On-Premise AIOps Infrastructure for a Software Editor SME: An Experience Report

Anes Bendimerad
Infologic R&D
26500 Bourg-Lès-Valence, France
abe@infologic.fr

Youcef Remil
Infologic R&D
26500 Bourg-Lès-Valence, France
yre@infologic.fr

Romain Mathonat
Infologic R&D
26500 Bourg-Lès-Valence, France
rma@infologic.fr

Mehdi Kaytoue
Infologic R&D
26500 Bourg-Lès-Valence, France
mka@infologic.fr

ABSTRACT

Information Technology has become a critical component in various industries, leading to an increased focus on software maintenance and monitoring. With the complexities of modern software systems, traditional maintenance approaches have become insufficient. The concept of AIOps has emerged to enhance predictive maintenance using Big Data and Machine Learning capabilities. However, exploiting AIOps requires addressing several challenges related to the complexity of data and incident management. Commercial solutions exist, but they may not be suitable for certain companies due to high costs, data governance issues, and limitations in covering private software. This paper investigates the feasibility of implementing on-premise AIOps solutions by leveraging open-source tools. We introduce a comprehensive AIOps infrastructure that we have successfully deployed in our company, and we provide the rationale behind different choices that we made to build its various components. Particularly, we provide insights into our approach and criteria for selecting a data management system and we explain its integration. Our experience can be beneficial for companies seeking to internally manage their software maintenance processes with a modern AIOps approach.

CCS CONCEPTS

• **Computing methodologies** → **Artificial intelligence**; • **Software and its engineering** → **Software performance**.

KEYWORDS

AIOps, Predictive Maintenance, AI, Enterprise Resource Planning

1 INTRODUCTION

In today's world, Information Technology (IT) has become a critical element for automation of business process across many industries. Specialized Information Systems and Enterprise Resource Planning (ERP) solutions are widely used by organizations to manage their internal operations, customer and supplier relationships, and Industry 4.0 factories that rely heavily on monitoring and observability. INFOLOGIC [20] is one of France's leading providers of ERP solutions for the agri-food, health nutrition, and cosmetic sectors. It has established a significant reputation over its 40 years in the market with a plethora of provided services and modules. INFOLOGIC also offers its customers continuous consultation, and proactive maintenance support. Its ERP system is currently deployed by hundreds of food industries in France, each of which has at least one server that is integral to their business and any disruption can result in critical losses. Therefore, it is crucial to ensure high availability and excellent maintenance for this ERP and its infrastructure.

Software maintenance and monitoring have garnered a notable interest among companies, owing to their increased dependence on software. Meanwhile, the maintenance of such systems is becoming increasingly complex as they always involve more interrelated software and business components, especially with the rise of microservices and virtualization. It follows that traditional maintenance approaches became obsolete and do not scale to today's large systems. These approaches focus on manually performing arduous tasks such as manually logging into devices, executing scripts for checking health status, resolving anomalies on a repetitive basis. Such setting triggers the need for autonomic and self-managing computing systems to address the key reasons for failures and to improve the efficiency and quality of IT services [38, 45, 61].

In this context, the term AIOps has been introduced by Gartner [58] in 2017 to address the ITOps (IT Operations) challenges with a data-driven approach that exploits AI. AIOps platforms combine big data and machine learning functionalities to intelligently improve, strengthen and automate a wide range of IT operations, from monitoring to incident management. However, building an effective AIOps solution requires addressing several challenges. Firstly, such solution must be able to handle large and diverse amount of data generated by IT, including metrics, logs, and network traffic. This involves continuously ingesting vast volumes of data from different sources while efficiently storing them to limit the cost of the infrastructure. The AIOps solution should also facilitate data access and efficient querying. Secondly, such solution must enable the monitoring of numerous applications and components while also being easily expandable to include new ones. Thirdly, it needs to simplify the integration of ML techniques into the incident management system, such as applying an anomaly detection technique [54], or an ML model for incident triage [41, 42, 60]. This necessitates being equipped with comprehensive MLOps tools to streamline data preparation, model training and testing, as well as model deployment and monitoring. Lastly, it is important to ensure that the solution is highly accessible and intuitive to facilitate the transition of engineers from traditional approaches to AIOps. In
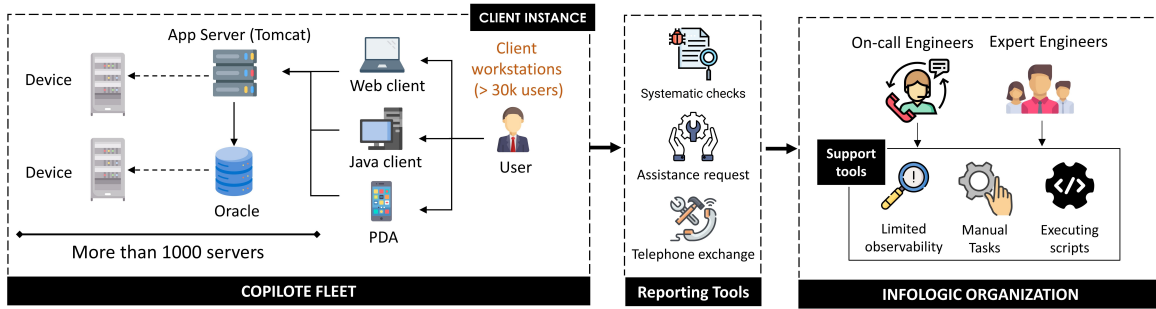
Figure 1: A simplified overview of former maintenance INFOLOGIC system before introducing AIOps and Predictive Maintenance.

fact, experience confirms that shifting the mindset and work habits within ITOps can pose a significant challenge [43].

Commercial solutions have been proposed to address these challenges by providing AIOps features such as automatic discovery, flexible data visualization, intelligent alert generation and incident diagnosis, automatic incident triage and resolution. Few examples of popular platforms are Dynatrace [11], Splunk [36], Datadog [9], New Relic [28], IBM Instana [21], and PagerDuty [31]. These platforms have witnessed an impressive success, with many of them already exceeding a billion dollars in annual revenue. While their effectiveness is undeniable, their cost can be prohibitive depending on the context, and their pricing structures often challenging to manage. Particularly at INFOLOGIC, with over a thousand servers hosting various applications, integrating a commercial solution can prove to be very expensive. Their cost is determined by multiple parameters, such as the number of monitored servers, the required storage size, the number of stored event logs. For example, our estimation of the cost for one of the aforementioned solutions ranges between $500k and $1 million per year to fully integrate it as a monitoring tool for INFOLOGIC. Additionally, while commercial AIOps solutions can monitor a broad range of popular applications, they do not cover private software such as INFOLOGIC's ERP, particularly its functional and business layer. Thus, integrating a commercial AIOps solution in our context would still require significant development efforts. Moreover, the use of such solutions may pose data governance issues as it often involves delegating data control to an external entity, which some of INFOLOGIC's clients may not be comfortable with, thus limiting our adoption of such tools.

Meanwhile, there has been a surge of open-source software solutions. Among them, numerous projects help to deal with some challenges related to AIOps, while not directly providing a complete AIOps platform. For instance, several open-source agents, such as Telegraf from InfluxData [37] and Beats from the ELK stack [5], enable seamless collection of data from applications. Also, open-source engines have been created to efficiently manage large volumes of data. Some of these solutions focus on specific types, such as metrics in the case of InfluxDB and TimescaleDB, while others are more generic, like ClickHouse and Elasticsearch. Therefore, it is becoming increasingly feasible for medium-size software companies, like ours, to develop in-house and modern AIOps platforms by combining and building on top of open-source components.

Contributions. We have thoroughly explored this possibility and investigated the different open source solutions and architectures that we can exploit to deploy an On-Premise AIOps system. In this paper, we share our experience by presenting our end-to-end on-premise infrastructure, while explaining the methodologies that we followed to make various decisions. Particularly, we provide insights into our approach and criteria for selecting a data management system and we explain how we integrated it. Although we have deployed this platform on our own servers, it can also be implemented on an IaaS Cloud solution which provides physical servers and manages them for its clients.

Outline. Sec 2 provides an overview of INFOLOGIC and its ERP software called Copilote, along with an introduction to AIOps capabilities. In Section 3, we delve into the maintenance and monitoring challenges faced in our software, as well as the obstacles encountered when deploying an on-premise AIOps solution. To address effectively these challenges, we establish a list of considered criteria that are provided in Sec 4. Sec 5 introduces our proposed AIOps infrastructure and elucidates the various choices we made. Finally, we present discussions and conclusions in Sec 6.

2 BACKGROUND

2.1 Infologic and Copilote

INFOLOGIC is a leading provider of ERP solutions for the agri-food, health nutrition, and cosmetic sectors in France. Its flagship product, Copilote, is an enterprise resource planning software designed to optimize and automate a large panel of business processes, including sales tracking, supply chain management, and customer relations. INFOLOGIC manages the complete Copilote implementation process, which includes integrating the software for the client, providing training, and maintaining the ERP instances and infrastructure in operation. This infrastructure comprises many dependent components and applications such as a Tomcat server, Oracle or Postgres databases, virtual machines, backup systems, and more. Currently, this ERP system is deployed by hundreds of food industries, ranging from small companies to large corporations. As the proper functioning of these businesses depends heavily on the reliable performance and accessibility of the ERP, it is crucial to ensure high availability and excellent maintenance for Copilote.

Before 2019, INFOLOGIC's service was mainly focused on systematic and corrective maintenance, that is, apply time-based checks

Figure 2: Main AIOps Features

and solving incident tickets that are created by clients. Fig 1 depicts the former maintenance system before introducing AIOps. The user can interact with the ERP through a web client, a Java client, or a PDA (Personal Digital Assistant). During these interactions, she may encounter some errors, outages, or difficulties of usage, for which she would submit an incident or assistance ticket for INFOLOGIC. OCEs (On-Call engineers) are the first to receive these tickets. If feasible, they can address and resolve them directly, otherwise, they route them manually to expert engineers. The scope of predictive maintenance was constrained due to a notable absence of tools and limited system observability. For instance, the lack of centralized monitoring data necessitated connecting to servers to access component-specific observations. Additionally, data retention was typically limited to one week in most cases, as storing more data using traditional database systems could incur high costs due to the high volume and velocity of the data.

Although this former approach may have been sufficient a decade ago, it no longer scales as the number of clients grows. In fact, the size of the environment production fleet has exceeded a thousand servers, each of them handling various applications. INFOLOGIC is confronted with an important challenge of enhancing operational efficiency, all while ensuring optimal customer satisfaction levels are upheld. Since 2019, INFOLOGIC has invested a lot on shifting the approach and focusing more on Predictive Maintenance rather than Corrective Maintenance. This can be achieved by following a data-driven philosophy where data is at the heart of any operation, starting from observability to automated recovery.

2.2 AIOps

The challenges encountered by INFOLOGIC have sparked our interest toward replacing manual maintenance routines with a unified and intelligent approach to conduct maintenance incidents from creation through diagnosis to resolution. The ultimate objective is to automate as many of the associated tasks as feasible, with a view to optimize the so called MTTD (Mean Time To Detect) and MTTR (Mean Time To Resolve). This approach is referred to as AIOps, which stands for AI for Operating Systems. AIOps leverages big data and machine learning to intelligently automate a wide range of IT and maintenance operations, and to accelerate the identification and resolution of IT issues and outages [43, 58]. According to research [58, 63, 65], a prototypical AIOps system ought to encompass six fundamental abilities that cover diverse tasks, as shown in Figure 2 and detailed below.

Perception. It corresponds to the capacity of gathering heterogeneous data types, including logs, metrics, and incident tickets, from multiple sources. The ingestion process needs to accommodate both real-time streaming and historical data analysis. Additionally, powerful data visualization and querying are necessary to ensure optimal response times when accessing the data.

Prevention. This process involves anticipating potential failure scenarios and forecasting high-severity outages before they occur in the system, utilizing simple alert rules, statistical models and predictive machine learning algorithms.

Detection. If errors or performance issues occur, the system must detect related anomalies. This involves analyzing extensive historical data to spot abnormal occurrences, while reducing false alarms and redundant events.

Location. The objective of this process is to identify and analyze potential root causes responsible for errors. This involves conducting a causality and correlation study, as well as providing tools that well characterize the context of those errors.

Action. Once the root cause of a problem is identified, the next step is to act to solve it. The AIOps platform must provide simplified actions based on the current context and past solutions outlined in the prescriptive maintenance protocol. Some actions can even be triggered automatically once a corresponding incident is detected.

Interaction. AIOps solutions should enable interactive analysis between the intelligence provided by integrated models and the expertise of users. This includes facilitating communication between different maintenance teams or customers, promoting efficient information sharing and problem resolution. This is achieved through various means, such as automatic incident documentation and summarization, or chatbot agents utilizing NLP techniques.

3 CHALLENGES

While adopting an AIOps solution can deliver significant benefits to our organization, it comes with numerous challenges that need to be carefully considered during the implementation of such solution.

Human challenge. One of the main challenges is the difficulty of shifting the mindset of IT professionals to adopt new ways of working [43]. A notable effort needs to be invested to simplify this transition. Thus, it is important to ensure interpretability and explainability of AIOps solutions to build trust in these approaches [52].

Novelty of AIOps. AIOps is still a recent and unstructured field [55]. Given its novelty and cross-disciplinary nature, AIOps contributions and methods are widely dispersed, lacking standardized taxonomic conventions for data management, as well as clear capabilities and implementation details. As such, discovering and comparing these methods can be a challenging endeavor. Therefore, one should determine the class of methods and tools that are needed given the defined goals, and establish a clear and objective procedure for comparing and selecting the right ones.

Handling large volumes of data. Data plays a crucial role in different aspects of AIOps. Ensuring high observability and real-time analysis requires efficient and continuous integration of large volumes of data from multiple sources. However, this integration poses several challenges. Foremost, data collection routines must not negatively impact the performance of the monitored applications, such as overwhelming their memory or network resources.

Additionally, the data management system used must exhibit high performance in terms of data ingestion, compression, and querying capabilities across various data types to support this process.

Data quality. Without effective data quality control, AIOps data may exhibit numerous defects. For example, network failures can lead to missing data, which undermines both data observability and representativeness. Message queues can help to address this issue [47]. Moreover, data related to incident management is largely sourced from human actions, such as incident tickets, corrective actions, and labels, which may be prone to human errors. This can lead to inaccuracies when using this data to train AIOps models.

Integrating ML models. AIOps presents challenges in integrating ML solutions. For instance, anomaly detection is a key task in AIOps, but data in this field is extremely imbalanced as anomalies are rare. Moreover, lack of labeled data is a common issue, limiting the scope of applications to unsupervised or semi-supervised methods. Another important fact is that AIOps data is prone to concept drift [51]. Software systems evolve for many reasons, such as the introduction of new components or an increase in the number of users. Any of these changes may impact significantly the characteristics of data. Therefore, it is essential to leverage MLOps tools [49] to monitor model quality in production, and update models accordingly.

Integrating expert knowledge. Although modern ML techniques can derive valuable insights from data, this data may not contain all the requisite knowledge to accurately perform the given task, particularly in AIOps where there is a clear lack of ground truth. Moreover, due to the evolving nature of a software system, certain knowledge and inferences may become obsolete, while new knowledge and ways of dealing with problems emerge. Hence, an AIOps platform should support an interactive process that enables the integration of knowledge from expert engineers and help to continuously improve the quality of algorithms' results. Active Learning [62, 64] can be especially beneficial in this context.

Security and Privacy. AIOps heavily relies on data, including sensitive information that could pose potential security threats or privacy concerns if not handled properly. An AIOps platform may need to track user interactions with the monitored software, or collect incident tickets containing confidential information about clients. Thus, we must ensure a high standard of security and comply with regulations such as the General Data Protection Regulation (GDPR) [15]. Privacy is one of the concerns that prompt us to implement an on-premise platform, to have full control over our data.

4 CRITERIA OF CHOICE

During the construction of our AIOps infrastructure, we had to make multiple decisions regarding the selection and configuration of architectural components, as well as choosing between using existing external tools or developing custom solutions. These decisions aimed at providing an AIOps solution that effectively address the challenges we previously identified while ensuring optimal efficiency. To make well-informed choices, we have established a list of objective criteria that guided the implementation of this project. These criteria were tailored to suit the varying circumstances and specific requirements of each component within the architecture.

License. This aspect concerns external projects that can be integrated into our infrastructure. A software solution is subject to a license that defines the terms and conditions of its usage. Several licenses are recognized as open-source, but some of them are more permissive than others. In our context, we are mainly interested in open-source solutions whose license permits both the usage and commercialization of the software. It should be noted that some solutions, such as Elasticsearch [12] and MongoDB [27], have licenses that are copyleft. Although we prefer a fully permissive license, such as Apache license or MIT license, we also consider solutions whose open-source usage is limited to the default distribution.

Notoriety and popularity. This is a relevant indicator which suggests that a project has been widely tested, validated, used, and accepted by the community. Moreover, popular projects tend to benefit from stability, continuous development, updates, and support. There exist metrics that can be used to estimate a project popularity, such as the overall number of downloads and deployment, the number of GitHub stars and forks, the number of active contributors. Some factors that are relatively subjective can also be considered, such as recommendations of experts.

Cost of implementation. This includes human and material resources that are required to implement a solution. When considering an external project that is already developed, the costs primarily involve deployment and adaptation to our ERP system. On the other hand, if we are developing a custom component from scratch, the main factor influencing this criterion is the cost of development.

Operational cost. Since we aim for an on-premise solution, this criterion involves the cost of infrastructure needed to operate the software, alongside the cost of its administration and monitoring. In contrast, in the case of a cloud-hosted solution, the primary consideration lies in the pricing associated with utilizing the service.

Flexibility and Adaptability to internal context. This criterion assesses the ease of integrating and adapting a tool to the existing technology stack and adhering to our company's standards. In our specific context, this factor holds significant importance since the AIOps infrastructure must seamlessly integrate with the existing ERP, which possesses its own characteristics and utilizes internal tools that can potentially be leveraged in the AIOps project. Furthermore, the monitoring of such a proprietary software necessarily implies specific developments to cover all its scope from technical to business layer.

Ease of usage. This is a crucial criterion as it heavily affects the adoption rate of a tool, which in turn, impacts the success probability of the project. For instance, when evaluating database systems, one key factor that influences this criterion is the query language used by the solution. Systems that use SQL such as ClickHouse and TimescaleDB are generally easier to adopt than those that employ specific query languages such as Elasticsearch and InfluxDB.

Security and privacy. Given that AIOps relies on critical data sources and applications, we need to ensure that all the architecture has strong security standards, especially when integrating external projects. Evaluating this aspect when integrating open-source tools into the project can be accomplished through various criteria, including code reviews and audits, community engagement on good practices, security features and their documentations, and security track record.

Limiting the number of tools. It is important to keep in mind that when the number of distinct tools employed in the infrastructure increases, so does the complexity and cost associated with its
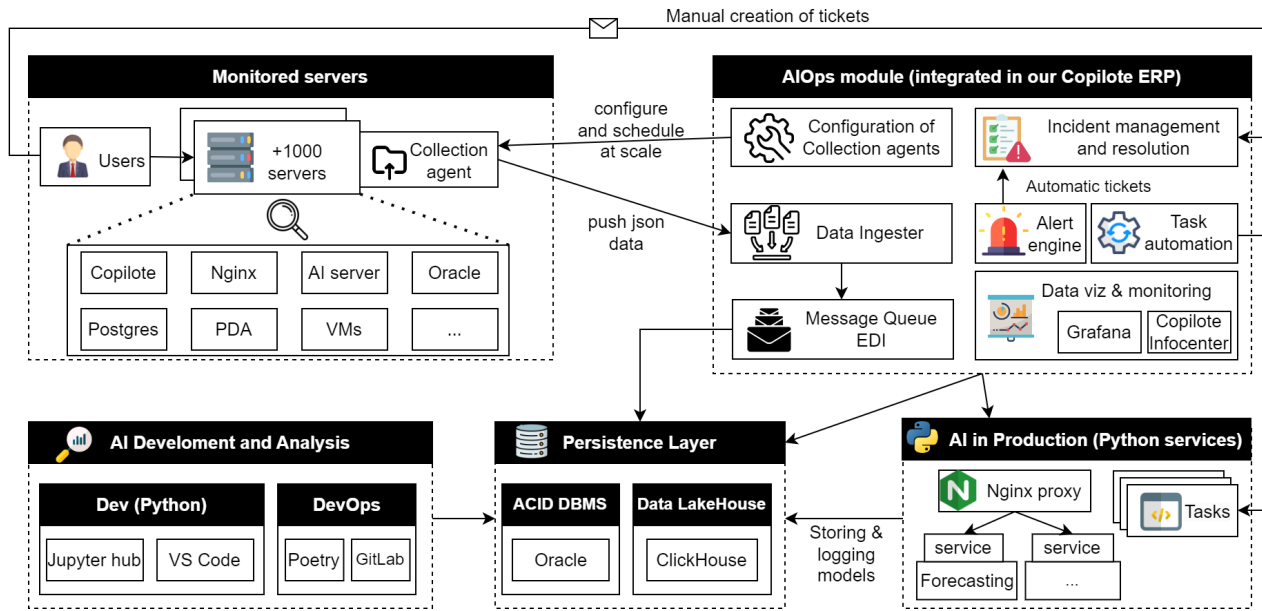
Figure 3: On-premise AIOps Infrastructure: Overview and main used tools.

maintenance. In fact, each tool may require additional specific skills and engender more complex dependencies within the architecture. **Performance and scalability.** Performance concerns the efficiency of a tool in executing its assigned task. It can be measured by the resources required to accomplish these tasks, such as time, memory, CPU, network resources. On the other hand, scalability corresponds to the capacity of the solution to scale to larger workloads, and this refers to both vertical and horizontal scalability.

5 ON PREMISE AIOPS INFRASTRUCTURE

Fig 3 provides an overview of the infrastructure and highlights the key tools used to implement different aspects of AIOps. Our monitored infrastructure comprises over a thousand servers that host the Copilote ERP and other applications that interact with it. To extract monitoring data, a collection agent is installed on each server and configured at scale using our new AIOps module, which is fully integrated into Copilote. This module ingests data from collection agents and transmit them to a Message Queue. Then, these data are stored in a relational ACID DBMS or a Data LakeHouse [48, 68], depending on the use case. Our AIOps module includes visualization tools to explore these data and create dashboards that facilitate monitoring. Moreover, a collaborative platform has been implemented to enable data scientists to analyze and experiment with collected data and develop AI algorithms. Effective algorithms can be then deployed in the AI in Production infrastructure as a web-service that serves REST queries on demand or a Python task that is orchestrated by the Copilote AIOps module. We have also introduced an internal alerting engine that can generate alerts using simple rules or advanced Machine Learning techniques, by interacting with the AI services and tasks. This engine can be configured to automatically create incident tickets that will be then handled by the incident management and resolution node. It is also possible to

manually create an incident ticket. In what follows, we describe in detail these main infrastructure components.

5.1 Collection agent

This is a lightweight application that is installed on monitored servers, and that can interact with various server components to gather their health status, performance metrics, logs, and more. Notable components that we need to monitor include our ERP software, related databases such as Oracle or PostgreSQL, operating systems, virtual machines, and more. The most critical component to monitor is the ERP software, including its technical, applicative, functional, and business layers. There exist popular open-source collection agents that can be considered in this study. For instance, InfluxData offers Telegraf [37] which is able to connect with more than 300 popular applications. The ELK stack provides the Beats suit [5] to collect and send metrics to Elasticsearch. Fluentd [14] is another agent that can collect events and send them to various types of destinations such as files or DBMS.

These tools offer numerous advantages, being open-source with substantial popularity indicated by at least 10k GitHub stars and thousands of forks. Moreover, their deployment and configuration is simple. However, they do not provide tools to gather specific data related to our ERP since it is a proprietary software, which represents a significant portion of our AIOps monitoring scope. Additionally, leveraging these tools does not allow us to directly utilize our internal data manipulation libraries, a factor that could simplify our work and enhance maintainability. Therefore, we built upon existing tools and libraries to implement our own collection agent. This choice also enables us to tailor the agent to specific network constraints that we have at Infologic.

For each installed agent, we can schedule a list of collection tasks that are continuously executed. A task can be a shell script, a Java
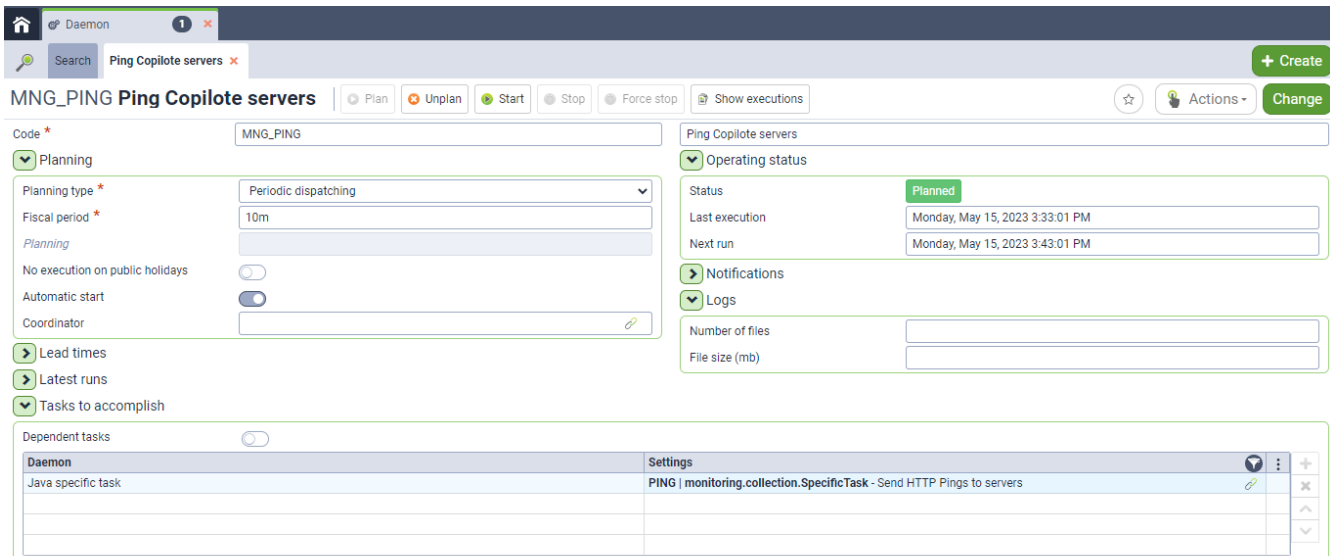
Figure 4: Agent configuration tool: Example of configuration of a task that pings Copilote servers every 10 minutes.

or Python code snippets, an SQL query, or an HTTP Rest query. An example of a simple shell task that is executed every 10 minutes with the agent to get the total number of processes in the Operating System is "ps -ef –no-headers | wc -l". The result of a task is then sent to our AIOps module that consolidates and stores the data. To address potential network failures, our collection agent is able to buffer the data collected for a certain duration and retry sending later. This ensures that our AIOps monitoring remains reliable and robust even in the face of potential network disruptions.

5.2 Agent configuration tool

We have introduced an agent configuration node within the AIOps module of our ERP. This node is used to perform the installation and configuration at scale of data collection agents on the monitored servers. Such configuration tool is crucial since we are applying an agent-based monitoring whereas collection agents need to be installed, configured, and continuously updated on a thousand servers. This tool makes it possible to flexibly configure a data collection task on a specific subset of servers. For example, to monitor database tables, we need to configure the collection of their sizes and number of rows on only DBMS servers. This is made possible thanks to the agent configuration tool. In Fig 4, we show an example of a task that is scheduled using our agent configuration tool. This task performs a basic ping on all our servers to check whether they are alive. The configuration tool offers buttons for task planning, unplanning, and other related actions. In the "Planning" section, we can observe that the task is scheduled to run every 10 minutes. Additionally, this tool provides information and logs about the previous task executions, facilitating monitoring and troubleshooting.

We developed our own configuration tool for the same reasons that motivated us to create our own agent. Moreover, this configuration tool was built by extending our ERP task management module, which is widely used in Copilote to schedule tasks in various business modules such as sales and supply chains. Utilizing this module

```java
public class PingDataIngestor extends AbstractIngestor
{

    /**
     * This function is executed whenever the ingester receives ping data
     * from any collection agent.
     * @param collectionAgent   The agent that sends the data
     * @param client            The client corresponding the the monitored server
     * @param receivedData      The data sent by the agent
     * @param metadata          Contextual information
     */
    @Override
    protected Optional<String> onData(
        Agent collectionAgent,
        Client client,
        JSONArray receivedData,
        JSONObject metadata)
    {
        // Here we can apply transformations on receivedData
        // Then we send results to a message queue using our EDI API
        return Optional.of("Data has been ingested successfully");
    }
}
```

Figure 5: An example of class that can be used to perform ingestion of ping data using Java internal APIs.

allowed us to leverage our strong internal expertise for its development and ensured direct compatibility with our monitoring and alerting system, which simplified its adoption. Here again, there are open-source solutions that can be used in this context, such as Ansible [1] from Red Hat. Some alternative tools provide an agentless approach to collect data, which eliminates the need of installing agents on monitored servers. Such approach generally requires less effort to set up a generic monitoring, it is known to be limited as it lacks the same level of interaction with applications like agent-based monitoring. In [56], an in-depth comparison of pros and cons of agent-based and agentless monitoring is detailed.

5.3 Data ingester

Our AIOps module provides a data ingester that receives monitoring data from collection agents. The JSON format has been mainly used to transmit data, as it is one of the simplest and standard data

format. Transmitted data are automatically compressed using the popular GZIP algorithm. Thanks to this ingester, data transformation and augmentation can be performed before sending it to the message queue. In Fig 5, we show the Java class and the method `onData(...)` which can be used to ingest data that we receive from collection agents. This method is executed when an agent sends new data provided by the variable `receivedData`. At this level, we can flexibly perform data transformation and then send the result to a message queue using our internal EDI (Electronic Data Interchange) API. Our framework also provides a generic data ingester that simply receives the data and sends it to a specified message queue without modifying it.

A notable open-source alternative to ingest data is Logstash [25]. While it has been mainly developed to ingest data coming from beats and going to Elasticsearch, it is now compliant with other databases and agents like Telegraf. Since the implementation of our data ingester had a limited and controlled cost, we opted for this solution to achieve full integration with our software to avoid adding an unnecessary dependency on an additional tool like Logstash.

5.4 Message queue system

A message queue system serves as an interface between the data ingester and data consumers, offering several benefits. Firstly, it ensures data durability, meaning data remains in the message queue until consumed by all the consumers that require it. Additionally, it simplifies the consumption of data by multiple destinations at once, as certain data needs to be stored in several database systems. Furthermore, asynchronous communication via the message queue notably improves ingestion performance by allowing the ingester to send data without waiting for it to be processed.

Our ERP already includes a built-in message queue module founded on the EDI standard [59], which is extensively used across various business modules of Copilote. Hence, we have leveraged this fully integrated feature for our AIOps module. An example of the EDI interface used for sending data to ClickHouse is illustrated in Fig 6. In "Emissions list for this partner", we define a list of "topics" that are consumed by this interface and stored on ClickHouse. These same topics can be consumed by other EDI interfaces as well, allowing them to be sent to additional destinations when needed. In this example, we have also attached a Java profile (displayed on the right-hand side of the screen), which is a code snippet that enables final data transformation specifically related to ClickHouse.

Had we not had the ready-to-use EDI module, we could have utilized one of the open-source message queue platforms such as Apache Kafka [3], RabbitMQ [34], or RocketMQ [35]. These systems vary in design and characteristics, suiting different project requirements. A thorough comparison of popular message queuing systems is provided in [47]. This paper highlights the distinct features of each system and summarizes their best-suited cases.

5.5 Persistence layer

Consumers from the message queue send data to the persistence layer for storage. Due to significant variations in data characteristics and significance, we could not rely on a "one size fits all" database in this context, as also supported by similar studies [65, 67]. For instance, some data related to the general characteristics of servers

and their clients require strong consistency but are of limited size. In this case, an ACID-compliant relational DBMS like Oracle is suitable. However, such a database system is not well-suited for storing thousands of performance metrics collected from each server at high frequency and retained for extended periods. For example, to monitor CPU usage alone, we collect dozens of metrics from a thousand servers at a frequency of 0.1 hertz. This data, stored for two years, results in a table with six billion rows and multiple columns. Log data presents an even greater challenge, with every web service call generating multiple lines of logs, considering over 100 million web service calls collected daily. We have devoted significant effort to identifying the appropriate data management tools for these wide variety and volume data sets. In the first version of our AIOps module, we used InfluxDB for time series data and Elasticsearch for textual data. However, we recently shifted to ClickHouse, which is highly efficient in managing various types of data, including metrics and text, and is easier to use than InfluxDB and Elasticsearch, thanks to its SQL syntax. Recent studies have shown the superiority of ClickHouse on online analytics [7, 53] compared to competitors. Therefore, Several companies, such as Contentsquare [8], eBay [30], Cloudflare [18], and Disney+ [10], have migrated to ClickHouse from other solutions.

To verify that ClickHouse is suitable for our context, we conducted a comparative study whose results led us to migrate our Data LakeHouse to this solution. This choice also simplifies administration and reduces costs by consolidating to one tool instead of two (InfluxDB and Elasticsearch). We have compared it with three other popular open source solutions which are InfluxDB, Elasticsearch, and TimescaleDB. In what follows, we detail further the position of each solution given specific criteria.

License. The four solutions we have considered all offer licenses that at least allow unrestricted usage of their default distributions. However, the extent of features included in these distributions varies between the solutions. For instance, the open-source version of InfluxDB lacks horizontal scalability, limiting its usage to a single node, while the three other alternatives can be freely used in distributed settings. Among them, ClickHouse and TimescaleDB have the most permissive licenses, being covered by Apache License 2.0.

Notoriety and popularity. All the four solutions have established a strong popularity over the open-source community. To compare their GitHub repositories' activity, we provide interesting metrics in Tab 1. The repository with the highest number of stars is Elasticsearch. Fig 7 illustrates the evolution of stars over time, showing a notable increase for ClickHouse, which has recently surpassed InfluxDB. Although all four repositories demonstrate considerable activity across the other metrics listed in Tab 1, TimescaleDB has relatively smaller values. In addition, ClickHouse stands out with a significantly higher number of commits in the past year, indicating a faster project evolution rate. These metrics provide insights into the popularity of these projects, but they do not necessarily imply that one project is universally more interesting than another.

Ease of usage and adaptability. This is an important factor that greatly influences a project's success. Our experience has shown that it is extremely helpful when the used query language is SQL, as it is widely popular. InfluxDB and Elasticsearch employ their own specific languages for data querying, which has hindered their adoption at Infologic. In contrast, ClickHouse, with its SQL-based query
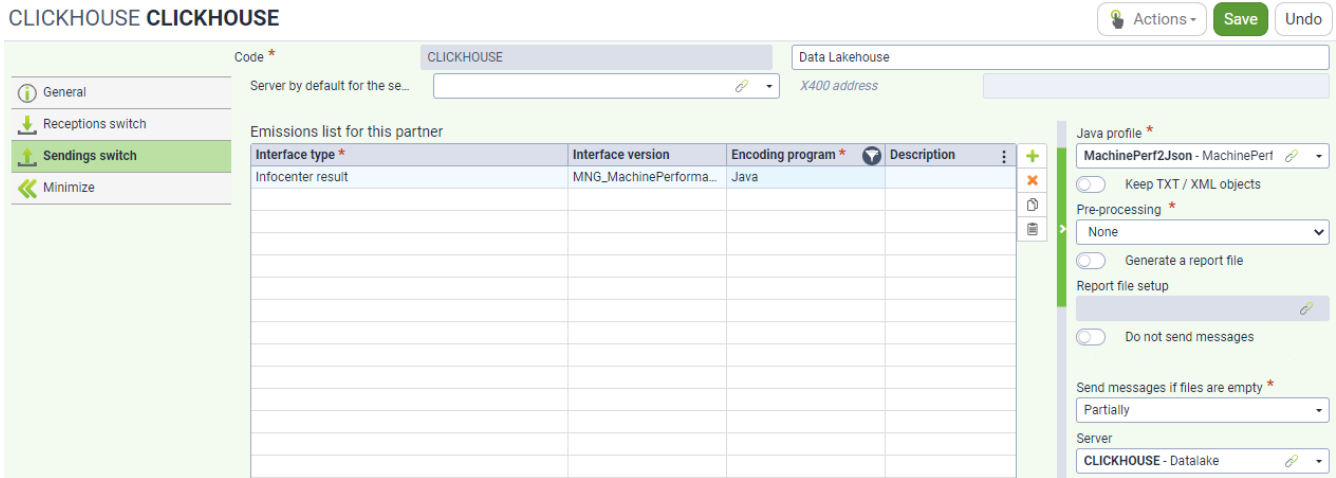
Figure 6: An EDI interface used to persist data on ClickHouse. In this simplified example, it consumes only one topic related to machine performance data.

Table 1: Metrics comparing GitHub repositories of the four DBMS considered in our study.

| Solution	Stars	Forks	Contributors	Last year commits
ClickHouse	29k	5.7k	1.2k	26k
Elasticsearch	64k	23k	1.8k	4.5k
InfluxDB	27k	3.5k	454	4.3k
TimescaleDB	15k	793	83	700

Figure 7: Evolution of number of GitHub stars for Elasticsearch, ClickHouse, InfluxDB, and TimescaleDB.

language, has witnessed a faster adoption within our company. Another aspect for evaluating the ease of usage and adaptabilty of a solution is the availability of libraries and tools that facilitate its integration with other projects. For instance, integrating a database source on a visualization tool like Grafana. All the four studied solutions benefit from extensive library support that ease their integration on different projects.

Performance. Numerous empirical studies have been performed to compare DBMS solutions for online analytics on large volumes of data. Mostafa et al. [53] provide a thorough study that evaluates ClickHouse, InfluxDB, TimescaleDB, and PostgreSQL in handling time series data. It compares their ingestion and storage capacity, and their performance on executing different types of queries. This study shows the superiority of ClickHouse on handling metrics compared to other solutions. While it also indicates satisfactory performance of InfluxDB, this study has not evaluated it on datasets with high time series cardinality. In fact, it is known that the Achilles heel of InfluxDB is when the number of time series exceeds a few millions, whereas the performance of the open-source InfluxDB version becomes poor. Furthermore, ClickBench [7] is a wide benchmark that showcases ClickHouse as one of the top performers in online analytics compared to over thirty other solutions, including Elasticsearch, TimescaleDB, and several paid cloud solutions. These findings are reinforced by the experiences of various companies. For example, Contentsquare [8] migrated its infrastructure from Elasticsearch to ClickHouse, reducing infrastructure costs by eleven times while storing six times more data and achieving a tenfold improvement in query execution speed for 99% of queries. Similarly, eBay [30] reduced its infrastructure footprint by 90% after migrating from Druid to ClickHouse. Cloudflare [18] improved its performance by migrating from CitusDB to ClickHouse, and Disney+ [10] chose ClickHouse over Elasticsearch, Apache Flink, and Hadoop after thorough evaluation.

Scope of usage. ClickHouse and Elasticsearch are generic databases capable of handling various types of data, including time series, text, logs, and unstructured data in JSON format. In contrast, InfluxDB and TimescaleDB are specifically designed for time series data. Despite this specialization, our performance study did not demonstrate superior performance of the specialized solutions. Therefore, it is not advantageous for us to invest in a dedicated time series database, which would increase the complexity of our system.

5.6 Data visualization and monitoring

Efficient data visualization tools are a critical component of a successful AIOps project. These tools must offer a user-friendly interface and provide a variety of visualization panels that are suitable for monitoring, such as time series, gauges, dynamic topologies. These tools help IT professionals to monitor the infrastructure, diagnose problems, identify component relationships, navigate to root causes, and make informed decisions. Dashboards, which are pre-configured sets of visualization panels that succinctly represent data in a single view, are one of the primary uses of these tools. Our AIOps module integrates two visualization tools. The first is
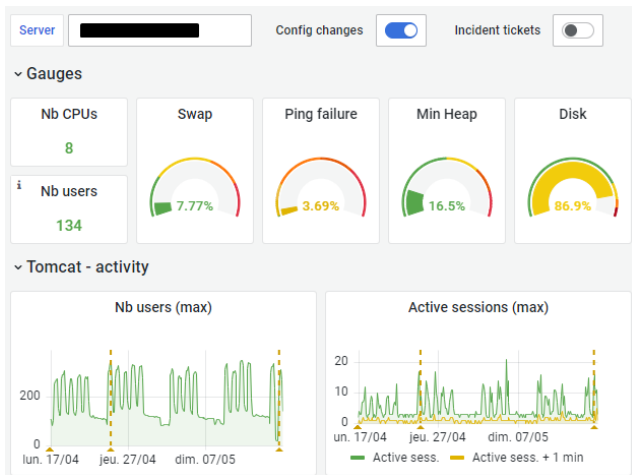
Figure 8: Example of a Grafana dashboard that visualizes data loaded from both ClickHouse and Oracle.

Copilote infocenter, our BI tool, which is fully integrated into our software and simplifies manipulation of data stored in the Oracle database of the ERP using a drag-and-drop approach. However, this tool currently cannot interact with external data sources such as ClickHouse. To overcome this limitation, we also use Grafana [17], an open-source tool that allows us to query data from different sources including Oracle and ClickHouse. Fig 8 shows a Grafana dashboard including data loaded from both ClickHouse and Oracle. Another notable advantage of Grafana dashboards and panels is their ability to be directly embedded in the web interface of our ERP. Other open-source solutions that could have been used instead of Grafana include SuperSet [4] and Metabase [26].

5.7 AI development and analysis

The application of AI in ITOps requires extensive exploration and experimentation to identify effective methods for addressing related challenges. As such, it is important to offer an efficient collaborative framework and toolset for data-scientists. This starts with providing a Version Control System to manage and version all source code related to experimental studies and development. Our team exploits GitLab [16], an open-source git platform that can be used for version control, CI/CD, automatic unit testing, and more. We use Python as a programming language as it is the language of choice for data science, but also, it is completely apt and well-equipped to run in production. We host a Jupyter Hub [23] to facilitate collaborative work around Python notebooks. Each data scientist can also benefit from using Visual Studio as an IDE, which offers integrated features such as auto-completion, git feature simplification, code formatting, and error detection. To perfectly isolate between different AI projects, a separate Python virtual environment is employed for each project with Poetry [32] as a virtual environment manager.

5.8 AI in Production

Once effective algorithms have been developed and tested, they can be transitioned into production. Our infrastructure supports the deployment of both HTTP RESTful query-response services

and scheduled data transformation tasks. FastAPI [13] has emerged as a popular choice for managing Python services in production environments. By using Nginx [29] as a proxy, we can redirect HTTP queries to the appropriate service, thereby facilitating the deployment of multiple services. Docker containers can be used to isolate services and streamline their deployment. Conversely, certain Python routines, such as data transformations (ETL), are more appropriately defined as tasks. These tasks are orchestrated using a built-in job scheduler in our ERP system, a tool that is also widely employed in several other modules of Copilote. Open-source alternatives for orchestrating Python jobs include Airflow [2] and Prefect [33]. Python tasks can be scheduled based on specific frequency or triggered by certain events. Orchestration tools also allow for the management of dependent job executions, where the execution of one job is contingent on the output of others.

5.9 Alert engine

Alerting is a core feature of predictive maintenance. It aims to generate timely alerts when anomalies or other relevant issues are detected (resp. predicted) in the current (resp. future) state of the system. These alerts can be created through both simple threshold-based rules and sophisticated ML techniques. To enable a full integration with our AIOps module and internal incident management tool, we have developed our own alert engine. This engine can access ERP data via Copilote Infocenter and can interact with external databases like ClickHouse. Additionally, it can trigger alerts based on computations produced by an AI service. This engine can be used to create, update, and fine-tune alerts. The definition of an alert needs to be expressed with a simple rule such as: if metric > threshold $\implies$ generate alert with a given severity level. This straightforward format is equally effective for complex alerts based on ML methods. For example, the metric used to trigger the alert could be a system anomaly score, computed beforehand with an LSTM auto-encoder [54]. One of our usage of ML methods is time series forecasting to predict the time remained before saturation of resources such as storage, memory, swap, and Oracle table space. We have employed a hybrid approach by combining simple Linear regression and a Deep Neural network approach called NHITS [40] which has shown its efficiency in long horizon time series forecasting. These methods are used to predict the number of remaining days before saturation, then, a simple alert rule is configured when this metric is lower than a specific threshold. Moreover, alert rules can be configured to initiate automatic actions, such as creating and assigning an incident ticket or applying a corrective action. Several open-source solutions, like Grafana [17] and Kibana [24], can also be used as alert engines. While Grafana has the capability to define alerts on various data sources, Kibana is limited only to Elasticsearch. It's worth noting that using such products involves the additional cost of interfacing alerts with our own system.

5.10 Incident Management and resolution

After a problem has been detected, an incident ticket is created to address this problem. This ticket can be created automatically with the alert engine, or manually by a Copilote user. An incident includes a title and description, as well as contextualized attributes

[n° 4020] Title		Predicion of saturation of a log storage space			
Created at	**Server**	**Version**	**Client**	**Level**	**Team**
31/03/23, 10h20	XXX	20_3	YYY	Critical	Infra
Storage		/data/path/to/application/production/log			

Description

The log storage on the server XXX is almost full (92%), it will be saturated in 3 days if no correction is applied. "`ls -ltrh`" shows that the following file is taking the majority of space: sales.log

Comment 1, 31/03/23 at 10h35 [Palliative action]

I've started by extending the size of the storage by 10% to avoid imminent saturation. Now I'll analyze the content of sales.log to understand the root cause.

Comment 2, 31/03/23 at 11h05 [Diagnosis]

90% of lines of sales.log starts with : "Creation of commands*". In fact, fine grained logs of command creation have been activated to diagnose the incident n°3530, but they have not been disabled afterward.

Comment 3, 31/03/23 at 11h15 [Corrective action]

I've removed all logs related to the command creation module, and I've disable the fine grained log generation for this module. The consumption of the storage space has significantly decreased. I close the ticket.

Figure 9: Example of an incident ticket.

to identify the concerned server, application, client, the time of incident, etc., as shown in Fig 9. Managing incidents includes several operations. It first starts by assigning the incident to the right team or person, as maintenance staff are organized on several teams that are specialized on different topics (database, application, infrastructure, etc.). This is called incident triage and it can be completely automated with efficient ML methods [41, 42, 60]. Additionally, incidents of each team need to be ranked with respect to their criticality. Then, for each incident, the assigned person starts by diagnosing the problem in order to understand its origin. It is noteworthy that in some urgent cases, the person starts by applying a palliative action to unlock the problem. To make the diagnosis process efficient, maintenance teams need to be equipped with effective tools that enable them to contextualize incidents and perform root cause analysis. Once the root cause is identified, corrective actions can be applied. The incident ticket also contains a list of timestamped comments that explain the operations that have been done to diagnose and solve the problem, as illustrated in Fig 9. In our Copilote AIOps module, we have implemented an incident management node that covers all the aforementioned operations and automate as many steps as possible. Alternatively, one could exploit iTop [22] which is an open-source incident management system.

6 CONCLUSION

Setting up an on-premise AIOps architecture is a strategic choice for medium-size software engineering companies when it comes to keep costs under control but also data. In this article, we presented our AIOps architecture and most importantly what led us there according to a set of adaptable criteria to guide decision-making, including data sovereignty, cost controls, ease of use, adaptability and stability among others. We believe that our methodology can

be adapted to many companies, but also simply brings elements of discussion in general. Shifting it towards the cloud is possible and is, again, a strategic decision[1].

We now run our approach in production. We drastically reduced the number of fatal incidents of all the components we monitor. For example, we almost eradicated Copilote service unavailability's due to disks/databases saturation, JVM memory leaks, UNIX swap leaks, and many more. Our infrastructure is stable: the only components we changed since the first production launches (we proceeded iteratively and added new features each time) are the NoSQL databases. The cost of handling our infrastructure notably decreased when we moved from using two solutions (InfluxDB and Elasticsearch) for our Data LakeHouse to only one (ClickHouse) and it came with a better adoption as *anyone knows SQL*. It is also easy to set up new data collections and alerts based on rules: anyone knowing either SQL or the Copilote BI module can set up an alert. For predictive and anomaly detection tasks, we still need dedicated studies and programs, generally in Python using well-known libraries such that scikit-learn [57] or specialized algorithms [40]. However, we are actively researching more efficient tools to facilitate the integration of ML/DM, including model and feature stores, as well as specialized tools for monitoring models in production (see e.g. MindsDB). Our second main line of research is to explore the application of Large Language Models [39, 44, 46, 50, 66] in various aspects of maintenance. We can analyze historical incident tickets to automatically diagnose new incidents, provide recommended corrections, and generate summaries. Currently, ChatGPT [6], a model owned by OpenAI, is remarkably more accurate than open-source alternatives. However, we are witnessing rapid progress in open-source LLMs [19] that can be trained and deployed on-premise, avoiding the need to share private data to third parties.

REFERENCES

[1] 2023. Ansible. https://www.ansible.com/
[2] 2023. Apache Airflow. https://airflow.apache.org/
[3] 2023. Apache Kafka. https://kafka.apache.org/
[4] 2023. Apache Superset. https://superset.apache.org/
[5] 2023. Beats from the ELK stack. https://www.elastic.co/fr/beats/
[6] 2023. ChatGPT, the OpenAI chat bot. https://chat.openai.com/
[7] 2023. ClickBench: a Benchmark For Analytical Databases. https://github.com/ClickHouse/ClickBench
[8] 2023. Contentsquare Migration from Elasticsearch to ClickHouse. https://clickhouse.com/blog/contentsquare-migration-from-elasticsearch-to-clickhouse
[9] 2023. Datadog. https://www.datadoghq.com/
[10] 2023. Disney+ClickHouse, Disney's Flexible ELT Pipelines in ClickHouse. https://clickhouse.com/blog/nyc-meetup-report-high-speed-content-distribution-analytics-for-streaming-platforms
[11] 2023. Dynatrace. https://www.dynatrace.com/
[12] 2023. Elasticsearch. https://www.elastic.co/
[13] 2023. FastAPI. https://fastapi.tiangolo.com/
[14] 2023. Fluentd. https://www.fluentd.org/
[15] 2023. GDPR. https://gdpr.eu/
[16] 2023. GitLab. https://gitlab.com/gitlab-org/gitlab
[17] 2023. Grafana. https://grafana.com/
[18] 2023. HTTP Analytics for 6M requests per second using ClickHouse. https://blog.cloudflare.com/http-analytics-for-6m-requests-per-second-using-clickhouse/
[19] 2023. HuggingFace Open LLM Leaderboard. https://huggingface.co/spaces/HuggingFaceH4/open_llm_leaderboard
[20] 2023. Infologic-Copilote. https://www.infologic-copilote.fr/
[21] 2023. Instana. https://www.instana.com/
[22] 2023. iTop. https://www.combodo.com/itop-193
[23] 2023. JupyterHub. https://jupyter.org/hub

---

[1]See for example https://world.hey.com/dhh/why-we-re-leaving-the-cloud-654b47e0

[24] 2023. Kibana from the ELK stack. https://www.elastic.co/fr/kibana/
[25] 2023. Logstash. https://www.elastic.co/fr/logstash/
[26] 2023. Metabase. https://www.metabase.com/
[27] 2023. MongoDB. https://www.mongodb.com/
[28] 2023. New Relic. https://newrelic.com/fr
[29] 2023. Nginx. https://www.nginx.com/
[30] 2023. Our Online Analytical Processing Journey with ClickHouse on Kubernetes. https://tech.ebayinc.com/engineering/ou-online-analytical-processing/
[31] 2023. PagerDuty. https://www.pagerduty.com/
[32] 2023. Poetry. https://python-poetry.org/
[33] 2023. Prefect. https://www.prefect.io/
[34] 2023. RabbitMQ. https://www.rabbitmq.com/
[35] 2023. RocketMQ. https://rocketmq.apache.org/
[36] 2023. Splunk. https://www.splunk.com/
[37] 2023. Telegraf from InfluxData. https://www.influxdata.com/time-series-platform/telegraf/
[38] Jasmin Bogatinovski, Sasho Nedelkoski, Alexander Acker, Florian Schmidt, Thorsten Wittkopp, Soeren Becker, Jorge Cardoso, and Odej Kao. 2021. Artificial Intelligence for IT Operations (AIOPS) Workshop White Paper. arXiv preprint arXiv:2101.06054 (2021).
[39] Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott M. Lundberg, Harsha Nori, Hamid Palangi, Marco Túlio Ribeiro, and Yi Zhang. 2023. Sparks of Artificial General Intelligence: Early experiments with GPT-4. CoRR abs/2303.12712 (2023).
[40] Cristian Challu, Kin G. Olivares, Boris N. Oreshkin, Federico Garza Ramírez, Max Mergenthaler Canseco, and Artur Dubrawski. 2023. NHITS: Neural Hierarchical Interpolation for Time Series Forecasting. In Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI 2023, Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence, IAAI 2023, Thirteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2023, Washington, DC, USA, February 7-14, 2023, Brian Williams, Yiling Chen, and Jennifer Neville (Eds.). AAAI Press, 6989–6997.
[41] Junjie Chen, Xiaoting He, Qingwei Lin, Yong Xu, Hongyu Zhang, Dan Hao, Feng Gao, Zhangwei Xu, Yingnong Dang, and Dongmei Zhang. 2019. An empirical investigation of incident triage for online service systems. In Proceedings of the 41st International Conference on Software Engineering: Software Engineering in Practice, ICSE (SEIP) 2019, Montreal, QC, Canada, May 25-31, 2019, Helen Sharp and Mike Whalen (Eds.). IEEE / ACM, 111–120.
[42] Junjie Chen, Xiaoting He, Qingwei Lin, Hongyu Zhang, Dan Hao, Feng Gao, Zhangwei Xu, Yingnong Dang, and Dongmei Zhang. 2019. Continuous Incident Triage for Large-Scale Online Service Systems. In 34th IEEE/ACM International Conference on Automated Software Engineering, ASE 2019, San Diego, CA, USA, November 11-15, 2019. IEEE, 364–375.
[43] Yingnong Dang, Qingwei Lin, and Peng Huang. 2019. AIOps: real-world challenges and research innovations. In 2019 IEEE/ACM 41st International Conference on Software Engineering: Companion Proceedings (ICSE-Companion). IEEE, 4–5.
[44] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers), Jill Burstein, Christy Doran, and Thamar Solorio (Eds.). Association for Computational Linguistics, 4171–4186.
[45] Mostafa Farshchi, Jean-Guy Schneider, Ingo Weber, and John Grundy. 2018. Metric selection and anomaly detection for cloud operations using log and metric correlation analysis. Journal of Systems and Software 137 (2018), 531–549.
[46] Luciano Floridi and Massimo Chiriatti. 2020. GPT-3: Its Nature, Scope, Limits, and Consequences. Minds Mach. 30, 4 (2020), 681–694.
[47] Guo Fu, Yanfeng Zhang, and Ge Yu. 2021. A Fair Comparison of Message Queuing Systems. IEEE Access 9 (2021), 421–432.
[48] Ahmed A Harby and Farhana Zulkernine. 2022. From Data Warehouse to Lakehouse: A Comparative Review. In 2022 IEEE International Conference on Big Data (Big Data). IEEE, 389–395.
[49] Dominik Kreuzberger, Niklas Kühl, and Sebastian Hirschl. 2023. Machine learning operations (mlops): Overview, definition, and architecture. IEEE Access (2023).
[50] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020, Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel R. Tetreault (Eds.). Association for Computational Linguistics, 7871–7880.
[51] Jie Lu, Anjin Liu, Fan Dong, Feng Gu, Joao Gama, and Guangquan Zhang. 2018. Learning under concept drift: A review. IEEE transactions on knowledge and data engineering 31, 12 (2018), 2346–2363.
[52] Yingzhe Lyu, Gopi Krishnan Rajbahadur, Dayi Lin, Boyuan Chen, and Zhen Ming Jiang. 2021. Towards a consistent interpretation of aiops models. ACM Transactions on Software Engineering and Methodology (TOSEM) 31, 1 (2021), 1–38.
[53] Jalal Mostafa, Sara Wehbi, Suren Chilingaryan, and Andreas Kopmann. 2022. SciTS: A Benchmark for Time-Series Databases in Scientific Experiments and Industrial Internet of Things. In Proceedings of the 34th International Conference on Scientific and Statistical Database Management. 1–11.
[54] Sasho Nedelkoski, Jorge Cardoso, and Odej Kao. 2019. Anomaly detection from system tracing data using multimodal deep learning. In 2019 IEEE 12th International Conference on Cloud Computing (CLOUD). IEEE, 179–186.
[55] Paolo Notaro, Jorge Cardoso, and Michael Gerndt. 2021. A systematic mapping study in AIOps. In Service-Oriented Computing–ICSOC 2020 Workshops: AIOps, CFTIC, STRAPS, AI-PA, AI-IOTS, and Satellite Events, Dubai, United Arab Emirates, December 14–17, 2020, Proceedings. Springer, 110–123.
[56] Himanshu Pandey and Er Kushagra Mittal. 2020. Analogy between Agent Less Monitoring and Agent Based Monitoring. Reliability: Theory & Applications 15, 3 (2020), 117–124.
[57] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. Journal of Machine Learning Research 12 (2011), 2825–2830.
[58] Pankaj Prasad and Charley Rich. 2018. Market Guide for AIOps Platforms. Retrieved March 12 (2018), 2018.
[59] G. Prem Premkumar, Keshavamurthy Ramamurthy, and Sree Nilakanta. 1994. Implementation of Electronic Data Interchange: An Innovation Diffusion Perspective. J. Manag. Inf. Syst. 11, 2 (1994), 157–186.
[60] Youcef Remil, Anes Bendimerad, Marc Plantevit, Céline Robardet, and Mehdi Kaytoue. 2021. Interpretable Summaries of Black Box Incident Triaging with Subgroup Discovery. In 8th IEEE International Conference on Data Science and Advanced Analytics, DSAA 2021, Porto, Portugal, October 6-9, 2021. IEEE, 1–10.
[61] Hansheng Ren, Bixiong Xu, Yujing Wang, Chao Yi, Congrui Huang, Xiaoyu Kou, Tony Xing, Mao Yang, Jie Tong, and Qi Zhang. 2019. Time-series anomaly detection service at microsoft. In Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. 3009–3017.
[62] Pengzhen Ren, Yun Xiao, Xiaojun Chang, Po-Yao Huang, Zhihui Li, Brij B. Gupta, Xiaojiang Chen, and Xin Wang. 2022. A Survey of Deep Active Learning. ACM Comput. Surv. 54, 9 (2022), 180:1–180:40.
[63] Laxmi Rijal, Ricardo Colomo-Palacios, and Mary Sánchez-Gordón. 2022. Aiops: A multivocal literature review. Artificial Intelligence for Cloud and Edge Computing (2022), 31–50.
[64] Burr Settles. 2009. Active learning literature survey. (2009).
[65] Shijun Shen, Jiuling Zhang, Daochao Huang, and Jun Xiao. 2020. Evolving from Traditional Systems to AIOps: Design, Implementation and Measurements. In 2020 IEEE International Conference on Advances in Electrical Engineering and Computer Applications (AEECA). IEEE, 276–280.
[66] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA, Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (Eds.). 5998–6008.
[67] Ajay Reddy Yeruva. 2023. Monitoring Data Center Site Infrastructure Using AIOPS Architecture. Eduvest-Journal of Universal Studies 3, 1 (2023), 265–277.
[68] Matei Zaharia, Ali Ghodsi, Reynold Xin, and Michael Armbrust. 2021. Lakehouse: A New Generation of Open Platforms that Unify Data Warehousing and Advanced Analytics. In 11th Conference on Innovative Data Systems Research, CIDR 2021, Virtual Event, January 11-15, 2021, Online Proceedings. www.cidrdb.org.