

Telemetry Based Anomaly Detection and Correlation in Data Center

Arun S Jois, Jayasimha S R

Abstract: Data center is a complex amalgamation of servers where there are thousands of services, storage, networking, routers, switches and softwares providing services 24x7 to customers. Services provided can range from websites, storage, cloud platform, Email marketing etc. A team is established to detect the anomaly generated from the monitoring system. Anomaly or issues in servers cause high downtime of service. Detecting these anomalies with high accuracy and performing Root cause analysis has been a major issue. The team often remediates the symptom than the anomaly. With the use of Artificial Neural networks a trained model can provide solutions with high accuracy and scalability which result in higher uptime and reduced MTTR for customers.

Keywords: AIOps, AI, TechOps, Root cause analysis

I. INTRODUCTION

There is a massive problem when dealing with a large collection of servers [1],[6]. There is a complex ecosystem of networking, server, software, tools and security systems in place for data centres to function. To maintain these, system infrastructure monitoring tools such as Zabbix, Nagios, Monit etc. are used [2]. These tools collect the required telemetry data to remediate issues. Monitoring tools have been established to listen to the server's internal data. These data can be regarding CPU utilization, Disk Usage, Network usage, no. of processes in the queue just to name a few [5]. The telemetry data is time series data of issues [3]. Teams looking over these systems try to solve the issues generated by monitoring tools. What often ends up happening is the team tries to remediate the anomaly rather than the actual problem. To manage such issues using AI is vital for customer service, reducing operational complexity and effective resource utilization [4],[1]. AI helps teams to reduce the noisy data, filter out unnecessary data, and correlate the anomalies given by monitoring tools. It also helps teams to operate at a higher scale in reducing MTTR(Mean Time To Recovery).

II. ARCHITECTURE

In a typical data centre the rack of servers are connected to router which is connected to the gateway shown in Fig 1. Each component is monitored by tools such as Zabbix or Nagios. These monitoring tools have agents through which they get the metrics to diagnose anomaly.

Every service running on server is surveilled for proper functioning. As these components are interconnected there are chances where a minor failure has set off a torrent of alerts on monitoring systems. The team maintaining data centre have to remediate the root issues rather than remediating the torrent of alerts.




Fig 1. Simple Architecture of Data Centre

A. Scenario-I

The time series of data coming in from server. In the server there are many services for example consider a LAMP server running HTTP, FTP, SSH, MYSQL etc. Suppose the HTTP goes down due to load in one of the servers, there will be an alert in the monitoring system. In case MYSQL is not responding there will be two alerts in the monitoring system. One will be of HTTP not able to connect to MYSQL, another of MYSQL not responding.

B. Scenario-II

The servers contain a lot of tooling and instrumentation. Main server is connected to database, storage, Email, cloud services, third party application server etc. Monitoring tools keep a check on all these systems. An Email server has a lot of email in the queue and has exceeded the number of Emails that can be sent in a time frame. Once the limit is reached there is an automatic block of outbound Emails. Email account can be blocked for several reasons: too many scripts running, account compromised or cron jobs. Huge array of possible error sources causes alert fatigue for a team handling such issues. This scenario offsets a chain reaction of errors which create torrents of alerts flooding through the system. An error on the server side, database error, user interface, etc. Now the team has to investigate the root of the problem which is outbound Emails and fix it. In such situations the teams gets highly counter productive. In these servers there are alerts coming from monitoring tools.

The alerts coming in are from heterogeneous ecosystems. The ecosystem contains Saas, Paas etc.

Revised Manuscript Received on May 16, 2020.

Arun S Jois, Department of MCA, RV College of Engineering@ Bengaluru, India.

Jayasimha S. R, Assistant Professor, Department of MCA, RV College of Engineering@, Bangalore, India.




Fig 2 Faulty Network in Data Centre

Fig. 2 explains a simple case of RCA. Consider the server’s storage distributed across servers using RAID. 2 routers are installed for separate Public and Private communication. Suppose one of the Private routers has gone down then there can be a lot of errors concerning distributed database and cloud storage. But with RCA the main issues can be fixed as early as possible.

III. THE CASE FOR AIOPS

In the world of TechOps there is a case for automation, which helps in accurately analyzing and providing robust solutions for the Tech Industry. Apart from aforementioned problems AIOPS helps in classifying and finding root cause of issues.

A. Artificial neural network

With recent success in application of neural networks on real world problems ANN stands out to solve AIOPS issues. A neural network is modeled after biological neural networks. Human brain has about 100 billion neurons. Brain processes all the different inputs from sensory organs, it can also be called the major integrator. Brain also regulates hormonal activities in human body. So a model based on Neurons in human brain is an ingenious idea. The neurons are connected in a complex network which enables humans to perform activities. A simple ANN has input, training and output layer as shown in Fig 3. Each node is representative of a neuron in the human brain.

B. Supervised Learning

As the knowledge of solving issues in the data centre is available, supervised learning can be applied to resolve existing problems. Supervised Learning will help in multiple ways. It can help in root cause analysis, understanding depth of problem set, isolating issues etc. It also helps to understand anomalies better. With a better understanding of anomaly we can automate IT infrastructure at scale. Supervised learning uses labelled data sets to properly understand correct and incorrect data sets. As the anomalies are known the machine is trained to detect known anomalies and remediate them.

C. Unsupervised Learning

There are situations and cases where clustering is helpful to diagnose anomaly, to classify the anomaly into categories based on telemetry data, causality analysis and severity of damage. Classifying the anomalies helps in categorizing them from which the issues can be solved based on priority. Clustering technique helps in mapping similar types of

events of anomalies, so the incoming data from production servers and monitoring systems are divided into clusters. This further helps in providing a bigger picture of situations and concerned team can handle the issues further. For example: all the Networking issues can be combined into a single cluster. This will be majorly handled by networking team.




Fig. 3. Artificial Neural Network

D. Root Cause Analysis

In most of the data centres it is vital for business improvement to perform RCA. It not only provides the main cause of anomalies but also provides an insight into efficient management of servers. RCA helps when there are hundreds of servers running in production environments to remediate the main cause rather than the symptoms.

E. Correlation

The anomalies detected can be correlated into a single anomaly. Consider a server with multiple services running. As these services are dependent on each other a failure of one service causes a huge failure of services. Using correlation the team can have an insight as to which issues are related with each other. Along with RCA solving root anomaly will result in fixing of all other related anomalies.

IV. METHODOLOGY

The ANN has 3 sets of tasks to perform input, training and output. Before the data is given as input from data lakes, data has to go through the process of data cleaning, selection and transformation. Log files are primary source of data along with production telemetry extracted from monitoring tools [7]. Logs contain the data about each software tool used in the server throughout data centres.

From ANN a graph is plotted with causality of anomalies which denotes the services that faced downtime due to errors on the Y-axis and MTTR(Mean time to recovery) on X-axis as depicted in Fig 4.

ANN clears the noise from monitoring system as to which cluster of issues should be handled faster and efficiently leading to minimum downtime for customers. The correlation of anomalies will reduce the time needed for remediating anomalies. It is clear from Fig. 4 that P1 cluster has to be given higher priority than P2.




Fig 4 Clustering Based on Anomalies

Correlation

By applying association rules we can associate a group of anomalies that occur with each other very frequently. A sample of typical server errors are listed in Table 1. By calculating Support and Confidence certain patterns of existing anomalies are traceable. For example PHP and MYSQL are comorbid errors. A need arises to find any related pattern regarding such errors.

TABLE -1 List of Regular Issues

No.	Issues
1	PHP,MYSQL,WordPress
2	PHP,MYSQL,WordPress,HTTP
3	PHP,WordPress
4	ASP.NET,MYSQL
5	MYSQL

$$\text{Support(PHP,MYSQL)} = 2/5$$

$$\text{Confidence(PHP,MYSQL} \rightarrow \text{WordPress)} = 2/2$$

For the association PHP,MYSQL implies Wordpress error a 100% Confidence is found. A new pattern of anomalies is found that was previously unknown. Using this method the correlation of anomalies reduces the noise in data centre and server maintenance.

V. CONCLUSION

AIOps is a promising technology to achieve a highly accurate, scalable and robust system to track and reduce anomalies, and noise on servers. With Machine Learning and Root Cause Analysis the team can become highly efficient in dealing with errors and lower downtime for customers. Reducing alert fatigue is highly possible.

REFERENCES

1. <https://www.technologyreview.com/2019/08/15/133696/the-aiops-mission-simplify-the-complex/>
2. Haban D., Wybraniec D. (1988) A Tool for Measuring and Monitoring Distributed Systems During Operation. In: Kastens U., Rammig F.J. (eds) Architektur und Betrieb von Rechensystemen. Informatik-Fachberichte, vol 168. Springer, Berlin, Heidelberg
3. <https://blog.zabbix.com/zabbix-time-series-data-and-timescaledb/6642/>
4. A. Levin et al., "AIOps for a Cloud Object Storage Service," 2019

5. IEEE International Congress on Big Data (BigDataCongress), Milan, Italy, 2019, pp. 165-169.
6. Aparna Datt, Anita Goel, S.C. Gupta, Analysis of Infrastructure Monitoring Requirements for OpenStack Nova, Procedia Computer Science, Volume 54 2015, Pages 127-136,
7. Y. Dang, Q. Lin and P. Huang, "AIOps: Real-World Challenges and Research Innovations," 2019 IEEE/ACM 41st International Conference on Software Engineering: Companion Proceedings (ICSE-Companion), Montreal, QC, Canada, 2019, pp. 4-5.
8. Min Du, Feifei Li, Guineng Zheng, and Vivek Srikumar. 2017. DeepLog: Anomaly Detection and Diagnosis from System Logs through Deep Learning. In Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (CCS '17). Association for Computing Machinery, New York, NY, USA, 1285-1298. DOI:https://doi.org/10.1145/3133956.3134015

AUTHORS PROFILE



Arun S Jois BCA, MCA from RV College of Engineering@ Bengaluru. The Author has over 121 contributions on GitHub. Has worked over various technologies ranging from PHP in Web Development to Angular and recently working on automation in TechOps. Author is interested in various field of Operating System, Application development of both desktop and mobile applications, Web Apps, etc.



Jayasimha S. R working as an assistant Professor in the department of MCA at RV College of Engineering@, Bangalore. He served in the institution from 2013 to till the date. Currently he submitted his PhD Thesis to the VTU. His area of interest is cloud computing. He published more than 20 papers in national, international conferences and international Scopus indexed journals.