UNICO I+D Project
6G-EDGEDT-03

6G-EDGEDT-03-E12

Initial design of monitoring solutions for Network Digital Twins

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL
SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia

Abstract

This Deliverable E12 presents the architecture of a monitoring solution designed to capture, process and exploit data from network elements to generate efficient digital twins capable of accurately depicting the actual state of the physical network they represent. It proposes harnessing the capabilities of cutting-edge tools employed in data processing pipelines in order to efficiently train and continuously update the digital twin. The design presented in this document serves as the basis for Deliverable E13 in the context of the 6G-EDGEDT project.

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

Document properties

Document number	6G-EDGEDT-03-E12
Document title	Initial design of monitoring solutions for Network Digital Twins
Document responsible	Antonio de la Oliva Delgado (UC3M)
Document editor	Constantine Ayimba (UC3M)
Editorial team	Antonio de la Oliva Delgado (UC3M) Constantine Ayimba (UC3M)
Target dissemination level	
Status of the document	Draft
Version	2
Delivery Date	31/12/2023
Actual Delivery Date	19/12/2023

Production properties

Reviewers	…

Document history

Revision	Date	Issued by	Description
…	…	…	…

Disclaimer

This document has been produced in the context of the Name Project. The research leading to these results has received funding from the Name Programme under grant agreement N°

Contents

List of Figures

List of Tables

List of Acronyms

AI: Artificial Intelligence

API: Application Programming Interface

ARIMA: Autoregressive Integrated Moving Average

DT: Digital Twin

HTTP: Hypertext Transfer Protocol

IPFIX: Internet Protocol Flow Information Export

JMX: Java Management Extensions

KPI: Key Performance Indicator

MANO: Management and Orchestration

ML: Machine Learning

MQTT: Message Queuing Telemetry Transport

sFLOW: sampled Flow

SLA: Service Level Agreement

SNMP: Simple Network Management Protocol

SQL: Structured Query Language

Resumen Ejecutivo

Este documento presenta la arquitectura inicial de una solución de monitoreo que permite el mantenimiento de un gemelo digital confiable en la nube perimetral. La arquitectura propuesta en este producto integra a la perfección la adquisición de datos desde entornos de nube perimetral, aprovechando la tecnología de vanguardia para potenciar un sistema de inteligencia artificial (IA). El objetivo es facilitar la generación y actualización continua de un Network Digital Twin de borde.

En resumen, la arquitectura presentada en este entregable 6G-EDGEDT-03-E12 introduce las siguientes capacidades:

- Extracción de datos históricos y en tiempo real de diversos entornos de nube perimetral, lo que garantiza una cobertura integral de métricas de red, indicadores de rendimiento y estados de dispositivos.
- Aprovecha los últimos avances en Inteligencia Artificial para permitir la creación de un gemelo digital capaz de reflejar de manera confiable el complejo entorno de red física.
- Utiliza un enfoque modular para facilitar la escalabilidad y hace que la solución esté preparada para el futuro, ya que es adaptable a avances más allá del estado actual del arte.

La arquitectura propuesta en este documento mejora la gestión de la red, haciéndola más proactiva. Al explotar las sinergias en la adquisición de datos en la nube y la inteligencia artificial sofisticada, permite la creación perfecta y la evolución continua de un Network Digital Twin, una herramienta indispensable para navegar las complejidades de las infraestructuras de red modernas.

Executive Summary

This document presents the initial architecture of a monitoring solution that enables the maintenance of a reliable edge cloud digital twin. The architecture proposed in this deliverable seamlessly integrates data procurement from edge cloud environments, leveraging cutting-edge technology to empower an artificial intelligence (AI) system. The aim is to facilitate the continuous generation and updates of an edge Network Digital Twin.

In summary, the architecture presented in this deliverable 6G-EDGEDT-03-E12 introduces the following capabilities:

- Extraction of real-time and historical data from diverse edge cloud environments, ensuring comprehensive coverage of network metrics, performance indicators, and device statuses.
- Leverages state of the art in Artificial Intelligence to enable the creation of a digital twin capable of reliably mirroring the complex physical network environment.
- Uses a modular approach to facilitate scalability and makes the solution futureproof in that it is adaptable to advances beyond the current state of the art.

The architecture proposed in this document enhances network management, making it more proactive. Exploiting synergies in edge cloud data procurement and sophisticated artificial intelligence, it allows for the seamless creation and continuous evolution of a Network Digital Twin—an indispensable tool for navigating the complexities of modern network infrastructures.

…

1. Introduction

Given the growing complexity in the topology of modern networks including semi-autonomous edge networks with enhanced capabilities to support niche services, making configuration changes directly on production setups may lead to adverse unintended consequences. However, to accommodate new low latency use-cases such configuration changes may be inevitable. In these situations, having an accurate, up to date, virtual representation (digital twin) of the edge network is vital. The latter facilitates the testing of configuration changes before they are deployed in production. The more up to date the digital twin is, the better it represents the current state of the network and the more reliable the predictions it can make regarding the impact on Service Level Agreements (SLAs) that proposed configuration changes may have on new and existing services.

In order that the digital twin is kept up to date, telemetry data collected from components of the real network need to regularly update the state of their virtual counterparts in the twin. Therefore, agile monitoring solutions capable of efficiently collecting, processing and updating digital twins with these data are needed particularly the stringent low latency requirements of the services deployed at the edge network. It is also especially crucial if digital twins are part of zero touch closed loop automated management and orchestration.

This work envisions the creation and implementation of a distributed monitoring and data management system. The latter should be capable of gathering data and key performance indicators (KPIs) from various data sources, such as L2, L3, MANO, and others. Data collected will be used to support the training of Digital Twins.

A telemetry collection solution that allows algorithms to predict network behaviour based on its current status is also presaged. To support diverse data sources, technology-specific plugins will be developed. These plugins will be responsible for translating and adapting communication protocols used by field buses or native monitoring functions into a common information model and distribution method. This ensures that data can be collected, shared, and distributed in a standardized manner.

A hierarchical monitoring system will be established to support distributed and synchronized storage as well as efficient data streaming. The latter will be designed to integrate and complement open-source software for time-series databases, distributed streaming buses, monitoring dashboards and other such systems.

User access control mechanisms will be achieved through open and secure interfaces that support both query and subscribe/notify patterns.

2. State of the art in monitoring solutions

Monitoring solutions can be broadly grouped into three categories; collectors, processors and exporters. Collectors extricate telemetry data from network nodes either by push operations whereby the nodes themselves regularly send data to the component or pull operations whereby the collector itself draws out the data from the nodes. Processors transform the collected data in a variety of ways such as filtering to remove duplicates, altering the format to make it suitable for distribution or cross-checking data from multiple sources to establish coherence among others. They can also send out alerts should anomalies are detected in the data. Exporters are responsible for the distribution of processed data to other entities.

Under the auspices of the EU H2020 5G-DIVE [1] and 5Growth [2] projects, solutions such as OpenFlowMon (Antonio Cobos, 2021) and DEEP (Carlos Guimarães, 2021) have been proposed which leverage open-source tools to monitor network resources at both compute and device level. Figure 2-1 and Figure 2-2 posit the use of monitoring in 5GROWTH (Papagianni, 2020) and 5G-DIVE (5G-DIVE, 2021) respectively. The solutions proposed in these projects are geared towards pre-emptive maintenance. Our architecture presented in this deliverable enhances the ideas proposed to produce robust digital twins capable of closed-loop autonomy thereby eliminating the need for a human in the loop.

[1] https://cordis.europa.eu/project/id/859881
[2] https://5growth.eu/

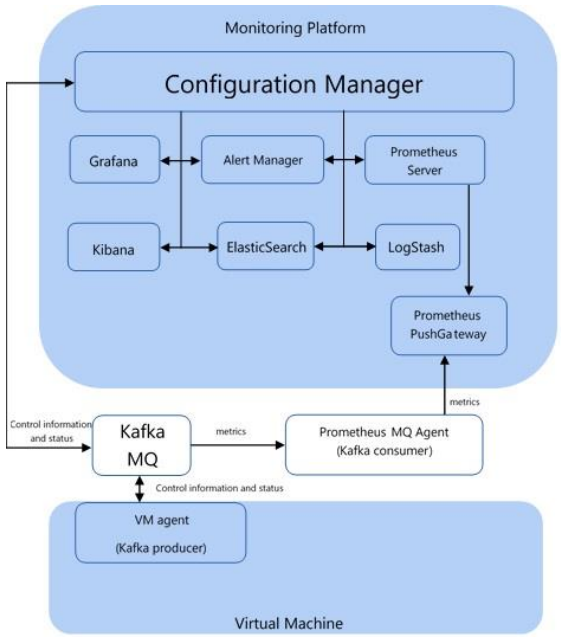

FIGURE 2-1 5GROWTH VOMS ARCHITECTURE (PAPAGIANNI, 2020)

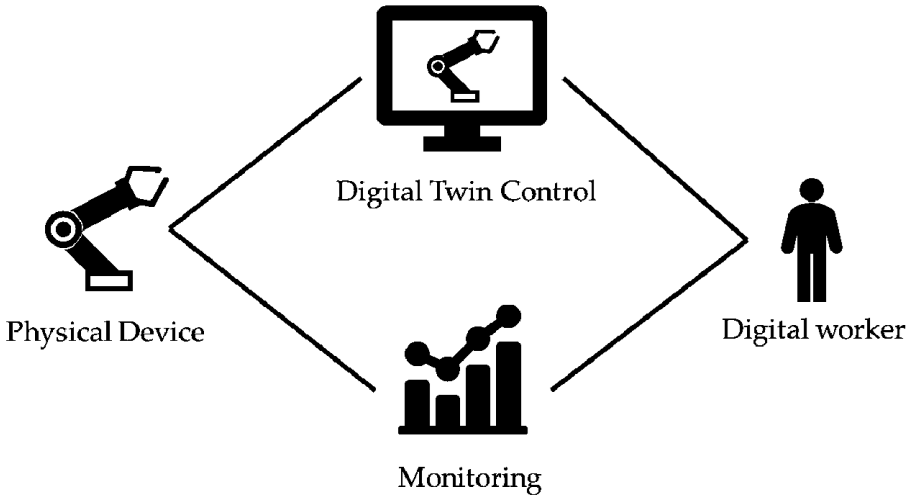

FIGURE 2-2 5G-DIVE MONITORING (5G-DIVE, 2021)

2.1. Collectors

A well-known open-source network monitoring and accounting tool is pmacct[3]. It is designed for collecting, processing, and storing traffic data and is particularly focused on network traffic monitoring, accounting, and billing. It is widely used by Internet Service Providers (ISPs), network administrators, and telecom operators to gain insights into network usage patterns, analyse traffic flows, and perform billing and accounting tasks. It supports multiple protocols for traffic acquisition, including NetFlow, sFlow, IPFIX etc.

Telegraf[4] is an open-source plugin-driven server agent designed for collecting, processing, and sending metrics and data from various sources to databases. It is built around a plugin system, allowing users to extend its capabilities by adding input, output, and processor plugins. There is a wide range of officially supported plugins for collecting metrics from various sources and sending data to various destinations. It can collect data from a multitude of sources, including system metrics (CPU, memory, disk), network devices, databases, web services, and many more. It supports protocols like SNMP, JMX, HTTP, MQTT, and more for data retrieval. It also includes processors that can manipulate and transform data before it is sent to a database. For example, filters can be used to aggregate, or modify metrics. Its plugin architecture makes Telegraf highly extensible. Users can create custom plugins to collect data from specific sources or send data to custom destinations. Telegraf is cross-platform and can be installed and run on a variety of operating systems, including Linux, Windows, macOS among others.

Prometheus[5] is an open-source monitoring and alerting toolkit designed for reliability and scalability. It is primarily used for collecting and processing metrics and time-series data from various systems, services, and applications. It uses a pull-based model, periodically scraping data from target systems and services. This allows it to collect metrics from a wide range of sources, including application endpoints, databases, hardware, and more. It includes a built-in alerting system that can be configured to trigger alerts based on defined thresholds and conditions. It can send alerts via various channels including email, Slack, or webhooks. Prometheus provides service discovery mechanisms to dynamically discover and monitor targets, making it suitable for changing environments and infrastructure. Prometheus is commonly used in cloud-native and containerized environments, where it can monitor the performance and health of microservices, containers, and Kubernetes clusters. It is also used in traditional infrastructure monitoring.

[3] http://www.pmacct.net/
[4] https://www.influxdata.com/time-series-platform/telegraf/
[5] https://prometheus.io/

2.2. Data Storage and processors

InfluxDB[6] is an open-source, high-performance, and scalable time-series database designed for efficiently storing, querying, and processing time-series data. It is optimized for storing time-series data, that is indexed and queried based on timestamps. This makes it particularly well-suited for applications involving metrics, monitoring, IoT (Internet of Things), and sensor data. It is designed to handle high write and query loads and can be easily scaled horizontally by adding more nodes to a cluster. This scalability is crucial for systems that generate large volumes of time-series data. InfluxDB is built for speed and can handle millions of data points per second. It uses an efficient storage format and indexing mechanisms to achieve fast read and write operations. To limit the use of resources, it allows users to define retention policies, which determine how long data is retained in the database before it is automatically deleted It incorporates a SQL-like query language (InfluxQL) specifically tailored for time-series data. It allows users to perform complex queries and aggregations on their data. It supports integrations and plugins to collect data from different sources, including Telegraf.

2.3. Exporters

OpenTelemetry[7] is a key solution that does not fit snugly into one category given that it performs all three functions of collecting, processing and exporting data. It is vendor neutral and can work with other monitoring solutions. It is an open-source project that provides a set of APIs, libraries, agents, and instrumentation to enable observability in software applications. Its aims to standardize how telemetry data (e.g., metrics, traces, and logs) is collected and exported from applications and services. It is mainly used to gain insights into the performance, behaviour, and health of distributed systems, making it a critical component of observability and monitoring solutions.

OpenTelemetry supports the collection of application and system-level metrics, such as CPU usage, memory consumption, request rates, and error rates. These metrics can be used to monitor the performance and health of applications. It supports various trace data exporters, allowing trace data to be sent to various backends and observability platforms, such as Jaeger[8], Zipkin[9], and Prometheus. It also supports vendor-agnostic formats like OpenTelemetry Protocol (OTLP)[10]. It is used in modern cloud-native, microservices, and containerized environments, where understanding the behaviour and performance of distributed systems is crucial. By providing standardized telemetry data collection and propagation, OpenTelemetry simplifies the task of building, monitoring, and

[6] https://www.influxdata.com/
[7] https://opentelemetry.io/
[8] https://www.jaegertracing.io/
[9] https://zipkin.io/
[10] https://opentelemetry.io/docs/specs/otel/protocol/

troubleshooting complex applications and services. It is often integrated with monitoring and observability platforms, making it easier to visualize and analyse telemetry data.

Apache Kafka[11] is an open-source stream processing platform and message broker that used for building real-time data pipelines and streaming applications. Kafka is designed to handle high volumes of data and real-time event streaming. It uses a publish-subscribe messaging model, where producers publish messages to topics, and consumers subscribe to topics to receive and process those messages. It is designed to be distributed and can scale horizontally across multiple servers or nodes. This allows it to handle large amounts of data and high-throughput workloads. It allows data durability by persisting messages to disk, meaning that should a consumer fail or a message not be immediately processed, it can be retrieved later. It is also fault-tolerant and can continue to operate even if some of its nodes or brokers fail. This reliability is crucial for mission-critical applications. Kafka enables real-time stream processing and is known for its high throughput and low latency, making it suitable for use cases where real-time data ingestion and processing are essential. It stores data as an immutable log of records, which allows for easy replay of events and can be used for auditing and tracking.

2.4. Technical gaps in state of the art

Several technical gaps exist in current state-of-the-art network monitoring solutions, creating challenges and opportunities for improvement. Here is an outline of some of these technical gaps:

1. *Scalability Issues*:
   Challenge: Many network monitoring solutions struggle to scale efficiently as network sizes and data volumes increase.
   Opportunity: Developing scalable architectures and algorithms to handle the growing complexity and volume of data in large-scale networks.
2. *Real-Time Monitoring Challenges:*
   Challenge: Achieving real-time monitoring, especially in environments with high-speed networks and diverse data sources, remains a challenge.
   Opportunity: Advancing technologies and methodologies for real-time data processing and analysis to provide instant insights into network performance.
3. *Diversity in Data Sources:*
   Challenge: Network environments generate diverse data from various sources, including IoT devices, cloud services, and different protocols, making it challenging to integrate and analyse all data effectively.
   Opportunity: Developing standardized protocols and approaches for collecting, aggregating, and analysing heterogeneous data sources.

[11] https://kafka.apache.org/

4. *Security and Privacy Concerns:*
   Challenge: Ensuring the security and privacy of sensitive network data is a significant concern, especially as monitoring solutions collect and process increasingly valuable information.
   Opportunity: Implementing robust security measures, encryption techniques, and privacy-preserving algorithms to protect sensitive network information.
5. *Dynamic Network Environments:*
   Challenge: Networks are becoming more dynamic with constant changes in configurations, devices, and traffic patterns, posing challenges for traditional monitoring solutions.
   Opportunity: Adapting monitoring solutions to dynamically changing environments through automated discovery, self-adjusting algorithms, and machine learning models.
6. *Interoperability:*
   Challenge: Lack of standardization and interoperability among different network monitoring tools and devices hinders seamless integration.
   Opportunity: Developing and adopting common standards to enable interoperability between diverse monitoring solutions and components.
7. *Complexity in Fault Detection:*
   Challenge: Identifying and diagnosing complex network faults, especially in distributed and cloud-based architectures, is a significant technical challenge.
   Opportunity: Utilizing advanced machine learning techniques, anomaly detection algorithms, and predictive analytics to enhance fault detection capabilities.
8. *Inadequate User Interfaces:*
   Challenge: User interfaces of some network monitoring tools may lack intuitiveness and comprehensive visualization, making it challenging for operators to interpret and act upon data effectively.
   Opportunity: Improving user interfaces with interactive and customizable dashboards, intuitive visualization, and user-friendly reporting tools.
9. *Resource Overhead:*
   Challenge: Some monitoring solutions introduce significant resource overhead, impacting network performance and scalability.
   Opportunity: Optimizing resource utilization through lightweight monitoring agents, efficient data compression, and distributed architectures.
10. *Lack of Predictive Analytics:*
    Challenge: Limited incorporation of predictive analytics in network monitoring solutions, hindering the ability to anticipate issues before they impact network performance.
    Opportunity: Integrating machine learning models and predictive analytics for proactive identification and mitigation of potential network issues.

Addressing these technical gaps in network monitoring solutions will contribute to the development of more robust, scalable, and efficient tools for managing modern and complex network infrastructures.

3. Initial design of monitoring solution for edge network DT

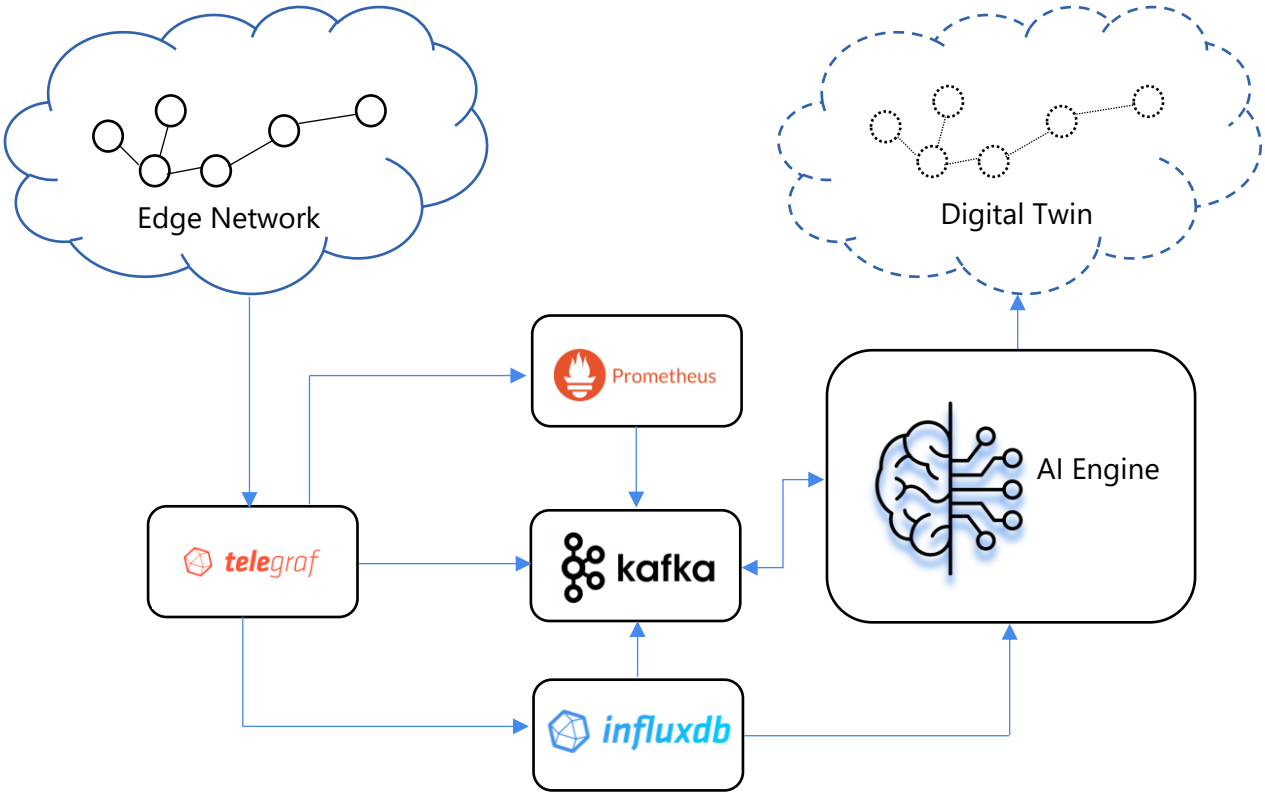

FIGURE 3-1 INITIAL ARCHITECTURE OF MONITORING SOLUTION

The system would be comprised of Telegraf plugins that pull telemetry data from the physical and/or virtual network elements. Application data will be forwarded to Prometheus to leverage its service monitoring and alerting capabilities. Prometheus will discover and scrape metrics from Telegraf agents by defining them as targets. Prometheus will serve as a key data pipeline to the Kafka streaming bus. Telegraf agents will also directly send metrics to influxdb.

The AI driven Digital Twin modelling component will glean network topology data from the Kafka bus. The status of the abstracted digital twin will constantly be updated by the network traffic time series data extracted from influxdb.

3.1. Detail of monitoring solution

This network monitoring solution utilizes a combination of tools and technologies to collect, process, store, and analyse data, ultimately feeding it into an AI engine to generate a digital twin of the network. Here's a breakdown of the components and their roles:

1. *Telegraf*: serves as the data collection agent responsible for gathering metrics and performance data from various sources within the network, including servers, applications, and network devices. It collects real-time metrics such as CPU usage, memory utilization,

network traffic, and other relevant statistics. It acts as the initial point of contact for ingesting data into the monitoring solution.

2. *Prometheus*: functions as a monitoring and alerting system. It stores and processes the collected metrics from Telegraf, allowing for service-level monitoring. It utilizes a powerful query language (PromQL) to analyse and retrieve specific metrics, and it also provides alerting capabilities to notify administrators or automated systems when predefined thresholds are exceeded.
3. *InfluxDB*: is utilized for efficient storage and processing of time-series data. Telegraf sends the collected metrics to InfluxDB, where they are stored and processed. InfluxDB's capabilities include data retention policies, continuous queries, and efficient storage structures, making it well-suited for handling time-series data and enabling historical analysis.
4. *Kafka*: acts as a data export and streaming platform. Telegraf can export data to Kafka, which acts as a centralized hub for efficient and real-time data streaming. This allows for seamless communication between different components of the monitoring solution and provides scalability and fault tolerance for handling large volumes of data.
5. *AI Engine*: processes the data from Prometheus, InfluxDB, and potentially Kafka to generate a digital twin of the network: By leveraging machine learning algorithms and analytics, the AI engine analyses both historical and real-time data from the monitored network. It identifies patterns, predicts potential issues, and creates a digital twin—a virtual representation of the network's structure, behaviour, and potential future states.

3.2. AI Engine

We now discuss the AI modelling engine that is the key component that facilitates the creation *of the digital twin*. It comprises modules for data collection, data pre-processing, feature extraction, time series analysis, machine learning, Digital Twin Modelling, continuous learning, reporting and alerting. The table below outlines the roles and functionality of each of these modules. It addresses most of the challenges of state-of-the-art monitoring solutions by exploiting their synergy.

TABLE 3-1 AI ENGINE COMPONENTS

Function	Functionality	Components
Data Ingestion	Gathers diverse metrics, including network traffic, latency, error rates, and device statuses.	Interfaces with data streams from monitoring tools like Telegraf, Prometheus, or other data repositories like InfluxDB
Pre-processing and normalization	Handles missing values, normalizes data units, and ensures data quality before feeding it into the modelling pipeline	Data pre-processing modules
Feature extraction	Determines which metrics are most influential in representing the network's state and behaviour.	Algorithms for feature selection and extraction.
Time-Series Analysis	Recognizes recurring patterns, anomalies, and trends in the network metrics over time.	Time-series analysis algorithms e.g. (ARIMA)
Machine Learning Models	Learns patterns, dependencies, and correlations within the network data to create predictive models.	Multiple models for different aspects (e.g., regression models, clustering algorithms).
Digital Twin Generation	Synthesizes information from various models into a comprehensive representation of the network's current and potential future state.	Integration layer
Continuous learning	Dynamically updates the digital twin based on new data, ensuring the model remains accurate and reflective of the evolving network	Online learning algorithms
Visualization and reporting	Creates dashboards, charts, and reports to enable easy interpretation of the digital twin's insights	Visualization tools and reporting interfaces
Alerting mechanism	Triggers alerts based on predefined thresholds, helping operators take timely actions	
Integration with Management systems	Ensures seamless integration with existing network infrastructure, facilitating coordination with network management systems	API Connectors

4. Summary and conclusions

This deliverable presents a network monitoring solution that provides a comprehensive approach to data collection, service monitoring, storage, and analysis, with the added capability of exporting data through Kafka. The integration of an AI engine enhances the solution by creating a digital twin, enabling predictive insights into the network's performance and behaviour.

This architecture also ensures that the AI engine continuously leverages network monitoring data to create and update an accurate digital twin, providing actionable insights and predictions for effective network management.

5. References

5G-DIVE, 2021. *5G-DIVE architecture and detailed analysis of vertical use cases.*. [Online] Available at: https://euprojects.netcom.it.uc3m.es/5g-dive/wp-content/uploads/2021/05/D1.1_Final_updated-with-review-comments.pdf [Accessed 28 11 2023].

Antonio Cobos, C. G. A. D. L. O. A. Z., 2021. *OpenFlowMon: A Fully Distributed Monitoring Framework for Virtualized Environments.* Heraklion, Greece , s.n.

Carlos Guimarães, M. G. L. C. A. Z. L. M. C. S. T. T. C. Z. S. H. A. M. A. d. l. O., 2021. DEEP: A Vertical-Oriented Intelligent and Automated Platform for the Edge and Fog. *IEEE Communications Magazine,* 59(6), pp. 66-72.

Papagianni, C., 2020. *5Growth: AI-driven 5G for Automation in Vertical Industries.* Dubrovnik, Croatia, s.n.