

ANIMA
Internet-Draft
Intended status: Informational
Expires: January 1, 2016

M. Behringer, Ed.
Cisco Systems
B. Carpenter
Univ. of Auckland
T. Eckert
Cisco
L. Ciavaglia
Alcatel Lucent
B. Liu
Huawei Technologies
J. Nobre
Federal University of Rio Grande do Sul
J. Strassner
Huawei Technologies
June 30, 2015

A Reference Model for Autonomic Networking
draft-behringer-anima-reference-model-03

Abstract

This document describes a reference model for Autonomic Networking. The goal is to define how the various elements in an autonomic context work together, to describe their interfaces and relations. While the document is written as generally as possible, the initial solutions are limited to the chartered scope of the WG.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

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

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 1, 2016.

Copyright Notice

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

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

Table of Contents

1.	Introduction	3
2.	The Network View	4
3.	The Autonomic Network Element	5
3.1.	Architecture	5
3.2.	Full AN Nodes	6
3.3.	Constrained AN Nodes (*)	6
4.	The Autonomic Networking Infrastructure	6
4.1.	Naming	6
4.1.1.	Naming requirements	6
4.1.2.	Proposed Mechanisms	7
4.2.	Addressing	8
4.2.1.	Requirements and Fundamental Concepts	9
4.2.2.	The Base Addressing Scheme	10
4.2.3.	Possible Sub-Schemes	11
4.2.4.	Address Hierarchy	12
4.3.	Discovery	13
4.4.	Signaling Between Autonomic Nodes	13
4.5.	Intent Distribution	14
4.6.	Routing	14
4.7.	The Autonomic Control Plane	14
5.	Security and Trust Infrastructure	15
5.1.	Public Key Infrastructure	15
5.2.	Domain Certificate	15
5.3.	The MASA	15
5.4.	Sub-Domains (*)	15
5.5.	Cross-Domain Functionality (*)	15
6.	Autonomic Service Agents (ASA)	16
6.1.	General Description of an ASA	16
6.2.	Specific ASAs for the Enrolment Process	16
6.2.1.	The Enrolment ASA	16
6.2.2.	The Enrolment Proxy ASA	16

6.2.3.	The Registrar ASA	16
7.	Management and Programmability	16
7.1.	How an AN Network Is Managed	16
7.2.	Intent (*)	17
7.3.	Aggregated Reporting (*)	18
7.4.	Feedback Loops to NOC(*)	19
7.5.	Control Loops (*)	19
7.5.1.	Types of Control (*)	20
7.5.2.	Types of Control Loops (*)	20
7.5.3.	Management of an Autonomic Control Loop (*)	21
7.5.4.	Elements of an Autonomic Control Loop (*)	22
7.6.	APIs (*)	22
7.6.1.	Dynamic APIs (*)	22
7.6.2.	APIs and Semantics(*)	23
7.6.3.	API Considerations (*)	23
7.7.	Data Model (*)	23
8.	Coordination Between Autonomic Functions (*)	24
8.1.	The Coordination Problem (*)	24
8.2.	A Coordination Functional Block (*)	25
9.	Security Considerations	26
9.1.	Threat Analysis	26
10.	IANA Considerations	27
11.	Acknowledgements	27
12.	References	27
	Authors' Addresses	28

1. Introduction

The document "Autonomic Networking - Definitions and Design Goals" [[RFC7575](#)] explains the fundamental concepts behind Autonomic Networking, and defines the relevant terms in this space. In [section 5](#) it describes a high level reference model. This document defines this reference model with more detail, to allow for functional and protocol specifications to be developed in an architecturally consistent, non-overlapping manner. While the document is written as generally as possible, the initial solutions are limited to the chartered scope of the WG.

As discussed in [[RFC7575](#)], the goal of this work is not to focus exclusively on fully autonomic nodes or networks. In reality, most networks will run with some autonomic functions, while the rest of the network is traditionally managed. This reference model allows for this hybrid approach.

This is a living document and will evolve with the technical solutions developed in the ANIMA WG. Sections marked with (*) do not represent current charter items. While this document must give a

long term architectural view, not all functions will be standardized at the same time.

2. The Network View

This section describes the various elements in a network with autonomic functions, and how these entities work together, on a high level. Subsequent sections explain the detailed inside view for each of the autonomic network elements, as well as the network functions (or interfaces) between those elements.

Figure 1 shows the high level view of an Autonomic Network. It consists of a number of autonomic nodes, which interact directly with each other. Those autonomic nodes provide a common set of capabilities across the network, called the "Autonomic Networking Infrastructure" (ANI). The ANI provides functions like naming, addressing, negotiation, synchronization, discovery and messaging.

Autonomic functions typically span several, possibly all nodes in the network. The atomic entities of an autonomic function are called the "Autonomic Service Agents" (ASA), which are instantiated on nodes.

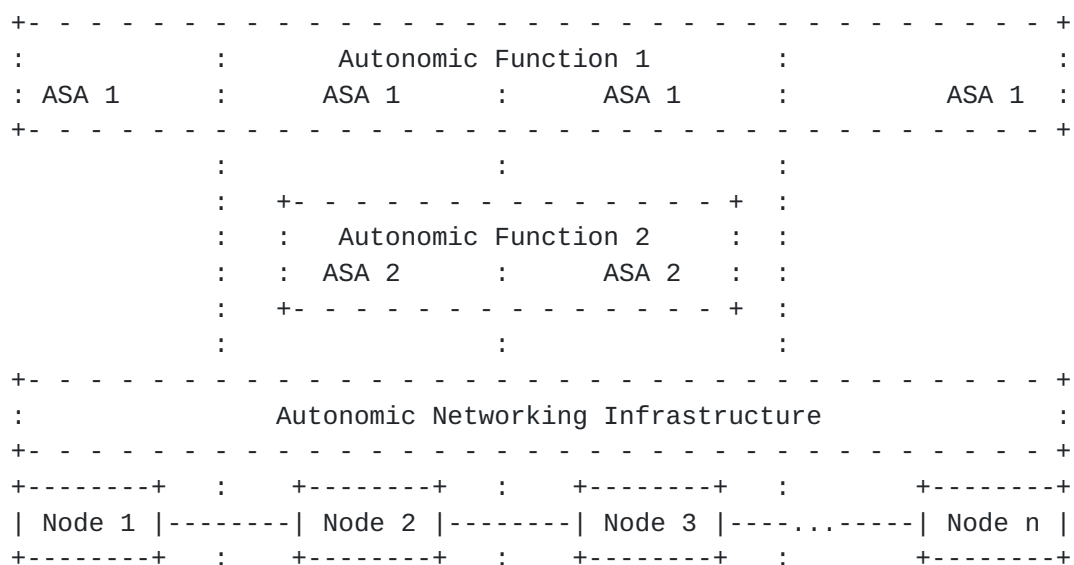


Figure 1: High level view of an Autonomic Network

In a horizontal view, autonomic functions span across the network, as well as the Autonomic Networking Infrastructure. In a vertical view, a node always implements the ANI, plus it may have one or several Autonomic Service Agents.

The Autonomic Networking Infrastructure (ANI) therefore is the foundation for autonomic functions. The current charter of the ANIMA

WG is to specify the ANI, using a few autonomic functions as use cases.

3. The Autonomic Network Element

3.1. Architecture

This section describes an autonomic network element and its internal architecture. The reference model explained in [\[I-D.irtf-nmrg-autonomic-network-definitions\]](#) shows the sources of information that an autonomic service agent can leverage: Self-knowledge, network knowledge (through discovery), Intent, and feedback loops. Fundamentally, there are two levels inside an autonomic node: the level of Autonomic Service Agents, and the level of the Autonomic Networking Infrastructure, with the former using the services of the latter. Figure 2 illustrates this concept.

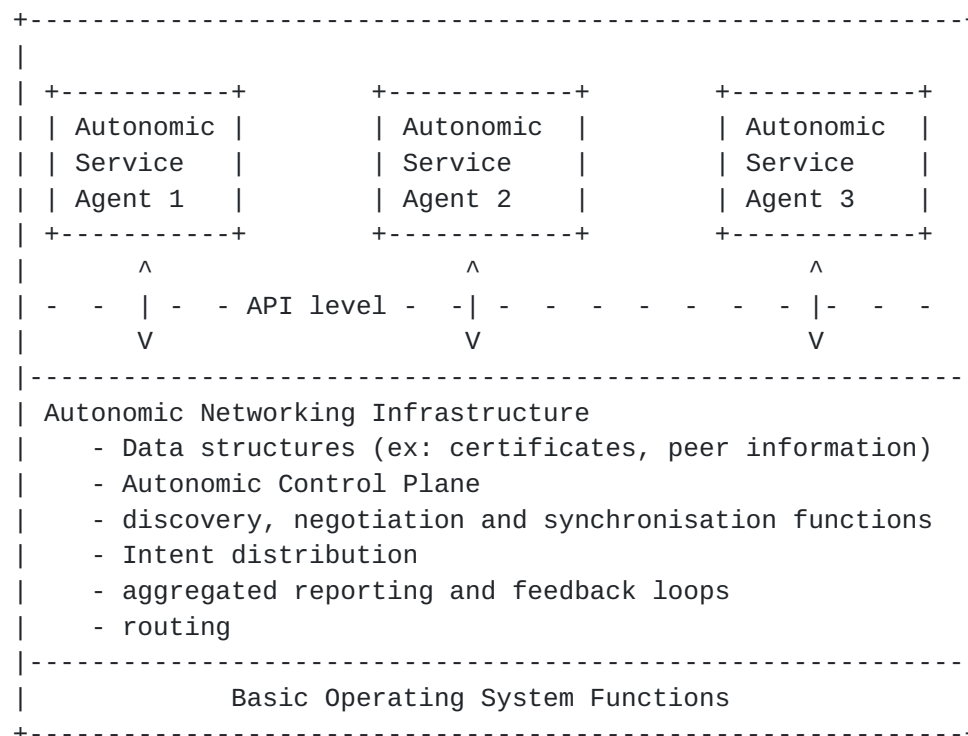


Figure 2: Model of an autonomic node

The Autonomic Networking Infrastructure (lower part of Figure 2) contains node specific data structures, for example trust information about itself and its peers, as well as a generic set of functions, independent of a particular usage. This infrastructure should be generic, and support a variety of Autonomic Service Agents (upper part of Figure 2). The Autonomic Control Plane is the summary of all

interactions of the Autonomic Networking Infrastructure with other nodes and services.

The use cases of "Autonomics" such as self-management, self-optimisation, etc, are implemented as Autonomic Service Agents. They use the services and data structures of the underlying autonomic networking infrastructure. The underlying Autonomic Networking Infrastructure should itself be self-managing.

The "Basic Operating System Functions" include the "normal OS", including the network stack, security functions, etc.

3.2. Full AN Nodes

Full AN nodes have the full Autonomic Networking Infrastructure, with the full functionality (details to be worked out). They support all the capabilities outlined in the rest of the document. [tbc]

3.3. Constrained AN Nodes (*)

Constrained nodes have a reduced ANI, with a well-defined minimal functionality (details to be worked out): They do need to be able to join the network, and communicate with at least a helper node which has full ANI functionality. Capabilities of constrained nodes need to be defined here. [tbc]

4. The Autonomic Networking Infrastructure

The Autonomic Networking Infrastructure provides a layer of common functionality across an Autonomic Network. It comprises "must implement" functions and services, as well as extensions.

An Autonomic Function, comprising of Autonomic Service Agents on nodes, can rely on the fact that all nodes in the network implement at least the "must implement" functions.

4.1. Naming

4.1.1. Naming requirements

- o Representing each device

Inside a domain, each autonomic device needs a domain specific identifier.

[Open Questions] Are there devices that don't need names? Do ASAs need names?

- o Uniqueness

The names MUST NOT collide within one autonomic domain.

It is acceptable that the names in different domains collide, since they could be distinguished by domains.

- o Semantic Encoding

It is RECOMMENDED that the names encode some semantics rather than meaningless strings. The semantics might be:

- + Location
- + Device type
- + Functional role
- + Ownership
- + etc.

This is for ease of management consideration that network administrators could easily recognize the device directly through the names.

- o Consistency

The devices' naming SHOULD follow the same pattern within a domain.

[4.1.2.](#) **Proposed Mechanisms**

—

- o Structured Naming Pattern

The whole name string could be divided into several fields, each of which representing a specific semantic as described above. For example: Location-DeviceType-FunctionalRole-DistinguisherNumber@NameofDomain.

The structure should be flexible that some fields are optional. When these optional fields are added, the name could still be recognized as the previous one. In above example, the "DistinguisherNumber" and "NameofDomain" are mandatory whereas others are optional. At initial stage, the devices might be only capable of self-generating the mandatory fields and the

"DeviceType" because of the lack of knowledge. Later, they might have learned the "Location" and "FunctionalRole" and added the fields into current name. However, the other devices could still recognize it according to the same "DistinguisherNumber".

- o Advertised Common Fields

Some fields in the structured name might be common among the domain (e.g. "Location" "NameofDomain"). Thus, these part of the names could be advertised through Intent Distribution [Section 4.5](#).

- o Self-generated Fields

The mandatory fields SHOULD be self-generated so that one device could name itself sufficiently without any advertised knowledges.

There should various methods for a device to extract/generate a proper word for each mandatory semantic fields (e.g. "DeviceType", "DistinguisherNum") from its self-knowledge.

Detailed design of specific naming patterns and methods are out of scope of this document.

[4.2](#). Addressing

Autonomic Service Agents (ASAs) need to communicate with each other, using the autonomic addressing of the node they reside on. This section describes the addressing approach of the Autonomic Networking Infrastructure, used by ASAs. It does NOT describe addressing approaches for the data plane of the network, which may be configured and managed in the traditional way, or negotiated as a service of an ASA. One use case for such an autonomic function is described in [\[I-D.jiang-auto-addr-management\]](#). The addressing of the Autonomic Networking Infrastructure is in scope for this section, the address space they negotiate for the data plane is not.

Autonomic addressing is a function of the Autonomic Networking Infrastructure (lower part of Figure 2). ASAs do not have their own addresses. They may use either API calls, or the autonomic addressing scheme of the Autonomic Networking Infrastructure.

4.2.1. Requirements and Fundamental Concepts

An autonomic addressing scheme has the following requirements:

- o Zero-touch for simple networks: Simple networks should have complete self-management of addressing, and not require any central address management, tools, or address planning.
- o Low-touch for complex networks: If complex networks require operator input for autonomic address management, it should be limited to high level guidance only, expressed in Intent.
- o Flexibility: The addressing scheme must be flexible enough for nodes to be able to move around, for the network to grow, split and merge.
- o Robustness: It should be as hard as possible for an administrator to negatively affect addressing (and thus connectivity) in the autonomic context.
- o Support for virtualization: Autonomic Nodes may support Autonomic Service Agents in different virtual machines or containers. The addressing scheme should support this architecture.
- o Simplicity: To make engineering simpler, and to give the human administrator an easy way to trouble-shoot autonomic functions.
- o Scale: The proposed scheme should work in any network of any size.
- o Upgradability: The scheme must be able to support different addressing concepts in the future.

These are the fundamental concepts of autonomic addressing:

- o IPv6 only: Autonomic processes SHOULD (as defined in [[RFC2119](#)]) use exclusively IPv6, for simplicity reasons.
- o 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.
- o Separation: Autonomic address space is used separately from user address space and other address realms. This supports the robustness requirement. Link-local is considered not part of user address space for this purpose.

- o Overlay network: Routeable addresses for AN nodes are used exclusively in a secure overlay network which is the basis of the ACP. This means that these addresses will be assigned to the loopback interface in most operating systems. All other interfaces exclusively use IPv6 link local for autonomic functions. The usage of IPv6 link local addressing is discussed in [[RFC7404](#)].
- o Use-ULA: For these overlay addresses of autonomic nodes, we use Unique Local Addresses (ULA), as specified in [[RFC4193](#)]. An alternative scheme was discussed, using assigned ULA addressing. The consensus was to use standard ULA, because it was deemed to be sufficient.
- o 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.

4.2.2. The Base Addressing Scheme

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



Figure 3: Base Addressing Scheme

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

- o "FD" identifies a locally defined ULA address.
- o The "global ID" is set here to be a hash of the domain name, which results in a pseudo-random 40 bit value. It is calculated as the first 40 bits of the MD5 hash of the domain name, in the example "example.com".
- o Type: Set to 000 (3 zero bits). This field allows different address sub-schemes in the future. The goal is to start with a minimal number of sub-scheme initially, but to allow for extensions later if and when required. This addresses the "upgradability" requirement. Assignment of types for this field should be maintained by IANA.

4.2.3. Possible Sub-Schemes

The sub-schemes listed here are not intended to be all supported initially, but are listed for discussion. The final document should define ideally only a single sub-scheme for now, and leave the other "types" for later assignment.

4.2.3.1. Sub-Scheme 1




Figure 4: Addressing Scheme 1

The fields are defined as follows: [Editor's note: The lengths of the fields is for discussion.]

- o Zone ID: If set to all zero bits: Flat addressing scheme. Any other value indicates a zone. See section [Section 4.2.4](#) on how this field is used in detail.
- o Device ID: A unique value for each device, typically assigned by a registrar.

The device ID is derived as follows: In an Autonomic Network, a registrar is enrolling new devices. As part of the enrolment process the registrar assigns a number to the device, which is unique for this registrar, but not necessarily unique in the domain. The 64 bit device ID is then composed as:

- o 48 bit: Registrar ID, a number unique inside the domain that identifies the registrar which assigned the name to the device. A MAC address of the registrar can be used for this purpose.
- o 16 bit: Device ID, a number which is unique for a given registrar, to identify the device. This can be a sequentially assigned number.

The "device ID" itself is unique in a domain (i.e., the Zone-ID is not required for uniqueness). Therefore, a device can be addressed either as part of a flat hierarchy (zone ID = 0), or with an aggregation scheme (any other zone ID). An address with zone-ID 0 (zero) could be interpreted as an identifier, with another zone-ID as a locator.

4.2.3.2. Sub-Scheme 2




Figure 5: Addressing Scheme 2

The fields are defined as follows: [Editor's note: The lengths of the fields is for discussion.]

- o Zone ID: As in sub-scheme 1.
- o Device ID: As in sub-scheme 1.
- o V: Virtualization bit(s): 1 or more bits that indicate a virtual context on an autonomic node.

In addition the scheme 1 ([Section 4.2.3.1](#)), this scheme allows the direct addressing of specific virtual containers / VMs on an autonomic node. An increasing number of hardware platforms have a distributed architecture, with a base OS for the node itself, and the support for hardware blades with potentially different OSs. The VMs on the blades could be considered as separate autonomic nodes, in which case it would make sense to be able to address them directly. Autonomic Service Agents (ASAs) could be instantiated in either the base OS, or one of the VMs on a blade. This addressing scheme allows for the easy separation of the hardware context.

The location of the V bit(s) at the end of the address allows to announce a single prefix for each autonomic node, while having separate virtual contexts addressable directly.

4.2.4. Address Hierarchy

The "zone ID" allows for the definition of a simple address hierarchy. If set to zero, the address scheme is flat. In this case, the addresses primarily act as identifiers for the nodes. Used like this, aggregation is not possible.

If aggregation is required, the 13 bit value allows for up to 8191 zones. (Theoretically, the 13 bits for the zone ID would allow also for two levels of zones, introducing a sub-hierarchy. We do not think this is required at this point, but a new type could be used in the future to support such a scheme.)

Another way to introduce hierarchy is to use sub-domains in the naming scheme. The node names "node17.subdomainA.example.com" and "node4.subdomainB.example.com" would automatically lead to different ULA prefixes, which can be used to introduce a routing hierarchy in the network, assuming that the subdomains are aligned with routing areas.

[4.3.](#) Discovery

Traditionally, most of the information a node requires is provided through configuration or northbound interfaces. An autonomic function should rely on such northbound interfaces minimally or not at all, and therefore it needs to discover peers and other resources in the network. This section describes various discovery functions in an autonomic network.

Discovering nodes and their properties and capabilities: A core function to establish an autonomic domain is the mutual discovery of autonomic nodes, primarily adjacent nodes and secondarily off-link peers. This may in principle either leverage existing discovery mechanisms, or use new mechanisms tailored to the autonomic context. An important point is that discovery must work in a network with no predefined topology, ideally no manual configuration of any kind, and with nodes starting up from factory condition or after any form of failure or sudden topology change.

Discovering services: Network services such as AAA should also be discovered and not configured. Service discovery is required for such tasks. An autonomic network can either leverage existing service discovery functions, or use a new approach, or a mixture.

Thus the discovery mechanism could either be fully integrated with autonomic signaling (next section) or could use an independent discovery mechanism such as DNS Service Discovery or Service Location Protocol. This choice could be made independently for each Autonomic Service Agent, although the infrastructure might require some minimal lowest common denominator (e.g., for discovering the security bootstrap mechanism, or the source of intent distribution, [Section 4.5](#)).

[4.4.](#) Signaling Between Autonomic Nodes

Autonomic nodes must communicate with each other, for example to negotiate and/or synchronize technical objectives (i.e., network parameters) of any kind and complexity. This requires some form of signaling between autonomic nodes. Autonomic nodes implementing a specific use case might choose their own signaling protocol, as long as it fits the overall security model. However, in the general case,

any pair of autonomic nodes might need to communicate, so there needs to be a generic protocol for this. A prerequisite for this is that autonomic nodes can discover each other without any preconfiguration, as mentioned above. To be generic, discovery and signaling must be able to handle any sort of technical objective, including ones that require complex data structures. The document "A Generic Discovery and Negotiation Protocol for Autonomic Networking"

[[I-D.carpenter-anima-gdn-protocol](#)] describes more detailed requirements for discovery, negotiation and synchronization in an autonomic network. It also defines a protocol, GDNP, for this purpose, including an integrated but optional discovery protocol.

[4.5.](#) Intent Distribution

Intent is the policy language of an Autonomic Network; see [Section 7.2](#) for general information on Intent. The distribution of Intent is also a function of the Autonomic Control Plane. It is expected that Intent will be expressed as quite complex human-readable data structures, and the distribution mechanism must be able to support that. Some Intent items will need to be flooded to most or all nodes, and other items of Intent may only be needed by a few nodes. Various methods could be used to distribute Intent across an autonomic domain. One approach is to treat it like any other technical objective needing to be synchronized across a set of nodes. In that case the autonomic signaling protocol could be used (previous section).

[4.6.](#) Routing

All autonomic nodes in a domain must be able to communicate with each other, and with autonomic nodes outside their own domain. Therefore, an Autonomic Control Plane relies on a routing function. For Autonomic Networks to be interoperable, they must all support one common routing protocol.

[4.7.](#) The Autonomic Control Plane

The totality of autonomic interactions forms the "Autonomic Control Plane". This control plane can be either implemented in the global routing table of a node, such as IGP's in today's networks; or it can be provided as an overlay network. The document "An Autonomic Control Plane" ([[I-D.behringer-anima-autonomic-control-plane](#)]) describes the details.

5. Security and Trust Infrastructure

An Autonomic Network is self-protecting. All protocols are secure by default, without the requirement for the administrator to explicitly configure security.

Autonomic nodes have direct interactions between themselves, which must be secured. Since an autonomic network does not rely on configuration, it is not an option to configure for example pre-shared keys. A trust infrastructure such as a PKI infrastructure must be in place. This section describes the principles of this trust infrastructure.

A completely autonomic way to automatically and securely deploy such a trust infrastructure is to set up a trust anchor for the domain, and then use an approach as in the document "Bootstrapping Key Infrastructures" [[I-D.pritikin-bootstrapping-keyinfrastructures](#)].

5.1. Public Key Infrastructure

An autonomic domain uses a PKI model. The root of trust is a certification authority (CA). A registrar acts as a registration authority (RA).

A minimum implementation of an autonomic domain contains one CA, one Registrar, and network elements.

5.2. Domain Certificate

We need to define how the fields in a domain certificate are to be used. [tbc]

5.3. The MASA

Explain briefly the function, point to [[I-D.pritikin-bootstrapping-keyinfrastructures](#)]. [tbc]

5.4. Sub-Domains (*)

Explain how sub-domains are handled. (tbc)

5.5. Cross-Domain Functionality (*)

Explain how trust is handled between different domains. (tbc)

6. Autonomic Service Agents (ASA)

This section describes how autonomic services run on top of the Autonomic Networking Infrastructure.

6.1. General Description of an ASA

general concepts, such as sitting on top of the ANI, etc. Also needs to explain that on a constrained node (see [Section 3.3](#)) not all ASAs may run, so we have two classes of ASAs: Ones that run on an unconstrained node, and limited function ASAs that run also on constrained nodes. We expect unconstrained nodes to support all ASAs.

6.2. Specific ASAs for the Enrolment Process

The following ASAs provide essential, required functionality in an autonomic network, and are therefore mandatory to implement on unconstrained autonomic nodes.

6.2.1. The Enrolment ASA

This section describes the function of an autonomic node to bootstrap into the domain with the help of an enrolment proxy (see previous section). [tbc]

6.2.2. The Enrolment Proxy ASA

This section describes the function of an autonomic node that helps a non-enrolled, adjacent devices to enrol into the domain. [tbc]

6.2.3. The Registrar ASA

This section describes the registrar function in an autonomic network. It explains the tasks of a registrar element, and how registrars are placed in a network, redundancy between several, etc. [tbc]

7. Management and Programmability

This section describes how an Autonomic Network is managed, and programmed.

7.1. How an AN Network Is Managed

Autonomic management usually co-exists with traditional management methods in most networks. Thus, autonomic behavior will be defined for individual functions in most environments. In fact, the co-

existence is twofold: autonomic functions can use traditional methods and protocols (e.g., SNMP and NETCONF) to perform management tasks; and autonomic functions can conflict with behavior enforced by the same traditional methods and protocols.

The autonomic intent is defined at a high level of abstraction. However, since it is necessary to address individual managed elements, autonomic management needs to communicate in lower-level interactions (e.g., commands and requests). For example, it is expected that the configuration of such elements be performed using NETCONF and YANG modules as well as the monitoring be executed through SNMP and MIBs.

Conflict can occur between autonomic default behavior, autonomic intent, traditional management methods. Conflict resolution is achieved in autonomic management through prioritization [[RFC7575](#)]. The rationale is that manual and node-based management have a higher priority over autonomic management. Thus, the autonomic default behavior has the lowest priority, then comes the autonomic Intent (medium priority), and, finally, the highest priority is taken by node-specific network management methods, such as the use of command line interfaces [[RFC7575](#)].

7.2. Intent (*)

This section describes Intent, and how it is managed. Intent and Policy-Based Network Management (PBNM) is already described inside the IETF (e.g., PCIM and SUPA) and in other SDOs (e.g., DMTF and TMF ZOOM).

Intent can be describe as an abstract, declarative, high-level policy used to operate an autonomic domain, such as an enterprise network [[RFC7575](#)]. Intent should be limited to high level guidance only, thus it does not directly define a policy for every network element separately. In an ideal autonomic domain, only one intent provided by human administrators is necessary to operate such domain [[RFC7576](#)]. However, it is als expected intent definition from autonomic function(s) and even from traditional network management elements (e.g., OSS).

Intent can be refined to lower level policies using different approaches, such as Policy Continuum model [ref]. This is expected in order to adapt the intent to the capabilities of managed devices. In this context, intent may contain role or function information, which can be translated to specific nodes [[RFC7575](#)]. One of the possible refinements of the intent is the refinement to Event Condition Action (ECA) rules. Such rules, which are more suitable to

individual entities, can be defined using different syntax and semantics.

Different parameters may be configured for intents. These parameters are usually provided by the human operator. Some of these parameters can influence the behavior of specific autonomic functions as well as the way the intent is used to manage the autonomic domain (towards intended operational point).

Some examples of parameters for intents are:

- o Model version: The version of the model used to define the intent.
- o Domain: The network scope in which the intent has effect.
- o Name: The name of the intent which describes the intent for human operators.
- o Version: The version of the intent, which is primarily used to control intent updates.
- o Signature: The signature is used as a security mechanism to provide authentication, integrity, and non-repudiation.
- o Timestamp: The timestamp of the creation of the intent using the format supported by the IETF [TBC].
- o Lifetime: The lifetime in which the intent may be observed. A special case of the lifetime is the definition of permanent intents.

Intent distribution is considered as one of the common control and management functions of an autonomic network [RFC7575]. Since distribution is fundamental for autonomic networking, it is necessary a mechanism to provision intent by all devices in a domain [[draft-carpenter-anima-gdn-protocol](#)]. The distribution of Intent is function of the Autonomic Control Plane and several methods can be used to distribute Intent across an autonomic domain [[draft-behringer-anima-reference-model](#)]. Intent distribution might not use the ANIMA signaling protocol itself [[draft-carpenter-anima-gdn-protocol](#)], but there is a proposal to extend such protocol for intent delivery [[draft-liu-anima-intent-distribution](#)].

7.3. Aggregated Reporting (*)

Autonomic Network should minimize the need for human intervention. In terms of how the network should behave, this is done through an autonomic intent provided by the human administrator. In an

analogous manner, the reports which describe the operational status of the network should aggregate the information produced in different network elements in order to present the effectiveness of autonomic intent enforcement. Therefore, reporting in an autonomic network should happen on a network-wide basis [[RFC7575](#)]. The information gathering and the reporting delivery should be done through the autonomic control plane.

Several events can occur in an autonomic network in the same way they can happen in a traditional network. These events can be produced considering traditional network management protocols, such as SNMP and syslog. However, when reporting to a human administrator, such events should be aggregated in order to avoid advertisement about individual managed elements. In this context, algorithms may be used to determine what should be reported (e.g., filtering) and in which way and how different events are related to each other. Besides that, an event in an individual element can be compensated by changes in other elements in order to maintain in a network-wide level which is described in the autonomic intent.

Reporting in an autonomic network may be in the same abstraction level of the intent. In this context, the visibility on current operational status of an autonomic network can be used to switch to different management modes. Despite the fact that autonomic management should minimize the need for user intervention, possibly there are some events that need to be addressed by human administrator actions. An alternative to model this is the use of exception-based management [[RFC7575](#)].

[7.4.](#) Feedback Loops to NOC(*)

Feedback loops are required in an autonomic network to allow the intervention of a human administrator or central control systems, while maintaining a default behaviour. Through a feedback loop an administrator can be prompted with a default action, and has the possibility to acknowledge or override the proposed default action.

[7.5.](#) Control Loops (*)

Control loops provide a generic mechanism for self-adaptation. That is, as user needs, business goals, and the ANI itself change, self-adaptation enables the ANI to change the services and resources it makes available to adapt to these changes. Self-adaptive systems move decision-making from static, pre-defined commands to dynamic processes computed at runtime.

Control loops operate to continuously capture data that enables the understanding of the system, and then provide actions to move the state of the system toward a common goal.

7.5.1. Types of Control (*)

There are two generic types of closed loop control. Feedback control adjusts the control loop based on measuring the output of the system being managed to generate an error signal (the deviation of the current state vs. its desired state). Action is then taken to reduce the deviation.

In contrast, feedforward control anticipates future effects on a controlled variable by measuring other variables whose values may be more timely, and adjusts the process based on those variables. In this approach, control is not error-based, but rather, based on knowledge.

Autonomic control loops MAY require both feedforward and feedback control.

7.5.2. Types of Control Loops (*)

There are many different types of control loops. In autonomics, the most commonly cited loop is called Monitor-Analyze-Plan-Execute (with Knowledge), called MAPE-K [Kephart03]. However, MAPE-K has a number of systemic problems, as described in [Strassner09]. Therefore, other autonomic architectures, such as AutoI [autoi] and FOCALE [Strassner07] and use control loops that evolved from the OODA (Observe-Orient-Decide-Act) control loop [Boyd95]. The reason for using this loop, and not the MAPE-K loop, is because the OODA loop contains a critical step not contained in other loops: orientation. Orientation determines how observations, decisions, and actions are performed.

Figure 6 shows a simplified model of a control loop containing both feedforward and feedback elements.




Figure 6: Control Loop with Feedforward and Feedback Elements

Note that Figure 6 is a STATIC model. Figure 7 is a dynamic version, called a Model-Reference Adaptive Control Loop (MRACL).

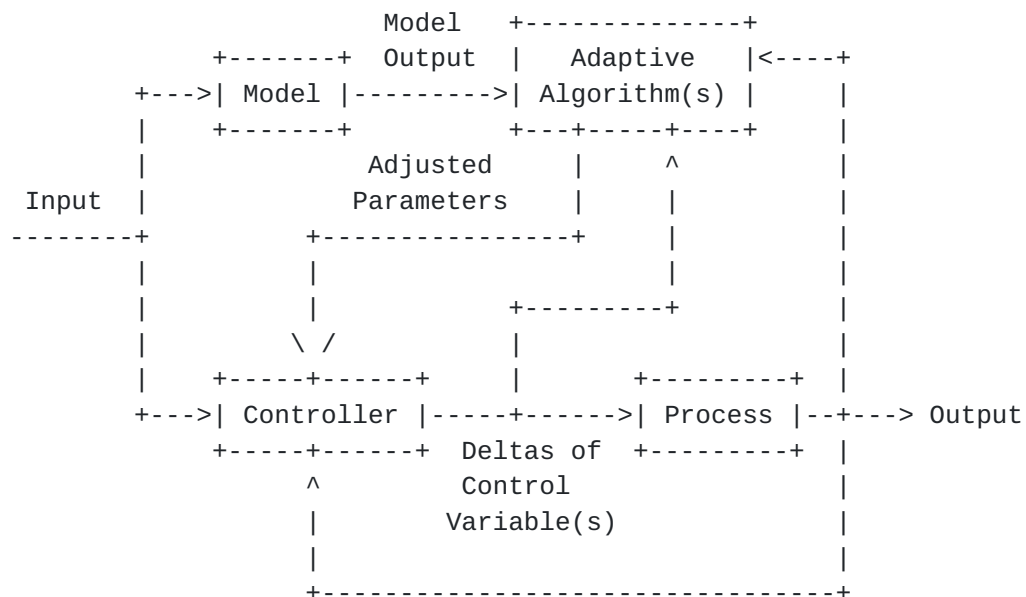


Figure 7: A Model-Reference Adaptive Control Loop

More complex adaptive control loops have been defined; these will be described in a future I-D, so that an appropriate gap analysis can be defined to recommend an architectural approach for ANIMA.

7.5.3. Management of an Autonomic Control Loop (*)

Both standard and adaptive control loops (e.g., as represented in Figures X and X1, respectively) enable intervention by a human administrator or central control systems, if required. Interaction mechanisms include changing the behaviour of one or more elements in

the control loop, as well as providing mechanisms to bypass parts of the control loop (e.g., skip the "decide" phase and go directly to the "action" phase of an OODA loop, as is done in FOCAL). This also enables the default behaviour to be changed if necessary.

7.5.4. Elements of an Autonomic Control Loop (*)

An autonomic control loop **MUST** be able to perform the following functions as part of its operation:

- o Observe and collect data from the system being managed
- o Orient these data, so that their meaning and significance can be understood in proper context
- o Analyze the collected data through filtering, correlation, and other mechanisms to define a model of past and current states
- o Plan different actions based on inferring trends, determining root causes, and similar processes
- o Decide which plan(s) to take
- o Execute the plan, and then repeat these steps

In addition, an autonomic control loop **SHOULD** be able to execute one or more machine learning algorithms that can learn from and make predictions on monitored data. This enables more efficient adaptivity. Note that machine learning is build from a model of exemplar inputs in order to make decisions and predictions. Supporting algorithms, such as those for data mining and analytics, **SHOULD** also be supported.

7.6. APIs (*)

Most APIs are static, meaning that they are pre-defined and represent an invariant mechanism for operating with data. An Autonomic Network **SHOULD** be able to use dynamic APIs in addition to static APIs. APIs **MUST** be able to express and preserve semantics across different domains.

7.6.1. Dynamic APIs (*)

A dynamic API is one that retrieves data using a generic mechanism, and then enables the client to navigate the retrieved data and operate on it. Such APIs typically use introspection and/or reflection (the former enables software to examine the type and properties of an object at runtime, while the latter enables a

program to manipulate the attributes, methods, and/or metadata of an object.

7.6.2. APIs and Semantics(*)

An API is NOT the same as an interface.

An interface is a boundary across which different components of a system exchange information. An API is a set of software (including tools, protocols, and programs) for building software applications. An API defines a set of data structures, inputs, outputs, and operations that can be used by a programmer to build an application.

An Autonomic API must pay particular attention to semantics. Previous designs have used the notion of a software contract to build high-quality APIs that are distributed and modular. A software contract [Meyer97] is based on the principle that a software-intensive system, such as an Autonomic Network, is a set of communicating components whose interaction is based on precisely-defined specification of the mutual obligations that interacting components must respect. For example, when a method executes, the following must hold:

- o pre-conditions must be satisfied before the method can start execution
- o post-conditions must be satisfied when the method has finished execution
- o invariant attributes must not change during the execution of the method

7.6.3. API Considerations (*)

APIs should perform one function well, not perform many different and unrelated functions. In software design, this is called the Single Responsibility Principle [srp]

7.7. Data Model (*)

The following definitions are taken from [supa-model]:

An information model is a representation of concepts of interest to an environment in a form that is independent of data repository, data definition language, query language, implementation language, and protocol. In contrast, a data model is a representation of concepts of interest to an environment in a form that is dependent on data

repository, data definition language, query language, implementation language, and protocol (typically, but not necessarily, all three).

The utility of an information model is to define objects and their relationships in a technology-neutral manner. This forms a consensual vocabulary that the ANI and ASAs can use. A data model is then a technology-specific mapping of all or part of the information model to be used by all or part of the system.

A system may have multiple data models. Operational Support Systems, for example, typically have multiple types of repositories, such as SQL and NoSQL, to take advantage of the different properties of each. If multiple data models are required by an Autonomic System, then an information model **SHOULD** be used to ensure that the concepts of each data model can be related to each other without technological bias.

A data model is essential for certain types of functions, such as a MRACL. More generally, a data model can be used to define the objects, attributes, methods, and relationships of a software system (e.g., the ANI, an autonomic node, or an ASA). A data model can be used to help design an API, as well as any language used to interface to the Autonomic Network.

8. Coordination Between Autonomic Functions (*)

8.1. The Coordination Problem (*)

Different autonomic functions may conflict in setting certain parameters. For example, an energy efficiency function may want to shut down a redundant link, while a load balancing function would not want that to happen. The administrator must be able to understand and resolve such interactions, to steer autonomic network performance to a given (intended) operational point.

Several interaction types may exist among autonomic functions, for example:

- o Cooperation: An autonomic function can improve the behavior or performance of another autonomic function, such as a traffic forecasting function used by a traffic allocation function.
- o Dependency: An autonomic function cannot work without another one being present or accessible in the autonomic network.
- o Conflict: A metric value conflict is a conflict where one metric is influenced by parameters of different autonomic functions. A parameter value conflict is a conflict where one parameter is modified by different autonomic functions.

Solving the coordination problem beyond one-by-one cases can rapidly become intractable for large networks. Specifying a common functional block on coordination is a first step to address the problem in a systemic way. The coordination life-cycle consists in three states:

- o At build-time, a "static interaction map" can be constructed on the relationship of functions and attributes. This map can be used to (pre-)define policies and priorities on identified conflicts.
- o At deploy-time, autonomic functions are not yet active/acting on the network. A "dynamic interaction map" is created for each instance of each autonomic functions and on a per resource basis, including the actions performed and their relationships. This map provides the basis to identify conflicts that will happen at run-time, categorize them and plan for the appropriate coordination strategies/mechanisms.
- o At run-time, when conflicts happen, arbitration is driven by the coordination strategies. Also new dependencies can be observed and inferred, resulting in an update of the dynamic interaction map and adaptation of the coordination strategies and mechanisms.

Multiple coordination strategies and mechanisms exists and can be devised. The set ranges from basic approaches such as random process or token-based process, to approaches based on time separation and hierarchical optimization, to more complex approaches such as multi-objective optimization, and other control theory approaches and algorithms family.

8.2. A Coordination Functional Block (*)

A common coordination functional block is a desirable component of the ANIMA reference model. It provides a means to ensure network properties and predictable performance or behavior such as stability, and convergence, in the presence of several interacting autonomic functions.

A common coordination function requires:

- o A common description of autonomic functions, their attributes and life-cycle.
- o A common representation of information and knowledge (e.g., interaction maps).

- o A common "control/command" interface between the coordination "agent" and the autonomic functions.

Guidelines, recommendations or BCPs can also be provided for aspects pertaining to the coordination strategies and mechanisms.

9. Security Considerations

9.1. Threat Analysis

This is a preliminary outline of a threat analysis, to be expanded and made more specific as the various Autonomic Networking specifications evolve.

Since AN will hand over responsibility for network configuration from humans or centrally established management systems to fully distributed devices, the threat environment is also fully distributed. On the one hand, that means there is no single point of failure to act as an attractive target for bad actors. On the other hand, it means that potentially a single misbehaving autonomic device could launch a widespread attack, by misusing the distributed AN mechanisms. For example, a resource exhaustion attack could be launched by a single device requesting large amounts of that resource from all its peers, on behalf of a non-existent traffic load. Alternatively it could simply send false information to its peers, for example by announcing resource exhaustion when this was not the case. If security properties are managed autonomically, a misbehaving device could attempt a distributed attack by requesting all its peers to reduce security protections in some way. In general, since autonomic devices run without supervision, almost any kind of undesirable management action could in theory be attempted by a misbehaving device.

If it is possible for an unauthorised device to act as an autonomic device, or for a malicious third party to inject messages appearing to come from an autonomic device, all these same risks would apply.

If AN messages can be observed by a third party, they might reveal valuable information about network configuration, security precautions in use, individual users, and their traffic patterns. If encrypted, AN messages might still reveal some information via traffic analysis, but this would be quite limited (for example, this would be highly unlikely to reveal any specific information about user traffic). AN messages are liable to be exposed to third parties on any unprotected Layer 2 link, and to insider attacks even on protected Layer 2 links.

10. IANA Considerations

This document requests no action by IANA.

11. Acknowledgements

Many people have provided feedback and input to this document: Sheng Jiang, Roberta Maglione, Jonathan Hansford.

12. References

- [I-D.behringer-anima-autonomic-addressing]
Behringer, M., "An Autonomic IPv6 Addressing Scheme", [draft-behringer-anima-autonomic-addressing-01](#) (work in progress), June 2015.
- [I-D.behringer-anima-autonomic-control-plane]
Behringer, M., Bjarnason, S., BL, B., and T. Eckert, "An Autonomic Control Plane", [draft-behringer-anima-autonomic-control-plane-02](#) (work in progress), March 2015.
- [I-D.carpenter-anima-gdn-protocol]
Carpenter, B. and B. Liu, "A Generic Discovery and Negotiation Protocol for Autonomic Networking", [draft-carpenter-anima-gdn-protocol-04](#) (work in progress), June 2015.
- [I-D.irtf-nmrg-autonomic-network-definitions]
Behringer, M., Pritikin, M., Bjarnason, S., Clemm, A., Carpenter, B., Jiang, S., and L. Ciavaglia, "Autonomic Networking - Definitions and Design Goals", [draft-irtf-nmrg-autonomic-network-definitions-07](#) (work in progress), March 2015.
- [I-D.jiang-auto-addr-management]
Jiang, S., Carpenter, B., and Q. Qiong, "Autonomic Networking Use Case for Auto Address Management", [draft-jiang-auto-addr-management-00](#) (work in progress), April 2014.
- [I-D.pritikin-bootstrapping-keyinfrastructures]
Pritikin, M., Behringer, M., and S. Bjarnason, "Bootstrapping Key Infrastructures", [draft-pritikin-bootstrapping-keyinfrastructures-01](#) (work in progress), September 2014.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", [RFC 4193](#), October 2005.
- [RFC7404] Behringer, M. and E. Vyncke, "Using Only Link-Local Addressing inside an IPv6 Network", [RFC 7404](#), November 2014.
- [RFC7575] Behringer, M., Pritikin, M., Bjarnason, S., Clemm, A., Carpenter, B., Jiang, S., and L. Ciavaglia, "Autonomic Networking: Definitions and Design Goals", [RFC 7575](#), June 2015.

Authors' Addresses

Michael H. Behringer (editor)
Cisco Systems
Building D, 45 Allée des Ormes
Mougins 06250
France

Email: mbehring@cisco.com

Brian Carpenter
Department of Computer Science
University of Auckland
PB 92019
Auckland 1142
New Zealand

Email: brian.e.carpenter@gmail.com

Toerless Eckert
Cisco

Email: eckert@cisco.com

Laurent Ciavaglia
Alcatel Lucent
Route de Villejust
Nozay 91620
France

Email: laurent.ciavaglia@alcatel-lucent.com

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

Email: leo.liubing@huawei.com

Jeferson Campos Nobre
Federal University of Rio Grande do Sul
Av. Bento Goncalves, 9500
Porto Alegre 91501-970
Brazil

Email: jcnobre@inf.ufrgs.br

John Strassner
Huawei Technologies
2330 Central Expressway
Santa Clara, CA 95050
USA

Email: john.sc.strassner@huawei.com

