

Internet Research Task Force (IRTF)
Request for Comments: 8316
Category: Informational
ISSN: 2070-1721

J. Nobre
University of Vale do Rio dos Sinos
L. Granville
Federal University of Rio Grande do Sul
A. Clemm
Huawei
A. Gonzalez Prieto
VMware
February 2018

Autonomic Networking Use Case for Distributed Detection of
Service Level Agreement (SLA) Violations

Abstract

This document describes an experimental use case that employs autonomic networking for the monitoring of Service Level Agreements (SLAs). The use case is for detecting violations of SLAs in a distributed fashion. It strives to optimize and dynamically adapt the autonomic deployment of active measurement probes in a way that maximizes the likelihood of detecting service-level violations with a given resource budget to perform active measurements. This optimization and adaptation should be done without any outside guidance or intervention.

This document is a product of the IRTF Network Management Research Group (NMRG). It is published for informational purposes.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Research Task Force (IRTF). The IRTF publishes the results of Internet-related research and development activities. These results might not be suitable for deployment. This RFC represents the consensus of the Network Management Research Group of the Internet Research Task Force (IRTF). Documents approved for publication by the IRSG are not candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8316>.

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.

Table of Contents

1. Introduction	3
2. Definitions and Acronyms	5
3. Current Approaches	6
4. Use Case Description	7
5. A Distributed Autonomic Solution	8
6. Intended User Experience	10
7. Implementation Considerations	11
7.1. Device-Based Self-Knowledge and Decisions	11
7.2. Interaction with Other Devices	11
8. Comparison with Current Solutions	12
9. Related IETF Work	12
10. IANA Considerations	13
11. Security Considerations	13
12. Informative References	13
Acknowledgements	16
Authors' Addresses	16

1. Introduction

The Internet has been growing dramatically in terms of size, capacity, and accessibility in recent years. Communication requirements of distributed services and applications running on top of the Internet have become increasingly demanding. Some examples are real-time interactive video or financial trading. Providing such services involves stringent requirements in terms of acceptable latency, loss, and jitter.

Performance requirements lead to the articulation of Service Level Objectives (SLOs) that must be met. Those SLOs are part of Service Level Agreements (SLAs) that define a contract between the provider and the consumer of a service. SLOs, in effect, constitute a service-level guarantee that the consumer of the service can expect to receive (and often has to pay for). Likewise, the provider of a service needs to ensure that the service-level guarantee and associated SLOs are met. Some examples of clauses that relate to SLOs can be found in [RFC7297].

Violations of SLOs can be associated with significant financial loss, which can be divided into two categories. First, there is the loss that can be incurred by the user of a service when the agreed service levels are not provided. For example, a financial brokerage's stock orders might suffer losses when it is unable to execute stock transactions in a timely manner. An electronic retailer may lose customers when its online presence is perceived by customers as sluggish. An online gaming provider may not be able to provide fair access to online players, resulting in frustrated players who are lost as customers. In each case, the failure of a service provider to meet promised service-level guarantees can have a substantial financial impact on users of the service. Second, there is the loss that is incurred by the provider of a service who is unable to meet promised SLOs. Those losses can take several forms, such as penalties for violating the service level agreement and even loss of future revenue due to reduced customer satisfaction (which, in many cases, is more serious). Hence, SLOs are a key concern for the service provider. In order to ensure that SLOs are not being violated, service levels need to be continuously monitored at the network infrastructure layer in order to know, for example, when mitigating actions need to be taken. To that end, service-level measurements must take place.

Network measurements can be performed using active or passive measurement techniques. In passive measurements, production traffic is observed, and no monitoring traffic is created by the measurement process itself. That is, network conditions are checked in a non-intrusive way. In the context of IP Flow Information Export

(IPFIX), several documents were produced that define how to export data associated with flow records, i.e., data that is collected as part of passive measurement mechanisms, generally applied against flows of production traffic (e.g., [RFC7011]). In addition, it is possible to collect real data traffic (not just summarized flow records) with time-stamped packets, possibly sampled (e.g., per [RFC5474]), as a means of measuring and inferring service levels. Active measurements, on the other hand, are more intrusive to the network in the sense that they involve injecting synthetic test traffic into the network to measure network service levels, as opposed to simply observing production traffic. The IP Performance Metrics (IPPM) Working Group produced documents that describe active measurement mechanisms such as the One-Way Active Measurement Protocol (OWAMP) [RFC4656], the Two-Way Active Measurement Protocol (TWAMP) [RFC5357], and the Cisco Service-Level Assurance Protocol [RFC6812]. In addition, there are some mechanisms that do not cleanly fit into either active or passive categories, such as Performance and Diagnostic Metrics (PDM) Destination Option techniques [RFC8250].

Active measurement mechanisms offer a high level of control over what and how to measure. They do not require inspecting production traffic. Because of this, active measurements usually offer better accuracy and privacy than passive measurement mechanisms. Traffic encryption and regulations that limit the amount of payload inspection that can occur are non-issues. Furthermore, active measurement mechanisms are able to detect end-to-end network performance problems in a fine-grained way (e.g., simulating the traffic that must be handled considering specific SLOs). As a result, active measurements are often preferred over passive measurement for SLA monitoring. Measurement probes must be hosted in network devices and measurement sessions must be activated to compute the current network metrics (for example, metrics such as the ones described in [RFC4148], although note that [RFC4148] was obsoleted by [RFC6248]). This activation should be dynamic in order to follow changes in network conditions, such as those related to routes being added or new customer demands.

While offering many advantages, active measurements are expensive in terms of network resource consumption. Active measurements generally involve measurement probes that generate synthetic test traffic that is directed at a responder. The responder needs to timestamp test traffic it receives and reflect it back to the originating measurement probe. The measurement probe subsequently processes the returned packets along with time-stamping information in order to compute service levels. Accordingly, active measurements consume substantial CPU cycles as well as memory of network devices to

generate and process test traffic. In addition, synthetic traffic increases network load. Thus, active measurements compete for resources with other functions, including routing and switching.

The resources required and traffic generated by the active measurement sessions are, in a large part, a function of the number of measured network destinations. (In addition, the amount of traffic generated for each measurement plays a role that, in turn, influences the accuracy of the measurement.) When more destinations are measured, a greater number of resources are consumed and more traffic is needed to perform the measurements. Thus, to have better monitoring coverage, it is necessary to deploy more sessions, which consequently increases consumed resources. Otherwise, enabling the observation of just a small subset of all network flows can lead to insufficient coverage.

Furthermore, while some end-to-end service levels can be determined by adding up the service levels observed across different path segments, the same is not true for all service levels. For example, the end-to-end delay or packet loss from a node A to a node C routed via a node B can often be computed simply by adding delays (or loss) from A to B and from B to C. This allows the decomposition of a large set of end-to-end measurements into a much smaller set of segment measurements. However, end-to-end jitter and mean opinion scores cannot be decomposed as easily and, for higher accuracy, must be measured end-to-end.

Hence, the decision about how to place measurement probes becomes an important management activity. The goal is to obtain the maximum benefits of service-level monitoring with a limited amount of measurement overhead. Specifically, the goal is to maximize the number of service-level violations that are detected with a limited number of resources.

The use case and the solution approach described in this document address an important practical issue. They are intended to provide a basis for further experimentation to lead to solutions for wider deployment. This document represents the consensus of the IRTF's Network Management Research Group (NMRG). It was discussed extensively and received three separate in-depth reviews.

2. Definitions and Acronyms

Active Measurements: Techniques to measure service levels that involve generating and observing synthetic test traffic

Passive Measurements: Techniques used to measure service levels based on observation of production traffic

Autonomic Network: A network containing exclusively autonomic nodes, requiring no configuration, and deriving all required information through self-knowledge, discovery, or intent.

Autonomic Service Agent (ASA): An agent implemented on an autonomic node that implements an autonomic function, either in part (in the case of a distributed function, as in the context of this document) or whole

Measurement Session: A communications association between a probe and a responder used to send and reflect synthetic test traffic for active measurements

Probe: The source of synthetic test traffic in an active measurement

Responder: The destination for synthetic test traffic in an active measurement

SLA: Service Level Agreement

SLO: Service Level Objective

P2P: Peer-to-Peer

(Note: The definitions for "Autonomic Network" and "Autonomic Service Agent" are borrowed from [RFC7575]).

3. Current Approaches

For feasible deployments of active measurement solutions to distribute the available measurement sessions along the network, the current best practice consists of relying entirely on the human administrator's expertise to infer the best location to activate such sessions. This is done through several steps. First, it is necessary to collect traffic information in order to grasp the traffic matrix. Then, the administrator uses this information to infer the best destinations for measurement sessions. After that, the administrator activates sessions on the chosen subset of destinations, taking the available resources into account. This practice, however, does not scale well because it is still labor intensive and error-prone for the administrator to determine which sessions should be activated given the set of critical flows that needs to be measured. Even worse, this practice completely fails in networks where the most critical flows change rapidly, resulting in dynamic changes to what would be the most important destinations. For example, this can be the case in modern cloud environments. This is because fast reactions are necessary to reconfigure the sessions, and administrators are just not quick enough in computing and

activating the new set of required sessions every time the network traffic pattern changes. Finally, the current practice for active measurements usually covers only a fraction of the network flows that should be observed, which invariably leads to the damaging consequence of undetected SLA violations.

4. Use Case Description

The use case involves a service-level provider that needs to monitor the network to detect service-level violations using active service-level measurements and wants to be able to do so with minimal human intervention. The goal is to conduct the measurements in an effective manner to maximize the percentage of detected service-level violations. The service-level provider has a bounded resource budget with regard to measurements that can be performed, specifically the number of measurements that can be conducted concurrently from any one network device and possibly the total amount of measurement traffic on the network. However, while at any one point in time the number of measurements conducted is limited, it is possible for a device to change which destinations to measure over time. This can be exploited to achieve a balance of eventually covering all possible destinations using a reasonable amount of "sampling" where measurement coverage of a destination cannot be continuous. The solution needs to be dynamic and able to cope with network conditions that may change over time. The solution should also be embeddable inside network devices that control the deployment of active measurement mechanisms.

The goal is to conduct the measurements in a smart manner that ensures that the network is broadly covered and that the likelihood of detecting service-level violations is maximized. In order to maximize that likelihood, it is reasonable to focus measurement resources on destinations that are more likely to incur a violation, while spending fewer resources on destinations that are more likely to be in compliance. In order to do this, there are various aspects that can be exploited, including past measurements (destinations close to a service-level threshold requiring more focus than destinations farther from it), complementation with passive measurements such as flow data (to identify network destinations that are currently popular and critical), and observations from other parts of the network. In addition, measurements can be coordinated among different network devices to avoid hitting the same destination at the same time and to share results that may be useful in future probe placement.

Clearly, static solutions will have severe limitations. At the same time, human administrators cannot be in the loop for continuous dynamic reconfigurations of measurement probes. Thus, an automated

solution, or ideally an autonomic solution, is needed so that network measurements are automatically orchestrated and dynamically reconfigured from within the network. This can be accomplished using an autonomic solution that is distributed, using ASAs that are implemented on nodes in the network.

5. A Distributed Autonomic Solution

The use of Autonomic Networking (AN) [RFC7575] can help such detection through an efficient activation of measurement sessions. Such an approach, along with a detailed assessment confirming its viability, is described in [P2PBNM-Nobre-2012]. The problem to be solved by AN in the present use case is how to steer the process of measurement session activation by a complete solution that sets all necessary parameters for this activation to operate efficiently, reliably, and securely, with no required human intervention other than setting overall policy.

When a node first comes online, it has no information about which measurements are more critical than others. In the absence of information about past measurements and information from measurement peers, it may start with an initial set of measurement sessions, possibly randomly seeding a set of starter measurements and perhaps taking a round-robin approach for subsequent measurement rounds. However, as measurements are collected, a node will gain an increasing amount of information that it can utilize to refine its strategy of selecting measurement targets going forward. For one, it may take note of which targets returned measurement results very close to service-level thresholds; these targets may require closer scrutiny compared to others. Second, it may utilize observations that are made by its measurement peers in order to conclude which measurement targets may be more critical than others and to ensure that proper overall measurement coverage is obtained (so that not every node incidentally measures the same targets, while other targets are not measured at all).

We advocate for embedding P2P technology in network devices in order to use autonomic control loops to make decisions about measurement sessions.

Specifically, we advocate for network devices to implement an autonomic function that monitors service levels for violations of SLOs and that determines which measurement sessions to set up at any given point in time based on current and past observations of the node and of other peer nodes.

By performing these functions locally and autonomically on the device itself, which measurements to conduct can be modified quickly based

on local observations while taking local resource availability into account. This allows a solution to be more robust and react more dynamically to rapidly changing service levels than a solution that has to rely on central coordination. However, in order to optimize decisions about which measurements to conduct, a node will need to communicate with other nodes. This allows a node to take into account other nodes' observations in addition to its own in its decisions.

For example, remote destinations whose observed service levels are on the verge of violating stated objectives may require closer monitoring than remote destinations that are comfortably within a range of tolerance. A distributed autonomic solution also allows nodes to coordinate their probing decisions to collectively achieve the best possible measurement coverage. Because the number of resources available for monitoring, exchanging measurement data, and coordinating with other nodes is limited, a node may be interested in identifying other nodes whose observations are similar to and correlated with its own. This helps a node prioritize and decide which other nodes to coordinate and exchange data with. All of this requires the use of a P2P overlay.

A P2P overlay is essential for several reasons:

- o It makes it possible for nodes (or more specifically, the ASAs that are deployed on those nodes) in the network to autonomically set up measurement sessions without having to rely on a central management system or controller to perform configuration operations associated with configuring measurement probes and responders.
- o It facilitates the exchange of data between different nodes to share measurement results so that each node can refine its measurement strategy based not just on its own observations, but also on observations from its peers.
- o It allows nodes to coordinate their measurements to obtain the best possible test coverage and avoid measurements that have a very low likelihood of detecting service-level violations.

The provisioning of the P2P overlay should be transparent for the network administrator. An Autonomic Control Plane such as defined in [ACP] provides an ideal candidate for the P2P overlay to run on.

An autonomic solution for the distributed detection of SLA violations provides several benefits. First, it provides efficiency; this solution should optimize the resource consumption and avoid resource starvation on the network devices. A device that is "self-aware" of

its available resources will be able to adjust measurement activities rapidly as needed, without requiring a separate control loop involving resource monitoring by an external system. Second, placing logic about where to conduct measurements into the node enables rapid control loops that allow devices to react instantly to observations and adjust their measurement strategy. For example, a device could decide to adjust the amount of synthetic test traffic being sent during the measurement itself depending on results observed so far on this and other concurrent measurement sessions. As a result, the solution could decrease the time necessary to detect SLA violations. Adaptivity features of an autonomic loop could capture the network dynamics faster than a human administrator or even a central controller. Finally, the solution could help to reduce the workload of human administrators.

In practice, these factors combine to maximize the likelihood of SLA violations being detected while operating within a given resource budget, allowing a continuous measurement strategy that takes into account past measurement results to be conducted, observations of other measures such as link utilization or flow data, measurement results shared between network devices, and future measurement activities coordinated among nodes. Combined, this can result in efficient measurement decisions that achieve a golden balance between offering broad network coverage and honing in on service-level "hot spots".

6. Intended User Experience

The autonomic solution should not require any human intervention in the distributed detection of SLA violations. By virtue of the solution being autonomic, human users will not have to plan which measurements to conduct in a network, which is often a very labor-intensive task that requires detailed analysis of traffic matrices and network topologies and is not prone to easy dynamic adjustment. Likewise, they will not have to configure measurement probes and responders.

There are some ways in which a human administrator may still interact with the solution. First, the human administrator will, of course, be notified and obtain reports about service-level violations that are observed. Second, a human administrator may set policies regarding how closely to monitor the network for service-level violations and how many resources to spend. For example, an administrator may set a resource budget that is assigned to network devices for measurement operations. With that given budget, the number of SLO violations that are detected will be maximized. Alternatively, an administrator may set a target for the percentage of SLO violations that must be detected, i.e., a target for the ratio

between the number of detected SLO violations and the number of total SLO violations that are actually occurring (some of which might go undetected). In that case, the solution will aim to minimize the resources spent (i.e., the amount of test traffic and number of measurement sessions) that are required to achieve that target.

7. Implementation Considerations

The active measurement model assumes that a typical infrastructure will have multiple network segments, multiple Autonomous Systems (ASes), and a reasonably large number of routers. It also considers that multiple SLOs can be in place at a given time. Since interoperability in a heterogeneous network is a goal, features found on different active measurement mechanisms (e.g., OWAMP, TWAMP, and Cisco Service Level Assurance Protocol) and device programmability interfaces (such as Juniper's Junos API or Cisco's Embedded Event Manager) could be used for the implementation. The autonomic solution should include and/or reference specific algorithms, protocols, metrics, and technologies for the implementation of distributed detection of SLA violations as a whole.

Finally, it should be noted that there are multiple deployment scenarios, including deployment scenarios that involve physical devices hosting autonomic functions or virtualized infrastructure hosting the same. Co-deployment in conjunction with Virtual Network Functions (VNFs) is a possibility for further study.

7.1. Device-Based Self-Knowledge and Decisions

Each device has self-knowledge about the local SLA monitoring. This could be in the form of historical measurement data and SLOs. Besides that, the devices would have algorithms that could decide which probes should be activated at a given time. The choice of which algorithm is better for a specific situation would be also autonomic.

7.2. Interaction with Other Devices

Network devices should share information about service-level measurement results. This information can speed up the detection of SLA violations and increase the number of detected SLA violations. For example, if one device detects that a remote destination is in danger of violating an SLO, other devices may conduct additional measurements to the same destination or other destinations in its proximity. For any given network device, the exchange of data may be more important with some devices (for example, devices in the same network neighborhood or devices that are "correlated" by some other means) than with others. Defining the network devices that exchange

measurement data (i.e., management peers) creates a new topology. Different approaches could be used to define this topology (e.g., correlated peers [P2PBNM-Nobre-2012]). To bootstrap peer selection, each device should use its known neighbors (e.g., FIB and RIB tables) as initial seeds to identify possible peers. It should be noted that a solution will benefit if topology information and network discovery functions are provided by the underlying autonomic framework. A solution will need to be able to discover measurement peers as well as measurement targets, specifically measurement targets that support active measurement responders and that will be able to respond to measurement requests and reflect measurement traffic as needed.

8. Comparison with Current Solutions

There is no standardized solution for distributed autonomic detection of SLA violations. Current solutions are restricted to ad hoc scripts running on a per-node fashion to automate some administrator actions. There are some proposals for passive probe activation (e.g., DECON [DECON] and CSAMP [CSAMP]), but these do not focus on autonomic features.

9. Related IETF Work

This section discusses related IETF work and is provided for reference. This section is not exhaustive; rather, it provides an overview of the various initiatives and how they relate to autonomic distributed detection of SLA violations.

1. LMAP: The Large-Scale Measurement of Broadband Performance Working Group standardizes the LMAP measurement system for performance management of broadband access devices. The autonomic solution could be relevant to LMAP because it deploys measurement probes and could be used for screening for SLA violations. Besides that, a solution to decrease the workload of human administrators in service providers is probably highly desirable.
2. IPFIX: IP Flow Information Export (IPFIX) Working Group (now concluded) aimed to standardize IP flows (i.e., netflows). IPFIX uses measurement probes (i.e., metering exporters) to gather flow data. In this context, the autonomic solution for the activation of active measurement probes could possibly be extended to also address passive measurement probes. Besides that, flow information could be used in making decisions regarding probe activation.

3. ALTO: The Application-Layer Traffic Optimization Working Group aims to provide topological information at a higher abstraction layer, which can be based upon network policy, and with application-relevant service functions located in it. Their work could be leveraged to define the topology for network devices that exchange measurement data.

10. IANA Considerations

This document has no IANA actions.

11. Security Considerations

The security of this solution hinges on the security of the network underlay, i.e., the Autonomic Control Plane. If the Autonomic Control Plane were to be compromised, an attacker could undermine the effectiveness of measurement coordination by reporting fraudulent measurement results to peers. This would cause measurement probes to be deployed in an ineffective manner that would increase the likelihood that violations of SLOs go undetected.

Likewise, the security of the solution hinges on the security of the deployment mechanism for autonomic functions (in this case, the autonomic function that conducts the service-level measurements). If an attacker were able to hijack an autonomic function, it could try to exhaust or exceed the resources that should be spent on autonomic measurements in order to deplete network resources, including network bandwidth due to higher-than-necessary volumes of synthetic test traffic generated by measurement probes. Again, it could also lead to reporting of misleading results; among other things, this could result in non-optimal selection of measurement targets and, in turn, an increase in the likelihood that service-level violations go undetected.

12. Informative References

- [ACP] Eckert, T., Ed., Behringer, M., Ed., and S. Bjarnason, "An Autonomic Control Plane (ACP)", Work in Progress, draft-ietf-anima-autonomic-control-plane-13, December 2017.
- [CSAMP] Sekar, V., Reiter, M., Willinger, W., Zhang, H., Kompella, R., and D. Andersen, "CSAMP: A System for Network-Wide Flow Monitoring", NSDI USENIX Symposium Networked Systems Design and Implementation, April 2008.

- [DECON] di Pietro, A., Huici, F., Costantini, D., and S. Niccolini, "DECON: Decentralized Coordination for Large-Scale Flow Monitoring", IEEE INFOCOM Workshops, DOI 10.1109/INFCOMW.2010.5466642, March 2010.
- [P2PBNM-Nobre-2012] Nobre, J., Granville, L., Clemm, A., and A. Gonzalez Prieto, "Decentralized Detection of SLA Violations Using P2P Technology, 8th International Conference Network and Service Management (CNSM)", 8th International Conference on Network and Service Management (CNSM), 2012, <http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6379997>.
- [RFC4148] Stephan, E., "IP Performance Metrics (IPPM) Metrics Registry", BCP 108, RFC 4148, DOI 10.17487/RFC4148, August 2005, <<https://www.rfc-editor.org/info/rfc4148>>.
- [RFC4656] Shalunov, S., Teitelbaum, B., Karp, A., Boote, J., and M. Zekauskas, "A One-way Active Measurement Protocol (OWAMP)", RFC 4656, DOI 10.17487/RFC4656, September 2006, <<https://www.rfc-editor.org/info/rfc4656>>.
- [RFC5357] Hedayat, K., Krzanowski, R., Morton, A., Yum, K., and J. Babiarez, "A Two-Way Active Measurement Protocol (TWAMP)", RFC 5357, DOI 10.17487/RFC5357, October 2008, <<https://www.rfc-editor.org/info/rfc5357>>.
- [RFC5474] Duffield, N., Ed., Chiou, D., Claise, B., Greenberg, A., Grossglauser, M., and J. Rexford, "A Framework for Packet Selection and Reporting", RFC 5474, DOI 10.17487/RFC5474, March 2009, <<https://www.rfc-editor.org/info/rfc5474>>.
- [RFC6248] Morton, A., "RFC 4148 and the IP Performance Metrics (IPPM) Registry of Metrics Are Obsolete", RFC 6248, DOI 10.17487/RFC6248, April 2011, <<https://www.rfc-editor.org/info/rfc6248>>.
- [RFC6812] Chiba, M., Clemm, A., Medley, S., Salowey, J., Thombare, S., and E. Yedavalli, "Cisco Service-Level Assurance Protocol", RFC 6812, DOI 10.17487/RFC6812, January 2013, <<https://www.rfc-editor.org/info/rfc6812>>.
- [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>>.

- [RFC7297] Boucadair, M., Jacquenet, C., and N. Wang, "IP Connectivity Provisioning Profile (CPP)", RFC 7297, DOI 10.17487/RFC7297, July 2014, <<https://www.rfc-editor.org/info/rfc7297>>.
- [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>>.
- [RFC8250] Elkins, N., Hamilton, R., and M. Ackermann, "IPv6 Performance and Diagnostic Metrics (PDM) Destination Option", RFC 8250, DOI 10.17487/RFC8250, September 2017, <<https://www.rfc-editor.org/info/rfc8250>>.

Acknowledgements

We wish to acknowledge the helpful contributions, comments, and suggestions that were received from Mohamed Boucadair, Brian Carpenter, Hanlin Fang, Bruno Klauser, Diego Lopez, Vincent Roca, and Eric Voit. In addition, we thank Diego Lopez, Vincent Roca, and Brian Carpenter for their detailed reviews.

Authors' Addresses

Jeferson Campos Nobre
University of Vale do Rio dos Sinos
Porto Alegre
Brazil

Email: jcnobre@unisinos.br

Lisandro Zambenedetti Granville
Federal University of Rio Grande do Sul
Porto Alegre
Brazil

Email: granville@inf.ufrgs.br

Alexander Clemm
Huawei USA - Futurewei Technologies Inc.
Santa Clara, California
United States of America

Email: ludwig@clemm.org, alexander.clemm@huawei.com

Alberto Gonzalez Prieto
VMware
Palo Alto, California
United States of America

Email: agonzalezpri@vmware.com