

Network Management Research Group  
Internet-Draft  
Intended status: Informational  
Expires: September 3, 2017

J. Nobre  
University of Vale do Rio dos Sinos  
L. Granville  
Federal University of Rio Grande do Sul  
A. Clemm  
Huawei  
A. Gonzalez Prieto  
Cisco Systems  
March 2, 2017

**Autonomic Networking Use Case for Distributed Detection of SLA  
Violations  
draft-irtf-nmrg-autonomic-sla-violation-detection-07**

**Abstract**

This document describes a use case for autonomic networking in distributed detection of Service Level Agreement (SLA) violations.

**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 September 3, 2017.

**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 (<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

<a href="#">1.</a>	Introduction . . . . .	<a href="#">2</a>
<a href="#">2.</a>	Definitions and Acronyms . . . . .	<a href="#">5</a>
<a href="#">3.</a>	Current Approaches . . . . .	<a href="#">5</a>
<a href="#">4.</a>	Use Case Description . . . . .	<a href="#">6</a>
<a href="#">5.</a>	A Distributed Autonomic Solution . . . . .	<a href="#">7</a>
<a href="#">6.</a>	Intended User and Administrator Experience . . . . .	<a href="#">8</a>
<a href="#">7.</a>	Analysis of Parameters and Information Involved . . . . .	<a href="#">8</a>
<a href="#">7.1.</a>	Device Based Self-Knowledge and Decisions . . . . .	<a href="#">8</a>
<a href="#">7.2.</a>	Interaction with other devices . . . . .	<a href="#">9</a>
<a href="#">8.</a>	Comparison with current solutions . . . . .	<a href="#">9</a>
<a href="#">9.</a>	Related IETF Work . . . . .	<a href="#">9</a>
<a href="#">10.</a>	Acknowledgements . . . . .	<a href="#">10</a>
<a href="#">11.</a>	IANA Considerations . . . . .	<a href="#">10</a>
<a href="#">12.</a>	Security Considerations . . . . .	<a href="#">10</a>
<a href="#">13.</a>	Informative References . . . . .	<a href="#">10</a>
	Authors' Addresses . . . . .	<a href="#">12</a>

## [1.](#) Introduction

The Internet has been growing dramatically in terms of size and capacity, and accessibility in the last 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, or jitter.

Performance requirements lead to the articulation of Service Level Objectives (SLOs) which 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 service level objectives can be found in [[RFC7297](#)]).

Violations of SLOs can be associated with significant financial loss, which can be divided into two categories. For one, 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 their 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. By the same token, there is the loss that is incurred by the provider of a service who is unable to meet promised service level objectives. Those losses can take several forms, such as penalties for not meeting the service and, in many cases more important, loss of revenue due to reduced customer satisfaction. Hence, service level objectives 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) WG, several documents were produced to define passive measurement mechanisms (e.g., flow records specification [[RFC3954](#)]). Active measurements, on the other hand, are intrusive in the sense that it involves injecting synthetic test traffic into the network to measure network service levels. The IP Performance Metrics (IPPM) WG produced documents that describe active measurement mechanisms, such as: One-Way Active Measurement Protocol (OWAMP) [[RFC4656](#)], Two-Way Active Measurement Protocol (TWAMP) [[RFC5357](#)], and Cisco Service Level Assurance Protocol (SLA) [[RFC6812](#)]. In addition, there are some mechanisms that do not fit into either active or passive categories, such as Performance and Diagnostic Metrics Destination Option (PDM) techniques [[draft-ietf-ippm-6man-pdm-option](#)].

Active measurement mechanisms offer a high level of control of 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 Service Level Objectives - 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 (e.g., considering those described in [[RFC4148](#)]). This activation should be dynamic in order to follow changes in network conditions, such as those related with 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. Active measurements thus compete for resources with other functions, including routing and switching.

The resources required and traffic generated by the active measurement sessions are to 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, which in turn influences the accuracy of the measurement.) The more destinations are being measured, the larger the amount of resources consumed and traffic needed to perform the measurements. Thus, to have a better monitoring coverage it is necessary to deploy more sessions which consequently turns increases consumed resources. Otherwise, enabling the observation of just a small subset of all network flows can lead to an 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 B to C. This allows to decompose a large set of end-to-end measurements into a much smaller set of segment measurements. However, end-to-end jitter and (for example) Mean Opinion Scores cannot be decomposed as easily and, for higher accuracy, must be measured end-to-end.

Hence, the decision how to place measurement probes becomes an important management activity. The goal is to obtain 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 amount of resources.



## 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

AN: Autonomic Network

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

SL0: Service Level Objective

P2P: Peer-to-Peer

## 3. Current Approaches

The current best practice in feasible deployments of active measurement solutions to distribute the available measurement sessions along the network consists in relying entirely on the human administrator expertise to infer which would be 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 which are the best destinations for measurement sessions. After that, the administrator activates sessions on the chosen subset of destinations considering the available resources. 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 whose critical flows are too short in time and dynamic in terms of traversing network path, like in modern cloud environments. That is so because fast reactions are necessary to reconfigure the sessions and administrators are not just enough in computing and activating the new set of required sessions every time the network traffic pattern changes. Finally, the current active measurements practice 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 who 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 maximizing the percentage of detected service level violations. The service level provider has a bounded resource budget with regards to measurements that can be performed, specifically, with regards to the number of measurements that can be conducted concurrently from any one network device. 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 be able to cope with network conditions which 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 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 less resources on destinations that are more likely to be in compliance. In order to do so, there are various aspects that can be exploited, including past measurements (destinations close to a service level threshold requiring more focus than destinations further from it), complementation with passive measurements such as flow data (to identify network destinations that are currently popular and critical), an 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 be able 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 measurement probe reconfigurations. Accordingly, an automated or, ideally, autonomic solution is needed in which network measurements are automatically orchestrated and dynamically reconfigured from within the network.



## 5. A Distributed Autonomic Solution

The use of Autonomic Networking (AN) can help such detection through an efficient activation of measurement sessions [[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.

We advocate for embedding Peer-to-Peer (P2P) technology in network devices in order to conduct the measurement session activation decisions using autonomic control loops. This requires the use of a P2P overlay. A P2P overlay is important for several reasons:

- o It makes it possible for nodes in the network to autonomically set up Measurement Sessions, without having to rely on central management system or controller to perform configuration operations associated with configuring measurement probes and responders.
- o It facilitates the exchange local data between different devices that is used to coordinate measurements and to share measurement results to refine measurement strategy.

The provisioning of the P2P overlay should be transparent for the network administrator. An Autonomic Control Plane such as defined in [[I-D.anima-autonomic-control-plane](#)] provides an ideal candidate for the P2P overlay's underlay.

An autonomic solution for the distributed detection of SLA violations provide several benefits. First, 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. Secondly, placing logic where to conduct measurements in the node enables rapid control loops in which devices are able 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 on 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 faster the network dynamics than an human administrator and even a central controller. Finally, the solution could help to reduce the workload



of human administrator, or, at least, to avoid their need to perform operational tasks.

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

## **6.    Intended User and Administrator Experience**

The autonomic solution should not require human intervention in the distributed detection of SLA violations. This also enables SLA monitoring of a network by less experienced human administrators. However, some information may be provided from the human administrator. For example, the human administrator may set a policy regarding the resource budget that is assigned to network devices for measurement operations, or set a target for the number or percentage of SLO violations that must be detected allowing the solution to minimize the resources required to achieve that target.

## **7.    Analysis of Parameters and Information Involved**

The active measurement model assumes that a typical infrastructure will have multiple network segments and Autonomous Systems (ASs), and a reasonably large number of several of routers and hosts. It also considers that multiple SLOs can be in place at a given time. Since interoperability in a heterogenous network is a goal, features found on different active measurement mechanisms (e.g. OWAMP, TWAMP, and IPSLA) 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.

### **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 in 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 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. The definition of 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 endpoints neighbors (e.g., FIB and RIB tables) as the initial seed to get possible peers.

## **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's actions. There some proposals for passive probe activation (e.g., DECON and CSAMP), but without the focus on autonomic features. It is also mentioning a proposal from Barford et al. to detect and localize links which cause anomalies along a network path.

## **9.   Related IETF Work**

The following paragraphs discuss related IETF work and are 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 aims at the standards for performance management. Since their mechanisms also consist in deploying measurement probes the autonomic solution could be relevant for LMAP specially considering SLA violation screening. 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) aims at the process of standardization of 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 be possibly extended to address also passive measurement probes. Besides that, flow information could be used in the decision making of probe activation. 3. [ALT0]: 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 for the definition of the topology regarding the network devices which exchange measurement data.

## **10. Acknowledgements**

We wish to acknowledge the helpful contributions, comments, and suggestions that were received from Mohamed Boucadair, Bruno Klauser, Eric Voit, and Hanlin Fang.

## **11. IANA Considerations**

This memo includes no request to IANA.

## **12. Security Considerations**

The bootstrapping of a new device follows the approach proposed on anima wg [[draft-anima-boot](#)], thus in order to exchange data a device should register first. This registration could be performed by a "Registrar" device or a cloud service provided by the organization to facilitate autonomic mechanisms. The new device sends its own credentials to the Registrar, and after successful authentication, receives domain information, to enable subsequent enrollment to the domain. The Registrar sends all required information: a device name, domain name, plus some parameters for the operation. Measurement data should be exchanged signed and encrypted among devices since these data could carry sensible information about network infrastructures. Some attacks should be considering when analyzing the security of the autonomic solution. Denial of service (DoS) attacks could be performed if the solution be tempered to active more local probe than the available resources allow. Besides that, results could be forged by a device (attacker) in order to this device be considered peer of a specific device (target). This could be done to gain information about a network.

## **13. Informative References**

[[draft-anima-boot](#)]

Pritikin, M., Richardson, M., Behringer, M., and S. Bjarnason, "[draft-ietf-anima-bootstrapping-keyinfra](#)", [draft-ietf-anima-bootstrapping-keyinfra-04](#) (work in progress), January 2017.



[[draft-ietf-ippm-6man-pdm-option](#)]

Elkins, N., Hamilton, R., and M. Ackermann, "[draft-ietf-ippm-6man-pdm-option](#)", [draft-ietf-ippm-6man-pdm-option-08](#) (work in progress), February 2017.

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

Behringer, M., Eckert, T., and S. Bjarnason, "An Autonomic Control Plane", [draft-ietf-anima-autonomic-control-plane-05](#) (work in progress), January 2017.

[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)", 2012, <[http://ieeexplore.ieee.org/xpls/abs\\_all.jsp?arnumber=6379997](http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6379997)>.

[RFC3954] Claise, B., Ed., "Cisco Systems NetFlow Services Export Version 9", [RFC 3954](#), DOI 10.17487/RFC3954, October 2004, <<http://www.rfc-editor.org/info/rfc3954>>.

[RFC4148] Stephan, E., "IP Performance Metrics (IPPM) Metrics Registry", [BCP 108](#), [RFC 4148](#), DOI 10.17487/RFC4148, August 2005, <<http://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, <<http://www.rfc-editor.org/info/rfc4656>>.

[RFC5357] Hedayat, K., Krzanowski, R., Morton, A., Yum, K., and J. Babiarz, "A Two-Way Active Measurement Protocol (TWAMP)", [RFC 5357](#), DOI 10.17487/RFC5357, October 2008, <<http://www.rfc-editor.org/info/rfc5357>>.

[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, <<http://www.rfc-editor.org/info/rfc6812>>.

[RFC7297] Boucadair, M., Jacquenet, C., and N. Wang, "IP Connectivity Provisioning Profile (CPP)", [RFC 7297](#), DOI 10.17487/RFC7297, July 2014, <<http://www.rfc-editor.org/info/rfc7297>>.



Authors' Addresses

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

Email: jcnobre@unisinos.br

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

Email: granville@inf.ufrgs.br

Alexander Clemm  
Huawei  
Santa Clara, California  
USA

Email: ludwig@clemm.org

Alberto Gonzalez Prieto  
Cisco Systems  
San Jose  
USA

Email: albertgo@cisco.com