

Conference Paper

Empowering IT Operations through Artificial Intelligence – A New Business Perspective

Logica Banica¹, Persefoni Polychronidou², Cristian Stefan³, and Alina Hagiu¹¹Faculty of Economics and Law, University of Pitesti, Targu din Vale Street, no. 1, 110040 Pitesti, Romania²Department of Accounting and Finance, Central Macedonia Institute of Technology, Terma Magnisias, 62124, Serres, Greece³Faculty of Engineering, Informatics and Geography, Spiru Haret University, 13 Ion Ghica Street, District 3, Bucharest, Romania

Abstract

This paper aims to describe the concept of applying Artificial Intelligence to IT Operations (AIOps) and its main components, Big Data, Machine Learning and Trend Analysis. The concept was implemented by developing a multi-layered fusion of the technologies that powers the components in AIOps platforms present on the IT market. The core of an AIOps platform is represented by the Big Data organization structure and by a massive parallel data processing platform like Apache Hadoop. The ML component of the platform is able to infer the future behaviour and the regular operations that are performed from the large volume of collected data, in order to develop the ability to automate the activities. AIOps platforms find their place especially in very complex IT infrastructures, ones that require constant monitoring and quick decisions in case of failures. The case study is based on the Moogsoft AIOps platform, and its features are presented in detail, using the Cloud trial version, clearly showing the potential of such an advanced tool for infrastructure monitoring and reporting. The experiment was focused on the way Moogsoft is monitoring computing resources, is handling events and records alerts for the defined timespan, alerts grouped by category (like web services, social media, networking). The platform is also able to display at any given moment the unresolved situations and their type of origin, and includes automated remediation tools. The study presents the features of this software category, consisting in benefits for the business environment and their integration into the Internet-of-Things model.

Keywords: Big Data, Machine Learning, AIOps, business performance.

1. Introduction

The application of high-performance IT solutions in the business environment in order to apply development strategies and achieve business goals on the one hand, and the big investments that companies make in IT on the other hand, are the key factors that encourage the development of intelligent software. Just as important as implementing front line software technologies and tools is to ensure that data and application privacy and security are guaranteed.

Corresponding Author:

Logica Banica

olga.banica@upit.ro

Received: 17 November 2019

Accepted: 6 January 2019

Published: 12 January 2020

Publishing services provided by
Knowledge E

© Logica Banica et al. This article is distributed under the terms of the [Creative Commons Attribution License](#), which permits unrestricted use and redistribution provided that the original author and source are credited.

Selection and Peer-review under the responsibility of the EBEEC Conference Committee.

 OPEN ACCESS

Though, the biggest challenge that the IT world is facing today is the processing of large, structured and unstructured data volumes stored in Big Data systems, by using artificial intelligence for analyses, simulations and forecasts.

Lately many concepts appeared on the market, the most important and powerful of them being, in our opinion, the Internet-of-Things (IoT), and, independently or influenced by it, many technologies have emerged, contributing to its implementation (e.g. Fog/Edge, microservices, containers, Big Data, analytics, machine learning and AIOps). Even if there are many standardization and fragmentation issues, we believe that their convergence is possible, in order to fulfil the functionalities of a theoretical multilayer model, analogous to the OSI (Open System Interconnection) model.

In this context, we consider that the new concept introduced by Gartner, the prestigious research company, namely Artificial Intelligence for IT Operations (AIOps), is best suited for using analytics and machine learning to analyse data collected from IoT devices.

This paper is organized into five parts, as follows: the second section, State of the Art, provides a brief overview of the current research regarding the AIOps concept, its development stage in the IT world, and the ways to integrate into the Internet-of-Things vision. The Methodology part will detail the AIOps components, focusing on Machine Learning (ML), their functions and ways of cooperation and integration into the concept. In the experimental section, we propose technologies and software tools launched by IT companies that can be used to implement AIOps for in-depth analysis of the collected data, forecasting, and building systems that can self-diagnose and react to various problems. The fifth section underlines the major directions addressed by the AIOps platforms and the advantages of using them in business environment. The Conclusions section contains the final remarks regarding the significance of this new concept for IT operations and Internet-of-Things (IoT), and about the development pace, as stand-alone systems or by integration in other technologies.

2. State-of-the-Art

According to Gartner, "AIOps platforms are software systems that combine Big Data and artificial intelligence (AI) or machine learning functionality to enhance and partially replace a broad range of IT Operations processes and tasks, including availability and performance monitoring, event correlation and analysis, IT service management, and automation" [1, 2].

Figure 1 shows the logical architecture of this new concept, that equally leverages all three components (Big Data, Artificial Intelligence and Machine Learning), their continuous improvement for integration into complex information systems.

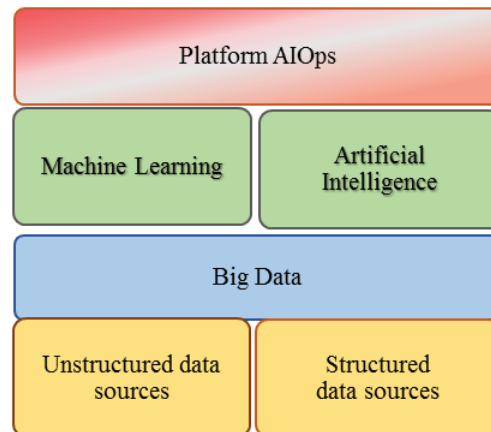


Figure 1: The components making up the AIOps platform (Source: [3]).

2.1. Big Data - component of AIOps

A definition of Big Data in correlation with the other two components of AIOps is formulated by Banica and Radulescu: “Big data is a massive collection of shareable data originating from any kind of private or public digital sources, which represents on its own a source for ongoing discovery, analysis, artificial intelligence and forecasting” [4].

Although the structured data category is more consistent and easier to process, the most important data volume comes from unstructured data sources, such as social media and mobile networks, and this volume is steadily increasing as the number of IoT devices grows.

The business environment’s interest for Big data is justified, and therefore the industry encourages the development of technologies that collect information from unconventional data sources, process them in order to analyse and detect behavioural patterns and trends, which could influence current decisions and short and medium-term forecasts.

Companies using Big Data store and process data by using Cloud Computing platforms and also rely on special Artificial Intelligence and Machine Learning applications to improve their business.

We will specify several examples [5, 6]:

i. Large IT companies, such as Google, Facebook and Amazon, were the first that showed interest in developing powerful cloud infrastructures dedicated to Big Data, in order to capture and analyse online customers opinions about the company's products and services in the shortest possible time.

ii. Detecting market trends in a significant volume of information can help companies design new products. For this purpose, one of the largest car manufacturers in the world, Ford Motor Co., collected and stored over 10,000 comments on social networks in order to improve the design a set of components for Ford cars and trucks.

iii. Manufacturing companies as well as retailers start monitoring social networks and analysing large unstructured data collections from different angles, one of the most important parameters being customer satisfaction. By combining this information with log files and sales data, companies can realize forecasts and anticipate customer behaviour.

iv. Banking sector is also using Big Data to understand the ways that clients interact with banks (websites, call centers, e-mail) and their preferences concerning financial services.

2.2. Artificial Intelligence and Machine Learning

Artificial Intelligence (AI) refers, in short, to systems that are able to learn (supervised or unsupervised), and then, based on their findings, to make decisions and react like intelligent people. AI's range of applications is vast and has many directions, but we will focus on business, and in particular on consumer marketing.

After the collection and processing of data obtained from integrated enterprise information systems (ERP) and external sources, a set of Business Intelligence and Analytics tools are used to help companies to understand customers and to model their behaviour on channels and categories, streamline sales processes and finally, to lead to optimal decisions in short time. Big Data Analytics has a different meaning in comparison with Machine Learning, by its contents. Big Data Analytics represents the process of collecting and analysing large data sets to discover hidden customer behaviour patterns, models and other information such as market trends that can help companies make business decisions and predict customer behaviour [7]. Machine Learning is a field of Artificial Intelligence and represents the ability of systems "to automatically learn and improve from experience without being explicitly programmed" [8].

Deep Learning is only achievable using Big Data sources, because the algorithms need to process huge amounts of samples in order to develop their "intelligence" and

hone their ability to obtain a better accuracy. The other component of Deep Learning is represented by the specially designed algorithms that have the power to analyse data and create “thinking patterns” [9]. Deep learning is a specific Machine Learning technique, based on Artificial Neural Network (ANN) with many layers, and has been successfully applied in numerous Big Data applications.

In this paper we will focus on the new AIOps concept, which is a software that combines Big Data and AI, machine learning, natural language processing (NLP) and other related technologies in order to deliver information of greater interest to IT, such as patterns, anomalies and predictions.

3. Methodology

Being a new concept, AIOps does not have a rigorous development methodology, and thus the IT companies are approaching solutions that are different regarding the components and technologies that are used. Thus, inside an AIOps platform Big Data processing, Machine Learning and predictive analytics will be implemented, according to the developer's point of view.

For example, for the robotization of operations in the automotive industry, many consider the Machine learning component more important, while for a Cloud computing company, the focus is on IT security, cyber-attack analysis, pattern discovery to prevent other attacks, on the predictability of hardware and software requirements, on anticipating load patterns and correlating with clients to schedule routine maintenance (backups, upgrades and new releases).

3.1. Step-by-step building AIOps

We will present several open-source, free software existing on the market for the development of the AIOps components, as well as the platforms launched in the past year to implement this concept.

In this section we will briefly present the new type of database - NoSQL, used for storing Big Data, the software that allows for massively parallel computing – Hadoop, and the Machine Learning concept as an AI component.

3.1.1. Big Data processing

The first step is to identify data sources, store and process data, and understand their significance. The AIOps core is Big Data, so access to all relevant, structured and unstructured data is required [10].

In order to facilitate the ingestion and quick processing of huge data volumes, a new database concept was introduced – called NoSQL (for Not Only SQL), that offers great performance advantages over traditional relational database systems (like Oracle and SQL Server). Dedicated platforms for Big Data have been created, the most important being the open source Apache Hadoop, distributed commercially by two major companies (Cloudera and Hortonworks). Hadoop can run on a single machine, though it was designed to be deployed on medium to large clusters of independent nodes (dedicated either to monitoring and management, or to data storage and processing). The NoSQL implementation in Hadoop is Hive Database, which can be loaded, interrogated with very complex queries, but not updated [5].

Figure 2 represents the structure of Hadoop based on two components: Hadoop Distributed File System (HDFS) and MapReduce engine. HDFS splits large data files into subsets of records which are managed by different nodes in the cluster and MapReduce distributes work around a cluster, so that operations can be run in parallel on different nodes [11].

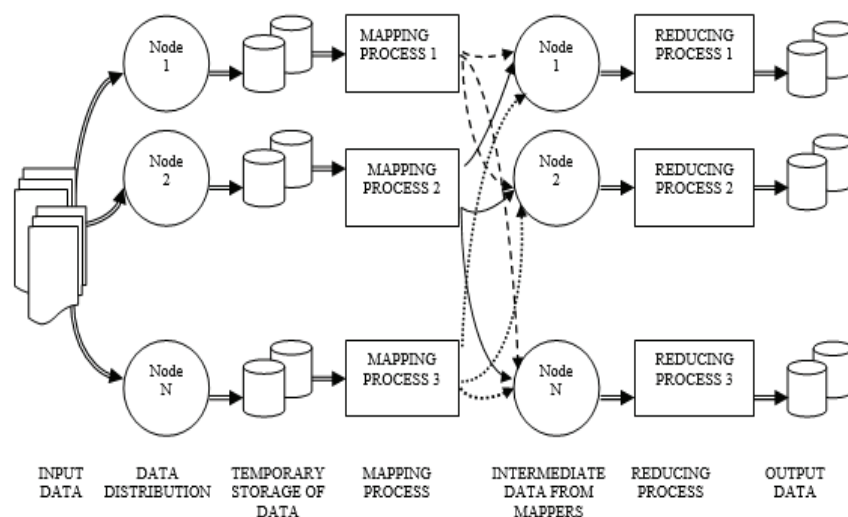


Figure 2: Big Data Processing with Hadoop (Source: [12]).

3.1.2. Machine Learning

Big Data includes important business information, but there are problems with extraction and processing data, due to its large volume, variety and collection speed.

There is a whole range of available analytics solutions for Big Data, one of them being machine learning, in variant of deep learning model.

Machine learning is a domain of Artificial Intelligence (AI), realized as a software application that "learns" the behavior and operations that take place on large volumes of data, then being able to perform complex, human-like tasks or that exceed human capabilities [13].

In order to deepen the acquired knowledge, a reward/punishment logic is implemented, that feedbacks the system at each action, and thus guiding it to the overall goal.

Depending on the algorithms used to perform basic functions, there are several categories of Machine Learning [14]:

- Supervised learning (Task driven): involves prior definition of the models, learning being based on input-output pairs, and training the system to obtain further variants with different inputs and outputs and establishing solutions;

- Unsupervised learning (Data driven): In this category, training data does not include goals, so the system does not have predefined targets, they occur as a result of processing unstructured input data and searching for desirable features for the system.

- Reinforcement Learning: It uses an algorithm whereby the decisions of system are taken based on the observations gathered from continuous interacting with the environment in order to establish measures that maximize the importance or minimize the risk of the last action performed by the system. The learning algorithm is also called "agent," as it learns continuously from the environment in an iterative way, exploring the entire range of possible states.

For developing and deploying a Machine Learning model, it is necessary to be followed several steps (Figure 3):

Briefly, before starting these steps, it is necessary to define the problem that machine learning must solve and then focus on selecting the model.

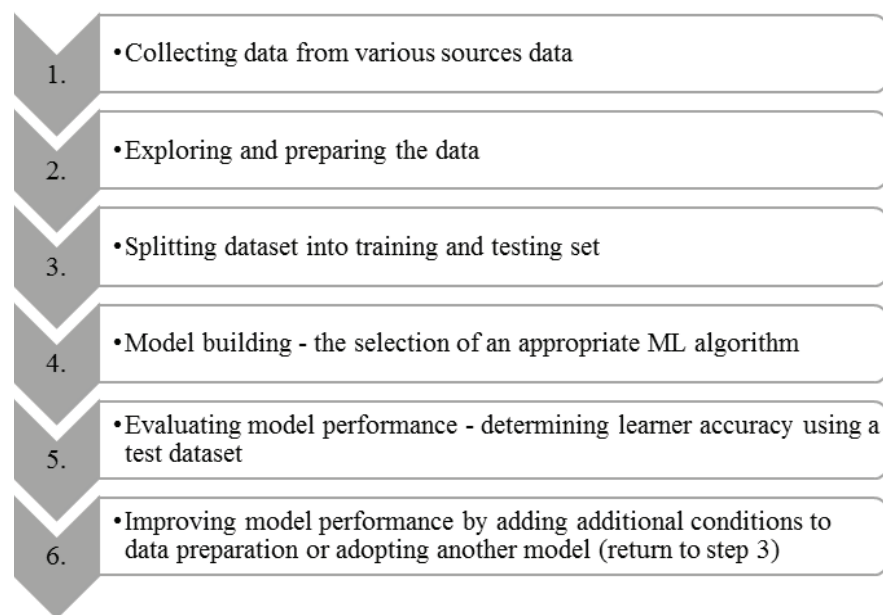


Figure 3: The phases of developing a Machine Learning model (Source: [15]).

3.1.3. Conceptul de AIOps

As we mentioned in the introduction, Artificial Intelligence for IT Operations (AIOps) is a recently emerged concept, that includes Big Data Analytics, Machine Learning (ML) and other Artificial Intelligence (AI) technologies in order to “automate the identification and solving common information technology (IT) problems”. In business, analytics technologies within the AIOps platform requires algorithms to encode the company's expertise, policies and objectives [16].

At the core of the AIOps concept is the algorithmic processing, as stated by Gartner, the company that introduced this new term. Thus, to achieve effective AIOps, it is necessary to use Machine learning for automating tasks and processes, followed by building new algorithms due to changes in the environment. Process automation leads to triggering actions when the Machine Learning and Analytics detect the fulfillment of certain conditions [17].

An AIOps platform also provides a number of visualization tools, such as dashboards, reports, graphics and predictions, so managers can follow the changes and events in the environment and make decisions that exceed the software capabilities. In complex informatic systems used by the business environment, sometimes events show up, which could be warnings (but must be tracked over time) or failure points that need to be solved quickly.

AIOps will address the events by: reporting the alerts, detecting the cause of the problem, and applying the remediation process, if the Machine Learning detects the appropriate method [18].

4. Solutions for Deploying AIOps

Unlike other software applications, the services of an AIOps platform can be deployed successively, as ML learns what to do in certain situations, thus assessing the impact and mitigating certain risks.

Although is an emerging field, there are an increasing number of AIOps platforms on IT market, dedicated for business requirements, such as:

- Moogsoft AIOps;
- Splunk's IT Service Intelligence (ITSI) tool;
- BMC's TrueSight platform;
- Cisco's Crosswork Situation Manager – AIOps.

Moogsoft is one of the most popular and widespread software products, and its features could be tested using Cloud Trial or On-Prem Trial, both delivered as pre-configured Docker Containers [19].

For Cloud Trial the user will receive a URL link to the software in Cloud. In order to use the On-Prem Trial, it is necessary to download a Docker trial image and run the AIOps on this Docker engine.

A container represents an operating environment that offers hardware capabilities for storage, connectivity, security and portability in different environments.

A container can include a single application or a component of an application/service, isolated from other software applications, which gives it increased security. The second option means that is possible to build a distributed application on multiple containers, each of them performing a specific function [20].

Docker is an open platform for running containerized applications in Linux. In order to create a dedicated environment called container, the user is required to specify the application and dependencies, environment parameters and specific configurations. When everything is in place, the deployment script is generated, that can be replicated wherever is needed [21].

We used Moogsoft's Cloud Trial to test some of the product features related to our computing system (events, alerts or situations) recorded over a 7-day period, as well as

average the mean time to acknowledge (MTTA) and the mean time to resolve (MTTR) in minutes [22].

Figure 4 shows the Workbench Summary screen, a system overview containing these statistics, automatically updated every five minutes.

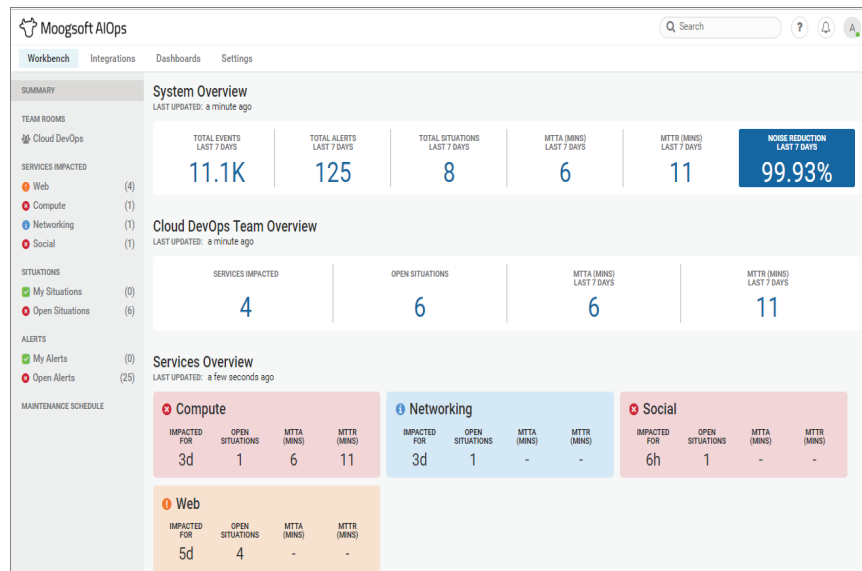


Figure 4: Workbench Summary screen (Source: developed by the authors using Moogsoft's Cloud Trial).

Workbench Summary offers a comprehensive overview about the Services Impacted, unsolved situations (Open Situations) and unsolved alerts (Open Alerts).

Figure 5 presents the Dashboards submenu, detailing some aspects of the impacted services (Web, Social, Networking, Compute) such as distribution, severity and number of situations.

Using an AIOps platform in a company is suitable if it identifies problems that occur frequently, affecting the fluent workflows and requiring redundant work for an accurate and quick response. Highly important is the relevance of data sources, which must provide the correct and necessary data for the proposed cases to be solved [23].

The AIOps software should be integrated with the existing monitoring systems and the staff should be trained to use it efficiently, on usecases like monitoring and reporting events, categorizing them (major, minor, warnings) and solving alerts.

5. Benefits of AIOps Platforms for Business

The business environment is currently experiencing a large flow of alert data, out of which a significant part is not so relevant, and the remainder needing to be placed in

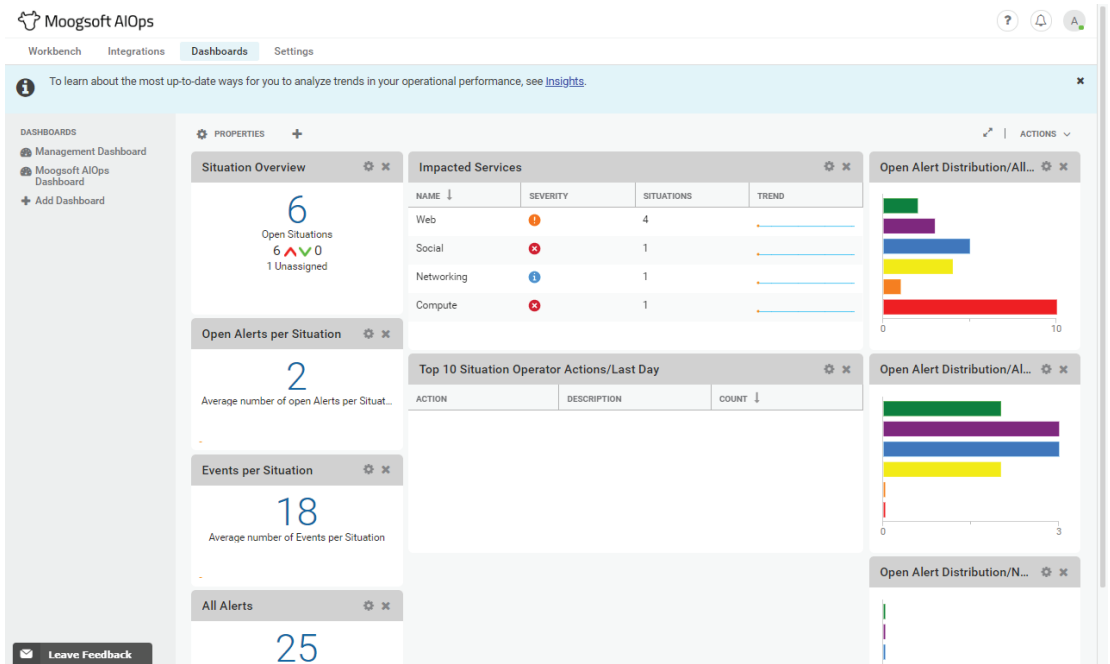


Figure 5: Detailing the Impacted Services in Dashboards submenu (Source: developed by the authors using Moogsoft's Cloud Trial).

a proper context, although all must be collected and processed in order to distinguish between useful data and background noise.

Some of the current management systems are outdated and fail to handle such amounts of data, and ultimately to help solve the problems, and thus intelligent software solutions are required, and AIOps is one of them.

Companies which implemented AIOps solutions recorded positive changes such as increased revenue, enhanced customer loyalty, lower costs and improved performance.

The concept has been launched and is widely popularized by Gartner and, essentially, it represents the combination of machine learning and data analysis to support the management systems of the companies [16].

There are three main directions addressed by an AIOps platform that bring important business benefits:

- i. Monitoring and automating of routine operations - Most computer applications frequently involve a number of routine operations that overload end-users. Also, by making a connection with IoT, we know that smart objects require the same type of microservices (ex: login, authentication, logout). In this category of routine operations, we can include the warnings sent by devices that do not require immediate actions, they should just be stored and analyzed after a while. There are only a few examples of operations that can be monitored and involve automated actions that can be accomplished with AIOps software.

ii. The second direction represents a step further from the first category of operations and refers to the identification of the fine line between warnings and alerts that an operator needs more time and experience to distinguish. In this category, we include detecting critical behavior of applications, possible viruses, worms, spyware, etc., and blocking them by using antivirus applications.

iii. Maintaining a permanent connection between Machine Learning and IT teams - The ML component of AIOps is able to learn from the past examples, though the IT team is required to provide relevant information that the platform can analyze and process. This way a human-computer interaction process occurs, process that emphasizes new cases and allows for knowledgebase expansion and applying the suitable action strategies.

AIOps applications could be integrated into the Internet-of-Things (IoT) model, providing key problem-solving capabilities for both Fog/ Edge and Cloud Application layers.

The Fog/Edge is an intermediate level of the IoT model composed of nodes that collect structured and unstructured information into NoSQL databases and run microservices in containers, required by the devices.

These are routine operations involving the use of dedicated platforms for Big Data and automated actions that can be accomplished with Machine Learning.

The Cloud Application layer includes Big Data analytical processing, that could be extended by using AIOps business intelligence tools and graphical viewing tools for results.

6. Conclusions

This paper aims to emphasize the new concept launched by Gartner, AIOps, and ways of implementing it in the business environment, as support of the IT teams.

Therefore, after defining the concept and its components, we presented the capabilities of AIOps platforms and examples of integrating them into the information systems of the companies. We used the Moogsoft Cloud Trial platform for this purpose.

Enterprise IT teams have a real support in this category of software, because AIOps platforms provide the ability to analyze data from different sources, identify sensitive points in the functioning of the digital infrastructure and quickly solve problems with the machine learning component, and reducing the mean time to resolution (MTTR).

Also, in our paper we tried to identify the possibility of using AIOps in the IoT model, at the Fog / Edge and Cloud Application layers.

Concerning the future work, the authors aim to expand the theoretical study and experimental research and test the features offered by other AIOps platforms, especially for the networking layer of the IT operations.

References

- [1] Gartner official website (2017). Market Guide for AIOps Platforms. Published: 03 August 2017 ID: G00322184, <https://www.gartner.com/en/documents/3772124/market-guide-for-aiops-platforms>
- [2] Lawrence, C. (2018). *How the Convergence of AI and IoT is Transforming Customer Support*. <https://www.momenta.partners/edge/how-the-convergence-of-ai-and-iot-is-transforming-customer-support>
- [3] Paskin, S. (2018). *What is AIOps? Artificial Intelligence for IT Operations Explained*, <https://www.bmc.com/blogs/what-is-aiops/>.
- [4] Banica, L. and Radulescu, M. (2015). Using Big Data in the Academic Environment. *In Procedia Economics and Finance*, Volume 33/2015, pp. 277-286
- [5] Banica, L. and Hagi, A. (2015). Big Data in Business Environment. *Scientific Bulletin -- Economic Sciences*, Volume 14/ Issue 1, pp.79-86
- [6] Rosenbush, S. and Totty, M. (2013). How Big Data Is Changing the Whole Equation for Business, <http://www.wsj.com/articles/SB10001424127887324178904578340071261396666>
- [7] Martinez, I. (2017). *Finding Hidden Customer Behavior Patterns Using Big Data Analytics*, <https://dzone.com/articles/find-hidden-customer-behavior-patterns-using-big-d>
- [8] Tiwari, K. (2017). *When Machine Learning meets Big Data*, <https://towardsdatascience.com/when-machine-learning-meets-big-data-4923091ba140>
- [9] Zhang, Q., Yang, L. T., Chen, Z., et al. (2018). A survey on deep learning for big data. *Information Fusion*, Volume 42, 2018, pp. 146-157
- [10] Null, C. (2018). *Essential guide to AIOps: Top tools and implementation tips*, <https://techbeacon.com/essential-guide-aiops-top-tools-implementation-tips>
- [11] Lo, F. (2014). *Big Data Technology*, available at <https://datajobs.com/what-is-hadoop-and-nosql>
- [12] Banica, L., Paun, V., Stefan, C. (2014). Big Data leverages Cloud Computing opportunities. *International Journal of Computers & Technology*, Volume 13, No.12, <http://cirworld.org/journals/index.php/ijct/article/view/3036>

- [13] Klass, L. (2018). Machine Learning - Definition and application examples, https://www.spotlightmetal.com/machine-learning--definition-and-application-examples-a-746226/?cmp=go-aw-art-trf-SLM_DSA-20180820&gclid=EAlalQobChMIhrWU7_iU3wIVy4eyCh2PfQVHEAAYAAEgLvAPD_BwE
- [14] Ayodele, T., O. (2010). *Types of Machine Learning Algorithms*, DOI: 10.5772/9385, https://www.researchgate.net/publication/221907660_Types_of_Machine_Learning_Algorithms
- [15] Sahu, Y. (2018). *Machine Learning Applications for Businesses*, <https://www.mangoblogger.com/blog/machine-learning-applications-for-businesses/>
- [16] Rouse, M. (2018). AIOps (artificial intelligence for IT operations), <https://searchitoperations.techtarget.com/definition/AIOps>
- [17] Harper_a, R. (2017). *Understanding the Machine Learning in AIOps, Part 1*, <https://www.moogsoft.com/blog/aiops/understanding-machine-learning-aiops/>
- [18] Harper_b, R. (2017). *Understanding the Machine Learning in AIOps -- Part 3: Fishing for Information in a Sea of Data*, <https://www.moogsoft.com/blog/aiops/understanding-machine-learning-aiops-part-3>
- [19] Casper, D. (2017). *Moogsoft AIOps On-Prem Trial*, <https://www.moogsoft.com/blog/moogsoft-aiops-on-prem-trial/>
- [20] Ismail, B.I., Mostajeran, E., Ab Karim, M., B., et al. (2015). *Evaluation of Docker as Edge Computing Platform*, DOI: 10.1109/ICOS.2015.7377291, <https://www.researchgate.net/publication/281445982>
- [21] Docker official website (2018). <https://www.docker.com/what-docker>
- [22] Moogsoft official website (2018). *Moogsoft AIOps Version 7.0.1*, <https://www.moogsoft.com/blog/moogsoft-aiops-on-prem-trial/>
- [23] Riley, C. (2018). *Tips for Putting AIOps Into Practice: What You Can Do Right Now*, <https://dzone.com/articles/tips-for-putting-aiops-into-practice-what-you-can>