

Crosslayer Network Outage Classification Using Machine Learning

Jan Marius Evang
marius@simula.no
SimulaMet, OsloMet
Oslo, Norway

Ahmed Elmokashfi
ahmed@simula.no
SimulaMet
Oslo, Norway

Azza H. Ahmed
azza@simula.no
SimulaMet, OsloMet
Oslo, Norway

Haakon Bryhni
haakonbryhni@simula.no
SimulaMet
Oslo, Norway

ABSTRACT

Network failures are common, difficult to troubleshoot, and small operators with limited resources need better tools for troubleshooting. In this paper, we analyse two years of outages from a small global network for high-quality services. Then, we develop a machine learning model for outage classification that can be set up with little effort and low risk. We use passive Bidirectional Forwarding Detection (BFD) data to classify Layer2 problems and add active packet loss data to classify other problems. The Layer2 problems were classified with a 99% accuracy and the other problems with 40%–100% accuracy. This is a significant improvement when we observe that only 35% of the customer cases we studied received any Reason for Outage (RFO) response from the Customer Support Centre.

CCS CONCEPTS

• **Networks** → **Public Internet**; **Network measurement**; • **Computing methodologies** → **Supervised learning**.

ACM Reference Format:

Jan Marius Evang, Azza H. Ahmed, Ahmed Elmokashfi, and Haakon Bryhni. 2022. Crosslayer Network Outage Classification Using Machine Learning. In *Applied Networking Research Workshop (ANRW '22)*, July 25–29, 2022, PHILADELPHIA, PA, USA. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3547115.3547193>

1 INTRODUCTION

Today’s Internet comprises a group of small, medium, large and extra large networks as far as geographic presence and traffic volume are concerned. The end-to-end network service is produced following a three-layer model that is similar to the lower levels of the OSI reference model [1].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
ANRW '22, July 25–29, 2022, PHILADELPHIA, PA, USA

© 2022 Association for Computing Machinery.
ACM ISBN 978-1-4503-9444-4/22/07...\$15.00
<https://doi.org/10.1145/3547115.3547193>

Table 1: Manually classified causes.

Class	Cause	Count
MultiLoss	Multiple Layer2 providers	870
CogentLoss	Cogent’s network	556
Customer	Customer’s equipment	411
TeliaLoss	Telia’s network	316
Layer3	Layer3 only	222
InternMaint	Internal maintenance	114
Optic	3dB Optical change	74
ProvMaint	Provider maintenance	70
EquinixLoss	Equinix Cloud Fabric	58
SubseaCable	Subsea cable outages	42
EquipFail	Equipment failure	40
FiberCut	Fiber cut in provider network	39
Layer1	Leased Layer1 lines	18
Metro	Metropolitan area links	18
DoS	Denial of Service attacks	4

A few large providers sell Layer2 capacity based on the global mesh of Layer1 optical fibres, which are used by Layer3 providers to compose end to end services.

This layered architecture is exposed to various types of faults, such as physical fiber faults, equipment faults, planned maintenance and malicious attacks. Our data shows that the Layer1/Layer2 service has a high number of faults (see Table 1). Smaller networks that lease Layer1/Layer2 services need to quickly attribute such faults and report them to the respective providers. This is important for two reasons. First, it can help shorten the resolution time. Second, faults must be reported during the incident to be acknowledged according to the Service Level Agreements (SLAs).

Unlike large networks with sizable organization and abundant resources, small and medium network operators have a much smaller Network Operations Centre (NOC) with limited resources and staff. A typical small-medium NOC either operates a single enterprise network or is a speciality Internet Service Provider (ISP) providing a service to select customers in a narrow business area or in a geographic area.

Smaller NOCs often have a small but highly demanding customer base, for instance their co-workers in an enterprise, people in their

own geographic area or specialized service providers. This makes detecting and isolating faults very important yet a demanding task.

The NOC usually has automatic network monitoring systems in operation, but they can suffer from large numbers of both false positives (alerts without a real fault) and false negatives (faults that do not generate an alarm). This often causes true positives to be overlooked [2]. In an outage event where one component in the network has failed, causing interruption to network traffic, an overwhelming amount of log messages and alerts will be arriving from different monitoring systems. This makes the NOC waste time and effort to find the real cause. In other cases, a problem may not be noticed until customers complain. Customer Support (CS), may not have enough information to respond to a customer case because the NOC is busy troubleshooting. Alarm Consolidation systems exist but they suffer from high complexity [3], narrow field [4] or high compute requirements [5]. In this work, we tackle these problems by developing a generic model to assist NOCs and CSes. We leverage supervised learning to assist in classifying different outages. For classification, we use the Support-Vector Machine model (SVM) [6]. Our system is two-stage. In the first, it discriminates Layer1/Layer2 problems from Layer3 ones. Here, we identify a set of easily to collect metrics that can help achieving this in an efficient manner. In the second, it classifies Layer3 problems based on their root causes.

The research in [7] claims that supervised learning for fault classification is often suffering from low quality of training data, but in our research we have access to precise outage data, including root cause data.

Our system requires minimal changes to the network, and has a minimal impact on networking equipment and computing power. We also demonstrate that our proposed system is implementable and can be used to assist an existing provider efficiently.

With the system developed here, the NOC will speed up troubleshooting, quickly create trouble tickets with the providers, and the CS will improve customer satisfaction by giving informed feedback to all customer support cases. Compared to similar systems such as [5], investments in time and equipment are small, changes to configuration is minimal, and causes are successfully predicted with an f1-score of 0.99 for Layer2 cases and f1-score of 0.66 for other cases (see Section 4.1). Without the tool, only 35% of the cases received any outage report from CS.

2 RELATED WORK

Various works have used machine learning and other statistical methods for attributing faults for specific network protocols, however, there is still lack of work that leverages logs from different layers, and predict causes across network layers.

Existing research such as [3] implements a complex system of user defined scenarios, while they do not require detailed knowledge of the underlying system, they cannot detect problems outside manually defined failure scenarios. Our labeled data and two-stage approach makes classification of known faults across all layers possible and efficient, and the feedback loop handles new fault classes. Moreover, several projects [4, 8–12] examine how very detailed measurements of optical signal strength can be used to gain knowledge about the underlying Layer1 links. However, these

methods require measurement of q-factor [13] telemetry [14] which is unavailable to higher layers providers.

The research in [15] also uses customer tickets for anomaly detection, but focuses only on Layer1 and last mile. The authors in [16] analyse Layers 2-7, while we analyse backbone Layers 1-3, and a common “No issue” class for any problems in other layers or outside the backbone network. Unlike [17], our system does not need any knowledge about the underlying network, only the manual feedback needs this.

Some commercial service providers have implemented systems for anomaly detection in system logs, for instance [18]. These systems have the advantage that they analyse the existing logs, and therefore are easy to start using, however, there is a high risk of exposing confidential information to a third party. In our system, only the feedback loop will have any confidentiality risk.

Finally, the authors in [5] analyse traffic by using a distributed Apache Storm [19] system in combination with data obtained from the Netflow [20] protocol. This puts extra stress on the networking equipment [21] and demands much more storage and CPU power, making it undesirable unless Netflow is already used for other purposes. BFD, on the other hand, is usually implemented in hardware.

The objective in this paper is to fill the gap and use simple data logs from various layers together with customer support data to classify outages in a fast and easily-implementable low-impact solution.

3 METHODOLOGY

3.1 Description of system

Our system consists of a data collection unit, a classification model (See Section 3.4), and alert and feedback units as shown in Figure 1.




Figure 1: Our proposed outage classification system.

We collected the measurements from a global network covering 12 cities around the world, which we depict in Figure 2. The network uses three different Layer2 service providers to interconnect its points of presence (PoPs). The first is Telia VPLS, which is full-mesh Layer2 switched network based on VPLS/ELAN [22] over their global backbone network. The VPLS service supports Q-in-Q switching [23], so individual point-to-point VLANs [24] are configured, with each VLAN having member ports from only two cities. The second is Cogent L2C, which is a point-to-point MPLS [25] based service where multiple point-to-point Layer2 links are provided over the same physical interface. The third is Equinix Cloud Exchange Fabric (ECXF), which is a service of multiple point-to-point Layer2 links over the same physical interface.

The network is set up with IS-IS + BFD [26, 27] as Interior Gateway Protocol (IGP). The IGP makes sure that in case of issues on one link or device, customer traffic is automatically re-routed to an alternative path.

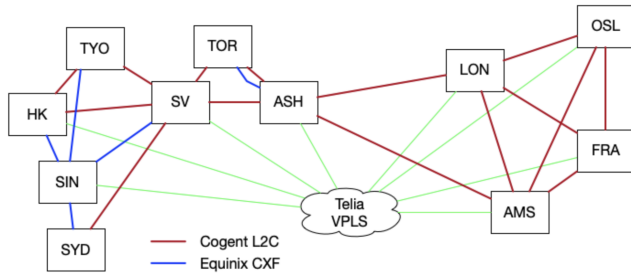


Figure 2: Network design.

3.2 Data description

In this work, we collected data from the monitored ISP over two years (2019-08-08 to 2021-10-01). Below, we list these measurements alongside their description.

Optical signal strength measurements. Every 30 minutes, optical received signal strength for the local link to the provider was read via SNMP [27]. There were 12 total outages on 5 interfaces, 88 drops in optical strength of more than 3dB on 14 interfaces and 53 increases of more than 3dB on 10 interfaces.

Interface error counters. Every 10 minutes, interface error counters for all devices were logged by SNMP polling. No interface errors were recorded for Layer2 switch ports, because any such errors would have been revealed and corrected during pre-production testing.

Buffer overflows/Tail-drops. Every 10 minutes, the buffer overflow/tail-drop counters were logged by SNMP polling. Only two interfaces showed tail drops, altogether 40 incidents. The NOC had especially amended this risk by over-provisioned the network to handle network traffic peaks without packet loss.

Layer2 packet loss data. In each of the 11 cities shown in Figure 2, 2 probe Virtual Machines (VMs) were set up. Our probe software is based on OpenNetNorad [28], which we rewrote in C to improve performance and reduce CPU consumption. This “pinger” transmits 100 UDP 64 byte packets every 0.5 seconds and waits for responses, and the “ponger” immediately returns any received packets to the sending IP address. The number of lost packets is then recorded. Packets are transmitted from 100 different UDP ports to detect any issues related to link aggregation or Equal Cost Multipath (ECMP) within the provider network. In addition to the probes measuring point-to-point (P2P) loss over the Layer2 links, a full mesh of probes (FM) were set up to measure the Layer3 service. One or more lost UDP packets in a 0.5sec interval generates one loss report. There were 196 million loss reports for 36 different pairs of probe VMs. These were pre-processed to 717352 unique events (see Section 3.3).

BFD traps. For each point-to-point link or VLAN, BFD (Bidirectional Forwarding Detection) [29] is configured to send one packet

every 100ms. If 3 packets in a row are lost, the link is declared down and an SNMP trap message is sent to a collector. SNMP trap data is passively collected and stored in a database for later processing. The IS-IS protocol also receives BFD events and takes care of re-routing traffic.

Software crash logs. There were 62 instances of software crash/core-dumps incidents on the routers and switches. Most of these did not cause any interruption to network traffic since the Forwarding Engines were still operational.

Configuration change logs. Configuration change logs indicate which piece of equipment was configured and when. Also a textual description of the work was performed.

Customer complaints data. The customers’ systems have strict network requirements for latency, packet loss and jitter. Customer cases were raised upon any violation of these requirements. The data was anonymized and made available for this work. During the period, there were 19399 customer cases, of which 8120 were related to the network. The complaints were reduced to 2855 unique cases on 21 different paths.

Customer service response data. For each customer case, CS analysed logs and provided a Reason For Outage (RFO) if possible. Out of 2855 cases, 1014 (35%) received RFO from CS, 109 of these were “no issue found”.

Manual analysis of customer reports. We looked at all available data for each customer reported case and determined the reason for the incident. In most cases the cause was in a Layer2 provider’s network. For other cases the cause could be determined more precisely from CS responses. The results are presented in Table 1.

In some cases, there were losses in multiple providers at the same time, which may be caused by either an (undetected) failure in the monitored network, a larger failure that impacted multiple providers, short traffic peaks that caused packet loss and therefore triggered a re-routing to another provider and subsequent loss there, or could be just a coincidence.

Multiple customer complaints received within a 5-minute interval were counted as one case. Still, a single root cause could cause multiple cases over a longer time. Some incidents were caused by planned or unplanned maintenance. These were recorded as cases, if they caused customer complaints even when the customer had been informed ahead of time.

The 713857 events that did not correspond to customer cases were not manually analysed.

3.3 Data preprocessing

The data used for the Machine Learning algorithm was BFD SNMP events (BFD), point-to-point UDP pings (P2P) and full-mesh UDP pings (FM). The other data was used only in the manual classification process of all the cases. The result of the manual classification was used to train the supervised machine learning system.

Due to small delays in detection and collection of test data, the resolution of the timestamps had to be reduced to match events from different sources. Each measuring point was added as a separate feature, with an aggregation of the number of such events per minute. One minute aggregation was chosen as a trade-off between fast detection and data size. For the BFD and P2P data, the measuring points were each link, for the FM data, the measuring points were

the unique pairs of PoPs. This resulted in a dataset of 717352 unique events and 2855 unique cases. The features were 47 BFD, 32 P2P and 125 FM. The classes with < 4 cases were omitted.

3.4 Model description

We tested both Multilayer Perception neural networks (MLP) and SVM. SVM had both shortest processing time and highest classification accuracy, and is used in this paper. The data was split 75:25 into a training dataset and a testing dataset, and we tuned the hyperparameters using grid search. The optimal kernel was the Radial Basis Function (RBF) kernel with $C = 150$ and $\gamma = 7.5 \times 10^{-5}$. SVM is in general resistant to overfitting and we verified this by ShuffleSplit [30] and saw that the f1-score remained the same.

The first stage classification used only BFD data for classifying the largest and most precisely defined classes, i.e. the Layer2 provider cases. The output was five classes. One per each Layer2 provider, A fourth class that involve cases where more than one Layer2 provider, and one “Layer3” class for cases which were not caused by Layer2 events. A large number of events were processed in the first stage, but since fewer features were used, processing requirements were greatly reduced. The second stage classification used BFD, P2P and FM data for the Layer3 class to give an indication of the root cause. Since a much smaller subset of events was processed in this stage, the addition of more features did not lead to a large increase in processing power requirement. See also Section 4.4.

The feedback loop is used by NOC/CS when a prediction has failed, to manually correct the case label in the data and re-train the model.

After training the two machine learning models on the case data, the trained models were applied to all events, to see what knowledge could be gained.

4 PERFORMANCE EVALUATION

4.1 Evaluation metrics

We used the precision, recall and f1-score to assess our classifier.

For each class, the precision is the number of correctly predicted cases divided by the total predictions in that class. Recall is the number of correctly predicted cases divided by the number of true cases in that class. F1-score is the harmonic mean of precision and recall [31].

To visually evaluate the output of the classification process, we plot the Confusion Matrices. These show how well the model was able to assign a correct “predicted label” to each class of “true labels”. The diagonals of the matrices show the correct predictions.

4.2 Accuracy and Feature importance

We performed the first classification stage initially by including all features, which resulted in a precision of 0.89, a recall of 0.89 and an f1-score of 0.92 (see the confusion matrix is in Figure 3a).

Using only BFD features showed much better scores for Layer2 cases, but did (as expected) not distinguish between Layer3 and Customer issues as seen in the confusion matrix in Figure 3b and scores in Table 2. Total f1-score was now 0.99 with a combined Customer+Layer3 class. Further, repeating the first stage while

Table 2: First stage evaluation, based on BFD.

class	precision	recall	f1-score
CogentLoss	1.00	0.99	0.99
TeliaLoss	1.00	1.00	1.00
MultiLoss	0.99	0.99	0.99
EquinixLoss	0.88	1.00	0.94
Customer	0.65	1.00	0.79
Layer3	0.00	0.00	0.00

Table 3: Second stage prediction scores (Based on BFD+P2P+FM)

class	precision	recall	f1-score
InternMaint	0.65	0.72	0.68
Optic	0.75	0.43	0.55
ProvMaint	0.33	0.55	0.41
SubseaCable	0.92	0.92	0.92
EquipFail	0.40	0.40	0.40
FiberCut	0.75	0.64	0.69
Layer1	0.83	1.00	0.91
Metro	1.00	0.43	0.60
DoS	1.00	1.00	1.00

including only FM and only P2P gave poor results with f1-score 0.35 for FM and and f1-score of 0.26 for P2P (see Figures 3c and 3d).

The BFD analysis contained only 4 misclassifications: 2 Multi-Loss events classified as EquinixLoss were caused by two unrelated coinciding loss events where the EquinixLoss event affected multiple Equinix links, and 2 CogentLoss events classified as MultiLoss were multiple coinciding Cogent events. The analysis including all features added the capability of distinguishing between Layer3 loss and Customer loss, at the expense of requiring more computing time and adding more “noise” to the various Layer2-classifications. Still, we see a relatively small number of misclassifications (14 misclassified and 429 correctly classified Layer2 events).

For the second stage, the events that were identified by the first stage classification were removed, and a new supervised classification was attempted for the remaining events. After hyperparameter tuning, this classification showed an f1-score of 0.66. The size of the dataset in this analysis is only 437 cases with 204 features, and the results were not as good as for the first stage, but a reasonable suggestion for a root cause might still provide valuable input to the NOC’s troubleshooting process. Figure 4 and Table 3 show the confusion matrix and classification score for stage 2, respectively. We can clearly see that determining the exact root cause can be hard for a few types of failures. For instance, ProvMaint and InternalMaint events may cause a wide variety of different error symptoms, that may be indistinguishable from the other classes. Interestingly, subsea cable cuts (f1-score 0.92) and fiber cuts (f1-score 0.69) had relatively good classification scores, even though these were thought to be difficult to distinguish. A point for future study might be to understand why.

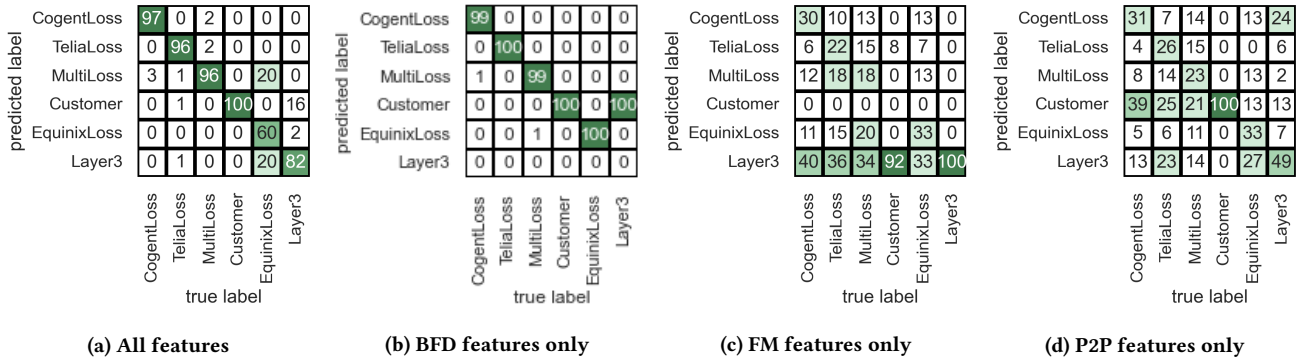


Figure 3: First stage classifications

predicted label	InternMaint	72	14	36	8	60	14	0	29	0
	Optic	3	43	0	0	0	7	0	0	0
	ProvMaint	11	36	54	0	0	7	0	29	0
	SubseaCable	0	0	9	92	0	0	0	0	0
	EquipFail	8	0	0	0	40	0	0	0	0
	FiberCut	6	7	0	0	0	64	0	0	0
	Layer1	0	0	0	0	0	7	100	0	0
	Metro	0	0	0	0	0	0	0	43	0
	DoS	0	0	0	0	0	0	0	0	100
		true label	InternMaint	Optic	ProvMaint	SubseaCable	EquipFail	FiberCut	Layer1	Metro

Figure 4: Confusion matrix for second stage classification

Table 4: First stage data of the Layer2 cases and predictions

class	support cases	extrapolated cases
CogentLoss	556 (19.4%)	49997 (7.0%)
TeliaLoss	316 (11.1%)	64859 (9.1%)
MultiLoss	870 (30.5%)	47227 (6.6%)
EquinixLoss	58 (2.0%)	15346 (2.1%)
Customer+Layer3	633 (22.2%)	536428 (75.1%)
Other	14.8%	

4.3 Extrapolation

Using the first stage model, BFD-trained on the cases with well known cause and symptoms, we ran a prediction on all the events where we did not get any customer complaints, to get an idea of how common the various types of problems are in these events. The very high f1-score of the model fitted on the complaint-data means that the predictions on the non-complaint-data will be highly relevant for our research. However, selection bias in that some hidden class of outages never leads to complaints might reduce the accuracy of the extrapolation.

For the first stage model, the results can be seen in Table 4. The most interesting observation is that the “Customer+Layer3” classification is much more common than in the cases where the

Table 5: Second stage data for the Layer3 cases and predictions

class	cases	predictions
InternMaint	114 (27.1%)	185169 (34.5%)
Optic	74 (17.6%)	216591 (40.4%)
ProvMaint	70 (16.7%)	26883 (5.0%)
SubseaCable	42 (10.0%)	19522 (3.6%)
EquipFail	40 (9.5%)	18308 (3.4%)
FiberCut	39 (9.3%)	58099 (10.8%)
Layer1	18 (4.3%)	10450 (1.9%)
Metro	19 (4.5%)	1329 (0.2%)
DoS	4 (1.0%)	77(0.01%)

customers filed complaints. (75.1% of the events, versus 22.2% of the cases). This means that the test network does a good job of hiding Layer3 problems from customers, and Layer2 problems are more likely to cause customer complaints, but still only 0.4% of all events caused customer cases.

The “MultiLoss” class is only 6.6% of the events in the non-complaint dataset, vs 30.5% of the complaint-cases. This indicates that the network is better at hiding Layer2 problems in a single provider, and problems affecting multiple providers are more likely to generate customer complaints.

Further, we used the model fitted on the Second stage data from the cases, and made a prediction using only the “Customer+Layer3” class from the first stage non-complaint events. Applying the second stage model to the non-complaint data gives an indication that Internal maintenance and Optic events are less likely to cause customer complaints than the other classes, but the size of the dataset and the lower accuracy of the model makes these results much less certain. See Table 5.

4.4 Processing performance

BFD is implemented in hardware on our routers and do not put any load on the routers’ CPU. To compare, Netflow would cause 15%-20% CPU impact according to [5], which matches our own experience. SNMP traps produced by the routers using the lowest priority processes, and all data is transmitted blindly using UDP, also reducing processing. Our ML processing on an M1 Pro 10 core

CPU took <1sec. The amount of stored data for the ML system is low. For each BFD trap we store timestamp+link-id and for each UDP measurement we store timestamp, source/destination address and loss percentage.

5 DISCUSSION

Our analysis of two years of outage data shows that a two-stage classification system is well suited to classify network outages, providing the NOC with useful predictions on where to start troubleshooting, and providing CS with RFO for all cases with a much better success rate than the observed 35% of CS responses during the period of the study. BFD data exhibits their high importance in the classification. In contrast, although the active P2P and FM raw data provides very precise measurements, they are not highly contributing to discriminating features in the classification model.

One important shortcoming is that we do not have latency measurements. But as our analysis reveals, BFD SNMP traps are very good indicator of problem types and location, so latency changes would probably not have a great impact on this result. Moreover, in this work, customer complaints are the only source for determining whether a packet loss event is regarded as an outage. Only the cases that are received as customer complaints are analysed in detail. This means that some outages may be overlooked if the customer did not complain, and some complaints may be groundless (i.e caused by other factors than the test network). A customer complaint is only counted as a network outage if the timestamp is reported as within 60 seconds of an internal packet loss or BFD trap event.

There are many features that show some correlation, which might disturb the machine learning classification model since one event is likely to affect multiple features. But since the features have a large geographic spread, and since there are many features, a certain degree of correlation should not cause problems for our analysis.

Model Drift (MD) is another consideration. During the 2 years of data collection, there were continuous changes to both the network topology, the routing protocols and the customer's monitoring system. MD may have degraded our analysis, in that patterns for the various classes of events change over time. However, this will also reflect more accurately a real-life situation. The results prove that our first stage analysis was not significantly affected by MD. In the future, we plan to gain insight into how our model may degrade, for instance by temporal cross validation, and how to rectify it through a system for retraining while running in the production. A future improvement, especially for the second stage, would be to also report the second ranked classification for an outage.

Another very important practical consideration is the difference in complexity of gathering the data for the first stage and the second stage. The passive BFD data used for the first stage is very easy to collect. Most networks already use the BFD protocol as a part of the IGP protocol, but very few actually gather the SNMP trap data from BFD. The changes and risk to the network will be small, the BFD events are already being detected, so the only change is to generate SNMP trap messages and set up one central location to store these (optionally a second location for redundancy.)

The active P2P and FM raw data are very accurate and provides very precise measurements, but showed less precision in case classification, and the system used to gather this data is much more expensive in management and computing power.

6 CONCLUSION

We have developed a system that Network Operations Centres and Support Centres for smaller operators can use in a failure situation. Using minimal resources, we passively collect BFD data and classify the Layer2 events to an f1-score of 0.99. By adding a second stage with active monitoring to collect UDP ping data we predict other types of root cases with a 0.66 f1-score. Our analysis interestingly shows that BFD features, which are the easiest to collect, give the best results for outage classification.

REFERENCES

- [1] J. Day and H. Zimmermann, "The osi reference model," *Proceedings of the IEEE*, vol. 71, no. 12, pp. 1334–1340, 1983.
- [2] B. AlAhmadi, L. Axon, and I. Martinovic, "99% false positives: A qualitative study of soc analysts' perspectives on security alarms," *USENIX Association*, 2021. [Online]. Available: <https://ora.ox.ac.uk/objects/uuid:0be05f6b-7470-4210-acb6-2018d5dc6ca0>
- [3] K. Appleby, G. S. Goldszmidt, and M. Steinder, "Yemanja—a layered fault localization system for multi-domain computing utilities," *Journal of Network and Systems Management*, vol. 10, pp. 171–194, 2004.
- [4] T. Christopoulos, O. Tsilipakos, G. Sinatkas, and E. E. Kriezis, "On the calculation of the quality factor in contemporary photonic resonant structures," *Opt. Express*, vol. 27, no. 10, pp. 14 505–14 522, May 2019. [Online]. Available: <http://opg.optica.org/oe/abstract.cfm?URI=oe-27-10-14505>
- [5] Y. Du, J. Liu, F. Liu, and L. Chen, "A real-time anomalies detection system based on streaming technology," in *2014 Sixth International Conference on Intelligent Human-Machine Systems and Cybernetics*, vol. 2, 2014, pp. 275–279.
- [6] R. Soentpiet et al., *Advances in kernel methods: support vector learning*. MIT press, 1999.
- [7] S. Ayoubi, N. Limam, M. A. Salahuddin, N. Shahriar, R. Boutaba, F. Estrada-Solano, and O. M. Caicedo, "Machine learning for cognitive network management," *IEEE Communications Magazine*, vol. 56, no. 1, pp. 158–165, 2018.
- [8] C. Natalino, A. di Giglio, M. Schiano, and M. Furdek, "Root cause analysis for autonomous optical networks: A physical layer security use case," in *2020 European Conference on Optical Communications (ECOC)*, 2020, pp. 1–4.
- [9] L. Shu, Z. Yu, Z. Wan, J. Zhang, S. Hu, and K. Xu, "Low-complexity dual-stage soft failure detection by exploiting digital spectrum information," in *45th European Conference on Optical Communication (ECOC 2019)*, 2019, pp. 1–4.
- [10] C. Delezoide, P. Ramantanis, L. Gifre, F. Boitier, and P. Layec, "Field trial of failure localization in a backbone optical network," in *2021 European Conference on Optical Communication (ECOC)*, 2021, pp. 1–4.
- [11] Ujjwal, J. Thangaraj, and A. A. Dias Barreto, "Accurate qot estimation for the optimized design of optical transport network based on advanced deep learning model," *Optical Fiber Technology*, vol. 70, p. 102895, 2022. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1068520022000785>
- [12] C. Miao, M. Chen, A. Gupta, Z. Meng, L. Ye, J. Xiao, J. Chen, Z. He, X. Luo, J. Wang, and H. Yu, "Detecting ephemeral optical events with OpTel," in *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*. Renton, WA: USENIX Association, Apr. 2022, pp. 339–353. [Online]. Available: <https://www.usenix.org/conference/nsdi22/presentation/miao>
- [13] "Recommendation O.201: Q-factor test equipment to estimate the transmission performance of optical channels," International Organization for Standardization, Geneva, CH, Standard, 2003.
- [14] [Online]. Available: <https://www.juniper.net/documentation/us/en/software/junos/interfaces-telemetry/index.html>
- [15] J. Hu, Z. Zhou, and X. Yang, "Characterizing Physical-Layer transmission errors in cable broadband networks," in *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*. Renton, WA: USENIX Association, Apr. 2022, pp. 845–859. [Online]. Available: <https://www.usenix.org/conference/nsdi22/presentation/hu>
- [16] J. Iurman, F. Brockners, and B. Donnet, "Towards cross-layer telemetry," in *Proceedings of the Applied Networking Research Workshop*, ser. ANRW '21. New York, NY, USA: Association for Computing Machinery, 2021, p. 15–21. [Online]. Available: <https://doi.org/10.1145/3472305.3472313>
- [17] R. R. Kompella, J. Yates, A. Greenberg, and A. C. Snoeren, "IP fault localization via risk modeling," in *2nd Symposium on Networked Systems*

- Design & Implementation (NSDI 05)*. Boston, MA: USENIX Association, May 2005. [Online]. Available: <https://www.usenix.org/conference/nsdi-05/ip-fault-localization-risk-modeling>
- [18] [Online]. Available: zerbiium.com
- [19] The Apache Software Foundation, "Apache storm," <https://storm.apache.org/>.
- [20] E. B. Claise, "Cisco systems netflow services export version 9," Internet Requests for Comments, RFC Editor, RFC 3954, 8 2004, <http://www.rfc-editor.org/rfc/rfc3954.txt>. [Online]. Available: <http://www.rfc-editor.org/rfc/rfc3954.txt>
- [21] "Netflow services," p. 74, 2003.
- [22] Metro Ethernet Forum, "Ethernet services definitions - phase 2," 4 2008.
- [23] "Provider bridges, ieee std. 802.1ad," 2005.
- [24] "Bridges and bridged networks, ieee std. 802.1q-2018," 2016.
- [25] E. Rosen, A. Viswanathan, and R. Callon, "Multiprotocol label switching architecture," Internet Requests for Comments, RFC Editor, RFC 3031, 1 2001, <http://www.rfc-editor.org/rfc/rfc3031.txt>. [Online]. Available: <http://www.rfc-editor.org/rfc/rfc3031.txt>
- [26] International Organization for Standardization, *Information technology — Telecommunications and information exchange between systems — Intermediate System to Intermediate System intra-domain routeing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service*, ISO/IEC10589:2002 ed. Vernier, Geneva, Switzerland: International Organization for Standardization, 2015. [Online]. Available: <https://www.iso.org/standard/30932.html>
- [27] D. Katz and D. Ward, "Bidirectional forwarding detection (bfd) for ipv4 and ipv6 (single hop)," Internet Requests for Comments, RFC Editor, RFC 5881, 6 2010.
- [28] Facebook Inc, "Opennetnorad," <https://github.com/fbsamples/OpenNetNorad>, 2017.
- [29] R. Presuhn, "Version 2 of the protocol operations for the simple network management protocol (snmp)," Internet Requests for Comments, RFC Editor, STD 62, 12 2002, <http://www.rfc-editor.org/rfc/rfc3416.txt>. [Online]. Available: <http://www.rfc-editor.org/rfc/rfc3416.txt>
- [30] Q.-S. Xu and Y.-Z. Liang, "Monte carlo cross validation," *Chemometrics and Intelligent Laboratory Systems*, vol. 56, no. 1, pp. 1–11, 2001. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S016974390001222>
- [31] N. Chinchor, "MUC-4 Evaluation Metrics," in *Proceedings of the 4th Conference on Message Understanding*, ser. MUC4 '92. USA: Association for Computational Linguistics, 1992, p. 22–29. [Online]. Available: <https://doi.org/10.3115/1072064.1072067>