

# Erratic server behavior detection using machine learning on streams of monitoring data

Martin Adam<sup>1,2,3,\*</sup>, Luca Magnoni<sup>2</sup>, Martin Pilát<sup>3</sup>, and Dagmar Adamová<sup>1</sup>

<sup>1</sup>Academy of Sciences of the Czech Republic (CZ)

<sup>2</sup>CERN, European Organization for Nuclear Research (CH)

<sup>3</sup>Charles University, Faculty of Mathematics and Physics (CZ)

## Abstract.

With the explosion of the number of distributed applications, a new dynamic server environment emerged grouping servers into clusters, utilization of which depends on the current demand for the application. To provide reliable and smooth services it is crucial to detect and fix possible erratic behavior of individual servers in these clusters. Use of standard techniques for this purpose requires manual work and delivers sub-optimal results. Using only application agnostic monitoring metrics our machine learning based method analyzes the recent performance of the inspected server as well as the state of the rest of the cluster, thus checking not only the behavior of the single server, but the load on the whole distributed application as well. We have implemented our method in a Spark job running in the CERN MONIT infrastructure. In this contribution we present results of testing multiple machine learning algorithms and pre-processing techniques to identify the servers erratic behavior. We also discuss the challenges of deploying our new method into production.

## 1 Introduction

In the last few decades the amount of digitally saved data has been growing exponentially [?]. In 2014 the world's technological capacity to store information has reached almost 5 zettabytes [1]. Handling this incredible amount of incoming data requires innovative techniques increasingly leveraging horizontal scaling; an approach utilizing many computers instead of one more powerful. Such novel approaches create additional concerns for the system administrators, particularly when it comes to noticing errors that pose a threat to the efficiency and availability of the application. Although traditional monitoring methods require lots of manual labor when applied to this problem, developers of open source monitoring system have been so far reluctant to include any advanced tools. In this project we set to explore the possibility of using machine learning to spot erratic servers within a cluster running a distributed application.

In an attempt to simplify administrators work, many applications offer a set of internal metrics describing their performance. Incorporating these metrics in the existing monitoring systems might be too time-consuming, considering that the lack of skilled administrators often leads to understaffed teams. Therefore this project avoids using application specific

---

\*e-mail: [madam@fzu.cz](mailto:madam@fzu.cz)

metrics considering only basic os-level performance metrics, which leads to a black-box approach of the servers.

In this article, we present a process of acquiring and processing a stream of raw monitoring data in the MONIT [2] infrastructure. The data is then analyzed offline by a concept application using a data processing pipeline to detect unexpected behavior considering both workload (current load on the distributed application) and workflow (usual node behavior running a specific application). In its core is a prediction model trained by machine learning. We also discuss the efficiency of such an approach, benchmark the core model and present plans for future development.

## 2 Monitoring Systems Overview

The primary goal of traditional monitoring systems is gathering and saving monitoring data, with presenting and especially analysing it as a secondary objective. Instead, some applications are making the footprint of the client as small as possible, such as Ganglia [3], others offer effortless extensibility, such as Munin [4]. The only common erratic behavior detection feature available is comparing a single metric to a pre-defined threshold.

A completely different approach is taken by Nagios [5], which specializes in monitoring a status of a service on a server, with collecting metrics as a secondary task. It however does not offer an easy way to correlate metrics within a server, let alone within a cluster of servers.

Academic works on anomaly detection in monitoring data are available, but custom systems are always used. Some works explore the subject from the angle of intrusion detection [6]. Others take on the challenge of predicting failures [7], but given the advanced difficulty of this task all the works make use of all the available data including application metrics or logs [8] making them very application specific. There are some works looking into the black box method [9], but they are usually only showcased on artificial datasets.

At CERN an extensive monitoring system has been build to gather, process and visualize metrics and logs from the 40,000 machines powering the entire scientific, administrative and computing infrastructure. This data pipeline, which handles more than 3TB of compressed data per day, is composed of several distributed services, each running on several tens of machines, providing key functionality such as Apache Flume (data collection and ingestion) [10], Apache Kafka (data transport) [11] and Apache Spark (data processing) [12]. Each of those services represents the perfect case study for our research, as an application that distributes its load evenly was needed. This project will be using the MONIT infrastructure as a platform to both collect and analyze data.

## 3 Data Gathering

The MONIT infrastructure offers OS metrics via its collectd stream [13]. The metrics are reported individually by type and also each host reports them at a different time. We therefore decided to aggregate the metrics individually over a fixed time window first and join them together later. The longest metric readout interval among these metrics is 5 minutes, so the averaging time window length was set as 20 minutes. This guarantees that in each time window we get at least 4 samples of each metric, providing a simple smoothing procedure.

Two Spark streaming jobs were deployed to the MONIT infrastructure realizing the described operations. First is consuming all the collectd data. To decompress and read 13MB/s it needs to run in cluster mode on 25 CPU cores and using over 111GB of memory. It selects only data for clusters of interest, applies the aggregation and writes the processed data back into Kafka.

The second job consumes only the processed data and therefore has a significantly smaller resource usage. As joining is a stateful operation, it still needs to keep a checkpoint saved for fault-tolerance, which can pose a challenge to the underlying storage system due to its size. Metrics are joined on hostname and time-window and the results are then written by a Flume agent to HDFS for later analyses.

### 3.1 Creating Anomalies

It became clear soon after we started collecting the data, that the number of naturally occurring anomalies is too small for a meaningful statistic. We selected a cluster of virtual servers running the Kafka service for MONIT in a development setting<sup>1</sup> and started creating our own anomalies. Alternating the already collected metrics might not produce plausible results, so creating an error on the server itself, while data was being collected and the rest of the cluster ran a normal workload, was our preferred option. Several classes of errors were thought of and implemented:

- Base line anomaly — the main application is shut down on one of the nodes in the cluster.
- Memory-leak type — the memory usage gradually increases up until a point where the OOM killer<sup>2</sup> engages. In this case the memory allocation is stopped before the OOM killer intervenes, to avoid the situation when the main application process is killed.
- CPU over-utilization — the CPU usage increases conflicting with the main application possibly rendering it un-operationable.
- Combined — both cpu and memory usage increases. Such an anomaly might indicate a rogue program being launched, or defective load balancing in the main application (the anomalous node is overloaded).

The service had to run smoothly for a period of time before we started inserting anomalies for a model to be able to learn. After at least 2 days of stable and uninterrupted production, anomalies were inserted once or twice per day. Table 1 presents the final number of collected anomalies per type.

anomaly type	count
service stop	6
memory overuse	2
cpu overuse	8
combined cpu and memory overuse	10

Table 1: Final anomaly count per type.

## 4 Analysing the Data

Collecting all the available collectd metrics resulted in status snapshots with 35 attributes. Input data with such high dimensionality might cause sub-optimal performance of some of

<sup>1</sup>This cluster was used for testing changes before applying them into production. Most of the time the cluster ran with the same setting as the production cluster, but meddling with it did not pose a risk to the stability of services provided by MONIT.

<sup>2</sup>Out Of Memory Killer (OOM) is a process that the Linux kernel employs when the system is critically low on memory. The kernel attempts to recover memory by terminating programs so that the system can continue running




Figure 1: Proportions of variance explained by each principal component.

the algorithms and because we were expecting a portion of the metrics to be correlated, the Principle Component Analysis was employed [14]. The number of dimensions to be used after the PCA transformation based on the variance explained by respective components (see Figure 1) was selected to be 3 or 5 for some algorithms.

#### 4.1 Unsupervised learning

Anomaly detection can be approached with unsupervised learning as a problem of characterizing a dataset based solely on the regular data points. New data are then determined to be anomalies based on how well they would fit the learnt characterization. Two algorithms provided by the Scikit-learn library [15] were selected, namely the One Class Support Vector Machine [16] and the Isolation Forest [17]. The hyperparameters for both algorithms were chosen using 5-fold cross validation on one random day with an anomaly and then fixed for the rest of the anomaly days. From Table 2 it is clear, that the Isolation Forest algorithm yields better results than the One Class SVN. PCA pre-processing helps the OC-SVN, which is more sensitive to high dimensionality, but it is not helpful in case of the Isolation Forest. Other available methods were not tested due to their similar properties to clustering algorithms, which modeled the data poorly<sup>3</sup>.

<sup>3</sup>Clustering algorithm analysis was performed in the data exploration step of this project. It was omitted from this paper due to insignificant results.

model	precision	recall
I-forest-nopca	0.79	0.37
I-forest-pca3	0.72	0.32
OC-SVN-nopca	0.29	0.16
OC-SVN-pca3	0.40	0.21

Table 2: Mean precision and recall for the unsupervised algorithms. Suffix -pca3 means the algorithms' input data were pre-processed by PCA and the first three components were used and the -nopca suffix denotes an algorithm trained without any PCA pre-processing.

## 4.2 Supervised learning

Although using supervised classification on labeled anomalies was agreed not to be explored in this work, supervised learning could be employed in a different way. With time series data, a regression task is hidden in predicting the next data point based on one or several historic ones. To represent the relevant history, input data are several metric vectors concatenated together, each representing a server status snapshot in time. The output would then be vector describing the next status snapshot<sup>4</sup>. The dimensionality of such input data can grow dramatically when trying to capture longer periods of time or without using any additional dimensionality reduction techniques. With this in mind, the decision was made to set the number of historical snapshots  $n$  to 4. With the 20 minute time window that gives the algorithm knowledge about an 80 minutes history of the server in question.

$$X_t = (x_t^1, x_t^2, x_t^3, \dots, x_t^N)$$
$$f(X_{t-n}, X_{t-(n-1)}, \dots, X_{t-1}) = \tilde{X}_t$$

To provide information about the rest of the cluster, each status snapshot was extended by the clusters average at that time window. When using the first three dimensions of PCA, this gave us a input vector with 24 attributes.

As a baseline to be able to benchmark the performance of more complex algorithms, we use the simplest future prediction possible: repeating the last seen state. Because the nature of our data is mostly stable, this approach yielded reasonably good base results.

To consider the standard approach to time series the ARIMAX<sup>5</sup> class of models was also evaluated.

Mainly two algorithms supported by the SparkML library [18] were employed. These would least complicate the future conversion to a production application. For initial code development and testing the **Linear Regression** model was used for its simplicity and speed. Then the same approach was applied to train a **Random Forest** model [19], which could model more elaborate functions and was the only available ensemble method supporting full parallelization. Hyper parameters for both models were selected using a grid search with 3-fold cross validation as the evaluating technique. The evaluation metric for the cross validation was root mean squared error, a standard regression evaluation metric.

The algorithms prediction was compared to the actual next status and a prediction error is computed.

$$error = \sqrt{(\tilde{X}_t - X_t)^2}$$

Information about the anomalousness of a specific data point was then extracted by comparing the prediction error to the average prediction error during that day.

$$is\_anomaly(d) = \begin{cases} True & \text{if } error(d) > mean\_error + 3 * std\_dev \\ False & \text{otherwise} \end{cases}$$

The results after applying this rule were not satisfactory, with the best being the Random Forest with a precision and recall of 0.81 and 0.38 respectively. When reviewing these results, we inspected the progression of the prediction error. The beginning of the anomaly was

<sup>4</sup>All of the used classifiers support only a one dimensional output, therefore each data dimension had its own prediction model

<sup>5</sup>AutoRegressive Integrated Moving Average with eXogeneous input

pred_time	hostname	anomaly
2019-02-09 12:00:00	anomalous.cern.ch	True
2019-02-09 12:20:00	anomalous.cern.ch	True
2019-02-09 12:40:00	anomalous.cern.ch	True
2019-02-09 13:00:00	anomalous.cern.ch	True
2019-02-09 16:00:00	host01.cern.ch	False
2019-02-09 16:00:00	host09.cern.ch	False
2019-02-09 16:00:00	anomalous.cern.ch	True

Table 3: All anomalies marked by applying the three-sigma rule to the first dimension error of the Random Forest regressor on a day with a service stop anomaly. The anomaly started at 12:00 and ended at 16:00.

usually denoted with a peak, but so was the end. Moreover most of the false positives (for the more advanced algorithms) came from data points representing non-anomalous servers in the same time the anomaly ended (see Table 3). Both the end peak and the false positives might be explained by the application at the anomalous host synchronizing itself with the rest of the cluster: the anomalous host has to consume all the data that arrived during the time its service was not operational and the data has to be streamed from other members of the cluster.

Table 3 would suggest, that if there is an anomaly, there should be more than one high errors in a row. Scanning for two notifications in a row lowered the number of notifications significantly, but most importantly filtered all of the false positives. To counter act the drop of recall, the  $3 * std\_dev$  threshold was relaxed to a single standard deviation. The final results are listed in Table 4. Note that the results can be improved by not applying the PCA dimensionality reduction. This however resulted in the need to train 35 models, because the algorithms implemented in SparkML support only a single dimensional output. If the full hyper parameter selection would to be done, the computational demand would be too high. Instead the hyperparameters were fixed after running cross-validation once on a single set of training data.

model	recall
Random Forest - nopca	0.78
Random Forest	0.70
Linear Regression	0.56
ARIMAX	0.56
BASELINE	0.48

Table 4: Summary results for the whole system with notification filtering and  $1 * std\_dev$  as the anomaly classification prediction error threshold. Notifying on the anomaly at the time of its end is not counted as valid notification.

## 5 Conclusions and Future Work

This project set to explore the possibility of using advanced data analytics to detect anomalies implying erratic behavior of nodes in a cluster running a distributed application. We first built

an infrastructure to collect, pre-process and store os-level monitoring data using the MONIT infrastructure at CERN. With continuous data gathering in place, we selected a cluster of ten nodes running Kafka, a distributed streaming application. We then began altering the server performance by running other programs or stopping the Kafka application service to induce anomalous behavior.

After testing multiple different approaches, best results were yielded by employing the Random Forest algorithm to predict the state of a server based on its short history augmented by the clusters average. To improve the usability for the final user a simple heuristic eradicating the false positives while keeping a high recall rate was employed. The final system with a Random Forest regressor, PCA dimensionality reduction and anomaly notification filtering heuristic yielded a recall of 0.7 and precision of 100%, with all the notifications were issued shortly after the anomaly start. Recall could be improved still to 0.78 by using a model trained on data without reduced dimensionality, which however increases the complexity and also results in a delay in detecting certain types of anomalies.

Work on this project will continue by developing a production version. This will pose many challenges, including making the process of extraction of anomalies from prediction errors online. Also the problems of automatic retraining of the models have to be solved. Those involve the problem of detecting days, when the whole application is experiencing difficulties and excluding those from the training set and others.

## References

- [1] M. Hilbert, *Information Quantity* (Springer International Publishing, Cham, 2017), pp. 1–4, ISBN 978-3-319-32001-4, [https://doi.org/10.1007/978-3-319-32001-4\\_511-1](https://doi.org/10.1007/978-3-319-32001-4_511-1)
- [2] A. Aimar, A.A. Corman, P. Andrade, S. Belov, J.D. Fernandez, B.G. Bear, M. Georgiou, E. Karavakis, L. Magnoni, R.R. Ballesteros et al., *Journal of Physics: Conference Series* **898**, 092033 (2017)
- [3] M.L. Massie, B.N. Chun, D.E. Culler, *Parallel Computing* **30**, 817 (2004)
- [4] P. Jung, *Linux J.* **2009** (2009)
- [5] *Nagios, the industry standard in it infrastructure monitoring*, [www.nagios.org/projects/nagios-core/](http://www.nagios.org/projects/nagios-core/), [Online, Accessed: 2019-03-06]
- [6] A.L. Buczak, E. Guven, *IEEE Communications Surveys & Tutorials* **18**, 1153 (2015)
- [7] F. Salfner, M. Lenk, M. Malek, *ACM Comput. Surv.* **42** (2010)
- [8] X. Gu, H. Wang, *Online Anomaly Prediction for Robust Cluster Systems* (2009), pp. 1000 – 1011
- [9] X. Pan, J. Tan, S. Kavulya, R. G, P. Narasimhan, Tech. rep. (2008)
- [10] *Distributed service for collecting, aggregating and moving log data, "Apache Flume" [software]*, <https://flume.apache.org/>
- [11] N. Garg, *Apache Kafka* (Packt Publishing, 2013), iISBN: 978-1782167938
- [12] M. Zaharia, R.S. Xin, P. Wendell, T. Das, M. Armbrust, A. Dave, X. Meng, J. Rosen, S. Venkataraman, M.J. Franklin et al., *Commun. ACM* **59**, 56 (2016)
- [13] *Collectd, the system statistics collection daemon*, [www.collectd.org](http://www.collectd.org), [Online, Accessed: 2019-03-06]
- [14] H. Abdi, L.J. Williams, *WIREs Comput. Stat.* **2**, 433 (2010)
- [15] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg et al., *Journal of Machine Learning Research* **12**, 2825 (2011)

- 
- [16] B. Schölkopf, R.C. Williamson, A.J. Smola, J. Shawe-Taylor, J.C. Platt, *Support vector method for novelty detection*, in *Advances in neural information processing systems* (2000), pp. 582–588
  - [17] F.T. Liu, K.M. Ting, Z.H. Zhou, *Isolation forest*, in *2008 Eighth IEEE International Conference on Data Mining* (IEEE, 2008), pp. 413–422
  - [18] X. Meng, J. Bradley, B. Yavuz, E. Sparks, S. Venkataraman, D. Liu, J. Freeman, D. Tsai, M. Amde, S. Owen et al., *J. Mach. Learn. Res.* **17**, 1235 (2016)
  - [19] L. Breiman, *Machine Learning* **45**, 5 (2001)