PROCEEDINGS OF SCIENCE

Collection and harmonization of system logs and prototypal Analytics services with the Elastic (ELK) suite at the INFN-CNAF computing centre

T. Diotalevi*
University of Bologna and INFN, Italy
E-mail: tommaso.diotalevi@studio.unibo.it

A. Falabella, B.Martelli, D. Michelotto, L. Morganti

INFN-CNAF, Italy

D. Bonacorsi, L. Giommi, S. Rossi Tisbeni

University of Bologna and INFN, Italy

The distributed Grid infrastructure for High Energy Physics experiments at the Large Hadron Collider (LHC) in Geneva comprises a set of computing centres, spread all over the world, as part of the Worldwide LHC Computing Grid (WLCG). In Italy, the Tier-1 functionalities are served by the INFN-CNAF data center, which provides also computing and storage resources to more than twenty non-LHC experiments. For this reason, a high amount of logs are collected each day from various sources, which are highly heterogeneous and difficult to harmonize. In this contribution, a working implementation of a system that collects, parses and displays the log information from CNAF data sources and the investigation of a Machine Learning based predictive maintenance system, is presented.

International Symposium on Grids & Clouds 2019, ISGC2019
31st March - 5th April, 2019
Academia Sinica, Taipei, Taiwan

*Speaker.

PoS(ISGC2019)027

1. Introduction

The Worldwide LHC Computing Grid (WLCG) is a global computing infrastructure with the mission of providing computing resources to store, distribute and analyse the data generated by the Large Hadron Collider (LHC), making the data equally available to all countries, regardless of their physical location. Computing infrastructures through which Grid services operate are hierarchically classified into *Tiers* according to the kind of services they offer.
Since 2003, the Italian Tier-1 for the High Energy Physics experiments is hosted at the Bologna INFN-CNAF Data Center [1], providing the resources, support and services needed for data storage and distribution, data processing and analysis as well as the production of simulated ('Monte Carlo') data, representing also a key computing facility for many non-LHC communities, making it one of the most important centers for distributed computing in Italy.
A key challenge, therefore, is the modernization of the center to be able to cope with the increasing flux of data expected in the near future e.g. with the new phase of operations of the High-Luminosity LHC [2]. These high-level standards of operation require a continuous work towards a full understanding of service behaviours and a constant seek for higher level of automation and optimization. Data centers worldwide are now witnessing the use of Artificial Intelligence (AI) solutions pushing them into a new phase, in which tasks traditionally managed by operators could be more efficiently managed by human-supervised machine decisions.
Besides, CNAF collects a large amount of logs every day from various sources, which are highly heterogeneous and difficult to harmonize: such log data is archived but almost never used, except for specific internal debugging and monitoring operations. This contribution, together with the implementation of a system that collects, parses and displays the log information from CNAF data sources, undertakes the investigation of a Machine-Learning-based predictive maintenance system, moving away from a preventive replacement of equipment, which is highly expensive and far from optimal efficiency.

2. The Elastic Stack

In order to create an indexed system with structured information from CNAF system logs, the Elastic Stack [3] has been chosen. The Elastic Stack is the union of three open source project developed by the Elastic company [4]: Elasticsearch [5], Logstash [6] and Kibana [7]. As shown in Figure 1, logs coming from different services and different nodes are firstly collected into a single stream by a platform called Beats [8], then they are aggregated and processed by Logstash, indexed and stored via the Elasticsearch search engine and finally visualized by the Kibana user interface. In the following, a brief description of these components is given.

2.1 Beats

Beats is the platform that handles the shipment of data coming from different sources. Different shippers exist depending on the kind of data that needs to be moved: Filebeat for log files, Metricbeat for metrics, etc. The installation of such services occurs directly on the servers that contain data and, for this reason, is fast and lightweight. In case the information collected from Beats is unstructured, data is shipped to Logstash for further transformation and parsing.

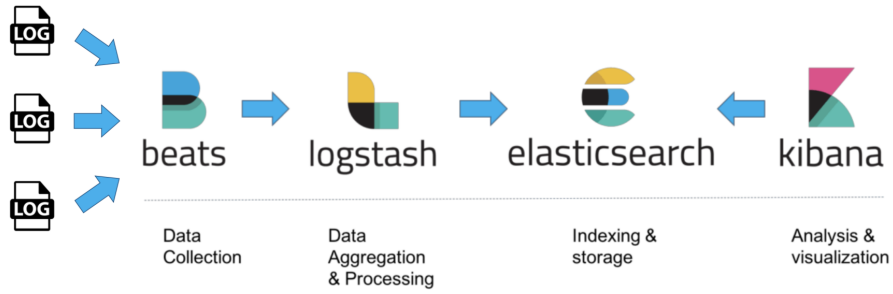

Figure 1: Workflow of the Elastic Stack log ingestion and processing.

2.2 Logstash

Logstash is an open source, server-side data processing pipeline that ingests data from a multitude of sources simultaneously, transforms them, and then sends them to a central processing engine (e.g. Elasticsearch). Since data is often scattered across many systems in many formats, Logstash allows to ingest logs, metrics, web applications, etc. in a continuous, streaming fashion.
As data travel from source to source, as shown in Figure 2, Logstash filters parse each event, identify named fields to build structure, and transform them to converge on a common format for easier, accelerated analysis. In this regard, Logstash dynamically transforms and prepares your data regardless of its format or complexity using different filters (e.g. grok filters [9] to derive structure from unstructured data, or geo filters to locate geographical coordinates from IP addresses). Data is then ready to be process by Elasticsearch, which is the embedded service provided by Elastic (this is not the only existing solution, e.g. Apache Kafka or MongoDB for logs data and InfluxDB for metrics data).

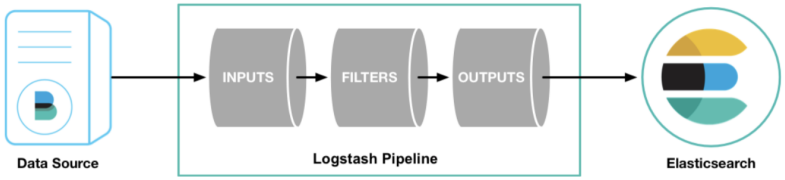

Figure 2: Logstash pipeline.

2.3 Elasticsearch

Elasticsearch is the central component of the Elastic Stack with a distributed search and analytics engine capable of solving a growing number of use cases. Using a RESTful API architecture [10], it allows to perform and combine many types of searches - structured, unstructured, geo, metric - using a Lucene querying syntax [11]. Created by Shay Banon in 2004, it was developed to provide a scalable search engine and a distributed indexing system, capable of replicating and routing data into different shards spread across the entire cluster.
Elasticsearch, currently, has been adopted and is running in several companies and organizations like Mozilla, GitHub, Netflix, Wikimedia, and in scientific endeavours (e.g. CERN).

2.4 Kibana

Kibana is an open source data visualization plugin for Elasticsearch. It provides visualization of the indexed content of an Elasticsearch cluster. Using the embedded user interface, users can create bar, line and scatter plots, or pie charts and maps based on large volumes of data. Kibana also provides several additional features, such as developer tools for advanced interactions with the Elastic Stack e.g. a console able to interact with the REST API of Elasticsearch in a cURL-like syntax. Some features are not included in the basic licence given by Elastic, but require a premium licence purchase: this allows users to explore anomalies in time series data with unsupervised Machine Learning features, which are described in more detail in the following sections.

3. Managing storage resources on Grid

The WLGC main mission, since the very beginning, is the capability to share data and resources over a wide-area network across several organizational domains. The current Grid infrastructure, therefore, has to face multiple heterogeneous storage and file system in order to manage, replicate and access files in a distributed system. Moreover, high performance disk-storage solutions have become increasingly important to deal with the large I/O throughput required by the High Energy Physics (HEP) community for both data analysis and Monte Carlo simulations.
Given such requirements, the HEP Grid community has developed and implemented the *Storage Resource Manager* (SRM) interface [12]. From one side a client can request the SRM to reserve space and manage files and directories; from the other side, SRM is able to dynamically decide which files to keep inside the storage space and which to remove e.g. to free some disk space. SRM relies on some basic concepts on space and files. A file can be volatile, permanent or durable. A *volatile* file is temporary with a lifetime associated, deleted by a garbage collector of the SRM. A *permanent* file can only be removed by the owner. A *durable* file has a lifetime associated with it, but can be removed by both the garbage collector and the owner.
In Grid, a file is identified in several ways: the *Logical File Name* (LFN) is the user friendly name, often user-defined. The *Storage* URL (SURL) points to the physical stored version of a file on a grid storage element. The *Transport* URL (TURL) is an URL issued by SRM for a SURL which can be used to retrieve or store data, e.g. with GridFTP. Finally the *file catalog* is the grid middleware server which maps LFNs to SURLs i.e. tracks the real physical locations of logical file names.
Some implementation of SRM service are dedicated to specific storage system e.g. SRM on CASTOR at CERN [13] and SRM on dCache [14].

3.1 The StoRM Storage Resource Manager

StoRM is a Storage Resource Manager that relies on a parallel file system or a standard Posix file system backend i.e. from high performance parallel file systems like GPFS (from IBM) [15]. StoRM provides advanced SRM functionalities to dynamically manage space and files according to the user requirements, specifying the desired lifetime and allowing for advance space reservation and a different quality of service provided by the underlying storage system. StoRM

takes advantage from the file system security mechanisms to ensure access to data in a secure way. Figure 3 shows the role of StoRM inside a site.
StoRM has a multilayer architecture made by two main components: the *front-end* (written in C/C++), and the *back-end* (written in Java).
The front-end(s) exposes the SRM web service interface, manages user authentication and stores the requests data into a database. The back-end, instead, is the core of StoRM service, taking care of executing all the synchronous and asynchronous SRM functionalities. It is responsible of the management of file and space metadata, it enforces authorization permissions on files and it interacts with other Grid services. Moreover, the back-end is able to use advanced functionalities provided by the file system (for example by GPFS and XFS) to accomplish space reservation requests.
Currently, StoRM is adopted in the context of WLCG infrastructure in various data centers, including the Italian Tier 1 at the CNAF-INFN institute. In order to satisfy the high availability and scalability requirements coming from the HEP community, StoRM can be deployed in a clustered configuration, with multiple instances of front-end and back-end services and with a dedicated DBMS installation: multiple front-end instances can be deployed on separate machines and configured to work on the same database and the same back-end service. In Figure 4, a sketch of this architecture is shown.

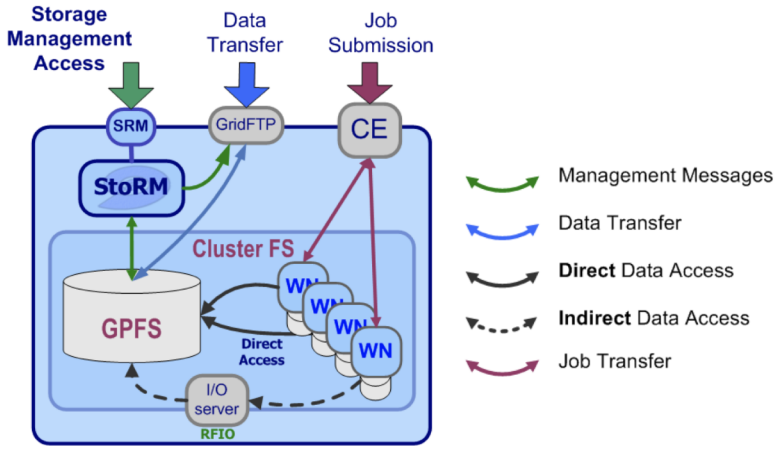

Figure 3: StoRM role in a cluster site.

4. StoRM Logging information

StoRM services produce a huge quantity of log information, with different levels of verbosity that are stored in different files.

4.1 StoRM Frontend Logging

The Frontend (FE) logs information on the service status and SRM requests received and managed by the process. Different levels of logging are supported such as ERROR, WARNING, INFO, DEBUG that can be set from dedicated parameters in a configuration file.

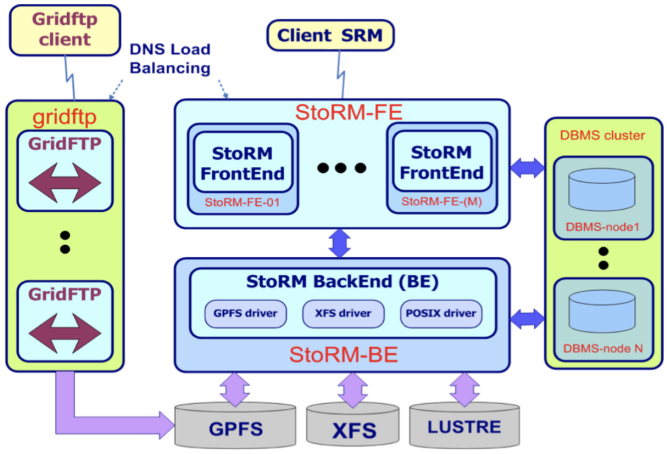

Figure 4: Schematic representation of the StoRM architecture.

storm-frontend-server.log

The Frontend log file named *storm-frontend-server.log* is placed in the /var/log/storm directory. At start-up time, the FE prints here the whole set of configuration parameters, this can be useful to check desired values. When a new SRM request is managed, the FE logs information about the user subject name, community group and role identifiers as well as several other details of the request parameters, as shown in Figure 5.

```
03/30 23:59:59.736 Thread 50 -  INFO [4306                              5a88]: Request 'PTG status' from Client IP='2001:6
30:58:        :b7d6' Client DN='/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=d      n/CN=5    7/CN=Robot: ATLAS Data Ma
nagement' # Requested token 'a5bd6                                  1f7'
```

Figure 5: *storm-frontend-server* log line example.

monitoring.log

The monitoring service, if enabled, provides information about the operations executed within a certain time frame, writing them to file as shown in Figure 6. This amount of time (called Monitoring Round) is configurable and its default value is 1 minute. At each Monitoring Round, a single row is printed in the log. This row reports both information about requests that have been performed in the last Monitoring Round and information considering the whole FE execution time (Aggregate Monitoring). Information reported are generated from both Synchronous (operations that return the control to the client when the request has been executed) and Asynchronous requests (that return the control as soon as the request has been accepted by the system) and tell: how many requests have been performed in the last Monitoring Round, how many of them were successful, how many failed, how many produced an error, the average execution time, the minimum execution time and the maximum execution time.

4.2 StoRM Backend Logging

The Backend (BE) log files provide information on the execution process of all SRM requests. Backend logging is based on *logback* framework. Logback provides a way to set the level of

```
03/30 23:59:47 : [#  4893 lifetime=81:33:01] S [OK:1615253,F:221593,E:0,m:0.000,M:104.146,Avg:0.025] A [
OK:192994,F:0,E:0,m:0.004,M:3.854,Avg:0.012] Last:(S [OK:280,F:46,E:0,m:0.000,M:0.108] A [OK:31,F:0,E:0,
m:0.007,M:0.014]) Tasks(max_active:64,active:1,max_pending:1000,pending:0)
```

Figure 6: *monitoring* log line example.

verbosity depending on the use case. The level supported are FATAL, ERROR, INFO, WARN, DEBUG.

storm-backend.log

This log file contains all the information about the SRM execution process, error or warning. At start-up time, the BE logs here all the properties value, which can be useful to check parameters effectively adopted by the system. After that, at the INFO level, the BE logs for each SRM operation who has requested the operation (user subject name), on which files (SURLs) and with which result. If ERROR or FATAL levels are set, the only event logged in the file are those due to error conditions. In Figure 7 an example of *storm-backend* log is shown.

```
23:59:59.590 - INFO [xml       50] - srmReleaseFiles: user </DC=ch/DC=cern/OU=Organic Units/
                        ATLAS Data Management> operation on [SURL: srm://storm-fe.cr.cnaf.infn.it/atlas
/atlasdatadisk/                                b4714abdb196aef1d5fbf] succesfully done with: [status: SRM
_SUCCESS: Released]
```

Figure 7: *storm-backend* log line example.

heartbeat.log

StoRM provides a bookkeeping framework that elaborates informations on SRM requests processed by the system to provide user-friendly aggregated data that can be used to get a quick view on system health. The *heartbeat.log* file contains information on the SRM requests processed by the system from its startup, adding new information at each beat. The beat time interval can be configured, by default is 60 seconds. At each beat, the heartbeat component logs an entry.

As shown in Figure 8, the information contained in this line are the lifetime from the last startup, the BE process free heap size in Bytes, the number of Synchronous SRM requests executed in the last beat, the number of *srmPrepareToGet* and *srmPrepareToPut* requests executed from start-up and the number of *srmPrepareToGet* (as well as the *srmPrepareToPut*) executed in the last beat, with the number of requests terminated with success (OK=10) and the average completion time in millisecond (M.Dur.=150).

This file can help you to understand if the system is receiving SRM requests, or if the system is overloaded by SRM requests, or if PTG and PTP are running without any problem, or if the interaction with the filesystem is exceptionally slow (in case the mean duration is much higher than usual).

storm-backend-metrics.log

A finer grained monitoring of incoming synchronous requests is provided by this log file, which contains metrics for individual types of synchronous requests. An example entry log of

```
[2019-03-30 23:58:51,510]:  [#..4893 lifetime=81:32.01] Heap Free:1676485032 SYNCH [460] ASynch [PTG:17478
8 PTP:210415] Last:( [#PTG=19 OK=19 M.Dur.=27] [#PTP=35 OK=35 M.Dur.=54] )
```

Figure 8: *heartbeat* log line example.

storm-backend-metrics.log is shown in Figure 9: the information stored are the type of operation, the number of operation in the last minute, the number of operations from last startup, the maximum (minimum and average) duration of last bunch and the highest duration of the 95% (and 99%) last bunch operations.

```
23:59:48.673 - ea [(count=3673866, m1_rate=657.734089405885, m5_rate=611.8485002820605, m15_rate=621.70750
3018252) (max=0.463702, min=0.018446999999999998, mean=0.053741594912771286, p95=0.11159, p99=0.1826509999
9999998)] duration_units=milliseconds, rate_units=events/minute
```

Figure 9: *storm-backend-metrics* log line example.

5. Building the infrastructure with the Elastic Stack

The installation and setup of the Elastic Stack suite has been performed on a test-bed environment in the INFN-CNAF data center using a VM in Openstack [16]. In particular the specifics of the VM test-bed used are the following:

- 2.2 GHz Dual Core VCPUs (Broadwell architecture) with 4MB cache;
- 4 GB of memory;
- 40 GB of disk;
- two volumes attached with a cumulative storage of 600 GB

5.1 Parsing the log information

The information coming from the log sources, described in the previous section, contains heterogeneous data that require a structured filtering in order to be correctly indexed. Inside the local cluster itself, Logstash allows the creation of a well defined pipeline collecting data from Filebeat, which is installed in each node that contains the log files.
The different choice of filters is essential in order to correctly parse the log content: in particular, the *grok* filter has been adopted. A *grok* filter, based on Regular Expressions, selects a specific portion of text (both numeric or literal depending on the information required) by creating a series of patterns. Some patterns are predefined, such as the IP address or the timestamp ISO format (which are universally adopted and recognized), but the majority are custom made and stored in a specific configuration file.
In Figure 10 an example of parsed log, with new structured information, is shown. In this example, the original message is contained inside the *message* field and the other information parsed from the initial log are listed alphabetically, such as the *status* of the StoRM operation (INFO) or the *action* containing the specific operation performed (*srmReleaseFiles*). Some fields

are automatically added by Filebeats, e.g. the information of the sender node name (*beat.name*), the index name to which the log belongs, the *offset* field corresponding to the file offset the reported line starts at and the type of document processed (*log* in this case).

⏲	@timestamp	November 15th 2018, 18:25:06.478
t	@version	1
t	_id	gzpnGGcBcvwUa1jlsGXn
t	_index	filebeat-2018.11.15
#	_score	-
t	_type	doc
t	action	srmReleaseFiles
t	beat.hostname	storm-atlas.cr.cnaf.infn.it
t	beat.name	storm-atlas.cr.cnaf.infn.it
t	beat.version	6.4.2
t	clientDN	/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=atl 1/CN=61 0/CN=Robot: ATLAS 1
t	host.name	storm-atlas.cr.cnaf.infn.it
t	input.type	log
t	message	18:25:06.478 - INFO [xmlrpc-488926] - srmReleaseFiles: user </DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=atl 1/CN=61 0 /CN=Robot: ATLAS 1> operation on [SURL: srm://storm-fe.cr.cnaf.infn.it/atlas/atlasdatadisk/rucio/data15_13TeV/85/6e/AOD.11227506 .pool.root.1] succesfully done with: [status: SRM_SUCCESS: Released]
#	offset	404,176,017
t	prospector.type	log
t	result	SRM_SUCCESS
t	source	/var/log/storm/storm-backend.log
t	status	INFO
t	surl	srm://storm-fe.cr.cnaf.infn.it/atlas/atlasdatadisk/rucio/data15_13TeV/85/6e/AOD.11227506 .pool.root.1
t	tags	beats_input_codec_plain_applied
t	timestamp	2018-11-15 18:25:06.478
t	token	xmlrpc-488926

Figure 10: Example of StoRM Backend metrics log line parsing.

5.2 Data visualization

Using the Kibana User Interface, the information extracted from the log parsing is then processed and plotted with several histograms and charts. In Figure 11, a gauge chart shows the count of the different status in the StoRM Backend log file on a specific time range defined by the user. Another possible plot, shown in Figure 12, represents the count of all possible StoRM Frontend operations in a time series format showing the user's most frequent (Ls, Connection, status of Prepare To Get and Prepare To Put). Using a heat map view, it is also possible to geographically locate the most frequent IP addresses of StoRM client requests, as shown in Figure 13.

6. Machine Learning with the Elastic Stack

In the Elasticsearch new release (version 6.0 and above), a new extension was released using Machine Learning capabilities for data search and analytics. This functionalities are embedded within the premium features of *X-Pack* and are mainly focused on the time series anomaly detection. Using unsupervised proprietary algorithms, the most straightforward use case of this technology is to identify when a metric value or event rate deviates from its "normal" behaviour. Therefore, the entry point into the Machine Learning features is a Single Metric job that allows the investigation

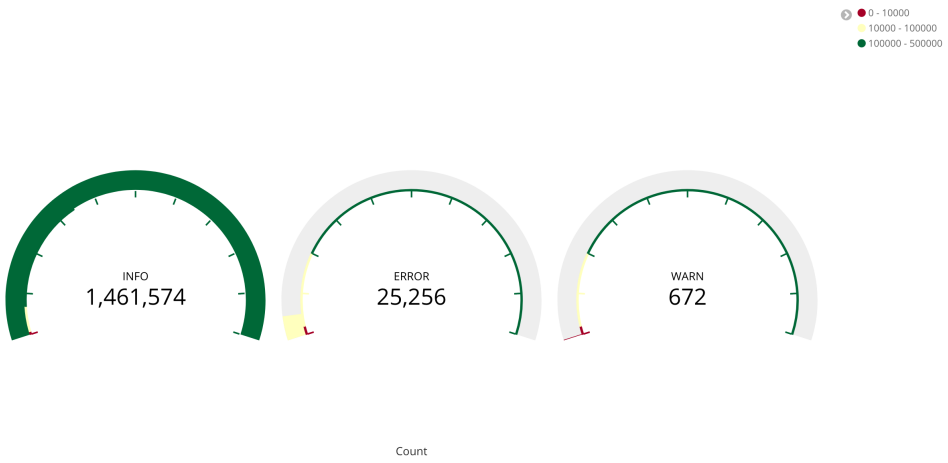

Figure 11: Gauge visualisation with the possible statuses of the StoRM Backend.

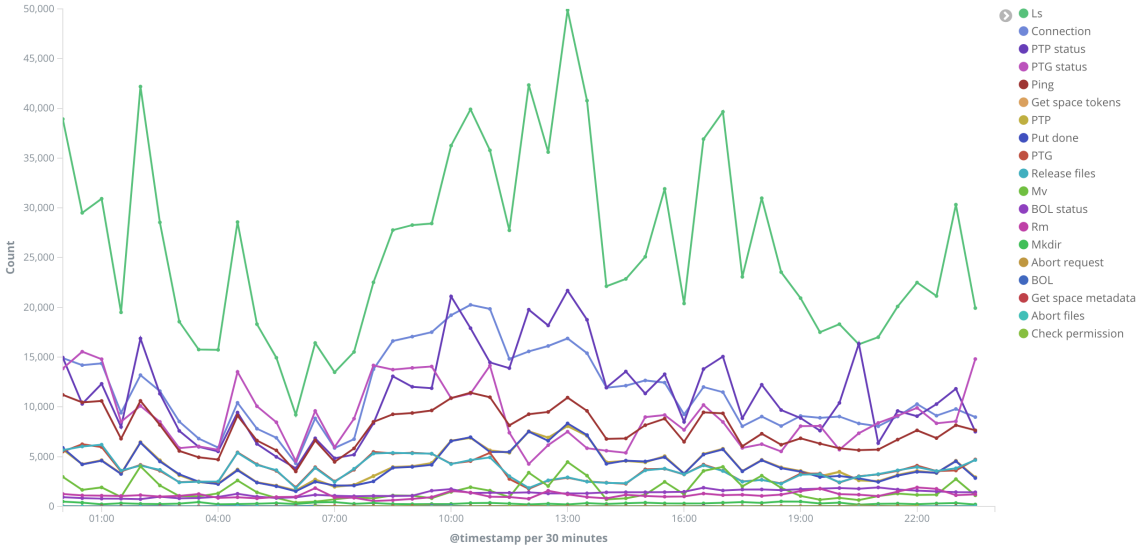

Figure 12: Count of different requests for the StoRM Frontend.

of data trend during normal activity in order to identify anomalies on univariate time series data. If the anomalies found are useful, it is then possible to run the analysis continually in real time and alert when an anomaly occurs.

The implementation of such functionalities is optimized to run in an Elasticsearch cluster, allowing the elaboration of millions of events in a small amount of time. An example is shown in Figure 14, using the information from the last continuous bunch of *srmPrepareToGet* synchronous operations from the *heartbeat* log file. The light blue area corresponds to the model bounds created, at first, by the training phase of a past time interval and then by the real-time training process. Data coming from the log is shown as a straight line and - during a normal activity - should be contained inside the model bounds; when an anomaly occurs, the metric value exceeds the model threshold and, therefore, produces a warning alert with a severity proportional to their relative discrepancy.

Figure 15, instead, shows an application of these functionalities with data coming from the parsed

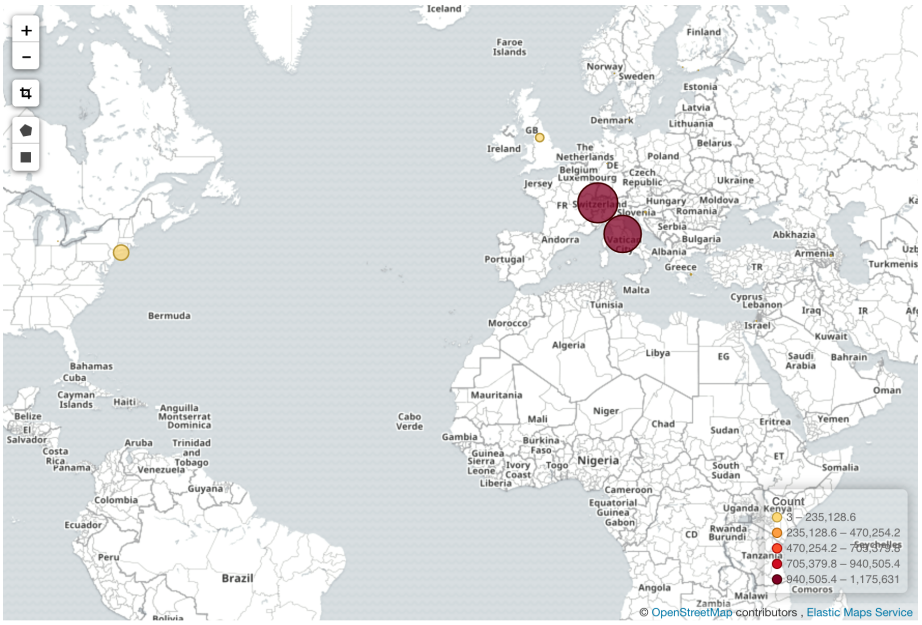

Figure 13: Geo heat map for the IP address location of the most frequent StoRM requests.

information of the StoRM Backend metrics log at the CNAF data center. Using the mean duration of the last bunch of operations described in each log entry, it is possible to notice that - starting at a certain point - the time duration increases by an order of magnitude causing, therefore, an increase also in the number of warnings (shown in figure as small coloured dots based on the level of alert). Figure 15 also depicts another important consideration: in fact - in the few hours preceding the sudden increase in the metric value - some anomalous isolated peaks may indicate and may be associated with the subsequent failure. The proactive identification of the anomaly, however, is not possible using this particular tool given in the Elastic Stack. The only tool provided by the Machine Learning functionality consists in a *forecast* button which purpose is the prediction of a time interval, given a specific portion of training previous data. This option, however, does not take under consideration any possible quick fluctuation of the examined metric since only the average behaviour of the trained model is used for making future predictions.

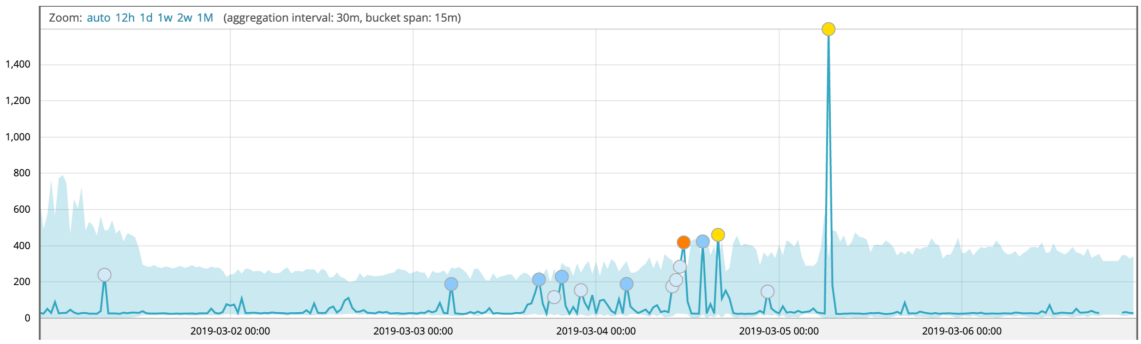

Figure 14: Time series anomaly detection for the last bunch of *ptg* operations, from *heartbeat* log.

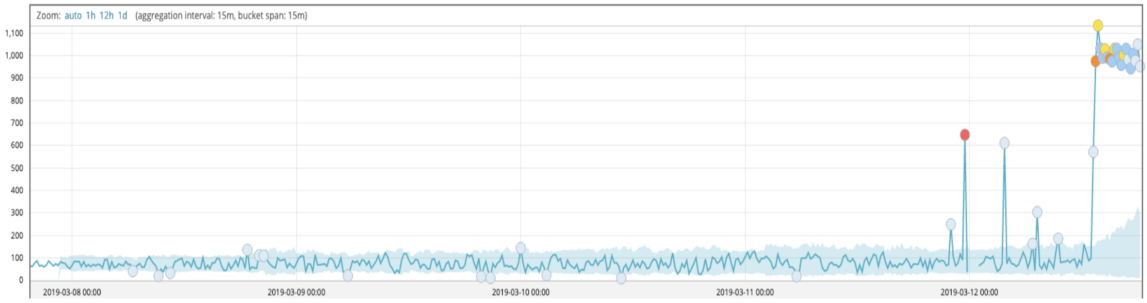

Figure 15: Time series anomaly detection for the mean duration of the continuous last bunch of operations, from *storm-backend-metrics* log.

6.1 Computing resources management

Another important aspect to consider is the workload on the VM resources used by the Elastic Stack test-bed environment itself. The performance of the computing resources provided are, in fact, continuously monitored and displayed using Graphana [17].
In Figure 16, the behaviour of the CPU usage over time is shown: a fair amount of CPU utilization is observed in the time range shown, even during the most intensive operations reaching a maximum value of about 90% but with an average value of 35-40%. The predominant task involving CPU time is the real-time Machine Learning anomaly detection analysis (over several different jobs submitted, using different test metrics); the main indication of such hypothesis is the sudden drop in CPU in the last quarter of the above-mentioned chart caused by the termination of each job due to licence expiration.
Memory usage, instead, is critical using the Elastic Stack as shown in Figure 17: during the entire time range, in fact, memory is always fully saturated. The only noticeable decrease - visible in figure - is mainly caused by the manual interruption of older indexes in order to optimize memory usage.
Concerning storage, two volumes are used: the first one, which is currently full, of 200GB and a second one, currently in use, of 400GB. Figure 18 shows the filling trend of the second disk (orange line) in the two months' observation window.

7. Conclusion

The Elastic Stack suite is a powerful tool for log analysis and ingestion as well as their indexing and visualisation for online monitoring. At the INFN-CNAF Tier-1 data center, this suite has been investigated for the creation of a centralised platform for logs coming from the StoRM service.
Using a test-bed environment on an isolated virtual machine, log files coming from different StoRM machines were collected and parsed using specific filters based on Regular Expressions. In this way, the information carried out by such logs is structured and indexable by the Elasticsearch engine providing fast search capabilities and visualisation of specific variables.
Moving towards a predictive failure system, the Elastic Stack suite provides a premium functionality with a Machine Learning approach which was investigated and tested. This system,

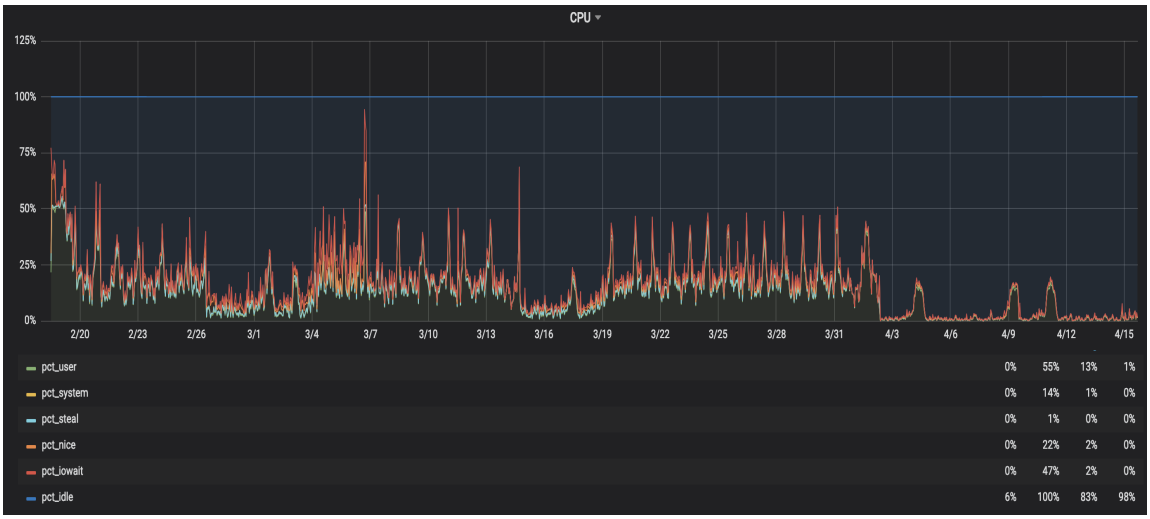

Figure 16: CPU usage on a two months time range.

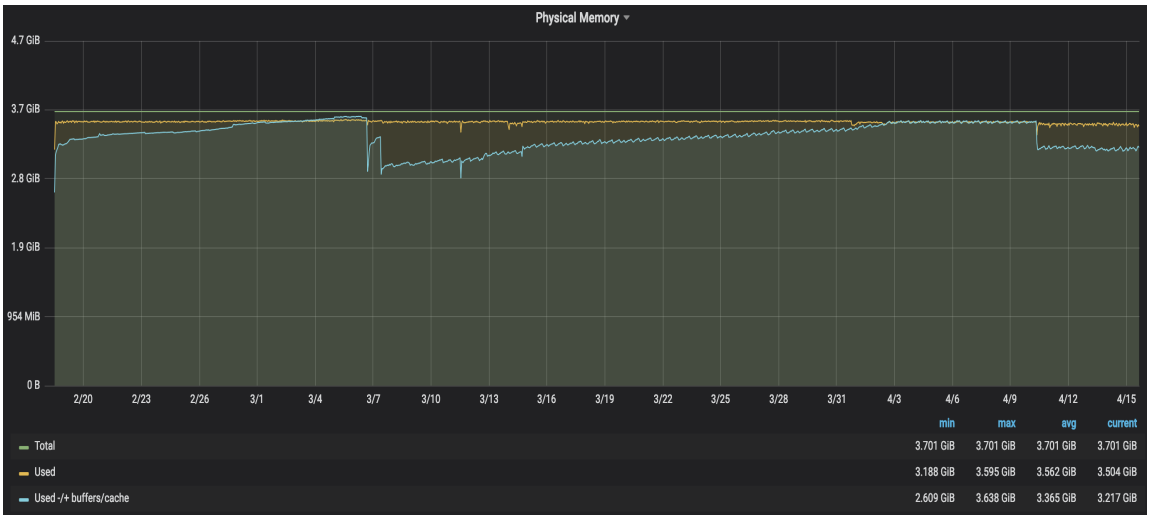

Figure 17: Memory usage on a two months time range.

using unsupervised learning techniques, is mainly used for an anomaly detection scenario by alerting operators when a certain metric exceeds a specified threshold given by the trained model, continuously updated with new information. Using several metrics coming from the StoRM Backend and Frontend instances during anomalous observation windows, several alerts were correctly identified by the system. However, for a predictive scenario and a proactive identification of failures, this may not be the optimal solution.

As a short-term goal at the INFN-CNAF computing center, some plans are currently being investigated by developers. The first step is the creation of a centralised log source stored inside a physical storage partition of the Tier-1. This unique log file will contain all the information appended from different services running and may be accessed from any virtual machine by a NFS mount point. A clean installation for the Elastic Stack suite will then be created on a

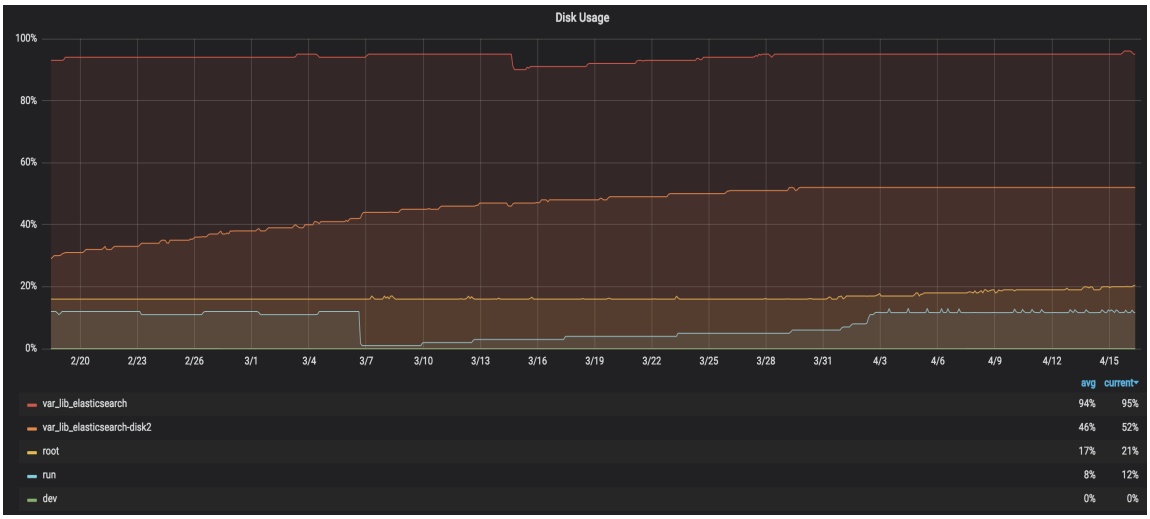

Figure 18: Storage usage on a two months time range.

dedicated physical cluster and - using the Elasticsearch engine to process all the log information - adopt other Machine Learning algorithms using conventional frameworks [18].
Another possible direction towards Big Data processing is also taken under consideration by adding an Apache Spark [19] cluster on a cloud machine at CNAF with three dedicated storage volumes of about 300GB each (installed via DODAS@CNAF).
Finally, also other logs must be analysed as well: Worker Nodes status, service machines, GPFS, GridFTP [20], xrootd [21], batch system and application level logging.

References

[1] INFN-CNAF Tier 1 Data Center. Website: https://www.cnaf.infn.it/en

[2] The High-Luminosity project at the LHC. Website: http://hilumilhc.web.cern.ch/

[3] The Elastic Stack Suite main site. Webiste: https://www.elastic.co/elk-stack

[4] Elastic company main site. Website: https://www.elastic.co

[5] Elasticsearch main site. Website: https://www.elastic.co/products/elasticsearch

[6] Logstash main site. Website: https://www.elastic.co/products/logstash

[7] Kibana main site. Website: https://www.elastic.co/products/kibana

[8] Beats main site. Website: https://www.elastic.co/products/beats

[9] Grok Filter reference guide. Webiste:
https://www.elastic.co/guide/en/logstash/current/plugins-filters-grok.html

[10] Fielding, Roy Thomas. *Architectural Styles and the Design of Network-based Software Architectures*. University of California, Irvine (2000).

[11] Apache Lucene reference guide. Website: https://lucene.apache.org/core/2_9_4/queryparsersyntax

[12] Magnoni, Luca and Zappi, Riccardo and Ghiselli, Antonia. *StoRM: A flexible solution for Storage Resource Manager in grid*. IEEE Nuclear Science Symposium conference record (2008).

[13] *CASTOR: CERN Advanced STORage manager*. Website: http://castor.web.cern.ch/castor

[14] dCache, the Book. Website: https://www.dcache.org/manuals

[15] GPFS reference guide. Website here

[16] OpenStack Documentation. Website: https://www.openstack.org

[17] Graphana Documentation. Website: https://grafana.com

[18] L.Giommi et al., Towards Predictive Maintenance with Machine Learning at the INFN-CNAF computing centre. In this conference.

[19] Apache Spark reference guide. Website: https://spark.apache.org

[20] GridFTP reference guide. Website: http://toolkit.globus.org/toolkit/docs/latest-stable/gridftp

[21] XRootD reference guide. Website: http://xrootd.org/