

Article

Combining Log Files and Monitoring Data to Detect Anomaly Patterns in a Data Center

Laura Viola ¹, Elisabetta Ronchieri ^{1,2} , Claudia Cavallaro ³ 

¹ Department of Statistical Sciences, University of Bologna; laura.viola5@studio.unibo.it

² INFN CNAF, Bologna, Italy; elisabetta.ronchieri@cnafe.infn.it

³ Department of Mathematics and Computer Science, University of Catania; claudia.cavallaro@unict.it

* Correspondence: elisabetta.ronchieri@cnafe.infn.it; Tel.: +39 0512095072

Abstract: Context: Anomaly detection in a data center is a challenging task, having to consider different services on various resources. Current literature shows the application of artificial intelligence techniques to either log files or monitoring data: the former created by services at run time, while the latter produced by specific sensors directly on the physical or virtual machine.

Objectives: We propose a model that exploits information both in log files and monitoring data to identify patterns and detect anomalies over time.

Methods: The key idea is to use on one side natural language processing solutions to detect problems at service level, extracting words that represent anomalies. Clustering and topic modeling techniques have been used to identify patterns and group them with respect to topics. On the other side time series anomaly detection technique has been applied to sensors data in order to combine problems found in the log files with problems stored in the monitoring data.

Results: We have tested our approach on a real data center equipped with log files and monitoring data that can characterize the behaviour of physical and virtual resources in production. We have observed a correspondence between anomalies in log files and monitoring data, e.g. an increase in memory usage or machine load. The results are extremely promising.

Conclusion: Our model requires to integrate site administrators' expertise in order to consider all critical scenario in the data center and understand results properly.

Keywords: log analysis; monitoring data; anomaly detection; natural language processing; topic modeling; clustering technique; time series anomaly detection

1. Introduction

A data center usually handles a huge amount of different resources, where a plethora of services can run, collecting information about computer systems and machine states. It is quite common to have services that store their events in specific files, known as log files, allowing site administrators to understand their current functioning and react proactively in case of issues. Furthermore, data center monitors the physical and virtual machines, tracing different devices readings (such as CPU load, memory, disk space and so on), information about network traffic and bandwidth usage, in order to send alerts to administrators if critical situations happen. Services produce log files, while monitoring sensors measure different metrics with respect to the resource level they refer to.

Logs are created by a service and contain semi-structured texts that are appended to a file with the .log extension. These files grow in size and can become very large, but they tend to include similar texts over time that may cover various aspects of each service, such as warning, error, information and so on. This leads to develop solutions that automate the processing of logs and support system administrators while analyze systems' health. Log files usually do not contain the same type of information: *syslog* event logs a system activity, while *cron* event logs the CRON entries that show up in *syslog*. Each file tends to describe a partial view of the whole machine. The stored information can contain: the

time and date of specific event to log exactly what happened; the process name and process identifier; the machine hostname and its internet protocol address. The services can use different key words to express normal or erroneous behaviour.

Monitoring data are created by sensors that run on the machine, containing e.g. a timestamp expressed in epoch time, metric's name, metric's value, and the machine hostname. These measurements cover all the running state of the machine, therefore site administrators can decide to store them with different frequency. An alarming email can be sent to the site administrators when a data center-based threshold is exceeded, such as when the used memory of a machine has reached 95% of the total memory.

Some studies have proposed approaches to handle the error problems from manual operation to automated operation [1]. They define pipelines that include the transformation of log files into a more readable format understandable by analysis tools, such as *.csv* and *.json*, the classification of observations with respect to the threshold values, and the extraction of any suspicious information, like the *anomaly* term. Quite often it is not feasible just selecting the file and retyping the file extension as *.csv* or *.json*, because the transformed file could be wrongly reformatted, e.g. containing multiple rows of heading variables. Furthermore, the file usually contains daily service information, so the number of data can be over 60 – 120 K rows, penalizing the usage of some analysis tools. Before starting any analysis it is essential to decide, which variables have to be included in the resultant data sets. Spreading the data across multiple columns is another aspect to consider in order to organise your data set into a manageable format. The resultant files can be used to identify trends and unusual activities that are beneficial for both short- and long-term data center management.

In our previous work [2] we have just considered log files to identify anomaly detection patterns by using natural language processing (NLP), autoencoder and invariant mining techniques. In this new study we have combined the knowledge present in monitoring information and log files in order to improve our analysis. We have decided to aggregate data at machine level by considering all the log files for the running services and the measurements of machine sensors in order to determine the variation of values nearby an anomaly.

Artificial intelligence techniques are promising solutions to automatically identify anomaly detection patterns and predict failures in a machine. In this paper, we describe how we have combined the results obtained after the application of NLP techniques and time series anomaly detection. Log files contain a considerable amount of texts, therefore we have used modern NLP methods [3], [4] to easily label the log entries and contribute to defining administrators' operations. Tools, like simple log file clustering tool [5], have the disadvantage of being based on searching only for the most frequent types of messages in the log file, neglecting the less frequent ones. For anomaly detection, however, it may be necessary to find rare types of messages and interesting patterns, so word n-grams solution has been also applied. The clustering methods in the log core analysis phase would lead to the risk of losing significant results on the anomalies to be considered individually: for this reason we have used topic modelling techniques. In this study for each machine we have also tried to overlap monitoring information with log files anomalies to determine when problems happened.

The proposed approach is repeatable in other contexts and domains. With minimum setup effort and the usage of artificial intelligence and statistical tools, it is possible to automatically extract relevant information about machine state. Through experiments, we illustrate the potential benefits of our approach by answering our research questions.

The reminder of this paper is as follows. Section 2 introduces the rational of our approach. Section 3 talks about feature creation, while Section 4 details machine learning techniques. Section 5 concludes the paper discussing the presented work.

2. Methodology for Anomaly Detection

The main objective of this study has been the development and validation of an anomaly detection model by considering a data center use case. A generalized workflow is shown in Figure 1 that schematizes the steps that have been undertaken to implement this model. Once data has been collected both at machine level with log files and monitoring data, we have dedicated effort to the following phases: data preprocessing, trends and conditions identification, and data transformation that have allowed us to build a set of datasets to train and build our anomaly detection models.

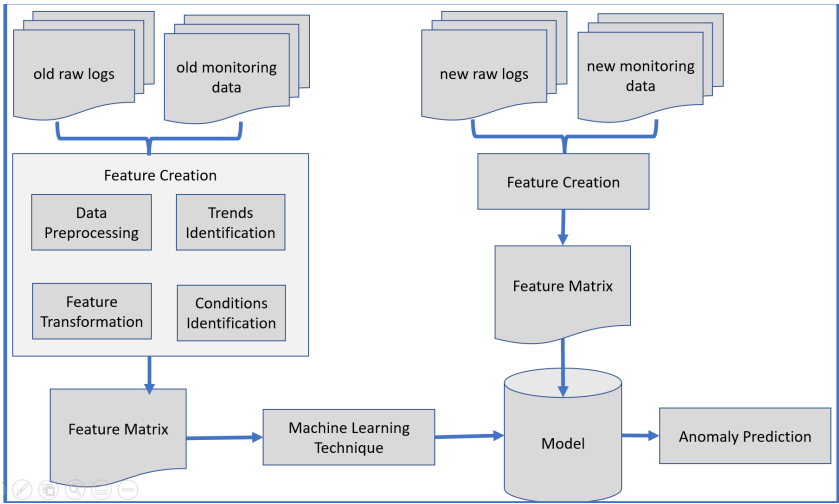


Figure 1. General methodology overview.

Source data Our work has started with the datasets selection. The log files are related to a set of services running on machines at INFN Tier-1 data center [6] that are used by the large hadron collider experiments [7]. Figure 2 shows the groups which make up the TIER-1 infrastructure and cooperate with each other.




Figure 2. The groups making up the Tier-1 infrastructure at CNAF-INFN.

Each log mainly belongs to Linux system services, such as the software utility *crond*, the free and open-source main transfer agent postfix and the standard for message logging *syslog*. Table 1 summarizes the first 30 filenames according to their frequency, that is the number of times a value of an unique filename occurs. The suffix-type frequencies of the various available files are summarized as follows: *.gz* 10,869, *.log* 3,562,759, *.manifest* 5, *.meta* 6, *.pending* 1, *.txt* 2. In this work we have considered the same log files used in the previous study [2].

Table 1. The top 30 log files per frequency.

filename	frequency	filename	frequency	filename	frequency
sudo.log	378,781	systemd.log	107,700	userhelper.log	21,380
puppet-agent.log	368,530	mmfs.log	72,620	nsdcd.log	20,544
run-parts.log	365,734	rsyslogd.log	70,210	neutron_linuxbridge.log	8,572
crontab.log	348,896	kernel.log	65,938	runuser.log	6,859
crond.log	347,708	logrotate.log	62,531	cvmfs_x509_validator.log	6,031
sshd.log	303,919	syslog.log	47,330	cvmfs_x509_helper.log	5,399
anacron.log	287,419	yum.log	43,301	srp_daemon.log	4,938
postfix.log	175,558	fusinv-agent.log	42,125	edg-mkgridmap.log	4,083
auditd.log	120,473	root.log	37,345	libvirt.log	33,28
smartd.log	109,441	gpfs.log	31,000	dbus.log	3,301

Each machine can run different services and is usually characterized by a certain amount of memory and disk space. To check each machine status, the data center exploits a monitoring system that is able to measure various metrics, show them through the *graphana* service and store them through an *influxd* service [8]. On the basis of site administrators' feedback, in this study we have considered a subset of metrics grouped in three categories as shown in Table 2: load average that refers to the average system load on a server for a specific period of time; memory that provides information about the consumption of memory; iostat average that captures the status of the device. Load averages are usually three numbers, showing the average load in the last minute, in the last five minutes, and the last fifteen minutes: if the one minute average is higher than the five or fifteen minute averages, then load is increasing; if the five minute average is lower than the five or fifteen minute averages, then load is decreasing. The memory category provides information about the normal and swap memory: when the swap memory is used, then it means that the normal memory is full. The used memory is one of the memory metrics for which an alarming email is automatically sent to site administrators when its value exceeds 95% of total memory; iostat averages give input and output devices utilization.

Table 2. Monitoring metrics.

category	metrics	category	metrics	category	metrics
Linux load average	1 minute	Memory	swap free	iostat average	cpu pct iowait
Linux load average	5 minutes	Memory	swap total	iostat average	cpu pct nice
Linux load average	15 minutes	Memory	swap used	iostat average	cpu pct steal
		Memory	available	iostat average	cpu pct system
		Memory	buffers	iostat average	cpu pct user
		Memory	cached	iostat average	cpu pct idle
		Memory	dirty		
		Memory	free		
		Memory	used		
		Memory	total		

Log files Each of these log files contains a different amount of lines. They contain numerical and textual data that describe system states and run time information. Each log entry includes a message that contains a natural-language text (i.e. a list of words) describing some events. Logs are generally generated by logging statements inserted, either by software developers in source code or by system administrators in configuration files, to record particular events and software behaviour. Each log entry in the log file represents a specific event. Figures 3 and 4 show two different log entry samples composed

of a log header and a log message: the former is generally composed of a timestamp, a custom-configuration information (such as the hostname in Figure 3, where the service runs and a log level verbosity in Figure 4) and the name of the service the message is associated to; the latter is just the message that contains information of the logged event.

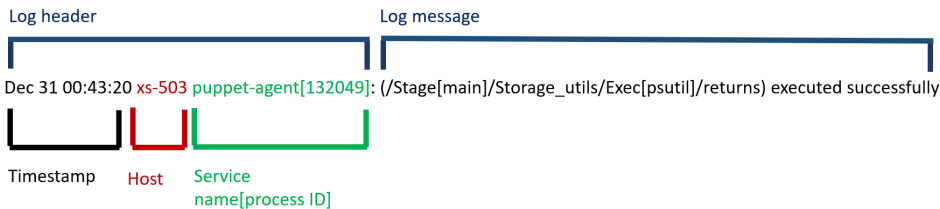


Figure 3. A log entry sample from the puppet-agent service.

Figures 5 and 6 show two examples of the log message of a log entry, characterized by a natural-language text which interpretation is difficult because there is not an official standard defining the message format. The text is usually composed of different fields called dynamic and static: a dynamic field is a string or a set of strings that are assigned at run time; a static field does not change during events. Such fields can be delimited by different separators, such as a comma, a white space, a parenthesis. Logging practice is scarcely well documented. This activity mainly depends on human expertise [9]. They often have to analyze a large volume of information that may be unrelated to the problematic scenarios and lead to overwhelming messages [10].

Monitoring metric files Each of these files contains a different amount of lines. They contain numerical and textual data that describe machine state for a specific metric over time. Each file includes the hostname of the machine, epoch time and measurements. They are produced by using the *influx* client that queries a specific database (defined for a set of machines) to extract a given metric in a certain period of time, and dumps values into a *.csv* file. Below it is reported an example of the query performed by command line:

```
$ influx -host=<service hostname> -port=<port number> \
-username="<user>" -password="<password>" \
-database="<database>" \
-execute="SELECT * FROM \"all_data\".\"load_avg.five\" \
WHERE time > now() - <number of days> GROUP BY \"host\" \" \" \
-format=csv > one_week.csv
```

3. Feature Creation

Data preprocessing During this phase activities are performed on log files and monitoring files.

The log files change format and turn into *.csv* files. In addition the following variables are added to each file: date, time, timestamp, hostname, internet protocol (ip) address, service name, process identifier (id), component name, message (msg). The hostname and ip couple are not always both available, especially when the service runs on a virtual machine. Each file is related to a particular service that runs on a well-known machine in




Figure 4. A log entry sample from the puppet-agent service.

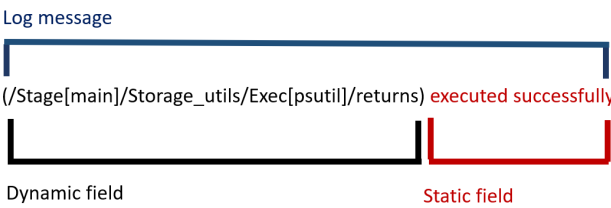


Figure 5. A log message fields with just one dynamic field and static field.




Figure 6. A log message fields with a sequence of dynamic and static fields.

the data center. Its location is got by a local database and included into the resulting file. During this phase we have tackled some site administration peculiarities: the same service called in lower or capital letters; the process identifier included in the service logging file; the service name included (or not) in the log message; the process identifier included (or not) in the log message; the logging filename included typo error. We have also identified some typo errors in the log messages. Before performing any cleaning operations, we have excluded meaningless services’ log files, especially those with a small number of events. In the remaining logs, the following changes have been applied in the message field: the removal of unwanted texts, such as punctuation, non-alphabets, and any other kind of characters that are not part of the language by involving regular expressions; the exclusion of non-English characters; the cancellation of stopwords, that is frequent general words (like *of, are, the, it, is*) with a low meaning. For stopwords, we have decided to keep negative terms, and other words that may refer to a problem, such as *up, again, too, ok, out, yet, and more*.

The monitoring files are created with multiple heading, therefore we have removed some rows in the files before performing any analysis.

According to the amount and types of services, in this phase we have started to trace the types of log events, such as abort, fail and invalid, and to identify anomaly key terms that can be used to classify the reason of the problems in the service. This part of the study has been applied to all the set of files examined contributing to better understand variation in the machine status. However, Table 3 summarizes few examples of message lines that describe a wrong service behaviour.

Table 3. Examples of message lines for the crond log file.

log event msg	log event type	anomaly key term
.. reset error counters	error	reset
.. failed create session connection time out	fail	connection

Identification of trends and anomaly conditions During this phase monitoring data have been used. Some metrics show a noisy behavior (as shown in Figure 7), making the graphs unreadable and difficult extract essential information about trends and large-scale deviations. At this moment of the study we have not applied any smooth functions to better identify trends in the metrics.

Features Transformation Data transformation includes the creation of a new dataset that includes a matrix, whose dimension and values change according to the machine learning technique used to build the anomaly detection model. The resulting dataset can contain either binary data or numerical data. In case of binary data there is a numerical value of 1 for features that are present in the log event and a numerical value of 0 otherwise.




Figure 7. Comparison of load average metrics.

In case of NLP usage, the tokenization phase (see Figure 4) determines the element of the matrix: the message string is split up in single words that can code relevant anomalies. Once transformed the message events into a matrix of features, in case of problem with large dimension, it is possible to apply rules to reduce data, e.g., according to the uniqueness of the messages. For NLP, low frequency words or words that are not important for the meaning of the anomalies are filtered out.

Table 4. Examples of message strings split up in single words for the crond log file.

log event msg	..	error	..	fail	..	time	..
.. reset error counters	..	1	..	0	..	0	..
.. failed create session connection time out	..	0	..	1	..	1	..

During the tokenization phase, we have developed a log parsing method that uses word n -gram techniques to identify anomaly features. The n -gram dictionaries have been built with $n = 1,2,3$. Figure 8 shows the first (and last) thirty subsequences of length 3 from *auditd* log messages sequences [11].

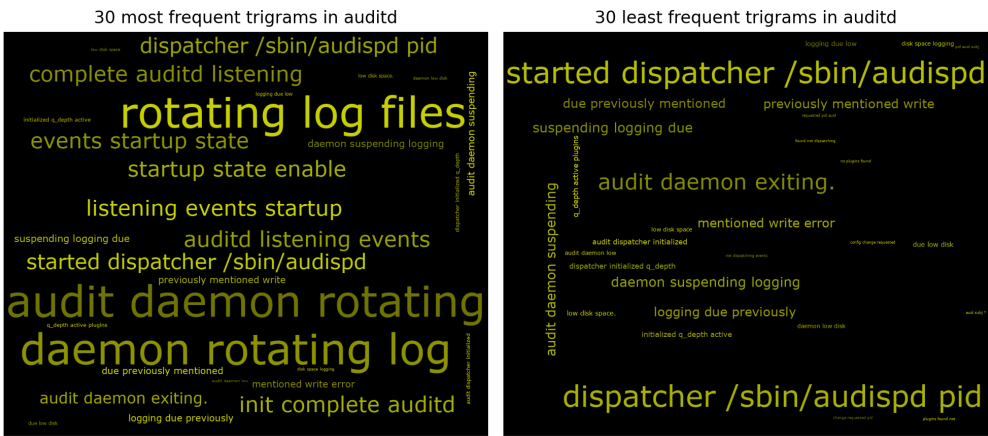


Figure 8. Word cloud for the *auditd* service.

The NLP solutions allow us to identify words (programmed by developers) that refer to the semantic area of the anomaly and present in the log files. The following anomalies have been determined: *abort, aborted, aborting, abortion, alert, cannot, can't, couldn't, deprecated, deprecate, disabled, error, exception, fail, fails, failure, failed, failing, fatality, fatal,*

invalid, impossible, huped, misconfiguration, problem, sslerror, suppressed, suppressing, suspended, suspend, suspending, suspension, stopped, stopping, stop, unable, unsupported, warning, warn, warned. They can be used to label each observation with 1 or 0 with respect to the presence of anomalies. According to this rule we have counted almost 27 thousand anomalies in all the log files that usually can contain just 1 or just 0 or both values 1 and 0. When the log file contains both 1 and 0 as anomaly values, the observations are typically unbalanced.

Figures 9, 10 and 11 show four services that run on a specific machine, called *tb-cloud01-net*. The colored vertical lines represent anomalies identified in the log files of the selected services. Figures 9 and 10 show the time series for the load averages at 15 minutes and the memory used respectively. Figure 10 highlights the variation in the memory used after an intervention in one of the listed services. Figure 11 shows as time series the iostat average for the *cpu_pct_iostat* metric.

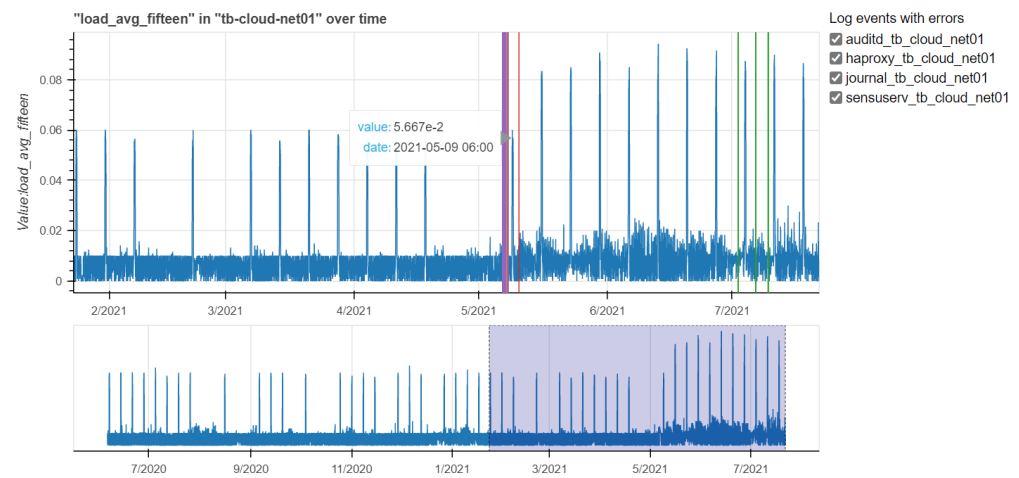


Figure 9. Overlapping log files and load averages at 15 minutes for the *tb-cloud01-net* machine.




Figure 10. Overlapping log files and memory used values for the *tb-cloud01-net* machine.

The term frequency and inverse document frequency (TF-IDF) technique has also been considered to extract features and group together observations [12]. TF-IDF is used for word count, where TF stands for term frequency matrix and it measures the association of a word with respect to a given document; and IDF stands for inverse document frequency and it represents the importance of one word. If a word appears in many log files, the importance of this word will be decreased. This approach applied alone could not be necessarily capable of providing useful information for the identification of anomalies, because the relevance itself does not take into account the semantic meaning of the word. Relevant or irrelevant words may not coincide with a normal or abnormal message. In




Figure 11. Overlapping log files and iostat average cpu pct_iowait values for the *tb-cloud01-net* machine.

addition, the feature matrix may have high dimensionality as it will have a number of columns equal to the number of unique words contained in the messages.

4. Machine Learning Techniques

In our previous work we have applied word2vec, autoencoder and invariant mining techniques [2]. The autoencoder was used to to learn a more efficient representation of data while minimizing the corresponding error. The invariant mining was used to collect the frequent patterns of error messages that generate a problem in a physical or virtual machine at INFN Tier-1 data center.

In this work we have mainly experimented topic modelling technique to find latent topics in mostly unstructured collections of log messages, identify word structures and anomalies that belong to specific topic. Topic modelling assumes that documents are a mixture of topics, while topics are a distribution of words [13]. Figures 12 and 13 show how terms are distributed in the different topics for the *screen* and *virtlogd* services. In both Figures we can observe that during the preprocessing activity we have cleaned log messages keeping meaningful information, such as the function name, the name of the machine, users' names (that for privacy reason we have omitted), and so on.

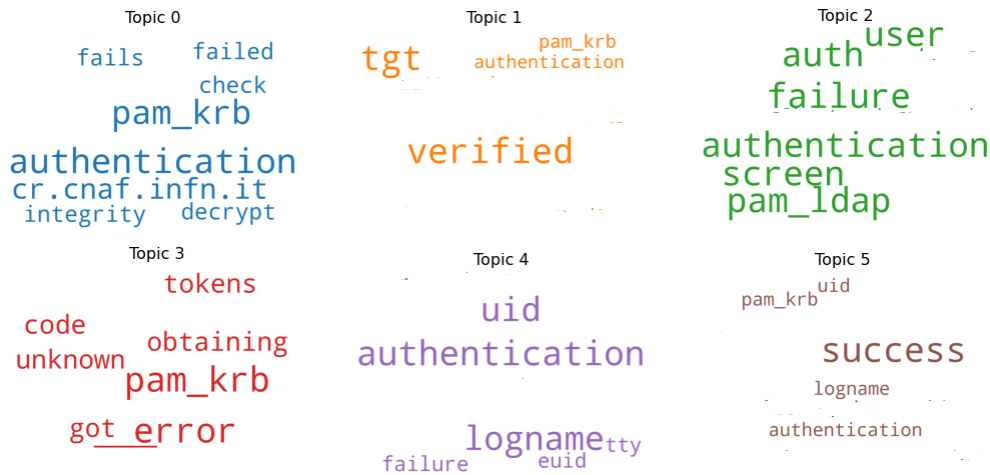


Figure 12. Identified topics for the *screen* service.

Furthermore, we have developed a clustering algorithm that aims at grouping log messages according to the words or sequences that compose them. Through this procedure



Figure 13. Identified topics for the *virtlogd* service.

it is possible to distinguish the various types of anomalous messages included in the log file. The messages that do not report anomalies are grouped in a single cluster with label 0. Figures 14 and 15 show how terms are distributed in the different topics for the *screen* and *virtlogd* services.



Figure 14. Grouping messages for the *screen* service.



Figure 15. Grouping messages for the *virtlogd* service.

With respect to time series, we have applied JumpStarter solution [14] to compute anomaly score. JumpStarter is a multivariate time series anomaly detection based on compressed sensing to compute anomaly score and use it to label observations. Figure 16 highlights with red dots on the memory used time series, anomalies computed with JumpStarter. The results of this technique requires a further investigation.




Figure 16. JumpStarter results for the memory used metric on the *tb-cloud01-net* machine.

5. Conclusions

In this paper we have discussed the specific findings from our study. The results include the data extraction, the exploration of server logs on physical and virtual resources, the monitoring data. They have been obtained applying natural language processing solutions and clustering technique.

The data extraction has been challenging, having access to a huge amount of data. It was hard to identify the proper use cases selecting log files and monitoring data on a given machine. The results require a further investigation with site administrators.

Acknowledgments: The authors gratefully acknowledge INFN Tier-1 site administrators who provide log files.

References

1. Farshchi, M.; Schneider, J.G.; Weber, I.; Grundy, J. Experience report: Anomaly detection of cloud application operations using log and cloud metric correlation analysis. *IEEE Transactions on Software Engineering*. IEEE, 2015. doi:10.1109/ISSRE.2015.7381796.
2. Cavallaro, C.; Ronchieri, E. Identifying anomaly detection patterns from log files: a dynamic approach. *Computational Science and Its Applications – ICCSA 2021*; Gervasi, O.; Murgante, B.; Misra, S.; Garau, C.; Blečić, I.; Taniar, D.; Apduhan, B.O.; Rocha, A.M.A.; Tarantino, E.; Torre, C.M., Eds. Springer International Publishing, 2021, pp. 517–532.
3. Bertero, C.; Roy, M.; Sauvinaud, C.; Trédan, G. Experience Report: Log Mining using Natural Language Processing and Application to Anomaly Detection. *28th International Symposium on Software Reliability Engineering (ISSRE 2017)*; , 2017; p. 10p.
4. Layer, L.; Abercrombie, D.R.; Bakhshiansohi, H.; Adelman-McCarthy, J.; Agarwal, S.; Hernandez, A.V.; Si, W.; Vlimant, J.R. Automatic log analysis with NLP for the CMS workflow handling. *24th International Conference on Computing in High Energy and Nuclear Physics (CHEP 2019)*, 2020, p. 7p. doi:https://doi.org/10.1051/epjconf/202024503006.
5. Vaarandi, R. Mining event logs with SLCT and LogHound. *NOMS 2008 - 2008 IEEE Network Operations and Management Symposium*. IEEE, 2008. doi:10.1109/noms.2008.4575281.
6. dell'Agnello, L.; Boccali, T.; Cesini, D.; Chiarelli, L.; Chierici, A.; Dal Pra, S.; Girolamo, D.; Falabella, A.; Fattibene, E.; Maron, G.; et al. INFN Tier-1: a distributed site. *EPJ Web of Conferences* **2019**, 214, 08002. doi:10.1051/epjconf/201921408002.
7. Breskin, R.V.A. *The CERN Large Hadron Collider: Accelerator And Experiments Volume 2: CMS, LHCb, LHCf, And Totem*; CERN, 2009.
8. Bovina, S.; Michelotto, D. The evolution of monitoring system: the INFN-CNAF case study. *J. Phys.: Conf. Ser.* **2017**, 898, 092029.
9. He, P.; Chen, Z.; He, S.; Lyu, M.R. Characterizing the natural language descriptions in software logging statements. *Proceedings of the 33rd ACM/IEEE international conference on automated software engineering ASE*, 2018, pp. 178–189. doi:10.1145/3238147.3238193.
10. Chen, B.; Jiang, Z.M.J. Characterizing and detecting anti-patterns in the logging code. *Proceedings of the IEEE/ACM 39th International Conference on Software Engineering (ICSE)*. IEEE Press, 2017, pp. 71–81. doi:10.1109/ICSE.2017.15.
11. Dai, H.; Li, H.; Chen, C.S.; Shang, W.; Chen, T.H. Logram: Efficient Log Parsing Using n-Gram Dictionaries. *IEEE Transactions on Software Engineering* **2020**, pp. 1–1. doi:10.1109/tse.2020.3007554.
12. Sandhu, A.; Mohammed, S. Detecting Anomalies in Logs by Combining NLP features with Embedding or TF-IDF. *TechRxiv* **2022**.
13. Blei, D.M.; Carin, L.; Dunson, D. Probabilistic topic models. *IEEE Signal Processing Magazine* **2010**, 55, 77–84.
14. Ma, M.; Zhang, S.; Chen, J.; Xu, J.; Li, H.; Lin, Y.; Nie, X.; Zhio, B.; Wang, Y.; Pei, D. Jump-Starting Multivariate Time Series Anomaly Detection for Online Service Systems. *Proceedings of the USENIX Annual Technical Conference - USENIX ATC 21*, 2021, pp. 413–426.