HPE POINTNEXT SERVICES

Technical white paper

Check if the document is available in the language of your choice.

ARTIFICIAL INTELLIGENCE FOR IT OPERATIONS

A study on high-performance compute

CONTENTS

INTRODUCTION

The use of artificial intelligence (AI) to simplify IT operations management and accelerate and automate problem resolution in complex modern IT environments is commonly referred to as AI for IT operations (AIOps). This term represents the evolution of IT operations analytics (ITOA), highlighting the new role of artificial intelligence in IT operations.

According to IDC, IT automation is the top use case where AI is being currently employed.[1] The global IT operations analytics market is expected to be worth $45 billion by 2025, growing at a 37.2% CAGR in the 2020–2025 period.[2]

But companies exploring AIOps adoption encounter questions such as the following:

1. Which AIOps use cases best fit my organization's needs and challenges?
2. Which applications and technologies are most appropriate for my use cases?
3. What infrastructure is required?
4. How to get started and build a road map for AIOps?

This paper covers AIOps use cases, technologies and approaches, as well as the road map for successful AIOps implementations. It also describes an HPE reference architecture blueprint for AIOps solutions and a case study for AIOps implementation in HPC systems including top-end exascale systems.

AIOPS AT A GLANCE

The goal of AIOps is to automate IT operations processes, which, in turn, allows IT organizations to increase service levels, reduce or stabilize costs for managing the growing complexity of IT environments, and limit risks associated with security and compliance. Although full automation can be considered as the ultimate objective of AIOps, a road map for AIOps implementation should be structured in various stages that allow incremental realization of benefits while organizations become familiar with its capabilities. Some examples related to how AIOps can improve incident management processes include:

- Identifying anomalies to spot problems and understand trends, event correlation, and log analytics
- Correlating events to further reduce noise and boost context
- Performing root-cause analysis, or orchestrate and automate workflows for commonly recurring events

In addition to incident management, other IT operation processes, such as capacity management, change management, and performance management benefit from AIOps adoption.

The core technology enabler for AIOps is **machine learning**. Different types of supervised, semisupervised, and unsupervised machine learning (ML) and deep learning (DL) models can be involved in AIOps use cases. For example, trend identification with ML models enable real-time anomaly detection and predictive capabilities, which differentiate AIOps from previous IT operations analytics approaches. However, it is critical to realize that model building is just a step of the ML model lifecycle, which also needs to consider data preparation, model training, model deployment, model monitoring, and model retraining/redeployment.

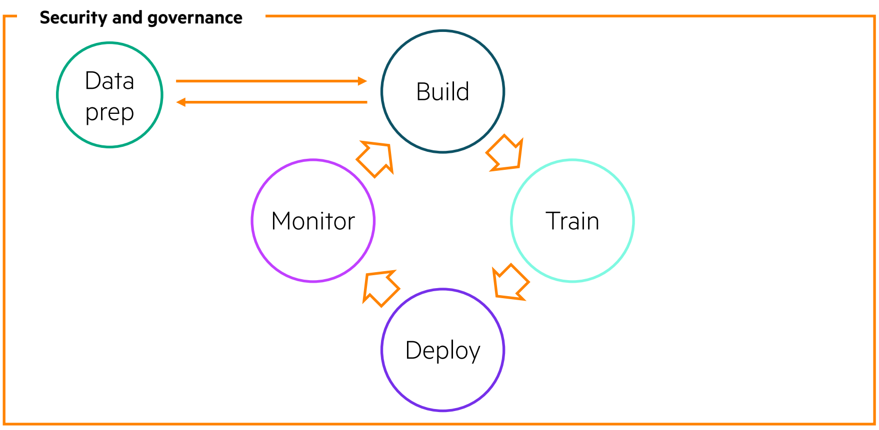

FIGURE 1. ML model lifecycle

[1] "Market Analysis Perspective, Worldwide Artificial Intelligence Software," IDC, 2020.
[2] "IT Operations Analytics Markets—Forecast to 2025," MarketsandMarkets, 2020.

AIOps also requires a **scalable data platform** that allows the ingestion, storage, and analysis of the variety, velocity, and volume of data generated by IT, at the right speed for each use case, without creating silos. In terms of variety, AIOps data sources include logs, metrics, and events from all IT technology layers, data center facilities telemetry, IT service management tool records (such as alerts from monitoring tools, incidents from service desk and incident management tools, and so on), configuration management database, text documents, and others. All these represent different types of data, for example, structured, unstructured, or semistructured, which require different strategies for its optimal ingestion, storage, and analysis. In terms of velocity, AIOps use cases will require both batch and real-time ingestion and analytics.

As an example of a technical solution, covering ML models and a scalable data platform, we will present HPE AIOps implementation for HPC systems, encompassing exascale systems in the next section. Although exascale computing represents an extreme case of the IT paradigm, the principles described are applicable to any multivendor, edge-to-cloud IT environment.

CASE STUDY: AIOPS FOR HPC SYSTEMS

Enterprise and HPC data centers are dealing with thousands of sensor metrics and associated data while a feasible target for exascale computers is in the range of 10 million data points per second.

Exascale systems will cost hundreds of millions of dollars and consist of millions of components resembling a complex system, which requires the use of advanced data analytic techniques for management and system optimization. These techniques require a much higher resolution and coverage of sensor data than the current monitoring solutions can provide.[3]

Most HPC data centers rely on set points and numerous dashboards to keep things running smoothly, but this approach typically results in thousands of false alarms that are often ignored with sometimes catastrophic consequences. The escalating volume and speed of data generation is making things more difficult, and outages are increasing. According to the Uptime Institute, nearly 33% of all data centers had an incident in 2018, up from 25% in 2017 while 80% of the incidents could have been prevented through robust anomaly detection systems.[4] Furthermore, Uptime found that more than a third of the respondents to their survey reported outages with costs exceeding $250,000, and 41 respondents reporting outages that cost more than $1,000,000.

Hewlett Packard Enterprise is using AI/ML to develop advanced, non-threshold-based, real-time analytics to reduce data center downtime via rapid and early anomaly detection that performs at scale, speed, and automatically. In addition, HPE is developing predictive capabilities to improve data center energy efficiency and sustainability with initial focus on power usage effectiveness (PUE), predictive scheduling of cooling for large jobs, water usage effectiveness (WUE), carbon usage effectiveness (CUE), and such. The effort encompasses both IT systems and the supporting facility.

To support this AIOps initiative, HPE is also developing a generic high-performance system monitoring framework. This next-generation system monitoring framework for HPC machines, called Kraken Mare, is developed under the HPE/DOE PathForward project.[5] It is designed to collect, move, and store vast amounts of data without any assumption of static data sources in a distributed, highly scalable way, and to provide different access patterns (such as streaming analytics or traditional data analysis using long-term storage) in a fault-tolerant way.

The synergistic operation of AIOps models and a high-performance, scalable monitoring infrastructure and data platform, such as Kraken Mare, is required to successfully address the new challenges requiring advanced data analytics techniques that can operate automatically and in real-time on massive amounts of data.

AIOps models

Our AIOps models currently include anomaly detection capabilities while predictive capabilities will be added in the near term. The anomaly detection will be used to help HPE customers prevent data center failures and improve data center uptime. The predictive capabilities will be used to improve energy efficiency and sustainability, via, amongst other things, optimization of the power usage effectiveness (PUE). For HPC systems, capabilities such as the predictive scheduling of cooling for particularly large jobs will also be developed. In addition, HPE is developing the ability to detect sensor drift, after which drift compensation and recalibration techniques can be applied in order to improve instrumentation resiliency and integrity. Statistical and machine learning-based approaches will be used, either stand-alone or in combination with anomaly detection models, to improve overall end-to-end performance.

[3] "Large-Scale System Monitoring Experiences and Recommendations," V. Ahlgren et al., 2018 IEEE International Conference on Cluster Computing (CLUSTER), Belfast, 2018, pp. 532-542, DOI: 10.1109/CLUSTER.2018.00069.

[4] "Data Center Outages are Common, Costly, and Preventable," Uptime Institute, 2018.

[5] PathForward is a project under the Exascale Compute Project (ECP) run by the U.S. Department of Energy (DOE) with the goal to accelerate technology development for upcoming exascale-class HPC systems.

An example end-to-end anomaly detection pipeline is shown in Figure 2. As shown, metrics are fed into metric topics in a data bus such as Kafka. The anomaly detectors read from the metric topics and conduct analytics to determine if there is any anomalous data in any of the metrics. Anomalies that are found are written to an anomaly topic. The anomaly processor reads the anomalies from the anomaly topic. The anomalies are analyzed to determine if an alert should be written to the alert topic. The alert processor reads the alerts from the alert topic, and then analyses the alerts to determine if a notification should be written to the notification topic. Notifications written to the notification topic are sent to the cluster manager, Alerta, and even a database (not shown in the figure). The end-to-end pipeline is designed to reduce false positives/negatives, improve the accuracy of the alert process, and create actionable alerts and notifications for the customer.

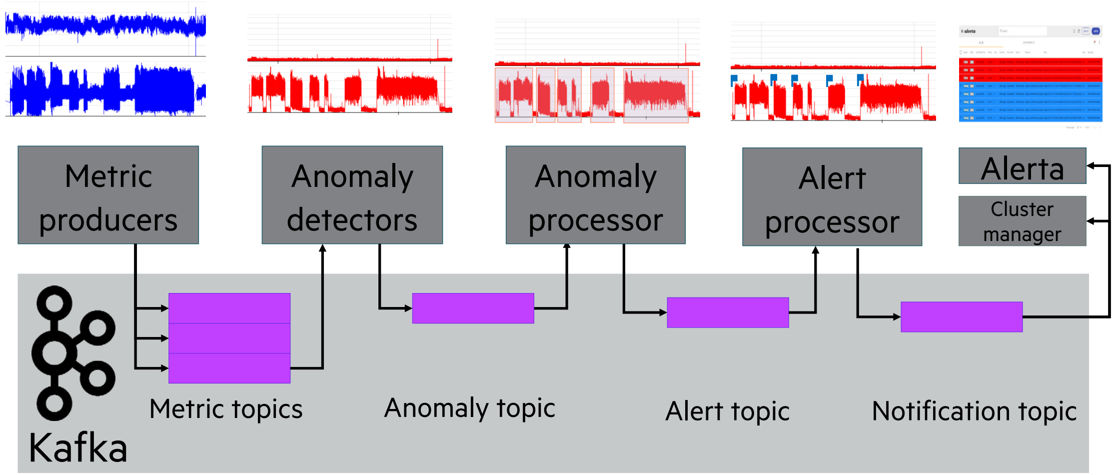

FIGURE 2. Anomaly detection workflow

The AIOps anomaly detection models for time-series data include the following features:

TABLE 1. Anomaly detection features

Feature	Description/challenge	Approaches
Univariate anomaly detection	Models operate on individual metrics.	• Baseline statistical methods, like Z-score-based anomaly detection with preprocessing methods to smooth input time-series data in real-time • Advanced statistical methods that compare the distribution of metric values in time windows with expected distribution • Forecasting models: ARIMA, deep learning models based on fully connected and Recurrent Neural Networks (RNNs)
Multivariate anomaly detection	More robust anomaly detection can be achieved when a given metric is conditioned on one or more metrics if they are correlated or interdependent.	• Auto-encoders • HPE package for training and testing models: preprocessing, on-the-fly data transformations, automated hyperparameter optimization
Real-time and automated anomaly detection	One anomaly detection algorithm does not fit all metric types (stationary, nonstationary, categorical, binary, multimodal) and anomaly incidents, so automatic model selection helps with matching metrics to the best model.	• Automated evaluation of anomaly detection pipelines: – Hyperparameter optimization, exploring different parameters and configurations of anomaly detection pipelines – Unsupervised evaluation with nominal data and artificially generated anomalies

Kraken Mare Data Platform

Kraken Mare is designed from the ground up to not only collect unprecedented volumes of data at scale but also to provide functionality that directly supports advanced analytics as provided by AIOps. The design fills the gap between a pure IoT model and a classical message bus. Based on exascale system sizes, Kraken Mare is designed to be able to collect and process at least 10 million data messages per second.

Kraken Mare provides an open and extensible common data platform providing data description, data acquisition, data flow, and data storage technologies. Kraken Mare enables the use of common methodologies and approaches to data analytics.

Kraken Mare components are based on a microservice architecture using Docker containers.

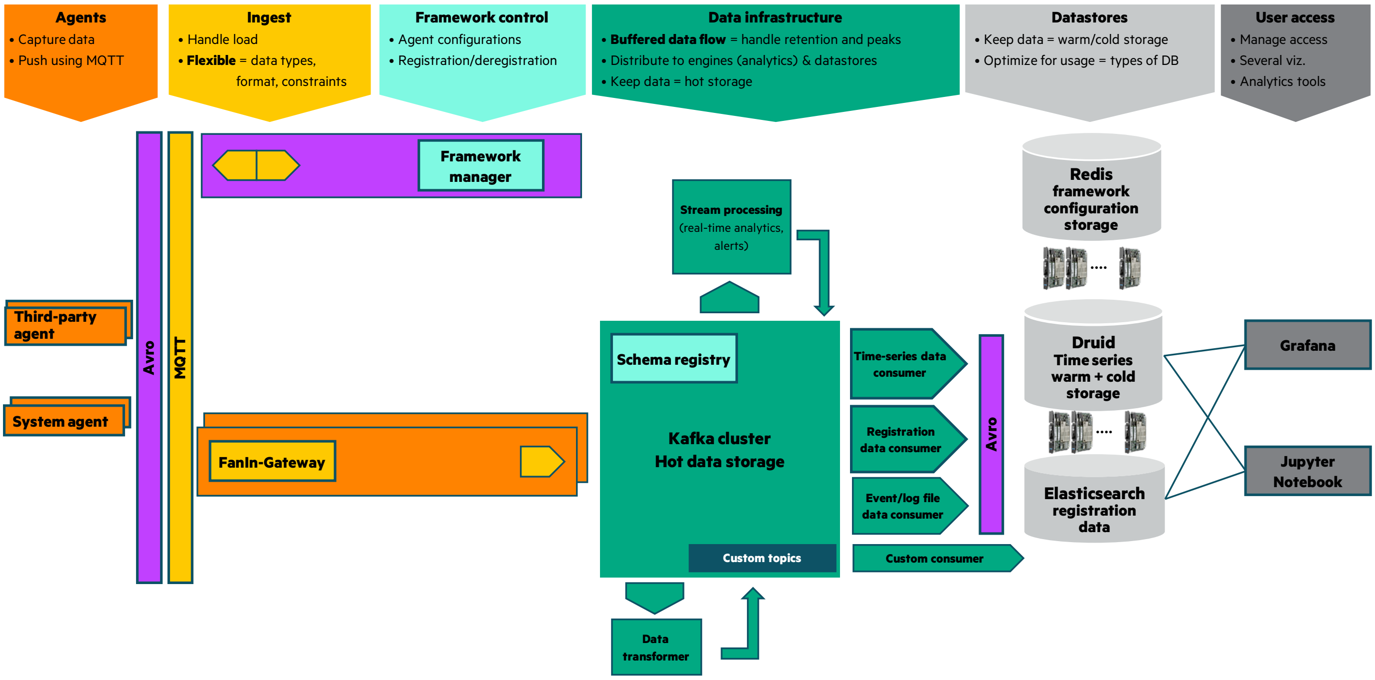

FIGURE 3. Kraken Mare component high-level overview

Kraken Mare is composed of open-source Big Data software (see Figure 3) implemented as microservices. The uniqueness comes from the communication functionality and data descriptions designed into the framework. A publish/subscribe communication paradigm was chosen in order to decouple the framework components enabling the flexibility needed to provide a standard monitoring framework and at the same time enable customizations for different customers and/or monitoring/analytics requirements. This design enables HPE to provide value-added features on top or as part of the framework distinguishing the solution from a pure component pipeline implementation. This approach makes it relatively easy, for example, to add third-party data sources to the monitoring framework and to add additional functionality, such as additional connectors to customer preferred ML frameworks, without requiring any framework changes.

The Kraken Mare prototype uses Kafka as the central messaging bus. Kafka is a distributed messaging system providing fast, highly scalable, and highly available messaging. The data collecting agents use the MQTT pub/sub protocol to submit data to a FanIn-Gateway, which represents the scalable unit in the design. The FanIn-Gateway assembles a larger message from received individual messages for better message efficiency and to protect the Kafka bus from too many individual data sources improving the scalability of the solution. Kraken Mare will be able to collect time-series data, text data (such as logfiles), and system metadata such as scheduling/job information. Kafka provides the short-term storage; Druid was chosen for the prototype to provide the long-term storage for time-series data, and Elasticsearch was chosen to provide the long term storage for text data (like logfiles) and agent metadata (This includes information about devices and sensors related to the agent). Due to the flexibility of the solution, individual long-term storage solutions can be picked based on customer and/or data requirements.

All framework components are loosely coupled enabling a truly scalable and distributed data management and control infrastructure via message busses and standardized communication messages. The communication message format can evolve over time, being backward and forward compatible. Communication messages include, for example, messages for collecting and managing data as well as control messages flowing from edge components to the central system and vice versa.

Kraken Mare scaling is done by increasing the number of FanIn-Gateways. During testing, the maximum observed message throughput was 1.15 million single metric messages per second when using 1021 nodes as data providers and one HPE ProLiant DL560 Gen10 (quad socket server board with 14-core Intel® Xeon® and 128 GB RAM) as dedicated FanIn-Gateway node. When combining this message rate of a single FanIn-Gateway with the Kafka soft limit of 100K producers, Kraken Mare could theoretically support up to 1.15 trillion messages per second.

Kraken Mare will be one of the metric producers for AIOps, and the technology will be incorporated into the next-generation Lustre monitoring system and leveraged into the next-generation HPE HPC monitoring framework.

BUILDING YOUR AIOPS SOLUTION

Building the road map

At HPE, we understand that AIOps initiatives are on a continuum that's driven by business needs and goals. Each organization has a unique path toward building a data foundation, developing advanced analytics solutions and experimenting with AIOps for select use cases.

Analysis framework

The first step that we advise customers to take in beginning their AIOps journey is a one-day **HPE AI Transformation Workshop** for key data, business, and IT stakeholders. Depending on your needs and goals, senior **HPE Pointnext Services** AI and data experts will be assigned to help you.

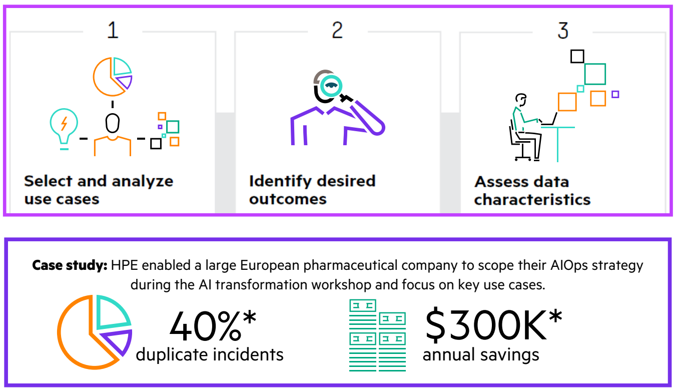

FIGURE 4. AIOps workshop analysis framework

We focus on selecting and analyzing AIOps use cases that are beneficial to customers with a focus on business-to-IT alignment. AIOps is a broad field with use cases covering event management, anomaly detection, noise reduction, thresholding, root cause determination, and such. It can be daunting for customers to begin this journey and our experience can help guide customers in focusing on the most important use cases and expand accordingly. We do this also by understanding what the desired outcomes of our approach are. Some of the outcomes that we aspire to for our customers include a more efficient IT operations team, improvement in services levels and reduction in service downtime. Finally, we assess the data characteristics, qualify data sources, and ascertain the data cleaning and aggregation requirements.

In terms of data sources, the AIOps domain deals with data coming directly from IT systems (logs and time series from network, storage, server, and other layers from the enterprise stack), as well as structured data from IT management systems, such as existing infrastructure monitoring tools, application performance and network performance monitoring tools, and IT operations management tools. Data sources can even encompass the supporting data center facilities that play a critical role in ensuring the proper functionality of the IT infrastructure. Breaking existing data silos to be able to aggregate and correlate the richness of these data sources is key for any AIOps initiative.

* Based on a customer project with a global pharmaceutical company, 2018.

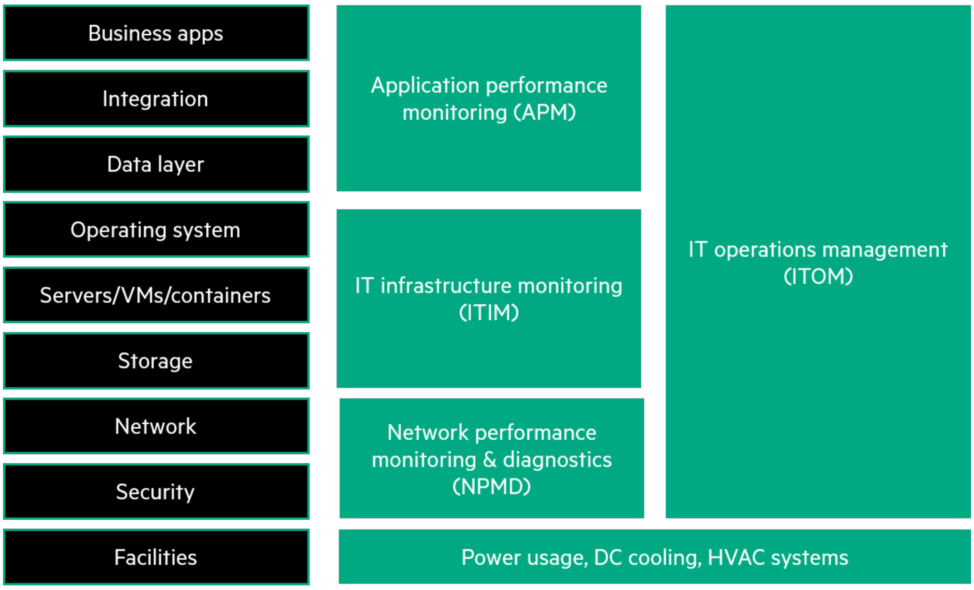

FIGURE 5. AIOps domain areas

The previously mentioned analysis framework is applied to identify our customers' most pertinent issues throughout the entire IT operations domain stack as shown in Figure 5.

Technical blueprint

HPE AIOps technical blueprint is an architectural framework created based on our extensive experience in this space to formulate tailored AIOps solutions. Figure 6 lays out all the necessary components required for a successful AIOps initiatives.

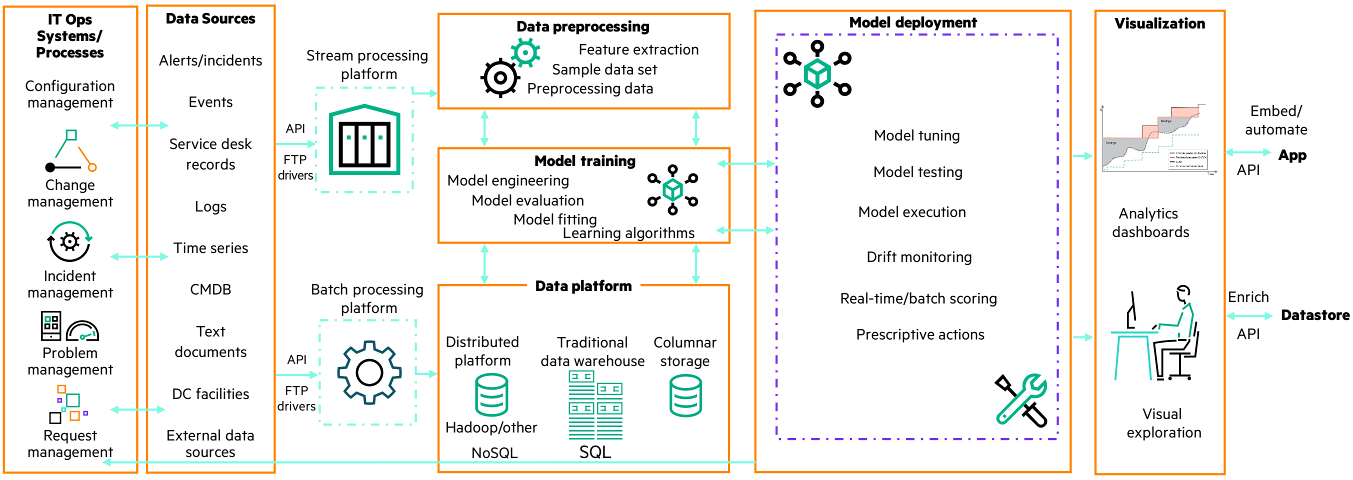

FIGURE 6. Technical blueprint

Upon completion of the HPE AI Transformation Workshop, HPE would have gathered enough information on our customer environments and map it to this blueprint once a customer is ready to go into deeper levels of discussion with regard to an AIOps solution implementation.

TABLE 2. Key questions to map out AIOps solution blueprint

Component	Description	Key questions that HPE seeks to answer
Data sources	All forms of relevant IT operations data, for example, system logs, application logs, metrics from the enterprise stack shown in Figure 6 are collected for analysis.	What is the customer's current IT system topology? What data sources exist within the IT environment that we can leverage to incorporate in our AIOps initiatives?
Data ingestion	Real-time streaming versus batch processing of raw IT operations data.	Is the data to be ingested in real-time or batch format? How can we maximize real-time ingestion for all data sources in the system?
Data preprocessing	Convert raw data into processed data that is acceptable for analysis by AI models used in AIOps use cases.	What is the ideal kind of preparation that is required for the data before AIOps analysis?
Model training	Build AI models using training data sets for AIOps use cases.	What algorithms do we train the data on?
Data storage	Data platforms used to manage and store structured, semistructured, and unstructured data.	Which data platform is ideal to manage and store the IT operations data? Should the data platform be able to handle time-series logs, structured, semistructured, and unstructured data?
Model deployment	Operationalize AI models for deployment in production environments.	What are the right tools to operationalize models that have been trained previously?
Visualization	Command and control dashboard for AIOps monitoring purposes.	Is the information being displayed of value to the IT operations team? How can we better enhance the UI experience for our customers?
Integration	Enable actions or automation through API integrations.	What integrations are needed to enable actions and automation?

SUMMARY

Unexpected system downtime in IT systems and facilities that support IT systems can be costly in terms of lost productivity and risk. The IT operations teams of many organizations are spending a lot of time performing manual monitoring, troubleshooting, and diagnostics of the increasingly complex IT environments. The efficiency and effectiveness of today's IT operations can be significantly improved with the adoption of AIOps, which represents one of the top domains where AI is being currently employed.

Hewlett Packard Enterprise offers you a strategic advantage with expertise, technology, and partnerships to deliver on the key aspects of successful AIOps implementations including the methodology to build a road map for AIOps use cases driven by business needs and goals; the design and implementation of a high-performance, scalable, and flexible data platform that is able to manage the volumes of the different data types, at the right speed for each use case; and the development and operationalization of machine learning models that can operate automatically and in real-time on massive amounts of data. In this technical paper, we have chosen to showcase these aspects with the example of our AIOps implementation for HPC systems. These methodologies, design principles, architecture modules, and ML models can be reused, customized, and expanded to the needs and objectives of any IT organization looking at AIOps adoption.

Contact your HPE representative for more information, reference configurations, use cases, demonstration, services, and support offerings.

LEARN MORE AT

hpe.com/services/ai-data

a50003274ENW, December 2020

Hewlett Packard Enterprise