

Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 28, 2018

B. Carpenter
Univ. of Auckland
L. Ciavaglia
Nokia
S. Jiang
Huawei Technologies Co., Ltd
P. Peloso
Nokia
October 25, 2017

Guidelines for Autonomic Service Agents
draft-carpen-ter-anima-asa-guidelines-03

Abstract

This document proposes guidelines for the design of Autonomic Service Agents for autonomic networks. It is based on the Autonomic Network Infrastructure outlined in the ANIMA reference model, making use of the Autonomic Control Plane and the Generic Autonomic Signaling 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 April 28, 2018.

Copyright Notice

Copyright (c) 2017 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
2.	Logical Structure of an Autonomic Service Agent	3
3.	Interaction with the Autonomic Networking Infrastructure	5
3.1.	Interaction with the security mechanisms	5
3.2.	Interaction with the Autonomic Control Plane	5
3.3.	Interaction with GRASP and its API	5
3.4.	Interaction with Intent mechanism	6
4.	Design of GRASP Objectives	6
5.	Life Cycle	7
5.1.	Installation phase	8
5.1.1.	Installation phase inputs and outputs	9
5.2.	Instantiation phase	9
5.2.1.	Operator's goal	10
5.2.2.	Instantiation phase inputs and outputs	10
5.2.3.	Instantiation phase requirements	11
5.3.	Operation phase	11
6.	Coordination	12
7.	Robustness	12
8.	Security Considerations	13
9.	IANA Considerations	14
10.	Acknowledgements	14
11.	References	14
11.1.	Normative References	14
11.2.	Informative References	14
Appendix A.	Change log [RFC Editor: Please remove]	16
	Authors' Addresses	16

[1.](#) Introduction

This document proposes guidelines for the design of Autonomic Service Agents (ASAs) in the context of an Autonomic Network (AN) based on the Autonomic Network Infrastructure (ANI) outlined in the ANIMA reference model [[I-D.ietf-anima-reference-model](#)]. This infrastructure makes use of the Autonomic Control Plane (ACP) [[I-D.ietf-anima-autonomic-control-plane](#)] and the Generic Autonomic Signaling Protocol (GRASP) [[I-D.ietf-anima-grasp](#)].

There is a considerable literature about autonomic agents with a variety of proposals about how they should be characterized. Some examples are [[DeMola06](#)], [[Huebscher08](#)], [[Movahedi12](#)] and [[GANA13](#)].

However, for the present document, the basic definitions and goals for autonomic networking given in [[RFC7575](#)] apply. According to [RFC 7575](#), an Autonomic Service Agent is "An agent implemented on an autonomic node that implements an autonomic function, either in part (in the case of a distributed function) or whole."

The reference model [[I-D.ietf-anima-reference-model](#)] expands this by adding that an ASA is "a process that makes use of the features provided by the ANI to achieve its own goals, usually including interaction with other ASAs via the GRASP protocol [[I-D.ietf-anima-grasp](#)] or otherwise. Of course it also interacts with the specific targets of its function, using any suitable mechanism. Unless its function is very simple, the ASA will need to be multi-threaded so that it can handle overlapping asynchronous operations. It may therefore be a quite complex piece of software in its own right, forming part of the application layer above the ANI."

A basic property of an ASA is that it is a relatively complex software component that will in many cases control and monitor simpler entities in the same host or elsewhere. For example, a device controller that manages tens or hundreds of simple devices might contain a single ASA.

The remainder of this document offers guidance on the design of ASAs.

2. Logical Structure of an Autonomic Service Agent

As mentioned above, all but the simplest ASAs will be multi-threaded programs.

A typical ASA will have a main thread that performs various initial housekeeping actions such as:

- o Obtain authorization credentials.
- o Register the ASA with GRASP.
- o Acquire relevant policy Intent.
- o Define data structures for relevant GRASP objectives.
- o Register with GRASP those objectives that it will actively manage.
- o Launch a self-monitoring thread.
- o Enter its main loop.

The logic of the main loop will depend on the details of the autonomic function concerned. Whenever asynchronous operations are required, extra threads will be launched. Examples of such threads include:

- o A background thread to repeatedly flood an objective to the AN, so that any ASA can receive the objective's latest value.
- o A thread to accept incoming synchronization requests for an objective managed by this ASA.
- o A thread to accept incoming negotiation requests for an objective managed by this ASA, and then to conduct the resulting negotiation with the counterpart ASA.
- o A thread to manage subsidiary non-autonomic devices directly.

These threads should all either exit after their job is done, or enter a wait state for new work, to avoid blocking other threads unnecessarily.

Note: If the programming environment does not support multi-threading, an 'event loop' style of implementation could be adopted, in which case each of the above threads would be implemented as an event handler called in turn by the main loop. In this case, the GRASP API ([Section 3.3](#)) must provide non-blocking calls. If necessary, the GRASP session identifier will be used to distinguish simultaneous negotiations.

According to the degree of parallelism needed by the application, some of these threads might be launched in multiple instances. In particular, if negotiation sessions with other ASAs are expected to be long or to involve wait states, the ASA designer might allow for multiple simultaneous negotiating threads, with appropriate use of queues and locks to maintain consistency.

The main loop itself could act as the initiator of synchronization requests or negotiation requests, when the ASA needs data or resources from other ASAs. In particular, the main loop should watch for changes in policy Intent that affect its operation. It should also do whatever is required to avoid unnecessary resource consumption, such as including an arbitrary wait time in each cycle of the main loop.

The self-monitoring thread is of considerable importance. Autonomic service agents must never fail. To a large extent this depends on careful coding and testing, with no unhandled error returns or exceptions, but if there is nevertheless some sort of failure, the

self-monitoring thread should detect it, fix it if possible, and in the worst case restart the entire ASA.

3. Interaction with the Autonomic Networking Infrastructure

3.1. Interaction with the security mechanisms

An ASA by definition runs in an autonomic node. Before any normal ASAs are started, such nodes must be bootstrapped into the autonomic network's secure key infrastructure in accordance with [\[I-D.ietf-anima-bootstrapping-keyinfra\]](#). This key infrastructure will be used to secure the ACP (next section) and may be used by ASAs to set up additional secure interactions with their peers, if needed.

Note that the secure bootstrap process itself may include special-purpose ASAs that run in a constrained insecure mode.

3.2. Interaction with the Autonomic Control Plane

In a normal autonomic network, ASAs will run as clients of the ACP. It will provide a fully secured network environment for all communication with other ASAs, in most cases mediated by GRASP (next section).

Note that the ACP formation process itself may include special-purpose ASAs that run in a constrained insecure mode.

3.3. Interaction with GRASP and its API

GRASP [\[I-D.ietf-anima-grasp\]](#) is expected to run as a separate process with its API [\[I-D.liu-anima-grasp-api\]](#) available in user space. Thus ASAs may operate without special privilege, unless they need it for other reasons. The ASA's view of GRASP is built around GRASP objectives ([Section 4](#)), defined as data structures containing administrative information such as the objective's unique name, and its current value. The format and size of the value is not restricted by the protocol, except that it must be possible to serialise it for transmission in CBOR [\[RFC7049\]](#), which is no restriction at all in practice.

The GRASP API offers the following features:

- o Registration functions, so that an ASA can register itself and the objectives that it manages.
- o A discovery function, by which an ASA can discover other ASAs supporting a given objective.

- o A negotiation request function, by which an ASA can start negotiation of an objective with a counterpart ASA. With this, there is a corresponding listening function for an ASA that wishes to respond to negotiation requests, and a set of functions to support negotiating steps.
- o A synchronization function, by which an ASA can request the current value of an objective from a counterpart ASA. With this, there is a corresponding listening function for an ASA that wishes to respond to synchronization requests.
- o A flood function, by which an ASA can cause the current value of an objective to be flooded throughout the AN so that any ASA can receive it.

For further details and some additional housekeeping functions, see [\[I-D.liu-anima-grasp-api\]](#).

This API is intended to support the various interactions expected between most ASAs, such as the interactions outlined in [Section 2](#). However, if ASAs require additional communication between themselves, they can do so using any desired protocol. One option is to use GRASP discovery and synchronization as a rendez-vous mechanism between two ASAs, passing communication parameters such as a TCP port number as the value of a GRASP objective. As noted above, either the ACP or in special cases the autonomic key infrastructure will be used to secure such communications.

[3.4.](#) Interaction with Intent mechanism

At the time of writing, the Intent mechanism for the ANI is undefined. It is expected to operate by an information distribution mechanism that can reach all autonomic nodes, and therefore every ASA. However, each ASA must be capable of operating "out of the box" in the absence of locally defined Intent, so every ASA implementation must include carefully chosen default values and settings for all parameters and choices that might depend on Intent.

[4.](#) Design of GRASP Objectives

The general rules for the format of GRASP Objective options, their names, and IANA registration are given in [\[I-D.ietf-anima-grasp\]](#). Additionally that document discusses various general considerations for the design of objectives, which are not repeated here. However, we emphasize that the GRASP protocol does not provide transactional integrity. In other words, if an ASA is capable of overlapping several negotiations for a given objective, then the ASA itself must use suitable locking techniques to avoid interference between these

negotiations. For example, if an ASA is allocating part of a shared resource to other ASAs, it needs to ensure that the same part of the resource is not allocated twice. This might impact the design of the objective as well as the logic flow of the ASA.

In particular, if 'dry run' mode is defined for the objective, its specification, and every implementation, must consider what state needs to be saved following a dry run negotiation, such that a subsequent live negotiation can be expected to succeed. It must be clear how long this state is kept, and what happens if the live negotiation occurs after this state is deleted. An ASA that requests a dry run negotiation must take account of the possibility that a successful dry run is followed by a failed live negotiation. Because of these complexities, the dry run mechanism should only be supported by objectives and ASAs where there is a significant benefit from it.

The actual value field of an objective is limited by the GRASP protocol definition to any data structure that can be expressed in Concise Binary Object Representation (CBOR) [[RFC7049](#)]. For some objectives, a single data item will suffice; for example an integer, a floating point number or a UTF-8 string. For more complex cases, a simple tuple structure such as [item1, item2, item3] could be used. Nothing prevents using other formats such as JSON, but this requires the ASA to be capable of parsing and generating JSON. The formats acceptable by the GRASP API will limit the options in practice. A fallback solution is for the API to accept and deliver the value field in raw CBOR, with the ASA itself encoding and decoding it via a CBOR library.

5. Life Cycle

Autonomic functions could be permanent, in the sense that ASAs are shipped as part of a product and persist throughout the product's life. However, a more likely situation is that ASAs need to be installed or updated dynamically, because of new requirements or bugs. Because continuity of service is fundamental to autonomic networking, the process of seamlessly replacing a running instance of an ASA with a new version needs to be part of the ASA's design.

The implication of service continuity on the design of ASAs can be illustrated along the three main phases of the ASA life-cycle, namely Installation, Instantiation and Operation.




Figure 1: Life cycle of an Autonomic Service Agent

5.1. Installation phase

Before being able to instantiate and run ASAs, the operator must first provision the infrastructure with the sets of ASA software corresponding to its needs and objectives. The provisioning of the infrastructure is realized in the installation phase and consists in installing (or checking the availability of) the pieces of software of the different ASA classes in a set of Installation Hosts.

There are 3 properties applicable to the installation of ASAs:

The dynamic installation property allows installing an ASA on demand, on any hosts compatible with the ASA.

The decoupling property allows controlling resources of a NE from a remote ASA, i.e. an ASA installed on a host machine different from the resources' NE.

The multiplicity property allows controlling multiple sets of resources from a single ASA.

These three properties are very important in the context of the installation phase as their variations condition how the ASA class could be installed on the infrastructure.

5.1.1. Installation phase inputs and outputs

Inputs are:

[ASA class of type_x] that specifies which classes ASAs to install,

[Installation_target_Infrastructure] that specifies the candidate
Installation Hosts,

[ASA class placement function, e.g. under which criteria/constraints
as defined by the operator]

that specifies how the installation phase shall meet the
operator's needs and objectives for the provision of the
infrastructure. In the coupled mode, the placement function is
not necessary, whereas in the decoupled mode, the placement
function is mandatory, even though it can be as simple as an
explicit list of Installation hosts.

The main output of the installation phase is an up-to-date directory
of installed ASAs which corresponds to [list of ASA classes]
installed on [list of installation Hosts]. This output is also
useful for the coordination function and corresponds to the static
interaction map (see next section).

The condition to validate in order to pass to next phase is to ensure
that [list of ASA classes] are well installed on [list of
installation Hosts]. The state of the ASA at the end of the
installation phase is: installed. (not instantiated). The following
commands or messages are foreseen: install(list of ASA classes,
Installation_target_Infrastructure, ASA class placement function),
and un-install (list of ASA classes).

5.2. Instantiation phase

Once the ASAs are installed on the appropriate hosts in the network,
these ASA may start to operate. From the operator viewpoint, an
operating ASA means the ASA manages the network resources as per the
objectives given. At the ASA local level, operating means executing
their control loop/algorithm.

But right before that, there are two things to take into
consideration. First, there is a difference between 1. having a
piece of code available to run on a host and 2. having an agent based
on this piece of code running inside the host. Second, in a coupled
case, determining which resources are controlled by an ASA is
straightforward (the determination is embedded), in a decoupled mode
determining this is a bit more complex (hence a starting agent will
have to either discover or be taught it).

The instantiation phase of an ASA covers both these aspects: starting the agent piece of code (when this does not start automatically) and determining which resources have to be controlled (when this is not obvious).

5.2.1. Operator's goal

Through this phase, the operator wants to control its autonomic network in two things:

- 1 determine the scope of autonomic functions by instructing which of the network resources have to be managed by which autonomic function (and more precisely which class e.g. 1. version X or version Y or 2. provider A or provider B),
- 2 determine how the autonomic functions are organized by instructing which ASAs have to interact with which other ASAs (or more precisely which set of network resources have to be handled as an autonomous group by their managing ASAs).

Additionally in this phase, the operator may want to set objectives to autonomic functions, by configuring the ASAs technical objectives.

The operator's goal can be summarized in an instruction to the ANIMA ecosystem matching the following pattern:

```
[ASA of type_x instances] ready to control  
[Instantiation\_target\_Infrastructure] with  
[Instantiation\_target\_parameters]
```

5.2.2. Instantiation phase inputs and outputs

Inputs are:

[ASA of type_x instances] that specifies which are the ASAs to be targeted (and more precisely which class e.g. 1. version X or version Y or 2. provider A or provider B),

[Instantiation_target_Infrastructure] that specifies which are the resources to be managed by the autonomic function, this can be the whole network or a subset of it like a domain a technology segment or even a specific list of resources,

[Instantiation_target_parameters] that specifies which are the technical objectives to be set to ASAs (e.g. an optimization target)

Outputs are:

[Set of ASAs - Resources relations] describing which resources are managed by which ASA instances, this is not a formal message, but a resulting configuration of a set of ASAs,

5.2.3. Instantiation phase requirements

The instructions described in [section 4.2](#) could be either:

sent to a targeted ASA In which case, the receiving Agent will have to manage the specified list of

[[Instantiation target Infrastructure](#)], with the [[Instantiation target parameters](#)].

broadcast to all ASAs In which case, the ASAs would collectively determine from the list which Agent(s) would handle which

[[Instantiation target Infrastructure](#)], with the [[Instantiation target parameters](#)].

This set of instructions can be materialized through a message that is named an Instance Mandate (description TBD).

The conclusion of this instantiation phase is a ready to operate ASA (or interacting set of ASAs), then this (or those) ASA(s) can describe themselves by depicting which are the resources they manage and what this means in terms of metrics being monitored and in terms of actions that can be executed (like modifying the parameters values). A message conveying such a self description is named an Instance Manifest (description TBD).

Though the operator may well use such a self-description "per se", the final goal of such a description is to be shared with other ANIMA entities like:

- o the coordination entities (see [[I-D.ciavaglia-anima-coordination](#)] - Autonomic Functions Coordination)
- o collaborative entities in the purpose of establishing knowledge exchanges (some ASAs may produce knowledge or even monitor metrics that other ASAs cannot make by themselves why those would be useful for their execution)

5.3. Operation phase

Note: This section is to be further developed in future revisions of the document, especially the implications on the design of ASAs.

During the Operation phase, the operator can:

Activate/Deactivate ASA: meaning enabling those to execute their autonomic loop or not.

Modify ASAs targets: meaning setting them different objectives.

Modify ASAs managed resources: by updating the instance mandate which would specify different set of resources to manage (only applicable to decouples ASAs).

During the Operation phase, running ASAs can interact the one with the other:

in order to exchange knowledge (e.g. an ASA providing traffic predictions to load balancing ASA)

in order to collaboratively reach an objective (e.g. ASAs pertaining to the same autonomic function targeted to manage a network domain, these ASA will collaborate - in the case of a load balancing one, by modifying the links metrics according to the neighboring resources loads)

During the Operation phase, running ASAs are expected to apply coordination schemes

then execute their control loop under coordination supervision/instructions

The ASA life-cycle is discussed in more detail in "A Day in the Life of an Autonomic Function" [[I-D.peloso-anima-autonomic-function](#)].

6. Coordination

Some autonomic functions will be completely independent of each other. However, others are at risk of interfering with each other - for example, two different optimization functions might both attempt to modify the same underlying parameter in different ways. In a complete system, a method is needed of identifying ASAs that might interfere with each other and coordinating their actions when necessary. This issue is considered in "Autonomic Functions Coordination" [[I-D.ciavaglia-anima-coordination](#)].

7. Robustness

It is of great importance that all components of an autonomic system are highly robust. In principle they must never fail. This section lists various aspects of robustness that ASA designers should consider.

1. If despite all precautions, an ASA does encounter a fatal error, it should in any case restart automatically and try again. To mitigate a hard loop in case of persistent failure, a suitable pause should be inserted before such a restart. The length of the pause depends on the use case.
2. If a newly received or calculated value for a parameter falls out of bounds, the corresponding parameter should be either left unchanged or restored to a safe value.
3. If a GRASP synchronization or negotiation session fails for any reason, it may be repeated after a suitable pause. The length of the pause depends on the use case.
4. If a session fails repeatedly, the ASA should consider that its peer has failed, and cause GRASP to flush its discovery cache and repeat peer discovery.
5. Any received GRASP message should be checked. If it is wrongly formatted, it should be ignored. Within a unicast session, an Invalid message (M_INVALID) may be sent. This function may be provided by the GRASP implementation itself.
6. Any received GRASP objective should be checked. If it is wrongly formatted, it should be ignored. Within a negotiation session, a Negotiation End message (M_END) with a Decline option (O_DECLINE) should be sent. An ASA may log such events for diagnostic purposes.
7. If an ASA receives either an Invalid message (M_INVALID) or a Negotiation End message (M_END) with a Decline option (O_DECLINE), one possible reason is that the peer ASA does not support a new feature of either GRASP or of the objective in question. In such a case the ASA may choose to repeat the operation concerned without using that new feature.
8. All other possible exceptions should be handled in an orderly way. There should be no such thing as an unhandled exception (but see point 1 above).

8. Security Considerations

ASAs are intended to run in an environment that is protected by the Autonomic Control Plane [[I-D.ietf-anima-autonomic-control-plane](#)], admission to which depends on an initial secure bootstrap process [[I-D.ietf-anima-bootstrapping-keyinfra](#)]. However, this does not relieve ASAs of responsibility for security. In particular, when ASAs configure or manage network elements outside the ACP, they must

use secure techniques and carefully validate any incoming information. As appropriate to their specific functions, ASAs should take account of relevant privacy considerations [[RFC6973](#)].

Authorization of ASAs is a subject for future study. At present, ASAs are trusted by virtue of being installed on a node that has successfully joined the ACP.

9. IANA Considerations

This document makes no request of the IANA.

10. Acknowledgements

TBD.

11. References

11.1. Normative References

- [I-D.ietf-anima-autonomic-control-plane]
Behringer, M., Eckert, T., and S. Bjarnason, "An Autonomic Control Plane (ACP)", [draft-ietf-anima-autonomic-control-plane-12](#) (work in progress), October 2017.
- [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-08](#) (work in progress), October 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), July 2017.
- [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>>.

11.2. Informative References

- [DeMola06]
De Mola, F. and R. Quitadamo, "An Agent Model for Future Autonomic Communications", Proceedings of the 7th WOA 2006 Workshop From Objects to Agents 51-59, September 2006.

- [GANA13] ETSI GS AFI 002, "Autonomic network engineering for the self-managing Future Internet (AFI): GANA Architectural Reference Model for Autonomic Networking, Cognitive Networking and Self-Management.", April 2013, <http://www.etsi.org/deliver/etsi_gs/AFI/001_099/002/01.01.01_60/gs_afi002v010101p.pdf>.
- [Huebscher08] Huebscher, M. and J. McCann, "A survey of autonomic computing--degrees, models, and applications", ACM Computing Surveys (CSUR) Volume 40 Issue 3 DOI: 10.1145/1380584.1380585, August 2008.
- [I-D.ciavaglia-anima-coordination] Ciavaglia, L. and P. Peloso, "Autonomic Functions Coordination", [draft-ciavaglia-anima-coordination-01](#) (work in progress), March 2016.
- [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](#) (work in progress), October 2017.
- [I-D.liu-anima-grasp-api] Carpenter, B., Liu, B., Wang, W., and X. Gong, "Generic Autonomic Signaling Protocol Application Program Interface (GRASP API)", [draft-liu-anima-grasp-api-05](#) (work in progress), October 2017.
- [I-D.peloso-anima-autonomic-function] Pierre, P. and L. Ciavaglia, "A Day in the Life of an Autonomic Function", [draft-peloso-anima-autonomic-function-01](#) (work in progress), March 2016.
- [Movahedi12] Movahedi, Z., Ayari, M., Langar, R., and G. Pujolle, "A Survey of Autonomic Network Architectures and Evaluation Criteria", IEEE Communications Surveys & Tutorials Volume: 14 , Issue: 2 DOI: 10.1109/SURV.2011.042711.00078, Page(s): 464 - 490, 2012.
- [RFC6973] Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", [RFC 6973](#), DOI 10.17487/RFC6973, July 2013, <<https://www.rfc-editor.org/info/rfc6973>>.

[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>>.

Appendix A. Change log [RFC Editor: Please remove]

[draft-carpenter-anima-asa-guidelines-03](#), 2017-10-25:

Added details on life cycle.

Added details on robustness.

Added co-authors.

[draft-carpenter-anima-asa-guidelines-02](#), 2017-07-01:

Expanded description of event-loop case.

Added note about 'dry run' mode.

[draft-carpenter-anima-asa-guidelines-01](#), 2017-01-06:

More sections filled in

[draft-carpenter-anima-asa-guidelines-00](#), 2016-09-30:

Initial version

Authors' Addresses

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

Email: brian.e.carpenter@gmail.com

Laurent Ciavaglia
Nokia
Villarceaux
Nozay 91460
FR

Email: laurent.ciavaglia@nokia.com

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

Email: jiangsheng@huawei.com

Pierre Peloso
Nokia
Villarceaux
Nozay 91460
FR

Email: pierre.peloso@nokia.com

