-channel mechanism

it supports, e.g., IKEv2 and IPsec ESP packets or DTLS packets to the IPv6 link-local address of an ACP neighbor on the interface. It also requires TCP/TLS packets for its BRSKI proxy functionality, if it does support BRSKI.

### 9.3.2. Admin Down State

Interfaces on most network equipment have at least two states: "up" and "down". These may have product specific names. "down" for example could be called "shutdown" and "up" could be called "no shutdown". The "down" state disables all interface operations down to the physical level. The "up" state enables the interface enough for all possible L2/L3 services to operate on top of it and it may also auto-enable some subset of them. More commonly, the operations of various L2/L3 services is controlled via additional node-wide or interface level options, but they all become only active when the interface is not "down". Therefore, an easy way to ensure that all L2/L3 operations on an interface are inactive is to put the interface into "down" state. The fact that this also physically shuts down the interface is in many cases just a side effect, but it may be important in other cases (see below, Section 9.3.2.2).

One of the common problems of remote management is for the operator or SDN controller to cut its own connectivity to the remote node by a configuration impacting its own management connection into the node. The ACP itself should have no dedicated configuration other than aforementioned enablement of the ACP on brownfield ACP nodes. This leaves configuration that cannot distinguish between ACP and Data-Plane as sources of configuration mistakes as these commands will impact the ACP even though they should only impact the Data-Plane.

The one ubiquitous type of commands that do this on many type of routers are interface "down" commands/configurations. When such a command is applied to the interface through which the ACP provides access for remote management it would cut the remote management connection through the ACP because, as outlined above, the "down" commands typically impact the physical layer too and not only the Data-Plane services.

To provide ACP/ANI resilience against such operator misconfiguration, this document recommends to separate the "down" state of interfaces into an "admin down" state where the physical layer is kept running and ACP/ANI can use the interface and a "physical down" state. Any existing "down" configurations would map to "admin down". In "admin down", any existing L2/L3 services of the Data-Plane should see no difference to "physical down" state. To ensure that no Data-Plane packets could be sent/received, packet filtering could be established automatically as described above in Section 9.3.1.

An example of non-ACP but ANI traffic that should be permitted to pass even in "admin-down" state is BRSKI enrollment traffic between BRSKI pledge and a BRSKI proxy.

As necessary (see discussion below) new configuration options could be introduced to issue "physical down". The options should be provided with additional checks to minimize the risk of issuing them in a way that breaks the ACP without automatic restoration. For example, they could be denied to be issued from a control connection (NETCONF/SSH) that goes across the interface itself ("do not disconnect yourself"). Or they could be performed only temporary and only be made permanent with additional later reconfirmation.

In the following sub-sections important aspects to the introduction of "admin down" state are discussed.

#### 9.3.2.1. Security

Interfaces are physically brought down (or left in default down state) as a form of security. "Admin down" state as described above provides also a high level of security because it only permits ACP/ANI operations which are both well secured. Ultimately, it is subject to security review for the deployment whether "admin down" is a feasible replacement for "physical down".

The need to trust the security of ACP/ANI operations needs to be weighed against the operational benefits of permitting this: Consider the typical example of a CPE (customer premises equipment) with no on-site network expert. User ports are in physical down state unless explicitly configured not to be. In a misconfiguration situation, the uplink connection is incorrectly plugged into such as user port. The device is disconnected from the network and therefore no diagnostics from the network side is possible anymore. Alternatively, all ports default to "admin down". The ACP (but not the Data-Plane) would still automatically form. Diagnostics from the network side is possible and operator reaction could include to either make this port the operational uplink port or to instruct re-cabling. Security wise, only ACP/ANI could be attacked, all other functions are filtered on interfaces in "admin down" state.

#### 9.3.2.2. Fast state propagation and Diagnostics

"Physical down" state propagates on many interface types (e.g., Ethernet) to the other side. This can trigger fast L2/L3 protocol reaction on the other side and "admin down" would not have the same (fast) result.

Bringing interfaces to "physical down" state is to the best of our knowledge always a result of operator action, but today, never the result of autonomic L2/L3 services running on the nodes. Therefore, one option is to change the operator action to not rely on link-state propagation anymore. This may not be possible when both sides are under different operator control, but in that case it is unlikely that the ACP is running across the link and actually putting the interface into "physical down" state may still be a good option.

Ideally, fast physical state propagation is replaced by fast software driven state propagation. For example, a DULL GRASP "admin-state" objective could be used to auto configure a Bidirectional Forwarding Protocol (BFD, [RFC5880]) session between the two sides of the link that would be used to propagate the "up" vs. admin down state.

Triggering physical down state may also be used as a mean of diagnosing cabling in the absence of easier methods. It is more complex than automated neighbor diagnostics because it requires coordinated remote access to both (likely) sides of a link to determine whether up/down toggling will cause the same reaction on the remote side.

See Section 9.1 for a discussion about how LLDP and/or diagnostics via GRASP could be used to provide neighbor diagnostics, and therefore hopefully eliminating the need for "physical down" for neighbor diagnostics - as long as both neighbors support ACP/ANI.

### 9.3.2.3. Low Level Link Diagnostics

"Physical down" is performed to diagnose low-level interface behavior when higher layer services (e.g., IPv6) are not working. Especially Ethernet links are subject to a wide variety of possible wrong configuration/cablings if they do not support automatic selection of variable parameters such as speed (10/100/1000 Mbps), crossover (Auto-MDIX) and connector (fiber, copper - when interfaces have multiple but can only enable one at a time). The need for low level link diagnostic can therefore be minimized by using fully auto configuring links.

In addition to "Physical down", low level diagnostics of Ethernet or other interfaces also involve the creation of other states on interfaces, such as physical Loopback (internal and/or external) or bringing down all packet transmissions for reflection/cable-length measurements. Any of these options would disrupt ACP as well.

In cases where such low-level diagnostics of an operational link is desired but where the link could be a single point of failure for the ACP, ASA on both nodes of the link could perform a negotiated

diagnostic that automatically terminates in a predetermined manner without dependence on external input ensuring the link will become operational again.

#### 9.3.2.4. Power Consumption Issues

Power consumption of "physical down" interfaces, may be significantly lower than those in "admin down" state, for example on long-range fiber interfaces. Bringing up interfaces, for example to probe reachability, may also consume additional power. This can make these type of interfaces inappropriate to operate purely for the ACP when they are not currently needed for the Data-Plane.

#### 9.3.3. Interface level ACP/ANI enable

The interface level configuration option "ACP enable" enables ACP operations on an interface, starting with ACP neighbor discovery via DULL GRAP. The interface level configuration option "ANI enable" on nodes supporting BRSKI and ACP starts with BRSKI pledge operations when there is no domain certificate on the node. On ACP/BRSKI nodes, "ACP enable" may not need to be supported, but only "ANI enable". Unless overridden by global configuration options (see later), "ACP/ANI enable" will result in "down" state on an interface to behave as "admin down".

#### 9.3.4. Which interfaces to auto-enable?

(Section 6.4) requires that "ACP enable" is automatically set on native interfaces, but not on non-native interfaces (reminder: a native interface is one that exists without operator configuration action such as physical interfaces in physical devices).

Ideally, ACP enable is set automatically on all interfaces that provide access to additional connectivity that allows to reach more nodes of the ACP domain. The best set of interfaces necessary to achieve this is not possible to determine automatically. Native interfaces are the best automatic approximation.

Consider an ACP domain of ACP nodes transitively connected via native interfaces. A Data-Plane tunnel between two of these nodes that are non-adjacent is created and "ACP enable" is set for that tunnel. ACP RPL sees this tunnel as just as a single hop. Routes in the ACP would use this hop as an attractive path element to connect regions adjacent to the tunnel nodes. In result, the actual hop-by-hop paths used by traffic in the ACP can become worse. In addition, correct forwarding in the ACP now depends on correct Data-Plane forwarding config including QoS, filtering and other security on the Data-Plane path across which this tunnel runs. This is the main issue why "ACP/ANI enable" should not be set automatically on non-native interfaces.

If the tunnel would connect two previously disjoint ACP regions, then it likely would be useful for the ACP. A Data-Plane tunnel could also run across nodes without ACP and provide additional connectivity for an already connected ACP network. The benefit of this additional ACP redundancy has to be weighed against the problems of relying on the Data-Plane. If a tunnel connects two separate ACP regions: how many tunnels should be created to connect these ACP regions reliably enough? Between which nodes? These are all standard tunneled network design questions not specific to the ACP, and there are no generic fully automated answers.

Instead of automatically setting "ACP enable" on these type of interfaces, the decision needs to be based on the use purpose of the non-native interface and "ACP enable" needs to be set in conjunction with the mechanism through which the non-native interface is created/configured.

In addition to explicit setting of "ACP/ANI enable", non-native interfaces also need to support configuration of the ACP RPL cost of the link - to avoid the problems of attracting too much traffic to the link as described above.

Even native interfaces may not be able to automatically perform BRSKI or ACP because they may require additional operator input to become operational. Example include DSL interfaces requiring PPPoE credentials or mobile interfaces requiring credentials from a SIM card. Whatever mechanism is used to provide the necessary config to the device to enable the interface can also be expanded to decide on whether or not to set "ACP/ANI enable".

The goal of automatically setting "ACP/ANI enable" on interfaces (native or not) is to eliminate unnecessary "touches" to the node to make its operation as much as possible "zero-touch" with respect to ACP/ANI. If there are "unavoidable touches" such a creating/configuring a non-native interface or provisioning credentials for a native interface, then "ACP/ANI enable" should be added as an option

to that "touch". If a wrong "touch" is easily fixed (not creating another high-cost touch), then the default should be not to enable ANI/ACP, and if it is potentially expensive or slow to fix (e.g., parameters on SIM card shipped to remote location), then the default should be to enable ACP/ANI.

#### 9.3.5. Node Level ACP/ANI enable

A node level command "ACP/ANI enable [up-if-only]" enables ACP or ANI on the node (ANI = ACP + BRSKI). Without this command set, any interface level "ACP/ANI enable" is ignored. Once set, ACP/ANI will operate an interface where "ACP/ANI enable" is set. Setting of interface level "ACP/ANI enable" is either automatic (default) or explicit through operator action as described in the previous section.

If the option "up-if-only" is selected, the behavior of "down" interfaces is unchanged, and ACP/ANI will only operate on interfaces where "ACP/ANI enable" is set and that are "up". When it is not set, then "down" state of interfaces with "ACP/ANI enable" is modified to behave as "admin down".

##### 9.3.5.1. Brownfield nodes

A "brownfield" node is one that already has a configured Data-Plane.

Executing global "ACP/ANI enable [up-if-only]" on each node is the only command necessary to create an ACP across a network of brownfield nodes once all the nodes have a domain certificate. When BRSKI is used ("ANI enable"), provisioning of the certificates only requires set-up of a single BRSKI registrar node which could also implement a CA for the network. This is the simplest way to introduce ACP/ANI into existing (== brownfield) networks.

The need to explicitly enable ACP/ANI is especially important in brownfield nodes because otherwise software updates may introduce support for ACP/ANI: Automatic enablement of ACP/ANI in networks where the operator does not only not want ACP/ANI but where the operator likely never even heard of it could be quite irritating to the operator. Especially when "down" behavior is changed to "admin down".

Automatically setting "ANI enable" on brownfield nodes where the operator is unaware of BRSKI and MASA operations could also be an unlikely but then critical security issue. If an attacker could impersonate the operator and register as the operator at the MASA or otherwise get hold of vouchers and can get enough physical access to the network so pledges would register to an attacking registrar, then the attacker could gain access to the ACP, and through the ACP gain access to the Data-Plane.

In networks where the operator explicitly wants to enable the ANI this could not happen, because the operator would create a BRSKI registrar that would discover attack attempts, and the operator would be setting up his registrar with the MASA. Nodes requiring "ownership vouchers" would not be subject to that attack. See [I-D.ietf-anima-bootstrapping-keyinfra] for more details. Note that a global "ACP enable" alone is not subject to these type of attacks, because it always depends on some other mechanism first to provision domain certificates into the device.

#### 9.3.5.2. Greenfield nodes

An ACP "greenfield" node is one that does not have any prior configuration and that can be bootstrapped into the ACP across the network. To support greenfield nodes, ACP as described in this document needs to be combined with a bootstrap protocol/mechanism that will enroll the node with the ACP keying material - ACP certificate and TA. For ANI nodes, this protocol/mechanism is BRSKI.

When such a node is powered on and determines it is in greenfield condition, it enables the bootstrap protocol(s)/mechanism(s), and once the ACP keying material is enrolled, greenfield state ends and the ACP is started. When BRSKI is used, the node's state reflects this by setting "ANI enable" upon determination of greenfield state at power on.

ACP greenfield nodes that in the absence of ACP would have their interfaces in "down" state SHOULD set all native interfaces into "admin down" state and only permit Data-Plane traffic required for the bootstrap protocol/mechanisms.

ACP greenfield state ends either through successful enrolment of ACP keying material (certificate, TA) or detection of a permitted termination of ACP greenfield operations.

ACP nodes supporting greenfield operations MAY want to provide backward compatibility with other forms of configuration/provisioning, especially when only a subset of nodes are expected to be deployed with ACP. Such an ACP node SHOULD observe attempts to

provision/configure the node via interfaces/methods that traditionally indicate physical possession of the node, such as a serial or USB console port or a USB memory stick with a bootstrap configuration. When such an operation is observed before enrollment of the ACP keying material has completed, the node SHOULD put itself into the state the node would have been in, if ACP/ANI was disabled at boot (terminate ACP greenfield operations).

When an ACP greenfield node enables multiple automated ACP or non-ACP enrollment/bootstrap protocols/mechanisms in parallel, care must be taken not to terminate any protocol/mechanism before another one has succeeded to enroll ACP keying material or has progressed to a point where it is permitted to be a termination reason for ACP greenfield operations.

Highly secure ACP greenfield nodes may not permit any reason to terminate ACP greenfield operations, including physical access.

Nodes that claim to support ANI greenfield operations SHOULD NOT enable in parallel to BRSKI any enrollment/bootstrap protocol/mechanism that allows Trust On First Use (TOFU, [RFC7435]) over interfaces other than those traditionally indicating physical possession of the node. Protocols/mechanisms with published default username/password authentication are considered to suffer from TOFU. Securing the bootstrap protocol/mechanism by requiring a voucher ([RFC8366]) can be used to avoid TOFU.

In summary, the goal of ACP greenfield support is to allow remote automated enrollment of ACP keying materials, and therefore automated bootstrap into the ACP and to prohibit TOFU during bootstrap with the likely exception (for backward compatibility) of bootstrapping via interfaces traditionally indicating physical possession of the node.

#### 9.3.6. Undoing ANI/ACP enable

Disabling ANI/ACP by undoing "ACP/ANI enable" is a risk for the reliable operations of the ACP if it can be executed by mistake or unauthorized. This behavior could be influenced through some additional (future) property in the certificate (e.g., in the acp-node-name extension field): In an ANI deployment intended for convenience, disabling it could be allowed without further constraints. In an ANI deployment considered to be critical more checks would be required. One very controlled option would be to not permit these commands unless the domain certificate has been revoked or is denied renewal. Configuring this option would be a parameter on the BRSKI registrar(s). As long as the node did not receive a domain certificate, undoing "ANI/ACP enable" should not have any additional constraints.

### 9.3.7. Summary

Node-wide "ACP/ANI enable [up-if-only]" commands enable the operation of ACP/ANI. This is only auto-enabled on ANI greenfield devices, otherwise it must be configured explicitly.

If the option "up-if-only" is not selected, interfaces enabled for ACP/ANI interpret "down" state as "admin down" and not "physical down". In "admin-down" all non-ACP/ANI packets are filtered, but the physical layer is kept running to permit ACP/ANI to operate.

(New) commands that result in physical interruption ("physical down", "loopback") of ACP/ANI enabled interfaces should be built to protect continuance or reestablishment of ACP as much as possible.

Interface level "ACP/ANI enable" control per-interface operations. It is enabled by default on native interfaces and has to be configured explicitly on other interfaces.

Disabling "ACP/ANI enable" global and per-interface should have additional checks to minimize undesired breakage of ACP. The degree of control could be a domain wide parameter in the domain certificates.

### 9.4. Partial or Incremental adoption

The ACP Zone Addressing Sub-Scheme (see Section 6.11.3) allows incremental adoption of the ACP in a network where ACP can be deployed on edge areas, but not across the core that is connecting those edges.

In such a setup, each edge network, such as a branch or campus of an enterprise network has a disjointed ACP to which one or more unique Zone-IDs are assigned: ACP nodes registered for a specific ACP zone have to receive ACP Zone Addressing Sub-scheme addresses, for example by virtue of configuring for each such zone one or more ACP Registrars with that Zone-ID. All the Registrars for these ACP Zones need to get ACP certificates from CAs relying on a common set of TA and of course the same ACP domain name.

These ACP zones can first be brought up as separate networks without any connection between them and/or they can be connected across a non-ACP enabled core network through various non-autonomic operational practices. For example, each separate ACP Zone can have an edge node that is a layer 3 VPN PE (MPLS or IPv6 layer 3 VPN), where a complete non-autonomic ACP-Core VPN is created by using the ACP VRFs and exchanging the routes from those ACP VRFs across the VPNs non-autonomic routing protocol(s).

While such a setup is possible with any ACP addressing sub-scheme, the ACP-Zone Addressing sub-scheme makes it easy to configure and scalable for any VPN routing protocols because every ACP zone would only need to indicate one or more /64 ACP Zone Addressing prefix routes into the ACP-Core VPN as opposed to routes for every individual ACP node as required in the other ACP addressing schemes.

Note that the non-autonomous ACP-Core VPN would require additional extensions to propagate GRASP messages when GRASP discovery is desired across the zones.

For example, one could set up on each Zone edge router a remote ACP tunnel to a GRASP hub. The GRASP hub could be implemented at the application level and could run in the NOC of the network. It would serve to propagate GRASP announcements between ACP Zones and/or generate GRASP announcements for NOC services.

Such a partial deployment may prove to be sufficient or could evolve to become more autonomous through future standardized or non-standardized enhancements, for example by allowing GRASP messages to be propagated across the layer 3 VPN, leveraging for example L3VPN Multicast support.

Finally, these partial deployments can be merged into a single contiguous complete autonomous ACP (given appropriate ACP support across the core) without changes in the crypto material, because the node's ACP certificates are from a single ACP.

#### 9.5. Configuration and the ACP (summary)

There is no desirable configuration for the ACP. Instead, all parameters that need to be configured in support of the ACP are limitations of the solution, but they are only needed in cases where not all components are made autonomic. Wherever this is necessary, it relies on pre-existing mechanisms for configuration such as CLI or YANG ([RFC7950]) data models.

The most important examples of such configuration include:

- \* When ACP nodes do not support an autonomic way to receive an ACP certificate, for example BRSKI, then such certificate needs to be configured via some pre-existing mechanisms outside the scope of this specification. Today, routers have typically a variety of mechanisms to do this.
- \* Certificate maintenance requires PKI functions. Discovery of these functions across the ACP is automated (see Section 6.2.5), but their configuration is not.

- \* When non-ACP capable nodes such as pre-existing NMS need to be physically connected to the ACP, the ACP node to which they attach needs to be configured with ACP-connect according to Section 8.1. It is also possible to use that single physical connection to connect both to ACP and the Data-Plane of the network as explained in Section 8.1.4.
- \* When devices are not autonomically bootstrapped, explicit configuration to enable the ACP needs to be applied. See Section 9.3.
- \* When the ACP needs to be extended across interfaces other than L2, the ACP as defined in this document cannot autodiscover candidate neighbors automatically. Remote neighbors need to be configured, see Section 8.2.

Once the ACP is operating, any further configuration for the Data-Plane can be configured more reliably across the ACP itself because the ACP provides addressing and connectivity (routing) independent of the Data-Plane itself. For this, the configuration methods simply need to also allow to operate across the ACP VRF - NETCONF, SSH or any other method.

The ACP also provides additional security through its hop-by-hop encryption for any such configuration operations: Some legacy configuration methods (SNMP, TFTP, HTTP) may not use end-to-end encryption, and most of the end-to-end secured configuration methods still allow for easy passive observation along the path about configuration taking place (transport flows, port numbers, IP addresses).

The ACP can and should equally be used as the transport to configure any of the aforementioned non-autonomic components of the ACP, but in that case, the same caution needs to be exercised as with Data-Plane configuration without ACP: Misconfiguration may cause the configuring entity to be disconnected from the node it configures - for example when incorrectly unconfiguring a remote ACP neighbor through which the configured ACP node is reached.

## 10. Summary: Benefits (Informative)

### 10.1. Self-Healing Properties

The ACP is self-healing:

- \* New neighbors will automatically join the ACP after successful validation and will become reachable using their unique ULA address across the ACP.

- \* When any changes happen in the topology, the routing protocol used in the ACP will automatically adapt to the changes and will continue to provide reachability to all nodes.
- \* The ACP tracks the validity of peer certificates and tears down ACP secure channels when a peer certificate has expired. When short-lived certificates with lifetimes in the order of OCSP/CRL refresh times are used, then this allows for removal of invalid peers (whose certificate was not renewed) at similar speeds as when using OCSP/CRL. The same benefit can be achieved when using CRL/OCSP, periodically refreshing the revocation information and also tearing down ACP secure channels when the peer's (long-lived) certificate is revoked. There is no requirement against ACP implementations to require this enhancement though to keep the mandatory implementations simpler.

The ACP can also sustain network partitions and mergers. Practically all ACP operations are link local, where a network partition has no impact. Nodes authenticate each other using the domain certificates to establish the ACP locally. Addressing inside the ACP remains unchanged, and the routing protocol inside both parts of the ACP will lead to two working (although partitioned) ACPs.

There are few central dependencies: A CRL may not be available during a network partition; a suitable policy to not immediately disconnect neighbors when no CRL is available can address this issue. Also, an ACP Registrar or Certification Authority might not be available during a partition. This may delay renewal of certificates that are to expire in the future, and it may prevent the enrollment of new nodes during the partition.

Highly resilient ACP designs can be built by using ACP Registrars with embedded sub-CA, as outlined in Section 9.2.4. As long as a partition is left with one or more of such ACP Registrars, it can continue to enroll new candidate ACP nodes as long as the ACP Registrar's sub-CA certificate does not expire. Because the ACP addressing relies on unique Registrar-IDs, a later re-merge of partitions will also not cause problems with ACP addresses assigned during partitioning.

After a network partition, a re-merge will just establish the previous status, certificates can be renewed, the CRL is available, and new nodes can be enrolled everywhere. Since all nodes use the same TA, a re-merge will be smooth.

Merging two networks with different TA requires the ACP nodes to trust the union of TA. As long as the routing-subdomain hashes are different, the addressing will not overlap. Accidentally, overlaps will only happen in the unlikely event of a 40-bit hash collision in SHA256 (see Section 6.11). Note that the complete mechanisms to merge networks is out of scope of this specification.

It is also highly desirable for implementation of the ACP to be able to run it over interfaces that are administratively down. If this is not feasible, then it might instead be possible to request explicit operator override upon administrative actions that would administratively bring down an interface across which the ACP is running. Especially if bringing down the ACP is known to disconnect the operator from the node. For example, any such down administrative action could perform a dependency check to see if the transport connection across which this action is performed is affected by the down action (with default RPL routing used, packet forwarding will be symmetric, so this is actually possible to check).

## 10.2. Self-Protection Properties

### 10.2.1. From the outside

As explained in Section 6, the ACP is based on secure channels built between nodes that have mutually authenticated each other with their domain certificates. The channels themselves are protected using standard encryption technologies such as DTLS or IPsec which provide additional authentication during channel establishment, data integrity and data confidentiality protection of data inside the ACP and in addition, provide replay protection.

Attacker will not be able to join the ACP unless they have a valid ACP certificate. On-path attackers without a valid ACP certificate cannot inject packets into the ACP due to ACP secure channels. They can also not decrypt ACP traffic except if they can crack the encryption. They can attempt behavioral traffic analysis on the encrypted ACP traffic.

The degree to which compromised ACP nodes can impact the ACP depends on the implementation of the ACP nodes and their impairment. When an attacker has only gained administrative privileges to configure ACP nodes remotely, the attacker can disrupt the ACP only through one of the few configuration options to disable it, see Section 9.3, or by configuring of non-autonomic ACP options if those are supported on the impaired ACP nodes, see Section 8. Injecting or extracting traffic into/from an impaired ACP node is only possible when an impaired ACP node supports ACP connect (see Section 8.1) and the attacker can control traffic into/from one of the ACP nodes interfaces, such as by having physical access to the ACP node.

The ACP also serves as protection (through authentication and encryption) for protocols relevant to OAM that may not have secured protocol stack options or where implementation or deployment of those options fail on some vendor/product/customer limitations. This includes protocols such as SNMP ([RFC3411]), NTP ([RFC5905]), PTP ([IEEE-1588-2008]), DNS ([RFC3596]), DHCPv6 ([RFC3315]), syslog ([RFC3164]), RADIUS ([RFC2865]), Diameter ([RFC6733]), TACACS ([RFC1492]), IPFIX ([RFC7011]), Netflow ([RFC3954]) – just to name a few. Not all of these protocol references are necessarily the latest version of protocols but versions that are still widely deployed.

Protection via the ACP secure hop-by-hop channels for these protocols is meant to be only a stopgap though: The ultimate goal is for these and other protocols to use end-to-end encryption utilizing the domain certificate and rely on the ACP secure channels primarily for zero-touch reliable connectivity, but not primarily for security.

The remaining attack vector would be to attack the underlying ACP protocols themselves, either via directed attacks or by denial-of-service attacks. However, as the ACP is built using link-local IPv6 addresses, remote attacks from the Data-Plane are impossible as long as the Data-Plane has no facilities to remotely send IPv6 link-local packets. The only exceptions are ACP connected interfaces which require higher physical protection. The ULA addresses are only reachable inside the ACP context, therefore, unreachable from the Data-Plane. Also, the ACP protocols should be implemented to be attack resistant and not consume unnecessary resources even while under attack.

#### 10.2.2. From the inside

The security model of the ACP is based on trusting all members of the group of nodes that receive an ACP certificate for the same domain. Attacks from the inside by a compromised group member are therefore the biggest challenge.

Group members must be protected against attackers so that there is no easy way to compromise them, or use them as a proxy for attacking other devices across the ACP. For example, management plane functions (transport ports) should only be reachable from the ACP but not the Data-Plane. Especially for those management plane functions that have no good protection by themselves because they do not have secure end-to-end transport and to whom ACP not only provides automatic reliable connectivity but also protection against attacks. Protection across all potential attack vectors is typically easier to do in devices whose software is designed from the ground up with ACP in mind than with legacy software based systems where the ACP is added on as another feature.

As explained above, traffic across the ACP should still be end-to-end encrypted whenever possible. This includes traffic such as GRASP, EST and BRSKI inside the ACP. This minimizes man in the middle attacks by compromised ACP group members. Such attackers cannot eavesdrop or modify communications, they can just filter them (which is unavoidable by any means).

See Appendix A.9.8 for further considerations how to avoid and deal with compromised nodes.

### 10.3. The Administrator View

An ACP is self-forming, self-managing and self-protecting, therefore has minimal dependencies on the administrator of the network. Specifically, since it is (intended to be) independent of configuration, there is only limited scope for configuration errors on the ACP itself. The administrator may have the option to enable or disable the entire approach, but detailed configuration is not possible. This means that the ACP must not be reflected in the running configuration of nodes, except a possible on/off switch (and even that is undesirable).

While configuration (except for Section 8 and Section 9.2) is not possible, an administrator must have full visibility of the ACP and all its parameters, to be able to do trouble-shooting. Therefore, an ACP must support all show and debug options, as for any other network function. Specifically, a network management system or controller must be able to discover the ACP, and monitor its health. This visibility of ACP operations must clearly be separated from visibility of Data-Plane so automated systems will never have to deal with ACP aspects unless they explicitly desire to do so.

Since an ACP is self-protecting, a node not supporting the ACP, or without a valid domain certificate cannot connect to it. This means that by default a traditional controller or network management system cannot connect to an ACP. See Section 8.1.1 for more details on how to connect an NMS host into the ACP.

## 11. Security Considerations

A set of ACP nodes with ACP certificates for the same ACP domain and with ACP functionality enabled is automatically "self-building": The ACP is automatically established between neighboring ACP nodes. It is also "self-protecting": The ACP secure channels are authenticated and encrypted. No configuration is required for this.

The self-protecting property does not include workarounds for non-autonomic components as explained in Section 8. See Section 10.2 for details of how the ACP protects itself against attacks from the outside and to a more limited degree from the inside as well.

However, the security of the ACP depends on a number of other factors:

- \* The usage of domain certificates depends on a valid supporting PKI infrastructure. If the chain of trust of this PKI infrastructure is compromised, the security of the ACP is also compromised. This is typically under the control of the network administrator.
- \* ACP nodes receive their certificates from ACP registrars. These ACP registrars are security critical dependencies of the ACP: Procedures and protocols for ACP registrars are outside the scope of this specification as explained in Section 6.11.7.1, only requirements against the resulting ACP certificates are specified.
- \* Every ACP registrar (for enrollment of ACP certificates) and ACP EST server (for renewal of ACP certificates) is a security critical entity and its protocols are security critical protocols. Both need to be hardened against attacks, similar to a CA and its protocols. A malicious registrar can enroll malicious nodes to an ACP network (if the CA delegates this policy to the registrar) or break ACP routing for example by assigning duplicate ACP address assignment to ACP nodes via their ACP certificates.
- \* ACP nodes that are ANI nodes rely on BRSKI as the protocol for ACP registrars. For ANI type ACP nodes, the security considerations of BRSKI apply. It enables automated, secure enrollment of ACP certificates.
- \* BRSKI and potentially other ACP registrar protocol options require that nodes have an (X.509v3 based) IDevID. IDevIDs are an option for ACP registrars to securely identify candidate ACP nodes that should be enrolled into an ACP domain.

- \* For IDevIDs to securely identify the node to which it IDevID is assigned, the node needs to (1) utilize hardware support such as a Trusted Platform Module (TPM) to protect against extraction/cloning of the private key of the IDevID and (2) a hardware/software infrastructure to prohibit execution of non-authenticated software to protect against malicious use of the IDevID.
- \* Like the IDevID, the ACP certificate should equally be protected from extraction or other abuse by the same ACP node infrastructure. This infrastructure for IDevID and ACP certificate is beneficial independent of the ACP registrar protocol used (BRSKI or other).
- \* Renewal of ACP certificates requires support for EST, therefore the security considerations of [RFC7030] related to certificate renewal/rekeying and TP renewal apply to the ACP. EST security considerations when using other than mutual certificate authentication do not apply nor do considerations for initial enrollment via EST apply, except for ANI type ACP nodes because BRSKI leverages EST.
- \* A malicious ACP node could declare itself to be an EST server via GRASP across the ACP if malicious software could be executed on it. CA should therefore authenticate only known trustworthy EST servers, such as nodes with hardware protections against malicious software. When Registrars use their ACP certificate to authenticate towards a CA, the id-kp-cmcRA [RFC6402] extended key usage attribute allows the CA to determine that the ACP node was permitted during enrollment to act as an ACP registrar. Without the ability to talk to the CA, a malicious EST server can still attract ACP nodes attempting to renew their keying material, but they will fail to perform successful renewal of a valid ACP certificate. The ACP node attempting to use the malicious EST server can then continue to use a different EST server, and log a failure against a malicious EST server.
- \* Malicious on-path ACP nodes may filter valid EST server announcements across the ACP, but such malicious ACP nodes could equally filter any ACP traffic such as the EST traffic itself. Either attack requires the ability to execute malicious software on an impaired ACP node though.
- \* In the absence of malicious software injection, an attacker that can misconfigure an ACP node which is supporting EST server functionality could attempt to configure a malicious CA. This would not result in the ability to successfully renew ACP certificates, but it could result in DoS attacks by becoming an EST server and making ACP nodes attempting their ACP certificate renewal via this impaired ACP node. This problem can be avoided when the EST server implementation can verify that the CA configured is indeed providing renewal for certificates of the node's ACP. The ability to do so depends on the EST-Server to CA protocol, which is outside the scope of this document.

In summary, attacks against the PKI/certificate dependencies of the ACP can be minimized by a variety of hardware/software components including options such as TPM for IDevID/ACP-certificate, prohibitions against execution of non-trusted software and design aspects of the EST Server functionality for the ACP to eliminate configuration level impairment.

Because ACP peers select one out of potentially more than one mutually supported ACP secure channel protocols via the approach described in Section 6.6, ACP secure channel setup is subject to downgrade attacks by MITM attackers. This can be discovered after such an attack by additional mechanisms described in Appendix A.9.9. Alternatively, more advanced channel selection mechanisms can be devised. [RFC-Editor: Please remove the following sentence]. See [ACPDRAFT] Appendix B.1. Both options are out of scope of this document.

The security model of the ACP as defined in this document is tailored for use with private PKI. The TA of a private PKI provide the security against maliciously created ACP certificates to give access to an ACP. Such attacks can create fake ACP certificates with correct looking AcpNodeNames, but those certificates would not pass the certificate path validation of the ACP domain membership check (see Section 6.2.3, point 2).

[RFC-Editor: please remove the following paragraph].

Using public CA is out of scope of this document. See [ACPDRAFT], Appendix B.3 for further considerations.

There is no prevention of source-address spoofing inside the ACP. This implies that if an attacker gains access to the ACP, it can spoof all addresses inside the ACP and fake messages from any other node. New protocol/services run across the ACP should therefore use end-to-end authentication inside the ACP. This is already done by GRASP as specified in this document.

The ACP is designed to enable automation of current network management and future autonomic peer-to-peer/distributed network automation. Any ACP member can send ACP IPv6 packet to other ACP members and announce via ACP GRASP services to all ACP members without dependency against centralized components.

The ACP relies on peer-to-peer authentication and authorization using ACP certificates. This security model is necessary to enable the autonomic ad-hoc any-to-any connectivity between ACP nodes. It provides infrastructure protection through hop by hop authentication and encryption - without relying on third parties. For any services

where this complete autonomic peer-to-peer group security model is appropriate, the ACP certificate can also be used unchanged. For example, for any type of Data-Plane routing protocol security.

This ACP security model is designed primarily to protect against attack from the outside, but not against attacks from the inside. To protect against spoofing attacks from compromised on-path ACP nodes, end-to-end encryption inside the ACP is used by new ACP signaling: GRASP across the ACP using TLS. The same is expected from any non-legacy services/protocols using the ACP. Because no group-keys are used, there is no risk for impacted nodes to access end-to-end encrypted traffic from other ACP nodes.

Attacks from impacted ACP nodes against the ACP are more difficult than against the Data-Plane because of the autoconfiguration of the ACP and the absence of configuration options that could be abused that allow to change/break ACP behavior. This is excluding configuration for workaround in support of non-autonomic components.

Mitigation against compromised ACP members is possible through standard automated certificate management mechanisms including revocation and non-renewal of short-lived certificates. In this version of the specification, there are no further optimization of these mechanisms defined for the ACP (but see Appendix A.9.8).

Higher layer service built using ACP certificates should not solely rely on undifferentiated group security when another model is more appropriate/more secure. For example, central network configuration relies on a security model where only few especially trusted nodes are allowed to configure the Data-Plane of network nodes (CLI, NETCONF). This can be done through ACP certificates by differentiating them and introduce roles. See Appendix A.9.5.

Operators and provisioning software developers need to be aware of how the provisioning/configuration of network devices impacts the ability of the operator / provisioning software to remotely access the network nodes. By using the ACP, most of the issues of configuration/provisioning caused loss of connectivity for remote provisioning/configuration will be eliminated, see Section 6. Only few exceptions such as explicit physical interface down configuration will be left Section 9.3.2.

Many details of ACP are designed with security in mind and discussed elsewhere in the document:

IPv6 addresses used by nodes in the ACP are covered as part of the node's domain certificate as described in Section 6.2.2. This allows even verification of ownership of a peer's IPv6 address when using a connection authenticated with the domain certificate.

The ACP acts as a security (and transport) substrate for GRASP inside the ACP such that GRASP is not only protected by attacks from the outside, but also by attacks from compromised inside attackers - by relying not only on hop-by-hop security of ACP secure channels, but adding end-to-end security for those GRASP messages. See Section 6.9.2.

ACP provides for secure, resilient zero-touch discovery of EST servers for certificate renewal. See Section 6.2.5.

ACP provides extensible, auto-configuring hop-by-hop protection of the ACP infrastructure via the negotiation of hop-by-hop secure channel protocols. See Section 6.6.

The ACP is designed to minimize attacks from the outside by minimizing its dependency against any non-ACP (Data-Plane) operations/configuration on a node. See also Section 6.13.2.

In combination with BRSKI, ACP enables a resilient, fully zero-touch network solution for short-lived certificates that can be renewed or re-enrolled even after unintentional expiry (e.g., because of interrupted connectivity). See Appendix A.2.

Because ACP secure channels can be long lived, but certificates used may be short lived, secure channels, for example built via IPsec need to be terminated when peer certificates expire. See Section 6.8.5.

Section 7.2 describes how to implement a routed ACP topology operating on what effectively is a large bridge-domain when using L3/L2 routers that operate at L2 in the Data-Plane. In this case, the ACP is subject to much higher likelihood of attacks by other nodes "stealing" L2 addresses than in the actual routed case. Especially when the bridged network includes non-trusted devices such as hosts. This is a generic issue in L2 LANs. L2/L3 devices often already have some form of "port security" to prohibit this. They rely on NDP or DHCP learning of which port/MAC-address and IPv6 address belong together and block MAC/IPv6 source addresses from wrong ports. This type of function needs to be enabled to prohibit DoS attacks and specifically to protect the ACP. Likewise the GRASP DULL instance needs to ensure that the IPv6 address in the locator-option matches the source IPv6 address of the DULL GRASP packet.

## 12. IANA Considerations

This document defines the "Autonomic Control Plane".

For the ANIMA-ACP-2020 ASN.1 module, IANA is asked to register value IANA1 for "id-mod-anima-acpnodename-2020" in the "SMI Security for PKIX Module Identifier" (1.3.6.1.5.5.7.0) registry.

For the otherName / AcpNodeName, IANA is asked to register a value for IANA2 for id-on-AcpNodeName in the "SMI Security for PKIX Other Name Forms" (1.3.6.1.5.5.7.8) registry.

The IANA is requested to register the value "AN\_ACP" (without quotes) to the GRASP Objectives Names Table in the GRASP Parameter Registry. The specification for this value is this document, Section 6.4.

The IANA is requested to register the value "SRV.est" (without quotes) to the GRASP Objectives Names Table in the GRASP Parameter Registry. The specification for this value is this document, Section 6.2.5.

Explanation: This document chooses the initially strange looking format "SRV.<service-name>" because these objective names would be in line with potential future simplification of the GRASP objective registry. Today, every name in the GRASP objective registry needs to be explicitly allocated with IANA. In the future, this type of objective names could be considered to be automatically registered in that registry for the same service for which a <service-name> is registered according to [RFC6335]. This explanation is solely informational and has no impact on the requested registration.

The IANA is requested to create an ACP Parameter Registry with currently one registry table - the "ACP Address Type" table.

"ACP Address Type" Table. The value in this table are numeric values 0...3 paired with a name (string). Future values MUST be assigned using the Standards Action policy defined by [RFC8126]. The following initial values are assigned by this document:

0: ACP Zone Addressing Sub-Scheme (ACP RFC Section 6.11.3)

1: ACP Vlong Addressing Sub-Scheme (ACP RFC Section 6.11.5) / ACP Manual Addressing Sub-Scheme (ACP RFC Section 6.11.4)

### 13. Acknowledgements

This work originated from an Autonomic Networking project at Cisco Systems, which started in early 2010. Many people contributed to this project and the idea of the Autonomic Control Plane, amongst which (in alphabetical order): Ignas Bagdonas, Parag Bhide, Balaji BL, Alex Clemm, Yves Hertoghs, Bruno Klauser, Max Pritikin, Michael Richardson, Ravi Kumar Vadapalli.

Special thanks to Brian Carpenter, Elwyn Davies, Joel Halpern and Sheng Jiang for their thorough reviews.

Many thanks Ben Kaduk, Roman Danyliv and Eric Rescorla for their thorough SEC AD reviews, Russ Housley and Erik Kline for their reviews and to Valery Smyslov, Tero Kivinen, Paul Wouters and Yoav Nir for review of IPsec and IKEv2 parameters and helping to understand those and other security protocol details better. Thanks for Carsten Borman for CBOR/CDDL help.

Further input, review or suggestions were received from: Rene Struik, Benoit Claise, William Atwood and Yongkang Zhang.

### 14. Contributors

For all things GRASP including validation code, ongoing document text support and technical input.

Brian Carpenter  
School of Computer Science  
University of Auckland  
PB 92019  
Auckland 1142  
New Zealand

Email: [brian.e.carpenter@gmail.com](mailto:brian.e.carpenter@gmail.com)

For RPL contributions and all things BRSKI/bootstrap including validation code, ongoing document text support and technical input.

Michael C. Richardson  
Sandelman Software Works

Email: [mcr+iETF@sandelman.ca](mailto:mcr+iETF@sandelman.ca)  
URI: <http://www.sandelman.ca/mcr/>

For the RPL technology choices and text.

Pascal Thubert  
Cisco Systems, Inc  
Building D  
45 Allee des Ormes - BP1200  
06254 MOUGINS - Sophia Antipolis  
France

Phone: +33 497 23 26 34  
Email: pthubert@cisco.com

## 15. Change log [RFC-Editor: Please remove]

This document was developed on <https://github.com/anima-wg/autonomic-control-plane/tree/master/draft-ietf-anima-autonomic-control-plane>. That github repository also contains the document review/reply emails.

### 15.1. Summary of changes since entering IESG review

This text replaces the prior changelog with a summary to provide guidance for further IESG review.

Please see revision -21 for the individual changelogs of prior versions .

#### 15.1.1. Reviews (while in IESG review status) / status

This document entered IESG review with version -13. It has since seen the following reviews:

IESG: Original owner/Yes: Terry Manderson (INT).

IESG: No Objection: Deborah Brungard (RTG), Alissa Cooper (GEN), Warren Kumari (OPS), Mirja Kuehlewind (TSV), Alexey Melnikov (ART), Adam Roach (ART).

IESG: No Objection, not counted anymore as they have left IESG: Ben Campbell (ART), Spencer Dawkins (TSV).

IESG: Open DISCUSS hopefully resolved by this version: Eric Rescorla (SEC, left IESG), Benjamin Kaduk (SEC).

Other: Michael Richardson (WG), Brian Carpenter (WG), Pascal Thubert (WG), Frank Xialiang (WG), Elwyn Davies (GEN), Joel Halpern (RTGdir), Yongkang Zhang (WG), William Atwood (WG).

## 15.1.2. BRSKI / ACP registrar related enhancements

Only after ACP entered IESG review did it become clear that the in-progress BRSKI document would not provide all the explanations needed for ACP registrars as expected earlier by ACP authors. Instead, BRSKI will only specify a subset of required ACP behavior related to certificate handling and registrar. There, it became clear that the ACP draft should specify generic ACP registrar behavior independent of BRSKI so ACP could be implemented with or without BRSKI and any manual/proprietary or future standardized BRSKI alternatives (for example via NETCONF) would understand the requirements for ACP registrars and its certificate handling.

This led to additional text about ACP registrars in the ACP document:

1. Defined relationship ACP / ANI (ANI = ACP + BRSKI).

6.1.4 (new) Overview of TA required for ACP.

6.1.5.5 Added explanations/requirements for Re-enrollment.

6.10.7 Normative requirements for ACP registrars (BRSKI or not).

10.2 Operational expectations against ACP registrars (BRSKI or not).

## 15.1.3. Normative enhancements since start of IESG review

In addition to above ACP registrar / BRSKI related enhancements there is a range of minor normative (also explanatory) enhancements since the start of IESG review:

6.1.1 Hex digits in ACP domain information field now upper-case (no specific reason except that both options are equally good, but capitalized ones are used in rfc5234).

6.1.5.3 Added explanations about CRLs.

6.1.5.6 Added explanations of behavior under failing certificates.

6.1.2 Allow ACP address '0' in ACP domain information field: presence of address indicates permission to build ACP secure channel to node, 0 indicates that the address of the node is assigned by (future) other means than certificate. Non-autonomic nodes have no address at all (that was in -13), and can only connect via ACP connect interfaces to ACP.

6.1.3 Distinction of real ACP nodes (with address) and those with domain certificate without address added as a new rule to ACP domain membership check.

6.6 Added throttling of secure-channel setup attempts.

6.11.1.14 Removed requirement to handle unknown destination ACP traffic in low-end nodes that would never be RPL roots.

6.12.5 Added recommendation to use IPv6 DAD.

6.1.1, 6.7.1.1, 6.7.2, 6.7.3, 6.8.2 Various refined additional certificate, secure channel protocol (IPsec/IKEv2 and DTLS) and ACP GRASP TLS protocol parameter requirements to ensure interoperating implementations (from SEC-AD review).

#### 15.1.4. Explanatory enhancements since start of IESG review

Beyond the functional enhancements from the previous two sections, the majority of changes since -13 are additional explanations from review feedback, textual nits and restructuring - with no functional requirement additions/changes.

1.1 Added "applicability and scope" section with summarized explanations.

2. Added in-band vs. out-of-band management definitions.

6.1.2 (was 6.1.1) expanded explanations of reasoning for elements of the ACP domain information field.

6.1.3 refined explanations of ACP domain membership check and justifications for it.

6.5 Elaborated step-by-step secure channel setup.

6.10 Additional explanations for addressing modes, additional table of addressing formats (thanks MichaelR).

6.10.5 introduced 'F' bit position as a better visual representation in the Vlong address space.

6.11.1.1 extensive overhaul to improve readability of use of RPL (from IESG feedback of non-routing/RPL experts).

6.12.2 Added caution about unconfiguring Data-Plane IPv6 addresses and impact to ACP (limitation of current ACP design, and pointint to more details in 10.2).

10.4 New explanations / summary of configurations for ACP (aka: all config is undesirable and only required for integrating with non-autonomic components, primarily ACP-connect and Registrars).

11. Textually enhanced / better structured security considerations section after IESG security review.

A. (new) Moved all explanations and discussions about futures from section 10 into this new appendix. This text should not be removed because it captures a lot of repeated asked questions in WG and during reviews and from users, and also captures ideas for some likely important followup work. But none of this is relevant to implementing (section 6) and operating (section 10) the ACP.

#### 15.2. draft-ietf-anima-autonomic-control-plane-30

-29 did pass all IESG DISCUSS. This version cleans up remaining comments.

Planned to be removed section Appendix A.6 was moved into new Appendix B.1 to be amended by further A.2, A.3 containing text felt to be unfit for publication in RFC (see below). Added reference to this last draft, and referencing those sections ([ACPDRAFT]).

Final discussion with responsible AD (Eric Vyncke): marked all references to [ACPDRAFT] as to be removed from RFC, as this would be too unconventional. Likewise also [ACPDRAFT] reference itself. Added explanation to appendix B.

Comments from Erik Kline:

2. Fine tuned ULA definition.

Comments Michael Richardson / Eric Vyncke.

6.2.4. / 11. Removed text arguing ability how to use public CA (or not). Replaced with reference to new [ACPDRAFT] section B.3 (not in RFC) that explains current state of understanding (unfinished).

B.3 New text detailing authors understanding of use of public CA (will not be in RFC).

Comments/proposals from Ben Kaduk:

Various: Replaced RFC4492 with RFC8422 which is superceding it.

6.1 Text fix for hash strength 384 bits (from SHA384); Text fix for ec\_point\_format extension.

6.2.1 Text fixup. Removed requirements for ECDH support in certificate, instead merely explaining the dependencies required IF this is desired (educational).

6.2.5.4. Fine tuning 2 sentences.

6.3.2. (ACP domain membership check) Add reference to ACPDRAFT B.2 explaining why ACP domain membership does not validate ACP address of the connection.

6.4. Downgraded SHOULD to MAY in new -29 suggestion how to deal with DoS attacks with many GRASP announcements. Will also separately ask TSV ADs.

6.4. Fixed extension points in CDDL objective-value definitions (with help from Carsten/Brian).

9.3.5.2. Added explanation when ACP greenfield state ends, and refined text explaining how to deal with this.

11. removed duplicate paragraph (first, kept paragraph was the fixed up, improved correct version).

11. Added references to ACPDRAFT B.1, B.2 as possible future solutions for downgrade attacks.

12. Fixed up text for IANA code point allocation request.

A.6 - removed.

A.9.9 - added one explanatory intro paragraph to makes it easier to distinguish this option from the B.1 considerations.

B.1 - new text suggested from Ben, replacing A.6 (will not be in RFC).

B.2 - new text discussing why there is no network layer address verification in ACP domain membership check (will not be in RFC).

B.4 - Text discussing DULL GRASP attacks via port sweeps and what do do against it.

Other.

1. Added sentence about FCC outage report from June as example for in-band management.

15. added reference to github where document was developed (removed in RFC, part of changelog).

### 15.3. draft-ietf-anima-autonomic-control-plane-29

Comments from Robert Wilton:

Improved several textual nits.

Discuss/Comments from Erik Kline:

Editorial suggestions and nits. Thanks!.

6.1.3 Added text about how/why rsub is irrelevant for domain membership check.

6.3 Added extension points to AN\_ACP DULL GRASP objective because for example ACP domain certificate could be a nice optional additional parameter and prior syntax would have forced us to encode into separate objective unnecessarily.

6.7 Using RFC8415 terminology for exponential backoff parameters.

6.11.2 Amended ACP Sub-Addressing table with future code points, explanations and prefix announced into RPL.

6.12.1.11. Reworked text to better explain how black hole route works and added explanation for prefix for manual address scheme.

8.1.3. Reworked explanation of RIOS for ACP connect interfaces for Type C vs. Type A/B hosts.

8.1.4. Added explanation how this "VRF select" option is required for auto-attachment of Type A/B hosts to ACP and other networks.

Discuss/Comments from Barry Leiba:

Various editorial nits - thanks.

6.1 New section pulling in TLS requirements, no need anymore to duplicat for ACP GRASP, EST, BRSKI (ACP/ANI nodes) and (if desired) OCSP/CRLDP. Added rule to start use secure channel only after negotiation has finished. Added rules not to optimize negotiation across multiple L2 interfaces to the same peer.

6.6 Changed role names in secure channel negotiation process: Alice/Bob -> Decider/Follower. Explanation enhancements. Added definition for ACP nodes with "0" address.

6.8.3 Improved explanation how IKEv2 forces preference of IPsec over GRE due to ACP IPsec profiles being Tunneled vs. Transport.

6.8.4 Limited mentioning of DTLS version requirements to this section.

6.9.2 Removed TLS requirements, they are now in 6.1.

6.10.6 Removed explanation of IANA allocation requirement. Redundant - already in IANA section, and was seen as confusing.

8.1.1 Clarified that there can be security impacts when weakening directly connected address RPF filtering for ACP connect interfaces.

Discuss/Comments from Ben Kaduk:

Many good editorial improvements - thanks!.

5. added explanation of what to do upon failed secure channel establishment.

6.1.1. refined/extended cert public key crypto algo and better distinguished algo for the keys of the cert and the key of the signer.

6.1.1. and following: explicitly defining "serialNumber" to be the X.520 subject name serialNumber, not the certificate serial Number.

6.1.1. emphasize additional authorization step for EST servers (id-kp-cmCRA).

6.1.2 changed AcpNodeName ABNF to again use 32HEXDIG instead of self-defined variation, because authors overlooked that ABNF is case agnostic (which is fine). Added recommendation to encode as lower case. Added full ABNF encoding for extensions (any characters as in "atoms" except the new "+" separator).

6.1.5.3. New text to explain reason for use of HTTPS (instead of HTTP) for CRLDP and when and how to use HTTPS then.

6.1.5.5. added text explaining why/how and when to maintain TA data upon failing cert renewal (one version with BRSKI, one version with other, see secure bootstrap protocols).

6.3. new text and requirement about the signaling of transport ports in DULL GRASP - benefits (no well-known ports required), and problems (additional DoS attack vector, albeit not worse than pre-existing ones, depending on setup of L2 subnets.).

6.7.3.1.1. Specified AUTH\_HMAC\_SHA2\_256\_128 (as the ESP authentication algorithm).

6.8.2. Added recommendations for TLS\_AES\_256\_GCM\_SHA384, TLS\_CHACHA20\_POLY1305\_SHA256 when supporting TLS 1.3.

8.2.2. Added explanation about downgrade attack across configured ACP tunnels and what to do against it.

9.3.5.2. Rewrote most of section as it originally was too centric on BRSKI. Should now well describe expectations against automated bootstrap. Introduces new requirement not to call node as in support of ANI if is ALSO has TOFU bootstrap.

11. Expanded text about malicious EST servers. Added paragraph about ACP secure channel downgrade attacks. Added paragraphs about private PKI as a core to allow security against fake certificates, added paragraph about considerations/problems when using public PI.

A.10.9 New appendix suggesting how to discover ACP secure channel negotiation downgrade attacks.

Discuss from Roman Danyliw:

6.1.5.1 - Added requirement to only announce SRV.est when a working CA connection.

15 - Amended security considerations with text about registrar dependencies, security of IDevID/ACP-certificate, EST-Server and GRASP for EST server discovery.

Other:

Conversion to XML v3. Solved empty () taxonomy xref problems. Various formatting fixes for v3.

Added contributors section.

15.4. draft-ietf-anima-autonomic-control-plane-28

IESG review Roman Danyliw:

6. Requested additional text elaborating misconfiguration plus attack vectors.

6.1.3.1 Added paragraph about unsecured NTP as basis for time in the absence of other options.

6.7.2 reworded text about additional secure channel protocol requirements.

6.7.3.1.2. Added requirement for ACP nodes supporting IKEv2 to support RFC8247 (not sure how that got dropped from prior versions.

Replaced minimum crypto requirements definition via specific AES options with more generic "symmetric key/hash strength" requirements.

6.10.7.3. Added example how to derive addressing scheme from IDevID (PID). Added explanation how to deal with non-persistent registrar address database (hint: it sucks or is wasteful, what did you expect).

8.1.1. Added explanation for 'Physical controlled/secured'.

8.1.5. Removed 'Physical controlled/secured' text, refer back to 8.1.1.

8.2.1. Fixed ABNF 'or' syntax line.

9.3.2. Added explanation of remote management problem with interface "down" type commands.

10.2.1. Added explanations for attacks from impaired ACP nodes.

11. Rewrote intro paragraph. Removed text referring to enrollment/registrars as they are out of scope of ACP (dependencies only).

11. Added note about need for new protocols inside ACP to use end-to-end authentication.

11. Rewrote paragraph about operator mistakes so as to be actionable. Operators must not make mistakes - but ACP minimizes the mistakes they can make.

ACP domain certificate -> ACP certificate.

Various other cosmetic edits (thanks!) and typo fixes (sorry for not running full spell check for every version. Will definitely do before RFC editor).

Other:

6.12.5.2.1./6.12.5.2.2. Added text explaining link breakage wrt. RTL (came about re-analyzing behavior after question about hop count).

Removed now unnecessary references for earlier rrc822Name otherName choice.

#### 15.5. draft-ietf-anima-autonomic-control-plane-27

Too many revisions with too many fixes. Lets do a one-word change revision for a change now if it helps to accelerate the review process.

Added "subjectAltName /" to make it unambiguous that AcpNodeName is indeed a SAN (from Russ).

#### 15.6. draft-ietf-anima-autonomic-control-plane-26

Russ Housley review of -25.

1.1 Explicit reference for TLS 1.2 RFC.

2. Changed term of "ACP Domain Information" to AcpNodeName (ASN.1) / acp-node-name (ABNF), also through rest of document.

2. Improved CA behavior definition. changed IDevID/LDevID to IDevID/LDevID certificate to be more unambiguous.

2. Changed definition of root CA to just refer to how its used in RFC7030 CA root key update, because thats the only thing relevant to ACP.

6.1.1 Moved ECDH requirement to end of text as it was not related to the subject of the initial paragrap. Likewise reference to CABFORUM.

6.1.1 Reduced cert key requirements to only be MUST for certs with 2048 RSA public key and P-256 curves. Reduced longer keys to SHOULD.

6.1.2 Changed text for conversion from rfc822Name to otherName / AcpNode, removed all the explanations of benefits coming with rfc822Name \*sob\* \*sob\* \*sob\*.

6.1.2.1 New ASN.1 definition for otherName / AcpNodeName.

6.1.3 Fixed up text. re the handling of missing connectivity for CRLDP / OCSP.

6.1.4 Fixed up text re. inability to use public CA to situation with otherName / AcpNodeName (no more ACME rfc822Name validation for us \*sob\* \*sob\* \*sob\*).

12. Added ASN.1 registration requests to IANA section.

Appenices. Minor changes for rfc822Name to otherName change.

Various minor verbal fixes/enhancements.

#### 15.7. draft-ietf-anima-autonomic-control-plane-25

Crypto parameter discuss from Valery Smyslov and Paul Wouters and resulting changes.

6.7.2 Moved Michael Richardson suggested diagnostic of signaling TA from IPsec section to this general requirements section and added explanation how this may be inappropriate if TA payload is considered secret by TA owner.

6.7.3.1 Added traffic selectors for native IPsec. Improved text explanation.

6.7.3.1.2 removed misleading text about signaling TA when using intermediate certs.

6.7.3.1.2 Removed requirement for 'PKCS #7 wrapped X.509 certificate' requirement on request of Valery Smyslov as it is not defined in RFC7296 and there are enough options mandated in RFC7296. Replaced with just informative text to educate readers who are not IPsec experts what the mandatory option in RFC7296 is that allows to signal certificates.

6.7.3.1.2 Added SHOULD requirement how to deal with CERTREQ so that 6.7.2 requirement for TA diagnostics will work in IKEv2 (ignoring CERTREQ is permitted by IKEv2). Added explanation how this will result in TA cert diagnostics.

6.7.3.1.2 Added requirement for IKEv2 to operate on link-local addresses for ACP so as to assume ACT cert as the only possible authenticator - to avoid potentially failing section from multiple available certs on a router.

6.7.3.1.2 fixed PKIX- style OID to ASN.1 object AlgorithmIdentifier (Paul).

6.7.3.2 Added IPsec traffic selectors for IPsec with GRE.

6.7.5 Added notion that IPsec/GRE MAY be preferred over IPsec/native. Luckily IPsec/native uses tunneling, whereas IPsec/GRE uses transport mode, and there is a long discuss whether it is permitted to even build IPsec connectings that only support transports instead of always being able to fall back to tunnel mode. Added explanatory paragraph why ACP nodes may prefer GRE over native (wonder how that was missing..).

9.1.1 Added section to explain need for secure channel peer diagnostics via signaling of TA. Four examples given.

Paul Wouters mentioned that ipkcs7 had to be used in some interop cases with windows CA, but that is an issue of ACP Registrar having to convert into PKCS#7 to talk to a windows CA, and this spec is not concerned with that, except to know that it is feasible, so not mentioned in text anywhere, just tracking discussion here in changelog.

Michael Richardson:

3.1.3 Added point in support of rfc822address that CA may not support to sign certificates with new attributes (such as new otherName).

Michael Richardson/Brian Carpenter fix:

6.1.5.1/6.3 Fixed GRASP examples.

Joe Halpern review:

1. Enhanced introduction text for in-band and of out-of-band, explaining how ACP is an in-band network aiming to achieve all possible benefits of an out-of-band network.

1. Comprehensive explanation for term Data-Plane as it is only logically following pre-established terminology on a fully autonomic node, when used for existing nodes augmented with ACP, Data-Plane has more functionality than usually associated with the term.

2. Removed explanatory text for Data-Plane, referring to section 1.

2. Reduced explanation in definition of in-band (management/signaling), out-of-band-signaling, now pointing to section 1.

5. Rewrote a lot of the steps (overview) as this text was not reviewed for long time. Added references to normative section for each step to hopefully avoid feedback of not explaining terms used (really not possible to give good summary without using forward references).

2. Separate out-of-band-management definition from virtual out-of-band-management definition (later one for ACP).

2. Added definitions for RPI and RPL.

6.1.1. added note about end-to-end authentication to distinguish channel security from overall ACP security model.

6.5 Fixed bugs in channel selection signaling step description (Alice vs. Bob).

6.7.1 Removed redundant channel selection explanation.

6.10.3 remove locator/identifier terminology from zone addressing scheme description (unnecessary), removed explanations (now in 9.4), simplified text, clarified requirement for Node-ID to be unique, recommend to use primarily zone 0.

6.10.3.1 Removed. Included a lot of insufficient suggestions for future standards extensions, most of it was wrong or would need to be revisited by WG anyhow. Idea now (just here for comment): Announce via GRASP Zone-ID (e.g. from per-zone edge-node/registrar) into a zone of the ACP so all nodes supporting the scheme can automatically self-allocate the Zone-ID.

6.11.1.1 (RPL overview), eliminated redundant text.

6.11.1.1.1 New subsection to better structure overview.

6.11.1.1.2 New subsection to better group overview, replaced TTL explanation (just the symptom) with hopefully better reconvergence text (intent of the profile) for the ACP RPL profile.

6.11.1.1.6 Added text to explain simple choice for rank\_factor.

6.11.1.1.3 moved explanation for RPI up into 6.11.1.1.

6.12.5.1 rewrote section for ACP Loopback Interface.

9.4 New informative/informational section for partial or incremental adoption of ACP to help understand why there is the Zone interface sub-scheme, and how to use it.

Unrelated fixes:

Ask to RFC editor to add most important abbreviations to RFC editor abbreviation list.

6.10.2 changed names in ACP addressing scheme table to be less suggestive of use.

Russ Hously review:

2. Fixed definition of "Enrollment", "Trust Anchor", "CA", and "root CA". Changed "Certificate Authority" to "Certification Authority" throughout the document (correct term according to X.509).

6.1 Fixed explanation of mutual ACP trust.

6.1.1 s/X509/X509v3/.

6.1.2 created bulleted lists for explanations and justifications for choices of ACP certificate encoding. No semantic changes, just to make it easier to refer to the points in discussions (rfcdiff seems to have a bug showing text differences due to formatting changes).

6.1.3 Moved content of rule #1 into previous rule #2 because certification chain validation does imply validation of lifetime. numbers of all rules reduced by 1, changed hopefully all references to the rule numbers in the document.

Rule #3, Hopefully fixed linguistic problem self-contradiction of MUST by lower casing MUST in the explanation part and rewriting the condition when this is not applicable.

6.1.4 Replaced redundant term "Trust Point" (TP) with Trust Anchor (TA). Replaced throughout document Trust Anchor with abbreviation TA.

Enhanced several sentences/rewrote paragraphs to make explanations clearer.

6.6 Added explanation how ACP nodes must throttle their attempts for connection making purely on the result of their own connection attempts, not based on those connections where they are responder.

15.8. draft-ietf-anima-autonomic-control-plane-24

Leftover from -23 review by Eric Vyncke:

Swapping sections 9 and 10, section 9 was meant to be at end of document and summarize. Its not meant to be misinterpreted as introducing any new information. This did happen because section 10 was added after section 9.

#### 15.9. draft-ietf-anima-autonomic-control-plane-23

Note: big rfcdiff of TOC is an rfcdiff bug, changes really minimal.

Review of IPsec security with Mcr and ipsec mailing list.

6.7.1 - new section: Moved general considerations for secure channel protocols here, refined them.

6.7.2 - new section: Moved common requirements for secure channel protocols here, refined them.

6.7.3.1.1. - improved requirements text related to RFC8221, better explanations re. HW acceleration issues.

6.7.3.1.2. - improved requirements text related to RFC8247, (some requirements still discussed to be redundant, will be finalized in next weeks.

Eric Vyncke review of -21:

Only noting most important changes, long list of smaller text/readability enhancements.

2. - New explanation of "normative", "informational" section title tags. alphabetic reordering of terms, refined definitions for CA, CRL. root CA.

6.1.1. - explanation when IDevID parameters may be copied into LDevID.

6.1.2. - Fixed hex digits in ACP domain information to lower case.

6.1.3.1. - New section on Realtime clock and Time Validation.

6.3 - Added explanation that DTLS means  $\geq$  version 1.2 (not only 1.2).

6.7 - New text in this main section explaining relationship of ACP secure channels and ACP virtual interfaces - with forward references to virtual interface section.

6.8.2 - reordered text and picture, no text change.

6.10.7.2 - describe first how Registrar-ID can be allocated for all type of registrars, then refined text for how to potentially use MAC addresses on physical registrars.

6.11.1.1 - Added text how this profile does not use Data-Plane artefacts (RPI) because hardware forwarding. This was previously hidden only later in the text.

6.11.1.13. - Rewrote RPL Data-Plane artefact text. Provide decoder ring for abbreviations and all relevant RFCs.

6.12.5.2. - Added more explicit text that secure channels are mapped into virtual interfaces, moved different type of interfaces used by ACP into separate subsections to be able to refer to them.

7.2 - Rewrote/refined text for ACP on L2, prior text was confusing and did not well explain why ACP for L2/L3 switches can be implemented without any L2 (HW) changes. Also missing explanation of only running GRASP untagged when VLANs are used.

8.1.1 - Added requirement for ACP Edge nodes to allow configurable filtering of IPv6 RPI headers.

11. - (security section). Moved explanation of address stealing from 7.2 to here.

15.10. draft-ietf-anima-autonomic-control-plane-22

Ben Kaduk review of -21:

RFC822 encoding of ACP domain information:

6.1.2 rewrote text for explaining / justifying use of rfc822name as identifier for node CP in certificate (was discussed in thread, but badly written in prior versions).

6.1.2 Changed EBNF syntax to use "+" after rfcSELF because that is the known primary name to extensions separator in many email systems ( "." was wrong in prior versions).

6.1.2 Rewrote/improved explanations for use of rfc822name field to explain better why it is PKIX compliant and the right thing to do.

Crypto parameters for IPsec:

6.1 - Added explanation of why manual keying for ACP is not feasible for ACP. Surprisingly, that text did not exist. Referred to by IPsec text (6.7.1), but here is the right place to describe the reasoning.

6.1.2 - Small textual refinement referring to requirements to authenticate peers (for the special cases of empty or '0' ACP address in ACP domain information field).

6.3 - To better justify Bens proposed change of secure channel protocol being IPsec vs. GRASP objective being IKEv2, better explained how protocol indicated in GRASP objective-value is name of protocol used to negotiate secure channel, use example of IKEv2 to negotiate IPsec.

6.7.1 - refinemenet similar to 6.3.

- moved new paragraph from Bens pull request up from 6.7.1.1 to 6.7.1 as it equally applies to GRE encapped IPsec (looks nicer one level up).

- created subsections 6.7.1.1 (IPsec/ESP) / 6.7.1.2 (IKEv2) to clearer distinguish between these two requirements blocks.

- Refined the text in these two sections to hopefully be a good answer to Valery's concern of not randomly mocking with existing requirements docs (rfc8247 / rfc8221).

6.7.1.1.1 - IPsec/ESP requirements section:

- MUST support rfc8221 mandatory EXCEPT for the superceeding requirements in this section. Previously, this was not quite clear from the text.

- Hopefully persuasive explanations about the requirements levels for ENCR\_AES\_GCM\_16, ENCR\_AES\_CBC, ENCR\_AES\_CCM\_8 and ENCR\_CHACHA20\_POLY1305: Restructured text for why not ENCR\_AES\_CBC (was in prior version, just not well structured), added new explanations for ENCR\_AES\_CCM\_8 and ENCR\_CHACHA20\_POLY130.

- In simple terms, requirements for ENCR\_AES\_CBC, ENCR\_AES\_CCM\_8, ENCR\_CHACHA20 are SHOULD when they are implementable with equal or faster performance than ENCR\_AES\_GCM\_16.

- Removed text about "additional rfc8221" requirements MAY be used. Now the logic is that all other requirements apply. Hopefully we have written enough so that we prohibited downgrades.

## 6.7.1.1.2 - RFC8247 requirements:

- Added mandate to support rfc8247, added explanation that there is no "stripping down" requirement, just additional stronger requirements to mandate correct use of ACP certificates during authentication.

- refined text on identifying ACP by IPv6 address to be clearer: Identifying in the context of IKEv2 and cases for '0' in ACP domain information.

- removed last two paragraphs about relationship to rfc8247, as this is now written in first paragraph of the section.

End of Ben Kaduk review related fixes.

Other:

Forgot to update example of ACP domain information to use capitalized hex-digits as required by HEXDIG used.

Added reference to RFC8316 (AN use-cases) to beginning of section 3 (ACP use cases).

Small Enhanced IPsec parameters description / requirements fixes (from Michael Richardson).

## 16. Normative References

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

Pritikin, M., Richardson, M., Eckert, T., Behringer, M., and K. Watsen, "Bootstrapping Remote Secure Key Infrastructures (BRSKI)", Work in Progress, Internet-Draft, draft-ietf-anima-bootstrapping-keyinfra-43, 7 August 2020, <<http://www.ietf.org/internet-drafts/draft-ietf-anima-bootstrapping-keyinfra-43.txt>>.

[I-D.ietf-anima-grasp]

Bormann, C., Carpenter, B., and B. Liu, "A Generic Autonomic Signaling Protocol (GRASP)", Work in Progress, Internet-Draft, draft-ietf-anima-grasp-15, 13 July 2017, <<http://www.ietf.org/internet-drafts/draft-ietf-anima-grasp-15.txt>>.

[IKEV2IANA]

IANA, "Internet Key Exchange Version 2 (IKEv2) Parameters", <<https://www.iana.org/assignments/ikev2-parameters/ikev2-parameters.xhtml>>.

- [RFC1034] Mockapetris, P.V., "Domain names - concepts and facilities", STD 13, RFC 1034, DOI 10.17487/RFC1034, November 1987, <<https://www.rfc-editor.org/info/rfc1034>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC3810] Vida, R., Ed. and L. Costa, Ed., "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, DOI 10.17487/RFC3810, June 2004, <<https://www.rfc-editor.org/info/rfc3810>>.
- [RFC4191] Draves, R. and D. Thaler, "Default Router Preferences and More-Specific Routes", RFC 4191, DOI 10.17487/RFC4191, November 2005, <<https://www.rfc-editor.org/info/rfc4191>>.
- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, DOI 10.17487/RFC4193, October 2005, <<https://www.rfc-editor.org/info/rfc4193>>.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <<https://www.rfc-editor.org/info/rfc4291>>.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, DOI 10.17487/RFC4301, December 2005, <<https://www.rfc-editor.org/info/rfc4301>>.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <<https://www.rfc-editor.org/info/rfc4861>>.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, DOI 10.17487/RFC4862, September 2007, <<https://www.rfc-editor.org/info/rfc4862>>.
- [RFC5234] Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, DOI 10.17487/RFC5234, January 2008, <<https://www.rfc-editor.org/info/rfc5234>>.

- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, DOI 10.17487/RFC5246, August 2008, <<https://www.rfc-editor.org/info/rfc5246>>.
- [RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, DOI 10.17487/RFC5280, May 2008, <<https://www.rfc-editor.org/info/rfc5280>>.
- [RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", RFC 6347, DOI 10.17487/RFC6347, January 2012, <<https://www.rfc-editor.org/info/rfc6347>>.
- [RFC6550] Winter, T., Ed., Thubert, P., Ed., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 6550, DOI 10.17487/RFC6550, March 2012, <<https://www.rfc-editor.org/info/rfc6550>>.
- [RFC6551] Vasseur, JP., Ed., Kim, M., Ed., Pister, K., Dejean, N., and D. Barthel, "Routing Metrics Used for Path Calculation in Low-Power and Lossy Networks", RFC 6551, DOI 10.17487/RFC6551, March 2012, <<https://www.rfc-editor.org/info/rfc6551>>.
- [RFC6552] Thubert, P., Ed., "Objective Function Zero for the Routing Protocol for Low-Power and Lossy Networks (RPL)", RFC 6552, DOI 10.17487/RFC6552, March 2012, <<https://www.rfc-editor.org/info/rfc6552>>.
- [RFC6553] Hui, J. and JP. Vasseur, "The Routing Protocol for Low-Power and Lossy Networks (RPL) Option for Carrying RPL Information in Data-Plane Datagrams", RFC 6553, DOI 10.17487/RFC6553, March 2012, <<https://www.rfc-editor.org/info/rfc6553>>.
- [RFC7030] Pritikin, M., Ed., Yee, P., Ed., and D. Harkins, Ed., "Enrollment over Secure Transport", RFC 7030, DOI 10.17487/RFC7030, October 2013, <<https://www.rfc-editor.org/info/rfc7030>>.
- [RFC7296] Kaufman, C., Hoffman, P., Nir, Y., Eronen, P., and T. Kivinen, "Internet Key Exchange Protocol Version 2 (IKEv2)", STD 79, RFC 7296, DOI 10.17487/RFC7296, October 2014, <<https://www.rfc-editor.org/info/rfc7296>>.

- [RFC7525] Sheffer, Y., Holz, R., and P. Saint-Andre, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", BCP 195, RFC 7525, DOI 10.17487/RFC7525, May 2015, <<https://www.rfc-editor.org/info/rfc7525>>.
- [RFC7676] Pignataro, C., Bonica, R., and S. Krishnan, "IPv6 Support for Generic Routing Encapsulation (GRE)", RFC 7676, DOI 10.17487/RFC7676, October 2015, <<https://www.rfc-editor.org/info/rfc7676>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8221] Wouters, P., Migault, D., Mattsson, J., Nir, Y., and T. Kivinen, "Cryptographic Algorithm Implementation Requirements and Usage Guidance for Encapsulating Security Payload (ESP) and Authentication Header (AH)", RFC 8221, DOI 10.17487/RFC8221, October 2017, <<https://www.rfc-editor.org/info/rfc8221>>.
- [RFC8247] Nir, Y., Kivinen, T., Wouters, P., and D. Migault, "Algorithm Implementation Requirements and Usage Guidance for the Internet Key Exchange Protocol Version 2 (IKEv2)", RFC 8247, DOI 10.17487/RFC8247, September 2017, <<https://www.rfc-editor.org/info/rfc8247>>.
- [RFC8422] Nir, Y., Josefsson, S., and M. Pegourie-Gonnard, "Elliptic Curve Cryptography (ECC) Cipher Suites for Transport Layer Security (TLS) Versions 1.2 and Earlier", RFC 8422, DOI 10.17487/RFC8422, August 2018, <<https://www.rfc-editor.org/info/rfc8422>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.
- [RFC8610] Birkholz, H., Vigano, C., and C. Bormann, "Concise Data Definition Language (CDDL): A Notational Convention to Express Concise Binary Object Representation (CBOR) and JSON Data Structures", RFC 8610, DOI 10.17487/RFC8610, June 2019, <<https://www.rfc-editor.org/info/rfc8610>>.

## 17. Informative References

- [ACPDRAFT] Eckert, T., Behringer, M., and S. Bjarnason, "An Autonomic Control Plane (ACP)", Work in Progress, Internet-Draft, draft-ietf-anima-autonomic-control-plane-30, <<https://tools.ietf.org/html/draft-ietf-anima-autonomic-control-plane-30.pdf>>. [RFC-Editor: Please remove this complete reference from the RFC] Refer to the IETF working group draft for the few sections removed from this document for various reasons. They capture the state of discussion about unresolved issues that may need to be revisited in future work.
- [AR8021] Group, W. -. H. L. L. P. W., "IEEE Standard for Local and metropolitan area networks - Secure Device Identity", December 2009, <<http://standards.ieee.org/findstds/standard/802.1AR-2009.html>>.
- [CABFORUM] CA/Browser Forum, "Certificate Contents for Baseline SSL", November 2019, <<https://cabforum.org/baseline-requirements-certificate-contents/>>.
- [FCC] FCC, "FCC STAFF REPORT ON NATIONWIDE T-MOBILE NETWORK OUTAGE ON JUNE 15, 2020 (PS Docket No. 20-183)", 2020, <<https://docs.fcc.gov/public/attachments/DOC-367699A1.docx>>. The FCC's Public Safety and Homeland Security Bureau issues a report on a nationwide T-Mobile outage that occurred on June 15, 2020. Action by: Public Safety and Homeland Security Bureau.
- [I-D.eckert-anima-noc-autoconfig]  
Eckert, T., "Autoconfiguration of NOC services in ACP networks via GRASP", Work in Progress, Internet-Draft, draft-eckert-anima-noc-autoconfig-00, 2 July 2018, <<http://www.ietf.org/internet-drafts/draft-eckert-anima-noc-autoconfig-00.txt>>.
- [I-D.ietf-acme-star]  
Sheffer, Y., Lopez, D., Dios, O., Pastor, A., and T. Fossati, "Support for Short-Term, Automatically-Renewed (STAR) Certificates in Automated Certificate Management Environment (ACME)", Work in Progress, Internet-Draft, draft-ietf-acme-star-11, 24 October 2019, <<http://www.ietf.org/internet-drafts/draft-ietf-acme-star-11.txt>>.
- [I-D.ietf-anima-prefix-management]  
Jiang, S., Du, Z., Carpenter, B., and Q. Sun, "Autonomic IPv6 Edge Prefix Management in Large-scale Networks", Work in Progress, Internet-Draft, draft-ietf-anima-prefix-

management-07, 18 December 2017, <<http://www.ietf.org/internet-drafts/draft-ietf-anima-prefix-management-07.txt>>.

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

Behringer, M., Carpenter, B., Eckert, T., Ciavaglia, L., and J. Nobre, "A Reference Model for Autonomic Networking", Work in Progress, Internet-Draft, draft-ietf-anima-reference-model-10, 22 November 2018, <<http://www.ietf.org/internet-drafts/draft-ietf-anima-reference-model-10.txt>>.

[I-D.ietf-roll-applicability-template]

Richardson, M., "ROLL Applicability Statement Template", Work in Progress, Internet-Draft, draft-ietf-roll-applicability-template-09, 3 May 2016, <<http://www.ietf.org/internet-drafts/draft-ietf-roll-applicability-template-09.txt>>.

[I-D.ietf-tls-dtls13]

Rescorla, E., Tschofenig, H., and N. Modadugu, "The Datagram Transport Layer Security (DTLS) Protocol Version 1.3", Work in Progress, Internet-Draft, draft-ietf-tls-dtls13-38, 29 May 2020, <<http://www.ietf.org/internet-drafts/draft-ietf-tls-dtls13-38.txt>>.

[IEEE-1588-2008]

IEEE, "IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems", December 2008, <<http://standards.ieee.org/findstds/standard/1588-2008.html>>.

[IEEE-802.1X]

Group, W. -. H. L. L. P. W., "IEEE Standard for Local and Metropolitan Area Networks: Port-Based Network Access Control", February 2010, <<http://standards.ieee.org/findstds/standard/802.1X-2010.html>>.

[LLDP]

Group, W. -. H. L. L. P. W., "IEEE Standard for Local and Metropolitan Area Networks: Station and Media Access Control Connectivity Discovery", June 2016, <<https://standards.ieee.org/findstds/standard/802.1AB-2016.html>>.

- [MACSEC] Group, W. - H. L. L. P. W., "IEEE Standard for Local and Metropolitan Area Networks: Media Access Control (MAC) Security", June 2006, <<https://standards.ieee.org/findstds/standard/802.1AE-2006.html>>.
- [RFC1112] Deering, S.E., "Host extensions for IP multicasting", STD 5, RFC 1112, DOI 10.17487/RFC1112, August 1989, <<https://www.rfc-editor.org/info/rfc1112>>.
- [RFC1492] Finseth, C., "An Access Control Protocol, Sometimes Called TACACS", RFC 1492, DOI 10.17487/RFC1492, July 1993, <<https://www.rfc-editor.org/info/rfc1492>>.
- [RFC1654] Rekhter, Y., Ed. and T. Li, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 1654, DOI 10.17487/RFC1654, July 1994, <<https://www.rfc-editor.org/info/rfc1654>>.
- [RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G. J., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, DOI 10.17487/RFC1918, February 1996, <<https://www.rfc-editor.org/info/rfc1918>>.
- [RFC2315] Kaliski, B., "PKCS #7: Cryptographic Message Syntax Version 1.5", RFC 2315, DOI 10.17487/RFC2315, March 1998, <<https://www.rfc-editor.org/info/rfc2315>>.
- [RFC2409] Harkins, D. and D. Carrel, "The Internet Key Exchange (IKE)", RFC 2409, DOI 10.17487/RFC2409, November 1998, <<https://www.rfc-editor.org/info/rfc2409>>.
- [RFC2865] Rigney, C., Willens, S., Rubens, A., and W. Simpson, "Remote Authentication Dial In User Service (RADIUS)", RFC 2865, DOI 10.17487/RFC2865, June 2000, <<https://www.rfc-editor.org/info/rfc2865>>.
- [RFC3164] Lonvick, C., "The BSD Syslog Protocol", RFC 3164, DOI 10.17487/RFC3164, August 2001, <<https://www.rfc-editor.org/info/rfc3164>>.
- [RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, DOI 10.17487/RFC3315, July 2003, <<https://www.rfc-editor.org/info/rfc3315>>.

- [RFC3411] Harrington, D., Presuhn, R., and B. Wijnen, "An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks", STD 62, RFC 3411, DOI 10.17487/RFC3411, December 2002, <<https://www.rfc-editor.org/info/rfc3411>>.
- [RFC3596] Thomson, S., Huitema, C., Ksinant, V., and M. Souissi, "DNS Extensions to Support IP Version 6", STD 88, RFC 3596, DOI 10.17487/RFC3596, October 2003, <<https://www.rfc-editor.org/info/rfc3596>>.
- [RFC3920] Saint-Andre, P., Ed., "Extensible Messaging and Presence Protocol (XMPP): Core", RFC 3920, DOI 10.17487/RFC3920, October 2004, <<https://www.rfc-editor.org/info/rfc3920>>.
- [RFC3954] Claise, B., Ed., "Cisco Systems NetFlow Services Export Version 9", RFC 3954, DOI 10.17487/RFC3954, October 2004, <<https://www.rfc-editor.org/info/rfc3954>>.
- [RFC4007] Deering, S., Haberman, B., Jinmei, T., Nordmark, E., and B. Zill, "IPv6 Scoped Address Architecture", RFC 4007, DOI 10.17487/RFC4007, March 2005, <<https://www.rfc-editor.org/info/rfc4007>>.
- [RFC4210] Adams, C., Farrell, S., Kause, T., and T. Mononen, "Internet X.509 Public Key Infrastructure Certificate Management Protocol (CMP)", RFC 4210, DOI 10.17487/RFC4210, September 2005, <<https://www.rfc-editor.org/info/rfc4210>>.
- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, DOI 10.17487/RFC4364, February 2006, <<https://www.rfc-editor.org/info/rfc4364>>.
- [RFC4429] Moore, N., "Optimistic Duplicate Address Detection (DAD) for IPv6", RFC 4429, DOI 10.17487/RFC4429, April 2006, <<https://www.rfc-editor.org/info/rfc4429>>.
- [RFC4541] Christensen, M., Kimball, K., and F. Solensky, "Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches", RFC 4541, DOI 10.17487/RFC4541, May 2006, <<https://www.rfc-editor.org/info/rfc4541>>.

- [RFC4604] Holbrook, H., Cain, B., and B. Haberman, "Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast", RFC 4604, DOI 10.17487/RFC4604, August 2006, <<https://www.rfc-editor.org/info/rfc4604>>.
- [RFC4607] Holbrook, H. and B. Cain, "Source-Specific Multicast for IP", RFC 4607, DOI 10.17487/RFC4607, August 2006, <<https://www.rfc-editor.org/info/rfc4607>>.
- [RFC4610] Farinacci, D. and Y. Cai, "Anycast-RP Using Protocol Independent Multicast (PIM)", RFC 4610, DOI 10.17487/RFC4610, August 2006, <<https://www.rfc-editor.org/info/rfc4610>>.
- [RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 4941, DOI 10.17487/RFC4941, September 2007, <<https://www.rfc-editor.org/info/rfc4941>>.
- [RFC4985] Santesson, S., "Internet X.509 Public Key Infrastructure Subject Alternative Name for Expression of Service Name", RFC 4985, DOI 10.17487/RFC4985, August 2007, <<https://www.rfc-editor.org/info/rfc4985>>.
- [RFC5790] Liu, H., Cao, W., and H. Asaeda, "Lightweight Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Version 2 (MLDv2) Protocols", RFC 5790, DOI 10.17487/RFC5790, February 2010, <<https://www.rfc-editor.org/info/rfc5790>>.
- [RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, DOI 10.17487/RFC5880, June 2010, <<https://www.rfc-editor.org/info/rfc5880>>.
- [RFC5905] Mills, D., Martin, J., Ed., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, DOI 10.17487/RFC5905, June 2010, <<https://www.rfc-editor.org/info/rfc5905>>.
- [RFC5912] Hoffman, P. and J. Schaad, "New ASN.1 Modules for the Public Key Infrastructure Using X.509 (PKIX)", RFC 5912, DOI 10.17487/RFC5912, June 2010, <<https://www.rfc-editor.org/info/rfc5912>>.

- [RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", RFC 6241, DOI 10.17487/RFC6241, June 2011, <<https://www.rfc-editor.org/info/rfc6241>>.
- [RFC6335] Cotton, M., Eggert, L., Touch, J., Westerlund, M., and S. Cheshire, "Internet Assigned Numbers Authority (IANA) Procedures for the Management of the Service Name and Transport Protocol Port Number Registry", BCP 165, RFC 6335, DOI 10.17487/RFC6335, August 2011, <<https://www.rfc-editor.org/info/rfc6335>>.
- [RFC6402] Schaad, J., "Certificate Management over CMS (CMC) Updates", RFC 6402, DOI 10.17487/RFC6402, November 2011, <<https://www.rfc-editor.org/info/rfc6402>>.
- [RFC6407] Weis, B., Rowles, S., and T. Hardjono, "The Group Domain of Interpretation", RFC 6407, DOI 10.17487/RFC6407, October 2011, <<https://www.rfc-editor.org/info/rfc6407>>.
- [RFC6554] Hui, J., Vasseur, JP., Culler, D., and V. Manral, "An IPv6 Routing Header for Source Routes with the Routing Protocol for Low-Power and Lossy Networks (RPL)", RFC 6554, DOI 10.17487/RFC6554, March 2012, <<https://www.rfc-editor.org/info/rfc6554>>.
- [RFC6724] Thaler, D., Ed., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, DOI 10.17487/RFC6724, September 2012, <<https://www.rfc-editor.org/info/rfc6724>>.
- [RFC6733] Fajardo, V., Ed., Arkko, J., Loughney, J., and G. Zorn, Ed., "Diameter Base Protocol", RFC 6733, DOI 10.17487/RFC6733, October 2012, <<https://www.rfc-editor.org/info/rfc6733>>.
- [RFC6762] Cheshire, S. and M. Krochmal, "Multicast DNS", RFC 6762, DOI 10.17487/RFC6762, February 2013, <<https://www.rfc-editor.org/info/rfc6762>>.
- [RFC6763] Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", RFC 6763, DOI 10.17487/RFC6763, February 2013, <<https://www.rfc-editor.org/info/rfc6763>>.
- [RFC6824] Ford, A., Raiciu, C., Handley, M., and O. Bonaventure, "TCP Extensions for Multipath Operation with Multiple Addresses", RFC 6824, DOI 10.17487/RFC6824, January 2013, <<https://www.rfc-editor.org/info/rfc6824>>.

- [RFC6830] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "The Locator/ID Separation Protocol (LISP)", RFC 6830, DOI 10.17487/RFC6830, January 2013, <<https://www.rfc-editor.org/info/rfc6830>>.
- [RFC7011] Claise, B., Ed., Trammell, B., Ed., and P. Aitken, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information", STD 77, RFC 7011, DOI 10.17487/RFC7011, September 2013, <<https://www.rfc-editor.org/info/rfc7011>>.
- [RFC7404] Behringer, M. and E. Vyncke, "Using Only Link-Local Addressing inside an IPv6 Network", RFC 7404, DOI 10.17487/RFC7404, November 2014, <<https://www.rfc-editor.org/info/rfc7404>>.
- [RFC7426] Haleplidis, E., Ed., Pentikousis, K., Ed., Denazis, S., Hadi Salim, J., Meyer, D., and O. Koufopavlou, "Software-Defined Networking (SDN): Layers and Architecture Terminology", RFC 7426, DOI 10.17487/RFC7426, January 2015, <<https://www.rfc-editor.org/info/rfc7426>>.
- [RFC7435] Dukhovni, V., "Opportunistic Security: Some Protection Most of the Time", RFC 7435, DOI 10.17487/RFC7435, December 2014, <<https://www.rfc-editor.org/info/rfc7435>>.
- [RFC7575] Behringer, M., Pritikin, M., Bjarnason, S., Clemm, A., Carpenter, B., Jiang, S., and L. Ciavaglia, "Autonomic Networking: Definitions and Design Goals", RFC 7575, DOI 10.17487/RFC7575, June 2015, <<https://www.rfc-editor.org/info/rfc7575>>.
- [RFC7576] Jiang, S., Carpenter, B., and M. Behringer, "General Gap Analysis for Autonomic Networking", RFC 7576, DOI 10.17487/RFC7576, June 2015, <<https://www.rfc-editor.org/info/rfc7576>>.
- [RFC7585] Winter, S. and M. McCauley, "Dynamic Peer Discovery for RADIUS/TLS and RADIUS/DTLS Based on the Network Access Identifier (NAI)", RFC 7585, DOI 10.17487/RFC7585, October 2015, <<https://www.rfc-editor.org/info/rfc7585>>.
- [RFC7721] Cooper, A., Gont, F., and D. Thaler, "Security and Privacy Considerations for IPv6 Address Generation Mechanisms", RFC 7721, DOI 10.17487/RFC7721, March 2016, <<https://www.rfc-editor.org/info/rfc7721>>.

- [RFC7761] Fenner, B., Handley, M., Holbrook, H., Kouvelas, I., Parekh, R., Zhang, Z., and L. Zheng, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", STD 83, RFC 7761, DOI 10.17487/RFC7761, March 2016, <<https://www.rfc-editor.org/info/rfc7761>>.
- [RFC7950] Bjorklund, M., Ed., "The YANG 1.1 Data Modeling Language", RFC 7950, DOI 10.17487/RFC7950, August 2016, <<https://www.rfc-editor.org/info/rfc7950>>.
- [RFC8028] Baker, F. and B. Carpenter, "First-Hop Router Selection by Hosts in a Multi-Prefix Network", RFC 8028, DOI 10.17487/RFC8028, November 2016, <<https://www.rfc-editor.org/info/rfc8028>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [RFC8316] Nobre, J., Granville, L., Clemm, A., and A. Gonzalez Prieto, "Autonomic Networking Use Case for Distributed Detection of Service Level Agreement (SLA) Violations", RFC 8316, DOI 10.17487/RFC8316, February 2018, <<https://www.rfc-editor.org/info/rfc8316>>.
- [RFC8366] Watsen, K., Richardson, M., Pritikin, M., and T. Eckert, "A Voucher Artifact for Bootstrapping Protocols", RFC 8366, DOI 10.17487/RFC8366, May 2018, <<https://www.rfc-editor.org/info/rfc8366>>.
- [RFC8368] Eckert, T., Ed. and M. Behringer, "Using an Autonomic Control Plane for Stable Connectivity of Network Operations, Administration, and Maintenance (OAM)", RFC 8368, DOI 10.17487/RFC8368, May 2018, <<https://www.rfc-editor.org/info/rfc8368>>.
- [RFC8398] Melnikov, A., Ed. and W. Chuang, Ed., "Internationalized Email Addresses in X.509 Certificates", RFC 8398, DOI 10.17487/RFC8398, May 2018, <<https://www.rfc-editor.org/info/rfc8398>>.
- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.

- [RFC8572] Watsen, K., Farrer, I., and M. Abrahamsson, "Secure Zero Touch Provisioning (SZTP)", RFC 8572, DOI 10.17487/RFC8572, April 2019, <<https://www.rfc-editor.org/info/rfc8572>>.
- [X.509] International Telecommunication Union, "Information technology - Open Systems Interconnection - The Directory: Public-key and attribute certificate frameworks", ITU-T Recommendation X.509, ISO/IEC 9594-8, October 2016, <<https://www.itu.int/rec/T-REC-X.509>>.
- [X.520] International Telecommunication Union, "Information technology - Open Systems Interconnection - The Directory: Selected attribute types", ITU-T Recommendation X.520, ISO/IEC 9594-6, October 2016, <<https://www.itu.int/rec/T-REC-X.520>>.

#### Appendix A. Background and Futures (Informative)

The following sections discuss additional background information about aspects of the normative parts of this document or associated mechanisms such as BRSKI (such as why specific choices were made by the ACP) and they provide discussion about possible future variations of the ACP.

##### A.1. ACP Address Space Schemes

This document defines the Zone, Vlong and Manual sub address schemes primarily to support address prefix assignment via distributed, potentially uncoordinated ACP registrars as defined in Section 6.11.7. This costs 48/46-bit identifier so that these ACP registrar can assign non-conflicting address prefixes. This design does not leave enough bits to simultaneously support a large number of nodes (Node-ID) plus a large prefix of local addresses for every node plus a large enough set of bits to identify a routing Zone. In result, Zone, Vlong 8/16 attempt to support all features, but via separate prefixes.

In networks that always expect to rely on a centralized PMS as described above (Section 9.2.5), the 48/46-bits for the Registrar-ID could be saved. Such variations of the ACP addressing mechanisms could be introduced through future work in different ways. If a new otherName was introduced, incompatible ACP variations could be created where every design aspect of the ACP could be changed. Including all addressing choices. If instead a new addressing sub-type would be defined, it could be a backward compatible extension of this ACP specification. Information such as the size of a zone-prefix and the length of the prefix assigned to the ACP node itself could be encoded via the extension field of the acp-node-name.

Note that an explicitly defined "Manual" addressing sub-scheme is always beneficial to provide an easy way for ACP nodes to prohibit incorrect manual configuration of any non-"Manual" ACP address spaces and therefore ensure that "Manual" operations will never impact correct routing for any non-"Manual" ACP addresses assigned via ACP certificates.

#### A.2. BRSKI Bootstrap (ANI)

BRSKI describes how nodes with an IDevID certificate can securely and zero-touch enroll with an LDevID certificate to support the ACP. BRSKI also leverages the ACP to enable zero-touch bootstrap of new nodes across networks without any configuration requirements across the transit nodes (e.g., no DHCP/DNS forwarding/server setup). This includes otherwise not configured networks as described in Section 3.2. Therefore, BRSKI in conjunction with ACP provides for a secure and zero-touch management solution for complete networks. Nodes supporting such an infrastructure (BRSKI and ACP) are called ANI nodes (Autonomic Networking Infrastructure), see [I-D.ietf-anima-reference-model]. Nodes that do not support an IDevID certificate but only an (insecure) vendor specific Unique Device Identifier (UDI) or nodes whose manufacturer does not support a MASA could use some future security reduced version of BRSKI.

When BRSKI is used to provision a domain certificate (which is called enrollment), the BRSKI registrar (acting as an enhanced EST server) must include the otherName / AcpNodeName encoded ACP address and domain name to the enrolling node (called pledge) via its response to the pledges EST CSR Attribute request that is mandatory in BRSKI.

The Certification Authority in an ACP network must not change the otherName / AcpNodeName in the certificate. The ACP nodes can therefore find their ACP address and domain using this field in the domain certificate, both for themselves, as well as for other nodes.

The use of BRSKI in conjunction with the ACP can also help to further simplify maintenance and renewal of domain certificates. Instead of relying on CRL, the lifetime of certificates can be made extremely small, for example in the order of hours. When a node fails to connect to the ACP within its certificate lifetime, it cannot connect to the ACP to renew its certificate across it (using just EST), but it can still renew its certificate as an "enrolled/expired pledge" via the BRSKI bootstrap proxy. This requires only that the BRSKI registrar honors expired domain certificates and that the pledge attempts to perform TLS authentication for BRSKI bootstrap using its expired domain certificate before falling back to attempting to use its IDevID certificate for BRSKI. This mechanism could also render CRLs unnecessary because the BRSKI registrar in conjunction with the CA would not renew revoked certificates - only a "Do-not-renew" list would be necessary on BRSKI registrars/CA.

In the absence of BRSKI or less secure variants thereof, provisioning of certificates may involve one or more touches or non-standardized automation. Node vendors usually support provisioning of certificates into nodes via PKCS#7 (see [RFC2315]) and may support this provisioning through vendor specific models via NETCONF ([RFC6241]). If such nodes also support NETCONF Zero-Touch ([RFC8572]) then this can be combined to zero-touch provisioning of domain certificates into nodes. Unless there are equivalent integration of NETCONF connections across the ACP as there is in BRSKI, this combination would not support zero-touch bootstrap across a not configured network though.

### A.3. ACP Neighbor discovery protocol selection

This section discusses why GRASP DULL was chosen as the discovery protocol for L2 adjacent candidate ACP neighbors. The contenders considered where GRASP, mDNS or LLDP.

#### A.3.1. LLDP

LLDP and Cisco's earlier Cisco Discovery Protocol (CDP) are example of L2 discovery protocols that terminate their messages on L2 ports. If those protocols would be chosen for ACP neighbor discovery, ACP neighbor discovery would therefore also terminate on L2 ports. This would prevent ACP construction over non-ACP capable but LLDP or CDP enabled L2 switches. LLDP has extensions using different MAC addresses and this could have been an option for ACP discovery as well, but the additional required IEEE standardization and definition of a profile for such a modified instance of LLDP seemed to be more work than the benefit of "reusing the existing protocol" LLDP for this very simple purpose.

### A.3.2. mDNS and L2 support

Multicast DNNS (mDNS) [RFC6762] with DNS Service Discovery (DNS-SD) Resource Records (RRs) as defined in [RFC6763] is a key contender as an ACP discovery protocol. because it relies on link-local IP multicast, it does operates at the subnet level, and is also found in L2 switches. The authors of this document are not aware of mDNS implementation that terminate their mDNS messages on L2 ports instead of the subnet level. If mDNS was used as the ACP discovery mechanism on an ACP capable (L3)/L2 switch as outlined in Section 7, then this would be necessary to implement. It is likely that termination of mDNS messages could only be applied to all mDNS messages from such a port, which would then make it necessary to software forward any non-ACP related mDNS messages to maintain prior non-ACP mDNS functionality. Adding support for ACP into such L2 switches with mDNS could therefore create regression problems for prior mDNS functionality on those nodes. With low performance of software forwarding in many L2 switches, this could also make the ACP risky to support on such L2 switches.

### A.3.3. Why DULL GRASP

LLDP was not considered because of the above mentioned issues. mDNS was not selected because of the above L2 mDNS considerations and because of the following additional points:

If mDNS was not already existing in a node, it would be more work to implement than DULL GRASP, and if an existing implementation of mDNS was used, it would likely be more code space than a separate implementation of DULL GRASP or a shared implementation of DULL GRASP and GRASP in the ACP.

### A.4. Choice of routing protocol (RPL)

This section motivates why RPL - "IPv6 Routing Protocol for Low-Power and Lossy Networks ([RFC6550] was chosen as the default (and in this specification only) routing protocol for the ACP. The choice and above explained profile was derived from a pre-standard implementation of ACP that was successfully deployed in operational networks.

Requirements for routing in the ACP are:

- \* Self-management: The ACP must build automatically, without human intervention. Therefore, routing protocol must also work completely automatically. RPL is a simple, self-managing protocol, which does not require zones or areas; it is also self-configuring, since configuration is carried as part of the protocol (see Section 6.7.6 of [RFC6550]).
- \* Scale: The ACP builds over an entire domain, which could be a large enterprise or service provider network. The routing protocol must therefore support domains of 100,000 nodes or more, ideally without the need for zoning or separation into areas. RPL has this scale property. This is based on extensive use of default routing.
- \* Low resource consumption: The ACP supports traditional network infrastructure, thus runs in addition to traditional protocols. The ACP, and specifically the routing protocol must have low resource consumption both in terms of memory and CPU requirements. Specifically, at edge nodes, where memory and CPU are scarce, consumption should be minimal. RPL builds a DODAG, where the main resource consumption is at the root of the DODAG. The closer to the edge of the network, the less state needs to be maintained. This adapts nicely to the typical network design. Also, all changes below a common parent node are kept below that parent node.
- \* Support for unstructured address space: In the Autonomic Networking Infrastructure, node addresses are identifiers, and may not be assigned in a topological way. Also, nodes may move topologically, without changing their address. Therefore, the routing protocol must support completely unstructured address space. RPL is specifically made for mobile ad-hoc networks, with no assumptions on topologically aligned addressing.
- \* Modularity: To keep the initial implementation small, yet allow later for more complex methods, it is highly desirable that the routing protocol has a simple base functionality, but can import new functional modules if needed. RPL has this property with the concept of "objective function", which is a plugin to modify routing behavior.
- \* Extensibility: Since the Autonomic Networking Infrastructure is a new concept, it is likely that changes in the way of operation will happen over time. RPL allows for new objective functions to be introduced later, which allow changes to the way the routing protocol creates the DAGs.
- \* Multi-topology support: It may become necessary in the future to support more than one DODAG for different purposes, using different objective functions. RPL allow for the creation of several parallel DODAGs, should this be required. This could be used to create different topologies to reach different roots.

- \* No need for path optimization: RPL does not necessarily compute the optimal path between any two nodes. However, the ACP does not require this today, since it carries mainly non-delay-sensitive feedback loops. It is possible that different optimization schemes become necessary in the future, but RPL can be expanded (see point "Extensibility" above).

#### A.5. ACP Information Distribution and multicast

IP multicast is not used by the ACP because the ANI (Autonomic Networking Infrastructure) itself does not require IP multicast but only service announcement/discovery. Using IP multicast for that would have made it necessary to develop a zero-touch auto configuring solution for ASM (Any Source Multicast - the original form of IP multicast defined in [RFC1112]), which would be quite complex and difficult to justify. One aspect of complexity where no attempt at a solution has been described in IETF documents is the automatic-selection of routers that should be PIM Sparse Mode (PIM-SM) Rendezvous Points (RPs) (see [RFC7761]). The other aspects of complexity are the implementation of MLD ([RFC4604]), PIM-SM and Anycast-RP (see [RFC4610]). If those implementations already exist in a product, then they would be very likely tied to accelerated forwarding which consumes hardware resources, and that in return is difficult to justify as a cost of performing only service discovery.

Some future ASA may need high performance in-network data replication. That is the case when the use of IP multicast is justified. Such an ASA can then use service discovery from ACP GRASP, and then they do not need ASM but only SSM (Source Specific Multicast, see [RFC4607]) for the IP multicast replication. SSM itself can simply be enabled in the Data-Plane (or even in an update to the ACP) without any other configuration than just enabling it on all nodes and only requires a simpler version of MLD (see [RFC5790]).

LSP (Link State Protocol) based IGP routing protocols typically have a mechanism to flood information, and such a mechanism could be used to flood GRASP objectives by defining them to be information of that IGP. This would be a possible optimization in future variations of the ACP that do use an LSP routing protocol. Note though that such a mechanism would not work easily for GRASP M\_DISCOVERY messages which are intelligently (constrained) flooded not across the whole ACP, but only up to a node where a responder is found. We do expect that many future services in ASA will have only few consuming ASA, and for those cases, M\_DISCOVERY is the more efficient method than flooding across the whole domain.

Because the ACP uses RPL, one desirable future extension is to use RPLs existing notion of DODAG, which are loop-free distribution trees, to make GRASP flooding more efficient both for M\_FLOOD and M\_DISCOVERY. See Section 6.13.5 how this will be specifically beneficial when using NBMA interfaces. This is not currently specified in this document because it is not quite clear yet what exactly the implications are to make GRASP flooding depend on RPL DODAG convergence and how difficult it would be to let GRASP flooding access the DODAG information.

#### A.6. CAs, domains and routing subdomains

There is a wide range of setting up different ACP solution by appropriately using CAs and the domain and rsub elements in the acp-node-name in the domain certificate. We summarize these options here as they have been explained in different parts of the document in before and discuss possible and desirable extensions:

An ACP domain is the set of all ACP nodes that can authenticate each other as belonging to the same ACP network using the ACP domain membership check (Section 6.2.3). GRASP inside the ACP is run across all transitively connected ACP nodes in a domain.

The rsub element in the acp-node-name permits the use of addresses from different ULA prefixes. One use case is to create multiple physical networks that initially may be separated with one ACP domain but different routing subdomains, so that all nodes can mutual trust their ACP certificates (not depending on rsub) and so that they could connect later together into a contiguous ACP network.

One instance of such a use case is an ACP for regions interconnected via a non-ACP enabled core, for example due to the absence of product support for ACP on the core nodes. ACP connect configurations as defined in this document can be used to extend and interconnect those ACP islands to the NOC and merge them into a single ACP when later that product support gap is closed.

Note that RPL scales very well. It is not necessary to use multiple routing subdomains to scale ACP domains in a way that would be required if other routing protocols where used. They exist only as options for the above mentioned reasons.

If different ACP domains are to be created that should not allow to connect to each other by default, these ACP domains simply need to have different domain elements in the acp-node-name. These domain elements can be arbitrary, including subdomains of one another: Domains "example.com" and "research.example.com" are separate domains if both are domain elements in the acp-node-name of certificates.

It is not necessary to have a separate CA for different ACP domains: an operator can use a single CA to sign certificates for multiple ACP domains that are not allowed to connect to each other because the checks for ACP adjacencies includes comparison of the domain part.

If multiple independent networks choose the same domain name but had their own CA, these would not form a single ACP domain because of CA mismatch. Therefore, there is no problem in choosing domain names that are potentially also used by others. Nevertheless it is highly recommended to use domain names that one can have high probability to be unique. It is recommended to use domain names that start with a DNS domain names owned by the assigning organization and unique within it. For example, "acp.example.com" if you own "example.com".

#### A.7. Intent for the ACP

Intent is the architecture component of autonomic networks according to [I-D.ietf-anima-reference-model] that allows operators to issue policies to the network. Its applicability for use is quite flexible and freeform, with potential applications including policies flooded across ACP GRASP and interpreted on every ACP node.

One concern for future definitions of Intent solutions is the problem of circular dependencies when expressing Intent policies about the ACP itself.

For example, Intent could indicate the desire to build an ACP across all domains that have a common parent domain (without relying on the rsub/routing-subdomain solution defined in this document). For example, ACP nodes with domain "example.com", "access.example.com", "core.example.com" and "city.core.example.com" should all establish one single ACP.

If each domain has its own source of Intent, then the Intent would simply have to allow adding the peer domains TA and domain names to the parameters for the ACP domain membership check (Section 6.2.3) so that nodes from those other domains are accepted as ACP peers.

If this Intent was to be originated only from one domain, it could likely not be made to work because the other domains will not build any ACP connection amongst each other, whether they use the same or different CA due to the ACP domain membership check.

If the domains use the same CA one could change the ACP setup to permit for the ACP to be established between two ACP nodes with different `acp-domain-names`, but only for the purpose of disseminating limited information, such as Intent, but not to set up full ACP connectivity, specifically not RPL routing and passing of arbitrary GRASP information. Unless the Intent policies permit this to happen across domain boundaries.

This type of approach where the ACP first allows Intent to operate and only then sets up the rest of ACP connectivity based on Intent policy could also be used to enable Intent policies that would limit functionality across the ACP inside a domain, as long as no policy would disturb the distribution of Intent. For example, to limit reachability across the ACP to certain type of nodes or locations of nodes.

#### A.8. Adopting ACP concepts for other environments

The ACP as specified in this document is very explicit about the choice of options to allow interoperable implementations. The choices made may not be the best for all environments, but the concepts used by the ACP can be used to build derived solutions:

The ACP specifies the use of ULA and deriving its prefix from the domain name so that no address allocation is required to deploy the ACP. The ACP will equally work not using ULA but any other /48 IPv6 prefix. This prefix could simply be a configuration of the ACP registrars (for example when using BRSKI) to enroll the domain certificates - instead of the ACP registrar deriving the /48 ULA prefix from the AN domain name.

Some solutions may already have an auto-addressing scheme, for example derived from existing unique device identifiers (e.g., MAC addresses). In those cases it may not be desirable to assign addresses to devices via the ACP address information field in the way described in this document. The certificate may simply serve to identify the ACP domain, and the address field could be omitted. The only fix required in the remaining way the ACP operate is to define another element in the domain certificate for the two peers to decide who is the Decider and who is the Follower during secure channel building. Note though that future work may leverage the `acp address` to authenticate "ownership" of the address by the device. If the address used by a device is derived from some pre-existing permanent local ID (such as MAC address), then it would be useful to store that address in the certificate using the format of the access address information field or in a similar way.

The ACP is defined as a separate VRF because it intends to support well managed networks with a wide variety of configurations. Therefore, reliable, configuration-indestructible connectivity cannot be achieved from the Data-Plane itself. In solutions where all transit connectivity impacting functions are fully automated (including security), indestructible and resilient, it would be possible to eliminate the need for the ACP to be a separate VRF. Consider the most simple example system in which there is no separate Data-Plane, but the ACP is the Data-Plane. Add BRSKI, and it becomes a fully autonomic network - except that it does not support automatic addressing for user equipment. This gap can then be closed for example by adding a solution derived from [I-D.ietf-anima-prefix-management].

TCP/TLS as the protocols to provide reliability and security to GRASP in the ACP may not be the preferred choice in constrained networks. For example, CoAP/DTLS (Constrained Application Protocol) may be preferred where they are already used, allowing to reduce the additional code space footprint for the ACP on those devices. Hop-by-hop reliability for ACP GRASP messages could be made to support protocols like DTLS by adding the same type of negotiation as defined in this document for ACP secure channel protocol negotiation. End-to-end GRASP connections can be made to select their transport protocol in future extensions of the ACP meant to better support constrained devices by indicating the supported transport protocols (e.g. TLS/DTLS) via GRASP parameters of the GRASP objective through which the transport endpoint is discovered.

The routing protocol RPL used for the ACP does explicitly not optimize for shortest paths and fastest convergence. Variations of the ACP may want to use a different routing protocol or introduce more advanced RPL profiles.

Variations such as what routing protocol to use, or whether to instantiate an ACP in a VRF or (as suggested above) as the actual Data-Plane, can be automatically chosen in implementations built to support multiple options by deriving them from future parameters in the certificate. Parameters in certificates should be limited to those that would not need to be changed more often than certificates would need to be updated anyhow; Or by ensuring that these parameters can be provisioned before the variation of an ACP is activated in a node. Using BRSKI, this could be done for example as additional follow-up signaling directly after the certificate enrollment, still leveraging the BRSKI TLS connection and therefore not introducing any additional connectivity requirements.

Last but not least, secure channel protocols including their encapsulations are easily added to ACP solutions. ACP hop-by-hop network layer secure channels could also be replaced by end-to-end security plus other means for infrastructure protection. Any future network OAM should always use end-to-end security anyhow and can leverage the domain certificates and is therefore not dependent on security to be provided for by ACP secure channels.

#### A.9. Further (future) options

##### A.9.1. Auto-aggregation of routes

Routing in the ACP according to this specification only leverages the standard RPL mechanism of route optimization, e.g. keeping only routes that are not towards the RPL root. This is known to scale to networks with 20,000 or more nodes. There is no auto-aggregation of routes for /48 ULA prefixes (when using rsub in the acp-node-name) and/or Zone-ID based prefixes.

Automatic assignment of Zone-ID and auto-aggregation of routes could be achieved for example by configuring zone-boundaries, announcing via GRASP into the zones the zone parameters (zone-ID and /48 ULA prefix) and auto-aggregating routes on the zone-boundaries. Nodes would assign their Zone-ID and potentially even /48 prefix based on the GRASP announcements.

##### A.9.2. More options for avoiding IPv6 Data-Plane dependencies

As described in Section 6.13.2, the ACP depends on the Data-Plane to establish IPv6 link-local addressing on interfaces. Using a separate MAC address for the ACP allows to fully isolate the ACP from the Data-Plane in a way that is compatible with this specification. It is also an ideal option when using Single-root input/output virtualization (SR-IOV - see [https://en.wikipedia.org/wiki/Single-root\\_input/output\\_virtualization](https://en.wikipedia.org/wiki/Single-root_input/output_virtualization)) in an implementation to isolate the ACP because different SR-IOV interfaces use different MAC addresses.

When additional MAC address(es) are not available, separation of the ACP could be done at different demux points. The same subnet interface could have a separate IPv6 interface for the ACP and Data-Plane and therefore separate link-local addresses for both, where the ACP interface is non-configurable on the Data-Plane. This too would be compatible with this specification and not impact interoperability.

An option that would require additional specification is to use a different Ethertype from 0x86DD (IPv6) to encapsulate IPv6 packets for the ACP. This would be a similar approach as used for IP

authentication packets in [IEEE-802.1X] which use the Extensible Authentication Protocol over Local Area Network (EAPoL) ethertype (0x88A2).

Note that in the case of ANI nodes, all the above considerations equally apply to the encapsulation of BRSKI packets including GRASP used for BRSKI.

A.9.3. ACP APIs and operational models (YANG)

Future work should define YANG ([RFC7950]) data model and/or node internal APIs to monitor and manage the ACP.

Support for the ACP Adjacency Table (Section 6.3) and ACP GRASP need to be included into such model/API.

A.9.4. RPL enhancements

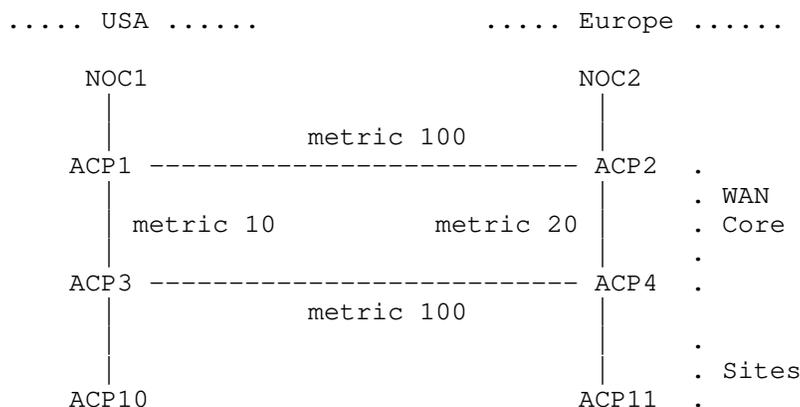


Figure 19: Dual NOC

The profile for RPL specified in this document builds only one spanning-tree path set to a root, typically a registrar in one NOC. In the presence of multiple NOCs, routing toward the non-root NOCs may be suboptimal. Figure 19 shows an extreme example. Assuming that node ACP1 becomes the RPL root, traffic between ACP11 and NOC2 will pass through ACP4-ACP3-ACP1-ACP2 instead of ACP4-ACP2 because the RPL calculated DODAG/routes are shortest paths towards the RPL root.

To overcome these limitations, extensions/modifications to the RPL profile can provide optimality for multiple NOCs. This requires utilizing Data-Plane artifact including IPinIP encap/decap on ACP routers and processing of IPv6 RPI headers. Alternatively, (Src,Dst) routing table entries could be used.

Flooding of ACP GRASP messages can be further constrained and therefore optimized by flooding only via links that are part of the RPL DODAG.

#### A.9.5. Role assignments

ACP connect is an explicit mechanism to "leak" ACP traffic explicitly (for example in a NOC). It is therefore also a possible security gap when it is easy to enable ACP connect on arbitrary compromised ACP nodes.

One simple solution is to define an extension in the ACP certificates ACP information field indicating the permission for ACP connect to be configured on that ACP node. This could similarly be done to decide whether a node is permitted to be a registrar or not.

Tying the permitted "roles" of an ACP node to the ACP certificate provides fairly strong protection against misconfiguration, but is still subject to code modifications.

Another interesting role to assign to certificates is that of a NOC node. This would allow to limit certain type of connections such as OAM TLS connections to only NOC initiator or responders.

#### A.9.6. Autonomic L3 transit

In this specification, the ACP can only establish autonomic connectivity across L2 hops and only explicitly configured options to tunnel across L3. Future work should specify mechanisms to automatically tunnel ACP across L3 networks. A hub&spoke option would allow to tunnel across the Internet to a cloud or central instance of the ACP, a peer-to-peer tunneling mechanism could tunnel ACP islands across an L3VPN infrastructure.

#### A.9.7. Diagnostics

Section 9.1 describes diagnostics options that can be done without changing the external, interoperability affecting characteristics of ACP implementations.

Even better diagnostics of ACP operations is possible with additional signaling extensions, such as:

1. Consider if LLDP should be a recommended functionality for ANI devices to improve diagnostics, and if so, which information elements it should signal (noting that such information is conveyed in an insecure manner). Includes potentially new information elements.
2. In alternative to LLDP, A DULL GRASP diagnostics objective could be defined to carry these information elements.
3. The IDevID certificate of BRSKI pledges should be included in the selected insecure diagnostics option. This may be undesirable when exposure of device information is seen as too much of a security issue (ability to deduce possible attack vectors from device model for example).
4. A richer set of diagnostics information should be made available via the secured ACP channels, using either single-hop GRASP or network wide "topology discovery" mechanisms.

#### A.9.8. Avoiding and dealing with compromised ACP nodes

Compromised ACP nodes pose the biggest risk to the operations of the network. The most common type of compromise is leakage of credentials to manage/configure the device and the application of malicious configuration including the change of access credentials, but not the change of software. Most of today's networking equipment should have secure boot/software infrastructure anyhow, so attacks that introduce malicious software should be a lot harder.

The most important aspect of security design against these type of attacks is to eliminate password based configuration access methods and instead rely on certificate based credentials handed out only to nodes where it is clear that the private keys cannot leak. This limits unexpected propagation of credentials.

If password based credentials to configure devices still need to be supported, they must not be locally configurable, but only be remotely provisioned or verified (through protocols like RADIUS or Diameter), and there must be no local configuration permitting to change these authentication mechanisms, but ideally they should be autoconfiguring across the ACP. See [I-D.eckert-anima-noc-autoconfig].

Without physical access to the compromised device, attackers with access to configuration should not be able to break the ACP connectivity, even when they can break or otherwise manipulate (spoof) the Data-Plane connectivity through configuration. To achieve this, it is necessary to avoid providing configuration options for the ACP, such as enabling/disabling it on interfaces. For example, there could be an ACP configuration that locks down the current ACP config unless factory reset is done.

With such means, the valid administration has the best chances to maintain access to ACP nodes, discover malicious configuration through ongoing configuration tracking from central locations for example, and to react accordingly.

The primary reaction is withdrawal/change of credentials, terminate malicious existing management sessions and fixing the configuration. Ensuring that management sessions using invalidated credentials are terminated automatically without recourse will likely require new work.

Only when these steps are not feasible would it be necessary to revoke or expire the ACP certificate credentials and consider the node kicked off the network - until the situation can be further rectified, likely requiring direct physical access to the node.

Without extensions, compromised ACP nodes can only be removed from the ACP at the speed of CRL/OCSP information refresh or expiry (and non-removal) of short lived certificates. Future extensions to the ACP could for example use GRASP flooding distribution of triggered updates of CRL/OCSP or explicit removal indication of the compromised nodes domain certificate.

#### A.9.9. Detecting ACP secure channel downgrade attacks

The following text proposes a mechanism to protect against downgrade attacks without introducing a new specialized UPFRONT GRASP secure channel mechanism. Instead, it relies on running GRASP after establishing a secure channel protocol to verify if the established secure channel option could have been the result of a MITM downgrade attack:

MITM attackers can force downgrade attacks for ACP secure channel selection by filtering/modifying DULL GRASP messages and/or actual secure channel data packets. For example, if at some point in time DTLS traffic could be easier decrypted than traffic of IKEv2, the MITM could filter all IKEv2 packets to force ACP nodes to use DTLS (assuming the ACP nodes in question supported both DTLS and IKEv2).

For cases where such MITM attacks are not capable to inject malicious traffic (but only to decrypt the traffic), a downgrade attack could be discovered after a secure channel connection is established, for example by use of the following type of mechanism:

After the secure channel connection is established, the two ACP peers negotiate via an appropriate (To Be Defined) GRASP negotiation which ACP secure channel protocol should have been selected between them (in the absence of a MITM attacker). This negotiation would have to

signal the DULL GRASP announced ACP secure channel options by each peer followed by an announcement of the preferred secure channel protocol by the ACP peer that is the Decider in the secure channel setup, e.g. the ACP peer that is deciding which secure channel protocol to pick. If that chosen secure channel protocol is different from the one that actually was chosen, then this mismatch is an indication that there is a MITM attacker or other similar issue (firewall prohibiting the use of specific protocols) that caused a non-preferred secure channel protocol to be chosen. This discovery could then result in mitigation options such as logging and ensuing investigations.

#### Appendix B. Unfinished considerations (To Be Removed From RFC)

[RFC-Editor: This whole appendix B. and its subsections to be removed for the RFC.

This appendix contains unfinished considerations that are removed from the RFC, they are maintained in this draft as a log of the state of discussion and point of reference. Together with this appendix, also the references pointing to it are marked to be removed from the RFC because no consensus could be reached that a self-reference to a draft version of the RFC is an appropriate breadcrumb to point to unfinished considerations.

The authors plan to move these considerations into a new target informational draft, please look for draft-eckert-anima-acp-considerations.

##### B.1. Considerations for improving secure channel negotiation

Proposed text from Benjamin Kaduk. It is suggested to replace the text of appendix A.6 in previous versions of this draft (up to version 29).

The discovery procedure in this specification for low-level ACP channel support by layer-2 peers involves DULL GRASP and attempting (usually in parallel) to establish all supported channel types, learning the peer ACP address and correspondingly the assignment of Decider and Follower roles, and tearing down all channels other than the one preferred by the Decider. This procedure, in general, becomes resource intensive as the number of possible secure channels grows; even worse, under some threat models, the security of the discovery result is only as strong as the weakest supported secure channel protocol. Furthermore, the unilateral determination of "best" channel type by the Decider does not result in the optimal outcome in all possible scenarios.

This situation is tolerable at present, with only two secure channels (DTLS and IPsec) defined, but long-term agility in the vein of [BCP201] will require the introduction of an alternate discovery/negotiation procedure. While IKEv2 is the IETF standard protocol for negotiating security associations, it currently does not have a defined mechanism flexible enough to negotiate the parameters needed for, e.g., an ACP DTLS channel, let alone for allowing ACP peers to indicate their preference metrics for channel selection. Such a mechanism or mechanisms could be defined, but if ACP agility requires introducing a new channel type, for example MacSec, IKEv2 would again need to be extended in order to negotiate an ACP MacSec association. Making ACP channel agility dependent on updates to IKEv2 is likely to result in obstacles due to different timescales of evolution, since IKEv2 implementations help form the core of Internet-scale security infrastructure and must accordingly be robust and thoroughly tested.

Accordingly, a dedicated ACP channel negotiation mechanism is appropriate as a way to provide long-term algorithm and secure-channel protocol agility. Such a mechanism is not currently defined, but one possible design is as follows. A new DULL GRASP objective is defined to indicate the GRASP-over-TLS channel, which is by definition preferred to other channel types (including DTLS and IPsec). When both peers advertise support for GRASP-over-TLS, GRASP-over-TLS must run to completion before other channel types are attempted. The GRASP-over-TLS channel performs the necessary negotiation by establishing a TLS connection between the peers and using that connection to secure a dedicated GRASP instance for negotiating supported channel types and preference metrics. This provides a rich language for determining what secure channel protocol to use for the ACP link while taking into account the capabilities and preferences of the ACP peers, all protected by the security of the TLS channel.

## B.2. ACP address verification

The `AcpNodeName` of most ACP nodes contains in the `acp-address` field the primary ACP address to be used by the node for end-to-end connections across ACP secure channels. Nevertheless, there is no verification of an ACP peers address specified in this document. This section explains the current understanding as to why this is not done.

Not all ACP nodes will have an actual IPv6 address in the `acp-address` field of their `AcpNodeName`. Those who do not include nodes that do not support ACP secure channels, such as pre-existing NOC equipment that only connects to the ACP via ACP connect interfaces. Likewise, future ACP node type that may want to have their Node-ID not be defined by an ACP registrar, but differently cannot have the ACP

address be provided in their ACP certificate where it would be defined by the registrar. In result, any scheme that would rely on verification of the acp-address in the ACP certificate would only apply to a subset of ACP nodes.

The transport stack network layer address used for ACP secure channels is not the acp-address. For automatically established ACP secure channels, it is a link-local IPv6 address. For explicitly configured ACP secure channels (to reach across non ACP L3 network segments), the address is any IPv4 or IPv6 address routable to that remote destination.

When the acp-address is actually used across the ACP, it can only be verified by a peer when the peer has the certificate of the peer. Unless further higher layer mechanisms are developed on top of the ACP (for example via ASA), the only mechanism to access a peers ACP certificate is for secure connections in which the peers certificates are exchanged and cryptographically verified, e.g. TLS and DTLS. Initially, it is expected that the ACP will carry many legacy network management control connections that unfortunately not end-to-end authenticated but that are solely protected by being carried across the ACP secure channels. ACP address verification therefore cannot be used for such connections without additional higher layer components.

For the remaining (TLS/DTLS) connections for which address verification can be used, the main question is: what additional benefit would address verification provide?

The main value that transport stack network layer address verification could provide for these type of connections would be the discovery of on-path transport proxies. For example, in case of BRSKI, pledges connect to an ACP registrar via an ASA implementing a TCP proxy because the pledge itself has at that point in time no ACP certificate valid to build ACP secure channels and hence needs to rely on such a proxy. This is one example where such a TCP proxy is required and not a form of attack.

In general, on path TCP proxies could be a form of attack, but it stands to reason, that an attacker that manages to enable a malicious TCP proxy could likely equally build a transparent proxy not changing the network layer addresses. Only when the attacker operates off-path would this option not be possible. Such attacks could indeed be possible: An impaired ACP node could announce itself as another service instance for a service whose utilization it wants to attack. It could then attempt to look like a valid server by simply TCP proxying the clients connections to a valid server and then attack the connections passing through it (passive decrypting or passive

fingerprint analysis). But like the BRSKI proxy, this behavior could be perfectly legitimate and not an attack. For example, TCP has in the past often suffered from performance issues across difficult (high capacity, high loss) paths, and TCP proxies where and are being used simply as a tool for isolating such path segments (such as a WAN), and providing caching and local-retransmit of in-transit packets, reducing the effective path segment capacity.

As explained elsewhere in this document already, considerations for these type of attack are therefore outside the scope of the ACP but fundamental to further design of the ASA infrastructure. Beyond distinguishing whether a TCP proxy would be beneficial or malicious, the even more fundamental question is how to determine from a multitude of service announcements which instance is the most trustworthy and functionally best. In the Internet/web, this question is NOT solved inside the network but through off-net human interaction ("trust me, the best search engine is www.<insert-your-personal-recommendation>.com").

### B.3. Public CA considerations

Public CAs are outside the scope of this document for the following reasons. This appendix describes the current state of understanding for those interested to consider utilizing public CA for the ACP in the future.

If public CA were to be used to enroll ACP nodes and act as TA, this would require a model in which the public CA would be able to assert the ownership of the information requested in the certificate, especially the AcpNodeName, for example mitigated by the domain registrar(s). Due to the use of a new, ACP unique encoding of the AcpNodeName, there is no mechanism for public CA to do so. More importantly though, isolation between ACPs of disjoint operated ACPs is achieved in the current ACP design through disjoint TA. A public CA is in general based on a single (set of) TA shared across all certificates signed by the CA.

Due to the fact that the ACP domain membership check also validates that a peers domain name in the AcpNodeName matches that of the ACP node itself, it would be possible to use the same TA across disjoint ACP domains, but the security and attack implications of such an approach are beyond the scope of this document.

The use of ULA addresses in the AcpNodeName is another novel aspect for certificates from a possible public CA. Typically, ULA addresses are not meant to be signed by a public CA when carried in an address field, because there is no ownership of a particular ULA address in the scope of the Internet, which is what public CA operate on.

Nevertheless, the ULA addresses used by the ACP are scoped to be valid only within the confines of a specific ACP as defined by the domain name in the `AcpNodeName`. However, this understanding has not been reviewed or accepted by any bodies responsible for policies of public CA.

Because in this specification, ACPs are isolated from each other primarily by their TA, when a public CA would intend to sign ACP certificates and using a single TA to sign TA of ACP certificates from different operators/domain, it could do so by ensuring that the domain name in the `AcpNodeName` was a globally owned DNS ACP domain name of the organization, and beyond that, it would need to validate that the ACP registrar of that domain who is mitigating the enrollment is authorized to vouch for the ownership of the `acp-` address within the scope of the ACP domain name.

#### B.4. Hardening DULL GRASP considerations

DULL GRASP suffers from similar type of DoS attacks as many link-local multicast discovery protocols, for example mDNS. Attackers on a subnet may be able to inject malicious DULL GRASP messages that are indistinguishable from non-malicious DULL GRASP messages to create Denial-of-Service (DoS) attacks that force ACP nodes to attempt many unsuccessful ACP secure channel connections.

When an ACP node sees multiple `AN_ACP` objectives for the same secure channel protocol on different transport addresses, it could prefer connecting via the well-known transport address if the secure channel method has one, such as UDP port 500 for IKEv2. For protocols such as (ACP secure channel over) DTLS for which there are no well defined port number, this heuristic does not provide benefits though.

DoS attack with port numbers can also be eliminated by relying on well known-port numbers implied by the GRASP method-name. For example, a future service name of "DTLSacp" could be defined to be associated only to a newly to be assigned well known UDP port for ACP over DTLS, and the port number in the GRASP transport address information would be ignored. Note that there is already a variety of ports assigned to specific protocols over DTLS by IANA, so especially for DTLS this would not be uncommon.

Authors' Addresses

Toerless Eckert (editor)  
Futurewei Technologies Inc. USA  
2330 Central Expy  
Santa Clara, 95050  
United States of America

Email: tte+ietf@cs.fau.de

Michael H. Behringer (editor)

Email: michael.h.behringer@gmail.com

Steinthor Bjarnason  
Arbor Networks  
2727 South State Street, Suite 200  
Ann Arbor, MI 48104  
United States

Email: sbjarnason@arbor.net

ANIMA WG  
Internet-Draft  
Intended status: Standards Track  
Expires: 15 May 2021

M. Pritikin  
Cisco  
M. Richardson  
Sandelman  
T.T.E. Eckert  
Futurewei USA  
M.H. Behringer

K.W. Watsen  
Watsen Networks  
11 November 2020

Bootstrapping Remote Secure Key Infrastructures (BRSKI)  
draft-ietf-anima-bootstrapping-keyinfra-45

Abstract

This document specifies automated bootstrapping of an Autonomic Control Plane. To do this a Secure Key Infrastructure is bootstrapped. This is done using manufacturer-installed X.509 certificates, in combination with a manufacturer's authorizing service, both online and offline. We call this process the Bootstrapping Remote Secure Key Infrastructure (BRSKI) protocol. Bootstrapping a new device can occur using a routable address and a cloud service, or using only link-local connectivity, or on limited/disconnected networks. Support for deployment models with less stringent security requirements is included. Bootstrapping is complete when the cryptographic identity of the new key infrastructure is successfully deployed to the device. The established secure connection can be used to deploy a locally issued certificate to the device as well.

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 <https://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 15 May 2021.

#### Copyright Notice

Copyright (c) 2020 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 (<https://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 . . . . .	5
1.1.	Prior Bootstrapping Approaches . . . . .	6
1.2.	Terminology . . . . .	8
1.3.	Scope of solution . . . . .	11
1.3.1.	Support environment . . . . .	11
1.3.2.	Constrained environments . . . . .	11
1.3.3.	Network Access Controls . . . . .	12
1.3.4.	Bootstrapping is not Booting . . . . .	12
1.4.	Leveraging the new key infrastructure / next steps . . .	12
1.5.	Requirements for Autonomic Network Infrastructure (ANI) devices . . . . .	13
2.	Architectural Overview . . . . .	13
2.1.	Behavior of a Pledge . . . . .	15
2.2.	Secure Imprinting using Vouchers . . . . .	16
2.3.	Initial Device Identifier . . . . .	17
2.3.1.	Identification of the Pledge . . . . .	18
2.3.2.	MASA URI extension . . . . .	19
2.4.	Protocol Flow . . . . .	20
2.5.	Architectural Components . . . . .	23
2.5.1.	Pledge . . . . .	23
2.5.2.	Join Proxy . . . . .	23
2.5.3.	Domain Registrar . . . . .	23
2.5.4.	Manufacturer Service . . . . .	23
2.5.5.	Public Key Infrastructure (PKI) . . . . .	24
2.6.	Certificate Time Validation . . . . .	24
2.6.1.	Lack of realtime clock . . . . .	24
2.6.2.	Infinite Lifetime of IDevID . . . . .	24
2.7.	Cloud Registrar . . . . .	25
2.8.	Determining the MASA to contact . . . . .	25
3.	Voucher-Request artifact . . . . .	26

3.1.	Nonceless Voucher Requests . . . . .	27
3.2.	Tree Diagram . . . . .	27
3.3.	Examples . . . . .	27
3.4.	YANG Module . . . . .	29
4.	Proxying details (Pledge - Proxy - Registrar) . . . . .	33
4.1.	Pledge discovery of Proxy . . . . .	34
4.1.1.	Proxy GRASP announcements . . . . .	35
4.2.	CoAP connection to Registrar . . . . .	37
4.3.	Proxy discovery and communication of Registrar . . . . .	37
5.	Protocol Details (Pledge - Registrar - MASA) . . . . .	38
5.1.	BRSKI-EST TLS establishment details . . . . .	40
5.2.	Pledge Requests Voucher from the Registrar . . . . .	41
5.3.	Registrar Authorization of Pledge . . . . .	43
5.4.	BRSKI-MASA TLS establishment details . . . . .	43
5.4.1.	MASA authentication of customer Registrar . . . . .	44
5.5.	Registrar Requests Voucher from MASA . . . . .	45
5.5.1.	MASA renewal of expired vouchers . . . . .	47
5.5.2.	MASA pinning of registrar . . . . .	48
5.5.3.	MASA checking of voucher request signature . . . . .	48
5.5.4.	MASA verification of domain registrar . . . . .	49
5.5.5.	MASA verification of pledge prior-signed-voucher-request . . . . .	50
5.5.6.	MASA nonce handling . . . . .	50
5.6.	MASA and Registrar Voucher Response . . . . .	50
5.6.1.	Pledge voucher verification . . . . .	53
5.6.2.	Pledge authentication of provisional TLS connection . . . . .	54
5.7.	Pledge BRSKI Status Telemetry . . . . .	55
5.8.	Registrar audit-log request . . . . .	56
5.8.1.	MASA audit log response . . . . .	58
5.8.2.	Calculation of domainID . . . . .	60
5.8.3.	Registrar audit log verification . . . . .	61
5.9.	EST Integration for PKI bootstrapping . . . . .	62
5.9.1.	EST Distribution of CA Certificates . . . . .	63
5.9.2.	EST CSR Attributes . . . . .	63
5.9.3.	EST Client Certificate Request . . . . .	64
5.9.4.	Enrollment Status Telemetry . . . . .	64
5.9.5.	Multiple certificates . . . . .	65
5.9.6.	EST over CoAP . . . . .	66
6.	Clarification of transfer-encoding . . . . .	66
7.	Reduced security operational modes . . . . .	66
7.1.	Trust Model . . . . .	66
7.2.	Pledge security reductions . . . . .	67
7.3.	Registrar security reductions . . . . .	68
7.4.	MASA security reductions . . . . .	69
7.4.1.	Issuing Nonceless vouchers . . . . .	69
7.4.2.	Trusting Owners on First Use . . . . .	70
7.4.3.	Updating or extending voucher trust anchors . . . . .	71

8.	IANA Considerations . . . . .	72
8.1.	The IETF XML Registry . . . . .	72
8.2.	YANG Module Names Registry . . . . .	72
8.3.	BRSKI well-known considerations . . . . .	72
8.3.1.	BRSKI .well-known registration . . . . .	72
8.3.2.	BRSKI .well-known registry . . . . .	73
8.4.	PKIX Registry . . . . .	73
8.5.	Pledge BRSKI Status Telemetry . . . . .	73
8.6.	DNS Service Names . . . . .	74
8.7.	GRASP Objective Names . . . . .	74
9.	Applicability to the Autonomic Control Plane (ACP) . . . . .	74
9.1.	Operational Requirements . . . . .	76
9.1.1.	MASA Operational Requirements . . . . .	76
9.1.2.	Domain Owner Operational Requirements . . . . .	77
9.1.3.	Device Operational Requirements . . . . .	77
10.	Privacy Considerations . . . . .	78
10.1.	MASA audit log . . . . .	78
10.2.	What BRSKI-EST reveals . . . . .	78
10.3.	What BRSKI-MASA reveals to the manufacturer . . . . .	79
10.4.	Manufacturers and Used or Stolen Equipment . . . . .	81
10.5.	Manufacturers and Grey market equipment . . . . .	83
10.6.	Some mitigations for meddling by manufacturers . . . . .	83
10.7.	Death of a manufacturer . . . . .	84
11.	Security Considerations . . . . .	85
11.1.	Denial of Service (DoS) against MASA . . . . .	86
11.2.	DomainID must be resistant to second-preimage attacks . . . . .	86
11.3.	Availability of good random numbers . . . . .	87
11.4.	Freshness in Voucher-Requests . . . . .	87
11.5.	Trusting manufacturers . . . . .	88
11.6.	Manufacturer Maintenance of trust anchors . . . . .	89
11.6.1.	Compromise of Manufacturer IDevID signing keys . . . . .	91
11.6.2.	Compromise of MASA signing keys . . . . .	91
11.6.3.	Compromise of MASA web service . . . . .	93
11.7.	YANG Module Security Considerations . . . . .	94
12.	Acknowledgements . . . . .	94
13.	References . . . . .	94
13.1.	Normative References . . . . .	94
13.2.	Informative References . . . . .	98
Appendix A.	IPv4 and non-ANI operations . . . . .	102
A.1.	IPv4 Link Local addresses . . . . .	102
A.2.	Use of DHCPv4 . . . . .	102
Appendix B.	mDNS / DNSSD proxy discovery options . . . . .	102
Appendix C.	Example Vouchers . . . . .	103
C.1.	Keys involved . . . . .	103
C.1.1.	Manufacturer Certificate Authority for IDevID signatures . . . . .	104
C.1.2.	MASA key pair for voucher signatures . . . . .	105
C.1.3.	Registrar Certificate Authority . . . . .	107

C.1.4. Registrar key pair . . . . .	108
C.1.5. Pledge key pair . . . . .	110
C.2. Example process . . . . .	111
C.2.1. Pledge to Registrar . . . . .	111
C.2.2. Registrar to MASA . . . . .	115
C.2.3. MASA to Registrar . . . . .	121
Appendix D. Additional References . . . . .	125
Authors' Addresses . . . . .	125

## 1. Introduction

The Bootstrapping Remote Secure Key Infrastructure (BRSKI) protocol provides a solution for secure zero-touch (automated) bootstrap of new (unconfigured) devices that are called pledges in this document. Pledges have an IDevID installed in them at the factory.

"BRSKI" is pronounced like "brewski", a colloquial term for beer in Canada and parts of the US-midwest. [brewski]

This document primarily provides for the needs of the ISP and Enterprise focused ANIMA Autonomic Control Plane (ACP) [I-D.ietf-anima-autonomic-control-plane]. This bootstrap process satisfies the [RFC7575] requirements of section 3.3 of making all operations secure by default. Other users of the BRSKI protocol will need to provide separate applicability statements that include privacy and security considerations appropriate to that deployment. Section 9 explains the detailed applicability for this the ACP usage.

The BRSKI protocol requires a significant amount of communication between manufacturer and owner: in its default modes it provides a cryptographic transfer of control to the initial owner. In its strongest modes, it leverages sales channel information to identify the owner in advance. Resale of devices is possible, provided that the manufacturer is willing to authorize the transfer. Mechanisms to enable transfers of ownership without manufacturer authorization are not included in this version of the protocol, but could be designed into future versions.

This document describes how pledges discover (or are discovered by) an element of the network domain to which the pledge belongs that will perform the bootstrap. This element (device) is called the registrar. Before any other operation, pledge and registrar need to establish mutual trust:

1. Registrar authenticating the pledge: "Who is this device? What is its identity?"

2. Registrar authorizing the pledge: "Is it mine? Do I want it? What are the chances it has been compromised?"
3. Pledge authenticating the registrar: "What is this registrar's identity?"
4. Pledge authorizing the registrar: "Should I join this network?"

This document details protocols and messages to answer the above questions. It uses a TLS connection and an PKIX-shaped (X.509v3) certificate (an IEEE 802.1AR [IDevID] IDevID) of the pledge to answer points 1 and 2. It uses a new artifact called a "voucher" that the registrar receives from a "Manufacturer Authorized Signing Authority" (MASA) and passes to the pledge to answer points 3 and 4.

A proxy provides very limited connectivity between the pledge and the registrar.

The syntactic details of vouchers are described in detail in [RFC8366]. This document details automated protocol mechanisms to obtain vouchers, including the definition of a 'voucher-request' message that is a minor extension to the voucher format (see Section 3) defined by [RFC8366].

BRSKI results in the pledge storing an X.509 root certificate sufficient for verifying the registrar identity. In the process a TLS connection is established that can be directly used for Enrollment over Secure Transport (EST). In effect BRSKI provides an automated mechanism for the "Bootstrap Distribution of CA Certificates" described in [RFC7030] Section 4.1.1 wherein the pledge "MUST [...] engage a human user to authorize the CA certificate using out-of-band" information. With BRSKI the pledge now can automate this process using the voucher. Integration with a complete EST enrollment is optional but trivial.

BRSKI is agile enough to support bootstrapping alternative key infrastructures, such as a symmetric key solutions, but no such system is described in this document.

### 1.1. Prior Bootstrapping Approaches

To literally "pull yourself up by the bootstraps" is an impossible action. Similarly the secure establishment of a key infrastructure without external help is also an impossibility. Today it is commonly accepted that the initial connections between nodes are insecure, until key distribution is complete, or that domain-specific keying material (often pre-shared keys, including mechanisms like SIM cards) is pre-provisioned on each new device in a costly and non-scalable

manner. Existing automated mechanisms are known as non-secured 'Trust on First Use' (TOFU) [RFC7435], 'resurrecting duckling' [Stajano99theresurrecting] or 'pre-staging'.

Another prior approach has been to try and minimize user actions during bootstrapping, but not eliminate all user-actions. The original EST protocol [RFC7030] does reduce user actions during bootstrap but does not provide solutions for how the following protocol steps can be made autonomic (not involving user actions):

- \* using the Implicit Trust Anchor [RFC7030] database to authenticate an owner specific service (not an autonomic solution because the URL must be securely distributed),
- \* engaging a human user to authorize the CA certificate using out-of-band data (not an autonomic solution because the human user is involved),
- \* using a configured Explicit TA database (not an autonomic solution because the distribution of an explicit TA database is not autonomic),
- \* and using a Certificate-Less TLS mutual authentication method (not an autonomic solution because the distribution of symmetric key material is not autonomic).

These "touch" methods do not meet the requirements for zero-touch.

There are "call home" technologies where the pledge first establishes a connection to a well known manufacturer service using a common client-server authentication model. After mutual authentication, appropriate credentials to authenticate the target domain are transferred to the pledge. This creates several problems and limitations:

- \* the pledge requires realtime connectivity to the manufacturer service,
- \* the domain identity is exposed to the manufacturer service (this is a privacy concern),
- \* the manufacturer is responsible for making the authorization decisions (this is a liability concern),

BRSKI addresses these issues by defining extensions to the EST protocol for the automated distribution of vouchers.

## 1.2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

The following terms are defined for clarity:

**ANI:** The Autonomic Network Infrastructure as defined by [I-D.ietf-anima-reference-model]. Section 9 details specific requirements for pledges, proxies and registrars when they are part of an ANI.

**Circuit Proxy:** A stateful implementation of the join proxy. This is the assumed type of proxy.

**drop-ship:** The physical distribution of equipment containing the "factory default" configuration to a final destination. In zero-touch scenarios there is no staging or pre-configuration during drop-ship.

**Domain:** The set of entities that share a common local trust anchor. This includes the proxy, registrar, Domain Certificate Authority, Management components and any existing entity that is already a member of the domain.

**domainID:** The domain IDentity is a unique value based upon the Registrar CA's certificate. Section 5.8.2 specifies how it is calculated.

**Domain CA:** The domain Certification Authority (CA) provides certification functionalities to the domain. At a minimum it provides certification functionalities to a registrar and manages the private key that defines the domain. Optionally, it certifies all elements.

**enrollment:** The process where a device presents key material to a network and acquires a network-specific identity. For example when a certificate signing request is presented to a certification authority and a certificate is obtained in response.

**imprint:** The process where a device obtains the cryptographic key material to identify and trust future interactions with a network. This term is taken from Konrad Lorenz's work in biology with new ducklings: during a critical period, the duckling would assume that anything that looks like a mother duck is in fact their

mother. An equivalent for a device is to obtain the fingerprint of the network's root certification authority certificate. A device that imprints on an attacker suffers a similar fate to a duckling that imprints on a hungry wolf. Securely imprinting is a primary focus of this document [imprinting]. The analogy to Lorenz's work was first noted in [Stajano99theresurrecting].

**IDevID:** An Initial Device Identity X.509 certificate installed by the vendor on new equipment. This is a term from 802.1AR [IDevID]

**IPIP Proxy:** A stateless proxy alternative.

**Join Proxy:** A domain entity that helps the pledge join the domain. A join proxy facilitates communication for devices that find themselves in an environment where they are not provided connectivity until after they are validated as members of the domain. For simplicity this document sometimes uses the term of 'proxy' to indicate the join proxy. The pledge is unaware that they are communicating with a proxy rather than directly with a registrar.

**Join Registrar (and Coordinator):** A representative of the domain that is configured, perhaps autonomically, to decide whether a new device is allowed to join the domain. The administrator of the domain interfaces with a "join registrar (and coordinator)" to control this process. Typically a join registrar is "inside" its domain. For simplicity this document often refers to this as just "registrar". Within [I-D.ietf-anima-reference-model] this is referred to as the "join registrar autonomic service agent". Other communities use the abbreviation "JRC".

**LDevID:** A Local Device Identity X.509 certificate installed by the owner of the equipment. This is a term from 802.1AR [LDevID]

**manufacturer:** the term manufacturer is used throughout this document to be the entity that created the device. This is typically the "original equipment manufacturer" or OEM, but in more complex situations it could be a "value added retailer" (VAR), or possibly even a systems integrator. In general, it a goal of BRSKI to eliminate small distinctions between different sales channels. The reason for this is that it permits a single device, with a uniform firmware load, to be shipped directly to all customers. This eliminates costs for the manufacturer. This also reduces the number of products supported in the field increasing the chance that firmware will be more up to date.

**MASA Audit-Log:** An anonymized list of previous owners maintained by

the MASA on a per device (per pledge) basis. Described in Section 5.8.1.

**MASA Service:** A third-party Manufacturer Authorized Signing Authority (MASA) service on the global Internet. The MASA signs vouchers. It also provides a repository for audit-log information of privacy protected bootstrapping events. It does not track ownership.

**nonced:** a voucher (or request) that contains a nonce (the normal case).

**nonceless:** a voucher (or request) that does not contain a nonce, relying upon accurate clocks for expiration, or which does not expire.

**offline:** When an architectural component cannot perform realtime communications with a peer, either due to network connectivity or because the peer is turned off, the operation is said to be occurring offline.

**Ownership Tracker:** An Ownership Tracker service on the global Internet. The Ownership Tracker uses business processes to accurately track ownership of all devices shipped against domains that have purchased them. Although optional, this component allows vendors to provide additional value in cases where their sales and distribution channels allow for accurate tracking of such ownership. Ownership tracking information is indicated in vouchers as described in [RFC8366]

**Pledge:** The prospective (unconfigured) device, which has an identity installed at the factory.

**(Public) Key Infrastructure:** The collection of systems and processes that sustain the activities of a public key system. The registrar acts as an [RFC5280] and [RFC5272] (see section 7) "Registration Authority".

**TOFU:** Trust on First Use. Used similarly to [RFC7435]. This is where a pledge device makes no security decisions but rather simply trusts the first registrar it is contacted by. This is also known as the "resurrecting duckling" model.

**Voucher:** A signed artifact from the MASA that indicates to a pledge the cryptographic identity of the registrar it should trust. There are different types of vouchers depending on how that trust is asserted. Multiple voucher types are defined in [RFC8366]

### 1.3. Scope of solution

#### 1.3.1. Support environment

This solution (BRSKI) can support large router platforms with multi-gigabit inter-connections, mounted in controlled access data centers. But this solution is not exclusive to large equipment: it is intended to scale to thousands of devices located in hostile environments, such as ISP provided CPE devices which are drop-shipped to the end user. The situation where an order is fulfilled from distributed warehouse from a common stock and shipped directly to the target location at the request of a domain owner is explicitly supported. That stock ("SKU") could be provided to a number of potential domain owners, and the eventual domain owner will not know a-priori which device will go to which location.

The bootstrapping process can take minutes to complete depending on the network infrastructure and device processing speed. The network communication itself is not optimized for speed; for privacy reasons, the discovery process allows for the pledge to avoid announcing its presence through broadcasting.

Nomadic or mobile devices often need to acquire credentials to access the network at the new location. An example of this is mobile phone roaming among network operators, or even between cell towers. This is usually called handoff. BRSKI does not provide a low-latency handoff which is usually a requirement in such situations. For these solutions BRSKI can be used to create a relationship (an LDevID) with the "home" domain owner. The resulting credentials are then used to provide credentials more appropriate for a low-latency handoff.

#### 1.3.2. Constrained environments

Questions have been posed as to whether this solution is suitable in general for Internet of Things (IoT) networks. This depends on the capabilities of the devices in question. The terminology of [RFC7228] is best used to describe the boundaries.

The solution described in this document is aimed in general at non-constrained (i.e., class 2+ [RFC7228]) devices operating on a non-Challenged network. The entire solution as described here is not intended to be useable as-is by constrained devices operating on challenged networks (such as 802.15.4 Low-power Lossy Networks (LLN)s).

Specifically, there are protocol aspects described here that might result in congestion collapse or energy-exhaustion of intermediate battery powered routers in an LLN. Those types of networks should

not use this solution. These limitations are predominately related to the large credential and key sizes required for device authentication. Defining symmetric key techniques that meet the operational requirements is out-of-scope but the underlying protocol operations (TLS handshake and signing structures) have sufficient algorithm agility to support such techniques when defined.

The imprint protocol described here could, however, be used by non-energy constrained devices joining a non-constrained network (for instance, smart light bulbs are usually mains powered, and speak 802.11). It could also be used by non-constrained devices across a non-energy constrained, but challenged network (such as 802.15.4). The certificate contents, and the process by which the four questions above are resolved do apply to constrained devices. It is simply the actual on-the-wire imprint protocol that could be inappropriate.

### 1.3.3. Network Access Controls

This document presumes that network access control has either already occurred, is not required, or is integrated by the proxy and registrar in such a way that the device itself does not need to be aware of the details. Although the use of an X.509 Initial Device Identity is consistent with IEEE 802.1AR [IDevID], and allows for alignment with 802.1X network access control methods, its use here is for pledge authentication rather than network access control. Integrating this protocol with network access control, perhaps as an Extensible Authentication Protocol (EAP) method (see [RFC3748]), is out-of-scope.

### 1.3.4. Bootstrapping is not Booting

This document describes "bootstrapping" as the protocol used to obtain a local trust anchor. It is expected that this trust anchor, along with any additional configuration information subsequently installed, is persisted on the device across system restarts ("booting"). Bootstrapping occurs only infrequently such as when a device is transferred to a new owner or has been reset to factory default settings.

## 1.4. Leveraging the new key infrastructure / next steps

As a result of the protocol described herein, the bootstrapped devices have the Domain CA trust anchor in common. An end entity certificate has optionally been issued from the Domain CA. This makes it possible to securely deploy functionalities across the domain, e.g:

- \* Device management.

- \* Routing authentication.
- \* Service discovery.

The major intended benefit is that it possible to use the credentials deployed by this protocol to secure the Autonomic Control Plane (ACP) ([I-D.ietf-anima-autonomic-control-plane]).

### 1.5. Requirements for Autonomic Network Infrastructure (ANI) devices

The BRSKI protocol can be used in a number of environments. Some of the options in this document are the result of requirements that are out of the ANI scope. This section defines the base requirements for ANI devices.

For devices that intend to become part of an Autonomic Network Infrastructure (ANI) ([I-D.ietf-anima-reference-model]) that includes an Autonomic Control Plane ([I-D.ietf-anima-autonomic-control-plane]), the BRSKI protocol MUST be implemented.

The pledge must perform discovery of the proxy as described in Section 4.1 using Generic Autonomic Signaling Protocol (GRASP)'s DULL [I-D.ietf-anima-grasp] M\_FLOOD announcements.

Upon successfully validating a voucher artifact, a status telemetry MUST be returned. See Section 5.7.

An ANIMA ANI pledge MUST implement the EST automation extensions described in Section 5.9. They supplement the [RFC7030] EST to better support automated devices that do not have an end user.

The ANI Join Registrar Autonomic Service Agent (ASA) MUST support all the BRSKI and above listed EST operations.

All ANI devices SHOULD support the BRSKI proxy function, using circuit proxies over the ACP. (See Section 4.3)

## 2. Architectural Overview

The logical elements of the bootstrapping framework are described in this section. Figure 1 provides a simplified overview of the components.

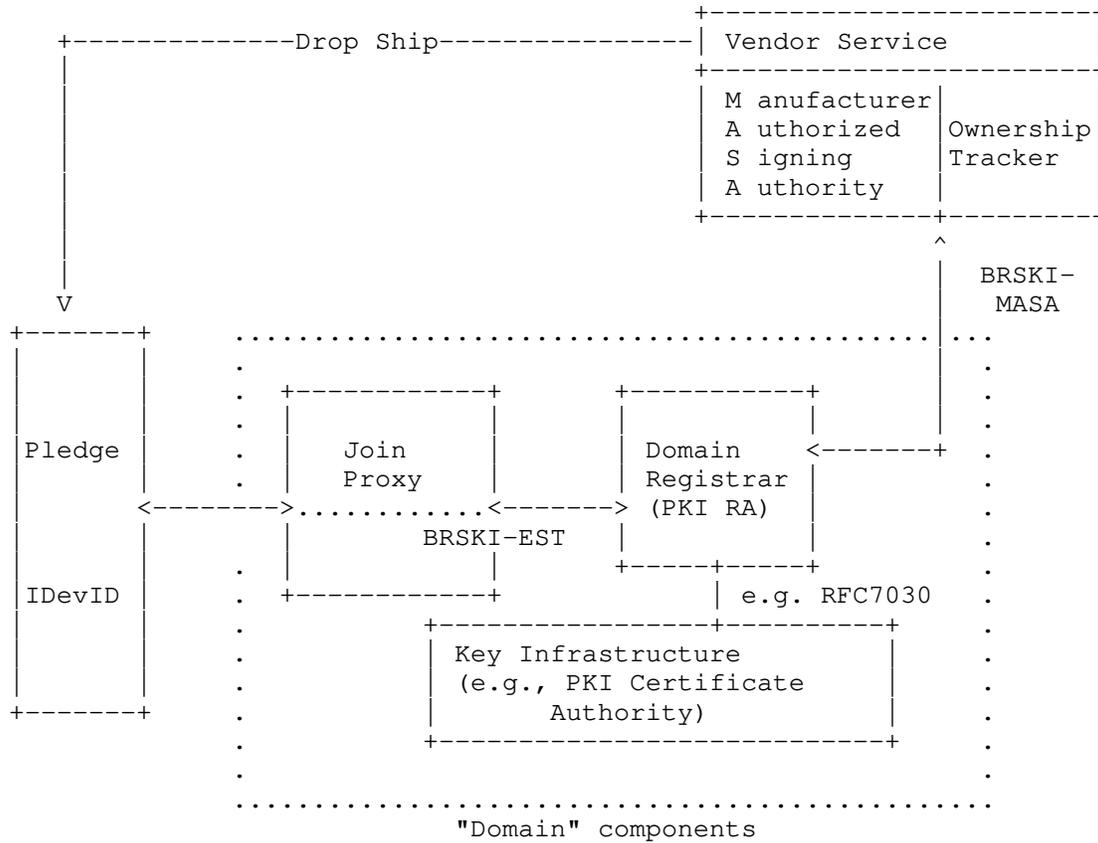


Figure 1: Architecture Overview

We assume a multi-vendor network. In such an environment there could be a Manufacturer Service for each manufacturer that supports devices following this document's specification, or an integrator could provide a generic service authorized by multiple manufacturers. It is unlikely that an integrator could provide Ownership Tracking services for multiple manufacturers due to the required sales channel integrations necessary to track ownership.

The domain is the managed network infrastructure with a Key Infrastructure the pledge is joining. The domain provides initial device connectivity sufficient for bootstrapping through a proxy. The domain registrar authenticates the pledge, makes authorization decisions, and distributes vouchers obtained from the Manufacturer Service. Optionally the registrar also acts as a PKI Certification Authority.

2.1. Behavior of a Pledge

The pledge goes through a series of steps, which are outlined here at a high level.

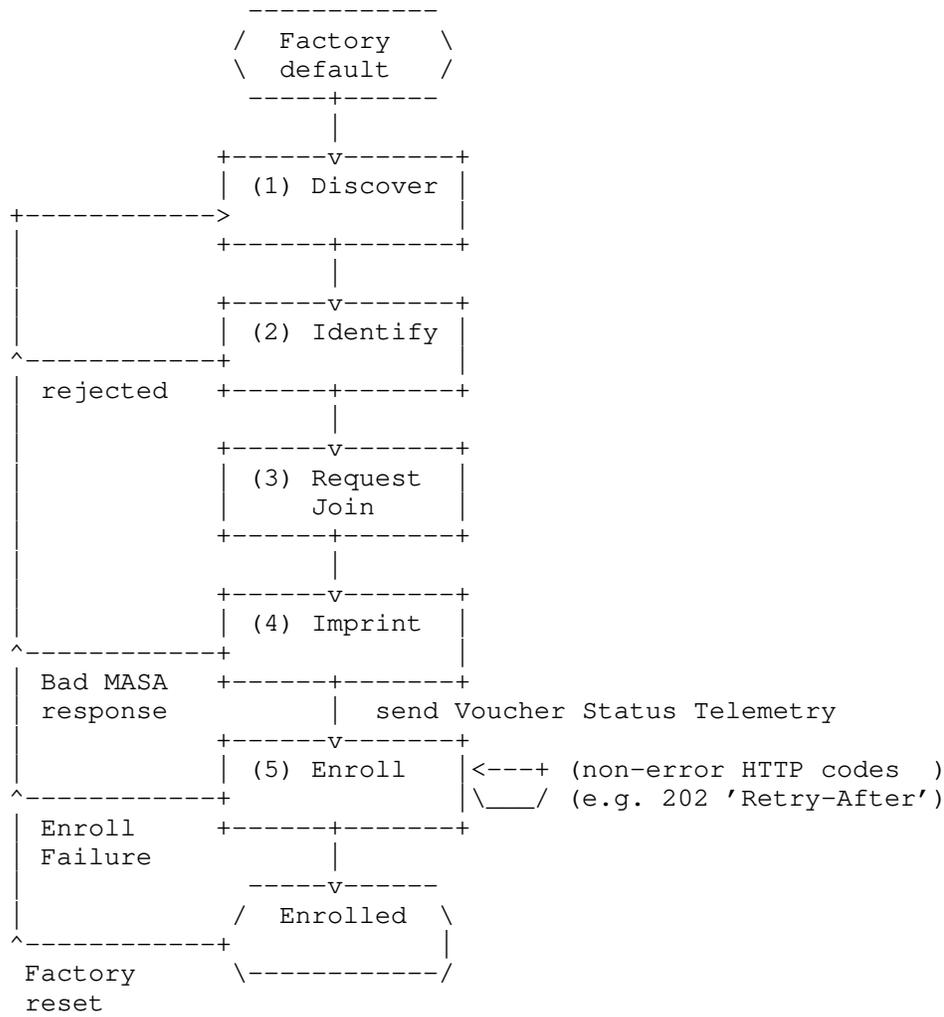


Figure 2: Pledge State Diagram

State descriptions for the pledge are as follows:

1. Discover a communication channel to a registrar.

2. Identify itself. This is done by presenting an X.509 IDevID credential to the discovered registrar (via the proxy) in a TLS handshake. (The registrar credentials are only provisionally accepted at this time).
3. Request to join the discovered registrar. A unique nonce is included ensuring that any responses can be associated with this particular bootstrapping attempt.
4. Imprint on the registrar. This requires verification of the manufacturer-service-provided voucher. A voucher contains sufficient information for the pledge to complete authentication of a registrar. This document details this step in depth.
5. Enroll. After imprint an authenticated TLS (HTTPS) connection exists between pledge and registrar. Enrollment over Secure Transport (EST) [RFC7030] can then be used to obtain a domain certificate from a registrar.

The pledge is now a member of, and can be managed by, the domain and will only repeat the discovery aspects of bootstrapping if it is returned to factory default settings.

This specification details integration with EST enrollment so that pledges can optionally obtain a locally issued certificate, although any Representational State Transfer (REST) (see [REST]) interface could be integrated in future work.

## 2.2. Secure Imprinting using Vouchers

A voucher is a cryptographically protected artifact (using a digital signature) to the pledge device authorizing a zero-touch imprint on the registrar domain.

The format and cryptographic mechanism of vouchers is described in detail in [RFC8366].

Vouchers provide a flexible mechanism to secure imprinting: the pledge device only imprints when a voucher can be validated. At the lowest security levels the MASA can indiscriminately issue vouchers and log claims of ownership by domains. At the highest security levels issuance of vouchers can be integrated with complex sales channel integrations that are beyond the scope of this document. The sales channel integration would verify actual (legal) ownership of the pledge by the domain. This provides the flexibility for a number of use cases via a single common protocol mechanism on the pledge and registrar devices that are to be widely deployed in the field. The MASA services have the flexibility to leverage either the currently defined claim mechanisms or to experiment with higher or lower security levels.

Vouchers provide a signed but non-encrypted communication channel among the pledge, the MASA, and the registrar. The registrar maintains control over the transport and policy decisions, allowing the local security policy of the domain network to be enforced.

### 2.3. Initial Device Identifier

Pledge authentication and pledge voucher-request signing is via a PKIX-shaped certificate installed during the manufacturing process. This is the 802.1AR Initial Device Identifier (IDevID), and it provides a basis for authenticating the pledge during the protocol exchanges described here. There is no requirement for a common root PKI hierarchy. Each device manufacturer can generate its own root certificate. Specifically, the IDevID enables:

1. Uniquely identifying the pledge by the Distinguished Name (DN) and subjectAltName (SAN) parameters in the IDevID. The unique identification of a pledge in the voucher objects are derived from those parameters as described below. Section 10.3 discusses privacy implications of the identifier.
2. Provides a cryptographic authentication of the pledge to the Registrar (see Section 5.3).
3. Secure auto-discovery of the pledge's MASA by the registrar (see Section 2.8).
4. Signing of voucher-request by the pledge's IDevID (see Section 3).
5. Provides a cryptographic authentication of the pledge to the MASA (see Section 5.5.5).

Section 7.2.13 (2009 edition) and section 8.10.3 (2018 edition) of [IDevID] discusses keyUsage and extendedKeyUsage extensions in the IDevID certificate. [IDevID] acknowledges that adding restrictions in the certificate limits applicability of these long-lived certificates. This specification emphasizes this point, and therefore RECOMMENDS that no key usage restrictions be included. This is consistent with [RFC5280] section 4.2.1.3, which does not require key usage restrictions for end entity certificates.

### 2.3.1. Identification of the Pledge

In the context of BRSKI, pledges have a 1:1 relationship with a "serial-number". This serial-number is used both in the "serial-number" field of voucher or voucher-requests (see Section 3) and in local policies on registrar or MASA (see Section 5).

There is a (certificate) serialNumber field is defined in [RFC5280] section 4.1.2.2. In the ASN.1, this is referred to as the CertificateSerialNumber. This field is NOT relevant to this specification. Do not confuse this field with the "serial-number" defined by this document, or by [IDevID] and [RFC4519] section 2.31.

The device serial number is defined in [RFC5280] section A.1 and A.2 as the X520SerialNumber, with the OID tag id-at-serialNumber.

The device serial number field (X520SerialNumber) is used as follows by the pledge to build the "serial-number" that is placed in the voucher-request. In order to build it, the fields need to be converted into a serial-number of "type string".

An example of a printable form of the "serialNumber" field is provided in [RFC4519] section 2.31 ("WI-3005"). That section further provides equality and syntax attributes.

Due to the reality of existing device identity provisioning processes, some manufacturers have stored serial-numbers in other fields. Registrar's SHOULD be configurable, on a per-manufacturer basis, to look for serial-number equivalents in other fields.

As explained in Section 5.5 the Registrar MUST extract the serial-number again itself from the pledge's TLS certificate. It can consult the serial-number in the pledge-request if there are any possible confusion about the source of the serial-number.

### 2.3.2. MASA URI extension

This document defines a new PKIX non-critical certificate extension to carry the MASA URI. This extension is intended to be used in the IDevID certificate. The URI is represented as described in Section 7.4 of [RFC5280].

The URI provides the authority information. The BRSKI `"/.well-known"` tree ([RFC5785]) is described in Section 5.

A complete URI MAY be in this extension, including the `'scheme'`, `'authority'`, and `'path'`. The complete URI will typically be used in diagnostic or experimental situations. Typically, (and in consideration to constrained systems), this SHOULD be reduced to only the `'authority'`, in which case a scheme of `"https://"` ([RFC7230] section 2.7.3) and `'path'` of `"/.well-known/brski"` is to be assumed.

The registrar can assume that only the `'authority'` is present in the extension, if there are no slash (`"/"`) characters in the extension.

Section 7.4 of [RFC5280] calls out various schemes that MUST be supported, including LDAP, HTTP and FTP. However, the registrar MUST use HTTPS for the BRSKI-MASA connection.

The new extension is identified as follows:

```
<CODE BEGINS>
MASAURLExtnModule-2016 { iso(1) identified-organization(3) dod(6)
internet(1) security(5) mechanisms(5) pkix(7)
id-mod(0) id-mod-MASAURLExtn2016(TBD) }

DEFINITIONS IMPLICIT TAGS ::= BEGIN

-- EXPORTS ALL --

IMPORTS
EXTENSION
FROM PKIX-CommonTypes-2009
  { iso(1) identified-organization(3) dod(6) internet(1)
  security(5) mechanisms(5) pkix(7) id-mod(0)
  id-mod-pkixCommon-02(57) }

id-pe FROM PKIX1Explicit-2009
  { iso(1) identified-organization(3) dod(6) internet(1)
  security(5) mechanisms(5) pkix(7) id-mod(0)
  id-mod-pkix1-explicit-02(51) } ;

MASACertExtensions EXTENSION ::= { ext-MASAURL, ... }
ext-MASAURL EXTENSION ::= { SYNTAX MASAURLSyntax
IDENTIFIED BY id-pe-masa-url }

id-pe-masa-url OBJECT IDENTIFIER ::= { id-pe TBD }

MASAURLSyntax ::= IA5String

END
<CODE ENDS>
```

Figure 3: MASAURL ASN.1 Module

The choice of id-pe is based on guidance found in Section 4.2.2 of [RFC5280], "These extensions may be used to direct applications to on-line information about the issuer or the subject". The MASA URL is precisely that: online information about the particular subject.

#### 2.4. Protocol Flow

A representative flow is shown in Figure 4

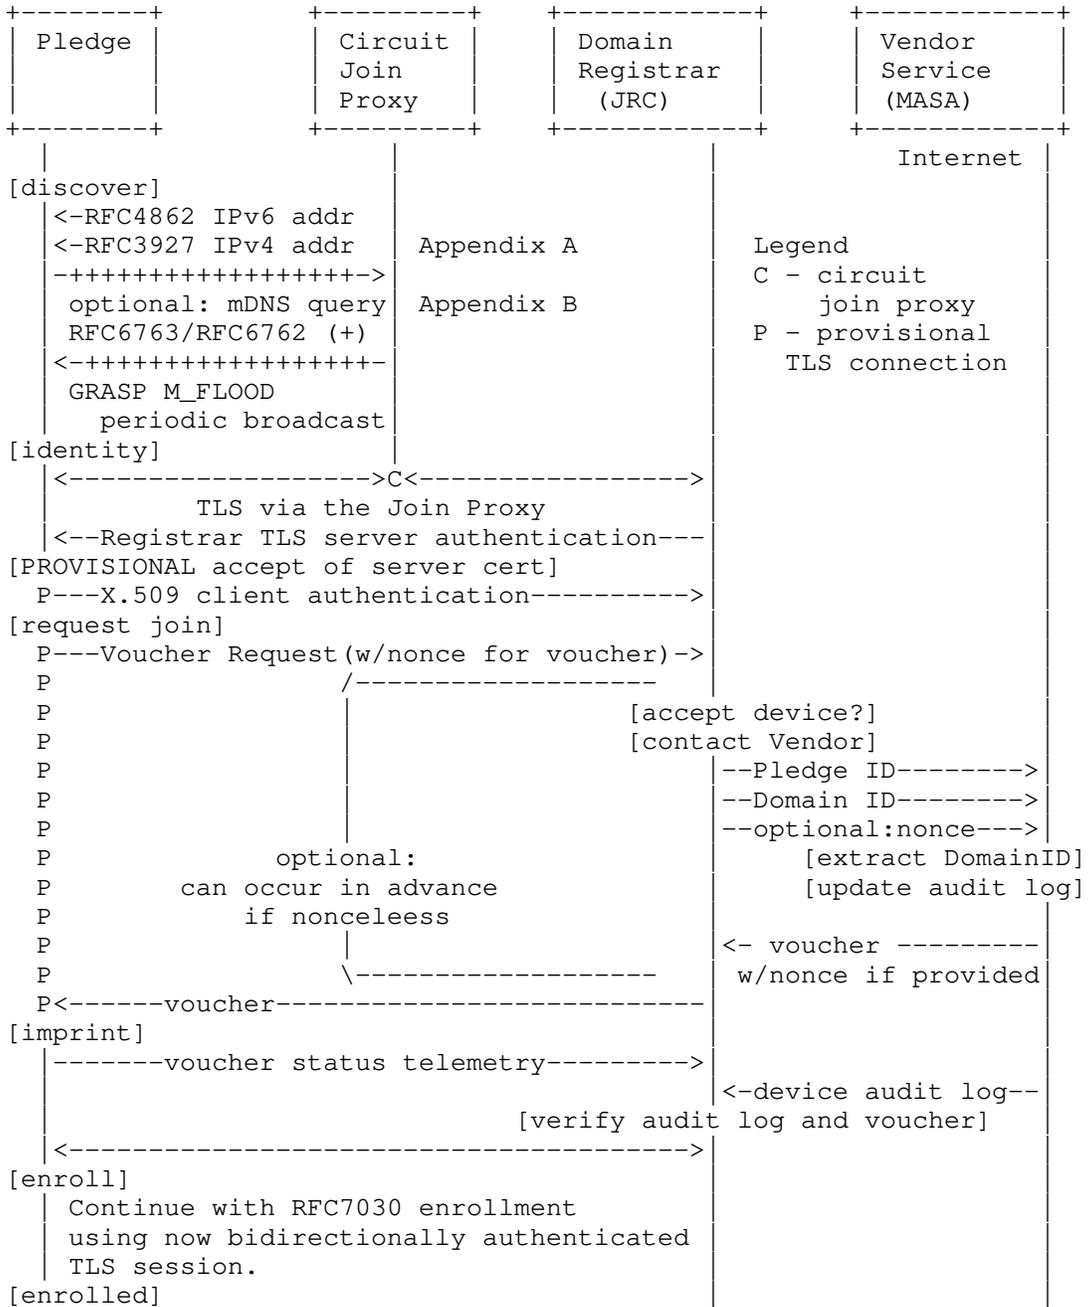


Figure 4: Protocol Time Sequence Diagram

On initial bootstrap, a new device (the pledge) uses a local service autodiscovery (GRASP or mDNS) to locate a join proxy. The join proxy connects the pledge to a local registrar (the JRC).

Having found a candidate registrar, the fledgling pledge sends some information about itself to the registrar, including its serial number in the form of a voucher request and its device identity certificate (IDevID) as part of the TLS session.

The registrar can determine whether it expected such a device to appear, and locates a MASA. The location of the MASA is usually found in an extension in the IDevID. Having determined that the MASA is suitable, the entire information from the initial voucher request (including device serial number) is transmitted over the internet in a TLS protected channel to the manufacturer, along with information about the registrar/owner.

The manufacturer can then apply policy based on the provided information, as well as other sources of information (such as sales records), to decide whether to approve the claim by the registrar to own the device; if the claim is accepted, a voucher is issued that directs the device to accept its new owner.

The voucher is returned to the registrar, but not immediately to the device -- the registrar has an opportunity to examine the voucher, the MASA's audit-logs, and other sources of information to determine whether the device has been tampered with, and whether the bootstrap should be accepted.

No filtering of information is possible in the signed voucher, so this is a binary yes-or-no decision. If the registrar accepts the voucher as a proper one for its device, the voucher is returned to the pledge for imprinting.

The voucher also includes a trust anchor that the pledge uses as representing the owner. This is used to successfully bootstrap from an environment where only the manufacturer has built-in trust by the device into an environment where the owner now has a PKI footprint on the device.

When BRSKI is followed with EST this single footprint is further leveraged into the full owner's PKI and a LDevID for the device. Subsequent reporting steps provide flows of information to indicate success/failure of the process.

## 2.5. Architectural Components

### 2.5.1. Pledge

The pledge is the device that is attempting to join. The pledge is assumed to talk to the Join Proxy using link-local network connectivity. In most cases, the pledge has no other connectivity until the pledge completes the enrollment process and receives some kind of network credential.

### 2.5.2. Join Proxy

The join proxy provides HTTPS connectivity between the pledge and the registrar. A circuit proxy mechanism is described in Section 4. Additional mechanisms, including a CoAP mechanism and a stateless IPIP mechanism are the subject of future work.

### 2.5.3. Domain Registrar

The domain's registrar operates as the BRSKI-MASA client when requesting vouchers from the MASA (see Section 5.4). The registrar operates as the BRSKI-EST server when pledges request vouchers (see Section 5.1). The registrar operates as the BRSKI-EST server "Registration Authority" if the pledge requests an end entity certificate over the BRSKI-EST connection (see Section 5.9).

The registrar uses an Implicit Trust Anchor database for authenticating the BRSKI-MASA connection's MASA TLS Server Certificate. Configuration or distribution of trust anchors is out-of-scope for this specification.

The registrar uses a different Implicit Trust Anchor database for authenticating the BRSKI-EST connection's Pledge TLS Client Certificate. Configuration or distribution of the BRSKI-EST client trust anchors is out-of-scope of this specification. Note that the trust anchors in/excluded from the database will affect which manufacturers' devices are acceptable to the registrar as pledges, and can also be used to limit the set of MASAs that are trusted for enrollment.

### 2.5.4. Manufacturer Service

The Manufacturer Service provides two logically separate functions: the Manufacturer Authorized Signing Authority (MASA) described in Section 5.5 and Section 5.6, and an ownership tracking/auditing function described in Section 5.7 and Section 5.8.

### 2.5.5. Public Key Infrastructure (PKI)

The Public Key Infrastructure (PKI) administers certificates for the domain of concern, providing the trust anchor(s) for it and allowing enrollment of pledges with domain certificates.

The voucher provides a method for the distribution of a single PKI trust anchor (as the "pinned-domain-cert"). A distribution of the full set of current trust anchors is possible using the optional EST integration.

The domain's registrar acts as an [RFC5272] Registration Authority, requesting certificates for pledges from the Key Infrastructure.

The expectations of the PKI are unchanged from EST [RFC7030]. This document does not place any additional architectural requirements on the Public Key Infrastructure.

## 2.6. Certificate Time Validation

### 2.6.1. Lack of realtime clock

Many devices when bootstrapping do not have knowledge of the current time. Mechanisms such as Network Time Protocols cannot be secured until bootstrapping is complete. Therefore bootstrapping is defined with a framework that does not require knowledge of the current time. A pledge MAY ignore all time stamps in the voucher and in the certificate validity periods if it does not know the current time.

The pledge is exposed to dates in the following five places: registrar certificate notBefore, registrar certificate notAfter, voucher created-on, and voucher expires-on. Additionally, CMS signatures contain a signingTime.

A pledge with a real time clock in which it has confidence, MUST check the above time fields in all certificates and signatures that it processes.

If the voucher contains a nonce then the pledge MUST confirm the nonce matches the original pledge voucher-request. This ensures the voucher is fresh. See Section 5.2.

### 2.6.2. Infinite Lifetime of IDevID

[RFC5280] explains that long lived pledge certificates "SHOULD be assigned the GeneralizedTime value of 99991231235959Z" for the notAfter field.

Some deployed IDevID management systems are not compliant with the 802.1AR requirement for infinite lifetimes, and put in typical  $\leq 3$  year certificate lifetimes. Registrars SHOULD be configurable on a per-manufacturer basis to ignore pledge lifetimes when the pledge did not follow the RFC5280 recommendations.

## 2.7. Cloud Registrar

There exist operationally open networks wherein devices gain unauthenticated access to the Internet at large. In these use cases the management domain for the device needs to be discovered within the larger Internet. The case where a device can boot and get access to larger Internet are less likely within the ANIMA ACP scope but may be more important in the future. In the ANIMA ACP scope, new devices will be quarantined behind a Join Proxy.

There are additionally some greenfield situations involving an entirely new installation where a device may have some kind of management uplink that it can use (such as via 3G network for instance). In such a future situation, the device might use this management interface to learn that it should configure itself to become the local registrar.

In order to support these scenarios, the pledge MAY contact a well known URI of a cloud registrar if a local registrar cannot be discovered or if the pledge's target use cases do not include a local registrar.

If the pledge uses a well known URI for contacting a cloud registrar a manufacturer-assigned Implicit Trust Anchor database (see [RFC7030]) MUST be used to authenticate that service as described in [RFC6125]. The use of a DNS-ID for validation is appropriate, and it may include wildcard components on the left-mode side. This is consistent with the human user configuration of an EST server URI in [RFC7030] which also depends on RFC6125.

## 2.8. Determining the MASA to contact

The registrar needs to be able to contact a MASA that is trusted by the pledge in order to obtain vouchers. There are three mechanisms described:

The device's Initial Device Identifier (IDevID) will normally contain the MASA URL as detailed in Section 2.3. This is the RECOMMENDED mechanism.

It can be operationally difficult to ensure the necessary X.509 extensions are in the pledge's IDevID due to the difficulty of aligning current pledge manufacturing with software releases and development. As a final fallback the registrar MAY be manually configured or distributed with a MASA URL for each manufacturer. Note that the registrar can only select the configured MASA URL based on the trust anchor -- so manufacturers can only leverage this approach if they ensure a single MASA URL works for all pledges associated with each trust anchor.

### 3. Voucher-Request artifact

Voucher-requests are how vouchers are requested. The semantics of the voucher-request are described below, in the YANG model.

A pledge forms the "pledge voucher-request", signs it with its IDevID and submits it to the registrar.

The registrar in turn forms the "registrar voucher-request", signs it with its Registrar keypair and submits it to the MASA.

The "proximity-registrar-cert" leaf is used in the pledge voucher-requests. This provides a method for the pledge to assert the registrar's proximity.

This network proximity results from the following properties in the ACP context: the pledge is connected to the Join Proxy (Section 4) using a Link-Local IPv6 connection. While the Join Proxy does not participate in any meaningful sense in the cryptography of the TLS connection (such as via a Channel Binding), the Registrar can observe that the connection is via the private ACP (ULA) address of the join proxy, and could not come from outside the ACP. The Pledge must therefore be at most one IPv6 Link-Local hop away from an existing node on the ACP.

Other users of BRSKI will need to define other kinds of assertions if the network proximity described above does not match their needs.

The "prior-signed-voucher-request" leaf is used in registrar voucher-requests. If present, it is the signed pledge voucher-request artifact. This provides a method for the registrar to forward the pledge's signed request to the MASA. This completes transmission of the signed "proximity-registrar-cert" leaf.

Unless otherwise signaled (outside the voucher-request artifact), the signing structure is as defined for vouchers, see [RFC8366].

### 3.1. Nonceless Voucher Requests

A registrar MAY also retrieve nonceless vouchers by sending nonceless voucher-requests to the MASA in order to obtain vouchers for use when the registrar does not have connectivity to the MASA. No "prior-signed-voucher-request" leaf would be included. The registrar will also need to know the serial number of the pledge. This document does not provide a mechanism for the registrar to learn that in an automated fashion. Typically this will be done via scanning of bar-code or QR-code on packaging, or via some sales channel integration.

### 3.2. Tree Diagram

The following tree diagram illustrates a high-level view of a voucher-request document. The voucher-request builds upon the voucher artifact described in [RFC8366]. The tree diagram is described in [RFC8340]. Each node in the diagram is fully described by the YANG module in Section 3.4. Please review the YANG module for a detailed description of the voucher-request format.

module: ietf-voucher-request

```

grouping voucher-request-grouping
  +-- voucher
    +-- created-on?          yang:date-and-time
    +-- expires-on?        yang:date-and-time
    +-- assertion?         enumeration
    +-- serial-number       string
    +-- idevid-issuer?     binary
    +-- pinned-domain-cert? binary
    +-- domain-cert-revocation-checks? boolean
    +-- nonce?             binary
    +-- last-renewal-date? yang:date-and-time
    +-- prior-signed-voucher-request? binary
    +-- proximity-registrar-cert? binary

```

Figure 5: YANG Tree diagram for Voucher-Request

### 3.3. Examples

This section provides voucher-request examples for illustration purposes. These examples show the JSON prior to CMS wrapping. JSON encoding rules specify that any binary content be base64 encoded ([RFC4648] section 4). The contents of the (base64) encoded certificates have been elided to save space. For detailed examples, see Appendix C.2. These examples conform to the encoding rules defined in [RFC7951].

Example (1) The following example illustrates a pledge voucher-request. The assertion leaf is indicated as 'proximity' and the registrar's TLS server certificate is included in the 'proximity-registrar-cert' leaf. See Section 5.2.

```
{
  "ietf-voucher-request:voucher": {
    "assertion": "proximity",
    "nonce": "62a2e7693d82fcda2624de58fb6722e5",
    "serial-number" : "JADA123456789",
    "created-on": "2017-01-01T00:00:00.000Z",
    "proximity-registrar-cert": "base64encodedvalue=="
  }
}
```

Figure 6: JSON representation of example Voucher-Request

Example (2) The following example illustrates a registrar voucher-request. The 'prior-signed-voucher-request' leaf is populated with the pledge's voucher-request (such as the prior example). The pledge's voucher-request is a binary CMS signed object. In the JSON encoding used here it must be base64 encoded. The nonce and assertion have been carried forward from the pledge request to the registrar request. The serial-number is extracted from the pledge's Client Certificate from the TLS connection. See Section 5.5.

```
{
  "ietf-voucher-request:voucher": {
    "assertion" : "proximity",
    "nonce": "62a2e7693d82fcda2624de58fb6722e5",
    "created-on": "2017-01-01T00:00:02.000Z",
    "idevid-issuer": "base64encodedvalue==",
    "serial-number": "JADA123456789",
    "prior-signed-voucher-request": "base64encodedvalue=="
  }
}
```

Figure 7: JSON representation of example Prior-Signed Voucher-Request

Example (3) The following example illustrates a registrar voucher-request. The 'prior-signed-voucher-request' leaf is not populated with the pledge's voucher-request nor is the nonce leaf. This form might be used by a registrar requesting a voucher when the pledge can not communicate with the registrar (such as when it is powered down, or

still in packaging), and therefore could not submit a nonce. This scenario is most useful when the registrar is aware that it will not be able to reach the MASA during deployment. See Section 5.5.

```
{
  "ietf-voucher-request:voucher": {
    "created-on": "2017-01-01T00:00:02.000Z",
    "idevid-issuer": "base64encodedvalue==",
    "serial-number": "JADA123456789"
  }
}
```

Figure 8: JSON representation of Offline Voucher-Request

### 3.4. YANG Module

Following is a YANG [RFC7950] module formally extending the [RFC8366] voucher into a voucher-request.

```
<CODE BEGINS> file "ietf-voucher-request@2018-02-14.yang"
module ietf-voucher-request {
  yang-version 1.1;

  namespace
    "urn:ietf:params:xml:ns:yang:ietf-voucher-request";
  prefix "vcr";

  import ietf-restconf {
    prefix rc;
    description "This import statement is only present to access
      the yang-data extension defined in RFC 8040.";
    reference "RFC 8040: RESTCONF Protocol";
  }

  import ietf-voucher {
    prefix vch;
    description "This module defines the format for a voucher,
      which is produced by a pledge's manufacturer or
      delegate (MASA) to securely assign a pledge to
      an 'owner', so that the pledge may establish a secure
      connection to the owner's network infrastructure";

    reference "RFC 8366: Voucher Artifact for
      Bootstrapping Protocols";
  }

  organization
```

"IETF ANIMA Working Group";

contact

"WG Web: <<https://datatracker.ietf.org/wg/anima/>>  
WG List: <<mailto:anima@ietf.org>>  
Author: Kent Watsen  
<<mailto:kent+ietf@watsen.net>>  
Author: Michael H. Behringer  
<<mailto:Michael.H.Behringer@gmail.com>>  
Author: Toerless Eckert  
<<mailto:tte+ietf@cs.fau.de>>  
Author: Max Pritikin  
<<mailto:pritikin@cisco.com>>  
Author: Michael Richardson  
<<mailto:mcr+ietf@sandelman.ca>>";

description

"This module defines the format for a voucher request.  
It is a superset of the voucher itself.  
It provides content to the MASA for consideration  
during a voucher request.

The key words 'MUST', 'MUST NOT', 'REQUIRED', 'SHALL', 'SHALL NOT', 'SHOULD', 'SHOULD NOT', 'RECOMMENDED', 'NOT RECOMMENDED', 'MAY', and 'OPTIONAL' in this document are to be interpreted as described in BCP 14 (RFC 2119) (RFC 8174) when, and only when, they appear in all capitals, as shown here.

Copyright (c) 2019 IETF Trust and the persons identified as authors of the code. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, is permitted pursuant to, and subject to the license terms contained in, the Simplified BSD License set forth in Section 4.c of the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>).

This version of this YANG module is part of RFC XXXX; see the RFC itself for full legal notices.";

```
revision "2018-02-14" {  
  description  
    "Initial version";  
  reference  
    "RFC XXXX: Bootstrapping Remote Secure Key Infrastructure";  
}
```

```
// Top-level statement
rc:yang-data voucher-request-artifact {
  uses voucher-request-grouping;
}

// Grouping defined for future usage
grouping voucher-request-grouping {
  description
    "Grouping to allow reuse/extensions in future work.";

  uses vch:voucher-artifact-grouping {
    refine "voucher/created-on" {
      mandatory false;
    }

    refine "voucher/pinned-domain-cert" {
      mandatory false;
      description "A pinned-domain-cert field
        is not valid in a voucher request, and
        any occurrence MUST be ignored";
    }

    refine "voucher/last-renewal-date" {
      description "A last-renewal-date field
        is not valid in a voucher request, and
        any occurrence MUST be ignored";
    }

    refine "voucher/domain-cert-revocation-checks" {
      description "The domain-cert-revocation-checks field
        is not valid in a voucher request, and
        any occurrence MUST be ignored";
    }

    refine "voucher/assertion" {
      mandatory false;
      description "Any assertion included in registrar voucher
        requests SHOULD be ignored by the MASA.";
    }

    augment "voucher" {
      description
        "Adds leaf nodes appropriate for requesting vouchers.";

      leaf prior-signed-voucher-request {
        type binary;
        description
          "If it is necessary to change a voucher, or re-sign and
```



#### 4. Proxying details (Pledge - Proxy - Registrar)

This section is normative for uses with an ANIMA ACP. The use of the GRASP mechanism is part of the ACP. Other users of BRSKI will need to define an equivalent proxy mechanism, and an equivalent mechanism to configure the proxy.

The role of the proxy is to facilitate communications. The proxy forwards packets between the pledge and a registrar that has been provisioned to the proxy via full GRASP ACP discovery.

This section defines a stateful proxy mechanism which is referred to as a "circuit" proxy. This is a form of Application Level Gateway ([RFC2663] section 2.9).

The proxy does not terminate the TLS handshake: it passes streams of bytes onward without examination. A proxy MUST NOT assume any specific TLS version. Please see [RFC8446] section 9.3 for details on TLS invariants.

A Registrar can directly provide the proxy announcements described below, in which case the announced port can point directly to the Registrar itself. In this scenario the pledge is unaware that there is no proxying occurring. This is useful for Registrars which are servicing pledges on directly connected networks.

As a result of the proxy Discovery process in Section 4.1.1, the port number exposed by the proxy does not need to be well known, or require an IANA allocation.

During the discovery of the Registrar by the Join Proxy, the Join Proxy will also learn which kinds of proxy mechanisms are available. This will allow the Join Proxy to use the lowest impact mechanism which the Join Proxy and Registrar have in common.

In order to permit the proxy functionality to be implemented on the maximum variety of devices the chosen mechanism should use the minimum amount of state on the proxy device. While many devices in the ANIMA target space will be rather large routers, the proxy function is likely to be implemented in the control plane CPU of such a device, with available capabilities for the proxy function similar to many class 2 IoT devices.

The document [I-D.richardson-anima-state-for-joinrouter] provides a more extensive analysis and background of the alternative proxy methods.

#### 4.1. Pledge discovery of Proxy

The result of discovery is a logical communication with a registrar, through a proxy. The proxy is transparent to the pledge. The communication between the pledge and Join Proxy is over IPv6 Link-Local addresses.

To discover the proxy the pledge performs the following actions:

1. MUST: Obtains a local address using IPv6 methods as described in [RFC4862] IPv6 Stateless Address AutoConfiguration. Use of [RFC4941] temporary addresses is encouraged. To limit pervasive monitoring ( [RFC7258]), a new temporary address MAY use a short lifetime (that is, set TEMP\_PREFERRED\_LIFETIME to be short). Pledges will generally prefer use of IPv6 Link-Local addresses, and discovery of proxy will be by Link-Local mechanisms. IPv4 methods are described in Appendix A
2. MUST: Listen for GRASP M\_FLOOD ([I-D.ietf-anima-grasp]) announcements of the objective: "AN\_Proxy". See section Section 4.1.1 for the details of the objective. The pledge MAY listen concurrently for other sources of information, see Appendix B.

Once a proxy is discovered the pledge communicates with a registrar through the proxy using the bootstrapping protocol defined in Section 5.

While the GRASP M\_FLOOD mechanism is passive for the pledge, the non-normative other methods (mDNS, and IPv4 methods) described in Appendix B are active. The pledge SHOULD run those methods in parallel with listening to for the M\_FLOOD. The active methods SHOULD back-off by doubling to a maximum of one hour to avoid overloading the network with discovery attempts. Detection of change of physical link status (Ethernet carrier for instance) SHOULD reset the back off timers.

The pledge could discover more than one proxy on a given physical interface. The pledge can have a multitude of physical interfaces as well: a layer-2/3 Ethernet switch may have hundreds of physical ports.

Each possible proxy offer SHOULD be attempted up to the point where a valid voucher is received: while there are many ways in which the attempt may fail, it does not succeed until the voucher has been validated.

The connection attempts via a single proxy SHOULD exponentially back-off to a maximum of one hour to avoid overloading the network infrastructure. The back-off timer for each MUST be independent of other connection attempts.

Connection attempts SHOULD be run in parallel to avoid head of queue problems wherein an attacker running a fake proxy or registrar could perform protocol actions intentionally slowly. Connection attempts to different proxies SHOULD be sent with an interval of 3 to 5s. The pledge SHOULD continue to listen to for additional GRASP M\_FLOOD messages during the connection attempts.

Each connection attempt through a distinct Join Proxy MUST have a unique nonce in the voucher-request.

Once a connection to a registrar is established (e.g. establishment of a TLS session key) there are expectations of more timely responses, see Section 5.2.

Once all discovered services are attempted (assuming that none succeeded) the device MUST return to listening for GRASP M\_FLOOD. It SHOULD periodically retry any manufacturer-specific mechanisms. The pledge MAY prioritize selection order as appropriate for the anticipated environment.

#### 4.1.1. Proxy GRASP announcements

A proxy uses the DULL GRASP M\_FLOOD mechanism to announce itself. This announcement can be within the same message as the ACP announcement detailed in [I-D.ietf-anima-autonomic-control-plane].

The formal Concise Data Definition Language (CDDL) [RFC8610] definition is:

```

<CODE BEGINS> file "proxygrasp.cddl"
flood-message = [M_FLOOD, session-id, initiator, ttl,
                 +[objective, (locator-option / [])]]

objective = ["AN_Proxy", objective-flags, loop-count,
            objective-value]

ttl          = 180000          ; 180,000 ms (3 minutes)
initiator    = ACP address to contact Registrar
objective-flags = sync-only   ; as in GRASP spec
sync-only    = 4              ; M_FLOOD only requires synchronization
loop-count   = 1              ; one hop only
objective-value = any         ; none

locator-option = [ O_IPv6_LOCATOR, ipv6-address,
                  transport-proto, port-number ]
ipv6-address   = the v6 LL of the Proxy
$transport-proto /= IPPROTO_TCP ; note this can be any value from the
                                ; IANA protocol registry, as per
                                ; [GRASP] section 2.9.5.1, note 3.
port-number    = selected by Proxy
<CODE ENDS>

```

Figure 10: CDDL definition of Proxy Discovery message

Here is an example M\_FLOOD announcing a proxy at fe80::1, on TCP port 4443.

```

[M_FLOOD, 12340815, h'fe800000000000000000000000000001', 180000,
  ["AN_Proxy", 4, 1, ""],
  [O_IPv6_LOCATOR,
   h'fe800000000000000000000000000001', IPPROTO_TCP, 4443]]

```

Figure 11: Example of Proxy Discovery message

On a small network the Registrar MAY include the GRASP M\_FLOOD announcements to locally connected networks.

The \$transport-proto above indicates the method that the pledge-proxy-registrar will use. The TCP method described here is mandatory, and other proxy methods, such as CoAP methods not defined in this document are optional. Other methods MUST NOT be enabled unless the Join Registrar ASA indicates support for them in its own announcement.

#### 4.2. CoAP connection to Registrar

The use of CoAP to connect from pledge to registrar is out of scope for this document, and is described in future work. See [I-D.ietf-anima-constrained-voucher].

#### 4.3. Proxy discovery and communication of Registrar

The registrar SHOULD announce itself so that proxies can find it and determine what kind of connections can be terminated.

The registrar announces itself using ACP instance of GRASP using M\_FLOOD messages, with the "AN\_join\_registrar" objective. A registrar may announce any convenient port number, including using a stock port 443. ANI proxies MUST support GRASP discovery of registrars.

The M\_FLOOD is formatted as follows:

```
[M_FLOOD, 51804321, h'fda379a6f6ee00000200000064000001', 180000,
  [{"AN_join_registrar", 4, 255, "EST-TLS"},
  [O_IPv6_LOCATOR,
   h'fda379a6f6ee00000200000064000001', IPPROTO_TCP, 8443]]]
```

Figure 12: An example of a Registrar announcement message

The formal CDDL definition is:

```
<CODE BEGINS> file "jrcgrasp.cddl"
flood-message = [M_FLOOD, session-id, initiator, ttl,
  +[objective, (locator-option / [])]]

objective = ["AN_join_registrar", objective-flags, loop-count,
  objective-value]

initiator = ACP address to contact Registrar
objective-flags = sync-only ; as in GRASP spec
sync-only = 4 ; M_FLOOD only requires synchronization
loop-count = 255 ; mandatory maximum
objective-value = text ; name of the (list of) of supported
; protocols: "EST-TLS" for RFC7030.
<CODE ENDS>
```

Figure 13: CDDL definition for Registrar announcement message

The M\_FLOOD message MUST be sent periodically. The default period SHOULD be 60 seconds, the value SHOULD be operator configurable but SHOULD NOT be smaller than 60 seconds. The frequency of sending MUST be such that the aggregate amount of periodic M\_FLOODs from all flooding sources cause only negligible traffic across the ACP.

Here are some examples of locators for illustrative purposes. Only the first one (\$transport-protocol = 6, TCP) is defined in this document and is mandatory to implement.

```
locator1 = [O_IPv6_LOCATOR, fd45:1345::6789, 6, 443]
locator2 = [O_IPv6_LOCATOR, fd45:1345::6789, 17, 5683]
locator3 = [O_IPv6_LOCATOR, fe80::1234, 41, nil]
```

A protocol of 6 indicates that TCP proxying on the indicated port is desired.

Registrars MUST announce the set of protocols that they support. They MUST support TCP traffic.

Registrars MUST accept HTTPS/EST traffic on the TCP ports indicated.

Registrars MUST support ANI TLS circuit proxy and therefore BRSKI across HTTPS/TLS native across the ACP.

In the ANI, the Autonomic Control Plane (ACP) secured instance of GRASP ([I-D.ietf-anima-grasp]) MUST be used for discovery of ANI registrar ACP addresses and ports by ANI proxies. The TCP leg of the proxy connection between ANI proxy and ANI registrar therefore also runs across the ACP.

## 5. Protocol Details (Pledge - Registrar - MASA)

The pledge MUST initiate BRSKI after boot if it is unconfigured. The pledge MUST NOT automatically initiate BRSKI if it has been configured or is in the process of being configured.

BRSKI is described as extensions to EST [RFC7030]. The goal of these extensions is to reduce the number of TLS connections and crypto operations required on the pledge. The registrar implements the BRSKI REST interface within the `"/.well-known/brski"` URI tree, as well as implementing the existing EST URIs as described in EST [RFC7030] section 3.2.2. The communication channel between the pledge and the registrar is referred to as "BRSKI-EST" (see Figure 1).

The communication channel between the registrar and MASA is a new communication channel, similar to EST, within the newly registered `"/.well-known/brski"` tree. For clarity this channel is referred to as "BRSKI-MASA". (See Figure 1).

The MASA URI is `"https://"` authority `"/.well-known/brski"`.

BRSKI uses existing CMS message formats for existing EST operations. BRSKI uses JSON [RFC8259] for all new operations defined here, and voucher formats. In all places where a binary value must be carried in a JSON string, the use of base64 format ([RFC4648] section 4) is to be used, as per [RFC7951] section 6.6.

While EST section 3.2 does not insist upon use of HTTP persistent connections ([RFC7230] section 6.3), BRSKI-EST connections SHOULD use persistent connections. The intention of this guidance is to ensure the provisional TLS state occurs only once, and that the subsequent resolution of the provision state is not subject to a MITM attack during a critical phase.

If non-persistent connections are used, then both the pledge and the registrar MUST remember the certificates seen, and also sent for the first connection. They MUST check each subsequent connections for the same certificates, and each end MUST use the same certificates as well. This places a difficult restriction on rolling certificates on the Registrar.

Summarized automation extensions for the BRSKI-EST flow are:

- \* The pledge either attempts concurrent connections via each discovered proxy, or it times out quickly and tries connections in series, as explained at the end of Section 5.1.
- \* The pledge provisionally accepts the registrar certificate during the TLS handshake as detailed in Section 5.1.
- \* The pledge requests a voucher using the new REST calls described below. This voucher is then validated.
- \* The pledge completes authentication of the server certificate as detailed in Section 5.6.1. This moves the BRSKI-EST TLS connection out of the provisional state.
- \* Mandatory bootstrap steps conclude with voucher status telemetry (see Section 5.7).

The BRSKI-EST TLS connection can now be used for EST enrollment.

The extensions for a registrar (equivalent to EST server) are:

- \* Client authentication is automated using Initial Device Identity (IDevID) as per the EST certificate based client authentication. The subject field's DN encoding MUST include the "serialNumber" attribute with the device's unique serial number as explained in Section 2.3.1
- \* The registrar requests and validates the voucher from the MASA.
- \* The registrar forwards the voucher to the pledge when requested.
- \* The registrar performs log verifications (described in Section 5.8.3) in addition to local authorization checks before accepting optional pledge device enrollment requests.

#### 5.1. BRSKI-EST TLS establishment details

The pledge establishes the TLS connection with the registrar through the circuit proxy (see Section 4) but the TLS handshake is with the registrar. The BRSKI-EST pledge is the TLS client and the BRSKI-EST registrar is the TLS server. All security associations established are between the pledge and the registrar regardless of proxy operations.

Use of TLS 1.3 (or newer) is encouraged. TLS 1.2 or newer is REQUIRED on the Pledge side. TLS 1.3 (or newer) SHOULD be available on the Registrar server interface, and the Registrar client interface, but TLS 1.2 MAY be used. TLS 1.3 (or newer) SHOULD be available on the MASA server interface, but TLS 1.2 MAY be used.

Establishment of the BRSKI-EST TLS connection is as specified in EST [RFC7030] section 4.1.1 "Bootstrap Distribution of CA Certificates" [RFC7030] wherein the client is authenticated with the IDevID certificate, and the EST server (the registrar) is provisionally authenticated with an unverified server certificate. Configuration or distribution of the trust anchor database used for validating the IDevID certificate is out-of-scope of this specification. Note that the trust anchors in/excluded from the database will affect which manufacturers' devices are acceptable to the registrar as pledges, and can also be used to limit the set of MASAs that are trusted for enrollment.

The signature in the certificate MUST be validated even if a signing key can not (yet) be validated. The certificate (or chain) MUST be retained for later validation.

A self-signed certificate for the Registrar is acceptable as the voucher can validate it upon successful enrollment.

The pledge performs input validation of all data received until a voucher is verified as specified in Section 5.6.1 and the TLS connection leaves the provisional state. Until these operations are complete the pledge could be communicating with an attacker.

The pledge code needs to be written with the assumption that all data is being transmitted at this point to an unauthenticated peer, and that received data, while inside a TLS connection, **MUST** be considered untrusted. This particularly applies to HTTP headers and CMS structures that make up the voucher.

A pledge that can connect to multiple Registrars concurrently **SHOULD** do so. Some devices may be unable to do so for lack of threading, or resource issues. Concurrent connections defeat attempts by a malicious proxy from causing a TCP Slowloris-like attack (see [slowloris]).

A pledge that can not maintain as many connections as there are eligible proxies will need to rotate among the various choices, terminating connections that do not appear to be making progress. If no connection is making progress after 5 seconds then the pledge **SHOULD** drop the oldest connection and go on to a different proxy: the proxy that has been communicated with least recently. If there were no other proxies discovered, the pledge **MAY** continue to wait, as long as it is concurrently listening for new proxy announcements.

## 5.2. Pledge Requests Voucher from the Registrar

When the pledge bootstraps it makes a request for a voucher from a registrar.

This is done with an HTTPS POST using the operation path value of `"/.well-known/brski/requestvoucher"`.

The pledge voucher-request Content-Type is:

`application/voucher-cms+json` [RFC8366] defines a "YANG-defined JSON document that has been signed using a CMS structure", and the voucher-request described in Section 3 is created in the same way. The media type is the same as defined in [RFC8366]. This is also used for the pledge voucher-request. The pledge **MUST** sign the request using the Section 2.3 credential.

Registrar implementations **SHOULD** anticipate future media types but of course will simply fail the request if those types are not yet known.

The pledge SHOULD include an [RFC7231] section 5.3.2 "Accept" header field indicating the acceptable media type for the voucher response. The "application/voucher-cms+json" media type is defined in [RFC8366] but constrained voucher formats are expected in the future. Registrars and MASA are expected to be flexible in what they accept.

The pledge populates the voucher-request fields as follows:

**created-on:** Pledges that have a realtime clock are RECOMMENDED to populate this field with the current date and time in yang:date-and-time format. This provides additional information to the MASA. Pledges that have no real-time clocks MAY omit this field.

**nonce:** The pledge voucher-request MUST contain a cryptographically strong random or pseudo-random number nonce (see [RFC4086] section 6.2). As the nonce is usually generated very early in the boot sequence there is a concern that the same nonce might be generated across multiple boots, or after a factory reset. Different nonces MUST be generated for each bootstrapping attempt, whether in series or concurrently. The freshness of this nonce mitigates against the lack of real-time clock as explained in Section 2.6.1.

**assertion:** The pledge indicates support for the mechanism described in this document, by putting the value "proximity" in the voucher-request, MUST include the "proximity-registrar-cert" field (below).

**proximity-registrar-cert:** In a pledge voucher-request this is the first certificate in the TLS server 'certificate\_list' sequence (see [RFC5246]) presented by the registrar to the pledge. That is, it is the end-entity certificate. This MUST be populated in a pledge voucher-request.

**serial-number** The serial number of the pledge is included in the voucher-request from the Pledge. This value is included as a sanity check only, but it is not to be forwarded by the Registrar as described in Section 5.5.

All other fields MAY be omitted in the pledge voucher-request.

An example JSON payload of a pledge voucher-request is in Section 3.3 Example 1.

The registrar confirms that the assertion is 'proximity' and that pinned 'proximity-registrar-cert' is the Registrar's certificate. If this validation fails, then there is an On-Path Attacker (MITM), and the connection MUST be closed after the returning an HTTP 401 error code.

### 5.3. Registrar Authorization of Pledge

In a fully automated network all devices must be securely identified and authorized to join the domain.

A Registrar accepts or declines a request to join the domain, based on the authenticated identity presented. For different networks, examples of automated acceptance may include:

- \* allow any device of a specific type (as determined by the X.509 IDevID),
- \* allow any device from a specific vendor (as determined by the X.509 IDevID),
- \* allow a specific device from a vendor (as determined by the X.509 IDevID) against a domain white list. (The mechanism for checking a shared white list potentially used by multiple Registrars is out of scope).

If validation fails the registrar SHOULD respond with the HTTP 404 error code. If the voucher-request is in an unknown format, then an HTTP 406 error code is more appropriate. A situation that could be resolved with administrative action (such as adding a vendor to a whitelist) MAY be responded with an 403 HTTP error code.

If authorization is successful the registrar obtains a voucher from the MASA service (see Section 5.5) and returns that MASA signed voucher to the pledge as described in Section 5.6.

### 5.4. BRSKI-MASA TLS establishment details

The BRSKI-MASA TLS connection is a 'normal' TLS connection appropriate for HTTPS REST interfaces. The registrar initiates the connection and uses the MASA URL obtained as described in Section 2.8. The mechanisms in [RFC6125] SHOULD be used in authentication of the MASA using a DNS-ID that matches that which is found in the IDevID. Registrars MAY include a mechanism to override the MASA URL on a manufacturer-by-manufacturer basis, and within that override it is appropriate to provide alternate anchors. This will typically be used by some vendors to establish explicit (or private) trust anchors for validating their MASA that is part of a sales channel integration.

Use of TLS 1.3 (or newer) is encouraged. TLS 1.2 or newer is REQUIRED. TLS 1.3 (or newer) SHOULD be available.

As described in [RFC7030], the MASA and the registrars SHOULD be prepared to support TLS client certificate authentication and/or HTTP Basic, Digest, or SCRAM authentication. This connection MAY also have no client authentication at all.

Registrars SHOULD permit trust anchors to be pre-configured on a per-vendor(MASA) basis. Registrars SHOULD include the ability to configure a TLS ClientCertificate on a per-MASA basis, or to use no client certificate. Registrars SHOULD also permit HTTP Basic and Digest authentication to be configured.

The authentication of the BRSKI-MASA connection does not change the voucher-request process, as voucher-requests are already signed by the registrar. Instead, this authentication provides access control to the audit-log as described in Section 5.8.

Implementors are advised that contacting the MASA is to establish a secured API connection with a web service and that there are a number of authentication models being explored within the industry. Registrars are RECOMMENDED to fail gracefully and generate useful administrative notifications or logs in the advent of unexpected HTTP 401 (Unauthorized) responses from the MASA.

#### 5.4.1. MASA authentication of customer Registrar

Providing per-customer options requires that the customer's registrar be uniquely identified. This can be done by any stateless method that HTTPS supports such as with HTTP Basic or Digest authentication (that is using a password), but the use of TLS Client Certificate authentication is RECOMMENDED.

Stateful methods involving API tokens, or HTTP Cookies, are not recommended.

It is expected that the setup and configuration of per-customer Client Certificates is done as part of a sales ordering process.

The use of public PKI (i.e. WebPKI) End-Entity Certificates to identify the Registrar is reasonable, and if done universally this would permit a MASA to identify a customers' Registrar simply by a FQDN.

The use of DANE records in DNSSEC signed zones would also permit use of a FQDN to identify customer Registrars.

A third (and simplest, but least flexible) mechanism would be for the MASA to simply store the Registrar's certificate pinned in a database.

A MASA without any supply chain integration can simply accept Registrars without any authentication, or can accept them on a blind Trust-on-First-Use basis as described in Section 7.4.2.

This document does not make a specific recommendation on how the MASA authenticates the Registrar as there are likely different tradeoffs in different environments and product values. Even within the ANIMA ACP applicability, there is a significant difference between supply chain logistics for \$100 CPE devices and \$100,000 core routers.

#### 5.5. Registrar Requests Voucher from MASA

When a registrar receives a pledge voucher-request it in turn submits a registrar voucher-request to the MASA service via an HTTPS interface ([RFC7231]).

This is done with an HTTP POST using the operation path value of `"/.well-known/brski/requestvoucher"`.

The voucher media type `"application/voucher-cms+json"` is defined in [RFC8366] and is also used for the registrar voucher-request. It is a JSON document that has been signed using a CMS structure. The registrar MUST sign the registrar voucher-request.

MASA implementations SHOULD anticipate future media ntypes but of course will simply fail the request if those types are not yet known.

The voucher-request CMS object includes some number of certificates that are input to the MASA as it populates the `'pinned-domain-cert'`. As the [RFC8366] is quite flexible in what may be put into the `'pinned-domain-cert'`, the MASA needs some signal as to what certificate would be effective to populate the field with: it may range from the End Entity (EE) Certificate that the Registrar uses, to the entire private Enterprise CA certificate. More specific certificates result in a tighter binding of the voucher to the domain, while less specific certificates result in more flexibility in how the domain is represented by certificates.

A Registrar which is seeking a nonceless voucher for later offline use benefits from a less specific certificate, as it permits the actual keypair used by a future Registrar to be determined by the pinned certificate authority.

In some cases, a less specific certificate, such a public WebPKI certificate authority, could be too open, and could permit any entity issued a certificate by that authority to assume ownership of a device that has a voucher pinned. Future work may provide a solution to pin both a certificate and a name that would reduce such risk of malicious ownership assertions.

The Registrar SHOULD request a voucher with the most specificity consistent with the mode that it is operating in. In order to do this, when the Registrar prepares the CMS structure for the signed voucher-request, it SHOULD include only certificates which are part of the chain that it wishes the MASA to pin. This MAY be as small as only the End-Entity certificate (with id-kp-cmcRA set) that it uses as it's TLS Server Certificate, or it MAY be the entire chain, including the Domain CA.

The Registrar SHOULD include an [RFC7231] section 5.3.2 "Accept" header field indicating the response media types that are acceptable. This list SHOULD be the entire list presented to the Registrar in the Pledge's original request (see Section 5.2) but MAY be a subset. The MASA is expected to be flexible in what it accepts.

The registrar populates the voucher-request fields as follows:

created-on: The Registrars SHOULD populate this field with the current date and time when the Registrar formed this voucher request. This field provides additional information to the MASA.

nonce: This value, if present, is copied from the pledge voucher-request. The registrar voucher-request MAY omit the nonce as per Section 3.1.

serial-number: The serial number of the pledge the registrar would like a voucher for. The registrar determines this value by parsing the authenticated pledge IDevID certificate. See Section 2.3. The registrar MUST verify that the serial number field it parsed matches the serial number field the pledge provided in its voucher-request. This provides a sanity check useful for detecting error conditions and logging. The registrar MUST NOT simply copy the serial number field from a pledge voucher request as that field is claimed but not certified.

idevid-issuer: The Issuer value from the pledge IDevID certificate is included to ensure unique interpretation of the serial-number. In the case of nonceless (offline) voucher-request, then an appropriate value needs to be configured from the same out-of-band source as the serial-number.

prior-signed-voucher-request: The signed pledge voucher-request SHOULD be included in the registrar voucher-request. The entire CMS signed structure is to be included, base64 encoded for transport in the JSON structure.

A nonceless registrar voucher-request MAY be submitted to the MASA. Doing so allows the registrar to request a voucher when the pledge is offline, or when the registrar anticipates not being able to connect to the MASA while the pledge is being deployed. Some use cases require the registrar to learn the appropriate IDevID SerialNumber field and appropriate 'Accept header field' values from the physical device labeling or from the sales channel (out-of-scope for this document).

All other fields MAY be omitted in the registrar voucher-request.

The "proximity-registrar-cert" field MUST NOT be present in the registrar voucher-request.

Example JSON payloads of registrar voucher-requests are in Section 3.3 Examples 2 through 4.

The MASA verifies that the registrar voucher-request is internally consistent but does not necessarily authenticate the registrar certificate since the registrar MAY be unknown to the MASA in advance. The MASA performs the actions and validation checks described in the following sub-sections before issuing a voucher.

#### 5.5.1. MASA renewal of expired vouchers

As described in [RFC8366] vouchers are normally short lived to avoid revocation issues. If the request is for a previous (expired) voucher using the same registrar (that is, a Registrar with the same Domain CA) then the request for a renewed voucher SHOULD be automatically authorized. The MASA has sufficient information to determine this by examining the request, the registrar authentication, and the existing audit-log. The issuance of a renewed voucher is logged as detailed in Section 5.6.

To inform the MASA that existing vouchers are not to be renewed one can update or revoke the registrar credentials used to authorize the request (see Section 5.5.4 and Section 5.5.3). More flexible methods will likely involve sales channel integration and authorizations (details are out-of-scope of this document).

### 5.5.2. MASA pinning of registrar

A certificate chain is extracted from the Registrar's signed CMS container. This chain may be as short as a single End-Entity Certificate, up to the entire registrar certificate chain, including the Domain CA certificate, as specified in Section 5.5.

If the domain's CA is unknown to the MASA, then it is to be considered a temporary trust anchor for the rest of the steps in this section. The intention is not to authenticate the message as having come from a fully validated origin, but to establish the consistency of the domain PKI.

The MASA MAY use the certificate farthest in the chain chain that it received from the Registrar from the end-entity, as determined by MASA policy. A MASA MAY have a local policy that it only pins the End-Entity certificate. This is consistent with [RFC8366]. Details of the policy will typically depend upon the degree of Supply Chain Integration, and the mechanism used by the Registrar to authenticate. Such a policy would also determine how the MASA will respond to a request for a nonceless voucher.

### 5.5.3. MASA checking of voucher request signature

As described in Section 5.5.2, the MASA has extracted Registrar's domain CA. This is used to validate the CMS signature ([RFC5652]) on the voucher-request.

Normal PKIX revocation checking is assumed during voucher-request signature validation. This CA certificate MAY have Certificate Revocation List distribution points, or Online Certificate Status Protocol (OCSP) information ([RFC6960]). If they are present, the MASA MUST be able to reach the relevant servers belonging to the Registrar's domain CA to perform the revocation checks.

The use of OCSP Stapling is preferred.

#### 5.5.4. MASA verification of domain registrar

The MASA MUST verify that the registrar voucher-request is signed by a registrar. This is confirmed by verifying that the id-kp-cmcRA extended key usage extension field (as detailed in EST RFC7030 section 3.6.1) exists in the certificate of the entity that signed the registrar voucher-request. This verification is only a consistency check that the unauthenticated domain CA intended the voucher-request signer to be a registrar. Performing this check provides value to the domain PKI by assuring the domain administrator that the MASA service will only respect claims from authorized Registration Authorities of the domain.

Even when a domain CA is authenticated to the MASA, and there is strong sales channel integration to understand who the legitimate owner is, the above id-kp-cmcRA check prevents arbitrary End-Entity certificates (such as an LDevID certificate) from having vouchers issued against them.

Other cases of inappropriate voucher issuance are detected by examination of the audit log.

If a nonceless voucher-request is submitted the MASA MUST authenticate the registrar as described in either EST [RFC7030] section 3.2.3, section 3.3.2, or by validating the registrar's certificate used to sign the registrar voucher-request using a configured trust anchor. Any of these methods reduce the risk of DDoS attacks and provide an authenticated identity as an input to sales channel integration and authorizations (details are out-of-scope of this document).

In the nonced case, validation of the Registrar's identity (via TLS Client Certificate or HTTP authentication) MAY be omitted if the device policy is to accept audit-only vouchers.

#### 5.5.5. MASA verification of pledge prior-signed-voucher-request

The MASA MAY verify that the registrar voucher-request includes the 'prior-signed-voucher-request' field. If so the prior-signed-voucher-request MUST include a 'proximity-registrar-cert' that is consistent with the certificate used to sign the registrar voucher-request. Additionally the voucher-request serial-number leaf MUST match the pledge serial-number that the MASA extracts from the signing certificate of the prior-signed-voucher-request. The consistency check described above is checking that the 'proximity-registrar-cert' SPKI fingerprint exists within the registrar voucher-request CMS signature's certificate chain. This is substantially the same as the pin validation described in in [RFC7469] section 2.6, paragraph three.

If these checks succeed the MASA updates the voucher and audit-log assertion leafs with the "proximity" assertion, as defined by [RFC8366] section 5.3.

#### 5.5.6. MASA nonce handling

The MASA does not verify the nonce itself. If the registrar voucher-request contains a nonce, and the prior-signed-voucher-request exists, then the MASA MUST verify that the nonce is consistent. (Recall from above that the voucher-request might not contain a nonce, see Section 5.5 and Section 5.5.4).

The MASA populates the audit-log with the nonce that was verified. If a nonceless voucher is issued, then the audit-log is to be populated with the JSON value "null".

#### 5.6. MASA and Registrar Voucher Response

The MASA voucher response to the registrar is forwarded without changes to the pledge; therefore this section applies to both the MASA and the registrar. The HTTP signaling described applies to both the MASA and registrar responses.

When a voucher request arrives at the registrar, if it has a cached response from the MASA for the corresponding registrar voucher-request, that cached response can be used according to local policy; otherwise the registrar constructs a new registrar voucher-request and sends it to the MASA.

Registrar evaluation of the voucher itself is purely for transparency and audit purposes to further inform log verification (see Section 5.8.3) and therefore a registrar could accept future voucher formats that are opaque to the registrar.

If the voucher-request is successful, the server (MASA responding to registrar or registrar responding to pledge) response MUST contain an HTTP 200 response code. The server MUST answer with a suitable 4xx or 5xx HTTP [RFC7230] error code when a problem occurs. In this case, the response data from the MASA MUST be a plaintext human-readable (UTF-8) error message containing explanatory information describing why the request was rejected.

The registrar MAY respond with an HTTP 202 ("the request has been accepted for processing, but the processing has not been completed") as described in EST [RFC7030] section 4.2.3 wherein the client "MUST wait at least the specified 'Retry-After' time before repeating the same request". (see [RFC7231] section 6.6.4) The pledge is RECOMMENDED to provide local feedback (blinking LED etc) during this wait cycle if mechanisms for this are available. To prevent an attacker registrar from significantly delaying bootstrapping the pledge MUST limit the 'Retry-After' time to 60 seconds. Ideally the pledge would keep track of the appropriate Retry-After header field values for any number of outstanding registrars but this would involve a state table on the pledge. Instead the pledge MAY ignore the exact Retry-After value in favor of a single hard coded value (a registrar that is unable to complete the transaction after the first 60 seconds has another chance a minute later). A pledge SHOULD only maintain a 202 retry-state for up to 4 days, which is longer than a long weekend, after which time the enrollment attempt fails and the pledge returns to discovery state.

A pledge that retries a request after receiving a 202 message MUST resend the same voucher-request. It MUST NOT sign a new voucher-request each time, and in particular, it MUST NOT change the nonce value.

In order to avoid infinite redirect loops, which a malicious registrar might do in order to keep the pledge from discovering the correct registrar, the pledge MUST NOT follow more than one redirection (3xx code) to another web origin. EST supports redirection but requires user input; this change allows the pledge to follow a single redirection without a user interaction.

A 403 (Forbidden) response is appropriate if the voucher-request is not signed correctly, stale, or if the pledge has another outstanding voucher that cannot be overridden.

A 404 (Not Found) response is appropriate when the request is for a device that is not known to the MASA.

A 406 (Not Acceptable) response is appropriate if a voucher of the desired type or using the desired algorithms (as indicated by the `Accept:` header fields, and algorithms used in the signature) cannot be issued such as because the MASA knows the pledge cannot process that type. The registrar SHOULD use this response if it determines the pledge is unacceptable due to inventory control, MASA audit-logs, or any other reason.

A 415 (Unsupported Media Type) response is appropriate for a request that has a `voucher-request` or `Accept:` value that is not understood.

The voucher response format is as indicated in the submitted `Accept` header fields or based on the MASA's prior understanding of proper format for this Pledge. Only the [RFC8366] `"application/voucher-cms+json"` media type is defined at this time. The syntactic details of vouchers are described in detail in [RFC8366]. Figure 14 shows a sample of the contents of a voucher.

```
{
  "ietf-voucher:voucher": {
    "nonce": "62a2e7693d82fcda2624de58fb6722e5",
    "assertion": "logged",
    "pinned-domain-cert": "base64encodedvalue==",
    "serial-number": "JADA123456789"
  }
}
```

Figure 14: An example voucher

The MASA populates the voucher fields as follows:

`nonce`: The nonce from the pledge if available. See Section 5.5.6.

`assertion`: The method used to verify the relationship between pledge and registrar. See Section 5.5.5.

`pinned-domain-cert`: A certificate. See Section 5.5.2. This figure is illustrative, for an example, see Appendix C.2 where an End Entity certificate is used.

`serial-number`: The serial-number as provided in the `voucher-request`. Also see Section 5.5.5.

`domain-cert-revocation-checks`: Set as appropriate for the pledge's

capabilities and as documented in [RFC8366]. The MASA MAY set this field to 'false' since setting it to 'true' would require that revocation information be available to the pledge and this document does not make normative requirements for [RFC6961] or equivalent integrations.

**expires-on:** This is set for nonceless vouchers. The MASA ensures the voucher lifetime is consistent with any revocation or pinned-domain-cert consistency checks the pledge might perform. See section Section 2.6.1. There are three times to consider: (a) a configured voucher lifetime in the MASA, (b) the expiry time for the registrar's certificate, (c) any certificate revocation information (CRL) lifetime. The expires-on field SHOULD be before the earliest of these three values. Typically (b) will be some significant time in the future, but (c) will typically be short (on the order of a week or less). The RECOMMENDED period for (a) is on the order of 20 minutes, so it will typically determine the lifespan of the resulting voucher. 20 minutes is sufficient time to reach the post-provisional state in the pledge, at which point there is an established trust relationship between pledge and registrar. The subsequent operations can take as long as required from that point onwards. The lifetime of the voucher has no impact on the lifespan of the ownership relationship.

Whenever a voucher is issued the MASA MUST update the audit-log sufficiently to generate the response as described in Section 5.8.1. The internal state requirements to maintain the audit-log are out-of-scope.

#### 5.6.1. Pledge voucher verification

The pledge MUST verify the voucher signature using the manufacturer-installed trust anchor(s) associated with the manufacturer's MASA (this is likely included in the pledge's firmware). Management of the manufacturer-installed trust anchor(s) is out-of-scope of this document; this protocol does not update these trust anchor(s).

The pledge MUST verify the serial-number field of the signed voucher matches the pledge's own serial-number.

The pledge MUST verify the nonce information in the voucher. If present, the nonce in the voucher must match the nonce the pledge submitted to the registrar; vouchers with no nonce can also be accepted (according to local policy, see Section 7.2)

The pledge MUST be prepared to parse and fail gracefully from a voucher response that does not contain a 'pinned-domain-cert' field. Such a thing indicates a failure to enroll in this domain, and the pledge MUST attempt joining with other available Join Proxy.

The pledge MUST be prepared to ignore additional fields that it does not recognize.

#### 5.6.2. Pledge authentication of provisional TLS connection

Following the process described in [RFC8366], the pledge should consider the public key from the pinned-domain-cert as the sole temporary trust anchor.

The pledge then evaluates the TLS Server Certificate chain that it received when the TLS connection was formed using this trust anchor. It is possible that the pinned-domain-cert matches the End-Entity Certificate provided in the TLS Server.

If a registrar's credentials cannot be verified using the pinned-domain-cert trust anchor from the voucher then the TLS connection is immediately discarded and the pledge abandons attempts to bootstrap with this discovered registrar. The pledge SHOULD send voucher status telemetry (described below) before closing the TLS connection. The pledge MUST attempt to enroll using any other proxies it has found. It SHOULD return to the same proxy again after unsuccessful attempts with other proxies. Attempts should be made repeated at intervals according to the backoff timer described earlier. Attempts SHOULD be repeated as failure may be the result of a temporary inconsistency (an inconsistently rolled registrar key, or some other mis-configuration). The inconsistency could also be the result an active MITM attack on the EST connection.

The registrar MUST use a certificate that chains to the pinned-domain-cert as its TLS server certificate.

The pledge's PKIX path validation of a registrar certificate's validity period information is as described in Section 2.6.1. Once the PKIX path validation is successful the TLS connection is no longer provisional.

The pinned-domain-cert MAY be installed as a trust anchor for future operations such as enrollment (e.g. [RFC7030] as recommended) or trust anchor management or raw protocols that do not need full PKI based key management. It can be used to authenticate any dynamically discovered EST server that contain the id-kp-cmcRA extended key usage extension as detailed in EST RFC7030 section 3.6.1; but to reduce system complexity the pledge SHOULD avoid additional discovery

operations. Instead the pledge SHOULD communicate directly with the registrar as the EST server. The 'pinned-domain-cert' is not a complete distribution of the [RFC7030] section 4.1.3 CA Certificate Response, which is an additional justification for the recommendation to proceed with EST key management operations. Once a full CA Certificate Response is obtained it is more authoritative for the domain than the limited 'pinned-domain-cert' response.

#### 5.7. Pledge BRSKI Status Telemetry

The domain is expected to provide indications to the system administrators concerning device lifecycle status. To facilitate this it needs telemetry information concerning the device's status.

The pledge MUST indicate its pledge status regarding the voucher. It does this by sending a status message to the Registrar.

The posted data media type: application/json

The client sends an HTTP POST to the server at the URI ".well-known/brski/voucher\_status".

The format and semantics described below are for version 1. A version field is included to permit significant changes to this feedback in the future. A Registrar that receives a status message with a version larger than it knows about SHOULD log the contents and alert a human.

The Status field indicates if the voucher was acceptable. Boolean values are acceptable, where "true" indicates the voucher was acceptable.

If the voucher was not acceptable the Reason string indicates why. In the failure case this message may be sent to an unauthenticated, potentially malicious registrar and therefore the Reason string SHOULD NOT provide information beneficial to an attacker. The operational benefit of this telemetry information is balanced against the operational costs of not recording that an voucher was ignored by a client the registrar expected to continue joining the domain.

The reason-context attribute is an arbitrary JSON object (literal value or hash of values) which provides additional information specific to this pledge. The contents of this field are not subject to standardization.

The version and status fields MUST be present. The Reason field SHOULD be present whenever the status field is false. The Reason-Context field is optional. In the case of a SUCCESS the Reason string MAY be omitted.

The keys to this JSON object are case-sensitive and MUST be lowercase. Figure 16 shows an example JSON.

```
<CODE BEGINS> file "voucherstatus.cddl"
voucherstatus-post = {
  "version": uint,
  "status": bool,
  ? "reason": text,
  ? "reason-context" : { $$arbitrary-map }
}
<CODE ENDS>
```

Figure 15: CDDL for voucher status POST

```
{
  "version": 1,
  "status":false,
  "reason":"Informative human readable message",
  "reason-context": { "additional" : "JSON" }
}
```

Figure 16: Example Status Telemetry

The server SHOULD respond with an HTTP 200 but MAY simply fail with an HTTP 404 error. The client ignores any response. Within the server logs the server SHOULD capture this telemetry information.

Additional standard JSON fields in this POST MAY be added, see Section 8.5. A server that sees unknown fields should log them, but otherwise ignore them.

## 5.8. Registrar audit-log request

After receiving the pledge status telemetry Section 5.7, the registrar SHOULD request the MASA audit-log from the MASA service.

This is done with an HTTP POST using the operation path value of `"/.well-known/brski/requestauditlog"`.

The registrar SHOULD HTTP POST the same registrar voucher-request as it did when requesting a voucher (using the same Content-Type). It is posted to the `/requestauditlog` URI instead. The `"idevid-issuer"`

and "serial-number" informs the MASA which log is requested so the appropriate log can be prepared for the response. Using the same media type and message minimizes cryptographic and message operations although it results in additional network traffic. The relying MASA implementation MAY leverage internal state to associate this request with the original, and by now already validated, voucher-request so as to avoid an extra crypto validation.

A registrar MAY request logs at future times. If the registrar generates a new request then the MASA is forced to perform the additional cryptographic operations to verify the new request.

A MASA that receives a request for a device that does not exist, or for which the requesting owner was never an owner returns an HTTP 404 ("Not found") code.

It is reasonable for a Registrar, that the MASA does not believe to be the current owner, to request the audit-log. There are probably reasons for this which are hard to predict in advance. For instance, such a registrar may not be aware that the device has been resold; it may be that the device has been resold inappropriately, and this is how the original owner will learn of the occurrence. It is also possible that the device legitimately spends time in two different networks.

Rather than returning the audit-log as a response to the POST (with a return code 200), the MASA MAY instead return a 201 ("Created") response ([RFC7231] sections 6.3.2 and 7.1), with the URL to the prepared (and idempotent, therefore cachable) audit response in the Location: header field.

In order to avoid enumeration of device audit-logs, MASA that return URLs SHOULD take care to make the returned URL unguessable. [W3C.WD-capability-urls-20140218] provides very good additional guidance. For instance, rather than returning URLs containing a database number such as <https://example.com/auditlog/1234> or the EUI of the device such as <https://example.com/auditlog/10-00-00-11-22-33>, the MASA SHOULD return a randomly generated value (a "slug" in web parlance). The value is used to find the relevant database entry.

A MASA that returns a code 200 MAY also include a Location: header for future reference by the registrar.

### 5.8.1. MASA audit log response

A log data file is returned consisting of all log entries associated with the device selected by the IDevID presented in the request. The audit log may be abridged by removal of old or repeated values as explained below. The returned data is in JSON format ([RFC8259]), and the Content-Type SHOULD be "application/json".

The following CDDL ([RFC8610]) explains the structure of the JSON format audit-log response:

```
<CODE BEGINS> file "auditlog.cddl"
audit-log-response = {
  "version": uint,
  "events": [ + event ]
  "truncation": {
    ? "nonced duplicates": uint,
    ? "nonceless duplicates": uint,
    ? "arbitrary": uint,
  }
}

event = {
  "date": text,
  "domainID": text,
  "nonce": text / null,
  "assertion": "verified" / "logged" / "proximity",
  ? "truncated": uint,
}
<CODE ENDS>
```

Figure 17: CDDL for audit-log response

An example:

```

{
  "version": "1",
  "events": [
    {
      "date": "2019-05-15T17:25:55.644-04:00",
      "domainID": "BduJhdHPpfhQLyponf48JzXSGZ8=",
      "nonce": "VOUFT-WwrEv0NuAQEHoV7Q",
      "assertion": "proximity",
      "truncated": "0"
    },
    {
      "date": "2017-05-15T17:25:55.644-04:00",
      "domainID": "BduJhdHPpfhQLyponf48JzXSGZ8=",
      "nonce": "f4G6Vilt8nKo/FieCVgpBg==",
      "assertion": "proximity"
    }
  ],
  "truncation": {
    "nonced duplicates": "0",
    "nonceless duplicates": "1",
    "arbitrary": "2"
  }
}

```

Figure 18: Example of audit-log response

The domainID is a binary SubjectKeyIdentifier value calculated according to Section 5.8.2. It is encoded once in base64 in order to be transported in this JSON container.

The date is in [RFC3339] format, which is consistent with typical JavaScript usage of JSON.

The truncation structure MAY be omitted if all values are zero. Any counter missing from the truncation structure is to be assumed to be zero.

The nonce is a string, as provided in the voucher-request, and used in the voucher. If no nonce was placed in the resulting voucher, then a value of null SHOULD be used in preference to omitting the entry. While the nonce is often created as a base64 encoded random series of bytes, this should not be assumed.

Distribution of a large log is less than ideal. This structure can be optimized as follows: Nonced or Nonceless entries for the same domainID MAY be abridged from the log leaving only the single most recent nonced or nonceless entry for that domainID. In the case of truncation the 'event' truncation value SHOULD contain a count of the

number of events for this domainID that were omitted. The log SHOULD NOT be further reduced but there could exist operational situation where maintaining the full log is not possible. In such situations the log MAY be arbitrarily abridged for length, with the number of removed entries indicated as 'arbitrary'.

If the truncation count exceeds 1024 then the MASA MAY use this value without further incrementing it.

A log where duplicate entries for the same domain have been omitted ("nonced duplicates" and/or "nonceless duplicates) could still be acceptable for informed decisions. A log that has had "arbitrary" truncations is less acceptable but manufacturer transparency is better than hidden truncations.

A registrar that sees a version value greater than 1 indicates an audit log format that has been enhanced with additional information. No information will be removed in future versions; should an incompatible change be desired in the future, then a new HTTP end point will be used.

This document specifies a simple log format as provided by the MASA service to the registrar. This format could be improved by distributed consensus technologies that integrate vouchers with technologies such as block-chain or hash trees or optimized logging approaches. Doing so is out of the scope of this document but is an anticipated improvement for future work. As such, the registrar SHOULD anticipate new kinds of responses, and SHOULD provide operator controls to indicate how to process unknown responses.

#### 5.8.2. Calculation of domainID

The domainID is a binary value (a BIT STRING) that uniquely identifies a Registrar by the "pinned-domain-cert".

If the "pinned-domain-cert" certificate includes the SubjectKeyIdentifier (Section 4.2.1.2 [RFC5280]), then it is to be used as the domainID. If not, the SPKI Fingerprint as described in [RFC7469] section 2.4 is to be used. This value needs to be calculated by both MASA (to populate the audit-log), and by the Registrar (to recognize itself in the audit log).

[RFC5280] section 4.2.1.2 does not mandate that the SubjectKeyIdentifier extension be present in non-CA certificates. It is RECOMMENDED that Registrar certificates (even if self-signed), always include the SubjectKeyIdentifier to be used as a domainID.

The domainID is determined from the certificate chain associated with the pinned-domain-cert and is used to update the audit-log.

### 5.8.3. Registrar audit log verification

Each time the Manufacturer Authorized Signing Authority (MASA) issues a voucher, it appends details of the assignment to an internal audit log for that device. The internal audit log is processed when responding to requests for details as described in Section 5.8. The contents of the audit log can express a variety of trust levels, and this section explains what kind of trust a registrar can derive from the entries.

While the audit log provides a list of vouchers that were issued by the MASA, the vouchers are issued in response to voucher-requests, and it is the contents of the voucher-requests which determines how meaningful the audit log entries are.

A registrar SHOULD use the log information to make an informed decision regarding the continued bootstrapping of the pledge. The exact policy is out of scope of this document as it depends on the security requirements within the registrar domain. Equipment that is purchased pre-owned can be expected to have an extensive history. The following discussion is provided to help explain the value of each log element:

**date:** The date field provides the registrar an opportunity to divide the log around known events such as the purchase date. Depending on context known to the registrar or administrator events before/after certain dates can have different levels of importance. For example for equipment that is expected to be new, and thus have no history, it would be a surprise to find prior entries.

**domainID:** If the log includes an unexpected domainID then the pledge could have imprinted on an unexpected domain. The registrar can be expected to use a variety of techniques to define "unexpected" ranging from white lists of prior domains to anomaly detection (e.g. "this device was previously bound to a different domain than any other device deployed"). Log entries can also be compared against local history logs in search of discrepancies (e.g. "this device was re-deployed some number of times internally but the external audit log shows additional re-deployments our internal logs are unaware of").

**nonce:** Nonceless entries mean the logged domainID could theoretically trigger a reset of the pledge and then take over management by using the existing nonceless voucher.

assertion: The assertion leaf in the voucher and audit log indicates why the MASA issued the voucher. A "verified" entry means that the MASA issued the associated voucher as a result of positive verification of ownership. However, this entry does not indicate whether the pledge was actually deployed in the prior domain, or not. A "logged" assertion informs the registrar that the prior vouchers were issued with minimal verification. A "proximity" assertion assures the registrar that the pledge was truly communicating with the prior domain and thus provides assurance that the prior domain really has deployed the pledge.

A relatively simple policy is to white list known (internal or external) domainIDs, and require all vouchers to have a nonce. An alternative is to require that all nonceless vouchers be from a subset (e.g. only internal) of domainIDs. If the policy is violated a simple action is to revoke any locally issued credentials for the pledge in question or to refuse to forward the voucher. The Registrar MUST then refuse any EST actions, and SHOULD inform a human via a log. A registrar MAY be configured to ignore (i.e. override the above policy) the history of the device but it is RECOMMENDED that this only be configured if hardware assisted (i.e. TPM anchored) Network Endpoint Assessment (NEA) [RFC5209] is supported.

#### 5.9. EST Integration for PKI bootstrapping

The pledge SHOULD follow the BRSKI operations with EST enrollment operations including "CA Certificates Request", "CSR Attributes" and "Client Certificate Request" or "Server-Side Key Generation", etc. This is a relatively seamless integration since BRSKI API calls provide an automated alternative to the manual bootstrapping method described in [RFC7030]. As noted above, use of HTTP persistent connections simplifies the pledge state machine.

Although EST allows clients to obtain multiple certificates by sending multiple Certificate Signing Requests (CSR) requests, BRSKI does not support this mechanism directly. This is because BRSKI pledges MUST use the CSR Attributes request ([RFC7030] section 4.5). The registrar MUST validate the CSR against the expected attributes. This implies that client requests will "look the same" and therefore result in a single logical certificate being issued even if the client were to make multiple requests. Registrars MAY contain more complex logic but doing so is out-of-scope of this specification. BRSKI does not signal any enhancement or restriction to this capability.

### 5.9.1. EST Distribution of CA Certificates

The pledge SHOULD request the full EST Distribution of CA Certificates message. See RFC7030, section 4.1.

This ensures that the pledge has the complete set of current CA certificates beyond the pinned-domain-cert (see Section 5.6.2 for a discussion of the limitations inherent in having a single certificate instead of a full CA Certificates response.) Although these limitations are acceptable during initial bootstrapping, they are not appropriate for ongoing PKIX end entity certificate validation.

### 5.9.2. EST CSR Attributes

Automated bootstrapping occurs without local administrative configuration of the pledge. In some deployments it is plausible that the pledge generates a certificate request containing only identity information known to the pledge (essentially the X.509 IDevID information) and ultimately receives a certificate containing domain specific identity information. Conceptually the CA has complete control over all fields issued in the end entity certificate. Realistically this is operationally difficult with the current status of PKI certificate authority deployments, where the CSR is submitted to the CA via a number of non-standard protocols. Even with all standardized protocols used, it could operationally be problematic to expect that service specific certificate fields can be created by a CA that is likely operated by a group that has no insight into different network services/protocols used. For example, the CA could even be outsourced.

To alleviate these operational difficulties, the pledge MUST request the EST "CSR Attributes" from the EST server and the EST server needs to be able to reply with the attributes necessary for use of the certificate in its intended protocols/services. This approach allows for minimal CA integrations and instead the local infrastructure (EST server) informs the pledge of the proper fields to include in the generated CSR (such as rfc822Name). This approach is beneficial to automated bootstrapping in the widest number of environments.

In networks using the BRSKI enrolled certificate to authenticate the ACP (Autonomic Control Plane), the EST CSR attributes MUST include the ACP Domain Information Fields defined in [I-D.ietf-anima-autonomic-control-plane] section 6.1.1.

The registrar MUST also confirm that the resulting CSR is formatted as indicated before forwarding the request to a CA. If the registrar is communicating with the CA using a protocol such as full CMC, which provides mechanisms to override the CSR attributes, then these mechanisms MAY be used even if the client ignores CSR Attribute guidance.

### 5.9.3. EST Client Certificate Request

The pledge MUST request a new client certificate. See RFC7030, section 4.2.

### 5.9.4. Enrollment Status Telemetry

For automated bootstrapping of devices, the administrative elements providing bootstrapping also provide indications to the system administrators concerning device lifecycle status. This might include information concerning attempted bootstrapping messages seen by the client. The MASA provides logs and status of credential enrollment. [RFC7030] assumes an end user and therefore does not include a final success indication back to the server. This is insufficient for automated use cases.

The client MUST send an indicator to the Registrar about its enrollment status. It does this by using an HTTP POST of a JSON dictionary with the of attributes described below to the new EST endpoint at `"/.well-known/brski/enrollstatus"`. (XXX ?)

When indicating a successful enrollment the client SHOULD first re-establish the EST TLS session using the newly obtained credentials. TLS 1.2 supports doing this in-band, but TLS 1.3 does not. The client SHOULD therefore always close the existing TLS connection, and start a new one.

In the case of a failed enrollment, the client MUST send the telemetry information over the same TLS connection that was used for the enrollment attempt, with a Reason string indicating why the most recent enrollment failed. (For failed attempts, the TLS connection is the most reliable way to correlate server-side information with what the client provides.)

The version and status fields MUST be present. The Reason field SHOULD be present whenever the status field is false. In the case of a SUCCESS the Reason string MAY be omitted.

The reason-context attribute is an arbitrary JSON object (literal value or hash of values) which provides additional information specific to the failure to unroll from this pledge. The contents of this field are not subject to standardization. This is represented by the group-socket "\$\$arbitrary-map" in the CDDL.

In the case of a SUCCESS the Reason string is omitted.

```
<CODE BEGINS> file "enrollstatus.cddl"
enrollstatus-post = {
  "version": uint,
  "status": bool,
  ? "reason": text,
  ? "reason-context" : { $$arbitrary-map }
}
<CODE ENDS>
```

Figure 19: CDDL for enrollment status POST

An example status report can be seen below. It is sent with with the media type: application/json

```
{
  "version": 1,
  "status":true,
  "reason":"Informative human readable message",
  "reason-context": { "additional" : "JSON" }
}
```

Figure 20: Example of enrollment status POST

The server SHOULD respond with an HTTP 200 but MAY simply fail with an HTTP 404 error.

Within the server logs the server MUST capture if this message was received over an TLS session with a matching client certificate.

#### 5.9.5. Multiple certificates

Pledges that require multiple certificates could establish direct EST connections to the registrar.

### 5.9.6. EST over CoAP

This document describes extensions to EST for the purposes of bootstrapping of remote key infrastructures. Bootstrapping is relevant for CoAP enrollment discussions as well. The definition of EST and BRSKI over CoAP is not discussed within this document beyond ensuring proxy support for CoAP operations. Instead it is anticipated that a definition of CoAP mappings will occur in subsequent documents such as [I-D.ietf-ace-coap-est] and that CoAP mappings for BRSKI will be discussed either there or in future work.

## 6. Clarification of transfer-encoding

[RFC7030] defines its endpoints to include a "Content-Transfer-Encoding" heading, and the payloads to be [RFC4648] Base64 encoded DER.

When used within BRSKI, the original RFC7030 EST endpoints remain Base64 encoded, but the new BRSKI end points which send and receive binary artifacts (specifically, `"/.well-known/brski/requestvoucher"`) are binary. That is, no encoding is used.

In the BRSKI context, the EST "Content-Transfer-Encoding" header field if present, SHOULD be ignored. This header field does not need to be included.

## 7. Reduced security operational modes

A common requirement of bootstrapping is to support less secure operational modes for support specific use cases. This section suggests a range of mechanisms that would alter the security assurance of BRSKI to accommodate alternative deployment architectures and mitigate lifecycle management issues identified in Section 10. They are presented here as informative (non-normative) design guidance for future standardization activities. Section 9 provides standardization applicability statements for the ANIMA ACP. Other users would be expected that subsets of these mechanisms could be profiled with an accompanying applicability statements similar to the one described in Section 9.

This section is considered non-normative in the generality of the protocol. Use of the suggested mechanisms here MUST be detailed in specific profiles of BRSKI, such as in Section 9.

### 7.1. Trust Model

This section explains the trust relationships detailed in Section 2.4:



Figure 10

**Pledge:** The pledge could be compromised and providing an attack vector for malware. The entity is trusted to only imprint using secure methods described in this document. Additional endpoint assessment techniques are RECOMMENDED but are out-of-scope of this document.

**Join Proxy:** Provides proxy functionalities but is not involved in security considerations.

**Registrar:** When interacting with a MASA a registrar makes all decisions. For Ownership Audit Vouchers (see [RFC8366]) the registrar is provided an opportunity to accept MASA decisions.

**Vendor Service, MASA:** This form of manufacturer service is trusted to accurately log all claim attempts and to provide authoritative log information to registrars. The MASA does not know which devices are associated with which domains. These claims could be strengthened by using cryptographic log techniques to provide append only, cryptographic assured, publicly auditable logs.

**Vendor Service, Ownership Validation:** This form of manufacturer service is trusted to accurately know which device is owned by which domain.

## 7.2. Pledge security reductions

The following is a list of alternative behaviours that the pledge can be programmed to implement. These behaviours are not mutually exclusive, nor are they dependent upon each other. Some of these methods enable offline and emergency (touch based) deployment use cases. Normative language is used as these behaviours are referenced in later sections in a normative fashion.

1. The pledge **MUST** accept nonceless vouchers. This allows for a use case where the registrar can not connect to the MASA at the deployment time. Logging and validity periods address the security considerations of supporting these use cases.

2. Many devices already support "trust on first use" for physical interfaces such as console ports. This document does not change that reality. Devices supporting this protocol MUST NOT support "trust on first use" on network interfaces. This is because "trust on first use" over network interfaces would undermine the logging based security protections provided by this specification.
3. The pledge MAY have an operational mode where it skips voucher validation one time. For example if a physical button is depressed during the bootstrapping operation. This can be useful if the manufacturer service is unavailable. This behavior SHOULD be available via local configuration or physical presence methods (such as use of a serial/craft console) to ensure new entities can always be deployed even when autonomic methods fail. This allows for unsecured imprint.
4. A craft/serial console could include a command such as "est-enroll [2001:db8:0:1]:443" that begins the EST process from the point after the voucher is validated. This process SHOULD include server certificate verification using an on-screen fingerprint.

It is RECOMMENDED that "trust on first use" or any method of skipping voucher validation (including use of craft serial console) only be available if hardware assisted Network Endpoint Assessment (NEA: [RFC5209]) is supported. This recommendation ensures that domain network monitoring can detect inappropriate use of offline or emergency deployment procedures when voucher-based bootstrapping is not used.

### 7.3. Registrar security reductions

A registrar can choose to accept devices using less secure methods. They MUST NOT be the default behavior. These methods may be acceptable in situations where threat models indicate that low security is adequate. This includes situations where security decisions are being made by the local administrator:

1. A registrar MAY choose to accept all devices, or all devices of a particular type, at the administrator's discretion. This could occur when informing all registrars of unique identifiers of new entities might be operationally difficult.
2. A registrar MAY choose to accept devices that claim a unique identity without the benefit of authenticating that claimed identity. This could occur when the pledge does not include an X.509 IDevID factory installed credential. New Entities without

an X.509 IDevID credential MAY form the Section 5.2 request using the Section 5.5 format to ensure the pledge's serial number information is provided to the registrar (this includes the IDevID AuthorityKeyIdentifier value, which would be statically configured on the pledge.) The pledge MAY refuse to provide a TLS client certificate (as one is not available.) The pledge SHOULD support HTTP-based or certificate-less TLS authentication as described in EST RFC7030 section 3.3.2. A registrar MUST NOT accept unauthenticated New Entities unless it has been configured to do so by an administrator that has verified that only expected new entities can communicate with a registrar (presumably via a physically secured perimeter.)

3. A registrar MAY submit a nonceless voucher-requests to the MASA service (by not including a nonce in the voucher-request.) The resulting vouchers can then be stored by the registrar until they are needed during bootstrapping operations. This is for use cases where the target network is protected by an air gap and therefore cannot contact the MASA service during pledge deployment.
4. A registrar MAY ignore unrecognized nonceless log entries. This could occur when used equipment is purchased with a valid history being deployed in air gap networks that required offline vouchers.
5. A registrar MAY accept voucher formats of future types that can not be parsed by the Registrar. This reduces the Registrar's visibility into the exact voucher contents but does not change the protocol operations.

#### 7.4. MASA security reductions

Lower security modes chosen by the MASA service affect all device deployments unless the lower-security behavior is tied to specific device identities. The modes described below can be applied to specific devices via knowledge of what devices were sold. They can also be bound to specific customers (independent of the device identity) by authenticating the customer's Registrar.

##### 7.4.1. Issuing Nonceless vouchers

A MASA has the option of not including a nonce in the voucher, and/or not requiring one to be present in the voucher-request. This results in distribution of a voucher that may never expire and in effect makes the specified Domain an always trusted entity to the pledge during any subsequent bootstrapping attempts. That a nonceless voucher was issued is captured in the log information so that the

registrar can make appropriate security decisions when a pledge joins the Domain. Nonceless vouchers are useful to support use cases where registrars might not be online during actual device deployment.

While a nonceless voucher may include an expiry date, a typical use for a nonceless voucher is for it to be long-lived. If the device can be trusted to have an accurate clock (the MASA will know), then a nonceless voucher CAN be issued with a limited lifetime.

A more typical case for a nonceless voucher is for use with offline onboarding scenarios where it is not possible to pass a fresh voucher-request to the MASA. The use of a long-lived voucher also eliminates concern about the availability of the MASA many years in the future. Thus many nonceless vouchers will have no expiry dates.

Thus, the long lived nonceless voucher does not require the proof that the device is online. Issuing such a thing is only accepted when the registrar is authenticated by the MASA and the MASA is authorized to provide this functionality to this customer. The MASA is RECOMMENDED to use this functionality only in concert with an enhanced level of ownership tracking, the details of which are out of scope for this document.

If the pledge device is known to have a real-time-clock that is set from the factory, use of a voucher validity period is RECOMMENDED.

#### 7.4.2. Trusting Owners on First Use

A MASA has the option of not verifying ownership before responding with a voucher. This is expected to be a common operational model because doing so relieves the manufacturer providing MASA services from having to track ownership during shipping and supply chain and allows for a very low overhead MASA service. A registrar uses the audit log information as a defense in depth strategy to ensure that this does not occur unexpectedly (for example when purchasing new equipment the registrar would throw an error if any audit log information is reported.) The MASA SHOULD verify the 'prior-signed-voucher-request' information for pledges that support that functionality. This provides a proof-of-proximity check that reduces the need for ownership verification. The proof-of-proximity comes from the assumption that the pledge and Join Proxy are on the same link-local connection.

A MASA that practices Trust-on-First-Use (TOFU) for Registrar identity may wish to annotate the origin of the connection by IP address or netblock, and restrict future use of that identity from other locations. A MASA that does this SHOULD take care to not create nuisance situations for itself when a customer has multiple registrars, or uses outgoing IPv4 NAT44 connections that change frequently.

#### 7.4.3. Updating or extending voucher trust anchors

This section deals with the problem of a MASA that is no longer available due to a failed business, or the situation where a MASA is uncooperative to a secondary sale.

A manufacturer could offer a management mechanism that allows the list of voucher verification trust anchors to be extended. [I-D.ietf-netconf-keystore] is one such interface that could be implemented using YANG. Pretty much any configuration mechanism used today could be extended to provide the needed additional update. A manufacturer could even decide to install the domain CA trust anchors received during the EST "cacerts" step as voucher verification anchors. Some additional signals will be needed to clearly identify which keys have voucher validation authority from among those signed by the domain CA. This is future work.

With the above change to the list of anchors, vouchers can be issued by an alternate MASA. This could be the previous owner (the seller), or some other trusted third party who is mediating the sale. If it was a third party, then the seller would need to have taken steps to introduce the third party configuration to the device prior disconnection. The third party (e.g. a wholesaler of used equipment) could however use a mechanism described in Section 7.2 to take control of the device after receiving it physically. This would permit the third party to act as the MASA for future onboarding actions. As the IDevID certificate probably can not be replaced, the new owner's Registrar would have to support an override of the MASA URL.

To be useful for resale or other transfers of ownership one of two situations will need to occur. The simplest is that the device is not put through any kind of factory default/reset before going through onboarding again. Some other secure, physical signal would be needed to initiate it. This is most suitable for redeploying a device within the same Enterprise. This would entail having previous configuration in the system until entirely replaced by the new owner, and represents some level of risk.

The second mechanism is that there would need to be two levels of factory reset. One would take the system back entirely to manufacturer state, including removing any added trust anchors, and the second (more commonly used) one would just restore the configuration back to a known default without erasing trust anchors. This weaker factory reset might leave valuable credentials on the device and this may be unacceptable to some owners.

As a third option, the manufacturer's trust anchors could be entirely overwritten with local trust anchors. A factory default would never restore those anchors. This option comes with a lot of power, but also a lot of responsibility: if access to the private part of the new anchors are lost the manufacturer may be unable to help.

## 8. IANA Considerations

This document requires the following IANA actions:

### 8.1. The IETF XML Registry

This document registers a URI in the "IETF XML Registry" [RFC3688]. IANA is asked to register the following:

```
URI: urn:ietf:params:xml:ns:yang:ietf-voucher-request
Registrant Contact: The ANIMA WG of the IETF.
XML: N/A, the requested URI is an XML namespace.
```

### 8.2. YANG Module Names Registry

This document registers a YANG module in the "YANG Module Names" registry [RFC6020]. IANA is asked to register the following:

```
name:          ietf-voucher-request
namespace:     urn:ietf:params:xml:ns:yang:ietf-voucher-request
prefix:        vch
reference:     THIS DOCUMENT
```

### 8.3. BRSKI well-known considerations

#### 8.3.1. BRSKI .well-known registration

To the Well-Known URIs Registry, at:  
"https://www.iana.org/assignments/well-known-uris/well-known-uris.xhtml", this document registers the well-known name "brski" with the following filled-in template from [RFC5785]:

```
URI suffix: brski
Change Controller: IETF
```

IANA is asked to change the registration of "est" to now only include RFC7030 and no longer this document. Earlier versions of this document used "/.well-known/est" rather than "/.well-known/brski".

### 8.3.2. BRSKI .well-known registry

IANA is requested to create a new Registry entitled: "BRSKI well-known URIs". The registry shall have at least three columns: URI, description, and reference. New items can be added using the Specification Required process. The initial contents of this registry shall be:

	URI	document	description
to MASA	requestvoucher	[THISRFC]	pledge to registrar, and from registrar
	voucher_status	[THISRFC]	pledge to registrar
	requestauditlog	[THISRFC]	registrar to MASA
	enrollstatus	[THISRFC]	pledge to registrar

### 8.4. PKIX Registry

IANA is requested to register the following:

This document requests a number for id-mod-MASAURLExtn2016(TBD) from the pkix(7) id-mod(0) Registry.

This document has received an early allocation from the id-pe registry (SMI Security for PKIX Certificate Extension) for id-pe-masa-url with the value 32, resulting in an OID of 1.3.6.1.5.5.7.1.32.

### 8.5. Pledge BRSKI Status Telemetry

IANA is requested to create a new Registry entitled: "BRSKI Parameters", and within that Registry to create a table called: "Pledge BRSKI Status Telemetry Attributes". New items can be added using the Specification Required process. The following items are to be in the initial registration, with this document (Section 5.7) as the reference:

- \* version
- \* Status
- \* Reason
- \* reason-context

## 8.6. DNS Service Names

IANA is requested to register the following Service Names:

Service Name: brski-proxy  
Transport Protocol(s): tcp  
Assignee: IESG <iesg@ietf.org>.  
Contact: IESG <iesg@ietf.org>  
Description: The Bootstrapping Remote Secure Key  
              Infrastructures Proxy  
Reference: [This document]

Service Name: brski-registrar  
Transport Protocol(s): tcp  
Assignee: IESG <iesg@ietf.org>.  
Contact: IESG <iesg@ietf.org>  
Description: The Bootstrapping Remote Secure Key  
              Infrastructures Registrar  
Reference: [This document]

## 8.7. GRASP Objective Names

IANA is requested to register the following GRASP Objective Names:

The IANA is requested to register the value "AN\_Proxy" (without quotes) to the GRASP Objectives Names Table in the GRASP Parameter Registry. The specification for this value is this document, Section 4.1.1.

The IANA is requested to register the value "AN\_join\_registrar" (without quotes) to the GRASP Objectives Names Table in the GRASP Parameter Registry. The specification for this value is this document, Section 4.3.

## 9. Applicability to the Autonomic Control Plane (ACP)

This document provides a solution to the requirements for secure bootstrap set out in Using an Autonomic Control Plane for Stable Connectivity of Network Operations, Administration, and Maintenance [RFC8368], A Reference Model for Autonomic Networking [I-D.ietf-anima-reference-model] and specifically the An Autonomic Control Plane (ACP) [I-D.ietf-anima-autonomic-control-plane], section 3.2 (Secure Bootstrap), and section 6.1 (ACP Domain, Certificate and Network).

The protocol described in this document has appeal in a number of other non-ANIMA use cases. Such uses of the protocol will be deploying into other environments with different tradeoffs of

privacy, security, reliability and autonomy from manufacturers. As such those use cases will need to provide their own applicability statements, and will need to address unique privacy and security considerations for the environments in which they are used.

The autonomic control plane (ACP) that is bootstrapped by the BRSKI protocol is typically used in medium to large Internet Service Provider organizations. Equivalent enterprises that have significant layer-3 router connectivity also will find significant benefit, particularly if the Enterprise has many sites. (A network consisting of primarily layer-2 is not excluded, but the adjacencies that the ACP will create and maintain will not reflect the topology until all devices participate in the ACP).

In the ACP, the Join Proxy is found to be proximal because communication between the pledge and the join proxy is exclusively on IPv6 Link-Local addresses. The proximity of the Join Proxy to the Registrar is validated by the Registrar using ANI ACP IPv6 Unique Local Addresses (ULA). ULAs are not routable over the Internet, so as long as the Join Proxy is operating correctly the proximity assertion is satisfied. Other uses of BRSKI will need make similar analysis if they use proximity assertions.

As specified in the ANIMA charter, this work "..focuses on professionally-managed networks." Such a network has an operator and can do things like install, configure and operate the Registrar function. The operator makes purchasing decisions and is aware of what manufacturers it expects to see on its network.

Such an operator is also capable of performing bootstrapping of a device using a serial-console (craft console). The zero-touch mechanism presented in this and the ACP document [I-D.ietf-anima-autonomic-control-plane] represents a significant efficiency: in particular it reduces the need to put senior experts on airplanes to configure devices in person.

There is a recognition as the technology evolves that not every situation may work out, and occasionally a human may still have to visit. In recognition of this, some mechanisms are presented in Section 7.2. The manufacturer MUST provide at least one of the one-touch mechanisms described that permit enrollment to be proceed without availability of any manufacturer server (such as the MASA).

The BRSKI protocol is going into environments where there have already been quite a number of vendor proprietary management systems. Those are not expected to go away quickly, but rather to leverage the secure credentials that are provisioned by BRSKI. The connectivity requirements of said management systems are provided by the ACP.

## 9.1. Operational Requirements

This section collects operational requirements based upon the three roles involved in BRSKI: The Manufacturer Authorized Signing Authority (MASA), the (Domain) Owner and the Device. It should be recognized that the manufacturer may be involved in two roles, as it creates the software/firmware for the device, and also may be the operator of the MASA.

The requirements in this section are presented using BCP14 ([RFC2119], [RFC8174]) language. These do not represent new normative statements, just a review of a few such things in one place by role. They also apply specifically to the ANIMA ACP use case. Other use cases likely have similar, but MAY have different requirements.

### 9.1.1. MASA Operational Requirements

The manufacturer **MUST** arrange for an online service to be available called the MASA. It **MUST** be available at the URL which is encoded in the IDevID certificate extensions described in Section 2.3.2.

The online service **MUST** have access to a private key with which to sign [RFC8366] format voucher artifacts. The public key, certificate, or certificate chain **MUST** be built in to the device as part of the firmware.

It is **RECOMMENDED** that the manufacturer arrange for this signing key (or keys) to be escrowed according to typical software source code escrow practices [softwareescrow].

The MASA accepts voucher requests from Domain Owners according to an operational practice appropriate for the device. This can range from any domain owner (first-come first-served, on a TOFU-like basis), to full sales channel integration where Domain Owners need to be positively identified by TLS Client Certificate pinned, or HTTP Authentication process. The MASA creates signed voucher artifacts according to its internally defined policies.

The MASA **MUST** operate an audit log for devices that is accessible. The audit log is designed to be easily cacheable and the MASA **MAY** find it useful to put this content on a CDN.

### 9.1.2. Domain Owner Operational Requirements

The domain owner MUST operate an EST ([RFC7030]) server with the extensions described in this document. This is the JRC or Registrar. This JRC/EST server MUST announce itself using GRASP within the ACP. This EST server will typically reside with the Network Operations Center for the organization.

The domain owner MAY operate an internal certificate authority (CA) that is separate from the EST server, or it MAY combine all activities into a single device. The determination of the architecture depends upon the scale and resiliency requirements of the organization. Multiple JRC instances MAY be announced into the ACP from multiple locations to achieve an appropriate level of redundancy.

In order to recognize which devices and which manufacturers are welcome on the domain owner's network, the domain owner SHOULD maintain a white list of manufacturers. This MAY extend to integration with purchasing departments to know the serial numbers of devices.

The domain owner SHOULD use the resulting overlay ACP network to manage devices, replacing legacy out-of-band mechanisms.

The domain owner SHOULD operate one or more EST servers which can be used to renew the domain certificates (LDevIDs) which are deployed to devices. These servers MAY be the same as the JRC, or MAY be a distinct set of devices, as appropriate for resiliency.

The organization MUST take appropriate precautions against loss of access to the certificate authority private key. Hardware security modules and/or secret splitting are appropriate.

### 9.1.3. Device Operational Requirements

Devices MUST come with built-in trust anchors that permit the device to validate vouchers from the MASA.

Device MUST come with (unique, per-device) IDevID certificates that include their serial numbers, and the MASA URL extension.

Devices are expected to find Join Proxies using GRASP, and then connect to the JRC using the protocol described in this document.

Once a domain owner has been validated with the voucher, devices are expected to enroll into the domain using EST. Devices are then expected to form ACPs using IPsec over IPv6 Link-Local addresses as described in [I-D.ietf-anima-autonomic-control-plane].

Once a device has been enrolled it SHOULD listen for the address of the JRC using GRASP, and it SHOULD enable itself as a Join Proxy, and announce itself on all links/interfaces using GRASP DULL.

Devices are expected to renew their certificates before they expire.

## 10. Privacy Considerations

### 10.1. MASA audit log

The MASA audit log includes the domainID for each domain a voucher has been issued to. This information is closely related to the actual domain identity. A MASA may need additional defenses against Denial of Service attacks (Section 11.1), and this may involve collecting additional (unspecified here) information. This could provide sufficient information for the MASA service to build a detailed understanding the devices that have been provisioned within a domain.

There are a number of design choices that mitigate this risk. The domain can maintain some privacy since it has not necessarily been authenticated and is not authoritatively bound to the supply chain.

Additionally the domainID captures only the unauthenticated subject key identifier of the domain. A privacy sensitive domain could theoretically generate a new domainID for each device being deployed. Similarly a privacy sensitive domain would likely purchase devices that support proximity assertions from a manufacturer that does not require sales channel integrations. This would result in a significant level of privacy while maintaining the security characteristics provided by Registrar based audit log inspection.

### 10.2. What BRSKI-EST reveals

During the provisional phase of the BRSKI-EST connection between the Pledge and the Registrar, each party reveals its certificates to each other. For the Pledge, this includes the serialNumber attribute, the MASA URL, and the identity that signed the IDevID certificate.

TLS 1.2 reveals the certificate identities to on-path observers, including the Join Proxy.

TLS 1.3 reveals the certificate identities only to the end parties, but as the connection is provisional, an on-path attacker (MITM) can see the certificates. This includes not just malicious attackers, but also Registrars that are visible to the Pledge, but which are not part of the intended domain.

The certificate of the Registrar is rather arbitrary from the point of view of the BRSKI protocol. As no [RFC6125] validations are expected to be done, the contents could be easily pseudonymized. Any device that can see a join proxy would be able to connect to the Registrar and learn the identity of the network in question. Even if the contents of the certificate are pseudonymized, it would be possible to correlate different connections in different locations belong to the same entity. This is unlikely to present a significant privacy concern to ANIMA ACP uses of BRSKI, but may be a concern to other users of BRSKI.

The certificate of the Pledge could be revealed by a malicious Join Proxy that performed a MITM attack on the provisional TLS connection. Such an attacker would be able to reveal the identity of the Pledge to third parties if it chose to so.

Research into a mechanism to do multi-step, multi-party authenticated key agreement, incorporating some kind of zero-knowledge proof would be valuable. Such a mechanism would ideally avoid disclosing identities until pledge, registrar and MASA agree to the transaction. Such a mechanism would need to discover the location of the MASA without knowing the identity of the pledge, or the identity of the MASA. This part of the problem may be unsolvable.

### 10.3. What BRSKI-MASA reveals to the manufacturer

With consumer-oriented devices, the "call-home" mechanism in IoT devices raises significant privacy concerns. See [livingwithIoT] and [IoTstrangeThings] for exemplars. The Autonomic Control Plane (ACP) usage of BRSKI is not targeted at individual usage of IoT devices, but rather at the Enterprise and ISP creation of networks in a zero-touch fashion where the "call-home" represents a different class of privacy and lifecycle management concerns.

It needs to be re-iterated that the BRSKI-MASA mechanism only occurs once during the commissioning of the device. It is well defined, and although encrypted with TLS, it could in theory be made auditable as the contents are well defined. This connection does not occur when the device powers on or is restarted for normal routines. (It is conceivable, but remarkably unusual, that a device could be forced to go through a full factory reset during an exceptional firmware update situation, after which enrollment would have to be repeated, and a new connection would occur)

The BRSKI call-home mechanism is mediated via the owner's Registrar, and the information that is transmitted is directly auditable by the device owner. This is in stark contrast to many "call-home" protocols where the device autonomously calls home and uses an undocumented protocol.

While the contents of the signed part of the pledge voucher request can not be changed, they are not encrypted at the registrar. The ability to audit the messages by the owner of the network is a mechanism to defend against exfiltration of data by a nefarious pledge. Both are, to re-iterate, encrypted by TLS while in transit.

The BRSKI-MASA exchange reveals the following information to the manufacturer:

- \* the identity of the device being enrolled. This is revealed by transmission of a signed voucher-request containing the serial-number. The manufacturer can usually link the serial number to a device model.
- \* an identity of the domain owner in the form of the domain trust anchor. However, this is not a global PKI anchored name within the WebPKI, so this identity could be pseudonymous. If there is sales channel integration, then the MASA will have authenticated the domain owner, either via pinned certificate, or perhaps another HTTP authentication method, as per Section 5.5.4.
- \* the time the device is activated,
- \* the IP address of the domain Owner's Registrar. For ISPs and Enterprises, the IP address provides very clear geolocation of the owner. No amount of IP address privacy extensions ([RFC4941]) can do anything about this, as a simple whois lookup likely identifies the ISP or Enterprise from the upper bits anyway. A passive attacker who observes the connection definitely may conclude that the given enterprise/ISP is a customer of the particular equipment vendor. The precise model that is being enrolled will remain private.

Based upon the above information, the manufacturer is able to track a specific device from pseudonymous domain identity to the next pseudonymous domain identity. If there is sales-channel integration, then the identities are not pseudonymous.

The manufacturer knows the IP address of the Registrar, but it can not see the IP address of the device itself. The manufacturer can not track the device to a detailed physical or network location, only to the location of the Registrar. That is likely to be at the Enterprise or ISPs headquarters.

The above situation is to be distinguished from a residential/individual person who registers a device from a manufacturer. Individuals do not tend to have multiple offices, and their registrar is likely on the same network as the device. A manufacturer that sells switching/routing products to enterprises should hardly be surprised if additional purchases switching/routing products are made. Deviations from a historical trend or an establish baseline would, however, be notable.

The situation is not improved by the enterprise/ISP using anonymization services such as ToR [Dingledine2004], as a TLS 1.2 connection will reveal the ClientCertificate used, clearly identifying the enterprise/ISP involved. TLS 1.3 is better in this regard, but an active attacker can still discover the parties involved by performing a Man-In-The-Middle-Attack on the first attempt (breaking/killing it with a TCP RST), and then letting subsequent connection pass through.

A manufacturer could attempt to mix the BRSKI-MASA traffic in with general traffic their site by hosting the MASA behind the same (set) of load balancers that the companies normal marketing site is hosted behind. This makes lots of sense from a straight capacity planning point of view as the same set of services (and the same set of Distributed Denial of Service mitigations) may be used. Unfortunately, as the BRSKI-MASA connections include TLS ClientCertificate exchanges, this may easily be observed in TLS 1.2, and a traffic analysis may reveal it even in TLS 1.3. This does not make such a plan irrelevant. There may be other organizational reasons to keep the marketing site (which is often subject to frequent re-designs, outsourcing, etc.) separate from the MASA, which may need to operate reliably for decades.

#### 10.4. Manufacturers and Used or Stolen Equipment

As explained above, the manufacturer receives information each time that a device which is in factory-default mode does a zero-touch bootstrap, and attempts to enroll into a domain owner's registrar.

The manufacturer is therefore in a position to decline to issue a voucher if it detects that the new owner is not the same as the previous owner.

1. This can be seen as a feature if the equipment is believed to have been stolen. If the legitimate owner notifies the manufacturer of the theft, then when the new owner brings the device up, if they use the zero-touch mechanism, the new (illegitimate) owner reveals their location and identity.
2. In the case of Used equipment, the initial owner could inform the manufacturer of the sale, or the manufacturer may just permit resales unless told otherwise. In which case, the transfer of ownership simply occurs.
3. A manufacturer could however decide not to issue a new voucher in response to a transfer of ownership. This is essentially the same as the stolen case, with the manufacturer having decided that the sale was not legitimate.
4. There is a fourth case, if the manufacturer is providing protection against stolen devices. The manufacturer then has a responsibility to protect the legitimate owner against fraudulent claims that the equipment was stolen. In the absence of such manufacturer protection, such a claim would cause the manufacturer to refuse to issue a new voucher. Should the device go through a deep factory reset (for instance, replacement of a damaged main board component, the device would not bootstrap.
5. Finally, there is a fifth case: the manufacturer has decided to end-of-line the device, or the owner has not paid a yearly support amount, and the manufacturer refuses to issue new vouchers at that point. This last case is not new to the industry: many license systems are already deployed that have significantly worse effect.

This section has outlined five situations in which a manufacturer could use the voucher system to enforce what are clearly license terms. A manufacturer that attempted to enforce license terms via vouchers would find it rather ineffective as the terms would only be enforced when the device is enrolled, and this is not (to repeat), a daily or even monthly occurrence.

### 10.5. Manufacturers and Grey market equipment

Manufacturers of devices often sell different products into different regional markets. Which product is available in which market can be driven by price differentials, support issues (some markets may require manuals and tech-support to be done in the local language), government export regulation (such as whether strong crypto is permitted to be exported, or permitted to be used in a particular market). When an domain owner obtains a device from a different market (they can be new) and transfers it to a different location, this is called a Grey Market.

A manufacturer could decide not to issue a voucher to an enterprise/ISP based upon their location. There are a number of ways which this could be determined: from the geolocation of the registrar, from sales channel knowledge about the customer, and what products are (un-)available in that market. If the device has a GPS the coordinates of the device could even be placed into an extension of the voucher.

The above actions are not illegal, and not new. Many manufacturers have shipped crypto-weak (exportable) versions of firmware as the default on equipment for decades. The first task of an enterprise/ISP has always been to login to a manufacturer system, show one's "entitlement" (country information, proof that support payments have been made), and receive either a new updated firmware, or a license key that will activate the correct firmware.

BRSKI permits the above process to automated (in an autonomic fashion), and therefore perhaps encourages this kind of differentiation by reducing the cost of doing it.

An issue that manufacturers will need to deal with in the above automated process is when a device is shipped to one country with one set of rules (or laws or entitlements), but the domain registry is in another one. Which rules apply is something will have to be worked out: the manufacturer could come to believe they are dealing with Grey market equipment, when it is simply dealing with a global enterprise.

### 10.6. Some mitigations for meddling by manufacturers

The most obvious mitigation is not to buy the product. Pick manufacturers that are up-front about their policies, who do not change them gratuitously.

Section 7.4.3 describes some ways in which a manufacturer could provide a mechanism to manage the trust anchors and built-in certificates (IDevID) as an extension. There are a variety of mechanism, and some may take a substantial amount of work to get exactly correct. These mechanisms do not change the flow of the protocol described here, but rather allow the starting trust assumptions to be changed. This is an area for future standardization work.

Replacement of the voucher validation anchors (usually pointing to the original manufacturer's MASA) with those of the new owner permits the new owner to issue vouchers to subsequent owners. This would be done by having the selling (old) owner to run a MASA.

The BRSKI protocol depends upon a trust anchor on the device and an identity on the device. Management of these entities facilitates a few new operational modes without making any changes to the BRSKI protocol. Those modes include: offline modes where the domain owner operates an internal MASA for all devices, resell modes where the first domain owner becomes the MASA for the next (resold-to) domain owner, and services where an aggregator acquires a large variety of devices, and then acts as a pseudonymized MASA for a variety of devices from a variety of manufacturers.

Although replacement of the IDevID is not required for all modes described above, a manufacturers could support such a thing. Some may wish to consider replacement of the IDevID as an indication that the device's warrantee is terminated. For others, the privacy requirements of some deployments might consider this a standard operating practice.

As discussed at the end of Section 5.8.1, new work could be done to use a distributed consensus technology for the audit log. This would permit the audit log to continue to be useful, even when there is a chain of MASA due to changes of ownership.

#### 10.7. Death of a manufacturer

A common concern has been that a manufacturer could go out of business, leaving owners of devices unable to get new vouchers for existing products. Said products might have been previously deployed, but need to be re-initialized, they might have been purchased used, or they might have kept in a warehouse as long-term spares.

The MASA was named the Manufacturer \*Authorized\* Signing Authority to emphasize that it need not be the manufacturer itself that performs this. It is anticipated that specialist service providers will come

to exist that deal with the creation of vouchers in much the same way that many companies have outsourced email, advertising and janitorial services.

Further, it is expected that as part of any service agreement that the manufacturer would arrange to escrow appropriate private keys such that a MASA service could be provided by a third party. This has routinely been done for source code for decades.

## 11. Security Considerations

This document details a protocol for bootstrapping that balances operational concerns against security concerns. As detailed in the introduction, and touched on again in Section 7, the protocol allows for reduced security modes. These attempt to deliver additional control to the local administrator and owner in cases where less security provides operational benefits. This section goes into more detail about a variety of specific considerations.

To facilitate logging and administrative oversight, in addition to triggering Registrar verification of MASA logs, the pledge reports on voucher parsing status to the registrar. In the case of a failure, this information is informative to a potentially malicious registrar. This is mandated anyway because of the operational benefits of an informed administrator in cases where the failure is indicative of a problem. The registrar is RECOMMENDED to verify MASA logs if voucher status telemetry is not received.

To facilitate truly limited clients EST RFC7030 section 3.3.2 requirements that the client MUST support a client authentication model have been reduced in Section 7 to a statement that the registrar "MAY" choose to accept devices that fail cryptographic authentication. This reflects current (poor) practices in shipping devices without a cryptographic identity that are NOT RECOMMENDED.

During the provisional period of the connection the pledge MUST treat all HTTP header and content data as untrusted data. HTTP libraries are regularly exposed to non-secured HTTP traffic: mature libraries should not have any problems.

Pledges might chose to engage in protocol operations with multiple discovered registrars in parallel. As noted above they will only do so with distinct nonce values, but the end result could be multiple vouchers issued from the MASA if all registrars attempt to claim the device. This is not a failure and the pledge choses whichever voucher to accept based on internal logic. The registrars verifying log information will see multiple entries and take this into account for their analytics purposes.

### 11.1. Denial of Service (DoS) against MASA

There are uses cases where the MASA could be unavailable or uncooperative to the Registrar. They include active DoS attacks, planned and unplanned network partitions, changes to MASA policy, or other instances where MASA policy rejects a claim. These introduce an operational risk to the Registrar owner in that MASA behavior might limit the ability to bootstrap a pledge device. For example this might be an issue during disaster recovery. This risk can be mitigated by Registrars that request and maintain long term copies of "nonceless" vouchers. In that way they are guaranteed to be able to bootstrap their devices.

The issuance of nonceless vouchers themselves creates a security concern. If the Registrar of a previous domain can intercept protocol communications then it can use a previously issued nonceless voucher to establish management control of a pledge device even after having sold it. This risk is mitigated by recording the issuance of such vouchers in the MASA audit log that is verified by the subsequent Registrar and by Pledges only bootstrapping when in a factory default state. This reflects a balance between enabling MASA independence during future bootstrapping and the security of bootstrapping itself. Registrar control over requesting and auditing nonceless vouchers allows device owners to choose an appropriate balance.

The MASA is exposed to DoS attacks wherein attackers claim an unbounded number of devices. Ensuring a registrar is representative of a valid manufacturer customer, even without validating ownership of specific pledge devices, helps to mitigate this. Pledge signatures on the pledge voucher-request, as forwarded by the registrar in the prior-signed-voucher-request field of the registrar voucher-request, significantly reduce this risk by ensuring the MASA can confirm proximity between the pledge and the registrar making the request. Supply chain integration ("know your customer") is an additional step that MASA providers and device vendors can explore.

### 11.2. DomainID must be resistant to second-preimage attacks

The domainID is used as the reference in the audit log to the domain. The domainID is expected to be calculated by a hash that is resistant to a second-preimage attack. Such an attack would allow a second registrar to create audit log entries that are fake.

### 11.3. Availability of good random numbers

The nonce used by the Pledge in the voucher-request SHOULD be generated by a Strong Cryptographic Sequence ([RFC4086] section 6.2). TLS has a similar requirement.

In particular implementations should pay attention to the advance in [RFC4086] section 3, particularly section 3.4. The random seed used by a device at boot MUST be unique across all devices and all bootstraps. Resetting a device to factory default state does not obviate this requirement.

### 11.4. Freshness in Voucher-Requests

A concern has been raised that the pledge voucher-request should contain some content (a nonce) provided by the registrar and/or MASA in order for those actors to verify that the pledge voucher-request is fresh.

There are a number of operational problems with getting a nonce from the MASA to the pledge. It is somewhat easier to collect a random value from the registrar, but as the registrar is not yet vouched for, such a registrar nonce has little value. There are privacy and logistical challenges to addressing these operational issues, so if such a thing were to be considered, it would have to provide some clear value. This section examines the impacts of not having a fresh pledge voucher-request.

Because the registrar authenticates the pledge, a full Man-in-the-Middle attack is not possible, despite the provisional TLS authentication by the pledge (see Section 5.) Instead we examine the case of a fake registrar (Rm) that communicates with the pledge in parallel or in close time proximity with the intended registrar. (This scenario is intentionally supported as described in Section 4.1.)

The fake registrar (Rm) can obtain a voucher signed by the MASA either directly or through arbitrary intermediaries. Assuming that the MASA accepts the registrar voucher-request (either because Rm is collaborating with a legitimate registrar according to supply chain information, or because the MASA is in audit-log only mode), then a voucher linking the pledge to the registrar Rm is issued.

Such a voucher, when passed back to the pledge, would link the pledge to registrar Rm, and would permit the pledge to end the provisional state. It now trusts Rm and, if it has any security vulnerabilities leveragable by an Rm with full administrative control, can be assumed to be a threat against the intended registrar.

This flow is mitigated by the intended registrar verifying the audit logs available from the MASA as described in Section 5.8. Rm might chose to collect a voucher-request but wait until after the intended registrar completes the authorization process before submitting it. This pledge voucher-request would be 'stale' in that it has a nonce that no longer matches the internal state of the pledge. In order to successfully use any resulting voucher the Rm would need to remove the stale nonce or anticipate the pledge's future nonce state. Reducing the possibility of this is why the pledge is mandated to generate a strong random or pseudo-random number nonce.

Additionally, in order to successfully use the resulting voucher the Rm would have to attack the pledge and return it to a bootstrapping enabled state. This would require wiping the pledge of current configuration and triggering a re-bootstrapping of the pledge. This is no more likely than simply taking control of the pledge directly but if this is a consideration the target network is RECOMMENDED to take the following steps:

- \* Ongoing network monitoring for unexpected bootstrapping attempts by pledges.
- \* Retrieval and examination of MASA log information upon the occurrence of any such unexpected events. Rm will be listed in the logs along with nonce information for analysis.

#### 11.5. Trusting manufacturers

The BRSKI extensions to EST permit a new pledge to be completely configured with domain specific trust anchors. The link from built-in manufacturer-provided trust anchors to domain-specific trust anchors is mediated by the signed voucher artifact.

If the manufacturer's IDevID signing key is not properly validated, then there is a risk that the network will accept a pledge that should not be a member of the network. As the address of the manufacturer's MASA is provided in the IDevID using the extension from Section 2.3, the malicious pledge will have no problem collaborating with it's MASA to produce a completely valid voucher.

BRSKI does not, however, fundamentally change the trust model from domain owner to manufacturer. Assuming that the pledge used its IDevID with RFC7030 EST and BRSKI, the domain (registrar) still needs to trust the manufacturer.

Establishing this trust between domain and manufacturer is outside the scope of BRSKI. There are a number of mechanisms that can adopted including:

- \* Manually configuring each manufacturer's trust anchor.
- \* A Trust-On-First-Use (TOFU) mechanism. A human would be queried upon seeing a manufacturer's trust anchor for the first time, and then the trust anchor would be installed to the trusted store. There are risks with this; even if the key to name mapping is validated using something like the WebPKI, there remains the possibility that the name is a look alike: e.g, dem0.example. vs demO.example.
- \* scanning the trust anchor from a QR code that came with the packaging (this is really a manual TOFU mechanism)
- \* some sales integration process where trust anchors are provided as part of the sales process, probably included in a digital packing "slip", or a sales invoice.
- \* consortium membership, where all manufacturers of a particular device category (e.g, a light bulb, or a cable-modem) are signed by a certificate authority specifically for this. This is done by CableLabs today. It is used for authentication and authorization as part of TR-79: [docsisroot] and [TR069].

The existing WebPKI provides a reasonable anchor between manufacturer name and public key. It authenticates the key. It does not provide a reasonable authorization for the manufacturer, so it is not directly useable on it's own.

#### 11.6. Manufacturer Maintenance of trust anchors

BRSKI depends upon the manufacturer building in trust anchors to the pledge device. The voucher artifact which is signed by the MASA will be validated by the pledge using that anchor. This implies that the manufacturer needs to maintain access to a signing key that the pledge can validate.

The manufacturer will need to maintain the ability to make signatures that can be validated for the lifetime that the device could be onboarded. Whether this onboarding lifetime is less than the device lifetime depends upon how the device is used. An inventory of devices kept in a warehouse as spares might not be onboarded for many decades.

There are good cryptographic hygiene reasons why a manufacturer would not want to maintain access to a private key for many decades. A manufacturer in that situation can leverage a long-term certificate authority anchor, built-in to the pledge, and then a certificate chain may be incorporated using the normal CMS certificate set. This may increase the size of the voucher artifacts, but that is not a significant issues in non-constrained environments.

There are a few other operational variations that manufacturers could consider. For instance, there is no reason that every device need have the same set of trust anchors pre-installed. Devices built in different factories, or on different days, or any other consideration could have different trust anchors built in, and the record of which batch the device is in would be recorded in the asset database. The manufacturer would then know which anchor to sign an artifact against.

Aside from the concern about long-term access to private keys, a major limiting factor for the shelf-life of many devices will be the age of the cryptographic algorithms included. A device produced in 2019 will have hardware and software capable of validating algorithms common in 2019, and will have no defense against attacks (both quantum and von-neuman brute force attacks) which have not yet been invented. This concern is orthogonal to the concern about access to private keys, but this concern likely dominates and limits the lifespan of a device in a warehouse. If any update to firmware to support new cryptographic mechanism were possible (while the device was in a warehouse), updates to trust anchors would also be done at the same time.

The set of standard operating procedures for maintaining high value private keys is well documented. For instance, the WebPKI provides a number of options for audits at [cabforumaudit], and the DNSSEC root operations are well documented at [dnssecroot].

It is not clear if Manufacturers will take this level of precaution, or how strong the economic incentives are to maintain an appropriate level of security.

This next section examines the risk due to a compromised manufacturer IDevID signing key. This is followed by examination of the risk due to a compromised MASA key. The third section sections below examines the situation where MASA web server itself is under attacker control, but that the MASA signing key itself is safe in a not-directly connected hardware module.

#### 11.6.1. Compromise of Manufacturer IDevID signing keys

An attacker that has access to the key that the manufacturer uses to sign IDevID certificates can create counterfeit devices. Such devices can claim to be from a particular manufacturer, but be entirely different devices: Trojan horses in effect.

As the attacker controls the MASA URL in the certificate, the registrar can be convinced to talk to the attackers' MASA. The Registrar does not need to be in any kind of promiscuous mode to be vulnerable.

In addition to creating fake devices, the attacker may also be able to issue revocations for existing certificates if the IDevID certificate process relies upon CRL lists that are distributed.

There does not otherwise seem to be any risk from this compromise to devices which are already deployed, or which are sitting locally in boxes waiting for deployment (local spares). The issue is that operators will be unable to trust devices which have been in an uncontrolled warehouse as they do not know if those are real devices.

#### 11.6.2. Compromise of MASA signing keys

There are two periods of time in which to consider: when the MASA key has fallen into the hands of an attacker, and after the MASA recognizes that the key has been compromised.

##### 11.6.2.1. Attacker opportunities with compromised MASA key

An attacker that has access to the MASA signing key could create vouchers. These vouchers could be for existing deployed devices, or for devices which are still in a warehouse. In order to exploit these vouchers two things need to occur: the device has to go through a factory default boot cycle, and the registrar has to be convinced to contact the attacker's MASA.

If the attacker controls a Registrar which is visible to the device, then there is no difficulty in delivery of the false voucher. A possible practical example of an attack like this would be in a data center, at an ISP peering point (whether a public IX, or a private peering point). In such a situation, there are already cables attached to the equipment that lead to other devices (the peers at the IX), and through those links, the false voucher could be delivered. The difficult part would be get the device put through a factory reset. This might be accomplished through social engineering of data center staff. Most locked cages have ventilation holes, and possibly a long "paperclip" could reach through to depress a factory

reset button. Once such a piece of ISP equipment has been compromised, it could be used to compromise equipment that was connected to (through long haul links even), assuming that those pieces of equipment could also be forced through a factory reset.

The above scenario seems rather unlikely as it requires some element of physical access; but were there a remote exploit that did not cause a direct breach, but rather a fault that resulted in a factory reset, this could provide a reasonable path.

The above deals with ANI uses of BRSKI. For cases where 802.11 or 802.15.4 is involved, the need to connect directly to the device is eliminated, but the need to do a factory reset is not. Physical possession of the device is not required as above, provided that there is some way to force a factory reset. With some consumer devices with low overall implementation quality, the end users might be familiar with needing to reset the device regularly.

The authors are unable to come up with an attack scenario where a compromised voucher signature enables an attacker to introduce a compromised pledge into an existing operator's network. This is the case because the operator controls the communication between Registrar and MASA, and there is no opportunity to introduce the fake voucher through that conduit.

#### 11.6.2.2. Risks after key compromise is known

Once the operator of the MASA realizes that the voucher signing key has been compromised it has to do a few things.

First, it MUST issue a firmware update to all devices that had that key as a trust anchor, such that they will no longer trust vouchers from that key. This will affect devices in the field which are operating, but those devices, being in operation, are not performing onboarding operations, so this is not a critical patch.

Devices in boxes (in warehouses) are vulnerable, and remain vulnerable until patched. An operator would be prudent to unbox the devices, onboard them in a safe environment, and then perform firmware updates. This does not have to be done by the end-operator; it could be done by a distributor that stores the spares. A recommended practice for high value devices (which typically have a <4hr service window) may be to validate the device operation on a regular basis anyway.

If the onboarding process includes attestations about firmware versions, then through that process the operator would be advised to upgrade the firmware before going into production. Unfortunately, this does not help against situations where the attacker operates their own Registrar (as listed above).

[RFC8366] section 6.1 explains the need for short-lived vouchers. The nonce guarantees freshness, and the short-lived nature of the voucher means that the window to deliver a fake voucher is very short. A nonceless, long-lived voucher would be the only option for the attacker, and devices in the warehouse would be vulnerable to such a thing.

A key operational recommendation is for manufacturers to sign nonceless, long-lived vouchers with a different key that they sign short-lived vouchers. That key needs significantly better protection. If both keys come from a common trust-anchor (the manufacturer's CA), then a compromise of the manufacturer's CA would compromise both keys. Such a compromise of the manufacturer's CA likely compromises all keys outlined in this section.

#### 11.6.3. Compromise of MASA web service

An attacker that takes over the MASA web service has a number of attacks. The most obvious one is simply to take the database listing customers and devices and to sell this data to other attackers who will now know where to find potentially vulnerable devices.

The second most obvious thing that the attacker can do is to kill the service, or make it operate unreliably, making customers frustrated. This could have a serious affect on ability to deploy new services by customers, and would be a significant issue during disaster recovery.

While the compromise of the MASA web service may lead to the compromise of the MASA voucher signing key, if the signing occurs offboard (such as in a hardware signing module, HSM), then the key may well be safe, but control over it resides with the attacker.

Such an attacker can issue vouchers for any device presently in service. Said device still needs to be convinced to do through a factory reset process before an attack.

If the attacker has access to a key that is trusted for long-lived nonceless vouchers, then they could issue vouchers for devices which are not yet in service. This attack may be very hard to verify and as it would involve doing firmware updates on every device in warehouses (a potentially ruinously expensive process), a manufacturer might be reluctant to admit this possibility.

### 11.7. YANG Module Security Considerations

As described in the Security Considerations section of [RFC8366] (section 7.4), the YANG module specified in this document defines the schema for data that is subsequently encapsulated by a CMS signed-data content type, as described in Section 5 of [RFC5652]. As such, all of the YANG modeled data is protected from modification.

The use of YANG to define data structures, via the 'yang-data' statement, is relatively new and distinct from the traditional use of YANG to define an API accessed by network management protocols such as NETCONF [RFC6241] and RESTCONF [RFC8040]. For this reason, these guidelines do not follow template described by Section 3.7 of [RFC8407].

### 12. Acknowledgements

We would like to thank the various reviewers for their input, in particular William Atwood, Brian Carpenter, Fuyu Eleven, Eliot Lear, Sergey Kasatkin, Anoop Kumar, Tom Petch, Markus Stenberg, Peter van der Stok, and Thomas Werner

Significant reviews were done by Jari Arko, Christian Huitema and Russ Housley.

Henk Birkholz contributed the CDDL for the audit log response.

This document started it's life as a two-page idea from Steinthor Bjarnason.

In addition, significant review comments were received by many IESG members, including Adam Roach, Alexey Melnikov, Alissa Cooper, Benjamin Kaduk, Eric Vyncke, Roman Danyliw, and Magnus Westerlund.

### 13. References

#### 13.1. Normative References

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

Eckert, T., Behringer, M., and S. Bjarnason, "An Autonomic Control Plane (ACP)", Work in Progress, Internet-Draft, draft-ietf-anima-autonomic-control-plane-30, 30 October 2020, <<http://www.ietf.org/internet-drafts/draft-ietf-anima-autonomic-control-plane-30.txt>>.

[I-D.ietf-anima-grasp]

Bormann, C., Carpenter, B., and B. Liu, "A Generic Autonomic Signaling Protocol (GRASP)", Work in Progress,

- Internet-Draft, draft-ietf-anima-grasp-15, 13 July 2017, <<http://www.ietf.org/internet-drafts/draft-ietf-anima-grasp-15.txt>>.
- [IDevID] "IEEE 802.1AR Secure Device Identifier", December 2009, <<http://standards.ieee.org/findstds/standard/802.1AR-2009.html>>.
- [ITU.X690.1994] International Telecommunications Union, "Information Technology - ASN.1 encoding rules: Specification of Basic Encoding Rules (BER), Canonical Encoding Rules (CER) and Distinguished Encoding Rules (DER)", ITU-T Recommendation X.690, 1994.
- [REST] Fielding, R.F., "Architectural Styles and the Design of Network-based Software Architectures", 2000, <<http://www.ics.uci.edu/~fielding/pubs/dissertation/top.htm>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC3339] Klyne, G. and C. Newman, "Date and Time on the Internet: Timestamps", RFC 3339, DOI 10.17487/RFC3339, July 2002, <<https://www.rfc-editor.org/info/rfc3339>>.
- [RFC3688] Mealling, M., "The IETF XML Registry", BCP 81, RFC 3688, DOI 10.17487/RFC3688, January 2004, <<https://www.rfc-editor.org/info/rfc3688>>.
- [RFC3748] Aboba, B., Blunk, L., Vollbrecht, J., Carlson, J., and H. Levkowitz, Ed., "Extensible Authentication Protocol (EAP)", RFC 3748, DOI 10.17487/RFC3748, June 2004, <<https://www.rfc-editor.org/info/rfc3748>>.
- [RFC3927] Cheshire, S., Aboba, B., and E. Guttman, "Dynamic Configuration of IPv4 Link-Local Addresses", RFC 3927, DOI 10.17487/RFC3927, May 2005, <<https://www.rfc-editor.org/info/rfc3927>>.
- [RFC4086] Eastlake 3rd, D., Schiller, J., and S. Crocker, "Randomness Requirements for Security", BCP 106, RFC 4086, DOI 10.17487/RFC4086, June 2005, <<https://www.rfc-editor.org/info/rfc4086>>.

- [RFC4519] Sciberras, A., Ed., "Lightweight Directory Access Protocol (LDAP): Schema for User Applications", RFC 4519, DOI 10.17487/RFC4519, June 2006, <<https://www.rfc-editor.org/info/rfc4519>>.
- [RFC4648] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", RFC 4648, DOI 10.17487/RFC4648, October 2006, <<https://www.rfc-editor.org/info/rfc4648>>.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, DOI 10.17487/RFC4862, September 2007, <<https://www.rfc-editor.org/info/rfc4862>>.
- [RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 4941, DOI 10.17487/RFC4941, September 2007, <<https://www.rfc-editor.org/info/rfc4941>>.
- [RFC5272] Schaad, J. and M. Myers, "Certificate Management over CMS (CMC)", RFC 5272, DOI 10.17487/RFC5272, June 2008, <<https://www.rfc-editor.org/info/rfc5272>>.
- [RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, DOI 10.17487/RFC5280, May 2008, <<https://www.rfc-editor.org/info/rfc5280>>.
- [RFC5652] Housley, R., "Cryptographic Message Syntax (CMS)", STD 70, RFC 5652, DOI 10.17487/RFC5652, September 2009, <<https://www.rfc-editor.org/info/rfc5652>>.
- [RFC6020] Bjorklund, M., Ed., "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)", RFC 6020, DOI 10.17487/RFC6020, October 2010, <<https://www.rfc-editor.org/info/rfc6020>>.
- [RFC6125] Saint-Andre, P. and J. Hodges, "Representation and Verification of Domain-Based Application Service Identity within Internet Public Key Infrastructure Using X.509 (PKIX) Certificates in the Context of Transport Layer Security (TLS)", RFC 6125, DOI 10.17487/RFC6125, March 2011, <<https://www.rfc-editor.org/info/rfc6125>>.

- [RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", RFC 6241, DOI 10.17487/RFC6241, June 2011, <<https://www.rfc-editor.org/info/rfc6241>>.
- [RFC6762] Cheshire, S. and M. Krochmal, "Multicast DNS", RFC 6762, DOI 10.17487/RFC6762, February 2013, <<https://www.rfc-editor.org/info/rfc6762>>.
- [RFC6763] Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", RFC 6763, DOI 10.17487/RFC6763, February 2013, <<https://www.rfc-editor.org/info/rfc6763>>.
- [RFC7030] Pritikin, M., Ed., Yee, P., Ed., and D. Harkins, Ed., "Enrollment over Secure Transport", RFC 7030, DOI 10.17487/RFC7030, October 2013, <<https://www.rfc-editor.org/info/rfc7030>>.
- [RFC7230] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", RFC 7230, DOI 10.17487/RFC7230, June 2014, <<https://www.rfc-editor.org/info/rfc7230>>.
- [RFC7231] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content", RFC 7231, DOI 10.17487/RFC7231, June 2014, <<https://www.rfc-editor.org/info/rfc7231>>.
- [RFC7469] Evans, C., Palmer, C., and R. Sleevi, "Public Key Pinning Extension for HTTP", RFC 7469, DOI 10.17487/RFC7469, April 2015, <<https://www.rfc-editor.org/info/rfc7469>>.
- [RFC7950] Bjorklund, M., Ed., "The YANG 1.1 Data Modeling Language", RFC 7950, DOI 10.17487/RFC7950, August 2016, <<https://www.rfc-editor.org/info/rfc7950>>.
- [RFC7951] Lhotka, L., "JSON Encoding of Data Modeled with YANG", RFC 7951, DOI 10.17487/RFC7951, August 2016, <<https://www.rfc-editor.org/info/rfc7951>>.
- [RFC8040] Bierman, A., Bjorklund, M., and K. Watsen, "RESTCONF Protocol", RFC 8040, DOI 10.17487/RFC8040, January 2017, <<https://www.rfc-editor.org/info/rfc8040>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

- [RFC8259] Bray, T., Ed., "The JavaScript Object Notation (JSON) Data Interchange Format", STD 90, RFC 8259, DOI 10.17487/RFC8259, December 2017, <<https://www.rfc-editor.org/info/rfc8259>>.
- [RFC8366] Watsen, K., Richardson, M., Pritikin, M., and T. Eckert, "A Voucher Artifact for Bootstrapping Protocols", RFC 8366, DOI 10.17487/RFC8366, May 2018, <<https://www.rfc-editor.org/info/rfc8366>>.
- [RFC8368] Eckert, T., Ed. and M. Behringer, "Using an Autonomic Control Plane for Stable Connectivity of Network Operations, Administration, and Maintenance (OAM)", RFC 8368, DOI 10.17487/RFC8368, May 2018, <<https://www.rfc-editor.org/info/rfc8368>>.
- [RFC8407] Bierman, A., "Guidelines for Authors and Reviewers of Documents Containing YANG Data Models", BCP 216, RFC 8407, DOI 10.17487/RFC8407, October 2018, <<https://www.rfc-editor.org/info/rfc8407>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.
- [RFC8610] Birkholz, H., Vigano, C., and C. Bormann, "Concise Data Definition Language (CDDL): A Notational Convention to Express Concise Binary Object Representation (CBOR) and JSON Data Structures", RFC 8610, DOI 10.17487/RFC8610, June 2019, <<https://www.rfc-editor.org/info/rfc8610>>.

### 13.2. Informative References

- [brewski] "Urban Dictionary: Brewski", October 2019, <<https://www.urbandictionary.com/define.php?term=brewski>>.
- [cabforumaudit] "Information for Auditors and Assessors", August 2019, <<https://cabforum.org/information-for-auditors-and-assessors/>>.
- [Dingledine2004] Dingledine, R., Mathewson, N., and P. Syverson, "Tor: the second-generation onion router", 2004, <<https://spec.torproject.org/tor-spec>>.

## [dnssecroot]

"DNSSEC Practice Statement for the Root Zone ZSK Operator", December 2017, <<https://www.iana.org/dnssec/dps/zsk-operator/dps-zsk-operator-v2.0.pdf>>.

## [docsisroot]

"CableLabs Digital Certificate Issuance Service", February 2018, <<https://www.cablelabs.com/resources/digital-certificate-issuance-service/>>.

## [I-D.ietf-ace-coap-est]

Stok, P., Kampanakis, P., Richardson, M., and S. Raza, "EST over secure CoAP (EST-coaps)", Work in Progress, Internet-Draft, draft-ietf-ace-coap-est-18, 6 January 2020, <<http://www.ietf.org/internet-drafts/draft-ietf-ace-coap-est-18.txt>>.

## [I-D.ietf-anima-constrained-voucher]

Richardson, M., Stok, P., and P. Kampanakis, "Constrained Voucher Artifacts for Bootstrapping Protocols", Work in Progress, Internet-Draft, draft-ietf-anima-constrained-voucher-09, 2 November 2020, <<http://www.ietf.org/internet-drafts/draft-ietf-anima-constrained-voucher-09.txt>>.

## [I-D.ietf-anima-reference-model]

Behringer, M., Carpenter, B., Eckert, T., Ciavaglia, L., and J. Nobre, "A Reference Model for Autonomic Networking", Work in Progress, Internet-Draft, draft-ietf-anima-reference-model-10, 22 November 2018, <<http://www.ietf.org/internet-drafts/draft-ietf-anima-reference-model-10.txt>>.

## [I-D.ietf-netconf-keystore]

Watsen, K., "A YANG Data Model for a Keystore", Work in Progress, Internet-Draft, draft-ietf-netconf-keystore-20, 20 August 2020, <<http://www.ietf.org/internet-drafts/draft-ietf-netconf-keystore-20.txt>>.

## [I-D.richardson-anima-state-for-joinrouter]

Richardson, M., "Considerations for stateful vs stateless join router in ANIMA bootstrap", Work in Progress, Internet-Draft, draft-richardson-anima-state-for-joinrouter-03, 22 September 2020, <<http://www.ietf.org/internet-drafts/draft-richardson-anima-state-for-joinrouter-03.txt>>.

- [imprinting] "Wikipedia article: Imprinting", July 2015, <[https://en.wikipedia.org/wiki/Imprinting\\_\(psychology\)](https://en.wikipedia.org/wiki/Imprinting_(psychology))>.
- [IoTstrangeThings] "IoT of toys stranger than fiction: Cybersecurity and data privacy update (accessed 2018-12-02)", March 2017, <<https://www.welivesecurity.com/2017/03/03/internet-of-things-security-privacy-iot-update/>>.
- [livingwithIoT] "What is it actually like to live in a house filled with IoT devices? (accessed 2018-12-02)", February 2018, <<https://www.siliconrepublic.com/machines/iot-smart-devices-reality>>.
- [minerva] Richardsdon, M., "Minerva reference implementation for BRSKI", 2020, <<https://minerva.sandelman.ca/>>.
- [minervagithub] Richardsdon, M., "GITHUB hosting of Minerva reference code", 2020, <<https://github.com/ANIMAgus-minerva>>.
- [openssl] "OpenSSL X509 utility", September 2019, <<https://www.openssl.org/docs/man1.1.1/man1/openssl-x509.html>>.
- [RFC2663] Srisuresh, P. and M. Holdrege, "IP Network Address Translator (NAT) Terminology and Considerations", RFC 2663, DOI 10.17487/RFC2663, August 1999, <<https://www.rfc-editor.org/info/rfc2663>>.
- [RFC5209] Sangster, P., Khosravi, H., Mani, M., Narayan, K., and J. Tardo, "Network Endpoint Assessment (NEA): Overview and Requirements", RFC 5209, DOI 10.17487/RFC5209, June 2008, <<https://www.rfc-editor.org/info/rfc5209>>.
- [RFC5785] Nottingham, M. and E. Hammer-Lahav, "Defining Well-Known Uniform Resource Identifiers (URIs)", RFC 5785, DOI 10.17487/RFC5785, April 2010, <<https://www.rfc-editor.org/info/rfc5785>>.
- [RFC6960] Santesson, S., Myers, M., Ankney, R., Malpani, A., Galperin, S., and C. Adams, "X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP", RFC 6960, DOI 10.17487/RFC6960, June 2013, <<https://www.rfc-editor.org/info/rfc6960>>.

- [RFC6961] Pettersen, Y., "The Transport Layer Security (TLS) Multiple Certificate Status Request Extension", RFC 6961, DOI 10.17487/RFC6961, June 2013, <<https://www.rfc-editor.org/info/rfc6961>>.
- [RFC7228] Bormann, C., Ersue, M., and A. Keranen, "Terminology for Constrained-Node Networks", RFC 7228, DOI 10.17487/RFC7228, May 2014, <<https://www.rfc-editor.org/info/rfc7228>>.
- [RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", BCP 188, RFC 7258, DOI 10.17487/RFC7258, May 2014, <<https://www.rfc-editor.org/info/rfc7258>>.
- [RFC7435] Dukhovni, V., "Opportunistic Security: Some Protection Most of the Time", RFC 7435, DOI 10.17487/RFC7435, December 2014, <<https://www.rfc-editor.org/info/rfc7435>>.
- [RFC7575] Behringer, M., Pritikin, M., Bjarnason, S., Clemm, A., Carpenter, B., Jiang, S., and L. Ciavaglia, "Autonomic Networking: Definitions and Design Goals", RFC 7575, DOI 10.17487/RFC7575, June 2015, <<https://www.rfc-editor.org/info/rfc7575>>.
- [RFC8340] Bjorklund, M. and L. Berger, Ed., "YANG Tree Diagrams", BCP 215, RFC 8340, DOI 10.17487/RFC8340, March 2018, <<https://www.rfc-editor.org/info/rfc8340>>.
- [slowloris] "Slowloris (computer security)", February 2019, <[https://en.wikipedia.org/wiki/Slowloris\\_\(computer\\_security\)](https://en.wikipedia.org/wiki/Slowloris_(computer_security))>.
- [softwareescrow] "Wikipedia article: Software Escrow", October 2019, <[https://en.wikipedia.org/wiki/Source\\_code\\_escrow](https://en.wikipedia.org/wiki/Source_code_escrow)>.
- [Stajano99theresurrecting] Stajano, F. and R. Anderson, "The resurrecting duckling: security issues for ad-hoc wireless networks", 1999, <<https://www.cl.cam.ac.uk/~fms27/papers/1999-StajanoAnd-duckling.pdf>>.
- [TR069] "TR-69: CPE WAN Management Protocol", February 2018, <<https://www.broadband-forum.org/standards-and-software/technical-specifications/tr-069-files-tools>>.

[W3C.WD-capability-urls-20140218]

Tennison, J., "Good Practices for Capability URLs", World Wide Web Consortium WD WD-capability-urls-20140218, 18 February 2014, <<https://www.w3.org/TR/2014/WD-capability-urls-20140218>>.

## Appendix A. IPv4 and non-ANI operations

The specification of BRSKI in Section 4 intentionally only covers the mechanisms for an IPv6 pledge using Link-Local addresses. This section describes non-normative extensions that can be used in other environments.

### A.1. IPv4 Link Local addresses

Instead of an IPv6 link-local address, an IPv4 address may be generated using [RFC3927] Dynamic Configuration of IPv4 Link-Local Addresses.

In the case that an IPv4 Link-Local address is formed, then the bootstrap process would continue as in the IPv6 case by looking for a (circuit) proxy.

### A.2. Use of DHCPv4

The Pledge MAY obtain an IP address via DHCP [RFC2131]. The DHCP provided parameters for the Domain Name System can be used to perform DNS operations if all local discovery attempts fail.

## Appendix B. mDNS / DNSSD proxy discovery options

Pledge discovery of the proxy (Section 4.1) MAY be performed with DNS-based Service Discovery [RFC6763] over Multicast DNS [RFC6762] to discover the proxy at "\_brski-proxy.\_tcp.local."

Proxy discovery of the registrar (Section 4.3) MAY be performed with DNS-based Service Discovery over Multicast DNS to discover registrars by searching for the service "\_brski-registrar.\_tcp.local."

To prevent unacceptable levels of network traffic, when using mDNS, the congestion avoidance mechanisms specified in [RFC6762] section 7 MUST be followed. The pledge SHOULD listen for an unsolicited broadcast response as described in [RFC6762]. This allows devices to avoid announcing their presence via mDNS broadcasts and instead silently join a network by watching for periodic unsolicited broadcast responses.

Discovery of registrar MAY also be performed with DNS-based service discovery by searching for the service "\_brski-registrar.\_tcp.example.com". In this case the domain "example.com" is discovered as described in [RFC6763] section 11 (Appendix A.2 suggests the use of DHCP parameters).

If no local proxy or registrar service is located using the GRASP mechanisms or the above mentioned DNS-based Service Discovery methods, the pledge MAY contact a well known manufacturer provided bootstrapping server by performing a DNS lookup using a well known URI such as "brski-registrar.manufacturer.example.com". The details of the URI are manufacturer specific. Manufacturers that leverage this method on the pledge are responsible for providing the registrar service. Also see Section 2.7.

The current DNS services returned during each query are maintained until bootstrapping is completed. If bootstrapping fails and the pledge returns to the Discovery state, it picks up where it left off and continues attempting bootstrapping. For example, if the first Multicast DNS \_bootstraps.\_tcp.local response doesn't work then the second and third responses are tried. If these fail the pledge moves on to normal DNS-based Service Discovery.

#### Appendix C. Example Vouchers

Three entities are involved in a voucher: the MASA issues (signs) it, the registrar's public key is mentioned in the voucher, and the pledge validates it. In order to provide reproduceable examples the public and private keys for an example MASA and registrar are first listed.

The keys come from an open source reference implementation of BRSKI, called "Minerva" [minerva]. It is available on github [minervagithub]. The keys presented here are used in the unit and integration tests. The MASA code is called "highway", the Registrar code is called "fountain", and the example client is called "reach".

The public key components of each are presented as both base64 certificates, as well as being decoded by openssl's x509 utility so that the extensions can be seen. This was version 1.1.1c of the [openssl] library and utility.

##### C.1. Keys involved

The Manufacturer has a Certificate Authority that signs the pledge's IDevID. In addition the Manufacturer's signing authority (the MASA) signs the vouchers, and that certificate must be distributed to the devices at manufacturing time so that vouchers can be validated.

## C.1.1. Manufacturer Certificate Authority for IDevID signatures

This private key is Certificate Authority that signs IDevID certificates:

```
<CODE BEGINS> file "vendor.key"
-----BEGIN EC PRIVATE KEY-----
MIGkAgEBBDcAYkoLWlIEA5SKKhMMdkTK7sJxk5ybKqYq9Yr5aR7tNwqXyLGS7z8G
8S4w/UJ58BqgBwYFK4EEACKhZANiAAQu5/yktJbFLjMC87h7b+yTreFuF8GwewKH
L4mS0r0dVAQubqDUQcTrjvpXrXCpTojiLCzgp8fzkcUDkZ9LD/M90LDipiLNIokP
juF8QkoAbT8pMrY83MS8y76wZ7AalNQ=
-----END EC PRIVATE KEY-----
<CODE ENDS>
```

This public key validates IDevID certificates:

file: examples/vendor.key

```
<CODE BEGINS> file "vendor.cert"
Certificate:
  Data:
    Version: 3 (0x2)
    Serial Number: 519772114 (0x1efb17d2)
    Signature Algorithm: ecdsa-with-SHA256
    Issuer: C = Canada, ST = Ontario, OU = Sandelman, CN = highway-test.ex
ample.com CA
    Validity
      Not Before: Feb 12 22:22:21 2019 GMT
      Not After : Feb 11 22:22:21 2021 GMT
    Subject: C = Canada, ST = Ontario, OU = Sandelman, CN = highway-test.e
xample.com CA
    Subject Public Key Info:
      Public Key Algorithm: id-ecPublicKey
      Public-Key: (384 bit)
      pub:
        04:2e:e7:fc:a4:b4:96:c5:2e:33:02:f3:b8:7b:6f:
        ec:93:ad:e1:6e:17:c1:b0:7b:02:87:2f:89:92:d2:
        bd:1d:54:04:2e:6e:a0:d4:41:c4:eb:8e:fa:57:ad:
        70:a9:4e:88:e2:2c:2c:e0:a7:c7:f3:91:c5:03:91:
        9f:4b:0f:f3:3d:d0:b0:e2:a6:22:cd:20:e9:0f:8e:
        e1:7c:42:4a:00:6d:3f:29:32:b6:3c:dc:c4:bc:cb:
        be:b0:67:b0:1a:94:d4
      ASN1 OID: secp384r1
      NIST CURVE: P-384
    X509v3 extensions:
      X509v3 Basic Constraints: critical
        CA:TRUE
      X509v3 Key Usage: critical
        Certificate Sign, CRL Sign
      X509v3 Subject Key Identifier:
```





## C.1.3. Registrar Certificate Authority

This Certificate Authority enrolls the pledge once it is authorized, and it also signs the Registrar's certificate.

```
<CODE BEGINS> file "ownerca_secp384rl.key"
-----BEGIN EC PRIVATE KEY-----
MIGkAgEBBDChnLI0MSOLf8XndiZqoZdqblcPR5YSoPGhPOuFxFxWy1gFi9HbWv8b/R
EGdRgGEVsjKgBwYFK4EEACKhZANiAAQbf1m6F8MavGaNjGzgw/oxcQ919iKRvbdW
gAfb37h6pUVNeYpGlxlZ1jGxj2l9Mr48yD5bY7VG9qjVb5v5wPPTuRQ/ckdRpHbd
0vC/9cqPMAF/+MJf0/UgAOSLi/IHbLQ=
-----END EC PRIVATE KEY-----
<CODE ENDS>
```

The public key is indicated in a pledge voucher-request to show proximity.

file: examples/ownerca\_secp384rl.key

```
<CODE BEGINS> file "ownerca_secp384rl.cert"
Certificate:
  Data:
    Version: 3 (0x2)
    Serial Number: 694879833 (0x296b0659)
    Signature Algorithm: ecdsa-with-SHA256
    Issuer: DC = ca, DC = sandelman, CN = fountain-test.example.com Unstru
ng Fountain Root CA
    Validity
      Not Before: Feb 25 21:31:45 2020 GMT
      Not After : Feb 24 21:31:45 2022 GMT
    Subject: DC = ca, DC = sandelman, CN = fountain-test.example.com Unstr
ung Fountain Root CA
    Subject Public Key Info:
      Public Key Algorithm: id-ecPublicKey
      Public-Key: (384 bit)
      pub:
        04:1b:7f:59:ba:17:c3:1a:bc:66:8d:8c:6c:e0:c3:
        fa:31:71:0f:65:f6:22:91:bd:b7:56:80:07:db:df:
        b8:7a:a5:45:4d:79:8a:46:97:19:59:96:31:b1:8f:
        69:7d:32:be:3c:c8:3e:5b:63:b5:46:f6:a8:d5:6f:
        9b:f9:c0:f3:d3:b9:14:3f:72:47:51:a4:76:dd:d2:
        f0:bf:f5:ca:8f:30:01:7f:f8:c2:5f:d3:f5:20:03:
        44:8b:8b:f2:07:6c:b4
      ASN1 OID: secp384rl
      NIST CURVE: P-384
    X509v3 extensions:
      X509v3 Basic Constraints: critical
      CA:TRUE
      X509v3 Key Usage: critical
      Certificate Sign, CRL Sign
```

X509v3 Subject Key Identifier:  
 B9:A5:F6:CB:11:E1:07:A4:49:2C:A7:08:C6:7C:10:BC:87:B3:74:26  
 X509v3 Authority Key Identifier:  
 keyid:B9:A5:F6:CB:11:E1:07:A4:49:2C:A7:08:C6:7C:10:BC:87:B3:74

:26

Signature Algorithm: ecdsa-with-SHA256  
 30:64:02:30:20:83:06:ce:8d:98:a4:54:7a:66:4c:4a:3a:70:  
 c2:52:36:5a:52:8d:59:7d:20:9b:2a:69:14:58:87:38:d8:55:  
 79:dd:fd:29:38:95:1e:91:93:76:b4:f5:66:29:44:b4:02:30:  
 6f:38:f9:af:12:ed:30:d5:85:29:7c:b1:16:58:bd:67:91:43:  
 c4:0d:30:f9:d8:1c:ac:2f:06:dd:bc:d5:06:42:2c:84:a2:04:  
 ea:02:a4:5f:17:51:26:fb:d9:2f:d2:5c

```
-----BEGIN CERTIFICATE-----
MIICazCCAFKgAwIBAgIEKWsGWTAKBggqhkJOPQQDAjBtMRIwEAYKczImiZPyLGQB
GRYCY2ExGTAXBgoJkiaJk/IsZAEZFglzYW5kZWxtYW4xPDA6BgNVBAMMM2ZvdW50
YWluLXRlc3QuZXBhbXBsZS5jb20gVW5zdHJ1bmcgRm91bnRhaW4gUm9vdCBDQTAe
Fw0yMDAyMjUyMTMxNDVaFw0yMjUyMTMxNDVaMG0xEjAQBgoJkiaJk/IsZAEZF
FgJjYTEZMBCGcmSjJomT8ixkARKWCXNhbmlbG1hbG1hbmR1bG1hbmR1bG1hbmR1
aW4tdGVzdC5leGFtcGxlLmNvbSBVbnN0cnVuZyBGb3VudGFpbSB290IENBMHYw
EAYHkoZIZjOCAQYFK4EEACIDYgAEG39ZuhfDGrxmjYxs4MP6MXEPZfYikb23VoAH
29+4eqVFTXmKRpcZWZYxsY9pfTK+PMg+W201Rvao1W+b+cDz07kUP3JHUAR23dLw
v/XKjzABf/jCX9P1IANEi4vyB2y0o2MwYTAPBgNVHRMBAf8EBTADAQH/MA4GA1Ud
DwEB/wQEAWIBBjAdBgNVHQ4EFgQUuaX2yxHhB6RJLkCIXnwQvIezdCYwHwYDVR0j
BBgwFoAUuaX2yxHhB6RJLkCIXnwQvIezdCYwCgYIKoZIzj0EAwIDZwAwZAIwIIMG
zo2YpFR6ZkxKOnDCUjZaUo1ZfSCbKmkUWIc42FV53f0pOJUekZN2tPvmKUS0AjBv
OPmvEu0w1YUpfLEWWL1nkUPEDTD52BysLwbdvNUGQiyEogTqAqRfF1Em+9kv0lw=
-----END CERTIFICATE-----
<CODE ENDS>
```

C.1.4. Registrar key pair

The Registrar is the representative of the domain owner. This key signs registrar voucher-requests, and terminates the TLS connection from the pledge.

```
<CODE BEGINS> file "jrc_prime256v1.key"
-----BEGIN EC PRIVATE KEY-----
MHcCAQEEIFZodk+PC5Mu24+ra0sbOjKzan+dW5rvDAR7YuJUOC1YoAoGCCqGSM49
AwEHoUQDQgAEImlVQcjs6n+Xd5l/28IFv6UiegQwSBztGj5dkK2MAjQIPV8l8lH+E
jLIOYdbJiI0VtEIf1/Jqt+TOBfinTNOLog==
-----END EC PRIVATE KEY-----
<CODE ENDS>
```

The public key is indicated in a pledge voucher-request to show proximity.

```

<CODE BEGINS> file "jrc_prime256v1.cert"
Certificate:
  Data:
    Version: 3 (0x2)
    Serial Number: 1066965842 (0x3f989b52)
    Signature Algorithm: ecdsa-with-SHA256
    Issuer: DC = ca, DC = sandelman, CN = fountain-test.example.com Unstru
ng Fountain Root CA
    Validity
      Not Before: Feb 25 21:31:54 2020 GMT
      Not After : Feb 24 21:31:54 2022 GMT
    Subject: DC = ca, DC = sandelman, CN = fountain-test.example.com
    Subject Public Key Info:
      Public Key Algorithm: id-ecPublicKey
      Public-Key: (256 bit)
      pub:
        04:96:65:50:72:34:ba:9f:e5:dd:e6:5f:f6:f0:81:
        6f:e9:48:9e:81:0c:12:07:3b:46:8f:97:64:2b:63:
        00:8d:02:0f:57:c9:7c:94:7f:84:8c:b2:0e:61:d6:
        c9:88:8d:15:b4:42:1f:d7:f2:6a:b7:e4:ce:05:f8:
        a7:4c:d3:8b:3a
      ASN1 OID: prime256v1
      NIST CURVE: P-256
    X509v3 extensions:
      X509v3 Extended Key Usage: critical
      CMC Registration Authority
      X509v3 Key Usage: critical
      Digital Signature
    Signature Algorithm: ecdsa-with-SHA256
      30:65:02:30:66:4f:60:4c:55:48:1e:96:07:f8:dd:1f:b9:c8:
      12:8d:45:36:87:9b:23:c0:bc:bb:f1:cb:3d:26:15:56:6f:5f:
      1f:bf:d5:1c:0e:6a:09:af:1b:76:97:99:19:23:fd:7e:02:31:
      00:bc:ac:c3:41:b0:ba:0d:af:52:f9:9c:6e:7a:7f:00:1d:23:
      c8:62:01:61:bc:4b:c5:c0:47:99:35:0a:0c:77:61:44:01:4a:
      07:52:70:57:00:75:ff:be:07:0e:98:cb:e5
-----BEGIN CERTIFICATE-----
MIIB/DCCAYKgAwIBAgIEP5ibUjAKBggqhkJOPQQDAjBtMRItwEAYKCZImiZPyLQQB
GRYCY2ExGTAXBgoJkiaJk/IsZAEZFglzYW5kZWxtYW4xPDA6BgNVBAMMM2ZvdW50
YWluLXRlc3QuZXRhbXBsZS5jb20gVW5zdHJ1bmcgRm91bnRhaW4gUm9vdCBDQTAe
Fw0yMDAyMjUyMTMxNTRaFw0yMjAyMjUyMTMxNTRaMFMxZjAQBgoJkiaJk/IsZAEZ
FgJjYTEZMBcGCgmSjomT8ixkARKWCXNhbMR1bG1hbJeiMCAGA1UEAwZm91bnRha
aW4tdGVzdC5leGFTcGx1LmNvbTBZMBMGBYqGSM49AgEGCCqGSM49AwEHA0IABJZl
UHI0up/l3eZf9vCBb+lInoEMEGc7Ro+XZCtjAI0CD1fJfJR/hIyyDmHWyYiNFbRC
H9fyarfkzgzX4p0zTizqjKjAoMByGA1UdJQEB/wQMMAoGCCsGAQUFBwMcMA4GA1Ud
DwEB/wQEAWIHgDAKBggqhkJOPQQDAgNoADB1AjBmT2BMVUgelgf43R+5yBKNRTaH
myPAvLvxyz0mFVZvXx+/1RwOagmvG3aXmRkj/X4CMQC8rMNBsLoNr1L5nG56fwAd
I8hiAWG8S8XAR5k1Cgx3YUQBSgdScFcAdf++Bw6Yy+U=
-----END CERTIFICATE-----
<CODE ENDS>

```

## C.1.5. Pledge key pair

The pledge has an IDevID key pair built in at manufacturing time:

```
<CODE BEGINS> file "idevid_00-D0-E5-F2-00-02.key"
-----BEGIN EC PRIVATE KEY-----
MHcCAQEIEIBHNh6r8QRevRuo+tEmBJeFjQKf6bpFA/9NGoltv+9sNoAoGCCqGSM49
AwEHoUQDQgAEA6N1Q4ezfMAKmoecrfb0OBMc1AyEH+BATkF58FsTSyBxs0SbSWLx
FjDOuwB9gLGn2TsTUJumJ6VPw5Z/TP4hJw==
-----END EC PRIVATE KEY-----
<CODE ENDS>
```

The certificate is used by the registrar to find the MASA.

```
<CODE BEGINS> file "idevid_00-D0-E5-F2-00-02.cert"
Certificate:
  Data:
    Version: 3 (0x2)
    Serial Number: 226876461 (0xd85dc2d)
    Signature Algorithm: ecdsa-with-SHA256
    Issuer: C = Canada, ST = Ontario, OU = Sandelman, CN = highway-test.ex
ample.com CA
    Validity
      Not Before: Feb  3 06:47:20 2020 GMT
      Not After : Dec 31 00:00:00 2999 GMT
    Subject: serialNumber = 00-D0-E5-F2-00-02
    Subject Public Key Info:
      Public Key Algorithm: id-ecPublicKey
      Public-Key: (256 bit)
      pub:
        04:03:a3:75:43:87:b3:7c:c0:0a:9a:87:9c:ad:f6:
        f4:38:13:1c:d4:0c:84:1f:e0:40:4e:41:79:f0:5b:
        13:4b:20:71:b3:44:9b:49:62:f1:16:30:ce:bb:00:
        7d:80:b1:a7:d9:3b:13:50:9b:a6:27:a5:4f:c3:96:
        7f:4c:fe:21:27
      ASN1 OID: prime256v1
      NIST CURVE: P-256
    X509v3 extensions:
      X509v3 Subject Key Identifier:
        45:88:CC:96:96:00:64:37:B0:BA:23:65:64:64:54:08:06:6C:56:AD
      X509v3 Basic Constraints:
        CA:FALSE
      1.3.6.1.5.5.7.1.32:
        ..highway-test.example.com:9443
    Signature Algorithm: ecdsa-with-SHA256
      30:65:02:30:23:e1:a9:2e:ef:22:12:34:5a:a5:c2:15:d6:28:
      bd:ed:3d:96:d6:ce:04:95:ef:a7:c8:dc:18:a8:31:c7:b8:04:
      34:f2:b7:4d:79:8a:67:22:24:03:4f:c5:cd:d6:06:ba:02:31:
      00:b3:8d:5c:0a:d0:fe:04:83:90:d3:4f:6d:72:97:b3:3e:02:
```

```

ea:f1:c8:5a:32:72:58:b7:45:02:50:78:bc:04:1d:23:5e:22:
6f:c3:7f:8c:7c:d7:9b:70:20:91:b4:e1:7f
-----BEGIN CERTIFICATE-----
MIIB5jCCAWygAwIBAgIEDYXcLTAKBggqhkJOPQODAjbDMQ8wDQYDVQQGEwZDYW5h
ZGExEDAOBgNVBAMwMB09udGFyaW8xEjAQBgNVBAsMCVhbmR1bG1hbG1hbjEjMCIGA1UE
AwwbaGlnaHdheS10ZXN0LmV4YW1wbGUuY29tIENBMCAXDTEwMDIwMzA2NDcyMfoY
DzI5OTkxMjMxMDAwMDAwWjAcMR0wGAYDVQQFDBEwMC1EMC1FNs1GMi0wMC0wMjBZ
MBMGByqGSM49AgEGCCqGSM49AwEHA0IABAOjdUOHs3zACpQHnK329DgTHNQmHb/g
QE5BefBbE0sgcbNEM0li8RYwzrsAfYcXp9k7E1CbpielT8OWf0z+ISejWTBxMB0G
A1UdDgQWBBRFiMyWlgBkN7C6I2VkJZFIbmxWrTajBgNVHRMEAjaAMCsGCCsGAQUF
BwEgBB8MHWhpZ2h3YXktZGVzdC5leGFtcGx1LmNvbTo5NDQzMAoGCCqGSM49BAMC
A2gAMGUCMCPHqS7vIhI0WqXCFdYove09ltbOBjXvp8jcGKgx7gENPK3TXmKZyIk
A0/FzdYGugIxALONXArQ/gSDkNNPbXKXsz4C6vHIWjJyWldFA1B4vAQdI14ib8N/
jHzXm3AgkbThfw==
-----END CERTIFICATE-----
<CODE ENDS>

```

## C.2. Example process

The JSON examples below are wrapped at 60 columns. This results in strings that have newlines in them, which makes them invalid JSON as is. The strings would otherwise be too long, so they need to be unwrapped before processing.

For readability, the output of the `asnlparse` has been truncated at 72 columns rather than wrapped.

### C.2.1. Pledge to Registrar

As described in Section 5.2, the pledge will sign a pledge voucher-request containing the registrar's public key in the `proximity-registrar-cert` field. The base64 has been wrapped at 60 characters for presentation reasons.

```
<CODE BEGINS> file "vr_00-D0-E5-F2-00-02.b64"
MIIG3wYJKoZIhvcNAQcCoIIg0DCCBswCAQEExDTALBg1ghkgBZQMEAgEwggOJBgkqhkiG9w0BBwGg
ggN6BIIDnsiaWV0Zi12b3VjaGVyLXJlcXVlc3Q6dm91Y2hlciI6eyJhc3NlcnRpb24iOiJwcm94
aW1pdHkiLCJjcmVhdGVkLW9uIjoimjAyMC0wMi0yNVQxODowND00C42NTItMDU6MDAiLCJzZXJp
YWwtbnVtYmVyIjoimDAtRDAtRTUtRjItMDAtMDIiLCJub25jZSI6ImFNamd1ZUtVVC0yMndWaW1q
NnoyN1EiLCJwcm94aW1pdHktcmVnaXN0cmFyLWN1cnQiOiJNSU1CL0RDQ0FZS2dBd0lCQWdJRVA1
aWJVakFLQmdncWhrak9QUVFEQWpCdE1SSXdfQV1LQ1pJbWlaUHM1MR1FCR1JZQ1kyRXhHVEFYQmdv
SmtpYUprL0lZWkFFWkZnbHpZVzVrWld4dFlXNHhQREe2QmdOVkJBtU1NMlp2ZFclMF1XbHVMWFJs
YzNRdVpYaGhiWEJzWlM1amIyMGdWVzV6ZEhKMWJtY2dSbTkxYm5SaGFxNGdVbTl2ZENCRFFUQWVG
dzB5TURBeU1qVXlNVE14TlRSYUZ3MH1NakF5TWpReU1UTXhOVFJhTUZNeEVqQVFCZ29Ka2lhSmsv
SXNaQUVaRmdKallURVpNqMNHQ2dtU0pvbVQ4aXhrQVJrV0NYTmh1bVJsYkcxagJqRW1NQ0FHQTFV
RUF3dlpabTkxYm5SaGFxNHRkR1Z6ZEM1bGVHRnRjR3hsTG1OdmJUQ1pNQ1hQnlxR1NNND1BZ0VH
Q0NxR1NNND1Bd0VIQTBJQUJKWmxVSEkwdXAvbDN1WmY5dkNCYitsSW5vRU1Fz2M3Um8rWFpDdGpB
STBDRDFmSmZKU19oSX15RG1IV31ZaU5GY1JDSdlmeWFyZmt6Z1g0cDB6VG16cWpLakFvTUJZROEx
VWRKUUVCL3dRTU1Bb0dDQ3NHQVFRkR1J3TWNNQTRHQTfVZER3RUlvd1FFQXdJSGdEQUtCZ2dxaGtq
T1BRUURBZ05vQURCbEFqQm1Umk1JN1VnZWxnZjQzUis1eUJLTlJUyUhteVBBdkx2eH16MG1GV1p2
WHgrLzFsd09hZ212RzNhWG1Sa2ovWDRDTRVfDOHJNTkZjZTG9OcJFMNW5HNTZmd0FkSThoaUFXRzhT
OFhBUjVrMUNneDNZVVFcu2dkU2NGY0FkZisrQnc2WXkrVT0ifX2gggHqMIIB5jCAAwYgAwIBAgIE
DYXcLTAKBggqhkiJOPQDAjBdMQ8wDQYDVQQGEWZDYW5hZGExEDAObGVBAGMB09udGFyaW8xEjAQ
BgNVBAsMCVNmhmR1bG1hbJEkMCIgA1UEAwwaG1naHdheS10ZXN0LmV4YW1wbGUuY29tIENBMCAX
DTIwMDIwMzA2NDcyMfoYDzI5OTkxMjMxMDAwMDAwWjAcMRowGAYDVQQQFDBEwMC1EMClFNS1GMi0w
MC0wMjBZMBMBGByqGSM49AgEGCCqGSM49AwEHA0IABAOjduOhs3zACpqHnK329DgTHNQMHb/gQE5B
efBbE0sgcbNEM0li8RYwzrsAfYcXp9k7E1CbpielT8OWf0z+ISejWTBxMB0GA1UdDgQWBBERFiMyW
lgBkn7C6I2VkJZFIbmxWrTAJBGNVHRMEAjAAMCsGCCsGAQUFBwEgBB8MHWhpZ2h3YXktZGVzdC5l
eGFtcGx1LmNvbTo5NDQzMAoGCCqGSM49BAMCA2gAMGUCMCPHqs7vIhI0WqXCFdYove091tbOBJXv
p8jcGKxxx7gENPK3TXmKZyIkA0/FzdYGugLxALONXArQ/gSDkNNPbXKXsz4C6vHIWjJyWLDFA1B4
vAQdI14ib8N/jHzXm3AgkbThfzGCATswggE3AgEBMGUwXTEPMA0GA1UEBhMGQ2F1YWRhMRawDgYD
VQQIDAdPbnRhcmlvMR1wEAYDVQQQLDAlTYW5kZWxtYW4xJDAiBgNVBAMMG2hpZ2h3YXktZGVzdC5l
eGFtcGx1LmNvbSBDQ0IEdYXcLTALBg1ghkgBZQMEAgGgaTAYBgkqhkiG9w0BQCMxYJKoZIhvcN
AQcBMBwGCSqGSIb3DQEJBTEPFw0yMDAyMjUyMzA0NDhaMC8GCSqGSIb3DQEJBDEiBCCx6IrwstHF
609Y0EqDK62QKby4duyyIWudvs15M16BBTAKBggqhkiJOPQDAgRHMEUCIBxwA1U1kIkuQDf/j7kZ
/MVefgr141+hKBFgrnNngjwpAIEAy8aXt0GSB9m1bmiEUpefCEhxSv2xLYurG1ugv0dfr/E=
<CODE ENDS>
```

The ASN1 decoding of the artifact:

```
file: examples/vr_00-D0-E5-F2-00-02.b64
```

```
0:d=0 hl=4 l=1759 cons: SEQUENCE
4:d=1 hl=2 l= 9 prim: OBJECT :pkcs7-signedData
15:d=1 hl=4 l=1744 cons: cont [ 0 ]
19:d=2 hl=4 l=1740 cons: SEQUENCE
23:d=3 hl=2 l= 1 prim: INTEGER :01
26:d=3 hl=2 l= 13 cons: SET
28:d=4 hl=2 l= 11 cons: SEQUENCE
30:d=5 hl=2 l= 9 prim: OBJECT :sha256
41:d=3 hl=4 l= 905 cons: SEQUENCE
45:d=4 hl=2 l= 9 prim: OBJECT :pkcs7-data
```

```

56:d=4  hl=4  l= 890 cons: cont [ 0 ]
60:d=5  hl=4  l= 886 prim: OCTET STRING      :{"ietf-voucher-request:v
950:d=3  hl=4  l= 490 cons: cont [ 0 ]
954:d=4  hl=4  l= 486 cons: SEQUENCE
958:d=5  hl=4  l= 364 cons: SEQUENCE
962:d=6  hl=2  l=   3 cons: cont [ 0 ]
964:d=7  hl=2  l=   1 prim: INTEGER           :02
967:d=6  hl=2  l=   4 prim: INTEGER           :0D85DC2D
973:d=6  hl=2  l=  10 cons: SEQUENCE
975:d=7  hl=2  l=   8 prim: OBJECT            :ecdsa-with-SHA256
985:d=6  hl=2  l=  93 cons: SEQUENCE
987:d=7  hl=2  l=  15 cons: SET
989:d=8  hl=2  l=  13 cons: SEQUENCE
991:d=9  hl=2  l=   3 prim: OBJECT            :countryName
996:d=9  hl=2  l=   6 prim: PRINTABLESTRING   :Canada
1004:d=7  hl=2  l=  16 cons: SET
1006:d=8  hl=2  l=  14 cons: SEQUENCE
1008:d=9  hl=2  l=   3 prim: OBJECT            :stateOrProvinceName
1013:d=9  hl=2  l=   7 prim: UTF8STRING       :Ontario
1022:d=7  hl=2  l=  18 cons: SET
1024:d=8  hl=2  l=  16 cons: SEQUENCE
1026:d=9  hl=2  l=   3 prim: OBJECT            :organizationalUnitName
1031:d=9  hl=2  l=   9 prim: UTF8STRING       :Sandelman
1042:d=7  hl=2  l=  36 cons: SET
1044:d=8  hl=2  l=  34 cons: SEQUENCE
1046:d=9  hl=2  l=   3 prim: OBJECT            :commonName
1051:d=9  hl=2  l=  27 prim: UTF8STRING       :highway-test.example.com
1080:d=6  hl=2  l=  32 cons: SEQUENCE
1082:d=7  hl=2  l=  13 prim: UTCTIME           :200203064720Z
1097:d=7  hl=2  l=  15 prim: GENERALIZEDTIME   :29991231000000Z
1114:d=6  hl=2  l=  28 cons: SEQUENCE
1116:d=7  hl=2  l=  26 cons: SET
1118:d=8  hl=2  l=  24 cons: SEQUENCE
1120:d=9  hl=2  l=   3 prim: OBJECT            :serialNumber
1125:d=9  hl=2  l=  17 prim: UTF8STRING       :00-D0-E5-F2-00-02
1144:d=6  hl=2  l=  89 cons: SEQUENCE
1146:d=7  hl=2  l=  19 cons: SEQUENCE
1148:d=8  hl=2  l=   7 prim: OBJECT            :id-ecPublicKey
1157:d=8  hl=2  l=   8 prim: OBJECT            :prime256v1
1167:d=7  hl=2  l=  66 prim: BIT STRING
1235:d=6  hl=2  l=  89 cons: cont [ 3 ]
1237:d=7  hl=2  l=  87 cons: SEQUENCE
1239:d=8  hl=2  l=  29 cons: SEQUENCE
1241:d=9  hl=2  l=   3 prim: OBJECT            :X509v3 Subject Key Ident
1246:d=9  hl=2  l=  22 prim: OCTET STRING     [HEX DUMP]:04144588CC9696
1270:d=8  hl=2  l=   9 cons: SEQUENCE
1272:d=9  hl=2  l=   3 prim: OBJECT            :X509v3 Basic Constraints
1277:d=9  hl=2  l=   2 prim: OCTET STRING     [HEX DUMP]:3000

```

```

1281:d=8 hl=2 l= 43 cons: SEQUENCE
1283:d=9 hl=2 l=  8 prim: OBJECT           :1.3.6.1.5.5.7.1.32
1293:d=9 hl=2 l= 31 prim: OCTET STRING     [HEX DUMP]:0C1D6869676877
1326:d=5 hl=2 l= 10 cons: SEQUENCE
1328:d=6 hl=2 l=  8 prim: OBJECT           :ecdsa-with-SHA256
1338:d=5 hl=2 l= 104 prim: BIT STRING
1444:d=3 hl=4 l= 315 cons: SET
1448:d=4 hl=4 l= 311 cons: SEQUENCE
1452:d=5 hl=2 l=  1 prim: INTEGER           :01
1455:d=5 hl=2 l= 101 cons: SEQUENCE
1457:d=6 hl=2 l= 93 cons: SEQUENCE
1459:d=7 hl=2 l= 15 cons: SET
1461:d=8 hl=2 l= 13 cons: SEQUENCE
1463:d=9 hl=2 l=  3 prim: OBJECT           :countryName
1468:d=9 hl=2 l=  6 prim: PRINTABLESTRING  :Canada
1476:d=7 hl=2 l= 16 cons: SET
1478:d=8 hl=2 l= 14 cons: SEQUENCE
1480:d=9 hl=2 l=  3 prim: OBJECT           :stateOrProvinceName
1485:d=9 hl=2 l=  7 prim: UTF8STRING       :Ontario
1494:d=7 hl=2 l= 18 cons: SET
1496:d=8 hl=2 l= 16 cons: SEQUENCE
1498:d=9 hl=2 l=  3 prim: OBJECT           :organizationalUnitName
1503:d=9 hl=2 l=  9 prim: UTF8STRING       :Sandelman
1514:d=7 hl=2 l= 36 cons: SET
1516:d=8 hl=2 l= 34 cons: SEQUENCE
1518:d=9 hl=2 l=  3 prim: OBJECT           :commonName
1523:d=9 hl=2 l= 27 prim: UTF8STRING       :highway-test.example.com
1552:d=6 hl=2 l=  4 prim: INTEGER           :0D85DC2D
1558:d=5 hl=2 l= 11 cons: SEQUENCE
1560:d=6 hl=2 l=  9 prim: OBJECT           :sha256
1571:d=5 hl=2 l= 105 cons: cont [ 0 ]
1573:d=6 hl=2 l= 24 cons: SEQUENCE
1575:d=7 hl=2 l=  9 prim: OBJECT           :contentType
1586:d=7 hl=2 l= 11 cons: SET
1588:d=8 hl=2 l=  9 prim: OBJECT           :pkcs7-data
1599:d=6 hl=2 l= 28 cons: SEQUENCE
1601:d=7 hl=2 l=  9 prim: OBJECT           :signingTime
1612:d=7 hl=2 l= 15 cons: SET
1614:d=8 hl=2 l= 13 prim: UTCTIME          :200225230448Z
1629:d=6 hl=2 l= 47 cons: SEQUENCE
1631:d=7 hl=2 l=  9 prim: OBJECT           :messageDigest
1642:d=7 hl=2 l= 34 cons: SET
1644:d=8 hl=2 l= 32 prim: OCTET STRING     [HEX DUMP]:B1E88AF0B2D1C5
1678:d=5 hl=2 l= 10 cons: SEQUENCE
1680:d=6 hl=2 l=  8 prim: OBJECT           :ecdsa-with-SHA256
1690:d=5 hl=2 l= 71 prim: OCTET STRING     [HEX DUMP]:304502201C7003

```

The JSON contained in the voucher request:



U2Eyb3ZXRFJEVfZGRE9ISk5Ua0p6VEc5T2NqRk1OVzVITlRabWQwRmtTVGhVYVVGWFJ6aFRPRmhC  
VWpWck1VIm5lRE5aV1ZGQ1UyZGtVMk5HWTBGa1ppc3JRbmMyV1hrc1ZUMGlmWDJnZ2dIcU1JSUI1  
akNDQVd5Z0F3SUJBZ0lFRFlYY0xUQUtCZ2dxaGtqT1BRUURBakJkTVE4d0RRWURWUVFHRXdaRFlX  
NWHaR0V4RURBT0JnTlZCQWdNQjA5dWRHRnlhVzh4RWRpBUUJnTlZCQXNNQ1ZOaGJtUmxiRzFoYmpF  
a01DSUdBMVVFQXd3YmFhbG5hSGRoZVMxMFPpYTjBmBY0WVcxd2JHVXVZMj10SUVOQk1DQVhEVEl3  
TURJd016QTJORG5TUZvWUR6STVPVgt4TWpNeE1EQXdNREF3V2pBY01Sb3dHQVlEVlFRRkRCRXdN  
QzFFtUMxRk5TMudNaTB3TUMwd01qQ1pNqk1HqnlxR1NNND1BZ0VHQ0Nxr1NNND1Bd0VIQTBjQUJB  
T2pkVU9IczN6QUNwcUhuSzMyOURnVEhOUUoQ19nUUU1QmVmQmJFMHNnY2JORW0wbGk4U1L3enJz  
QWZZQ3hwOWs3RTFDYnBpZwXUOE9XZjB6K01TZWpXVEJYTUlwR0ExVWREZ1FXQkJSRmlNeVdsZ0Jr  
TjdDnkkyVmtaRlFJQm14V3JUQUpCZ05WSFJNRUFqQUFNQ3NHQ0NzR0FRVUZCd0VnQkI4TUhXaHBA  
MmgzWVhrdGRHVnpkQzVsZUdGdGNHeGxMbU52Y1RvNU5EUXpNQW9HQ0Nxr1NNND1CQU1DQJnQU1H  
VUNNQ1BocVM3dk1oSTBXcVhDRmRZb3ZlMD1sdGJPQkpYdnA4amNHS2d4eDdnRU5QSzNUWG1LWnlJ  
a0EwL0Z6ZF1HdWdJeEFMT05YXJRL2dTRGtOT1BiWEtYc3o0QzZ2SElXakp5V0xkRkFsQjR2QVfK  
STE0aWI4T19qSHpYbTnBZ2tiVGhmekdQVRzdz2dnRtNBZ0VCTUdVd1hURVBNQTBHQTFVFRUJ0TUdR  
MkZ1WVdSaE1SQXdEZ1lEVlFRSURBZFBib1JoY21sdk1SSXdFQV1EVlFRTERBbFRZVzVrW1d4dFlX  
NHhKREFPqmdOVkBTU1HMmhWwjJoM1lYa3RkR1Z6ZEM1bGVHRnRjR3hsTG1OdmJTQkRRUU1FRFlY  
Y0xUQUxkZ2xnaGtnQ1pRTUUVBZ0dnYVRBWUJna3Foa2lHOXcwQkNRTXhdD1lKS29aSWH2Y05BUWNC  
TUJ3R0NTcUdTSWIZRFFFskJURVbGdzB5TURBeU1qVX1NekEwTkRoYU1DOEdDU3FHU0liMORRRUpC  
REVPqkNDeDZJcndzdEhGNjA5WTBfCURLNjJRS2J5NGR1eXlJV3VkdNMXNU0xNkJCVEFLQmdncWhr  
ak9QUVFEQWdSSE1FVUNJQnh3QTFVbGtJa3VRRGYvaJdrWi9NVmVmZ3IxNDEraEtCRmdybk5uZ2p3  
cEFpRUF50GFYdDBHU0I5bTFibW1FVXB1ZkNFaHhTdjJ4TF1lckdsdWd2MGRmci9FPSP9faCCBG8w  
ggH8MIIBGqADAgECAgQ/mJtSMAoGCCqGSM49BAMCMG0xJQAQBoJkiaJk/IsZAEZFgJjYTEZMBcG  
CgmSjomT8ixkARKWCXNhbMR1bG1hbje8MDoGA1UEAwZm91bnRhaW4tdGVzdC5leGftcGx1LmNv  
bSBVbnN0cnVuZyBGB3VudGFpbiBSb290IENBMB4XDTEwMDIyNTIxMzE1NFoXDTEyMDIyNDIxMzE1  
NFowUzESMBAGCgmSjomT8ixkARKWAmNhMRkWFwYKcZImiZPyLgQBGRYJc2FuzGVsbWVfMTwwOgYD  
VQDDb1mb3VudGFpbi10ZXN0LmV4YW1wbGUuY29tMFkwEwYHKoZIzj0CAQYIKoZIzj0DAQcDQgAE  
lmVQcjs6n+Xd5l/28IFv6UiegQwSBztGj5dkK2MAjQIPV8181H+EjLlIOYdbJiI0VtEiF1/Jqt+TO  
BfinTNOLOqMqMCgWfGyDVR0LAQH/BAwwCgYIKwYBBQUHAXwDgYDVR0PAQH/BAQDAgeAMAoGCCqG  
SM49BAMCA2gAMGUCMGPYExVSB6WB/jdH7nIEo1FNoeI8C8u/HLPsYVVM9fH7/VHA5qCa8bdpeZ  
GSP9fgIXALysw0Gwug2vUvmcbnp/AB0jyGIBYbxLxcBHmTUKDHdhRAFKB1JwVwB1/74HDpjL5TCC  
AmswggHyOAMCAQICBC1rBlkwCgYIKoZIzj0EAwIwbTESMBAGCgmSjomT8ixkARKWAmNhMRkWFwYK  
CZImiZPyLgQBGRYJc2FuzGVsbWVfMTwwOgYDQVQDDNmb3VudGFpbi10ZXN0LmV4YW1wbGUuY29t  
IFVuc3Rydw5nIEZvdW50YW1uIFJvb3QgQ0EwHhcNMjAwMjI1MjEzMTQ1WhcNMjAwMjI1MjEzMTQ1  
WjBtMRIwEAYKcZImiZPyLgQBGRYCY2ExGTAXBgoJkiaJk/IsZAEZFglzYW5kZWxtYW4xPDA6BgNV  
BAMM2ZvdW50YW1uLXRlc3QuZlZhbG9uZS5jb20gVW5zdHJ1bmcgRm91bnRhaW4gUm9vdCBDQTB2  
MBAGByqGSM49AgEGBSuBBAAiA2IABBT/WboXwxq8Zo2MbODD+jfxD2X2Ipg9t1aAB9vfuHq1RU15  
ikaXGVMWmBGPax0yvjzIPltjtUb2qNVvm/nA8905FD9yR1Gkdt3S8L/1yo8wAX/4w1/T9SADRIuL  
8gdstKNjMGEwDwYDVR0TAQH/BAUwAwEB/zAOBgNVHQ8BAf8EBAMCAQYwHQYDVR0OBBYEFm19ssR  
4QekSSynCMZ8ELyHs3QmMB8GA1UdIwQYMBaAFm19ssR4QekSSynCMZ8ELyHs3QmMAoGCCqGSM49  
BAMCA2cAMGQCMCCDBs6NmKRUMZMSjppw1I2W1KNWX0gmyppFFiHONhVed39KtiVHpgTdrT1Zi1E  
tAIwbzj5rxLtMNWFKXyXfli9Z5FDxA0w+dgcrC8G3bzVBkIshKIE6gKkXxdRjvvZL9JcMYIBSzcC  
AUcCAQEwdTBTMRIwEAYKcZImiZPyLgQBGRYCY2ExGTAXBgoJkiaJk/IsZAEZFglzYW5kZWxtYW4x  
PDA6BgNVBAMM2ZvdW50YW1uLXRlc3QuZlZhbG9uZS5jb20gVW5zdHJ1bmcgRm91bnRhaW4gUm9v  
dCBDQQIEP5ibUjALBg1ghkgBZQMEAgGgATAYBgkqhkiG9w0BCQMxCwYJKoZIhvcNAQcCBMwGCCqG  
SIb3DQEJBTBEPf0yMDAyMjUyMzA0NDlaMC8GCScqGSIB3DQEJBDiBCA9gYXR1sS0giII3PwvOK/N  
5RUBwjSL/cDcrH/Bd+E1ajAKBggqhkiJOPQDdAgRHMEUCIFieXZa07P9eZMPCVn2laB4czw7I0s0P  
s9+frcJtEBTTAiEAhCcB//qmgqcEA+90mquvVNENmFH9dxCH8Ihhz6SCVDI=  
<CODE ENDS>

The ASN1 decoding of the artifact:

file: examples/parboiled\_vr\_00\_D0-E5-02-00-2D.b64

```

0:d=0 hl=4 l=4087 cons: SEQUENCE
4:d=1 hl=2 l= 9 prim: OBJECT           :pkcs7-signedData
15:d=1 hl=4 l=4072 cons: cont [ 0 ]
19:d=2 hl=4 l=4068 cons: SEQUENCE
23:d=3 hl=2 l= 1 prim: INTEGER         :01
26:d=3 hl=2 l= 13 cons: SET
28:d=4 hl=2 l= 11 cons: SEQUENCE
30:d=5 hl=2 l= 9 prim: OBJECT          :sha256
41:d=3 hl=4 l=2572 cons: SEQUENCE
45:d=4 hl=2 l= 9 prim: OBJECT          :pkcs7-data
56:d=4 hl=4 l=2557 cons: cont [ 0 ]
60:d=5 hl=4 l=2553 prim: OCTET STRING  :{"ietf-voucher-request:v
2617:d=3 hl=4 l=1135 cons: cont [ 0 ]
2621:d=4 hl=4 l= 508 cons: SEQUENCE
2625:d=5 hl=4 l= 386 cons: SEQUENCE
2629:d=6 hl=2 l= 3 cons: cont [ 0 ]
2631:d=7 hl=2 l= 1 prim: INTEGER       :02
2634:d=6 hl=2 l= 4 prim: INTEGER       :3F989B52
2640:d=6 hl=2 l= 10 cons: SEQUENCE
2642:d=7 hl=2 l= 8 prim: OBJECT        :ecdsa-with-SHA256
2652:d=6 hl=2 l= 109 cons: SEQUENCE
2654:d=7 hl=2 l= 18 cons: SET
2656:d=8 hl=2 l= 16 cons: SEQUENCE
2658:d=9 hl=2 l= 10 prim: OBJECT       :domainComponent
2670:d=9 hl=2 l= 2 prim: IA5STRING     :ca
2674:d=7 hl=2 l= 25 cons: SET
2676:d=8 hl=2 l= 23 cons: SEQUENCE
2678:d=9 hl=2 l= 10 prim: OBJECT       :domainComponent
2690:d=9 hl=2 l= 9 prim: IA5STRING     :sandelman
2701:d=7 hl=2 l= 60 cons: SET
2703:d=8 hl=2 l= 58 cons: SEQUENCE
2705:d=9 hl=2 l= 3 prim: OBJECT        :commonName
2710:d=9 hl=2 l= 51 prim: UTF8STRING   :fountain-test.example.co
2763:d=6 hl=2 l= 30 cons: SEQUENCE
2765:d=7 hl=2 l= 13 prim: UTCTIME      :200225213154Z
2780:d=7 hl=2 l= 13 prim: UTCTIME      :220224213154Z
2795:d=6 hl=2 l= 83 cons: SEQUENCE
2797:d=7 hl=2 l= 18 cons: SET
2799:d=8 hl=2 l= 16 cons: SEQUENCE
2801:d=9 hl=2 l= 10 prim: OBJECT       :domainComponent
2813:d=9 hl=2 l= 2 prim: IA5STRING     :ca
2817:d=7 hl=2 l= 25 cons: SET
2819:d=8 hl=2 l= 23 cons: SEQUENCE
2821:d=9 hl=2 l= 10 prim: OBJECT       :domainComponent

```

```

2833:d=9 hl=2 l= 9 prim: IA5STRING :sandelman
2844:d=7 hl=2 l= 34 cons: SET
2846:d=8 hl=2 l= 32 cons: SEQUENCE
2848:d=9 hl=2 l= 3 prim: OBJECT :commonName
2853:d=9 hl=2 l= 25 prim: UTF8STRING :fountain-test.example.co
2880:d=6 hl=2 l= 89 cons: SEQUENCE
2882:d=7 hl=2 l= 19 cons: SEQUENCE
2884:d=8 hl=2 l= 7 prim: OBJECT :id-ecPublicKey
2893:d=8 hl=2 l= 8 prim: OBJECT :prime256v1
2903:d=7 hl=2 l= 66 prim: BIT STRING
2971:d=6 hl=2 l= 42 cons: cont [ 3 ]
2973:d=7 hl=2 l= 40 cons: SEQUENCE
2975:d=8 hl=2 l= 22 cons: SEQUENCE
2977:d=9 hl=2 l= 3 prim: OBJECT :X509v3 Extended Key Usag
2982:d=9 hl=2 l= 1 prim: BOOLEAN :255
2985:d=9 hl=2 l= 12 prim: OCTET STRING [HEX DUMP]:300A06082B0601
2999:d=8 hl=2 l= 14 cons: SEQUENCE
3001:d=9 hl=2 l= 3 prim: OBJECT :X509v3 Key Usage
3006:d=9 hl=2 l= 1 prim: BOOLEAN :255
3009:d=9 hl=2 l= 4 prim: OCTET STRING [HEX DUMP]:03020780
3015:d=5 hl=2 l= 10 cons: SEQUENCE
3017:d=6 hl=2 l= 8 prim: OBJECT :ecdsa-with-SHA256
3027:d=5 hl=2 l= 104 prim: BIT STRING
3133:d=4 hl=4 l= 619 cons: SEQUENCE
3137:d=5 hl=4 l= 498 cons: SEQUENCE
3141:d=6 hl=2 l= 3 cons: cont [ 0 ]
3143:d=7 hl=2 l= 1 prim: INTEGER :02
3146:d=6 hl=2 l= 4 prim: INTEGER :296B0659
3152:d=6 hl=2 l= 10 cons: SEQUENCE
3154:d=7 hl=2 l= 8 prim: OBJECT :ecdsa-with-SHA256
3164:d=6 hl=2 l= 109 cons: SEQUENCE
3166:d=7 hl=2 l= 18 cons: SET
3168:d=8 hl=2 l= 16 cons: SEQUENCE
3170:d=9 hl=2 l= 10 prim: OBJECT :domainComponent
3182:d=9 hl=2 l= 2 prim: IA5STRING :ca
3186:d=7 hl=2 l= 25 cons: SET
3188:d=8 hl=2 l= 23 cons: SEQUENCE
3190:d=9 hl=2 l= 10 prim: OBJECT :domainComponent
3202:d=9 hl=2 l= 9 prim: IA5STRING :sandelman
3213:d=7 hl=2 l= 60 cons: SET
3215:d=8 hl=2 l= 58 cons: SEQUENCE
3217:d=9 hl=2 l= 3 prim: OBJECT :commonName
3222:d=9 hl=2 l= 51 prim: UTF8STRING :fountain-test.example.co
3275:d=6 hl=2 l= 30 cons: SEQUENCE
3277:d=7 hl=2 l= 13 prim: UTCTIME :200225213145Z
3292:d=7 hl=2 l= 13 prim: UTCTIME :220224213145Z
3307:d=6 hl=2 l= 109 cons: SEQUENCE
3309:d=7 hl=2 l= 18 cons: SET

```

```

3311:d=8  hl=2 l= 16 cons: SEQUENCE
3313:d=9  hl=2 l= 10 prim: OBJECT           :domainComponent
3325:d=9  hl=2 l=  2 prim: IA5STRING         :ca
3329:d=7  hl=2 l= 25 cons: SET
3331:d=8  hl=2 l= 23 cons: SEQUENCE
3333:d=9  hl=2 l= 10 prim: OBJECT           :domainComponent
3345:d=9  hl=2 l=  9 prim: IA5STRING         :sandelman
3356:d=7  hl=2 l= 60 cons: SET
3358:d=8  hl=2 l= 58 cons: SEQUENCE
3360:d=9  hl=2 l=  3 prim: OBJECT           :commonName
3365:d=9  hl=2 l= 51 prim: UTF8STRING        :fountain-test.example.co
3418:d=6  hl=2 l= 118 cons: SEQUENCE
3420:d=7  hl=2 l= 16 cons: SEQUENCE
3422:d=8  hl=2 l=  7 prim: OBJECT           :id-ecPublicKey
3431:d=8  hl=2 l=  5 prim: OBJECT           :secp384r1
3438:d=7  hl=2 l= 98 prim: BIT STRING
3538:d=6  hl=2 l= 99 cons: cont [ 3 ]
3540:d=7  hl=2 l= 97 cons: SEQUENCE
3542:d=8  hl=2 l= 15 cons: SEQUENCE
3544:d=9  hl=2 l=  3 prim: OBJECT           :X509v3 Basic Constraints
3549:d=9  hl=2 l=  1 prim: BOOLEAN            :255
3552:d=9  hl=2 l=  5 prim: OCTET STRING      [HEX DUMP]:30030101FF
3559:d=8  hl=2 l= 14 cons: SEQUENCE
3561:d=9  hl=2 l=  3 prim: OBJECT           :X509v3 Key Usage
3566:d=9  hl=2 l=  1 prim: BOOLEAN            :255
3569:d=9  hl=2 l=  4 prim: OCTET STRING      [HEX DUMP]:03020106
3575:d=8  hl=2 l= 29 cons: SEQUENCE
3577:d=9  hl=2 l=  3 prim: OBJECT           :X509v3 Subject Key Ident
3582:d=9  hl=2 l= 22 prim: OCTET STRING      [HEX DUMP]:0414B9A5F6CB11
3606:d=8  hl=2 l= 31 cons: SEQUENCE
3608:d=9  hl=2 l=  3 prim: OBJECT           :X509v3 Authority Key Ide
3613:d=9  hl=2 l= 24 prim: OCTET STRING      [HEX DUMP]:30168014B9A5F6
3639:d=5  hl=2 l= 10 cons: SEQUENCE
3641:d=6  hl=2 l=  8 prim: OBJECT           :ecdsa-with-SHA256
3651:d=5  hl=2 l= 103 prim: BIT STRING
3756:d=3  hl=4 l= 331 cons: SET
3760:d=4  hl=4 l= 327 cons: SEQUENCE
3764:d=5  hl=2 l=  1 prim: INTEGER            :01
3767:d=5  hl=2 l= 117 cons: SEQUENCE
3769:d=6  hl=2 l= 109 cons: SEQUENCE
3771:d=7  hl=2 l= 18 cons: SET
3773:d=8  hl=2 l= 16 cons: SEQUENCE
3775:d=9  hl=2 l= 10 prim: OBJECT           :domainComponent
3787:d=9  hl=2 l=  2 prim: IA5STRING         :ca
3791:d=7  hl=2 l= 25 cons: SET
3793:d=8  hl=2 l= 23 cons: SEQUENCE
3795:d=9  hl=2 l= 10 prim: OBJECT           :domainComponent
3807:d=9  hl=2 l=  9 prim: IA5STRING         :sandelman

```

```
3818:d=7 hl=2 l= 60 cons: SET
3820:d=8 hl=2 l= 58 cons: SEQUENCE
3822:d=9 hl=2 l=  3 prim: OBJECT           :commonName
3827:d=9 hl=2 l= 51 prim: UTF8STRING      :fountain-test.example.co
3880:d=6 hl=2 l=  4 prim: INTEGER         :3F989B52
3886:d=5 hl=2 l= 11 cons: SEQUENCE
3888:d=6 hl=2 l=  9 prim: OBJECT           :sha256
3899:d=5 hl=2 l= 105 cons: cont [ 0 ]
3901:d=6 hl=2 l= 24 cons: SEQUENCE
3903:d=7 hl=2 l=  9 prim: OBJECT           :contentType
3914:d=7 hl=2 l= 11 cons: SET
3916:d=8 hl=2 l=  9 prim: OBJECT           :pkcs7-data
3927:d=6 hl=2 l= 28 cons: SEQUENCE
3929:d=7 hl=2 l=  9 prim: OBJECT           :signingTime
3940:d=7 hl=2 l= 15 cons: SET
3942:d=8 hl=2 l= 13 prim: UTCTIME         :200225230449Z
3957:d=6 hl=2 l= 47 cons: SEQUENCE
3959:d=7 hl=2 l=  9 prim: OBJECT           :messageDigest
3970:d=7 hl=2 l= 34 cons: SET
3972:d=8 hl=2 l= 32 prim: OCTET STRING    [HEX DUMP]:3D818C51D6C4B4
4006:d=5 hl=2 l= 10 cons: SEQUENCE
4008:d=6 hl=2 l=  8 prim: OBJECT           :ecdsa-with-SHA256
4018:d=5 hl=2 l= 71 prim: OCTET STRING    [HEX DUMP]:30450220589E5D
```

The JSON contained in the voucher request. Note that the previous voucher request is in the prior-signed-voucher-request attribute.

```

{"ietf-voucher-request:voucher":{"assertion":"proximity","created-on":"2020-02-25T23:04:49.054Z","serial-number":"00-D0-E5-F2-00-02","nonce":"aMjgueKUT-22wVimj6z27Q","prior-signed-voucher-request":"MIIG3wYJKoZIHvcNAQcCoIIG0DCCBswCAQExDTALBg1ghkgBZQMEAgEwggOJBgkqhkiG9w0BBwGgggN6BIIDdnsiaWV0Zi12b3VjaGVyLXJlcXVlc3Q6dm91Y2hlcilEeyJhc3NlcnRpb24iOiJwcm94aW1pdHkiLCJjcmVhdGvklW9uIjoimjAyMC0wMi0yNVQxODowND00C42NTItMDU6MDAiLCJzZXJpYWwtbnVtYmVyIjoimDAtRDAtRTUtRjItMDAtMDIiLCJub25jZSI6ImFNamd1ZUtVVC0yMndWaW1qNnoyN1EiLCJwcm94aW1pdHktcmVnaXN0cmFyLWV1cnQiOiJNSU1CL0RDQ0FZS2dBd0lCQWdJRVA1aWJVakFLQmdncWhrak9QUVFEQWpCdE1SSXdfQV1LQ1pJbWlaUH1MR1FCR1JZQ1kyRXhHVEFYQmdvSmtpYUprL0lZwKFFWkZnbHpvZVzVrWld4dFlXNHhQREE2QmdOVk1JBTU1N1p2ZFc1MFlXbHVmWFJsYzNRdVpYaGhiWEJzW1M1amIyMGdWVzV6ZEhKMWJtY2dSbTkkxYm5SaGFxNGdVbT12ZENCRFFUQWVGdzB5TURBeU1qVXlNVE14TlRSYU13MH1NakF5TWpReU1UTXh0VFJhTUZNeEVqQVFCZ29Ka2lhSmsvSXNaQUVaRmdKallURVpNQmNHQ2dtU0pvbVQ4aXhrQVJrV0NYTmhibVJsYkcxagJqRW1NQ0FHQTFVRU1F3d1pabTkkxYm5SaGFxNHRkR1Z6ZEM1bGVHRnRjR3hsTG1OdmJUQ1pNqk1HQnlxR1NNND1BZ0VHQ0Nxr1NNND1Bd0VIQTBJQUJKWmxVSEkwdXAvbDN1WmY5dkNCYitsSW5vRU1FZ2M3Um8rWFpDdGpBSTBDRDFmSmZKU19oSX15RG1IV3lZaU5GYlJDSdlmeWFyZmt6Z1g0cDB6VG16cWpLakFvTUJZR0ExVWRKUUVCL3dRTU1Bb0dDQ3NHQVFVRk1J3TWNNQTRHQTfVZER3RU1vd1FFQXdxJSGdEQUtCZ2dxag1BRUURBZ05vQURCbEFqQm1Umk1JNV1VnZWxnZjQzUis1eUJLTlJUyUhteVBBdkx2eH16MG1GV1p2WHgrLzFSd09hZ212RzNhWG1Sa2ovWDRDTRVFDHJNTk1JzTG90c1FjMmNw5HNTZmd0FkSThoaUFXRzhtOFhBUjVrMUNneDNzVVFcu2dkU2NGY0FkZisrQnc2WXkrVT0ifX2gggHqMIIB5jCCAWygAwIBAgIEDYXcLTAKBgqghk1jOPQQDAjBdMQ8wDQYDVQQGEWZDYW5hZGExEDAOBgNVBAgMB09udGFyaW8xZjAqBGNVBAAsMVCVhbmRlbG1hb1JkMCIgA1UEAwwbaGlnaHdheS10ZXN0LmV4YW1wbGUuY29tIENBMCAXDTEwMDIwMzA2NDcyMFoYDzI5OTkxMjMxMDAwMDAwWjAcMR0wGAYDVQQFDBEwMC1EMC1FN1GMi0wMC0wMjBZMBMGByqGSM49AgEGCCqGSM49AwEHA0IABAOjdUOHs3zACpqHnK329DgTHNQmhb/gQE5BeFBbE0sgcbNEm0li8RYwzrsAfYcXp9k7E1CbpielT8OWf0z+ISejWTBxMB0GA1UdDgQWBBERFiMyW1gBkN7C6I2VkJZFIbmXWrtAJBgNVHRMEAjAAMCsGCCsGAQUFBwEgBB8MHWhpZ2h3YXktZGVzZC5leGFtcGx1LmNvbTo5NDQzMAoGCCqGSM49BAMCA2gAMGUcmCPhqs7vIhI0WqXCFdYove091tbOBjXvp8jcGKgx7gENPK3TXmKZyIka0/FzdYGugIxaLONXArQ/gSDkNNPbXKXsz4C6vHIWjJyWldFA1B4vAQdI14ib8N/jHzXm3AgkbThfzGCATswggE3AgEBMGUwXTEPMA0GA1UEBhMGQ2FuYWRhMRAdGyYDVQIDAdPbnRhcmlvMRIwEAYDVQQLDAlTYW5kZWxtYW4xJDAlBgNVBAMMG2hpZ2h3YXktZGVzZC5leGFtcGx1LmNvbSBdQQIEDYXcLTALBg1ghkgBZQMEAgGgaTAYBgkqhkiG9w0BCQMxCwYJKoZIhvcNAQcBMBwGCSqGSIb3DQEJBTEPFw0yMDAyMjUyMzA0NDhaMC8GCSqGSIb3DQEJBDEiBCCx6IrwstHF609Y0EqDK62QKby4duyyIWudvs15M16BBTAKBggqhk1jOPQQDAgRHMEUCIBxwA1UlkIkuQDf/j7kZ/MVefgr141+hKBFgrnNngjwpAiEAY8aXt0GSB9mlbm1EUpefCEhxSv2xLYurGlugv0dfr/E="}}

```

C.2.3. MASA to Registrar

The MASA will return a voucher to the registrar, to be relayed to the pledge.



```

56:d=4  hl=4  l= 873 cons: cont [ 0 ]
60:d=5  hl=4  l= 869 prim: OCTET STRING      :{"ietf-voucher:voucher":
933:d=3  hl=4  l= 483 cons: cont [ 0 ]
937:d=4  hl=4  l= 479 cons: SEQUENCE
941:d=5  hl=4  l= 356 cons: SEQUENCE
945:d=6  hl=2  l=   3 cons: cont [ 0 ]
947:d=7  hl=2  l=   1 prim: INTEGER           :02
950:d=6  hl=2  l=   4 prim: INTEGER           :1B995F54
956:d=6  hl=2  l=  10 cons: SEQUENCE
958:d=7  hl=2  l=   8 prim: OBJECT            :ecdsa-with-SHA256
968:d=6  hl=2  l=  93 cons: SEQUENCE
970:d=7  hl=2  l=  15 cons: SET
972:d=8  hl=2  l=  13 cons: SEQUENCE
974:d=9  hl=2  l=   3 prim: OBJECT            :countryName
979:d=9  hl=2  l=   6 prim: PRINTABLESTRING  :Canada
987:d=7  hl=2  l=  16 cons: SET
989:d=8  hl=2  l=  14 cons: SEQUENCE
991:d=9  hl=2  l=   3 prim: OBJECT            :stateOrProvinceName
996:d=9  hl=2  l=   7 prim: UTF8STRING       :Ontario
1005:d=7  hl=2  l=  18 cons: SET
1007:d=8  hl=2  l=  16 cons: SEQUENCE
1009:d=9  hl=2  l=   3 prim: OBJECT            :organizationalUnitName
1014:d=9  hl=2  l=   9 prim: UTF8STRING       :Sandelman
1025:d=7  hl=2  l=  36 cons: SET
1027:d=8  hl=2  l=  34 cons: SEQUENCE
1029:d=9  hl=2  l=   3 prim: OBJECT            :commonName
1034:d=9  hl=2  l=  27 prim: UTF8STRING       :highway-test.example.com
1063:d=6  hl=2  l=  30 cons: SEQUENCE
1065:d=7  hl=2  l=  13 prim: UTCTIME           :190212222241Z
1080:d=7  hl=2  l=  13 prim: UTCTIME           :210211222241Z
1095:d=6  hl=2  l=  95 cons: SEQUENCE
1097:d=7  hl=2  l=  15 cons: SET
1099:d=8  hl=2  l=  13 cons: SEQUENCE
1101:d=9  hl=2  l=   3 prim: OBJECT            :countryName
1106:d=9  hl=2  l=   6 prim: PRINTABLESTRING  :Canada
1114:d=7  hl=2  l=  16 cons: SET
1116:d=8  hl=2  l=  14 cons: SEQUENCE
1118:d=9  hl=2  l=   3 prim: OBJECT            :stateOrProvinceName
1123:d=9  hl=2  l=   7 prim: UTF8STRING       :Ontario
1132:d=7  hl=2  l=  18 cons: SET
1134:d=8  hl=2  l=  16 cons: SEQUENCE
1136:d=9  hl=2  l=   3 prim: OBJECT            :organizationalUnitName
1141:d=9  hl=2  l=   9 prim: UTF8STRING       :Sandelman
1152:d=7  hl=2  l=  38 cons: SET
1154:d=8  hl=2  l=  36 cons: SEQUENCE
1156:d=9  hl=2  l=   3 prim: OBJECT            :commonName
1161:d=9  hl=2  l=  29 prim: UTF8STRING       :highway-test.example.com
1192:d=6  hl=2  l=  89 cons: SEQUENCE

```

```

1194:d=7 hl=2 l= 19 cons: SEQUENCE
1196:d=8 hl=2 l= 7 prim: OBJECT           :id-ecPublicKey
1205:d=8 hl=2 l= 8 prim: OBJECT           :prime256v1
1215:d=7 hl=2 l= 66 prim: BIT STRING
1283:d=6 hl=2 l= 16 cons: cont [ 3 ]
1285:d=7 hl=2 l= 14 cons: SEQUENCE
1287:d=8 hl=2 l= 12 cons: SEQUENCE
1289:d=9 hl=2 l= 3 prim: OBJECT           :X509v3 Basic Constraints
1294:d=9 hl=2 l= 1 prim: BOOLEAN         :255
1297:d=9 hl=2 l= 2 prim: OCTET STRING    [HEX DUMP]:3000
1301:d=5 hl=2 l= 10 cons: SEQUENCE
1303:d=6 hl=2 l= 8 prim: OBJECT           :ecdsa-with-SHA256
1313:d=5 hl=2 l= 105 prim: BIT STRING
1420:d=3 hl=4 l= 315 cons: SET
1424:d=4 hl=4 l= 311 cons: SEQUENCE
1428:d=5 hl=2 l= 1 prim: INTEGER         :01
1431:d=5 hl=2 l= 101 cons: SEQUENCE
1433:d=6 hl=2 l= 93 cons: SEQUENCE
1435:d=7 hl=2 l= 15 cons: SET
1437:d=8 hl=2 l= 13 cons: SEQUENCE
1439:d=9 hl=2 l= 3 prim: OBJECT           :countryName
1444:d=9 hl=2 l= 6 prim: PRINTABLESTRING :Canada
1452:d=7 hl=2 l= 16 cons: SET
1454:d=8 hl=2 l= 14 cons: SEQUENCE
1456:d=9 hl=2 l= 3 prim: OBJECT           :stateOrProvinceName
1461:d=9 hl=2 l= 7 prim: UTF8STRING      :Ontario
1470:d=7 hl=2 l= 18 cons: SET
1472:d=8 hl=2 l= 16 cons: SEQUENCE
1474:d=9 hl=2 l= 3 prim: OBJECT           :organizationalUnitName
1479:d=9 hl=2 l= 9 prim: UTF8STRING      :Sandelman
1490:d=7 hl=2 l= 36 cons: SET
1492:d=8 hl=2 l= 34 cons: SEQUENCE
1494:d=9 hl=2 l= 3 prim: OBJECT           :commonName
1499:d=9 hl=2 l= 27 prim: UTF8STRING      :highway-test.example.com
1528:d=6 hl=2 l= 4 prim: INTEGER         :1B995F54
1534:d=5 hl=2 l= 11 cons: SEQUENCE
1536:d=6 hl=2 l= 9 prim: OBJECT           :sha256
1547:d=5 hl=2 l= 105 cons: cont [ 0 ]
1549:d=6 hl=2 l= 24 cons: SEQUENCE
1551:d=7 hl=2 l= 9 prim: OBJECT           :contentType
1562:d=7 hl=2 l= 11 cons: SET
1564:d=8 hl=2 l= 9 prim: OBJECT           :pkcs7-data
1575:d=6 hl=2 l= 28 cons: SEQUENCE
1577:d=7 hl=2 l= 9 prim: OBJECT           :signingTime
1588:d=7 hl=2 l= 15 cons: SET
1590:d=8 hl=2 l= 13 prim: UTCTIME         :200225230449Z
1605:d=6 hl=2 l= 47 cons: SEQUENCE
1607:d=7 hl=2 l= 9 prim: OBJECT           :messageDigest

```

1618:d=7 hl=2 l= 34 cons: SET  
1620:d=8 hl=2 l= 32 prim: OCTET STRING [HEX DUMP]:8942CA3867D9AC  
1654:d=5 hl=2 l= 10 cons: SEQUENCE  
1656:d=6 hl=2 l= 8 prim: OBJECT :ecdsa-with-SHA256  
1666:d=5 hl=2 l= 71 prim: OCTET STRING [HEX DUMP]:30450221008A11

#### Appendix D. Additional References

RFC EDITOR Please remove this section before publication. It exists just to include references to the things in the YANG descriptions which are not otherwise referenced in the text so that xml2rfc will not complain.

[ITU.X690.1994]

#### Authors' Addresses

Max Pritikin  
Cisco

Email: [pritikin@cisco.com](mailto:pritikin@cisco.com)

Michael C. Richardson  
Sandelman Software Works

Email: [mcr+iETF@sandelman.ca](mailto:mcr+iETF@sandelman.ca)  
URI: <http://www.sandelman.ca/>

Toerless Eckert  
Futurewei Technologies Inc. USA  
2330 Central Expy  
Santa Clara, CA 95050  
United States of America

Email: [tte+iETF@cs.fau.de](mailto:tte+iETF@cs.fau.de)

Michael H. Behringer

Email: [Michael.H.Behringer@gmail.com](mailto:Michael.H.Behringer@gmail.com)

Kent Watsen  
Watsen Networks

Email: [kent+iETF@watsen.net](mailto:kent+iETF@watsen.net)

anima Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: May 6, 2021

M. Richardson  
Sandelman Software Works  
P. van der Stok  
vanderstok consultancy  
P. Kampanakis  
Cisco Systems  
November 02, 2020

Constrained Voucher Artifacts for Bootstrapping Protocols  
draft-ietf-anima-constrained-voucher-09

Abstract

This document defines a strategy to securely assign a pledge to an owner, using an artifact signed, directly or indirectly, by the pledge's manufacturer. This artifact is known as a "voucher".

This document builds upon the work in [RFC8366], encoding the resulting artifact in CBOR. Use with two signature technologies are described.

Additionally, this document explains how constrained vouchers may be transported as an extension to the [I-D.ietf-ace-coap-est] protocol.

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 <https://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 May 6, 2021.

Copyright Notice

Copyright (c) 2020 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 (<https://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 . . . . .	3
2. Terminology . . . . .	4
3. Requirements Language . . . . .	4
4. Survey of Voucher Types . . . . .	4
5. Discovery and URI . . . . .	5
6. Artifacts . . . . .	7
6.1. Voucher Request artifact . . . . .	7
6.1.1. Tree Diagram . . . . .	7
6.1.2. SID values . . . . .	8
6.1.3. YANG Module . . . . .	9
6.1.4. Example voucher request artifact . . . . .	13
6.2. Voucher artifact . . . . .	13
6.2.1. Tree Diagram . . . . .	14
6.2.2. SID values . . . . .	14
6.2.3. YANG Module . . . . .	15
6.2.4. Example voucher artifacts . . . . .	17
6.3. Signing voucher and voucher-request artifacts . . . . .	18
6.3.1. CMS signing . . . . .	18
6.3.2. COSE signing . . . . .	19
7. Design Considerations . . . . .	20
8. Security Considerations . . . . .	20
8.1. Clock Sensitivity . . . . .	20
8.2. Protect Voucher PKI in HSM . . . . .	20
8.3. Test Domain Certificate Validity when Signing . . . . .	20
9. IANA Considerations . . . . .	20
9.1. Resource Type Registry . . . . .	20
9.2. The IETF XML Registry . . . . .	21
9.3. The YANG Module Names Registry . . . . .	21
9.4. The RFC SID range assignment sub-registry . . . . .	21
9.5. The SMI Security for S/MIME CMS Content Type Registry . . . . .	22
9.6. Media-Type Registry . . . . .	22
9.6.1. application/voucher-cms+cbor . . . . .	22
9.6.2. application/voucher-cose+cbor . . . . .	22
9.7. CoAP Content-Format Registry . . . . .	23
10. Acknowledgements . . . . .	23
11. Changelog . . . . .	24

12. References	24
12.1. Normative References	24
12.2. Informative References	26
Appendix A. EST messages to EST-coaps	26
A.1. enrollstatus	26
A.2. voucher_status	28
Appendix B. CMS signed messages	28
B.1. signed requestvoucher	28
B.2. requestauditing	30
B.3. CMS signed voucher-request example	31
Appendix C. COSE examples	34
C.1. Pledge, Registrar and MASA keys	38
C.1.1. Pledge private key	38
C.1.2. Registrar private key	38
C.1.3. MASA private key	39
C.2. Pledge, Registrar and MASA certificates	39
C.2.1. Pledge IDevID certificate	39
C.2.2. Registrar Certificate	41
C.2.3. MASA Certificate	43
C.3. COSE signed voucher request from pledge to Registrar	45
C.4. COSE signed voucher request from Registrar to MASA	47
C.5. COSE signed voucher from MASA to Pledge via Registrar	49
Authors' Addresses	51

## 1. Introduction

Enrollment of new nodes into constrained networks with constrained nodes present unique challenges.

There are bandwidth and code space issues to contend. A solution such as [I-D.ietf-anima-bootstrapping-keyinfra] may be too large in terms of code space or bandwidth required.

This document defines a constrained version of [RFC8366]. Rather than serializing the YANG definition in JSON, it is serialized into CBOR ([RFC7049]).

This document follows a similar, but not identical structure as [RFC8366] and supplements the brski part to [I-D.ietf-ace-coap-est].

There are three constrained situations described in this document: 1. CMS signed CBOR encoded vouchers transported using CoAP, protected by DTLS (coaps). 2. COSE signed CBOR encoded vouchers transported using CoAP, protected by EDHOC or DTLS. 3. COSE signed CBOR encoded vouchers, integrated into the key exchange as described by [I-D.selander-ace-ake-authz]

Additional sections have been added concerning:

1. Addition of voucher-request specification as defined in [I-D.ietf-anima-bootstrapping-keyinfra],
2. Addition to [I-D.ietf-ace-coap-est] of voucher transport requests over CoAP.

The CBOR definitions for this constrained voucher format are defined using the mechanism describe in [I-D.ietf-core-yang-cbor] using the SID mechanism explained in [I-D.ietf-core-sid]. As the tooling to convert YANG documents into an list of SID keys is still in its infancy, the table of SID values presented here should be considered normative rather than the output of the pyang tool.

Two methods of signing the resulting CBOR object are described in this document:

1. One is CMS [RFC5652].
2. The other is COSE\_Sign1 [RFC8152] objects.

## 2. Terminology

The following terms are defined in [RFC8366], and are used identically as in that document: artifact, imprint, domain, Join Registrar/Coordinator (JRC), Manufacturer Authorized Signing Authority (MASA), pledge, Trust of First Use (TOFU), and Voucher.

## 3. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

## 4. Survey of Voucher Types

[RFC8366] provides for vouchers that assert proximity, that authenticate the registrar and that include different amounts of anti-replay protection.

This document does not make any extensions to the types of vouchers.

Time based vouchers are included in this definition, but given that constrained devices are extremely unlikely to have accurate time, their use is very unlikely. Most users of these constrained vouchers will be online and will use live nonces to provide anti-replay protection.

[RFC8366] defined only the voucher artifact, and not the Voucher Request artifact, which was defined in [I-D.ietf-anima-bootstrapping-keyinfra].

This document defines both a constrained voucher and a constrained voucher-request. They are presented in the order "voucher-request", followed by a "voucher" response as this is the time order that they occur.

This document defines both CMS-signed voucher requests and responses, and COSE signed voucher requests and responses. The use of CMS signatures implies the use of PKIX format certificates. The pinned-domain-cert present in a voucher, is the certificate of the Registrar.

The constrained voucher and constrained voucher request MUST be signed.

The use of the two signing formats permit the use of both PKIX format certificates, and raw public keys (RPK).

When RPKs are used, the voucher produced by the MASA pins the raw public key of the Registrar: the pinned-domain-subject-public-key-info in a voucher, is the raw public key of the Registrar. This is described in the YANG definition for the constrained voucher.

## 5. Discovery and URI

This section describes the BRSKI extensions to EST-coaps [I-D.ietf-ace-coap-est] to transport the voucher between registrar, proxy and pledge over CoAP. The extensions are targeted to low-resource networks with small packets. Saving header space is important and the EST-coaps URI is shorter than the EST URI.

The presence and location of (path to) the management data are discovered by sending a GET request to `"/.well-known/core"` including a resource type (RT) parameter with the value `"ace.est"` [RFC6690]. Upon success, the return payload will contain the root resource of the EST resources. It is up to the implementation to choose its root resource; throughout this document the example root resource `/est` is used.

The EST-coaps server URIs differ from the EST URI by replacing the scheme `https` by `coaps` and by specifying shorter resource path names:

```
coaps://www.example.com/est/short-name
```

Figure 5 in section 3.2.2 of [RFC7030] enumerates the operations and corresponding paths which are supported by EST. Table 1 provides the mapping from the BRSKI extension URI path to the EST-coaps URI path.

BRSKI	EST-coaps
/requestvoucher	/rv
/voucher_status	/vs
/enrollstatus	/es
/requestauditlog	/ra

Table 1: BRSKI path to EST-coaps path

/requestvoucher, /voucher\_status and /enrollstatus are needed between pledge and Registrar.

When discovering the root path for the EST resources, the server MAY return the full resource paths and the used content types. This is useful when multiple content types are specified for EST-coaps server. For example, the following more complete response is possible.

[ EDNOTE: spell out where voucher artifacts are used in BRSKI flows since the APIs ]

[ EDNOTE: The /requestauditlog and /voucher-status are exchanged by the Registrar and MASA. The JRC will likely talk to MASA over a normal (not constrained) medium. Do we need /ra and /vs? Do we need to remove them from the example too? Also what happens to the voucher-request and response in this case? Is MASA supposed to support constrained vouchers? ]

```
REQ: GET /.well-known/core?rt=brski*
```

```
RES: 2.05 Content
</b>; rt="brski"
</b/rv>; rt="brski.rv";ct=TBD2 TBD3
</b/vs>; rt="brski.vs";ct=50 60
</b/es>; rt="brski.es";ct=50 60
```

The return of the content-types allows the client to choose the most appropriate one from multiple content types.

ct=TBD2 stands for Content-Format "application/voucher-cms+cbor, and ct=TBD3 stands for Content-Format "application/voucher-cose+cbor".

Content-Formats TBD2 and TBD3 are defined in this document.

The Content-Format ("application/json") 50 MAY be supported. Content-Formats ("application/cbor") 60, TBD2, and TBD3 MUST be supported by the Registrar.

The Pledge and MASA need to support one or more formats for the voucher. The MASA needs to support whatever formats that the pledge's produced by that manufacturer supports.

## 6. Artifacts

This section describes the abstract (tree) definition as explained in [I-D.ietf-netmod-yang-tree-diagrams] first. This provides a high-level view of the contents of each artifact.

Then the assigned SID values are presented. These have been assigned using the rules in [I-D.ietf-core-sid], with an allocation that was made via the <http://comi.space> service.

### 6.1. Voucher Request artifact

#### 6.1.1. Tree Diagram

The following diagram is largely a duplicate of the contents of [RFC8366], with the addition of proximity-registrar-subject-public-key-info, proximity-registrar-cert, and prior-signed-voucher-request.

prior-signed-voucher-request is only used between the Registrar and the MASA. proximity-registrar-subject-public-key-info replaces proximity-registrar-cert for the extremely constrained cases.

```
module: ietf-constrained-voucher-request
```

```
  grouping voucher-request-constrained-grouping
```

```
    +-- voucher
```

```
      +-- created-on?
```

```
        | yang:date-and-time
```

```
      +-- expires-on?
```

```
        | yang:date-and-time
```

```
      +-- assertion
```

```
        | enumeration
```

```
      +-- serial-number
```

```
        | string
```

```
      +-- idevid-issuer?
```

```
        | binary
```

```
      +-- pinned-domain-cert?
```

```
        | binary
```

```
      +-- domain-cert-revocation-checks?
```

```
        | boolean
```

```
      +-- nonce?
```

```
        | binary
```

```
      +-- last-renewal-date?
```

```
        | yang:date-and-time
```

```
      +-- proximity-registrar-subject-public-key-info?
```

```
        | binary
```

```
      +-- proximity-registrar-sha256-of-subject-public-key-info?
```

```
        | binary
```

```
      +-- proximity-registrar-cert?
```

```
        | binary
```

```
      +-- prior-signed-voucher-request?
```

```
        | binary
```

#### 6.1.2. SID values

## SID Assigned to

```

-----
2501 data /ietf-constrained-voucher-request:voucher
2502 data .../assertion
2503 data .../created-on
2504 data .../domain-cert-revocation-checks
2505 data .../expires-on
2506 data .../idevid-issuer
2507 data .../last-renewal-date
2508 data /ietf-constrained-voucher-request:voucher/nonce
2509 data .../pinned-domain-cert
2510 data .../prior-signed-voucher-request
2511 data .../proximity-registrar-cert
2512 data mity-registrar-sha256-of-subject-public-key-info
2513 data .../proximity-registrar-subject-public-key-info
2514 data .../serial-number

```

WARNING, obsolete definitions

## 6.1.3. YANG Module

In the constrained-voucher-request YANG module, the voucher is "augmented" within the "used" grouping statement such that one continuous set of SID values is generated for the constrained-voucher-request module name, all voucher attributes, and the constrained-voucher-request attribute. Two attributes of the voucher are "refined" to be optional.

```

<CODE BEGINS> file "ietf-constrained-voucher-request@2019-09-01.yang"
module ietf-constrained-voucher-request {
  yang-version 1.1;

  namespace
    "urn:ietf:params:xml:ns:yang:ietf-constrained-voucher-request";
  prefix "constrained";

  import ietf-restconf {
    prefix rc;
    description
      "This import statement is only present to access
       the yang-data extension defined in RFC 8040.";
    reference "RFC 8040: RESTCONF Protocol";
  }

  import ietf-voucher {
    prefix "v";
  }
}

```

organization

```
"IETF ANIMA Working Group";
```

contact

```
"WG Web: <http://tools.ietf.org/wg/anima/>  
WG List: <mailto:anima@ietf.org>  
Author: Michael Richardson  
<mailto:mcr+ietf@sandelman.ca>  
Author: Peter van der Stok  
<mailto:consultancy@vanderstok.org>  
Author: Panos Kampanakis  
<mailto:pkampana@cisco.com>";
```

description

```
"This module defines the format for a voucher request,  
which is produced by a pledge to request a voucher.  
The voucher-request is sent to the potential owner's  
Registrar, which in turn sends the voucher request to  
the manufacturer or delegate (MASA).
```

A voucher is then returned to the pledge, binding the pledge to the owner. This is a constrained version of the voucher-request present in draft-ietf-anima-bootstrap-keyinfra.txt.

This version provides a very restricted subset appropriate for very constrained devices. In particular, it assumes that nonce-ful operation is always required, that expiration dates are rather weak, as no clocks can be assumed, and that the Registrar is identified by a pinned Raw Public Key.

The key words 'MUST', 'MUST NOT', 'REQUIRED', 'SHALL', 'SHALL NOT', 'SHOULD', 'SHOULD NOT', 'RECOMMENDED', 'MAY', and 'OPTIONAL' in the module text are to be interpreted as described in RFC 2119.";

```
revision "2019-09-01" {  
  description  
    "Initial version";  
  reference  
    "RFC XXXX: Voucher Profile for Constrained Devices";  
}
```

```
rc:yang-data voucher-request-constrained-artifact {  
  // YANG data template for a voucher.  
  uses voucher-request-constrained-grouping;  
}
```

```
// Grouping defined for future usage
grouping voucher-request-constrained-grouping {
  description
    "Grouping to allow reuse/extensions in future work.";

  uses v:voucher-artifact-grouping {

    refine voucher/created-on {
      mandatory false;
    }

    refine voucher/pinned-domain-cert {
      mandatory false;
    }

  }

  augment "voucher" {
    description "Base the constrained voucher-request upon the
      regular one";

    leaf proximity-registrar-subject-public-key-info {
      type binary;
      description
        "The proximity-registrar-subject-public-key-info replaces
        the proximit-registrar-cert in constrained uses of
        the voucher-request.
        The proximity-registrar-subject-public-key-info is the
        Raw Public Key of the Registrar. This field is encoded
        as specified in RFC7250, section 3.
        The ECDSA algorithm MUST be supported.
        The EdDSA algorithm as specified in
        draft-ietf-tls-rfc4492bis-17 SHOULD be supported.
        Support for the DSA algorithm is not recommended.
        Support for the RSA algorithm is MAY, but due to
        size is discouraged.";
    }

    leaf proximity-registrar-sha256-of-subject-public-key-info {
      type binary;
      description
        "The proximity-registrar-sha256-of-subject-public-key-info
        is an alternative to
        proximity-registrar-subject-public-key-info.
        and pinned-domain-cert. In many cases the
        public key of the domain has already been transmitted
        during the key agreement protocol, and it is wasteful
        to transmit the public key another two times.
        The use of a hash of public key info, at 32-bytes for
```

```
    sha256 is a significant savings compared to an RSA
    public key, but is only a minor savings compared to
    a 256-bit ECDSA public-key.
    Algorithm agility is provided by extensions to this
    specifications which define new leaf for other hash
    types.";
}

leaf proximity-registrar-cert {
  type binary;
  description
    "An X.509 v3 certificate structure as specified by
    RFC 5280,
    Section 4 encoded using the ASN.1 distinguished encoding
    rules (DER), as specified in ITU-T X.690.

    The first certificate in the Registrar TLS server
    certificate_list sequence (see [RFC5246]) presented by
    the Registrar to the Pledge. This MUST be populated in a
    Pledge's voucher request if the proximity assertion is
    populated.";
}

leaf prior-signed-voucher-request {
  type binary;
  description
    "If it is necessary to change a voucher, or re-sign and
    forward a voucher that was previously provided along a
    protocol path, then the previously signed voucher
    SHOULD be included in this field.

    For example, a pledge might sign a proximity voucher,
    which an intermediate registrar then re-signs to
    make its own proximity assertion. This is a simple
    mechanism for a chain of trusted parties to change a
    voucher, while maintaining the prior signature
    information.

    The pledge MUST ignore all prior voucher information
    when accepting a voucher for imprinting. Other
    parties MAY examine the prior signed voucher
    information for the purposes of policy decisions.
    For example this information could be useful to a
    MASA to determine that both pledge and registrar
    agree on proximity assertions. The MASA SHOULD
    remove all prior-signed-voucher-request information when
    signing a voucher for imprinting so as to minimize the
    final voucher size.";
```

```

    }
  }
}
}
<CODE ENDS>

```

#### 6.1.4. Example voucher request artifact

Below is a CBOR serialization of the constrained-voucher-request is shown in diagnostic CBOR notation. The enum value of the assertion field is calculated to be zero by following the algorithm described in section 9.6.4.2 of [RFC7950].

```

{
  2501: {
    +2 : "2016-10-07T19:31:42Z", / SID= 2503, created-on /
    +4 : "2016-10-21T19:31:42Z", / SID= 2505, expires-on /
    +1 : 2, / SID= 2502, assertion /
    / "proximity" /
    +13: "JADA123456789", / SID= 2514, serial-number /
    +5 : h'01020D0F', / SID= 2506, idevid-issuer /
    +10: h'cert.der', / SID=2511, proximity-registrar-cert/
    +3 : true, / SID= 2504, domain-cert
    -revocation-checks/
    +6 : "2017-10-07T19:31:42Z", / SID= 2507, last-renewal-date /
    +12: h'key_info' / SID= 2513, proximity-registrar
    -subject-public-key-info /
  }
}

```

#### 6.2. Voucher artifact

The voucher's primary purpose is to securely assign a pledge to an owner. The voucher informs the pledge which entity it should consider to be its owner.

This document defines a voucher that is a CBOR encoded instance of the YANG module defined in Section 5.3 that has been signed with CMS or with COSE.

## 6.2.1. Tree Diagram

The following diagram is largely a duplicate of the contents of [RFC8366], with only the addition of pinned-domain-subject-public-key-info.

```
module: ietf-constrained-voucher
```

```

grouping voucher-constrained-grouping
  +-- voucher
    +-- created-on?
       | yang:date-and-time
    +-- expires-on?
       | yang:date-and-time
    +-- assertion enumeration
    +-- serial-number string
    +-- idevid-issuer? binary
    +-- pinned-domain-cert? binary
    +-- domain-cert-revocation-checks? boolean
    +-- nonce? binary
    +-- last-renewal-date?
       | yang:date-and-time
    +-- pinned-domain-subject-public-key-info? binary
    +-- pinned-sha256-of-subject-public-key-info? binary

```

## 6.2.2. SID values

SID Assigned to

```

-----
2451 data /ietf-constrained-voucher:voucher
2452 data /ietf-constrained-voucher:voucher/assertion
2453 data /ietf-constrained-voucher:voucher/created-on
2454 data .../domain-cert-revocation-checks
2455 data /ietf-constrained-voucher:voucher/expires-on
2456 data /ietf-constrained-voucher:voucher/idevid-issuer
2457 data .../last-renewal-date
2458 data /ietf-constrained-voucher:voucher/nonce
2459 data .../pinned-domain-cert
2460 data .../pinned-domain-subject-public-key-info
2461 data .../pinned-sha256-of-subject-public-key-info
2462 data /ietf-constrained-voucher:voucher/serial-number

```

WARNING, obsolete definitions

## 6.2.3. YANG Module

In the constrained-voucher YANG module, the voucher is "augmented" within the "used" grouping statement such that one continuous set of SID values is generated for the constrained-voucher module name, all voucher attributes, and the constrained-voucher attribute. Two attributes of the voucher are "refined" to be optional.

```
<CODE BEGINS> file "ietf-constrained-voucher@2019-09-01.yang"
module ietf-constrained-voucher {
  yang-version 1.1;

  namespace
    "urn:ietf:params:xml:ns:yang:ietf-constrained-voucher";
  prefix "constrained";

  import ietf-restconf {
    prefix rc;
    description
      "This import statement is only present to access
       the yang-data extension defined in RFC 8040.";
    reference "RFC 8040: RESTCONF Protocol";
  }

  import ietf-voucher {
    prefix "v";
  }

  organization
    "IETF ANIMA Working Group";

  contact
    "WG Web: <http://tools.ietf.org/wg/anima/>
    WG List: <mailto:anima@ietf.org>
    Author: Michael Richardson
            <mailto:mcr+ietf@sandelman.ca>
    Author: Peter van der Stok
            <mailto:consultancy@vanderstok.org>
    Author: Panos Kampanakis
            <mailto:pkampana@cisco.com>";

  description
    "This module defines the format for a voucher, which is produced
    by a pledge's manufacturer or delegate (MASA) to securely assign
    one or more pledges to an 'owner', so that the pledges may
    establish a secure connection to the owner's network
    infrastructure.
```

This version provides a very restricted subset appropriate

for very constrained devices.

In particular, it assumes that nonce-ful operation is always required, that expiration dates are rather weak, as no clocks can be assumed, and that the Registrar is identified by a pinned Raw Public Key.

The key words 'MUST', 'MUST NOT', 'REQUIRED', 'SHALL', 'SHALL NOT', 'SHOULD', 'SHOULD NOT', 'RECOMMENDED', 'MAY', and 'OPTIONAL' in the module text are to be interpreted as described in RFC 2119."

```
revision "2019-09-01" {
  description
    "Initial version";
  reference
    "RFC XXXX: Voucher Profile for Constrained Devices";
}

rc:yang-data voucher-constrained-artifact {
  // YANG data template for a voucher.
  uses voucher-constrained-grouping;
}

// Grouping defined for future usage
grouping voucher-constrained-grouping {
  description
    "Grouping to allow reuse/extensions in future work.";

  uses v:voucher-artifact-grouping {

    refine voucher/created-on {
      mandatory false;
    }

    refine voucher/pinned-domain-cert {
      mandatory false;
    }

    augment "voucher" {
      description "Base the constrained voucher
                  upon the regular one";
      leaf pinned-domain-subject-public-key-info {
        type binary;
        description
          "The pinned-domain-subject-public-key-info replaces the
           pinned-domain-cert in constrained uses of
           the voucher. The pinned-domain-subject-public-key-info
           is the Raw Public Key of the Registrar."
      }
    }
  }
}
```



```
{
  2451: {
    +2 : "2016-10-07T19:31:42Z", / SID = 2453, created-on /
    +4 : "2016-10-21T19:31:42Z", / SID = 2455, expires-on /
    +1 : 0, / SID = 2452, assertion /
    / "verified" /
    +11: "JADA123456789", / SID = 2462, serial-number /
    +5 : h'01020D0F', / SID = 2456, idevid-issuer /
    +8 : h'cert.der', / SID = 2459, pinned-domain-cert/
    +3 : true, / SID = 2454, domain-cert
    / -revocation-checks /
    +6 : "2017-10-07T19:31:42Z", / SID = 2457, last-renewal-date /
    +9 : h'key-info' / SID = 2460, pinned-domain
    / -subject-public-key-info /
  }
}
```

The signing of the example is shown in Appendix B.3.

### 6.3. Signing voucher and voucher-request artifacts

#### 6.3.1. CMS signing

The IETF evolution of PKCS#7 is CMS [RFC5652]. The CMS signed voucher is much like the equivalent voucher defined in [RFC8366].

A different eContentType of TBD1 is used to indicate that the contents are in a different format than in [RFC8366]. The id-ct-animaJSONVoucher allocated by [RFC8366] indicates a voucher and voucher-request encoded in JSON, and the new value TBD1 indicates that the voucher and voucher-request are encoded in CBOR.

The ContentInfo structure contains a payload consisting of the CBOR encoded voucher. The [I-D.ietf-core-yang-cbor] use of delta encoding creates a canonical ordering for the keys on the wire. This canonical ordering is not important as there is no expectation that the content will be reproduced during the validation process.

Normally the recipient is the pledge and the signer is the MASA.

[I-D.ietf-anima-bootstrapping-keyinfra] supports both signed and unsigned voucher requests from the pledge to the JRC. In this specification, voucher-request artifact is not signed from the pledge to the registrar. [EDNOTE: Confirm that voucher requests do not need

to be signed ] From the JRC to the MASA, the voucher-request artifact MUST be signed by the domain owner key which is requesting ownership.

The considerations of [RFC5652] section 5.1, concerning validating CMS objects which are really PKCS7 objects (cmsVersion=1) applies.

The CMS structure SHOULD also contain all the certificates leading up to and including the signer's trust anchor certificate known to the recipient. The inclusion of the trust anchor is unusual in many applications, but without it third parties can not accurately audit the transaction.

The CMS structure MAY also contain revocation objects for any intermediate certificate authorities (CAs) between the voucher-issuer and the trust anchor known to the recipient. However, the use of CRLs and other validity mechanisms is discouraged, as the pledge is unlikely to be able to perform online checks, and is unlikely to have a trusted clock source. As described below, the use of short-lived vouchers and/or pledge provided nonce provides a freshness guarantee.

### 6.3.2. COSE signing

The COSE-Sign1 structure is discussed in section 4.2 of [RFC8152]. The CBOR object that carries the body, the signature, and the information about the body and signature is called the COSE\_Sign1 structure. It is used when only one signature is used on the body. Support for ECDSA with sha256 (secp256k1 and prime256v1 curves) is compulsory.

The supported COSE-sign1 object structure is shown in Figure 1. Support for EdDSA is encouraged. [EDNOTE: Expand and add a reference why. ]

```
COSE_Sign1(
  [
    h'A101382E',      # { "alg": EC256K1 }
    {
      "kid" : h'789'  # hash256(public key)
    },
    h'123', #voucher-request binary content
    h'456', #voucher-request binary public signature
  ]
)
```

Figure 1: cose-sign1 example

The [COSE-registry] specifies the integers that replace the strings and the mnemonics in Figure 1. The value of the "kid" parameter is

an example value. Usually a hash of the public key is used to identify the public key. The public key and its hash are derived from the relevant certificate (Pledge, Registrar or MASA certificate).

In Appendix C a binary cose-sign1 object is shown based on the voucher-request example of Section 6.1.4.

## 7. Design Considerations

The design considerations for the CBOR encoding of vouchers is much the same as for [RFC8366].

One key difference is that the names of the leaves in the YANG does not have a material effect on the size of the resulting CBOR, as the SID translation process assigns integers to the names.

## 8. Security Considerations

### 8.1. Clock Sensitivity

TBD.

### 8.2. Protect Voucher PKI in HSM

TBD.

### 8.3. Test Domain Certificate Validity when Signing

TBD.

## 9. IANA Considerations

### 9.1. Resource Type Registry

Additions to the sub-registry "CoAP Resource Type", within the "CoRE parameters" registry are specified below. These can be registered either in the Expert Review range (0-255) or IETF Review range (256-9999).

ace.rt.rv needs registration with IANA  
ace.rt.vs needs registration with IANA  
ace.rt.es needs registration with IANA  
ace.rt.ra needs registration with IANA



### 9.5. The SMI Security for S/MIME CMS Content Type Registry

This document registers an OID in the "SMI Security for S/MIME CMS Content Type" registry (1.2.840.113549.1.9.16.1), with the value:

Decimal	Description	References
46	id-ct-animaCBORVoucher	[ThisRFC]

### 9.6. Media-Type Registry

This section registers the 'application/voucher-cms+cbor' media type and the 'application/voucher-cose+cbor' in the "Media Types" registry. These media types are used to indicate that the content is a CBOR voucher either signed with a cms structure or a COSE\_Sign1 structure [RFC8152].

#### 9.6.1. application/voucher-cms+cbor

Type name: application  
 Subtype name: voucher-cms+cbor  
 Required parameters: none  
 Optional parameters: none  
 Encoding considerations: CMS-signed CBOR vouchers are CBOR encoded.  
 Security considerations: See Security Considerations, Section  
 Interoperability considerations: The format is designed to be broadly interoperable.  
 Published specification: THIS RFC.  
 Applications that use this media type: ANIMA, 6tisch, and other zero-touch imprinting systems  
 Additional information:  
 Magic number(s): None  
 File extension(s): .vch  
 Macintosh file type code(s): none  
 Person & email address to contact for further information: IETF ANIMA WG  
 Intended usage: LIMITED  
 Restrictions on usage: NONE  
 Author: ANIMA WG  
 Change controller: IETF  
 Provisional registration? (standards tree only): NO

#### 9.6.2. application/voucher-cose+cbor

Type name: application  
 Subtype name: voucher-cose+cbor  
 Required parameters: none  
 Optional parameters: cose-type  
 Encoding considerations: COSE\_Sign1 CBOR vouchers are COSE objects signed with one signer.  
 Security considerations: See Security Considerations, Section  
 Interoperability considerations: The format is designed to be broadly interoperable.  
 Published specification: THIS RFC.  
 Applications that use this media type: ANIMA, 6tisch, and other zero-touch imprinting systems  
 Additional information:  
   Magic number(s): None  
   File extension(s): .vch  
   Macintosh file type code(s): none  
 Person & email address to contact for further information: IETF ANIMA WG  
 Intended usage: LIMITED  
 Restrictions on usage: NONE  
 Author: ANIMA WG  
 Change controller: IETF  
 Provisional registration? (standards tree only): NO

### 9.7. CoAP Content-Format Registry

Additions to the sub-registry "CoAP Content-Formats", within the "CoRE Parameters" registry are needed for two media types. These can be registered either in the Expert Review range (0-255) or IETF Review range (256-9999).

Media type	mime type	Encoding	ID	References
application/voucher-cms+cbor	- -	CBOR	TBD2	[This RFC]
application/voucher-cose+cbor	"COSE-Sign1"	CBOR	TBD3	[This RFC]

### 10. Acknowledgements

We are very grateful to Jim Schaad for explaining COSE and CMS choices. Also thanks to Jim Schaad for correctinging earlier version of the COSE Sign1 objects.

Michel Veillette did extensive work on pyang to extend it to support the SID allocation process, and this document was among the first users.

We are grateful for the suggestions done by Esko Dijk.

## 11. Changelog

-08 Examples for cose\_sign1 are completed and improved.

-06 New SID values assigned; regenerated examples

-04 voucher and request-voucher MUST be signed examples for signed request are added in appendix IANA SID registration is updated SID values in examples are aligned signed cms examples aligned with new SIDs

-03

Examples are inverted.

-02

Example of requestvoucher with unsigned application/cbor is added attributes of voucher "refined" to optional CBOR serialization of vouchers improved Discovery port numbers are specified

-01

application/json is optional, application/cbor is compulsory Cms and cose mediatypes are introduced

## 12. References

### 12.1. Normative References

[I-D.ietf-ace-coap-est]

Stok, P., Kampanakis, P., Richardson, M., and S. Raza, "EST over secure CoAP (EST-coaps)", draft-ietf-ace-coap-est-18 (work in progress), January 2020.

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

Pritikin, M., Richardson, M., Eckert, T., Behringer, M., and K. Watsen, "Bootstrapping Remote Secure Key Infrastructures (BRSKI)", draft-ietf-anima-bootstrapping-keyinfra-44 (work in progress), September 2020.

[I-D.ietf-core-sid]

Veillette, M., Pelov, A., and I. Petrov, "YANG Schema Item Identifier (YANG SID)", draft-ietf-core-sid-14 (work in progress), July 2020.

- [I-D.ietf-core-yang-cbor]  
Veillette, M., Petrov, I., and A. Pelov, "CBOR Encoding of Data Modeled with YANG", draft-ietf-core-yang-cbor-13 (work in progress), July 2020.
- [I-D.selander-ace-ake-authz]  
Selander, G., Mattsson, J., Vucinic, M., Richardson, M., and A. Schellenbaum, "Lightweight Authorization for Authenticated Key Exchange.", draft-selander-ace-ake-authz-01 (work in progress), March 2020.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC3688] Mealling, M., "The IETF XML Registry", BCP 81, RFC 3688, DOI 10.17487/RFC3688, January 2004, <<https://www.rfc-editor.org/info/rfc3688>>.
- [RFC5652] Housley, R., "Cryptographic Message Syntax (CMS)", STD 70, RFC 5652, DOI 10.17487/RFC5652, September 2009, <<https://www.rfc-editor.org/info/rfc5652>>.
- [RFC6020] Bjorklund, M., Ed., "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)", RFC 6020, DOI 10.17487/RFC6020, October 2010, <<https://www.rfc-editor.org/info/rfc6020>>.
- [RFC7049] Bormann, C. and P. Hoffman, "Concise Binary Object Representation (CBOR)", RFC 7049, DOI 10.17487/RFC7049, October 2013, <<https://www.rfc-editor.org/info/rfc7049>>.
- [RFC7950] Bjorklund, M., Ed., "The YANG 1.1 Data Modeling Language", RFC 7950, DOI 10.17487/RFC7950, August 2016, <<https://www.rfc-editor.org/info/rfc7950>>.
- [RFC8152] Schaad, J., "CBOR Object Signing and Encryption (COSE)", RFC 8152, DOI 10.17487/RFC8152, July 2017, <<https://www.rfc-editor.org/info/rfc8152>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

- [RFC8366] Watsen, K., Richardson, M., Pritikin, M., and T. Eckert, "A Voucher Artifact for Bootstrapping Protocols", RFC 8366, DOI 10.17487/RFC8366, May 2018, <<https://www.rfc-editor.org/info/rfc8366>>.

## 12.2. Informative References

- [COSE-registry] IANA, ., "CBOR Object Signing and Encryption (COSE) registry", 2017, <<https://www.iana.org/assignments/cose/cose.xhtml>>.
- [I-D.ietf-netmod-yang-tree-diagrams] Bjorklund, M. and L. Berger, "YANG Tree Diagrams", draft-ietf-netmod-yang-tree-diagrams-06 (work in progress), February 2018.
- [RFC6690] Shelby, Z., "Constrained RESTful Environments (CoRE) Link Format", RFC 6690, DOI 10.17487/RFC6690, August 2012, <<https://www.rfc-editor.org/info/rfc6690>>.
- [RFC7030] Pritikin, M., Ed., Yee, P., Ed., and D. Harkins, Ed., "Enrollment over Secure Transport", RFC 7030, DOI 10.17487/RFC7030, October 2013, <<https://www.rfc-editor.org/info/rfc7030>>.

## Appendix A. EST messages to EST-coaps

This section extends the examples from Appendix A of [I-D.ietf-ace-coap-est]. The CoAP headers are only worked out for the enrollstatus example.

### A.1. enrollstatus

A coaps enrollstatus message can be :

```
GET coaps://[192.0.2.1:8085]/est/es
```

The corresponding coap header fields are shown below.

```

Ver = 1
T = 0 (CON)
Code = 0x01 (0.01 is GET)
Options
Option (Uri-Path)
  Option Delta = 0xb (option nr = 11)
  Option Length = 0x3
  Option Value = "est"
Option (Uri-Path)
  Option Delta = 0x0 (option nr = 11)
  Option Length = 0x2
  Option Value = "es"
Payload = [Empty]

```

The Uri-Host and Uri-Port Options are omitted because they coincide with the transport protocol destination address and port respectively.

A 2.05 Content response with an unsigned voucher status (ct=60) will then be:

```
2.05 Content (Content-Format: application/cbor)
```

With CoAP fields and payload:

```

Ver=1
T=2 (ACK)
Code = 0x45 (2.05 Content)
Options
Option1 (Content-Format)
  Option Delta = 0xC (option nr 12)
  Option Length = 0x2
  Option Value = 60 (application/cbor)

Payload (CBOR diagnostic) =
{
  "version": "1",
  "Status": 1, / 1 = Success, 0 = Fail /
  "Reason": "Informative human readable message",
  "reason-context": "Additional information"
}

```

The binary payload is:

```

A46776657273696F6E6131665374617475730166526561736F6E7822
496E666F726D61746976652068756D616E207265616461626C65206D
6573736167656e726561736F6E2D636F6E74657874
764164646974696F6E616C20696E666F726D6174696F6E

```

The binary payload disassembles to the above CBOR diagnostic code.

#### A.2. voucher\_status

A coaps voucher\_status message can be:

```
GET coaps://[2001:db8::2:1]:61616]/est/vs
```

A 2.05 Content response with a non signed CBOR voucher status (ct=60) will then be:

```
2.05 Content (Content-Format: application/cbor)
```

```
Payload =
```

```
A46776657273696F6E6131665374617475730166526561736F6E7822  
496E666F726D61746976652068756D616E207265616461626C65206D  
6573736167656e726561736F6E2D636F6E74657874  
764164646974696F6E616C20696E666F726D6174696F6E
```

#### Appendix B. CMS signed messages

Signed request-voucher-request payloads are sent from pledge to Registrar, as explained in Section 5.2 of [I-D.ietf-anima-bootstrapping-keyinfra].

##### B.1. signed requestvoucher

A CMS signed requestvoucher message from JRC to MASA is shown below. It would be CoAP POSTED to /est/rv.

```
POST coaps://[2001:db8::2:1]:61616]/est/rv  
(Content-Format: application/voucher-cms+cbor)
```

The payload would be in binary, but is presented in base64 in this document.

MIIDugYJKoZIhvcNAQcCoIIDqzCCA6cCAQEExDTALBglghkgBZQMEAgEwCwYJ  
KoZIhvcNAQcBoIICQTCCAj0wggHioAMCAQICCH52Yde1TkYyMAoGCCqGSM49  
BAMCMF0xCzAJBgNVBAYTA1VTMQswCQYDVQQIDAJDQTEUMBIGAlUECgwLRXhh  
bXBsZSBJamMxZjAUBGNVBAsMDWN1cnRpZmljYXRpb24xEzARBgNVBAMMCjgw  
Mi4xQVIgQ0EwIBcNMTkwMTMxMTEyOTE2WhgPOTk5OTEyMzEyMzU5NTlaMFwx  
CzAJBgNVBAYTA1VTMQswCQYDVQQIDAJDQTELMAGAlUEBwwCTEExFDASBgNV  
BAoMC2V4YW1wbGUgSW5jMQwwCgYDVQQLDANJb1QxDzANBgNVBAUTBld0MTIz  
NDBZMBMGByqGSM49AgEGCCqGSM49AwEHA0IABMi0IfEcJeR+OsVxI78tn9xJ  
TwKLw1HMgMA/FQv1DP+VjXVBnYGmoxf+ueQvpXPdfYC+RUmGPgWorI7Vjj1  
n9mjgYowgYcwCQYDVR0TBAlwADAdBgNVHQ4EFgQUlmanhxa/f9DnUtCsdgd3  
rWzdAqAwHwYDVR0jBBgwFoAUaNF1Uf1Rv8gqQx0Nnwi8LSBbEWAwdgYDVR0P  
AQH/BAQDAgWgMCoGAlUdEQQjMCGgHwYIKwYBBQUHCASeqEzARBgkrBgEEAbQ7  
CgEEBAECAwQwCgYIKoZIzj0EAwIDSQAwRgIhAMDYGZbSUH1pPzxI6qXulJG9  
ptshQJnZgRfGOzYTdM2GaiEAp3SYn0wyG1zyXYMqTTNqCK1n3yDxUGQhGIoK  
3m00k jYxggE/MIIBOwIBATBpMF0xCzAJBgNVBAYTA1VTMQswCQYDVQQIDAJD  
QTEUMBIGAlUECgwLRXhhbXBsZSBJamMxZjAUBGNVBAsMDWN1cnRpZmljYXRp  
b24xEzARBgNVBAMMCjgwMi4xQVIgQ0EwIBcNMTkwMTMxMTEyOTE2WhgPOTk5  
OTEyMzEyMzU5NTlaMFwxZzAaBjQcZSIB3DQEJAzELBglghkiG9w0BBwEwHAYJKoZIhvcNAQkFMQ8X  
DTE5MDQwODEwNDgzNlowLwYJKoZIhvcNAQkEMSIEIIEEdCTOLs2Zy7w3LgvSZ  
XZEadz3Lbzn0FBs6FMFN91RaMAoGCCqGSM49BAMCBECwRQIgASjDsIpr0tW/  
n6dRRHqvvsqgZlHbtFnErUbWfhS0KD4CIQDDUEqc5wTmRGf0adEQVQzqmIgh  
MEgF10vqXv02gL1jLw==

A 2.04 Changed response returning CBOR voucher signed with a cms structure(ct=TBD2) will then be:

2.04 Changed (Content-Format: application/voucher-cms+cbor)

```

MIIDuwYJKoZIhvcNAQcCoIIdrDCCA6gCAQEExDTALBglghkgBZQMEAgEwCwYJ
KoZIhvcNAQcBoIICQTCCAj0wggHioAMCAQICCH52Yde1TkYyMAoGCCqGSM49
BAMCMF0xCzAJBgNVBAYTA1VTMQswCQYDVQQIDAJDQTEUMBIGAlUECgwLRXhh
bXBsZSBjb24xZmMxZjAUBG9NBAsMDWN1cnRpb24xZmMxZjAUBG9NBAMMCjgw
Mi4xQVIGQ0EwIBcNMTkwMTMxMTEyOTE2WhgPOTk5OTEyMzEyMzU5NTlaMFwx
CzAJBgNVBAYTA1VTMQswCQYDVQQIDAJDQTELMakGAlUEBwwCTEExFDASBgNV
BAoMC2V4YW1wbGUgSW5jMQwwCgYDVQLDANJb1QxDzANBgNVBAUTBld0MTIz
NDBZMBMGByqGSM49AgEGCCqGSM49AwEHA0IABMi0IFecJeR+OsVxI78tn9xJ
TwKLw1HMgMA/FQv1DP+VjXVBnYGmoxf+ueQvpXPdfYC+RUmGPgWorI7Vj1
n9mjgYowgYcwcQYDVR0TBAlwADAdBgNVHQ4EFgQUlmanhxa/f9DnUtCsdgd3
rWzdAqAwHwYDVR0jBBgwFoAUaNF1Uf1Rv8gqQx0Nnwi8LSBbEWAwdgYDVR0P
AQH/BAQDAgWgMCoGAlUdEQQjMCGgHwYIKwYBBQUHCAAgEzARBgkrBgEEAbQ7
CgEEBAECAwQwCgYIKoZIzj0EAwIDSQAwRgIhAMDYgzbsUH1pPzxI6qXulJG9
ptshQJnZgRfGOzYTdM2GaiEAp3SYn0wyG1zyXYMqTTNqCK1n3yDxUGQhGIoK
3m00kjYxggFAMIIBPAIBATBpMF0xCzAJBgNVBAYTA1VTMQswCQYDVQQIDAJD
QTEUMBIGAlUECgwLRXhhbXBsZSBjb24xZmMxZjAUBG9NBAsMDWN1cnRpb24x
ZmMxZjAUBG9NBAMMCjgwMi4xQVIGQ0ECCH52Yde1TkYyMAoGCCqGSM49BAMC
BEGwRgIhALOd2FKbe9FGkn4Pg7FIgF+//cQv/N+v7tDZMzGBAFN0AIEAu5BI0oQ4o0wZcrDyKoU2GbeX
hlG/g+OgTUftYMJ32so=

```

B.2. requestauditing

A coaps requestauditing message contains the signed CBOR voucher :

```

POST coaps://[2001:db8::2:1]:61616]/est/ra
(Content-Format: application/voucher-cms+cbor)
Payload =
TO BE FILLED

```

A 2.05 Content response returning a log of the voucher (ct=60) will then be:

```

    2.05 Content (Content-Format: application/cbor)
    Payload =
{
  "version": "1",
  "events": [
    {
      "date": "<date/time of the entry>",
      "domainID": "<domainID extracted from voucher-request>",
      "nonce": "<any nonce if supplied (or the exact string 'NULL')>",
      "assertion": "<the value from the voucher assertion leaf>",
      "truncated": "<the number of domainID entries truncated>"
    },
    {
      "date": "<date/time of the entry>",
      "domainID": "<anotherDomainID extracted from voucher-request>",
      "nonce": "<any nonce if supplied (or the exact string 'NULL')>",
      "assertion": "<the value from the voucher assertion leaf>"
    }
  ],
  "truncation": {
    "nonced duplicates": "<total number of entries truncated>",
    "nonceless duplicates": "<total number of entries truncated>",
    "arbitrary": "<number of domainID entries removed entirely>"
  }
}

```

[EDNOTE: Change JSON to CBOR; Serialize CBOR payload to binary]

### B.3. CMS signed voucher-request example

The voucher-request example, visualized in CBOR diagnostic notation in Section 6.1.4 is shown as a hexadecimal dump of the binary file.

```

a11909c5a90274323031362d31302d30375431393a33313a34325a0474323031
362d31302d32315431393a33313a34325a01020d6d4a414441313233343536
373839054401020d0f0a4401020d0f03f50674323031372d31302d30375431
393a33313a34325a0c4401020d0f

```

The voucher-request example has been signed by using the WT1234 certificate and key pair shown in Appendix C of [I-D.ietf-ace-coap-est]. The CMS signing of the binary voucher-request leads to a binary signed voucher-request, shown with a hexadecimal representation shown in the payload of the request part of Appendix B.1 and Appendix B.2.

The breakdown of the CMS signed binary voucher-request file is visualized below:

```
CMS_ContentInfo:
  contentType: pkcs7-signedData (1.2.840.113549.1.7.2)
  d.signedData:
    version: 1
    digestAlgorithms:
      algorithm: sha256 (2.16.840.1.101.3.4.2.1)
      parameter: <ABSENT>
    encapContentInfo:
      eContentType: pkcs7-data (1.2.840.113549.1.7.1)
      eContent: <ABSENT>
    certificates:
      d.certificate:
        cert_info:
          version: 2
          serialNumber: 9112578475118446130
          signature:
            algorithm: ecdsa-with-SHA256 (1.2.840.10045.4.3.2)
            parameter: <ABSENT>
            issuer: C=US, ST=CA, O=Example Inc, OU=certification,
                  CN=802.1AR CA
          validity:
            notBefore: Jan 31 11:29:16 2019 GMT
            notAfter: Dec 31 23:59:59 9999 GMT
            subject: C=US, ST=CA, L=LA, O=example Inc,
                  OU=IoT/serialNumber=Wt1234
        key:
          algor:
            algorithm: id-ecPublicKey (1.2.840.10045.2.1)
            parameter: OBJECT:prime256v1 (1.2.840.10045.3.1.7)
          public_key: (0 unused bits)
            0000 - 04 c8 b4 21 f1 1c 25 e4-7e 3a c5 71 23 bf
            000e - 2d 9f dc 49 4f 02 8b c3-51 cc 80 c0 3f 15
            001c - 0b f5 0c ff 95 8d 75 41-9d 81 a6 a2 45 df
            002a - fa e7 90 be 95 cf 75 f6-02 f9 15 26 18 f8
            0038 - 16 a2 b2 3b 56 38 e5 9f-d9
          issuerUID: <ABSENT>
          subjectUID: <ABSENT>
        extensions:
          object: X509v3 Basic Constraints (2.5.29.19)
          critical: BOOL ABSENT
          value:
            0000 - 30
            0002 - <SPACES/NULS>

          object: X509v3 Subject Key Identifier (2.5.29.14)
          critical: BOOL ABSENT
          value:
            0000 - 04 14 96 60 0d 87 16 bf-7f d0 e7 52 d0
```

```
000d - ac 76 07 77 ad 66 5d 02-a0

object: X509v3 Authority Key Identifier (2.5.29.35)
critical: BOOL ABSENT
value:
  0000 - 30 16 80 14 68 d1 65 51-f9 51 bf c8 2a
  000d - 43 1d 0d 9f 08 bc 2d 20-5b 11 60

object: X509v3 Key Usage (2.5.29.15)
critical: TRUE
value:
  0000 - 03 02 05 a0

object: X509v3 Subject Alternative Name (2.5.29.17)
critical: BOOL ABSENT
value:
  0000 - 30 21 a0 1f 06 08 2b 06-01 05 05 07 08
  000d - 04 a0 13 30 11 06 09 2b-06 01 04 01 b4
  001a - 3b 0a 01 04 04 01 02 03-04

sig_alg:
  algorithm: ecdsa-with-SHA256 (1.2.840.10045.4.3.2)
  parameter: <ABSENT>
signature: (0 unused bits)
  0000 - 30 46 02 21 00 c0 d8 19-96 d2 50 7d 69 3f 3c
  000f - 48 ea a5 ee 94 91 bd a6-db 21 40 99 d9 81 17
  001e - c6 3b 36 13 74 cd 86 02-21 00 a7 74 98 9f 4c
  002d - 32 1a 5c f2 5d 83 2a 4d-33 6a 08 ad 67 df 20
  003c - f1 50 64 21 18 8a 0a de-6d 34 92 36

crls:
  <EMPTY>

signerInfos:
  version: 1
  d.issuerAndSerialNumber:
    issuer: C=US, ST=CA, O=Example Inc, OU=certification,
           CN=802.1AR CA
    serialNumber: 9112578475118446130
  digestAlgorithm:
    algorithm: sha256 (2.16.840.1.101.3.4.2.1)
    parameter: <ABSENT>
  signedAttrs:
    object: contentType (1.2.840.113549.1.9.3)
    value.set:
      OBJECT:pkcs7-data (1.2.840.113549.1.7.1)

    object: signingTime (1.2.840.113549.1.9.5)
    value.set:
      UTCTIME:Jul  3 08:53:30 2019 GMT
```

```

object: messageDigest (1.2.840.113549.1.9.4)
value.set:
  OCTET STRING:
    0000 - d4 b0 5c dd c8 b4 91 28-4a 18 ca 25 9d
    000d - be d0 60 23 cf ad a0 aa-c2 95 ac e9 3f
    001a - 0b 4f 44 9e 25
    0020 - <SPACES/NULS>
signatureAlgorithm:
  algorithm: ecdsa-with-SHA256 (1.2.840.10045.4.3.2)
  parameter: <ABSENT>
signature:
  0000 - 30 46 02 21 00 e5 e1 7f-23 c3 aa 14 9f 35 64
  000f - 1e c4 4a 0f 68 fe b0 16-3b e6 7c 06 51 af bf
  001e - 5a a0 99 59 e0 28 1f 02-21 00 b4 07 2f 7c c4
  002d - f9 26 0c 6d 47 a7 93 56-de b8 da f7 23 f0 af
  003c - 2b 59 16 cc 36 63 e7 91-89 39 df df
unsignedAttrs:
  <EMPTY>

```

#### Appendix C. COSE examples

These examples are generated on a pie 4 and a PC running BASH. Keys and Certificates have been generated with openssl with the following shell script:

```

#!/bin/bash
#try-cert.sh
export dir=./brski/intermediate
export cadir=./brski
export cnfdir=./conf
export format=pem
export default_crl_days=30
sn=8

DevID=pledge.1.2.3.4
serialNumber="serialNumber=$DevID"
export hwType=1.3.6.1.4.1.6715.10.1
export hwSerialNum=01020304
export subjectAltName="otherName:1.3.6.1.5.5.7.8.4;SEQ:hmodname"
echo $hwType - $hwSerialNum
echo $serialNumber

# remove all files
rm -r ./brski/*
#
# initialize file structure
# root level
cd $cadir

```

```
mkdir certs crl csr newcerts private
chmod 700 private
touch index.txt
touch serial
echo 11223344556600 >serial
echo 1000 > crlnumber
# intermediate level
mkdir intermediate
cd intermediate
mkdir certs crl csr newcerts private
chmod 700 private
touch index.txt
echo 11223344556600 >serial
echo 1000 > crlnumber
cd ../../

# file structure is cleaned start filling

echo "#####"
echo "create registrar keys and certificates "
echo "#####"

echo "create root registrar certificate using ecDSA with sha256"
openssl ecparam -name prime256v1 -genkey \
  -noout -out $cadir/private/ca-regis.key

openssl req -new -x509 \
  -key $cadir/private/ca-regis.key \
  -out $cadir/certs/ca-regis.crt \
  -extensions v3_ca \
  -days 365 \
  -subj "/C=NL/ST=NB/L=Helmond/O=vanderstok\
  "/OU=consultancy/CN=registrar.stok.nl"

# Combine authority certificate and key
echo "Combine authority certificate and key"
openssl pkcs12 -passin pass:watnietweet \
  -passout pass:watnietweet \
  -inkey $cadir/private/ca-regis.key \
  -in $cadir/certs/ca-regis.crt -export \
  -out $cadir/certs/ca-regis-comb.pfx

# converteer authority pkcs12 file to pem
echo "converteer authority pkcs12 file to pem"
openssl pkcs12 -passin pass:watnietweet -passout pass:watnietweet\
```

```
-in $cadir/certs/ca-regis-comb.pfx \<\  
-out $cadir/certs/ca-regis-comb.crt -nodes  
  
#show certificate in registrar combined certificate  
openssl x509 -in $cadir/certs/ca-regis-comb.crt -text  
  
#  
# Certificate Authority for MASA  
#  
echo "#####"  
echo "create MASA keys and certificates "  
echo "#####"  
  
echo "create root MASA certificate using ecdsa with sha 256 key"  
openssl ecparam -name prime256v1 -genkey -noout \  
-out $cadir/private/ca-masa.key  
  
openssl req -new -x509 \  
-days 365 -key $cadir/private/ca-masa.key \  
-out $cadir/certs/ca-masa.crt \  
-extensions v3_ca\  
-subj "/C=NL/ST=NB/L=Helmond/O=vanderstok/  
OU=manufacturer/CN=masa.stok.nl"  
  
# Combine authority certificate and key  
echo "Combine authority certificate and key for masa"  
openssl pkcs12 -passin pass:watnietweet \  
-passout pass:watnietweet\  
-inkey $cadir/private/ca-masa.key \  
-in $cadir/certs/ca-masa.crt -export \  
-out $cadir/certs/ca-masa-comb.pfx  
  
# converteer authority pkcs12 file to pem for masa  
echo "converteer authority pkcs12 file to pem for masa"  
openssl pkcs12 -passin pass:watnietweet \  
-passout pass:watnietweet\  
-in $cadir/certs/ca-masa-comb.pfx \  
-out $cadir/certs/ca-masa-comb.crt -nodes  
  
#show certificate in pledge combined certificate  
openssl x509 -in $cadir/certs/ca-masa-comb.crt -text  
  
#  
# Certificate for Pledge derived from MASA certificate  
#  
echo "#####"  
echo "create pledge keys and certificates "
```

```
echo "#####"

# Pledge derived Certificate

echo "create pledge derived certificate using ecDSA with sha256"
openssl ecparam -name prime256v1 -genkey \
  -noout -out $dir/private/pledge.key

echo "create pledge certificate request"
openssl req -nodes -new -sha256 \
  -key $dir/private/pledge.key -out $dir/csr/pledge.csr \
  -subj \
  "/C=NL/ST=NB/L=Helmond/O=vanderstok/OU=manufacturing\
  /CN=uuid:$DevID/$serialNumber"

# Sign pledge derived Certificate
echo "sign pledge derived certificate "
openssl ca -config $cnfdir/openssl-pledge.cnf \
  -extensions 8021ar_idevid\
  -days 365 -in $dir/csr/pledge.csr -out $dir/certs/pledge.crt

# Add pledge key and pledge certificate to pkcs12 file
echo "Add pledge key and pledge certificate to pkcs12 file"
openssl pkcs12 -passin pass:watnietweet\
  -passout pass:watnietweet\
  -inkey $dir/private/pledge.key \
  -in $dir/certs/pledge.crt -export \
  -out $dir/certs/pledge-comb.pfx

# converteer pledge pkcs12 file to pem
echo "converteer pledge pkcs12 file to pem"
openssl pkcs12 -passin pass:watnietweet \
  -passout pass:watnietweet\
  -in $dir/certs/pledge-comb.pfx \
  -out $dir/certs/pledge-comb.crt -nodes

#show certificate in pledge-comb.crt
openssl x509 -in $dir/certs/pledge-comb.crt -text

#show private key in pledge-comb.crt
openssl ecparam -name prime256v1 \
  -in $dir/certs/pledge-comb.crt -text
```

The xxxx-comb certificates have been generated as required by libcoap for the DTLS connection generation.

## C.1. Pledge, Registrar and MASA keys

This first section documents the public and private keys used in the subsequent test vectors below. These keys come from test code and are not used in any production system, and should only be used only to validate implementations.

## C.1.1. Pledge private key

```
-----BEGIN PRIVATE KEY-----
MIGHAgEAMBMGBYqGSM49AgEGCCqGSM49AwEHBG0wawIBAQQgIpP20ud7muTl460b
xFzupPkaMoaCIIIFwSOf0hvhQByhRANCAASKnIauvAtx6ZFWQniQQqvP0Zpdaudy
Ve6Vrc80AjjWRGnN3oyQ0rnr5dXynfG2xq8+cY+uGwTrAJYp9OyoZCAs
-----END PRIVATE KEY-----
Private-Key: (256 bit)
priv:
  22:93:f6:d2:e7:7b:9a:e4:e5:e3:ad:1b:c4:5c:ee:
  a4:f9:1a:32:86:82:20:82:05:c1:23:9f:d2:1b:e1:
  40:1c
pub:
  04:8a:9c:86:ae:bc:0b:71:e9:91:56:42:78:90:3a:
  ab:cf:d1:9a:5d:6a:e7:72:55:ee:95:ad:cf:34:02:
  3c:96:44:69:cd:de:8c:90:d2:b9:eb:e5:d5:f2:9d:
  f1:b6:c6:af:3e:71:8f:ae:1b:04:eb:00:96:29:f4:
  ec:a8:64:20:2c
ASN1 OID: prime256v1
NIST CURVE: P-256
```

## C.1.2. Registrar private key

```
-----BEGIN PRIVATE KEY-----
MIGHAgEAMBMGBYqGSM49AgEGCCqGSM49AwEHBG0wawIBAQQgHCCOKLhln+l8pLnx
gWtMUm7zRY4ugkznuFimYDKbrNihRANCAARqJKniS+I00XrUfnYMLLXh3E7hFa2J
ESrUpqZLsb9o+Rd9cOkQnLSMmw3H3yZBGZ0MLb/yHtWEA4rIP0eBvhOO
-----END PRIVATE KEY-----
Private-Key: (256 bit)
priv:
  1c:20:8e:28:b8:65:9f:e9:7c:a4:b9:f1:81:6b:4c:
  52:6e:f3:45:8e:2e:82:4c:e7:b8:58:a6:60:32:9b:
  ac:d8
pub:
  04:6a:24:a9:e2:4b:e2:34:d1:7a:d4:7e:76:0c:94:
  b5:e1:dc:4e:e1:15:ad:89:11:2a:d4:a6:a6:4b:b1:
  bf:68:f9:17:7d:70:e9:10:9c:b4:8c:9b:0d:c7:df:
  26:41:19:9d:0c:2d:bf:f2:1e:d5:84:03:8a:c8:3f:
  47:81:be:13:8e
ASN1 OID: prime256v1
NIST CURVE: P-256
```

## C.1.3. MASA private key

```
-----BEGIN PRIVATE KEY-----
MIGHAgEAMBMGBYqGSM49AgEGCCqGSM49AwEHBG0wawIBAQQgQODnSgB7xR/aa3Ea
JrPGz9lZhJ1aEc/56OEPiBr86SKhRANCAASB9HLsnEeyjtHrODNBANNi9khQ2gLQ
VrIie8hLgFmVdwfQwliMPPI8WwCDeVTaDdGwr6HC6M0sO9CGRZ+JcwrL
-----END PRIVATE KEY-----
Private-Key: (256 bit)
priv:
  40:e0:e7:4a:00:7b:c5:1f:da:6b:71:1a:26:b3:c6:
  cf:d9:59:84:9d:5a:11:cf:f9:e8:e1:0f:88:1a:fc:
  e9:22
pub:
  04:81:f4:72:ec:9c:47:b2:8e:d1:eb:38:33:41:00:
  d3:62:f6:48:50:da:02:d0:56:b2:22:7b:c8:4b:80:
  59:95:77:07:d0:c3:58:8c:3c:f2:3c:5b:00:83:79:
  54:da:0d:d1:b0:af:a1:c2:e8:cd:2c:3b:d0:86:45:
  9f:89:73:0a:cb
ASN1 OID: prime256v1
NIST CURVE: P-256
```

## C.2. Pledge, Registrar and MASA certificates

Below the certificates that accompany the keys. The certificate description is followed by the hexadecimal DER of the certificate

## C.2.1. Pledge IDevID certificate

## Certificate:

## Data:

```
Version: 3 (0x2)
Serial Number: 4822678189204992 (0x11223344556600)
Signature Algorithm: ecdsa-with-SHA256
Issuer: C=NL, ST=NB, L=Helmond, O=vanderstok,
        OU=manufacturer, CN=masa.stok.nl
Validity
  Not Before: Sep  9 07:42:03 2020 GMT
  Not After : Dec 31 23:59:59 9999 GMT
Subject: C=NL, ST=NB, L=Helmond, O=vanderstok,
        OU=manufacturing,
CN=uuid:pledge.1.2.3.4/serialNumber=pledge.1.2.3.4
Subject Public Key Info:
  Public Key Algorithm: id-ecPublicKey
  Public-Key: (256 bit)
  pub:
    04:8a:9c:86:ae:bc:0b:71:e9:91:56:42:78:90:3a:
    ab:cf:d1:9a:5d:6a:e7:72:55:ee:95:ad:cf:34:02:
    3c:96:44:69:cd:de:8c:90:d2:b9:eb:e5:d5:f2:9d:
    f1:b6:c6:af:3e:71:8f:ae:1b:04:eb:00:96:29:f4:
    ec:a8:64:20:2c
  ASN1 OID: prime256v1
  NIST CURVE: P-256
X509v3 extensions:
  X509v3 Basic Constraints:
    CA:FALSE
  X509v3 Subject Key Identifier:
59:B1:E1:19:F4:68:53:E9:0E:7C:9F:29:D0:FB:5B:1F:AC:C3:82:49
  X509v3 Authority Key Identifier:
    keyid:
22:BC:B8:20:D9:C5:6D:2D:5B:B3:BB:64:8B:E0:8B:A7:86:5E:CE:B4

  X509v3 Key Usage: critical
    Digital Signature, Key Encipherment
Signature Algorithm: ecdsa-with-SHA256
30:45:02:20:4d:fd:a8:83:78:31:d2:62:a4:e5:48:a2:e0:a7:
3b:c5:14:e9:7e:46:13:45:bc:30:fd:1d:e5:d6:63:3e:d8:f4:
02:21:00:a8:e5:1e:c2:79:77:90:fc:40:a8:7a:bf:b1:bd:81:
8b:ee:d7:56:1a:04:4d:8f:c8:3d:76:5f:4d:6e:36:a2:c2
```

This is the hexadecimal representation in (request-)voucher examples referred to as pledge-cert-hex.

30820254308201faa003020102020711223344556600300a06082a8648ce3d04  
0302306f310b3009060355040613024e4c310b300906035504080c024e423110  
300e06035504070c0748656c6d6f6e6431133011060355040a0c0a76616e6465  
7273746f6b31153013060355040b0c0c6d616e75666163747572657231153013  
06035504030c0c6d6173612e73746f6b2e6e6c3020170d323030393039303734  
3230335a180f39393939313233313233353935395a308190310b300906035504  
0613024e4c310b300906035504080c024e423110300e06035504070c0748656c  
6d6f6e6431133011060355040a0c0a76616e64657273746f6b31163014060355  
040b0c0d6d616e75666163747572696e67311c301a06035504030c1375756964  
3a706c656467652e312e322e332e34311730150603550405130e706c65646765  
2e312e322e332e343059301306072a8648ce3d020106082a8648ce3d03010703  
4200048a9c86aebc0b71e991564278903aabcfd19a5d6ae77255ee95adcf3402  
3c964469cdde8c90d2b9ebe5d5f29df1b6c6af3e718fae1b04eb009629f4eca8  
64202ca35d305b30090603551d1304023000301d0603551d0e0416041459b1e1  
19f46853e90e7c9f29d0fb5b1facc38249301f0603551d2304183016801422bc  
b820d9c56d2d5bb3bb648be08ba7865eceb4300e0603551d0f0101ff04040302  
05a0300a06082a8648ce3d040302034800304502204dfda8837831d262a4e548  
a2e0a73bc514e97e461345bc30fd1de5d6633ed8f4022100a8e51ec2797790fc  
40a87abfblbd818beed7561a044d8fc83d765f4d6e36a2c2

C.2.2. Registrar Certificate

## Certificate:

## Data:

```
Version: 3 (0x2)
Serial Number:
  39:73:74:f3:fa:81:2a:0d:37:10:3b:68:c1:84:81:c5:01:bc:7c:fe
Signature Algorithm: ecdsa-with-SHA256
Issuer: C=NL, ST=NB, L=Helmond, O=vanderstok,
  OU=consultancy, CN=registrar.stok.nl
Validity
  Not Before: Sep  9 07:42:03 2020 GMT
  Not After : Sep  9 07:42:03 2021 GMT
Subject: C=NL, ST=NB, L=Helmond, O=vanderstok,
  OU=consultancy, CN=registrar.stok.nl
Subject Public Key Info:
  Public Key Algorithm: id-ecPublicKey
  Public-Key: (256 bit)
  pub:
    04:6a:24:a9:e2:4b:e2:34:d1:7a:d4:7e:76:0c:94:
    b5:e1:dc:4e:e1:15:ad:89:11:2a:d4:a6:a6:4b:b1:
    bf:68:f9:17:7d:70:e9:10:9c:b4:8c:9b:0d:c7:df:
    26:41:19:9d:0c:2d:bf:f2:1e:d5:84:03:8a:c8:3f:
    47:81:be:13:8e
  ASN1 OID: prime256v1
  NIST CURVE: P-256
X509v3 extensions:
  X509v3 Subject Key Identifier:
25:CD:93:71:B5:A1:5F:6D:1E:E8:C3:7A:51:13:BE:0B:8F:13:2C:C2
  X509v3 Authority Key Identifier:
  keyid:
25:CD:93:71:B5:A1:5F:6D:1E:E8:C3:7A:51:13:BE:0B:8F:13:2C:C2

  X509v3 Basic Constraints:
  CA:TRUE
Signature Algorithm: ecdsa-with-SHA256
  30:46:02:21:00:a6:6d:9e:24:f9:de:08:b7:f0:cf:43:c3:c0:
  ee:57:cc:b6:60:de:ae:2e:70:cc:61:a1:a2:b3:35:35:02:5b:
  ba:02:21:00:bf:fd:74:6a:99:eb:da:01:77:fc:6c:37:95:75:
  8a:f4:a0:9f:99:8e:bc:4a:90:62:49:f0:7a:c9:65:96:dc:75
```

This the hexadecimal representation, in (request-)voucher examples referred to as regis-cert-hex

30820239308201dea0030201020214397374f3fa812a0d37103b68c18481c501  
bc7cfe300a06082a8648ce3d0403023073310b3009060355040613024e4c310b  
300906035504080c024e423110300e06035504070c0748656c6d6f6e64311330  
11060355040a0c0a76616e64657273746f6b31143012060355040b0c0b636f6e  
73756c74616e6379311a301806035504030c117265676973747261722e73746f  
6b2e6e6c301e170d3230303930393037343230335a170d323130393039303734  
3230335a3073310b3009060355040613024e4c310b300906035504080c024e42  
3110300e06035504070c0748656c6d6f6e6431133011060355040a0c0a76616e  
64657273746f6b31143012060355040b0c0b636f6e73756c74616e6379311a30  
1806035504030c117265676973747261722e73746f6b2e6e6c3059301306072a  
8648ce3d020106082a8648ce3d030107034200046a24a9e24be234d17ad47e76  
0c94b5e1dc4ee115ad89112ad4a6a64bb1bf68f9177d70e9109cb48c9b0dc7df  
2641199d0c2dbff21ed584038ac83f4781be138ea350304e301d0603551d0e04  
16041425cd9371b5a15f6d1ee8c37a5113be0b8f132cc2301f0603551d230418  
3016801425cd9371b5a15f6d1ee8c37a5113be0b8f132cc2300c0603551d1304  
0530030101ff300a06082a8648ce3d0403020349003046022100a66d9e24f9de  
08b7f0cf43c3c0ee57ccb660deae2e70cc61a1a2b33535025bba022100bffd74  
6a99ebda0177fc6c3795758af4a09f998ebc4a906249f07ac96596dc75

C.2.3. MASA Certificate

## Certificate:

## Data:

Version: 3 (0x2)

Serial Number:

70:5a:34:7e:67:d2:4d:70:b0:c6:ca:60:ff:fb:75:d9:46:82:e6:0e

Signature Algorithm: ecdsa-with-SHA256

Issuer: C=NL, ST=NB, L=Helmond, O=vanderstok,

OU=manufacturer, CN=masa.stok.nl

Validity

Not Before: Sep 9 07:42:03 2020 GMT

Not After : Sep 9 07:42:03 2021 GMT

Subject: C=NL, ST=NB, L=Helmond, O=vanderstok,

OU=manufacturer, CN=masa.stok.nl

Subject Public Key Info:

Public Key Algorithm: id-ecPublicKey

Public-Key: (256 bit)

pub:

04:81:f4:72:ec:9c:47:b2:8e:d1:eb:38:33:41:00:

d3:62:f6:48:50:da:02:d0:56:b2:22:7b:c8:4b:80:

59:95:77:07:d0:c3:58:8c:3c:f2:3c:5b:00:83:79:

54:da:0d:d1:b0:af:a1:c2:e8:cd:2c:3b:d0:86:45:

9f:89:73:0a:cb

ASN1 OID: prime256v1

NIST CURVE: P-256

X509v3 extensions:

X509v3 Subject Key Identifier:

22:BC:B8:20:D9:C5:6D:2D:5B:B3:BB:64:8B:E0:8B:A7:86:5E:CE:B4

X509v3 Authority Key Identifier:

keyid:

22:BC:B8:20:D9:C5:6D:2D:5B:B3:BB:64:8B:E0:8B:A7:86:5E:CE:B4

X509v3 Basic Constraints:

CA:TRUE

Signature Algorithm: ecdsa-with-SHA256

30:45:02:20:04:ac:8d:48:62:a2:a5:04:4f:61:fd:38:83:53:

9f:00:e7:d6:4b:4d:30:1b:84:29:d4:2d:35:58:b0:a0:0c:7d:

02:21:00:8c:f1:f4:f9:a2:11:fe:64:46:a9:87:9f:58:ca:ea:

da:4f:0a:42:32:c2:6a:e8:c5:9d:62:c0:67:f0:b8:44:43

This is the hexadecimal representation, in (request-)voucher examples referred to as masa-cert-hex.

```
30820230308201d6a0030201020214705a347e67d24d70b0c6ca60fffb75d946
82e60e300a06082a8648ce3d040302306f310b3009060355040613024e4c310b
300906035504080c024e423110300e06035504070c0748656c6d6f6e64311330
11060355040a0c0a76616e64657273746f6b31153013060355040b0c0c6d616e
7566616374757265723115301306035504030c0c6d6173612e73746f6b2e6e6c
301e170d3230303930393037343230335a170d3231303930393037343230335a
306f310b3009060355040613024e4c310b300906035504080c024e423110300e
06035504070c0748656c6d6f6e6431133011060355040a0c0a76616e64657273
746f6b31153013060355040b0c0c6d616e756661637475726572311530130603
5504030c0c6d6173612e73746f6b2e6e6c3059301306072a8648ce3d02010608
2a8648ce3d0301070342000481f472ec9c47b28ed1eb38334100d362f64850da
02d056b2227bc84b8059957707d0c3588c3cf23c5b00837954da0dd1b0afalc2
e8cd2c3bd086459f89730acba350304e301d0603551d0e0416041422bcb820d9
c56d2d5bb3bb648be08ba7865eceb4301f0603551d2304183016801422bcb820
d9c56d2d5bb3bb648be08ba7865eceb4300c0603551d13040530030101ff300a
06082a8648ce3d0403020348003045022004ac8d4862a2a5044f61fd3883539f
00e7d64b4d301b8429d42d3558b0a00c7d0221008cf1f4f9a211fe6446a9879f
58caeada4f0a4232c26ae8c59d62c067f0b84443
```

### C.3. COSE signed voucher request from pledge to Registrar

In this example the voucher request has been signed by the pledge, and has been sent to the JRC over CoAPS. This example uses the proximity-registrar-cert mechanism to request a voucher that pins the certificate of the registrar.

```
POST coaps://registrar.example.com/est/rv
(Content-Format: application/voucher-cose+cbor)
signed_request_voucher
```

The payload `signed_request_voucher` is shown as hexadecimal dump (with lf added):

```

d28444a101382ea1045820f8926f5ba385b7bccf23592b97a73c1b00bffc01023
0f647f06960870b1fd6ee5902aca11909c5a61909c77818323032302d31302d35
5431333a34363a31332d30303a30301909c97818323032322d31302d355431333
a34363a31332d30303a30301909c6021909cc5029c7bafb81a2c6160d3357d229
11f5101909d26e706c656467652e312e322e332e341909cf59023d30820239308
201dea0030201020214397374f3fa812a0d37103b68c18481c501bc7cfe300a06
082a8648ce3d0403023073310b3009060355040613024e4c310b3009060355040
80c024e423110300e06035504070c0748656c6d6f6e6431133011060355040a0c
0a76616e64657273746f6b31143012060355040b0c0b636f6e73756c74616e637
9311a301806035504030c117265676973747261722e73746f6b2e6e6c301e170d
3230303930393037343230335a170d3231303930393037343230335a3073310b3
009060355040613024e4c310b300906035504080c024e423110300e0603550407
0c0748656c6d6f6e6431133011060355040a0c0a76616e64657273746f6b31143
012060355040b0c0b636f6e73756c74616e6379311a301806035504030c117265
676973747261722e73746f6b2e6e6c3059301306072a8648ce3d020106082a864
8ce3d030107034200046a24a9e24be234d17ad47e760c94b5e1dc4ee115ad8911
2ad4a6a64bb1bf68f9177d70e9109cb48c9b0dc7df2641199d0c2dbff21ed5840
38ac83f4781be138ea350304e301d0603551d0e0416041425cd9371b5a15f6dle
e8c37a5113be0b8f132cc2301f0603551d2304183016801425cd9371b5a15f6dl
ee8c37a5113be0b8f132cc2300c0603551d13040530030101ff300a06082a8648
ce3d0403020349003046022100a66d9e24f9de08b7f0cf43c3c0ee57ccb660dea
e2e70cc61a1a2b33535025bba022100bffd746a99ebda0177fc6c3795758af4a0
9f998ebc4a906249f07ac96596dc7558473045022100fc28be418e5f25152590e
872b4bbdbe334cd31dlebb0a806e7a172cad5cfff604022056ee414ddac438e7f5
1dda9ddf6ec6e31a78cdde6574717fe46dd3a7c60f5bb5

```

The representation of signed\_voucher\_request in CBOR diagnostic format is:

```

Diagnose(signed_request_voucher) =
18 ([
h'A101382E',      # {"alg": -47}
{4:h'F8926F5BA385B7BCCF23592B97A73C1B00BFFC010230F647F06960870B1F
D6EE'},
h'request_voucher'
h'3045022100FC28BE418E5F25152590E872B4BBDDBE334CD31D1EBB0A806E7A17
2CAD5CFF604022056EE414DDAC438E7F51DDA9DDF6EC6E31A78CDDE6574717FE4
6DD3A7C60F5BB5'] )

```

```

Diagnose(request_voucher) =
{2501: {2503: "2020-10-5T13:46:13-00:00",
2505: "2022-10-5T13:46:13-00:00",
2502: 2,
2508: h'29C7BAFB81A2C6160D3357D22911F510',
2514: "pledge.1.2.3.4",
2511: h'regis-cert-hex'}},

```

## C.4. COSE signed voucher request from Registrar to MASA

In this example the voucher request has been signed by the JRC using the private key from Appendix C.1.2. Contained within this voucher request is the voucher request from the pledge to JRC.

```
POST coaps://masa.example.com/est/rv
(Content-Format: application/voucher-cose+cbor)
signed_masa_request_voucher
```

The payload `signed_masa_voucher_request` is shown as hexadecimal dump (with lf added):

```
d28444a101382ea1045820b86ae808f79af17e5948cbda731f158d04bd091c73f
485f2409eac08ee7ddb5c5903fea11909c5a61909c77818323032302d31302d35
5431333a34363a31332d30303a30301909c97818323032322d31302d355431333
a34363a31332d30303a30301909cc5029c7bafb81a2c6160d3357d22911f51019
09d26e706c656467652e312e322e332e341909ca586b433d4e4c2c2053543d4e4
22c204c3d48656c6d6f6e642c204f3d76616e64657273746f6b2c204f553d6d61
6e75666163747572696e672c20434e3d757569643a706c656467652e312e322e3
32e342c2073657269616c4e756d6265723d706c656467652e312e322e332e3419
09ce590323d28444a101382ea1045820f8926f5ba385b7bccf23592b97a73c1b0
0bffc010230f647f06960870b1fd6ee5902aca11909c5a61909c7781832303230
2d31302d355431333a34363a31332d30303a30301909c97818323032322d31302
d355431333a34363a31332d30303a30301909c6021909cc5029c7bafb81a2c616
0d3357d22911f5101909d26e706c656467652e312e322e332e341909cf59023d3
0820239308201dea0030201020214397374f3fa812a0d37103b68c18481c501bc
7cfe300a06082a8648ce3d0403023073310b3009060355040613024e4c310b300
906035504080c024e423110300e06035504070c0748656c6d6f6e643113301106
0355040a0c0a76616e64657273746f6b31143012060355040b0c0b636f6e73756
c74616e6379311a301806035504030c117265676973747261722e73746f6b2e6e
6c301e170d3230303930393037343230335a170d3231303930393037343230335
a3073310b3009060355040613024e4c310b300906035504080c024e423110300e
06035504070c0748656c6d6f6e6431133011060355040a0c0a76616e646572737
46f6b31143012060355040b0c0b636f6e73756c74616e6379311a301806035504
030c117265676973747261722e73746f6b2e6e6c3059301306072a8648ce3d020
106082a8648ce3d030107034200046a24a9e24be234d17ad47e760c94b5e1dc4e
e115ad89112ad4a6a64bb1bf68f9177d70e9109cb48c9b0dc7df2641199d0c2db
ff21ed584038ac83f4781be138ea350304e301d0603551d0e0416041425cd9371
b5a15f6d1ee8c37a5113be0b8f132cc2301f0603551d2304183016801425cd937
1b5a15f6d1ee8c37a5113be0b8f132cc2300c0603551d13040530030101ff300a
06082a8648ce3d0403020349003046022100a66d9e24f9de08b7f0cf43c3c0ee5
7ccb660deae2e70cc61a1a2b33535025bba022100bffd746a99ebda0177fc6c37
95758af4a09f998ebc4a906249f07ac96596dc7558473045022100fc28be418e5
f25152590e872b4bbdbe334cd31d1ebb0<<a806e7a172cad5cff604022056ee41
4ddac438e7f51dda9ddf6ec6e31a78cdde6574717fe46dd3a7c60f5bb55847304
5022047b5314c72cbb2d1212e51198061167c79e1002874cd2665a5b643fa6436
3c30022100ce49ac309f760bd0e75660a7e29edee82f0251724c124150f5326b9
b2654927c
```

The representation of signed\_masa\_voucher\_request in CBOR diagnostic format is:

```

Diagnose(signed_masa_request_voucher) =

18([
h'A101382E',      # {"alg": -47}
{4:h'B86AE808F79AF17E5948CBDA731F158D04BD091C73F485F2409EAC08EE7D
DB5C'},
h'masa_request_voucher',
h'3045022047B5314C72CBB2D1212E51198061167C79E1002874CD2665A5B643F
A64363C30022100CE49AC309F760BD0E75660A7E29EDEE82F0251724C124150F5
326B9B2654927C' ])

Diagnose(masa_request_voucher) =
{2501:
  {2503: "2020-10-5T13:46:13-00:00",
    2505: "2022-10-5T13:46:13-00:00",
    2508: h'29C7BAFB81A2C6160D3357D22911F510',
    2514: "pledge.1.2.3.4",
    2506:h'433D4E4C2C2053543D4E422C204C3D48656C6D6F6E642C
204F3D76616E64657273746F6B2C204F553D6D616E75666163747572696E672C2
0434E3D757569643A706C656467652E312E322E332E342C2073657269616C4E75
6D6265723D706C656467652E312E322E332E34',
    2510: h'request_voucher' }},

```

#### C.5. COSE signed voucher from MASA to Pledge via Registrar

The resulting voucher is created by the MASA and returned via the JRC to the Pledge. It is signed by the MASA's private key Appendix C.1.3 and can be verified by the pledge using the MASA's public key contained within the MASA certificate.

This is the raw binary signed\_voucher, encoded in hexadecimal (with lf added):

```

d28444a101382ea1045820ab59b0679fcf65d5223d4ce4266a27a9c7432702466
ff5f3648e822a64d61b145902b0a1190993a71909957818323032302d31302d35
5431333a34363a31342d30303a30301909977818323032322d31302d355431333
a34363a31342d30303a30301909940319099a5029c7bafb81a2c6160d3357d229
11f51019099e6e706c656467652e312e322e332e3419099b59023d30820239308
201dea0030201020214397374f3fa812a0d37103b68c18481c501bc7cfe300a06
082a8648ce3d0403023073310b3009060355040613024e4c310b3009060355040
80c024e423110300e06035504070c0748656c6d6f6e6431133011060355040a0c
0a76616e64657273746f6b31143012060355040b0c0b636f6e73756c74616e637
9311a301806035504030c117265676973747261722e73746f6b2e6e6c301e170d
3230303930393037343230335a170d3231303930393037343230335a3073310b3
009060355040613024e4c310b300906035504080c024e423110300e0603550407
0c0748656c6d6f6e6431133011060355040a0c0a76616e64657273746f6b31143
012060355040b0c0b636f6e73756c74616e6379311a301806035504030c117265
676973747261722e73746f6b2e6e6c3059301306072a8648ce3d020106082a864
8ce3d030107034200046a24a9e24be234d17ad47e760c94b5e1dc4ee115ad8911
2ad4a6a64bb1bf68f9177d70e9109cb48c9b0dc7df2641199d0c2dbff21ed5840
38ac83f4781be138ea350304e301d0603551d0e0416041425cd9371b5a15f6dle
e8c37a5113be0b8f132cc2301f0603551d2304183016801425cd9371b5a15f6dl
ee8c37a5113be0b8f132cc2300c0603551d13040530030101ff300a06082a8648
ce3d0403020349003046022100a66d9e24f9de08b7f0cf43c3c0ee57ccb660dea
e2e70cc61a1a2b33535025bba022100bffd746a99ebda0177fc6c3795758af4a0
9f998ebc4a906249f07ac96596dc751909960058483046022100d07cad5c2836
e7845d6d2e2652527386bd40258d20ab24b6bbce5515df915e9022100aba68a07
b2295c4b49d53f73ea370ca66f761ad5d8d8c11c19a2d505729285cb

```

The representation of signed\_voucher in CBOR diagnostic format is:

```

Diagnose (signed_voucher) =
18 ([
h'A101382E',      # {"alg": -47}
{4: h'AB59B0679FCF65D5223D4CE4266A27A9C7432702466FF5F3648E822A64D61
B14'},
h'voucher',
h'3046022100D07CAD5C2836E7845D6D2E2652527386BD40258D20AB24B6BBCE
5515DF915E9022100ABA68A07B2295C4B49D53F73EA370CA66F761AD5D8D8C11C
19A2D505729285CB' ])

```

Diagnose (voucher) =

```

{2451:
  {2453: "2020-10-5T13:46:14-00:00",
    2455: "2022-10-5T13:46:14-00:00",
    2452: 3,
    2458: h'29C7BAFB81A2C6160D3357D22911F510',
    2462: "pledge.1.2.3.4",
    2459: h'regis-cert-hex',
    2454: 0}}

```

Authors' Addresses

Michael Richardson  
Sandelman Software Works

Email: [mcr+ietf@sandelman.ca](mailto:mcr+ietf@sandelman.ca)

Peter van der Stok  
vanderstok consultancy

Email: [consultancy@vanderstok.org](mailto:consultancy@vanderstok.org)

Panos Kampanakis  
Cisco Systems

Email: [pkampana@cisco.com](mailto:pkampana@cisco.com)

ANIMA WG  
INTERNET-DRAFT  
Intended Status: Standard Track  
Expires: June 14, 2020

B. Liu  
Huawei Technologies  
X. Xiao  
A. Hecker  
MRC, Huawei Technologies  
S. Jiang  
Huawei Technologies  
Z. Despotovic  
MRC, Huawei Technologies  
December 12, 2019

Information Distribution in Autonomic Networking  
draft-liu-anima-grasp-distribution-13

Abstract

This document proposes a solution for information distribution in autonomic networks. Information distribution is categorized into two different modes: 1) instantaneous distribution; 2) publication for retrieval. In the former case, the information is sent, propagates and is disposed of after reception. In the latter case, information needs to be stored in the network.

The capabilities to distribute information are basic and fundamental needs for an autonomous network (cf. ANI [I-D.ietf-anima-reference-model]). This document describes typical use cases of information distribution in ANI and requirements to ANI, such that rich information distribution can be natively supported. The document proposes extensions to the autonomic nodes and suggests an implementation based on GRASP extensions as a protocol on the wire.

Status of this Memo

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

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/lid-abstracts.html>

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

#### Copyright and License Notice

Copyright (c) 2019 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	3
2.	Terminology	3
3.	Requirements of Enriched Information Distribution	4
4.	Node Behaviors	5
4.1	Instant Information Distribution (IID) Sub-module	5
4.2	Asynchronous Information Distribution (AID) Sub-module	6
4.3	Summary	10
5.	Extending GRASP for Information Distribution	10
5.1	Realizing Instant P2P Transmission	10
5.2	Realizing Instant Selective Flooding	11
5.3	Realizing Subscription as An Event	11
5.4	Un_Subscription Objective Option	12
5.5	Publishing Objective Option	12
6.	Security Considerations	13
7.	IANA Considerations	13
8.	References	13
8.1	Normative References	13
8.2	Informative References	13
	Authors' Addresses	14
	Appendix A. Real-world Use Cases of Information Distribution	15
	Appendix B. Information Distribution Module in ANI	18
	Appendix C. Asynchronous Information Distribution Integrated with GRASP APIs	18

## 1 Introduction

In an autonomic network, autonomic functions (AFs) running on autonomic nodes constantly exchange information, e.g. AF control/management signaling or AF data exchange. This document discusses the information distribution capability of such exchanges between AFs.

Depending on the number of participants, the information can be distributed in in the following scenarios:

- 1) Point-to-point (P2P) Communication: information is exchanged between parties, i.e. two nodes.
- 2) One-to-Many Communication: information exchanges involve an information source and multiple receivers.

The approaches to information distribution can be chiefly categorized into two basic modes:

- 1) An instantaneous mode (push): a source sends the actual content (e.g. control/management signaling, synchronization data and so on) to all interested receiver(s) immediately. Generally, some preconfiguration is required, as nodes interested in this information must be already known to all nodes in the sense that any receiving node must be able to decide, to which nodes this data is to be sent.
- 2) An asynchronous mode (delayed pull): here, a source publishes the content in some form in the network, which may later be looked for, found and retrieved by some endpoints in the AN. Here, depending on the size of the content, either the whole content or only its metadata might be published into the AN. In the latter case the metadata (e.g. a content descriptor, e.g. a key, and a location in the ANI) may be used for the actual retrieval. Importantly, the source, i.e. here publisher, needs to be able to determine the node, where the information (or its metadata) can be stored.

To avoid repetitive implementations by each AF developer, this document opts for a common support for information distribution implemented as a basic ANI capability, therefore available to all AFs. In fact, GRASP already provides part of the capabilities.

Regardless, an AF may still define and implement its own information distribution capability. Such a capability may then be advertised using the common information distribution capability defined in this document. Overall, ANI nodes and AFs may decide, which of the

information distribution mechanisms they want to use for which type of information, according to their own preferences (e.g. semantic routing table, etc.)

This document first analyzes requirements for information distribution in autonomic networks (Section 3) and then discuss the relevant node behavior (Section 4). After that, the required GRASP extensions are formally introduced (Section 5).

## 2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

### 3. Requirements for Information Distribution in ANI

The question of information distribution in an autonomic network can be discussed through particular use cases or more generally. Depending on the situation it can be quite simple or might require more complex provisions.

Indeed, in the simplest case, the information can be sent:

- 1) at once (in one packet, in one flow),
- 2) straightaway (send-and-forget),
- 3) to all nodes.

Presuming 1), 2) and 3) hold, information distribution in smaller or scarce topologies can be implemented using broadcast, i.e. unconstrained flooding. For reasons well-understood, this approach has its limits in larger and denser networks. In this case, a graph can be constructed such that it contains every node exactly once (e.g. a spanning tree), still allowing to distribute any information to all nodes straightaway. Multicast tree construction protocols could be used in this case. There are reasonable use cases for such scenarios, as presented in Appendix B.

A more complex scenario arises, if only 1) and 2) hold, but the information only concerns a subset of nodes. Then, some kind of selection becomes required, to which nodes the given information should be distributed. Here, a further distinction is necessary; notably, if the selection of the target nodes is with respect to the nature or position of the node, or whether it is with respect to the information content. If the first, some knowledge about the node types, its topological position, etc (e.g. the routing information within ANI) can be used to distinguish nodes accordingly. For instance, edge nodes and forwarding nodes can be distinguished in this way. If the distribution scope is primarily to be defined by the information elements, then a registration / join / subscription or label distribution mechanism is unavoidable. This would be the case, for instance, if the AFs can be dynamically deployed on nodes, and the information is majorily destined to the AFs. Then, depending on the current AF deployment, the distribution scope must be adjusted as well.

If only 1) holds, but the information content might be required again and again, or might not yet be fully available, then more complex mechanisms might be required to store the information within the network for later, for further redistribution, and for notification of interested nodes. Examples for this include distribution of reconfiguration information for different AF instances, which might not require an immediate action, but only an eventual update of the parameters. Also, in some situations, there could be a significant

delay between the occurrence of a new event and the full content availability (e.g. if the processing requires a lot of time).

Finally, none of the three might hold. Then, along with the subscription and notification, the actual content might be different from its metadata, i.e. some description of the content and, possibly, its location. The fetching can then be implemented in different, appropriate ways, if necessary as a complex transport session.

In essence, as flooding is usually not an option, and the interest of nodes for particular information elements can change over time, ANI should support autonomies also for the information distribution.

This calls for autonomic mechanisms in the ANI, allowing participating nodes to 1) advertise or publish 2) look for or subscribe to 3) store 4) fetch/retrieve 5) instantaneously push information elements.

In the following cases, situations depicting diverse information distribution needs are discussed.

1) Long Communication Intervals. The actual sending of the information is not necessarily instantaneous with some event. Advanced AFs may involve into longer jobs/tasks (e.g. database lookup, authentication etc.) when processing requests, and might not be able to reply immediately. Instead of actively waiting for the reply, a better way for an interested AF might be to get notified, when the reply is finally available.

2) Common Interest Distribution. AFs may share interest in common information. For example, the network intent will be distributed to network nodes enrolled, which is usually one-to-many scenario. Intent distribution can also be performed by an instant flooding (e.g. via GRASP) to every network node. However, because of network dynamics, not every node can be just ready at the moment when the network intent is broadcast. Also, a flooding often does not cover all network nodes as there is usually a limitation on the hop number. In fact, nodes may join in the network sequentially. In this situation, an asynchronous communication model could be a better choice where every (newly joining) node can subscribe the intent information and will get notified if it is ready (or updated).

3) Distributed Coordination. With computing and storage resources on autonomic nodes, alive AFs not only consume but also generate data information. An example is AFs coordinating with each other as distributed schedulers, responding to service requests and

distributing tasks. It is critical for those AFs to make correct decisions based on local information, which might be asymmetric as well. AFs may also need synthetic/aggregated data information (e.g. statistic info, like average values of several AFs, etc.) to make decisions. In these situations, AFs will need an efficient way to form a global view of the network (e.g. about resource consumption, bandwidth and statistics). Obviously, purely relying on instant communication model is inefficient, while a scalable, common, yet distributed data layer, on which AFs can store and share information in an asynchronous way, should be a better choice.

Therefore, for ANI, in order to support various communication scenarios, an information distribution module is required, and both instantaneous and asynchronous communication models should be supported. Some real-world use cases are introduced in Appendix A.

#### 4. Node Behaviors

In this section, how a node should behave in order to support the two identified modes of information distribution is discussed. An ANI is a distributed system, so the information distribution module must be implemented in a distributed way as well.

##### 4.1 Instant Information Distribution (IID) Sub-module

In this case, An information sender directly specifies the information receiver(s). The instant information distribution sub-module will be the main element.

###### 4.1.1 Instant P2P Communication

IID sub-module performs instant information transmission for ASAs running in an ANI. In specific, IID sub-module will have to retrieve the address of the information receiver specified by an ASA, then deliver the information to the receiver. Such a delivery can be done either in a connectionless or a connection-oriented way.

Current GRASP provides the capability to support instant P2P synchronization for ASAs. A P2P synchronization is a use case of P2P information transmission. However, as mention in Section 3, there are some scenarios where one node needs to transmit some information to another node(s). This is different to synchronization because after transmitting the information, the local status of the information does not have to be the same as the information sent to the receiver. This is not directly support by existing GRASP.

#### 4.1.2 Instant Flooding Communication

IID sub-module finishes instant flooding for ASAs in an ANI. Instant flooding is for all ASAs in an ANI. An information sender has to specify a special destination address of the information and broadcast to all interfaces to its neighbors. When another IID sub-module receives such a broadcast, after checking its TTL, it further broadcast the message to the neighbors. In order to avoid flooding storms in an ANI, usually a TTL number is specified, so that after a pre-defined limit, the flooding message will not be further broadcast again.

In order to avoid unnecessary flooding, a selective flooding can be done where an information sender wants to send information to multiple receivers at once. When doing this, sending information needs to contain criteria to judge on which interfaces the distributed information should and should not be sent. Specifically, the criteria contain:

- o Matching Condition: a set of matching rules such as addresses of recipients, node features and so on.
- o Action: what the node needs to do when the Matching Condition is fulfilled. For example, the action could be forwarding or discarding the distributed message.

Sent information must be included in the message distributed from the sender. The receiving node reacts by first checking the carried Matching Condition in the message to decide who should consume the message, which could be either the node itself, some neighbors or both. If the node itself is a recipient, Action field is followed; if a neighbor is a recipient, the message is sent accordingly.

An exemplary extension to support selective flooding on GRASP is described in Section 5.

#### 4.2 Asynchronous Information Distribution (AID) Sub-module

In asynchronous information distribution, sender(s) and receiver(s) are not immediately specified while they may appear in an asynchronous way. Firstly, AID sub-module enables that the information can be stored in the network; secondly, AID sub-module provides an information publication and subscription (Pub/Sub) mechanism for ASAs.

As sketched in the previous section, in general each node requires two modules: 1) Information Storage (IS) module and 2) Event Queue (EQ) module in the information distribution module. Details of the

two modules are described in the following sections.

#### 4.2.1 Information Storage

IS module handles how to save and retrieve information for ASAs across the network. The IS module uses a syntax to index information, generating the hash index value (e.g. a hash value) of the information and mapping the hash index to a certain node in ANI. Note that, this mechanism can use existing solutions. Specifically, storing information in an ANIMA network will be realized in the following steps.

- 1) ASA-to-IS Negotiation. An ASA calls the API provided by information distribution module (directly supported by IS sub-module) to request to store the information somewhere in the network. The IS module performs various checks of the request (e.g. permitted information size).
- 2) Storing Peer Mapping. The information block will be handled by the IS module in order to calculate/map to a peer node in the network. Since ANIMA network is a peer-to-peer network, a typical way is to use distributed hash table (DHT) to map information to a unique index identifier. For example, if the size of the information is reasonable, the information block itself can be hashed, otherwise, some meta-data of the information block can be used to generate the mapping.
- 3) Storing Peer Negotiation Request. Negotiation request of storing the information will be sent from the IS module to the IS module on the destination node. The negotiation request contains parameters about the information block from the source IS module. According to the parameters as well as the local available resource, the requested storing peer will send feedback the source IS module.
- 4) Storing Peer Negotiation Response. Negotiation response from the storing peer is sent back to the source IS module. If the source IS module gets confirmation that the information can be stored, source IS module will prepare to transfer the information block; otherwise, a new storing peer must be discovered (i.e. going to step 7).
- 5) Information Block Transfer. Before sending the information block to the storing peer that already accepts the request, the IS module of the source node will check if the information block can be afforded by one GRASP message. If so, the information block will be directly sent by calling a GRASP API. Otherwise, a bulk data transmission is needed. For that, there are multiple ways to

do it.

The first option is to utilize one of existing protocols that is independent of the GRASP stack. For example, a session connectivity can be established to the storing peer, and over the connection the bulky data can be transmitted part by part. In this case, the IS module should support basic TCP-based session protocols such as HTTP(s) or native TCP.

The second option is to directly use GRASP itself for bulky data transferring. [I-D.carpenter-anima-grasp-bulk-04].

6) Information Writing. Once the information block (or a smaller block) is received, the IS module of the storing peer will store the data block in the local storage is accessible.

7) (Optional) New Storing Peer Discovery. If the previously selected storing peer is not available to store the information block, the source IS module will have to identify a new destination node to start a new negotiation. In this case, the discovery can be done by using discovery GRASP API to identify a new candidate, or more complex mechanisms can be introduced.

Similarly, Getting information from an ANI will be realized in the following steps.

1) ASA-to-IS Request. An ASA accesses the IS module via the APIs exposed by the information distribution module. The key/index of the interested information will be sent to the IS module. An assumption here is that the key/index should be known to an ASA before an ASA can ask for the information. This relates to the publishing/subscribing of the information, which are handled by other modules (e.g. Event Queue with Pub/Sub supported by GRASP).

2) Storing Peer Mapping. IS module maps the key/index of the requested information to a peer that stores the information, and prepares the information request. The mapping here follows the same mechanism when the information is stored.

3) Retrieval Negotiation Request. The source IS module sends a request to the storing peer and asks if such an information object is available.

4) Retrieval Negotiation Response. The storing peer checks the key/index of the information in the request, and replies to the source IS module. If the information is found and the information block can be afforded within one GRASP message, the information will be sent together with the response to the source IS module.

5) (Optional) New Destination Request. If the information is not found after the source IS module gets the response from the originally identified storing peer, the source IS module will have to discover the location of the requested information.

IS module can reuse distributed databases and key value stores like NoSQL, Cassandra, DHT technologies. storage and retrieval of information are all event-driven responsible by the EQ module.

4.2.2 Event Queue The Event Queue (EQ) module is to help ASAs to publish information to the network and subscribe to interested information in asynchronous scenarios. In an ANI, information generated on network nodes is an event labeled with an event ID, which is semantically related to the topic of the information. Key features of EQ module are summarized as follows.

1) Event Group: An EQ module provides isolated queues for different event groups. If two groups of AFs could have completely different purposes, the EQ module allows to create multiple queues where only AFs interested in the same topic will be aware of the corresponding event queue.

2) Event Prioritization: Events can have different priorities in ANI. This corresponds to how much important or urgent the event implies. Some of them are more urgent than regular ones. Prioritization allows AFs to differentiate events (i.e. information) they publish or subscribe to.

3) Event Matching: an information consumer has to be identified from the queue in order to deliver the information from the provider. Event matching keeps looking for the subscriptions in the queue to see if there is an exact published event there. Whenever a match is found, it will notify the upper layer to inform the corresponding ASAs who are the information provider and subscriber(s) respectively.

The EQ module on every network node operates as follows.

1) Event ID Generation: If information of an ASA is ready, an event ID is generated according to the content of the information. This is also related to how the information is stored/saved by the IS module introduced before. Meanwhile, the type of the event is also specified where it can be of control purpose or user plane data.

2) Priority Specification: According to the type of the event, the ASA may specify its priority to say how this event is to be processed. By considering both aspects, the priority of the event will be determined.

3) Event Enqueue: Given the event ID, event group and its priority, a queue is identified locally if all criteria can be satisfied. If there is such a queue, the event will be simply added into the queue, otherwise a new queue will be created to accommodate such an event.

4) Event Propagation: The published event will be propagated to the other network nodes in the ANIMA domain. A propagation algorithm can be employed to optimize the propagation efficiency of the updated event queue states.

5) Event Match and Notification: While propagating updated event states, EQ module in parallel keeps matching published events and its interested consumers. Once a match is found, the provider and subscriber(s) will be notified for final information retrieval.

The category of event priority is defined as the following. In general, there are two event types:

1) Network Control Event: This type of events are defined by the ANI for operational purposes on network control. A pre-defined priority levels for required system messages is suggested. For highest level to lowest level, the priority value ranges from NC\_PRIOR\_HIGH to NC\_PRIOR\_LOW as integer values. The NC\_PRIOR\_\* values will be defined later according to the total number system events required by the ANI;

2) Custom ASA Event: This type of events are defined by the ASAs of users. This specifies the priority of the message within a group of ASAs, therefore it is only effective among ASAs that join the same message group. Within the message group, a group header/leader has to define a list of priority levels ranging from CUST\_PRIOR\_HIGH to CUST\_PRIOR\_LOW. Such a definition completely depends on the individual purposes of the message group.

When a system message is delivered, its event type and event priority value have to be both specified;

Event contains the address where the information is stored, after a subscriber is notified, it directly retrieves the information from the given location.

#### 4.3 Summary

In summary, the general requirements for the information distribution module on each autonomic node are realized by two sub-modules handling instant communications and asynchronous communications, respectively. For instantaneous mode, node requirements are simple,

calling for support for additional signaling. With minimum efforts, reusing the existing GRASP is possible.

For asynchronous mode, information distribution module uses new primitives on the wire, and implements an event queue and an information storage mechanism. An architectural consideration on ANI with the information distribution module is briefly discussed in Appendix B.

## 5. Extending GRASP for Information Distribution

### 5.1 Realizing Instant P2P Transmission

This could be a new message in GRASP. In fragmentary CDDL, an Un-solicited Synchronization message follows the pattern:

```
unsolicited_synch-message = [M_UNOLSOLIDSYNCH, session-id,
                             objective]
```

A node MAY actively send a unicast Un-solicited Synchronization message with the Synchronization data, to another node. This MAY be sent to port GRASP\_LISTEN\_PORT at the destination address, which might be obtained by GRASP Discovery or other possible ways. The synchronization data are in the form of GRASP Option(s) for specific synchronization objective(s).

### 5.2 Realizing Instant Selective Flooding

Since normal flooding is already supported by GRASP, this section only defines the selective flooding extension.

In fragmentary CDDL, the selective flooding follows the pattern:

```
selective-flood-option = [O_SELECTIVE_FLOOD, +O_MATCH-CONDITION,
                          match-object, action]
```

```
O_MATCH-CONDITION = [O_MATCH-CONDITION, Obj1, match-rule, Obj2]
Obj1 = text
```

```
match-rule = GREATER / LESS / WITHIN / CONTAIN
```

```
Obj2 = text
```

```
match-object = NEIGHBOR / SELF
```

```
action = FORWARD / DROP
```

The option field encapsulates a match-condition option which represents the conditions regarding to continue or discontinue flood the current message. For the match-condition option, the Obj1 and Obj2 are to objects that need to be compared. For example, the Obj1 could be the role of the device and Obj2 could be "RSG". The match rules between the two objects could be greater, less than, within, or contain. The match-object represents of which Obj1 belongs to, it could be the device itself or the neighbor(s) intended to be flooded. The action means, when the match rule applies, the current device just continues flood or discontinues.

### 5.3 Realizing Subscription as An Event

In fragmentary CDDL, a Subscription Objective Option follows the pattern:

```
subscription-objection-option = [SUBSCRIPTION, 2, 2, subobj]
objective-name = SUBSCRIPTION

objective-flags = 2

loop-count = 2

subobj = text
```

This option MAY be included in GRASP M\_Synchronization, when included, it means this message is for a subscription to a specific object.

### 5.4 Un\_Subscription Objective Option

In fragmentary CDDL, a Un\_Subscribe Objective Option follows the pattern:

```
Unsubscribe-objection-option = [UNSUBSCRIB, 2, 2, unsubobj]
objective-name = SUBSCRIPTION
objective-flags = 2
loop-count = 2
unsubobj = text
```

This option MAY be included in GRASP M\_Synchronization, when included, it means this message is for a un-subscription to a specific object.

### 5.5 Publishing Objective Option

In fragmentary CDDL, a Publish Objective Option follows the pattern:

```
publish-objection-option = [PUBLISH, 2, 2, pubobj] objective-name
= PUBLISH
objective-flags = 2
loop-count = 2
pubobj = text
```

This option MAY be included in GRASP M\_Synchronization, when included, it means this message is for a publish of a specific object data.

## 6. Security Considerations

The distribution source authentication could be done at multiple layers:

- o Outer layer authentication: the GRASP communication is within ACP (Autonomic Control Plane, [I-D.ietf-anima-autonomic-control-plane]). This is the default GRASP behavior.
- o Inner layer authentication: the GRASP communication might not be within a protected channel, then there should be embedded protection in distribution information itself. Public key infrastructure might be involved in this case.

## 7. IANA Considerations

TBD.

## 8. References

### 8.1 Normative References

```
[I-D.ietf-anima-grasp]
    Bormann, D., Carpenter, B., and B. Liu, "A Generic
    Autonomic Signaling Protocol (GRASP)", draft-ietf-
    animagrasp-15 (Standard Track), October 2017.
```

### 8.2 Informative References

```
[RFC7575] Behringer, M., "Autonomic Networking: Definitions and
    Design Goals", RFC 7575, June 2015
```

```
[I-D.ietf-anima-autonomic-control-plane]
    Eckert, T., Behringer, M., and S. Bjarnason, "An Autonomic
    Control Plane (ACP)", draft-behringer-anima-autonomic-
    control-plane-13, December 2017.
```

```
[I-D.ietf-anima-stable-connectivity-10]
```

Eckert, T., Behringer, M., "Using Autonomic Control Plane for Stable Connectivity of Network OAM", draft-ietf-anima-stable-connectivity-10, February 2018.

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

Behringer, M., Carpenter, B., Eckert, T., Ciavaglia, L., Pierre P., Liu, B., Nobre, J., and J. Strassner, "A Reference Model for Autonomic Networking", draft-ietf-anima-reference-model-05, October 2017.

[I-D.du-anima-an-intent]

Du, Z., Jiang, S., Nobre, J., Ciavaglia, L., and M. Behringer, "ANIMA Intent Policy and Format", draft-duanima-an-intent-05 (work in progress), February 2017.

[I-D.ietf-anima-grasp-api]

Carpenter, B., Liu, B., Wang, W., and X. Gong, "Generic Autonomic Signaling Protocol Application Program Interface (GRASP API)", draft-ietf-anima-grasp-api-00 (work in progress), December 2017.

[I-D.carpenter-anima-grasp-bulk-04]

Carpenter, B., Jiang, S., Liu, B., "Transferring Bulk Data over the GeneRIC Autonomic Signaling Protocol (GRASP)", draft-carpenter-anima-grasp-bulk-04 (work in progress), July 3, 2019

[3GPP.29.500]

3GPP, "Technical Realization of Service Based Architecture", 3GPP TS 29.500 15.1.0, 09 2018

[3GPP.23.501]

3GPP, "System Architecture for the 5G System", 3GPP TS 23.501 15.2.0, 6 2018, <<http://www.3gpp.org/ftp/Specs/html-info/23501.htm>>.

[3GPP.23.502]

3GPP, "Procedures for the 5G System", 3GPP TS 23.502 15.2.0, 6 2018, <<http://www.3gpp.org/ftp/Specs/html-info/23502.htm>>.

[5GAA.use.cases]

White Paper "Toward fully connected vehicles: Edge computing for advanced automotive communications", 5GAA <<http://5gaa.org/news/toward-fully-connected-vehicles-edge-computing-for-advanced-automotive-communications/>>

## Authors' Addresses

Bing Liu  
Huawei Technologies  
Q27, Huawei Campus

No.156 Beiqing Road  
Hai-Dian District, Beijing 100095  
P.R. China

Email: leo.liubing@huawei.com

Sheng Jiang  
Huawei Technologies  
Q27, Huawei Campus  
No.156 Beiqing Road  
Hai-Dian District, Beijing 100095  
P.R. China

Email: jiangsheng@huawei.com

Xun Xiao  
Munich Research Center  
Huawei technologies  
Riesstr. 25, 80992, Muenchen, Germany

Emails: xun.xiao@huawei.com

Artur Hecker  
Munich Research Center  
Huawei technologies  
Riesstr. 25, 80992, Muenchen, Germany

Emails: artur.hecker@huawei.com

Zoran Despotovic  
Munich Research Center  
Huawei technologies  
Riesstr. 25, 80992, Muenchen, Germany

Emails: zoran.despotovic@huawei.com

## Appendix A. Real-world Use Cases of Information Distribution

The requirement analysis in Section 3 shows that generally information distribution should be better of as an infrastructure layer module, which provides to upper layer utilizations. In this section, we review some use cases from the real-world where an information distribution module with powerful functions do plays a critical role there.

#### A.1 Service-Based Architecture (SBA) in 3GPP 5G

In addition to Internet, the telecommunication network (i.e. carrier mobile wireless networks) is another world-wide networking system. The architecture of the 5G mobile networks from 3GPP has been defined to follow a service-based architecture (SBA) where any network function (NF) can be dynamically associated with any other NF(s) when needed to compose a network service. Note that one NF can simultaneously associate with multiple other NFs, instead of being physically wired as in the previous generations of mobile networks. NFs communicate with each other over service-based interface (SBI), which is also standardized by 3GPP [3GPP.23.501].

In order to realize an SBA network system, detailed requirements are further defined to specify how NFs should interact with each other with information exchange over the SBI. We now list three requirements that are related to information distribution here.

- 1) NF Pub/Sub: Any NF should be able to expose its service status to the network and any NF should be able to subscribe the service status of an NF and get notified if the status is available. A concrete example is that a session management function (SMF) can subscribe to the REGISTER notification from an access management function (AMF) if there is a new user equipment trying to access the mobile network [3GPP.23.502].

- 2) Network Exposure Function (NEF): A particular network function that is required to manage the event exposure and distributions. Specifically, SBA requires such a functionality to register network events from the other NFs (e.g. AMF, SMF and so on), classify the events and properly handle event distributions accordingly in terms of different criteria (e.g. priorities) [3GPP.23.502].

- 3) Network Repository Function (NRF): A particular network function where all service status information is stored for the whole network. An SBA network system requires all NFs to be stateless so as to improve the resilience as well as agility of providing network services. Therefore, the information of the available NFs and the service status generated by those NFs will

be globally stored in NRF as a repository of the system. This clearly implies storage capability that keeps the information in the network and provides those information when needed. A concrete example is that whenever a new NF comes up, it first of all registers itself at NRF with its profile. When a network service requires a certain NF, it first inquires NRF to retrieve the availability information and decides whether or not there is an available NF or a new NF must be instantiated [3GPP.23.502].

(Note: 3GPP CT adopted HTTP2.0/JSON to be the protocol communicating between NFs, but autonomic networks can also load HTTP2.0 with in ACP.)

## A.2 Vehicle-to-Everything

Connected car is one of scenarios interested in automotive manufacturers, carriers and vendors. 5G Automotive Alliance - an industry collaboration organization defines many promising use cases where services from car industry should be supported by the 5G mobile network. Here we list two examples as follows [5GAA.use.cases].

1) Software/Firmware Update: Car manufacturers expect that the software/firmware of their car products can be remotely updated/upgraded via 5G network, instead of onsite visiting their 4S stores/dealers offline as nowadays. This requires the network to provide a mechanism for vehicles to receive the latest software updates during a certain period of time. In order to run such a service for a car manufacturer, the network shall not be just like a network pipe anymore. Instead, information data have to be stored in the network, and delivered in a publishing/subscribing fashion. For example, the latest release of a software will be first distributed and stored at the access edges of the mobile network, after that, the updates can be pushed by the car manufacturer or pulled by the car owner as needed.

2) Real-time HD Maps: Autonomous driving clearly requires much finer details of road maps. Finer details not only include the details of just static road and streets, but also real-time information on the road as well as the driving area for both local urgent situations and intelligent driving scheduling. This asks for situational awareness at critical road segments in cases of changing road conditions. Clearly, a huge amount of traffic data that are real-time collected will have to be stored and shared across the network. This clearly requires the storage capability, data synchronization and event notifications in urgent cases from the network, which are still missing at the infrastructure layer.

## A.3 Summary

Through the general analysis and the concrete examples from the real-world, we realize that the ways information are exchanged in the coming new scenarios are not just short and instant anymore. More advanced as well as diverse information distribution capabilities are required and should be generically supported from the infrastructure layer. Upper layer applications (e.g. ASAs in ANIMA) access and utilize such a unified mechanism for their own services.

#### Appendix B. Information Distribution Module in ANI

This section describes how the information distribution module fits into the ANI and what extensions of GRASP are required [I-D.ietf-anima-grasp].

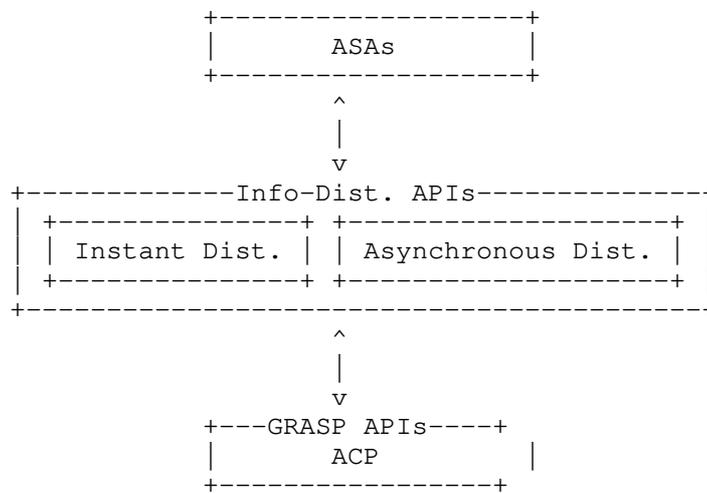


Figure 1. Information Distribution Module and GRASP Extension.

As the Fig 1 shows, the information distribution module two sub-modules for instant and asynchronous information distributions, respectively, and provides APIs to ASAs. Specific Behaviors of modules are described in Section 5.

#### Appendix C. Asynchronous ID Integrated with GRASP APIs

Actions triggered to the information distribution module will eventually invoke underlying GRASP APIs. Moreover, EQ and IS modules are usually correlated. When an AF(ASA) publishes information, not only such an event is translated and sent to EQ module, but also the information is indexed and stored simultaneously. Similarly, when an AF(ASA) subscribes information, not only subscribing event is

triggered and sent to EQ module, but also the information will be retrieved by IS module at the same time.

o Storing and publishing information: This action involves both IS and EQ modules where a node that can store the information will be discovered first and related event will be published to the network. For this, GRASP APIs `discover()`, `synchronize()` and `flood()` are combined to compose such a procedure. In specific, `discover()` call will specific its objective being to "store\_data" and the return parameters could be either an `ASA_locator` who will accept to store the data, or an error code indicating that no one could afford such data; after that, `synchronize()` call will send the data to the specified `ASA_locator` and the data will be stored at that node, with return of processing results like `store_data_ack`; meanwhile, such a successful event (i.e. data is stored successfully) will be flooded via a `flood()` call to interesting parties (such a multicast group existed).

o Subscribing and getting information: This action involves both IS and EQ modules as well where a node that is interested in a topic will subscribe the topic by triggering EQ module and if the topic is ready IS module will retrieve the content of the topic (i.e. the data). GRASP APIs such as `register_objective()`, `flood()`, `synchronize()` are combined to compose the procedure. In specific, any subscription action received by EQ module will be translated to `register_objective()` call where the interested topic will be the parameter inside of the call; the registration will be (selectively) flooded to the network by an API call of `flood()` with the option we extended in this draft; once a matched topic is found (because of the previous procedure), the node finding such a match will call API `synchronize()` to send the stored data to the subscriber.

6tisch Working Group  
Internet-Draft  
Intended status: Informational  
Expires: March 10, 2019

M. Richardson  
Sandelman Software Works  
J. Latour  
CIRA Labs  
F. Khan  
Twelve Dot Systems  
September 06, 2018

BRSKI enrollment for Smart Pledges  
draft-richardson-anima-smartpledge-00

Abstract

This document details the mechanism used for initial enrollment by a smartphone into a BRSKI based enrollment system.

There are two key differences in assumption from [I-D.ietf-anima-bootstrapping-keyinfra]: that the intended registrar has Internet, and that the Pledge has no user-interface.

This variation on BRSKI is intended to be used in the situation where the registrar device is new out of the box and is the intended gateway to the Internet (such as a home gateway), but has not yet been configured. This work is also intended as a transition to the Wi-Fi Alliance work on the Device Provisioning Protocol (DPP).

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 <https://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 March 10, 2019.

## Copyright Notice

Copyright (c) 2018 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 (<https://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 . . . . .	3
1.1.	Additional Motivation . . . . .	4
2.	Terminology . . . . .	4
3.	Requirements Language . . . . .	5
4.	Assumptions and Required Setup . . . . .	5
5.	Protocol Overview . . . . .	6
5.1.	Scan the QR code . . . . .	6
5.2.	Enroll with the manufacturer . . . . .	6
5.3.	Connect to BRSKI join network . . . . .	6
5.4.	Connect to Adolescent Registrar (AR) . . . . .	7
5.5.	Pledge Requests Voucher-Request from the Adolescent Registrar . . . . .	7
5.6.	AR processing of voucher-request, request. . . . .	7
5.7.	Smartpledge validates connection . . . . .	8
5.8.	Smart-Pledge connects to MASA . . . . .	8
5.9.	MASA processing . . . . .	9
5.10.	Smartpledge processing of voucher . . . . .	9
5.11.	Adolescent Registrar (AR) receives voucher . . . . .	9
5.12.	Adolescent Registrar (AR) grows up . . . . .	9
5.13.	Smartpledge enrolls . . . . .	10
5.14.	Validation of connection . . . . .	10
6.	Protocol Details . . . . .	10
6.1.	Quick Response Code (QR code) . . . . .	11
6.1.1.	The SmartPledge Attribute . . . . .	11
6.1.2.	Link-Layer Address Attribute . . . . .	11
6.1.3.	ESSID Name Attribute . . . . .	12
6.2.	Artifacts . . . . .	12
6.2.1.	Voucher-Request Challenge . . . . .	12
6.2.2.	Additions to Voucher-Request . . . . .	12
6.3.	Enrollment using EST . . . . .	12
7.	Smart Pledge enrollment with manufacturer . . . . .	12

8.	Threat Analysis . . . . .	12
8.1.	Wrong Administrator . . . . .	12
8.2.	Rogue Administrator . . . . .	13
8.3.	Attack from Internal device . . . . .	13
8.4.	Attack from camera enabled robot . . . . .	13
8.5.	Attack from manipulator enabled robot . . . . .	13
9.	Security Considerations . . . . .	13
10.	IANA Considerations . . . . .	13
11.	Acknowledgements . . . . .	13
12.	References . . . . .	13
12.1.	Normative References . . . . .	13
12.2.	Informative References . . . . .	14
	Authors' Addresses . . . . .	14

## 1. Introduction

The problem of bootstrapping a new device is described at length in [I-D.ietf-anima-bootstrapping-keyinfra] (aka BRSKI). The problem that BRSKI solves is the case of a smart, properly configured network with a minimum of network connectivity (or previously pre-prewired with nonceless vouchers), and a relatively stupid new device (the Pledge), which lacks a network interface.

The BRSKI problem is one of trust: how does the new device trust that it has found the correct network to join, and how does the new network become convinced that the new device is a device that is intended to join. BRSKI solves the problem well for the case where the network is well connected and easily talk to the device's Manufacturer Authorized Signing Authority (MASA), while providing appropriate proxy mechanisms to enable the new pledge to communicate it's proximity assertion to the MASA as well.

This document is about a variation of the problem: when the new device being introduce has no network connectivity, but a new device is intended to serve as the Registrar for the network. This new device is likely a home (or small office) gateway, and until it is properly configured there will be no direct network connectivity.

There are a number of protocols that permit an ISP to consider a new router brought into a home to be a new pledge to the ISPs' network, and for that new device to integrated into the ISP's (autonomic) network. BRSKI can be used itself, and there are ways to use the Broadband Forum's TR-069 to bootstrap the device in this way. This document is not about the situation where the router device is intended to belong to the ISP, but about the situation where the home user intends to own and control the device.

### 1.1. Additional Motivation

The Wi-Fi Alliance has released the Device Provisioning Protocol [dpp]. The specification is not public. The specification relies on being able to send and receive 802.11 Public Action frames, as well as Generic Advertisement Service (GAS) Public Action frames. Access to send new layer-2 frames is generally restricted in most smartphone operating systems (iOS, Android). At present there are no known public APIs that a generic application writer could use, and therefore the smart-phone side of the DPP can only be implemented at present by the vendors of those operating systems.

As both dominant vendors have competing proprietary mechanisms, it is unclear if generic applications will be produced soon. It is therefore impossible for a vendor on a smart-appliance to independently produce an application that can do proper DPP in 2018.

In addition to the above concern, DPP is primarily concerned about provisioning WiFi credentials to devices. It assumes that the WiFi Access Point is already provisioned and functioning correctly.

The smartpledge enrollment described in this document is about securely initializing the administrative connection with a device that is the WiFi Access Point.

## 2. Terminology

The following terminology is copied from [I-D.ietf-anima-bootstrapping-keyinfra]

**enrollment:** The process where a device presents key material to a network and acquires a network specific identity. For example when a certificate signing request is presented to a certification authority and a certificate is obtained in response.

**pledge:** The prospective device, which has an identity installed at the factory.

**IDevID:** a manufacturer signed keypair (different from the QRkey) which is generated at the factory. This is the 802.1AR artifact which is mandated by [I-D.ietf-anima-bootstrapping-keyinfra].

The following new terminology has been added

**smartpledge:** The prospective administrator device, usually a smartphone equipped with a QR capable camera, wifi and 3G connectivity.

adolescent router (AR): a home router or device containing a registrar. The device does not yet have network connectivity, and has no administrator. It is considered not a "baby" device in the same way that the pledge is, but it is not yet an adult. A better term would be welcome.

SelfDevID: a public/private key pair generated by the smartpledge, formed into a self-signed PKIX certificate. The private key part remains always on the smartpledge, but like other secondary device keys, should be encrypted for backup purposes. {EDNOTE: any references to Apple or Android APIs/specifications here?}

QRkey: a unique, raw ECDSA or EdDSA key pair generated in (or for) the adolescent router at the factory, and stored in the configuration portion of the firmware. The public portion is printed in a QRcode. This key is not formed into a certificate of any kind.

### 3. Requirements Language

In this document, the key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in BCP 14, RFC 2119 [RFC2119] and indicate requirement levels for compliant STuPiD implementations.

### 4. Assumptions and Required Setup

The first assumption is that intended device owner is active and is present. The device owner has a smart-phone that is capable of using Wi-Fi or being wired into the adolescent router (AR).

The smartpledge application generates a self-signed certificate with public/private keypair that it knows. It may generate a unique certificate for each manufacturer. This certificate is called the SelfDevID.

The second assumption is that the device has a QR code printed on the outside of the unit, and/or provided with the packaging/documentation. The QR code is as specified in section 5.3 of [dpp], with the additions specified in Section 6.1

The third assumption is that the AR, at manufacturing time, has the anchor for it's MASA (same assumption as for BRSKI pledge's). In addition, like the BRSKI pledge, the AR has an IDevID certificate (and associated private key) signed by the manufacturer.

The fourth assumption is that the key in the "K:" attribute Section 6.1 is a different public key pair. It MUST be different from the key used in the IDevID. This key is called the DPP-Keypair.

## 5. Protocol Overview

This is the overview of the process. {EDNOTE: there are many details here that belong in the next section. The goal in this section is to consisely explain the interaction among the components. Clearly this text currently fails in that regard}

### 5.1. Scan the QR code

The operator of the smartphone invokes the smartpledge application, and scans the QR code on the AR. The smartpledge learns the ESSID, Public-Key, mac-address, smartpledge URL, and link-local address of the AR.

### 5.2. Enroll with the manufacturer

The smartpledge uses it's 3G, or other WiFi internet access to connect to the manufacturer with TLS. The smartpledge does an HTTP POST to the provided URL using it's generated certificate as it's ClientCertificate. As described in Section 7, the manufacturer MAY respond with a 302 result code, and have the end user go through a web browser based process to enroll. After that process, a redirection will occur using OAUTH2.

The result should finally be a 201 result code, and at that URL is a new certificate signed by the manufacturer.

### 5.3. Connect to BRSKI join network

The application then reconnects the Wi-Fi interface of the smartphone to the ESSID of the AR. This involves normal 802.11 station attachment. The ESSID explicitly has no WPA or other security required on it.

There will be no DHCPv4. A IPv6 Router Solicitation may elicit an answer (confirming the device is there), but it is acceptable for there to be no prefix information. An IPv6 Neighbour Discovery is done for the IPv6 Link-Local address of the AR. Receipt of an answer confirms that the ESSID is correct and present.

(XXX - not using GRASP here. Could use GRASP, but QR code is better)

#### 5.4. Connect to Adolescent Registrar (AR)

The smartpledge application then makes a direct (no proxy) TLS connection to port 443 of the AR, on the IPv6 Link-Local address given. This is as in section 5.1 of [I-D.ietf-anima-bootstrapping-keyinfra]. The smartpledge uses it's SelfDevID as the TLS ClientCertificate, as the smartpledge does not have a manufacturer signed IDevID.

Additionally, the AR will use it's IDevID certificate as the ServerCertificate of the TLS connection. As with other BRSKI IDevID, it will have a MASA URL extension, as described in [I-D.ietf-anima-bootstrapping-keyinfra] section 2.3.2.

#### 5.5. Pledge Requests Voucher-Request from the Adolescent Registrar

The smartpledge generates a random nonce `_SPnonce_`. To this is adds `SOMETHING-that-is-time-unique`, to create a `_voucher-request challenge_`. This is placed in the `voucher-challenge-nonce` field.

Using the public-key of the AR that was scanned from the QR code, the smartpledge encrypts the challenge using CMS (or COSE?).

NOTE: DPP has a round with the SHA256 of the device's key to make sure that the correct device has been chosen. The TLS connection effectively provides the same privacy that the Bx keys provided.

The resulting object is POST'ed to the new BRSKI endpoint:

```
/.well-known/est/requestvoucherrequest
```

```
[or should it be named: /.well-known/est/requestvoucherchallenge
```

```
]
```

#### 5.6. AR processing of voucher-request, request.

The AR processes this POST. First it uses the private key that is associated with it's QR printed public key to decrypt the voucher-request challenge. Included in this challenge is a nonce, and also the link-local address of the smartpledge.

The AR SHOULD verify that the link-local address matches the originating address of the connection on which the request is received.

The AR then forms a voucher-request identically to as described in section 5.2 of [I-D.ietf-anima-bootstrapping-keyinfra]. Note that

the AR uses it's IDevID to sign the voucher-request. This is the same key used to terminate the TLS connection. It MUST be different from the public key printed in the QR code.

In addition to the randomly generated nonce that the AR generates to place in the the voucher-request, into the nonce field, it also includes the `_SPnonce_` in a new `_voucher-challenge-nonce_` field.  
{EDNOTE: hash of nonce?}

This voucher-request is then `_returned_` during the POST operation to the smartpledge. (This is in constrast that in ANIMA the voucher-request is sent by the device to the Registrar, or the MASA)

#### 5.7. Smartpledge validates connection

The smartpledge then examines the resulting voucher-request. The smartpledge validates that the voucher-request is signed by the same public key as was seen in the TLS ServerCertificate.

The smartpledge then examines the contents of the voucher-request, and looks for the `_voucher-challenge-nonce_`. As this nonce was encrypted to the AR, the only way that the resulting nonce could be correct is if the correct private key was present on the AR to decrypt it. Successful verification of the `_voucher-challenge-nonce_` (or the hash of it, see below) results in the smartpledge moving it's end of the connection from provisional to validated.

#### 5.8. Smart-Pledge connects to MASA

The smartpledge application then examines the MASA URL provided in the TLS ServerCertificate of the AR. The smartpledge application then connects to that URL using it's 3G/LTE connection, taking on the role of Registrar.

A wrapped voucher-request is formed by the smartpledge in the same way as described in section 5.4 of [I-D.ietf-anima-bootstrapping-keyinfra]. The inner prior-signed-voucher-request is filled in with the voucher-request that was created by the AR in the previous step.

The pinned-domain-cert of this voucher-request is set to be the SelfDevID certificate of the smartpledge. The voucher-request is to be signed by the SelfDevID.

The voucher-request is POST'ed to the MASA using the same URL that is used for Registrar/MASA operation:

```
/.well-known/est/requestvoucher
```

### 5.9. MASA processing

The MASA processing occurs as specified in section 5.5 of [I-D.ietf-anima-bootstrapping-keyinfra] as before. The MASA MUST also copy the voucher-challenge-nonce into the resulting voucher.

### 5.10. Smartpledge processing of voucher

The smartpledge will receive a voucher that contains it's IDevID as the pinned-domain-cert, and the voucher-challenge-nonce that it created will also be present. The smartpledge SHOULD verify the signature on the artifact, but may be unable to validate that the certificate used has a relationship to the TLS ServerCertificate used by the MASA. (This limitation exists in ANIMA as well).

The smartpledge will then POST the resulting voucher to the AR using the URL

```
/.well-known/est/voucher
```

### 5.11. Adolescent Registrar (AR) receives voucher

When the AR receives the voucher, it validates that it is signed by it's manufacturer. This process is the same as section 5.5.1 of [I-D.ietf-anima-bootstrapping-keyinfra]. Note that this is the future Registrar that is performing what in ANIMA is a pledge operation.

Inside the voucher, the pinned-domain-cert is examined. It should match the TLS ClientCertificate that the smartpledge used to connect. This is the SelfDevID.

At this point the AR has validated the identity of the smartpledge, and the AR moves it's end of the connection from provisional to validated.

### 5.12. Adolescent Registrar (AR) grows up

The roles are now slightly changed. The AR generates a new key pair as it's Domain CA key. It MAY generate intermediate CA certificates and a seperate Registrar certificate, but this is discouraged for home network use.

The AR is now considered a full registrar.

### 5.13. Smartpledge enrolls

The smartpledge MUST now request the full list of CA Certificates, as per [RFC7030] section 4.1. As the Registrar's CA certificate has just been generated, the smartpledge has no other way of knowing it.

The smartpledge MUST now also generate a CSR request as per [I-D.ietf-anima-bootstrapping-keyinfra] section 5.8.3. The smartpledge MAY reuse the SelfDevID key pair for this purpose. (XXX - maybe there are good reasons not to reuse)

The Registrar SHOULD grant administrator privileges to the smartpledge via the certificate that is issued. This may be done via special attributes in the issued certificate, or it may pin the certificate into a database. Which method to use is a local matter.

The EST connection MUST remain open at this point.

### 5.14. Validation of connection

The smartpledge MUST now open a new HTTPS connection to the Registrar (AR), using it's newly issued certificate. (XXX should this be on a different IP, or a different port? If so, how is this indicated?)

The smartpledge MUST validate that the new connection has a certificate that is validated by the Registrar's new CA certificate.

The registrar MUST validate that the smartpledge's ClientCertificate is validated by the Registrar's CA. The smartpledge SHOULD perform a POST operation on this new connection to the [I-D.ietf-anima-bootstrapping-keyinfra] Enrollment Status Telemetry mechanism, see section 5.8.3. The EST connection MAY not be closed.

Should the validations above fail, then the original EST connection MUST be used to GET a value from the

```
/.well-known/est/enrollstatus
```

from the Registrar. The contents of this value SHOULD then be send to the MASA, using a POST to the enrollstatus, and including the reply from the AR in a new attribute, "adolescent-registrar-reason".

## 6. Protocol Details

### 6.1. Quick Response Code (QR code)

Section 5.3 of [dpp] describes the contents for an [iso18004] image. It specifies content that starts with DPP:, and the contains a series of semi-colon (;) delimited section with a single letter and colon. This markup is terminated with a double semi-colon.

Although no amending formula is defined in DPP 1.0, this document is defining two extensions. This requires amending the ABNF from section 5.2.1 as follows:

```
dpp-qr = "DPP:" [channel-list ";"] [channel-list ";"]
        [mac ";"] [information ";"] public-key
        [";" llv6-addr ] [";" smartpledge ] [";" essid ] ";;"
llv6-addr = "L:" 8*hex-octet
essid     = "E:" *(%x20-3A / %x3C-7E) ; semicolon not allowed
smartpledge = "S:" *(%x20-3A / %x3C-7E) ; semicolon not allowed
```

While the ABNF defined in the [dpp] document assumes a specific order (C:, M:, I:, K:) the tags can come in any order. However, in order to make interoperation with future DPP-only clients as seamless as possible, the extensions suggested here are placed at the end of the list. This is consistent with the Postel Principle.

It is intended that parts of this protocol could be performed by an actual DPP implementation, should it become possible to implement DPP using current smartphone operating systems in an unprivileged way.

#### 6.1.1. The SmartPledge Attribute

The `_smartpledge_` attribute indicates that the device is capable of the protocol specified in this document. The contents of the `smartpledge` attribute contains part of a URL which identifies the manufacturer of this device, along with a unique token enabling service access to the device.

Short URLs are essential to fit into typical QR code space.

The `smartpledge` application prepends the text "https://" to the value provided to form the full address of the `smartpledge` enrollment.

#### 6.1.2. Link-Layer Address Attribute

The `_llv6-addr_` attribute is optional. When present, it specifies the IPv6 Link-Local address at which the adolescent router is listening. If not specified, then the link-local address may be formed according to the historical (privacy-violating) process described in [RFC4291] Appendix A. The `_llv6-addr_` attribute is

present so that devices that have implemented [RFC7217] stable addresses can express that address clearly.

#### 6.1.3. ESSID Name Attribute

The `_essid_` attribute provides the name of the 802.11 network to which the `_smartpledge_` SHOULD join in order to reach the AR. If this attribute is absent, then it defaults to "BRSKI".

### 6.2. Artifacts

#### 6.2.1. Voucher-Request Challenge

The `smartpledge` generates a random nonce `_SPnonce_`. To this is added SOMETHING-that-is-time-unique, to create a `_voucher-request-challenge_`. This is placed in the `voucher-challenge-nonce` field.

Using the public-key of the AR that was scanned from the QR code, the `smartpledge` encrypts the challenge using CMS (or COSE?).

#### 6.2.2. Additions to Voucher-Request

QUESTION: should the `_voucher-challenge-nonce_` be provided directly in the `voucher-request`, or should only a hash of the nonce be used? The nonce is otherwise not disclosed, and a MITM on the initial TLS connection would get to see the nonce. A hash of the nonce validates the nonce as easily.

### 6.3. Enrollment using EST

TBD

## 7. Smart Pledge enrollment with manufacturer

TBD.

## 8. Threat Analysis

The following attacks have been considered.

### 8.1. Wrong Administrator

Neighbours with similar setups wind up managing each other's network (by mistake).

## 8.2. Rogue Administrator

Uninitialized networks can be adopted by 'wardrivers' who search for networks that have no administrator.

## 8.3. Attack from Internal device

A compromised device inside the home can be used by an attack to take control of the home router.

## 8.4. Attack from camera enabled robot

A robot (such as a home vacuum cleaner) could be compromised, and then used by an attacker to observe and/or scan the router QRcode.

## 8.5. Attack from manipulator enabled robot

A robot (for instance, a toy) could be compromised, and then used by an attacker to push the WPA and/or factory reset button on the router.

## 9. Security Considerations

Go through the list of attacks above, and explain how each has been mitigated.

Go through the list of concerns in ANIMA and EST-RFC7030 and indicate if there are additional concerns, or if a concern does not apply.

## 10. IANA Considerations

TBD.

## 11. Acknowledgements

This work was supported by the Canadian Internet Registration Authority <https://cira.ca/blogs/cira-labs/about-cira-labs>.

## 12. References

### 12.1. Normative References

- [dpp] "Device Provisioning Protocol Specification", n.d., <[https://www.wi-fi.org/downloads-registered-guest/Device\\_Provisioning\\_Protocol\\_Draft\\_Technical\\_Specification\\_Package\\_v0\\_0\\_23\\_0.zip/31255](https://www.wi-fi.org/downloads-registered-guest/Device_Provisioning_Protocol_Draft_Technical_Specification_Package_v0_0_23_0.zip/31255)>.

- [I-D.ietf-anima-bootstrapping-keyinfra]  
Pritikin, M., Richardson, M., Behringer, M., Bjarnason, S., and K. Watsen, "Bootstrapping Remote Secure Key Infrastructures (BRSKI)", draft-ietf-anima-bootstrapping-keyinfra-16 (work in progress), June 2018.
- [iso18004]  
"Information technology --- Automatic identification and data capture techniques --- Bar code symbology --- QR Codes (ISO/IEC 18004:2015)", n.d.,  
<<https://github.com/yansikeim/QR-Code/blob/master/ISO%20IEC%2018004%202015%20Standard.pdf>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC7030] Pritikin, M., Ed., Yee, P., Ed., and D. Harkins, Ed., "Enrollment over Secure Transport", RFC 7030, DOI 10.17487/RFC7030, October 2013, <<https://www.rfc-editor.org/info/rfc7030>>.

## 12.2. Informative References

- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <<https://www.rfc-editor.org/info/rfc4291>>.
- [RFC7217] Gont, F., "A Method for Generating Semantically Opaque Interface Identifiers with IPv6 Stateless Address Autoconfiguration (SLAAC)", RFC 7217, DOI 10.17487/RFC7217, April 2014, <<https://www.rfc-editor.org/info/rfc7217>>.

## Authors' Addresses

Michael Richardson  
Sandelman Software Works

Email: [mcr+ietf@sandelman.ca](mailto:mcr+ietf@sandelman.ca)

Jacques Latour  
CIRA Labs

Email: [Jacques.Latour@cira.ca](mailto:Jacques.Latour@cira.ca)

Faud Khan  
Twelve Dot Systems

Email: [faud.khan@twelvedot.com](mailto:faud.khan@twelvedot.com)

anima Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: 26 March 2021

M. Richardson  
Sandelman Software Works  
P. van der Stok  
vanderstok consultancy  
P. Kampanakis  
Cisco Systems  
22 September 2020

Constrained Join Proxy for Bootstrapping Protocols  
draft-vanderstok-anima-constrained-join-proxy-04

Abstract

This document defines a protocol to securely assign a pledge to an owner, using an intermediary node between pledge and owner. This intermediary node is known as a "constrained Join Proxy".

This document extends the work of [I-D.ietf-anima-bootstrapping-keyinfra] by replacing the Circuit-proxy by a stateless constrained (CoAP) Join Proxy. It transports join traffic from the pledge to the Registrar without requiring per-client state.

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 <https://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 26 March 2021.

Copyright Notice

Copyright (c) 2020 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 (<https://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 . . . . .	3
2. Terminology . . . . .	4
3. Requirements Language . . . . .	4
4. Join Proxy functionality . . . . .	4
5. Join Proxy specification . . . . .	5
5.1. Statefull Join Proxy . . . . .	5
5.2. Stateless Join Proxy . . . . .	6
5.3. Stateless Message structure . . . . .	8
6. Comparison of stateless and statefull modes . . . . .	9
7. Discovery . . . . .	10
7.1. Pledge discovery of Join Proxy . . . . .	10
7.1.1. CoAP discovery . . . . .	10
7.1.2. Autonomous Network . . . . .	11
7.1.3. 6tisch discovery . . . . .	11
7.2. Join Proxy discovers EST server . . . . .	11
7.2.1. Autonomous Network . . . . .	11
7.2.2. CoAP discovery . . . . .	11
8. Security Considerations . . . . .	12
9. IANA Considerations . . . . .	12
9.1. Resource Type registry . . . . .	12
10. Acknowledgements . . . . .	12
11. Contributors . . . . .	13
12. Changelog . . . . .	13
12.1. 01 to 02 . . . . .	13
12.2. 00 to 01 . . . . .	13
12.3. 00 to 00 . . . . .	13
13. References . . . . .	13
13.1. Normative References . . . . .	13
13.2. Informative References . . . . .	15
Appendix A. Stateless Proxy payload examples . . . . .	16
A.1. cacerts . . . . .	18
A.2. serverkeygen . . . . .	19
Authors' Addresses . . . . .	20

## 1. Introduction

Enrolment of new nodes into constrained networks with constrained nodes present is described in [I-D.ietf-anima-bootstrapping-keyinfra] ("BRSKI") and makes use of Enrolment over Secure Transport (EST) [RFC7030] with [RFC8366] vouchers to securely enroll devices. BRSKI connects new devices ("pledges") to extended EST servers ("Registrars") via a Join Proxy.

The specified solutions use https and may be too large in terms of code space or bandwidth required. Constrained devices in constrained networks [RFC7228] typically implement the IPv6 over Low-Power Wireless personal Area Networks (6LoWPAN) [RFC4944] and Constrained Application Protocol (CoAP) [RFC7252].

CoAP can be run with the Datagram Transport Layer Security (DTLS) [RFC6347] as a security protocol for authenticity and confidentiality of the messages. This is described as the "coaps" scheme. A constrained version of EST, using Coap and DTLS, is described in [I-D.ietf-ace-coap-est].

DTLS is a client-server protocol relying on the underlying IP layer to perform the routing between the DTLS Client and the DTLS Server. However, the new "joining" device will not be IP routable until it is authenticated to the network. A new "joining" device can only initially use a link-local IPv6 address to communicate with a neighbour node using neighbour discovery [RFC6775] until it receives the necessary network configuration parameters. However, before the device can receive these configuration parameters, it needs to authenticate itself to the network to which it connects. IPv6 routing is necessary to establish a connection between joining device and the extended EST server.

This document specifies a new form of Join Proxy and protocol to act as intermediary between joining device and EST server to establish a connection between joining device and EST server.

This document is very much inspired by text published earlier in [I-D.kumar-dice-dtls-relay].

[I-D.richardson-anima-state-for-joinrouter] outlined the various options for building a join proxy.

[I-D.ietf-anima-bootstrapping-keyinfra] adopted only the Circuit Proxy method (1), leaving the other methods as future work. The document standardizes the CoAP/DTLS (method 4).

2. Terminology

The following terms are defined in [RFC8366], and are used identically as in that document: artifact, imprint, domain, Join Registrar/Coordinator (JRC), Manufacturer Authorized Signing Authority (MASA), pledge, Trust of First Use (TOFU), and Voucher.

3. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

4. Join Proxy functionality

As depicted in the Figure 1, the joining Device, or pledge (P), in an LLN mesh is more than one hop away from the EST server (E) and not yet authenticated into the network.

At this stage, it can only communicate one-hop to its nearest neighbour, the Join Proxy (J) using their link-local IPv6 addresses. However, the Pledge (P) needs to communicate with end-to-end security with a Registrar hosting the EST server (E) to authenticate and get the relevant system/network parameters. If the Pledge (P) initiates a DTLS connection to the EST server whose IP address has been pre-configured, then the packets are dropped at the Join Proxy (J) since the Pledge (P) is not yet admitted to the network or there is no IP routability to Pledge (P) for any returned messages.

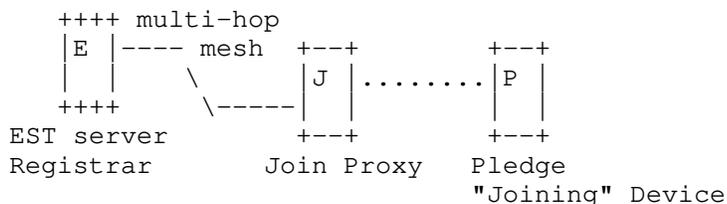


Figure 1: multi-hop enrolment.

Furthermore, the Pledge (P) may wish to establish a secure connection to the EST server (E) in the network assuming appropriate credentials are exchanged out-of-band, e.g. a hash of the Pledge (P)'s raw public key could be provided to the EST server (E). However, the Pledge (P) may be unaware of the IP address of the EST-server (E) to initiate a DTLS connection and perform authentication with.

A DTLS connection is required between Pledge and EST server. To overcome the problems with non-routability of DTLS packets and/or discovery of the destination address of the EST Server to contact, the Join Proxy is introduced. This Join Proxy functionality is configured into all authenticated devices in the network which may act as the Join Proxy for newly joining nodes. The Join Proxy allows for routing of the packets from the Pledge using IP routing to the intended EST Server.

## 5. Join Proxy specification

A Join Proxy can operate in two modes:

- \* Statefull mode
- \* Stateless mode

### 5.1. Statefull Join Proxy

In stateful mode, the joining node forwards the DTLS messages to the BRSKI Registrar.

Assume that the Pledge does not know the IP address of the EST Server it needs to contact. In that situation, the Join Proxy must know the (configured or discovered) IP address of a BRSKI Registrar.

(Discovery can be based upon [I-D.ietf-anima-bootstrapping-keyinfra] section 4.3, or via DNS-SD service discovery [RFC6763]) The Pledge initiates its request as if the Join Proxy is the intended Registrar. The Join Proxy changes the IP packet (without modifying the DTLS message) as in the previous case by modifying both the source and destination addresses to forward the message to the intended EST Server. The Join Proxy maintains a 4-tuple array to translate the DTLS messages received from the EST Server and forward it to the EST Client. This is a form of Network Address translation, where the Join Proxy acts as a forward proxy. In Figure 2 the various steps of the message flow are shown:

EST Client (P)	Join Proxy (J)	EST Server (E)	Message	
			Src_IP:port	Dst_IP:port
	--ClientHello-->		IP_P:p_P	IP_Ja:5684
	--ClientHello-->		IP_Jb:p_Jb	IP_E:5684
		<--ServerHello--	IP_E:5684	IP_Jb:p_Jb
		:		
<--ServerHello--		:	IP_Ja:5684	IP_P:p_P
:		:		
:		:		
:		:		
--Finished-->		:	IP_P:p_P	IP_Ja:5684
	--Finished-->		IP_Jb:p_Jb	IP_E:5684
		<--Finished--	IP_E:5684	IP_Jb:p_Jb
<--Finished--		:	IP_Ja:5684	IP_P:p_P
:		:		

IP\_P:p\_P = Link-local IP address and port of Pledge (DTLS Client)  
 IP\_E:5684 = Global IP address and coaps port of EST Server  
 IP\_Ja:5684 = Link-local IP address and coaps port of Join Proxy  
 IP\_Jb:p\_Rb = Global IP address and port of Join proxy

Figure 2: constrained statefull joining message flow with EST server address known to Join Proxy.

### 5.2. Stateless Join Proxy

The Join Proxy is stateless to minimize the requirements on the constrained Join Proxy device. Stateless operation requires no memory in the Join Proxy device, but may also reduce the CPU impact as the device does not need to search through a state table.

When a joining device as a client attempts a DTLS connection to the EST server, it uses its link-local IP address as its IP source address. This message is transmitted one-hop to a neighbouring (join proxy) node. Under normal circumstances, this message would be dropped at the neighbour node since the joining device is not yet IP routable or it is not yet authenticated to send messages through the network. However, if the neighbour device has the Join Proxy functionality enabled, it routes the DTLS message to a specific Registrar. Additional security mechanisms need to exist to prevent this routing functionality being used by rogue nodes to bypass any network authentication procedures.

If an untrusted DTLS Client that can only use link-local addressing wants to contact a trusted end-point Registrar, it sends the DTLS message to the Join Proxy.

The Join Proxy extends this message into a new type of message called Join Proxy (JPY) message and sends it on to the Registrar.

The JPY message payload consists of two parts:

- \* Header (H) field: consisting of the source link-local address and port of the Pledge (P), and
- \* Contents (C) field: containing the original DTLS message.

On receiving the JPY message, the BRSKI Registrar retrieves the two parts.

The BRSKI server transiently stores the Header field information. The Registrar server uses the Contents field to execute the Registrar server functionality. However, when the Registrar replies, it also extends its DTLS message with the header field in a JPY message and sends it back to the Join Proxy. The Registrar SHOULD NOT assume that it can decode the Header Field, it should simply repeat it when responding. The Header contains the original source link-local address and port of the DTLS Client from the transient state stored earlier (which can now be discarded) and the Contents field contains the DTLS message.

On receiving the JPY message, the Join Proxy retrieves the two parts. It uses the Header field to route the DTLS message retrieved from the Contents field to the Pledge.

The Figure 3 depicts the message flow diagram:

EST Client (P)	Join Proxy (J)	EST server (E)	Message	
			Src_IP:port	Dst_IP:port
	--ClientHello-->		IP_P:p_P	IP_Ja:5684
	--JPY[H(IP_P:p_P),--> C(ClientHello)]		IP_Jb:p_Jb	IP_E:5684
	<--JPY[H(IP_P:p_P),-- C(ServerHello)]		IP_E:5684	IP_Jb:p_Jb
	<--ServerHello--		IP_Ja:5684	IP_P:p_P
	:		:	:
	:		:	:
	--Finished-->		IP_P:p_P	IP_Ja:5684
	--JPY[H(IP_P:p_P),--> C(Finished)]		IP_Jb:p_Jb	IP_E:5684
	<--JPY[H(IP_P:p_P),-- C(Finished)]		IP_E:5684	IP_Jb:p_Jb
	<--Finished--		IP_Ja:5684	IP_P:p_P
	:		:	:

IP\_P:p\_P = Link-local IP address and port of the Pledge  
 IP\_E:5684 = Global IP address and coaps port of EST Server  
 IP\_Ja:5684 = Link-local IP address and coaps port of Join Proxy  
 IP\_Jb:p\_Jb = Global IP address and port of Join Proxy

JPY[H(),C()] = Join Proxy message with header H and content C

Figure 3: constrained stateless joining message flow.

### 5.3. Stateless Message structure

The JPY message is constructed as a payload with media-type application/multipart-core specified in [I-D.ietf-core-multipart-ct].

Header and Contents fields use different media formats:

1. header field: application/cbor containing a CBOR array [RFC7049] with the pledge IPv6 Link Local address as a 16-byte binary value, the pledge's UDP port number, if different from 5684, as a CBOR integer, and the proxy's ifindex or other identifier for the physical port on which the pledge is connected. Header is not DTLS encrypted.
2. Content field: Any of the media types specified in [I-D.ietf-ace-coap-est] and [I-D.ietf-anima-constrained-voucher] dependent on the function that is requested:

- \* application/pkcs7-mime; smime-type=server-generated-key
- \* application/pkcs7-mime; smime-type=certs-only
- \* application/voucher-cms+cbor
- \* application/voucher-cose+cbor
- \* application/pkcs8
- \* application/csrattrs
- \* application/pkcs10
- \* application/pkix-cert

(XXX- add CDDL for CBOR array above)

The content fields are DTLS encrypted. In CBOR diagnostic notation the payload JPY[H(IP\_P:p\_P), with cf is content-format of DTLS-content, will look like:

```
[ 60: [IP_p, p_P, ident]
  cf: h'DTLS-content']
```

Examples are shown in Appendix A.

## 6. Comparison of stateless and statefull modes

The stateful and stateless mode of operation for the Join Proxy have their advantages and disadvantages. This section should enable to make a choice between the two modes based on the available device resources and network bandwidth.

Properties	Stateful mode	Stateless mode
State Information	The Join Proxy needs additional storage to maintain mapping between the address and port number of the pledge and those of the EST-server.	No information is maintained by the Join Proxy
Packet size	The size of the forwarded message is the same as the original message.	Size of the forwarded message is bigger than the original, it includes additional source and destination addresses.
Specification complexity	The Join Proxy needs additional functionality to maintain state information, and modify the source and destination addresses of the DTLS handshake messages	New JPY message to encapsulate DTLS message The EST server and the Join Proxy have to understand the JPY message in order to process it.

Figure 4: Comparison between stateful and stateless mode

## 7. Discovery

It is assumed that Join Proxy seamlessly provides a coaps connection between Pledge and coaps EST-server. An additional Registrar is needed to connect the Pledge to an http EST server, see section 8 of [I-D.ietf-ace-coap-est]. In particular this section replaces section 4.2 of [I-D.ietf-anima-bootstrapping-keyinfra].

Three discovery cases are discussed: coap discovery, 6tisch discovery and GRASP discovery.

### 7.1. Pledge discovery of Join Proxy

The pledge and Join Proxy are assumed to communicate via Link-Local addresses.

#### 7.1.1. CoAP discovery

The discovery of the coaps EST server, using coap discovery, by the Join Proxy follows section 6 of [I-D.ietf-ace-coap-est].

### 7.1.2. Autonomous Network

In the context of autonomous networks, the Join Proxy uses the DULL GRASP M\_FLOOD mechanism to announce itself. Section 4.1.1 of [I-D.ietf-anima-bootstrapping-keyinfra] discusses this in more detail. The Registrar announces itself using ACP instance of GRASP using M\_FLOOD messages. Autonomous Network Join Proxies MUST support GRASP discovery of EST-server as described in section 4.3 of [I-D.ietf-anima-bootstrapping-keyinfra] .

### 7.1.3. 6tisch discovery

The discovery of EST server by the pledge uses the enhanced beacons as discussed in [I-D.ietf-6tisch-enrollment-enhanced-beacon].

## 7.2. Join Proxy discovers EST server

### 7.2.1. Autonomous Network

The pledge MUST listen for GRASP M\_FLOOD [I-D.ietf-anima-grasp] announcements of the objective: "AN\_Proxy". See section Section 4.1.1 [I-D.ietf-anima-bootstrapping-keyinfra] for the details of the objective.

### 7.2.2. CoAP discovery

In the context of a coap network without Autonomous Network support, discovery follows the standard coap policy. The Pledge can discover a Join Proxy by sending a link-local multicast message to ALL CoAP Nodes with address FF02::FD. Multiple or no nodes may respond. The handling of multiple responses and the absence of responses follow section 4 of [I-D.ietf-anima-bootstrapping-keyinfra].

The presence and location of (path to) the Join Proxy resource are discovered by sending a GET request to `"/.well-known/core"` including a resource type (rt) parameter with the value `"brski-proxy"` [RFC6690]. Upon success, the return payload will contain the root resource of the Join Proxy resources. It is up to the implementation to choose its root resource; throughout this document the example root resource `/jp` is used. The example below shows the discovery of the presence and location of Join Proxy resources.

```
REQ: GET coap://[FF02::FD]/.well-known/core?rt=brski-proxy
```

```
RES: 2.05 Content  
</jp>; rt="brski-proxy";ct=62
```

Port numbers, not returned in the example, are assumed to be the default numbers 5683 and 5684 for coap and coaps respectively (sections 12.6 and 12.7 of [RFC7252]). Discoverable port numbers MAY be returned in the <href> of the payload (see section 5.1 of [I-D.ietf-ace-coap-est]).

## 8. Security Considerations

It should be noted here that the contents of the CBOR map used to convey return address information is not protected. However, the communication is between the Proxy and a known registrar are over the already secured portion of the network, so are not visible to eavesdropping systems.

All of the concerns in [I-D.ietf-anima-bootstrapping-keyinfra] section 4.1 apply. The pledge can be deceived by malicious AN\_Proxy announcements. The pledge will only join a network to which it receives a valid [RFC8366] voucher.

If the proxy/Registrar was not over a secure network, then an attacker could change the cbor array, causing the pledge to send traffic to another node. If the such scenario needed to be supported, then it would be reasonable for the Proxy to encrypt the CBOR array using a locally generated symmetric key. The Registrar would not be able to examine the result, but it does not need to do so. This is a topic for future work.

## 9. IANA Considerations

This document needs to create a registry for key indices in the CBOR map. It should be given a name, and the amending formula should be IETF Specification.

### 9.1. Resource Type registry

This specification registers a new Resource Type (rt=) Link Target Attributes in the "Resource Type (rt=) Link Target Attribute Values" subregistry under the "Constrained RESTful Environments (CoRE) Parameters" registry.

rt="brski-proxy". This EST resource is used to query and return the supported EST resource of a Join Proxy placed between Pledge and EST server.

## 10. Acknowledgements

Many thanks for the comments by Brian Carpenter.

## 11. Contributors

Sandeep Kumar, Sye loong Keoh, and Oscar Garcia-Morchon are the co-authors of the draft-kumar-dice-dtls-relay-02. Their draft has served as a basis for this document. Much text from their draft is copied over to this draft.

## 12. Changelog

### 12.1. 01 to 02

- \* extended the discovery section
- \* removed inconsistencies from the the flow diagrams
- \* Improved readability of the examples.
- \* stateful configurations reduced to one

### 12.2. 00 to 01

- \* Added Contributors section
- \* Adapted content-formats to est-coaps formats
- \* Aligned examples with est-coaps examples
- \* Added statefull Proxy to stateless proxy

### 12.3. 00 to 00

- \* added payload examples in appendix
- \* discovery for three cases: AN, 6tisch and coaps

## 13. References

### 13.1. Normative References

[I-D.ietf-6tisch-enrollment-enhanced-beacon]  
Dujovne, D. and M. Richardson, "IEEE 802.15.4 Information Element encapsulation of 6TiSCH Join and Enrollment Information", Work in Progress, Internet-Draft, draft-ietf-6tisch-enrollment-enhanced-beacon-14, 21 February 2020, <<http://www.ietf.org/internet-drafts/draft-ietf-6tisch-enrollment-enhanced-beacon-14.txt>>.

- [I-D.ietf-ace-coap-est]  
Stok, P., Kampanakis, P., Richardson, M., and S. Raza,  
"EST over secure CoAP (EST-coaps)", Work in Progress,  
Internet-Draft, draft-ietf-ace-coap-est-18, 6 January  
2020, <<http://www.ietf.org/internet-drafts/draft-ietf-ace-coap-est-18.txt>>.
- [I-D.ietf-anima-bootstrapping-keyinfra]  
Pritikin, M., Richardson, M., Eckert, T., Behringer, M.,  
and K. Watsen, "Bootstrapping Remote Secure Key  
Infrastructures (BRSKI)", Work in Progress, Internet-  
Draft, draft-ietf-anima-bootstrapping-keyinfra-44, 21  
September 2020, <<http://www.ietf.org/internet-drafts/draft-ietf-anima-bootstrapping-keyinfra-44.txt>>.
- [I-D.ietf-anima-constrained-voucher]  
Richardson, M., Stok, P., and P. Kampanakis, "Constrained  
Voucher Artifacts for Bootstrapping Protocols", Work in  
Progress, Internet-Draft, draft-ietf-anima-constrained-  
voucher-08, 13 July 2020, <<http://www.ietf.org/internet-drafts/draft-ietf-anima-constrained-voucher-08.txt>>.
- [I-D.ietf-anima-grasp]  
Bormann, C., Carpenter, B., and B. Liu, "A Generic  
Autonomic Signaling Protocol (GRASP)", Work in Progress,  
Internet-Draft, draft-ietf-anima-grasp-15, 13 July 2017,  
<<http://www.ietf.org/internet-drafts/draft-ietf-anima-grasp-15.txt>>.
- [I-D.ietf-core-multipart-ct]  
Fossati, T., Hartke, K., and C. Bormann, "Multipart  
Content-Format for CoAP", Work in Progress, Internet-  
Draft, draft-ietf-core-multipart-ct-04, 21 August 2019,  
<<http://www.ietf.org/internet-drafts/draft-ietf-core-multipart-ct-04.txt>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate  
Requirement Levels", BCP 14, RFC 2119,  
DOI 10.17487/RFC2119, March 1997,  
<<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer  
Security Version 1.2", RFC 6347, DOI 10.17487/RFC6347,  
January 2012, <<https://www.rfc-editor.org/info/rfc6347>>.
- [RFC7049] Bormann, C. and P. Hoffman, "Concise Binary Object  
Representation (CBOR)", RFC 7049, DOI 10.17487/RFC7049,  
October 2013, <<https://www.rfc-editor.org/info/rfc7049>>.

- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8366] Watsen, K., Richardson, M., Pritikin, M., and T. Eckert, "A Voucher Artifact for Bootstrapping Protocols", RFC 8366, DOI 10.17487/RFC8366, May 2018, <<https://www.rfc-editor.org/info/rfc8366>>.

### 13.2. Informative References

- [duckling] Stajano, F. and R. Anderson, "The resurrecting duckling: security issues for ad-hoc wireless networks", 1999, <<https://www.cl.cam.ac.uk/~fms27/papers/1999-StajanoAnd-duckling.pdf>>.
- [I-D.kumar-dice-dtls-relay] Kumar, S., Keoh, S., and O. Garcia-Morchon, "DTLS Relay for Constrained Environments", Work in Progress, Internet-Draft, draft-kumar-dice-dtls-relay-02, 20 October 2014, <<http://www.ietf.org/internet-drafts/draft-kumar-dice-dtls-relay-02.txt>>.
- [I-D.richardson-anima-state-for-joinrouter] Richardson, M., "Considerations for stateful vs stateless join router in ANIMA bootstrap", Work in Progress, Internet-Draft, draft-richardson-anima-state-for-joinrouter-02, 25 January 2018, <<http://www.ietf.org/internet-drafts/draft-richardson-anima-state-for-joinrouter-02.txt>>.
- [pledge] Dictionary.com, ., "Dictionary.com Unabridged", 2015, <<http://dictionary.reference.com/browse/pledge>>.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007, <<https://www.rfc-editor.org/info/rfc4944>>.
- [RFC6690] Shelby, Z., "Constrained RESTful Environments (CoRE) Link Format", RFC 6690, DOI 10.17487/RFC6690, August 2012, <<https://www.rfc-editor.org/info/rfc6690>>.
- [RFC6763] Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", RFC 6763, DOI 10.17487/RFC6763, February 2013, <<https://www.rfc-editor.org/info/rfc6763>>.

- [RFC6775] Shelby, Z., Ed., Chakrabarti, S., Nordmark, E., and C. Bormann, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6775, DOI 10.17487/RFC6775, November 2012, <<https://www.rfc-editor.org/info/rfc6775>>.
- [RFC7030] Pritikin, M., Ed., Yee, P., Ed., and D. Harkins, Ed., "Enrollment over Secure Transport", RFC 7030, DOI 10.17487/RFC7030, October 2013, <<https://www.rfc-editor.org/info/rfc7030>>.
- [RFC7228] Bormann, C., Ersue, M., and A. Keranen, "Terminology for Constrained-Node Networks", RFC 7228, DOI 10.17487/RFC7228, May 2014, <<https://www.rfc-editor.org/info/rfc7228>>.
- [RFC7252] Shelby, Z., Hartke, K., and C. Bormann, "The Constrained Application Protocol (CoAP)", RFC 7252, DOI 10.17487/RFC7252, June 2014, <<https://www.rfc-editor.org/info/rfc7252>>.

#### Appendix A. Stateless Proxy payload examples

Examples are extensions of two examples shown in [I-D.ietf-ace-coap-est]. The following content formats are used:

- \* 60: application/cbor
- \* 62: application/multipart
- \* 281: application/pkcs7-mime; smime-type=certs-only
- \* 284: application/pkcs8
- \* 286: application/pkcs10

For presentation purposes the payloads are abbreviated as follows:

cacrts request payload:

<cacrts request payload> = <empty>

cacrts response payload:

```
<cacrts response payload> =
DTLS_encrypt (
3082027b06092a864886f70d010702a082026c308202680201013100300b
06092a864886f70d010701a082024e3082024a308201f0a0030201020209
009189bcd9c99244b300a06082a8648ce3d0403023067310b3009060355
040613025553310b300906035504080c024341310b300906035504070c02
4c4131143012060355040a0c0b4578616d706c6520496e63311630140603
55040b0c0d63657274696669636174696f6e3110300e06035504030c0752
6f6f74204341301e170d3139303130373130343034315a170d3339303130
323130343034315a3067310b3009060355040613025553310b3009060355
04080c024341310b300906035504070c024c4131143012060355040a0c0b
4578616d706c6520496e6331163014060355040b0c0d6365727469666963
6174696f6e3110300e06035504030c07526f6f742043413059301306072a
8648ce3d020106082a8648ce3d03010703420004814994082b6e8185f3df
53f5e0bee698973335200023ddf78cd17a443ffd8ddd40908769c55652ac
2ccb75c4a50a7c7ddb7c22dae6c85cca538209fdbbf104c9a38184308181
301d0603551d0e041604142495e816ef6ffcaaf356ce4adffe33cf492abb
a8301f0603551d230418301680142495e816ef6ffcaaf356ce4adffe33cf
492abba8300f0603551d130101ff040530030101ff300e0603551d0f0101
ff040403020106301e0603551d1104173015811363657274696679406578
616d706c652e636f6d300a06082a8648ce3d0403020348003045022100da
e37c96f154c32ec0b4af52d46f3b7ecc9687ddf267bceec368f7b7f135327
2f022047a28ae5c7306163b3c3834bab3c103f743070594c089aaa0ac870
cd13b902caa1003100
)
```

serverkeygen request payload:

```
<serverkeygen request payload> =
DTLS_encrypt (
3081cf3078020100301631143012060355040a0c0b736b67206578616d70
6c653059301306072a8648ce3d020106082a8648ce3d030107034200041b
b8c1117896f98e4506c03d70efbe820d8e38ea97e9d65d52c8460c5852c5
1dd89a61370a2843760fc859799d78cd33f3c1846e304f1717f8123f1a28
4cc99fa000300a06082a8648ce3d04030203470030440220387cd4e9cf62
8d4af77f92ebed4890d9d141dca86cd2757dd14cbd59cdf6961802202f24
5e828c77754378b66660a4977f113cacdaa0cc7bad7d1474a7fd155d090d
)
```

serverkeygen response payload:

```

<serverkeygen response payload> =
DTLS_encrypt (
  84                                     # array(4)
  19 011C                               # unsigned(284)
  58 8A                                 # bytes(138)
  308187020100301306072a8648ce3d020106082a8648ce3d030107046d30
  6b02010104200b9a67785b65e07360b6d28cfc1d3f3925c0755799deeca7
  45372b01697bd8a6a144034200041bb8c1117896f98e4506c03d70efbe82
  0d8e38ea97e9d65d52c8460c5852c51dd89a61370a2843760fc859799d78
  cd33f3c1846e304f1717f8123f1a284cc99f
  19 0119                               # unsigned(281)
  59 01D3                               # bytes(467)
  308201cf06092a864886f70d010702a08201c0308201bc0201013100300b
  06092a864886f70d010701a08201a23082019e30820143a0030201020208
  126de8571518524b300a06082a8648ce3d04030230163114301206035504
  0a0c0b736b67206578616d706c65301e170d313930313039303835373038
  5a170d3339303130343038353730385a301631143012060355040a0c0b73
  6b67206578616d706c653059301306072a8648ce3d020106082a8648ce3d
  030107034200041bb8c1117896f98e4506c03d70efbe820d8e38ea97e9d6
  5d52c8460c5852c51dd89a61370a2843760fc859799d78cd33f3c1846e30
  4f1717f8123f1a284cc99fa37b307930090603551d1304023000302c0609
  6086480186f842010d041f161d4f70656e53534c2047656e657261746564
  204365727469666963617465301d0603551d0e04160414494be598dc8dbc
  0dbc071c486b777460e5cce621301f0603551d23041830168014494be598
  dc8dbc0dbc071c486b777460e5cce621300a06082a8648ce3d0403020349
  003046022100a4b167d0f9add9202810e6bf6a290b8cfd9c9b9c9fea2cc1
  c8fc3a464f79f2c202210081d31ba142751a7b4a34fd1a01fcfb08716b9e
  b53bdaadc9ae60b08f52429c0fa1003100
)

```

#### A.1. cacerts

The request from Join Proxy to EST-server looks like:

```

Get coaps://192.0.2.1/est/crts
(Accept: 62)
(Content-format: 62)
payload =
  82                                     # array(2)
  18 3C                                 # unsigned(60)
  83                                     # array(3)
  69                                     # text(9)
      464538303A3A414238 # "FE80::AB8"
  19 237D                               # unsigned(9085)
  65                                     # text(5)
      6964656E74         # "ident"

```

In CBOR Diagnostic:

```
payload = [60, ["FE80::AB8", 9085, "ident"]]
```

The response will then be:

```
2.05 Content
(Content-format: 62)
  Payload =
84                                     # array(4)
18 3C                                 # unsigned(60)
83                                     # array(3)
69                                     # text(9)
    464538303A3A414238                # "FE80::AB8"
19 237D                               # unsigned(9085)
65                                     # text(5)
    6964656E74                        # "ident"
19 0119                               # unsigned(281)
59 027F                               # bytes(639)
<cacrts response payload>
]
```

In CBOR diagnostic:

```
payload = [60, ["FE80::AB8", 9085, "ident"],
           62, h'<cacrts response payload>']
```

#### A.2. serverkeygen

The request from Join Proxy to EST-server looks like:

```
Get coaps://192.0.2.1/est/skg
(Accept: 62)
(Content-Format: 62)
  Payload =
83                                     # array(4)
18 3C                                 # unsigned(60)
83                                     # array(3)
69                                     # text(9)
    464538303A3A414238                # "FE80::AB8"
19 237D                               # unsigned(9085)
65                                     # text(5)
    6964656E74                        # "ident"
19 011E                               # unsigned(286)
58 D2                                 # bytes(210)
<serverkeygen request payload>
```

In CBOR diagnostic:

```
payload = [60, ["FE80::AB8", 9085, "ident"],
           286, h'<serverkeygen request payload>']
```

The response will then be:

```
2.05 Content
(Content-format: 62)
  Payload =
83                                     # array(4)
18 3C                                 # unsigned(60)
83                                     # array(3)
69                                     # text(9)
    464538303A3A414238                # "FE80::AB8"
19 237D                               # unsigned(9085)
65                                     # text(5)
    6964656E74                        # "ident"
19 011E                               # unsigned(286)
59 0269                               # bytes(617)
<serverkeygen response payload>
```

In CBOR diagnostic:

```
payload = [60, ["FE80::AB8", 9085, "ident"],
           286, h'<serverkeygen response payload>']
```

#### Authors' Addresses

Michael Richardson  
Sandelman Software Works

Email: [mcr+ietf@sandelman.ca](mailto:mcr+ietf@sandelman.ca)

Peter van der Stok  
vanderstok consultancy

Email: [consultancy@vanderstok.org](mailto:consultancy@vanderstok.org)

Panos Kampanakis  
Cisco Systems

Email: [pkampana@cisco.com](mailto:pkampana@cisco.com)