

Network Management Research Group
Internet-Draft
Intended status: Informational
Expires: August 2, 2017

J. Nobre
University of Vale do Rio dos Sinos
L. Granville
Federal University of Rio Grande do Sul
A. Clemm
Huawei
A. Prieto
Cisco Systems
January 29, 2017

Autonomic Networking Use Case for Distributed Detection of SLA Violations

draft-irtf-nmrg-autonomic-sla-violation-detection-05

Abstract

This document describes a use case for autonomic networking in distributed detection of Service Level Agreement (SLA) violations. It is one of a series of use cases intended to illustrate requirements for autonomic networking.

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 August 2, 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

1.	Introduction	2
2.	Definitions and Acronyms	4
3.	Current Approaches	4
4.	Problem Statement	5
5.	Benefits of an Autonomic Solution	5
6.	Intended User and Administrator Experience	6
7.	Analysis of Parameters and Information Involved	6
7.1.	Device Based Self-Knowledge and Decisions	6
7.2.	Interaction with other devices	7
8.	Comparison with current solutions	7
9.	Related IETF Work	7
10.	Acknowledgements	8
11.	IANA Considerations	8
12.	Security Considerations	8
13.	References	8
13.1.	Normative References	8
13.2.	Informative References	9
	Authors' Addresses	9

[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. Those requirements lead to the articulation of Service Level Objectives (SLOs) which are to be met. Those SLOs become part of Service Level Agreements (SLAs) that articulate a contract between the provider and the consumer of a service. To fulfill a service, it needs to be ensured that the SLOs are met. Examples of service fulfillment clauses can be found on [[RFC7297](#)]). Violations of SLOs can be associated with significant financial loss, which can be divided in two types. First, there is the loss incurred by the service users (e.g., the trader whose orders are not executed in a timely manner) and the loss incurred by the service provider in terms of penalties for not meeting the service and loss of revenues due to reduced customer satisfaction. Thus, the service level requirements of critical network services have become a

key concern for network administrators. To ensure that SLAs are not being violated, service levels need to be constantly monitored at the network infrastructure layer. To that end, network measurements must take place.

Network measurement mechanisms are performed through either active or passive measurement techniques. In passive measurements, production traffic is observed. Network conditions are checked in a non intrusive way because no monitoring traffic is created by the measurement process itself. 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 measurement, on the other hand, is intrusive because it injects synthetic traffic into the network to measure the network performance. 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](#)]. Besides that, 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. It also does not require inspecting production traffic. Because of this, it usually offers 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.

The activation of active measurement sessions (hosted in senders and responders considering the architecture described by Cisco [[RFC6812](#)]) is expensive in terms of the resource consumption, e.g., CPU cycle and memory footprint, and monitoring functions compete for resources with other functions, including routing and switching. Besides that, the activated sessions also increase the network load because of the injected traffic. The resources required and traffic generated by the active measurement sessions are a function of the number of

measured network destinations, i.e., with more destinations the larger will be the resources and the traffic needed to deploy the sessions. Thus, to have a better monitoring coverage it is necessary to deploy more sessions what consequently turns increases consumed resources. Otherwise, enabling the observation of just a small subset of all network flows can lead to an insufficient coverage. Hence, the decision how to place measurement probes becomes an important management activity, so that with a limited amount of measurement overhead the maximum benefits in terms of service level monitoring are obtained.

2. Definitions and Acronyms

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

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

SLA: Service Level Parameter

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 compute 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. Problem Statement

The problem to solve involves automating the placement of active measurement probes in the most effective manner possible. Specifically, assuming a bounded resource budget that is available for measurements, the problem becomes how to place those measurement probes such that the likelihood of detecting service level violations is maximized, and subsequently performing the required configurations. The method should be embeddable as management software inside network devices that controls the deployment of active measurement mechanisms. The method shall furthermore be dynamic and be able to adapt to changing network conditions.

5. Benefits of an Autonomic Solution

The use case considered here is the distributed autonomic detection of SLA violations. The use of Autonomic Networking (AN) properties 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, while allowing for their input.

We advocate for embedding Peer-to-Peer (P2P) technology in network devices in order to improve the measurement session activation decisions using autonomic control loops. The provisioning of the P2P management overlay should be transparent for the network administrator. It would be possible to control the measurement session activation using local data and logic and to share measurement results among different network devices.

An autonomic solution for the distributed detection of SLA violations can provide several benefits. First, efficiency: this solution could optimize the resource consumption and avoid resource starvation on the network devices. In practice, the solution should maximize the benefits of SLA monitoring (i.e., maximize the likelihood of SLA violations being detected) by operating within a given resource budget. This optimization comes from different sources: taking into account past measurement results, taking into account other observations (such as, observations of link utilizations and passive measurements, where available) sharing of measurement results between network devices, better efficiency in the probe activation decisions, etc. Second, effectiveness: the number of detected SLA violations

could be increased. This increase is related with a better coverage of the network. Third, 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. Finally, the solution could help to reduce the workload of human administrator, or, at least, to avoid their need to perform operational tasks.

6. Intended User and Administrator Experience

The autonomic solution should not require the human intervention in the distributed detection of SLA violations. Besides that, it could enable the control of SLA monitoring by less experienced human administrators. However, some information may be provided from the human administrator. For example, the human administrator may provide the SLOs regarding the SLA being monitored. The configuration and bootstrapping of network devices using the autonomic solution should be minimal for the human administrator. Probably it would be necessary just to inform the address of a device which is already using the solution and the devices themselves could exchange configuration data.

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 in 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 programability interfaces (e.g., Cisco's EEM and onePK) 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. In any case, it is necessary to assure that the results from remote devices have local relevancy. 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 standartized 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. [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 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 enrolment 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. References

13.1. Normative References

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

Pritikin, M., Richardson, M., Behringer, M., and S. Bjarnason, "[draft-ietf-anima-bootstrapping-keyinfra](#)", [draft-ietf-anima-bootstrapping-keyinfra-03](#) (work in progress), June 2016.

[[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-06](#) (work in progress), September 2016.

[P2PBNM-Nobre-2012]

Nobre, J., Granville, L., Clemm, A., and A. 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)>.

[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>>.

13.2. Informative References

[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>>.

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
Santa Clara, California
USA

Email: ludwig@clemm.org

Alberto Gonzalez Prieto
Cisco Systems
San Jose
USA

Email: albertgo@cisco.com

