



HAL
open science

QoS Fault Detection and Localization Mechanisms (FDLM) in Multi-domain Networks Adapted to Export Methods

Aymen Belghith, Siwar Ben Adj Said, Bernard Cousin, Samer Lahoud

► **To cite this version:**

Aymen Belghith, Siwar Ben Adj Said, Bernard Cousin, Samer Lahoud. QoS Fault Detection and Localization Mechanisms (FDLM) in Multi-domain Networks Adapted to Export Methods. 4th International Conference on Advanced Communication Technology 2012 (ICACT 2012), Feb 2012, Pyeongchang, South Korea. pp.848 - 853. hal-01183602

HAL Id: hal-01183602

<https://hal.science/hal-01183602>

Submitted on 10 Aug 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

QoS Fault Detection and Localization Mechanisms (FDLM) in Multi-domain Networks Adapted to Export Methods

Aymen Belghith, Siwar Ben Hadj Said, Bernard Cousin, and Samer Lahoud

Université de Rennes 1 - IRISA,

Campus de Beaulieu, 35042 Rennes Cedex, France

Email: aymen.belghith@univ-rennes1.fr, siwar.benhadjsaid@telecom-bretagne.eu, bernard.cousin@irisa.fr, samer.lahoud@irisa.fr

Abstract—Monitoring the quality of service in a multi-domain network allows providers to ensure the control of multi-domain service performance. A multi-domain service is a service that crosses multiple domains. In this paper, we propose several mechanisms for fault detection and fault localization. A fault is detected when an end-to-end contract is not respected. Faulty domains are domains that do not fulfill their Quality of Service (QoS) requirements. Our three proposed fault detection and localization mechanisms (FDLM) depend on the export method used. These export methods define how the measurement results are exported for analysis. We consider the periodic export, the triggered export, and a combined method. For each FDLM, we propose two sub-schemes that use different fault detection strategies. In this paper, we describe these mechanisms and evaluate their performance using Network Simulator (NS-2).

I. INTRODUCTION

Network monitoring is necessary to guarantee precise and efficient management of a network communication system. It is required to control the Quality of Service (QoS) provided by the network. The performance requirements of the services are typically specified through a contract called Service Level Agreement (SLA). In order to guarantee the performance of the services, the network performance has to be verified by performing network monitoring. The Internet is composed of several autonomously managed routing domains. Generally, all equipments and data traffic in a domain are under the sole responsibility of the domain authority. Many monitoring architectures have been proposed for mono-domain networks such as AQUILA [1] and contractual SLA [2] monitoring architectures or proposed for multi-domain networks such as INTERMON [3], ENTHRONE [4], and EuQoS [5] monitoring architectures. A monitoring architecture can use standard monitoring protocols such as Real-time Traffic Flow Measurement (RTFM) [6], IP Flow Information eXport (IPFIX) [7], and Packet Sampling (PSAMP) [8]. In this paper, we want to monitor data traffic which crosses several domains and then we are interested in multi-domain monitoring.

The heterogeneity, interoperability, and confidentiality aspects of the different domains makes the multi-domain network monitoring a challenging problem. However, we note that all the above monitoring architectures do not take into account all these multi-domain network aspects. Therefore, we

have proposed in [9] a configurable monitoring architecture that solves heterogeneity, interoperability and confidentiality problems.

Once the monitoring is established, QoS measurements have to be performed and then the measurement results have to be exported for analysis purposes in order to detect the faulty domains. Faulty domains are domains that do not fulfill their per-domain contract. However, the existing multi-domain network architectures do not specify how the faults can be detected and localized.

We indicate that, in intra-domain networks, some works have already been done on anomaly diagnosis and fault detection such as in [10] and [11]. In [10], authors propose a path monitoring framework for detecting path anomalies concerning the delay without monitoring all the paths as they suppose that the faults are rare. Therefore, a path selection heuristic was proposed in order to select the paths to monitor. In [11], a path selection algorithm is proposed in order to fairly select the paths that are probed. Once the paths are probed, the measurement results are compared with thresholds that are specified by the user in order to label these paths with one of these three states: "functioning normally", "nearly in anomalous state", and "anomalous state".

In this paper, we propose three fault detection and localization mechanisms (FDLM) in a multi-domain networks. These mechanisms are respectively adapted to periodic export, triggered export, and combined export methods. Indeed, an FDLM has to take into account the characteristics of the export method used in order to efficiently analyze and then detect faulty domains. A fault detection and localization mechanism provides the main functionality of the network monitoring as it checks if the multi-domain service requirements are provided and it determines, when these requirements are not respected, the faulty domain.

Our proposed mechanisms should adapt to any compatible multi-domain network architecture like the architecture model defined by the IPSphere forum [12]. This model allows providers to overcome scalability and interoperability issues. The IPSphere forum has defined the role of each system entity: Administrative Owner (AO), Element Owner (EO), and customer. The AO is the entity that is responsible for providing

end-to-end services over a multi-domain network. We add monitoring functionality to the AO in order to guarantee end-to-end services. The EO is the entity that manages the resources of a network domain. Each service provided by the AO uses the resources of one or several EOs.

The exported measurement results allows the AO to detect and localize the faulty domains. The measurement results analysis depend on the kind of the metric to measure. There are two possible kinds of metrics: those that need aggregation and those that do not need aggregation. The aggregation of the measurement results represents the computation of the end-to-end results using the per-domain results exported by the different domains participating in the monitoring of the multi-domain service. In this paper, we consider the following metrics: the One-Way Delay (OWD) [13] that is a metric that needs aggregation and the throughput [14] that does not need aggregation.

This paper is organized as follows. We present our fault detection and localization mechanisms in section II. Section III presents the performance criteria and performance evaluations and comparisons of our proposed mechanisms. The evaluation of our proposed mechanisms, through extensive simulations, consists of studying the detection delay, export throughput, and detection efficiency. Conclusions are provided in section IV.

II. MECHANISMS OF FAULT DETECTION AND LOCALIZATION

In this section, we describe our three proposed mechanisms of fault detection and localization. These mechanisms depend on the export method used. The export methods used are periodic export, triggered export and combined export methods.

When the periodic export method is used, each domain periodically exports the measurement results. When the triggered export method is used, each domain exports the measurement results immediately at the violation of the per-domain contract. When the combined export method is used, each domain exports the measurement results periodically as well as immediately at the violation of the per-domain contract.

In this paper, a fault is detected when an end-to-end contract is not respected. At the establishment of an end-to-end contract between the AO and the client, the AO has to negotiate per-domain contracts with the EOs that can participate in the monitoring of the multi-domain service. Contract negotiation mechanisms will be studied in future work.

A. Fault detection and localization mechanisms based on periodic export method (FDLM-P)

We propose two sub-schemes for our proposed fault detection and localization mechanism based on periodic export (FDLM-P). In the first sub-scheme, called FDLM-P-strict, the number of the measurement results to export periodically is constant and does not depend on the number of the generated faults. In the second sub-scheme, called FDLM-P-adjustable, the number of the measurement results to export periodically is variable and is equal to the number of the faults detected locally during the export period.

1) *FDLM-P-strict*: Let N_{exp} be the maximum number of the measurement results that are exported periodically when FDLM-P-strict is used. N_{exp} is constant and does not depend on the number of the generated faults. Let N_{meas} be the number of available measurement results for an export period. N_{exp} has to be equal to N_{meas} if a domain decides to export all measurement results. N_{exp} is equal to the export period, called P_{exp} , divided by the measurement period, called P_{meas} ($N_{meas} = \frac{P_{exp}}{P_{meas}}$).

Therefore, the number of the measurement results that are exported periodically, called N_{strict} is equal to:

$$N_{strict} = \min(N_{exp}, N_{meas}) \quad (1)$$

When FDLM-P-strict is used, each domain periodically exports to the AO a number of measurement results equal to N_{strict} (see (1)). When the metric does not need aggregation, the AO easily detects and localizes the faulty domain as a fault in a domain implies that the end-to-end contract is not respected. When the metric needs aggregation, the AO has to aggregate the measurement results in order to check if the end-to-end contract is respected. Obviously, the aggregation is exact only when the domains export all the measurement results ($N_{exp} = N_{meas}$). When N_{exp} is lower than N_{meas} , there is at least a missing domain measurement result. Let $R(i, t)$ be a measurement result obtained at time t and exported by domain i . We suppose that domain I does not export its measurement result that was obtained at time T . Therefore, $R(I, T)$ has to be estimated using the other measurement results already exported by this domain. These measurement results belong to the last export period, called P_{last} , i.e. the estimation uses values among $R(I, t)$ where $t \in P_{last}$. We propose four variants for estimating the missing measurement results: FDLM-P-strict-max, FDLM-P-strict-min, FDLM-P-strict-avg, and FDLM-P-strict-avg-mm. We propose that the same variant is used by all domains in order to coherently estimate the end-to-end measurement result.

a) *FDLM-P-strict-max*: Each domain exports the N_{exp} greatest values. When the AO does not receive all measurement results, at time T for example, it estimates missing measurement result(s) by taking the minimum of the exported measurement results. For example, if domain I does not export $R(I, T)$ because this measurement result value is lower than the N_{exp} greatest values, the AO supposes that:

$$R(I, T) = \min_{t \in P_{last}} (R(I, t)) \quad (2)$$

We note that, when FDLM-P-strict-max is used, the estimated end-to-end measurement result value is lower than or equal to the real value.

b) *FDLM-P-strict-min*: Each domain exports the N_{exp} lowest values. When the AO does not receive all measurement results, it estimates the missing measurement result(s) by taking the maximum value of the exported measurement results. For example, if domain I does not export $R(I, T)$ because this measurement result value is greater than the N_{exp} lowest values, the AO supposes that:

$$R(I, T) = \max_{t \in P_{I,ast}} (R(I, t)) \quad (3)$$

We note that, when FDLM-P-strict-min is used, the estimated end-to-end measurement result value is greater than or equal to the real value.

c) *FDLM-P-strict-avg*: Each domain exports N_{exp} measurement results randomly. When the AO does not receive all measurement results, it estimates missing measurement result(s) by taking the average of the exported measurement results. For example, if domain I does not export $R(I, T)$, the AO supposes that:

$$R(I, T) = \sum_{t \in P_{I,ast}} \frac{R(I, t)}{N_{strict}} \quad (4)$$

As the measurement results are exported randomly, the estimation method cannot be accurate. So, in order to improve the estimation of the missing values, we propose variant FDLM-P-strict-avg-mm.

d) *FDLM-P-strict-avg-mm*: Each domain exports N_{exp} measurement results that belong to two sets V_{max} and V_{min} . V_{max} and V_{min} contain the $\frac{N_{exp}}{2}$ greatest and the $(\frac{N_{exp}}{2} + N_{exp} \bmod 2)$ lowest values, respectively. When the AO does not receive all measurement results, it estimates missing measurement result(s) by taking the average between the maximum value of V_{min} and the minimum value of V_{max} . For example, if domain I does not export $R(I, T)$, the AO supposes that:

$$R(I, T) = \frac{\min(V_{max}) + \max(V_{min})}{2} \quad (5)$$

2) *FDLM-P-adjustable*: Recall that, when FDLM-P-adjustable is used, the number of the measurement results to export periodically is variable and is equal to the number of the faults detected locally during the export period. We propose that this sub-scheme is especially used to detect faults related to metrics that do not need aggregation. In fact, a fault committed by a domain implies that the end-to-end contract is not respected and therefore all faults detected in a domain have to be exported. Let $N_{adjustable}$ be the number of exported measurement results and N_{fault} be the number of detected faults. $N_{adjustable}$ is equal to N_{fault} for each export period ($N_{adjustable} = N_{fault}$).

We note that FDLM-P-adjustable can be used to detect and localize faults related to metrics that need aggregation. In this case, the AO automatically considers that the domain that exports a measurement result as a faulty domain as this exported measurement result does not respect the per-domain contract.

B. Fault detection and localization mechanisms based on triggered export method (FDLM-T)

We propose two sub-schemes for our proposed fault detection and localization mechanism based on triggered export (FDLM-T). Recall that when the triggered export method is used, each domain exports the measurement results that do not fulfill its per-domain contract. In the first sub-scheme,

called FDLM-T-unsolicited, the AO settles for already exported measurement results and does not aggregate them. In the second sub-scheme, called FDLM-T-on-demand, the AO requests the EOs, that fulfill their per-domain contracts, to send some additional measurement results. For example, we suppose that a service crosses three domains A, B, and C. At time T , only domain A exports measurement results. In order to have exact aggregation, the AO requests domain B and domain C to export the measurement results already obtained at this violation time (T).

1) *FDLM-T-unsolicited*: We propose that sub-scheme FDLM-T-unsolicited is used to detect faults related to metrics that do not need aggregation. In this case, FDLM-T-unsolicited is exact as the domains export the measurement results that do not respect their contracts per domain. Moreover, for this kind of metrics, a fault at a domain implies that the end-to-end contract is violated. We note that FDLM-T-unsolicited can also be used to detect faults related to metrics that need aggregation. In this case, the AO is not tolerant and automatically considers that the domain that exports measurement results as a faulty domain.

2) *FDLM-T-on-demand*: We propose that sub-scheme FDLM-T-on-demand is used to detect faults related to metrics that need aggregation. In this case, the AO verifies if the end-to-end contract is respected. Indeed, the AO can be tolerant as a fault at a domain does not imply that the end-to-end contract is not respected.

C. Fault detection and localization mechanisms based on triggered export combined with periodic export (FDLM-TP)

We propose two sub-schemes for our proposed FDLM based on triggered export combined with periodic export (FDLM-TP).

In the first sub-scheme, called FDLM-TP-without-verification, the AO settles for already exported measurement results and does not aggregate them. However, the AO estimates the end-to-end measurement results using the values that were exported periodically, specifically in the previous export period. In order to improve the export throughput of FDLM-TP, we propose that the domains periodically export a single value that represents the average of all obtained measurement results. In fact, it is useless to send all obtained measurement results as the estimation uses values that do not belong to the current export period.

In the second sub-scheme, called FDLM-TP-with-verification, the AO requests the EOs, that fulfill their per-domain contracts, to send some additional measurement results for exact aggregation like for FDLM-T-on-demand.

D. FDLM pre-selection

We notice that it is useless to use FDLM-P-strict and FDLM-T-on-demand for detecting faults concerning the metrics that need no aggregation at the AO. Consequently, for this kind of metrics, we propose that only FDLM-P-adjustable, FDLM-T-unsolicited, or FDLM-TP-without-verification are used.

For metrics that need aggregation, we propose that the AO is tolerant, i.e. it does not automatically consider that a per-contract fault implies an end-to-end fault. However, the AO checks if aggregated measurements fulfill the end-to-end contract. Consequently, only FDLM-P-strict and FDLM-T-on-demand are used for this kind of metrics.

We note that FDLM-TP-with-verification is exact without needing measurement results exported periodically. Therefore, this sub-scheme will have the same performance of FDLM-T-on-demand but with additional throughput. Therefore, it is useless to study the performance of FDLM-TP-with-verification.

III. PERFORMANCE EVALUATION OF THE PROPOSED MECHANISMS OF FAULT DETECTION AND LOCALIZATION

A. Performance criteria

We evaluate the following performance criteria:

- The fault detection delay: represents the difference between the time of detection of a fault and the time of the occurrence of this fault. As the AO collects the measurement results of each domain, it can immediately localize the faulty domain when receiving measurement results. Therefore, for our detection and localization schemes, the fault localization delay, that represents the difference between the time of localization of a fault and the time of the occurrence of this fault, is equal to the fault detection delay.
- The export throughput: represents the throughput of messages used to export the measurement results. Two kinds of export messages are used. The first kind of export message includes the results of the delay and the throughput measurements. The second kind of export messages contains a request of a measurement result obtained at a specified instant. The message is used only by sub-scheme FDLM-T-on-demand. We note that this message is used by FDLM-TP-with-demand. However, as it is mentioned in the previous section, this sub-scheme will provide worse performance results than FDLM-T-on-demand and therefore this sub-scheme will not be evaluated.
- The detection efficiency: consists of two criteria: the ratio of the detection of a real fault, called fault detection ratio, and the ratio of the detection of a fault while the end-to-end contract is well respected, called false alarms ratio.

B. Simulation model

1) *Simulation scenario*: In this section, we consider a multi-domain network topology formed by four domains and fourteen measurement points (see Fig. 1). Each domain may contain numerous measurement points but we consider only measurement points that are located at the border of the domains for confidentiality reasons [9]. Domain A, domain B, domain C, and domain D contains three measurement points (a_1 , a_2 , and a_3), four measurement points (b_1 , b_2 , b_3 , and b_4), four measurement points (c_1 , c_2 , c_3 , and c_4), and three measurement points (d_1 , d_2 , and d_3), respectively. The main simulation parameters are presented in Table I. We evaluate




Fig. 1. Multi-domain network monitoring scenario.

TABLE I
SIMULATION PARAMETERS.

Simulation parameters	Values
Number of domains	4
Number of measurement points	14
Simulation time	1500 s
Measurement period (P_{meas})	0.2 s
Export period (P_{exp})	1 s, 2 s, and 4 s
Global delay threshold of end-to-end contracts (T_{global_delay})	0.045 s
Global throughput threshold of end-to-end contracts ($T_{global_throughput}$)	2 Gbit/s

our proposed mechanisms using Network Simulator (NS-2). We note that we have added to this simulator classes and methods in order to implement the functionalities of the AO, the EOs, and the measurement points (MPs).

Recall that N_{meas} indicates the number of available measurement results for an export period when FDLM-P-strict is used. In our simulation, N_{meas} is equal to 5, 10 and 20. For example, when we decide to decrease the measurement period in order to have more accurate measurements (for example P_{meas} becomes equal to 0.1 s), we can choose export periods equal to 0.5 s, 1 s and 2 s in order to provide the same simulation results described in this paper (as N_{meas} values are still equal to 5, 10 and 20).

In our performance study, we consider the delay and the throughput as metrics that need/or do not need aggregation at the AO, respectively.

2) *Delay and throughput thresholds*: We propose that all end-to-end contracts require a global delay lower than T_{global_delay} and a global throughput greater than $T_{global_throughput}$. Evidently, each domain must provide a throughput greater than $T_{global_throughput}$ (so $T_{delay_per_domain} = T_{global_throughput}$).

When negotiating contracts with EOs, we propose that the global delay is fairly distributed between domains involved in the multi-domain service monitoring. The delay threshold of a domain is computed as follows:

$$T_{delay_per_domain} = \frac{T_{global_delay}}{N_{Domain}} \quad (6)$$

where N_{Domain} represents the number of the domains that the monitoring service crosses.

TABLE II
AVERAGE EXPORT THROUGHPUT AND THE AVERAGE DELAY TO DETECT FAULTS WHEN FDLM-P-ADJUSTABLE IS USED.

Measure and export periods	Average export throughput (bit/s)	Average delay to detect faults (s)
$(P_{meas} = 0.2, P_{exp} = 1)$	120.68	0.438
$(P_{meas} = 0.2, P_{exp} = 2)$	93.81	0.938
$(P_{meas} = 0.2, P_{exp} = 4)$	80.28	1.954

TABLE III
DETECTION EFFICIENCY (CONCERNING THROUGHPUT AND DELAY) WHEN FDLM-P-ADJUSTABLE IS USED.

Efficiency of FDLM-P-adjustable	when metrics do not need aggregation
Detected faults ratio	1
False alarms ratio	0

3) *Fault generation model*: In our scenario, we study the performance of different fault detection and localization mechanisms for only one service and for a measurement result generation model called M1. This model generates faults with a probability equal to $3/23$ (called p_1) either for throughput or for delay. For that, we assume, for example, that the delay per domain is uniformly distributed in $[T_{delay_per_domain} * 1/3, (T_{delay_per_domain}) * 11/10]$ and that the throughput per domain is uniformly distributed in $[T_{throughput_per_domain} * 9/10, T_{throughput_per_domain} * 5/3]$. In this paper, we suppose that measurement results (and therefore the faults) are generated independently.

C. Evaluation of the FDLM-P scheme

Recall that when scheme FDLM-P is used, sub-schemes FDLM-P-adjustable and FDLM-P-strict are used to detect faults concerning the throughput and the delay, respectively.

1) Performance evaluation of FDLM-P-adjustable:

a) *Export throughput and detection delay evaluation for FDLM-P-adjustable*: The average export throughput and the average delay to detect faults concerning the throughput are presented in Table II.

We can verify that when P_{exp} increases, the export throughput decreases. Indeed, when P_{exp} increases, the number of exported messages and then the total size of the headers decreases. Recall that when FDLM-P-adjustable is used, the total number of measurement results to be exported is independent of P_{exp} as it depends only on the number of the generated faults. Therefore, there is less export message headers to send.

We also verify that the average delay to detect faults increases when P_{exp} increases because each domain has to wait the expiration of the export period to send its measurement results to the AO for analysis.

b) *Detection efficiency evaluation for FDLM-P-adjustable*: Table III represents the detection efficiency results of FDLM-P-adjustable.

We verify that sub-scheme FDLM-P-adjustable is exact when it is applied to detect faults related to metrics that do not need aggregation. Indeed, this sub-scheme exports all the faults detected locally to the AO. Moreover, with this kind of




Fig. 2. Average export throughput vs N_{exp} when FDLM-P-strict and FDLM-P-adjustable are used.

metrics, a fault detected in a domain induces an end-to-end fault.

2) *Performance evaluation of FDLM-P-strict*: In this section, FDLM-P-strict is used to detect and locate faults concerning the delay.

a) Throughput and delay evaluation for FDLM-P-strict:

Fig. 2 represents the average export throughput as a function of N_{exp} . Recall that N_{exp} represents the maximum number of measurement results to be exported during an export period. The export throughput is independent of the variant of FDLM-P-strict as it depends only on N_{exp} .

We verify that the export throughput increases when N_{exp} increases without exceeding N_{meas} . When N_{exp} is greater than N_{meas} , the export throughput remains constant. Recall that N_{meas} represents the total number of measurement results obtained during an export period. We also verify that the export throughput increases when P_{exp} decreases. Indeed, when P_{exp} increases, the total number of the messages exported decreases and therefore the number of export message headers decreases.

Now we compare the two sub-schemes of FDLM-P. We note that the export throughput generated by FDLM-P-strict is greater than that generated by FDLM-P-adjustable for most of the values of N_{exp} (see Fig. 2). The export throughput depends on the number of generated faults (for FDLM-P-adjustable) and N_{exp} (for FDLM-P-strict). For example, for an export period equal to 4 s (resp. equal to 2 s), the export throughput generated by FDLM-strict-P is greater than that generated by FDLM-P-adjustable when N_{exp} is greater than 2 (resp. greater than 1).


b) *Detection delay and export throughput evaluation for FDLM-P-strict*: The average delay to detect faults is presented in Table IV. The detection delay is independent of the sub-scheme variant.

We note that the average delay to detect faults, when FDLM-P-strict is used, is equal to that when FDLM-P-adjustable is

TABLE IV

AVERAGE DELAY TO DETECT FAULTS (CONCERNING THE DELAY AND THE THROUGHPUT) WHEN FDLM-P-STRICT IS USED.

Measure and export periods	Average delay to detect faults (s)
$(P_{meas} = 0.2, P_{exp} = 1)$	0.438
$(P_{meas} = 0.2, P_{exp} = 2)$	0.938
$(P_{meas} = 0.2, P_{exp} = 4)$	1.954


Fig. 3. Detected faults ratio vs N_{exp} when FDLM-P-strict is used.

used. This is due to the fault detection delay depends only on the export period. Note that in both sub-schemes of FDLM-P, the AO must wait the reception of the export messages from all domains that participate in the multi-domain monitoring before determining the faulty domain.

c) *Detected faults ratio evaluation for FDLM-P-strict:*

Fig. 3 represents the detected faults ratio as a function of N_{exp} . We note that variant FDLM-P-strict-max allows the AO to detect all the generated faults when N_{exp} is greater than the number of generated faults in a given export period. This value depends on the fault generation model. Evidently, when probability $p1$ increases, this value increases. When variant FDLM-P-strict-max is used, the AO can detect all the faults when at least one domain exports a measurement result immediately after the fault generated instant. For example, when P_{exp} is equal to 4 s (resp. 2 s), FDLM-P-strict-max can detect all the generated faults when N_{exp} is greater than or equal to 3 (resp. greater than or equal to 2).

We observe that variant FDLM-P-strict-min starts detecting detecting some faults only when N_{exp} is very close to the total number of measurement results. For example, for values of N_{meas} equal to 5, 10 and 20, the AO starts detecting faults when N_{exp} exceeds 4, 8, and 17, respectively. Therefore, the different domains have to export about 80% of all their measurement results obtained in each export period in order to enable the AO to detect some faults. In fact, when FDLM-P-strict-min is used, each domain sends the minimum values of its measurement results. So, the AO starts to detect faults from a certain value of N_{exp} where domains, at least, start

Fig. 4. False alarms ratio vs N_{exp} when FDLM-P-strict is used.

sending values that violate the per-domain contract. Recall that an end-to-end contract is violated if and only if at least one per-domain contract has been violated.

Variant FDLM-P-strict-avg improves the detected faults ratio compared to FDLM-P-strict-min. For instance, for values of N_{meas} equal to 5, 10, and 20, the AO starts detecting faults when N_{exp} exceeds 2, 5, and 10, respectively. In fact, with this variant, the AO aggregates average values and therefore aggregated values are closer to real values compared with FDLM-P-strict-min.

Finally, we note that FDLM-P-strict-avg-mm outperforms FDLM-P-strict-avg. For example, for values of N_{meas} equal to 5, 10, and 20, the AO starts detecting faults when the N_{exp} exceeds 2, 2, and 4, respectively. Indeed, when the variant FDLM-P-strict-avg-mm is used, the AO aggregates values that are closer to the real values than that obtained by variant FDLM-P-strict-avg.

d) *False alarms ratio evaluation for FDLM-P-strict:*

Fig. 4 represents the false alarms ratio as a function of N_{exp} . We remark that variant FDLM-P-strict-max generates a huge number of false alarms. Indeed, when this variant is used, the AO aggregates received values which are greater than or equal to the real values. Therefore, the AO generates many false alarms. For example, when each domain sends only two measurement results per export period (i.e. $N_{exp} = 2$), FDLM-P-strict-max generates false alarms equal to 3.5, 17.5, and 26 times the total number of faults that are really generated for export periods equal to 1 s, 2 s, and 4 s, respectively.

We note that the greater P_{exp} is, the greater is the number of false alarms generated by FDLM-P-strict-max. In fact, when the export period is larger, the estimation is less accurate.

We remark that variant FDLM-P-strict-avg generates much less false alarms than FDLM-P-strict-max. Indeed, each estimated value with FDLM-P-strict-avg is lower or equal to the estimated value with FDLM-P-strict-max. Moreover, FDLM-P-strict-avg uses average values for aggregation and then




Fig. 5. False alarms ratio vs N_{exp} when FDLM-P-strict-min, FDLM-P-strict-avg, and FDLM-P-strict-avg-mm are used.

TABLE V
AVERAGE DELAY TO DETECT FAULTS AND AVERAGE EXPORT THROUGHPUT WHEN FDLM-T-UNSOLICITED IS USED.

FDLM sub-scheme	Average delay to detect faults (s)	Average export throughput (bit/s)
FDLM-T-unsolicited	0.03	133.30

estimated values are closer to the real values. In our scenario, when FDLM-P-strict-avg is used, the false alarms ratio is lower than 0.45 times the total number of faults that was really generated (see Fig. 5).

We note that variant FDLM-P-strict-avg-mm outperforms FDLM-P-strict-avg as its estimation method is more accurate. For example, in our scenario, the false alarms ratio is lower than 0.1 times the total number of faults that was really generated.

Finally, we notice that FDLM-P-strict-min does not generate false alarms whatever the values of N_{exp} . Indeed, when this variant is used, the estimated values are lower or equal to the real values. For this reason, when the aggregated value exceeds the threshold defined in the end-to-end contract, the AO is sure that this end-to-end contract was violated.

D. Evaluation of the FDLM-T scheme

In this section, we evaluate the performance of the two sub-schemes of FDLM-T: FDLM-T-on-demand and FDLM-T-unsolicited. Recall that when FDLM-T is used, sub-schemes FDLM-T-unsolicited and FDLM-T-on-demand are respectively used for the detection of faults concerning the throughput and the delay.

1) *Performance evaluation of FDLM-T-unsolicited:* The simulation results show that FDLM-T-unsolicited detects violations of the throughput within 0.03 s (see Table V). This good performance is explained by the fact that each EO exports the measurement results immediately when it violates its per-domain contract. Therefore the AO rapidly detects faults.

TABLE VI
DETECTION EFFICIENCY WHEN FDLM-T-UNSOLICITED IS USED.

Efficiency of FDLM-T-unsolicited	when metrics do not need aggregation
Detected faults ratio	1
False alarms ratio	0

TABLE VII
AVERAGE DELAY TO DETECT FAULTS AND AVERAGE EXPORT THROUGHPUT WHEN FDLM-T-ON-DEMAND IS USED.

FDLM sub-scheme	Average delay to detect faults (s)	Average export throughput (bit/s)
FDLM-T-on-demand	0.076	308.84

We notice that FDLM-T-unsolicited generates more export throughput than FDLM-P-adjustable (see Table V and Table II). In fact, when FDLM-P-adjustable is used, each EO exports a single message, that contains all the faults produced at an export period, periodically. However, when FDLM-T-unsolicited is used, EOs send export messages whenever fault happens and therefore there are more extra packet headers.

Table VI presents the detection efficiency of FDLM-T-unsolicited. We verify that this sub-scheme is exact when it is applied to detect faults concerning metrics that do not need aggregation at the AO.

2) Performance evaluation of FDLM-T-on-demand:

FDLM-T-on-demand presents an average delay to detect faults concerning delay equal to 0.076 s (see Table VII). This detection delay is low because domains export measurement results that violate the per-domain contract immediately. We verify that this delay is greater than that provided by FDLM-T-unsolicited. In fact, when FDLM-T-on-demand is used, the AO does not immediately consider the domain that export messages as a faulty domain. In order to check the end-to-end value, the AO has to request the EOs, that fulfill their per-domain contracts, to send additional measurement results at violation instants for exact aggregation. Therefore, the export throughput generated by FDLM-T-on-demand is greater than that generated by FDLM-T-unsolicited due to these extra export messages (see Table VII). Moreover, FDLM-T-on-demand is an exact scheme as the AO uses an exact aggregation method. We note that it is useless, for fault detection, to export measurement results when all the per-domain contracts are respected as an end-to-end fault means that at least there is a fault in a domain.

E. Evaluation of the FDLM-TP scheme

In this section, we study the performance of FDLM-TP-without-verification and FDLM-TP-with-verification. Recall that when FDLM-TP is used, the EOs export the measurement results concerning the throughput only using the triggered export method (for this kind of metric, this scheme is exact). However, the measurement results concerning the delay are exported using the periodic export method as well as the triggered export method.

TABLE VIII
AVERAGE EXPORT THROUGHPUT AND DETECTION EFFICIENCY WHEN
FDLM-TP-WITHOUT-VERIFICATION IS USED.

(Measure period, export period)	Export throughput	Detected faults ratio	False alarms ratio
($P_{meas} = 0.2, P_{exp} = 1$)	237.44 bit/s	0.26	0.15
($P_{meas} = 0.2, P_{exp} = 2$)	185.24 bit/s	0.18	0.014
($P_{meas} = 0.2, P_{exp} = 4$)	159.34 bit/s	0.17	0.014

1) *Performance evaluation of FDLM-TP-without-verification:* The simulation results show that FDLM-TP-without-verification detects violations of delay and throughput within 0.03 s. The detection delays provided by FDLM-TP-without-verification and by FDLM-T-unsolicited are the same because, in both sub-schemes, each EO exports the measurement results immediately at the per-domain contract violation instant and then the AO automatically considers that the domain that exports measurement results as a faulty domain.

The export throughput of FDLM-TP-without-verification is presented in Table VIII. We show that the export throughput of FDLM-TP-without-verification is greater than that of FDLM-P-strict and FDLM-T-unsolicited (see Table II and Table V, respectively).

Table VIII presents the detection efficiency of FDLM-TP-without-verification. We note that this sub-scheme does not detect all faults and generates false alarms. The false alarms ratio of FDLM-TP-without-verification is lower than that of FDLM-P-strict-max, FDLM-P-strict-avg, and FDLM-P-strict-min-avg when these variants do not export all measurement results (see Fig. 4). We note that the detected faults ratio of FDLM-TP-without-verification is greater than that of FDLM-P-strict-min when this variant exports a number of measurement results lower than $\frac{N_{meas}}{2}$ (see Fig. 3).

IV. CONCLUSION

In this paper, we have proposed fault detection and localization mechanisms that depend on export methods. We have shown, through extensive simulations, that FDLM-TP-with-verification, FDLM-T-on-demand, and FDLM-P-strict when N_{exp} is equal to N_{meas} , are exact as they can detect all faults without generating any false alarm. When N_{exp} is lower than N_{meas} , variants are used to estimate end-to-end measurement results. We have concluded that FDLM-P-strict-max needs lower number of exported results to detect all generated faults. However, this variant generates a great number of false alarms. FDLM-P-strict-min does not generate false alarms but needs a great number of exported results to detect all generated faults. FDLM-P-strict-avg and FDLM-P-strict-avg-mm can detect faults using lower N_{exp} value compared with FDLM-P-strict-min and generates lower number of false alarms compared with FDLM-P-strict-max.

We have also shown that the better detection delay is provided by FDLM-T-unsolicited and FDLM-T-without-verification because the AO does not wait for the expiration of the export period to detect faults like in FDLM-P-adjustable

and FDLM-P-strict neither for further informations from the EOs in order to exactly aggregate measurement results like in FDLM-T-on-demand and FDLM-TP-with-verification. We have observed that the export throughput of FDLM-P-strict and FDLM-TP increases when P_{exp} decreases while the throughput of FDLM-P-adjustable and FDLM-T depends only on the number of generated faults.

ACKNOWLEDGMENT

This work has been performed within a collaboration with Alcatel-Lucent Bell Labs France, under the grant n. 09CT310-01.

REFERENCES

- [1] F. Strohmeier, H. Dörken, and B. Hechenleitner, "AQUILA distributed QoS measurement", International Conference on Advances in Communications and Control, Crete, Greece, 2001.
- [2] C. Molina-Jimenez, S. Shrivastava, J. Crowcroft, and P. Gevros, "On the monitoring of Contractual Service Level Agreements", the first IEEE International Workshop on Electronic Contracting, WEC, San Diego, CA, USA, 2004.
- [3] E. Boschi, S. D'Antonio, P. Malone, and C. Schmolli, "INTERMON: An architecture for inter-domain monitoring, modelling and simulation", NETWORKING 2005, Pages 1397 - 1400, Springer Berlin / Heidelberg, 2005.
- [4] A. Mehaoua et al., "Service-driven inter-domain QoS monitoring system for large-scale IP and DVB networks", Computer Communications, Volume 29, 2006.
- [5] M. Dabrowski, P. Owezarski, W. Burakowski, and A. Beben, "Overview of monitoring and measurement system in EuQoS multi-domain network", International Conference on Telecommunications and Multimedia (TEMU'06), Greece, 2006.
- [6] N. Brownlee, C. Mills, and G. Ruth, "Traffic Flow Measurement: Architecture", RFC 2722, October 1999.
- [7] G. Sadasivan, N. Brownlee, B. Claise, and J. Quittek, "Architecture for IP Flow Information Export", RFC 5470, May 2009.
- [8] B. Claise, Ed. A. Johnson, and J. Quittek, "Packet Sampling (PSAMP) Protocol Specifications", RFC 5476, March 2009.
- [9] Aymen Belghith, Bernard Cousin, Samer Lahoud, and Siwar Ben Hadj Said, "Proposal for the Configuration of multi-domain Network Monitoring Architecture", the 25th International Conference on Information Networking, ICOIN 2011, Kuala Lumpur, Malaysia, January 2011.
- [10] Naidu, K.V.M., Panigrahi, D., Rastogi, R.: Detecting anomalies using end-to-end path measurements. In: the 27th Conference on Computer Communications, INFOCOM 2008, Phoenix, AZ, USA, 2008.
- [11] Barford, P., Duffield, N., Ron, A., Sommers, J.: Network performance anomaly detection and localization. In: the 28th Conference on Computer Communications, INFOCOM 2009, Rio de Janeiro, Brazil, 2009.
- [12] J.-M. Uzé, "IPSphere Forum: status on technical specifications", TER-ENA Networking Conference 2007, Copenhagen, Denmark, 2007.
- [13] G. Almes, S. Kalidindi, and M. Zekauskas, "A One-way Delay Metric for IPPM", RFC 2679, September 1999.
- [14] P. Chimento and J. Ishac, "Defining Network Capacity", RFC 5136, February 2008.