



TALON

Autonomous and Self-organized Artificial Intelligent Orchestrator for a
Greener Industry 4.0

Deliverable

D3.4 Self-Healing, Self-Recovery & Self-Correcting Mechanisms
Toolkit

Actual submission date: 07/09/2024

Project Number:	101070181		
Project Acronym:	TALON		
Project Title:	Autonomous and Self-organized Artificial Intelligent Orchestrator for a Greener Industry 4.0		
Start date:	October 1st, 2022	Duration:	36 months
<i>D3.4 Self-Healing, Self-Recovery & Self-Correcting Mechanisms Toolkit</i>			
Work Package:	WP3		
Lead partner:	UNIVERSITAT POLITECNICA DE VALENCIA (UPV)		
Author(s):	Foivos Psarommatis Giannakopoulos (UPV); Stylianos Trevlakis (IC); Lamprini Mitsiou (IC); Vasileios Kouvakis (IC); Theodoros Tsiftsis (IC)		
Reviewers:	Ilias Siniosoglou (MINDS) Thomas Lagkas (DUTH)		
Due date:	31/08/2024		
Deliverable Type:	DEM	Dissemination Level:	PU
Version number:	1.0		

Revision History

Version	Date	Author	Description
0.1	01/08/2024	UPV	ToC release
0.2	06/08/2024	UBI	UBI contribution
0.3	26/08/2024	UPV	Contributions in all sections
0.4	27/08/2024	IC	IC contributions
0.5	27/08/2024	UPV	Revision
0.6	31/08/2024	IC, MINDS	Revision
0.7	05/09/2024	UPV	Revision
1.0	07/09/2024	ENG	Final coordinator review before submission

Table of Contents

Table of Contents	3
Table of Figures	6
Abbreviation list	7
1 Introduction	8
2 Background and Motivation	9
2.1 Network Resilience	9
2.2 Self-Healing and Self-Correcting Mechanisms	9
2.2.1 AI and Machine Learning Algorithms	10
2.2.2 Edge Computing	10
2.2.3 Predictive Maintenance	10
2.2.4 Automated Incident Response	10
2.2.5 Self-Optimizing Networks (SON)	10
2.2.6 Integration with Cloud Services	11
2.2.7 Security Enhancements	11
2.2.8 Human-AI Collaboration	11
2.3 Importance of AI in Network Management	11
3 System Architecture	13
3.1 Overall tool description and architecture	13
3.1.1 Edge Devices	13
3.1.2 Kubernetes Cluster	14
3.1.3 AI Models	14
3.1.4 Decision Support System (DSS)	14
3.1.5 Monitoring Tools	14
3.1.6 Overall architecture	15
3.2 Edge-to-Cloud Network Monitoring and Data Persistence	17
3.2.1 Data Collection at the Edge	17
3.2.2 Real-Time Data Analysis and Processing	17
3.2.3 Data Persistence and Historical Analysis	17
3.2.4 Synchronization Between Edge and Cloud	17
3.2.5 Scalability and Flexibility	17
3.2.6 Enhanced Reliability Through Redundancy	17
3.2.7 Security and Privacy Considerations	18

3.2.8	Continuous Improvement and Adaptation	18
3.3	Traffic Management	18
4	AI Analytics for Network Reliability	20
4.1	Data Processing Techniques	20
4.2	Identifying Network Failures	20
4.2.1	Data Collection and Preprocessing	20
4.2.2	Model Selection and Training	21
4.2.3	Model Deployment	21
4.2.4	Evaluation and Monitoring	21
4.2.5	Common network failures and problems.....	23
4.2.6	Scenarios of Network Failures	24
5	Proactive Mechanisms for Failure Mitigation	26
5.1	Design of AI-Supported Mechanisms	26
5.1.1	Data Acquisition and Integration.....	26
5.1.2	Machine Learning Model Selection	26
5.1.3	Anomaly Detection	26
5.1.4	Predictive Maintenance	26
5.1.5	Automated Response Systems	27
5.1.6	Continuous Learning and Adaptation	27
5.1.7	Integration with Existing Network Management Systems	27
5.1.8	Security and Privacy Considerations	27
5.1.9	Performance Metrics and Evaluation	27
5.1.10	Scalability and Flexibility	28
5.2	Development of Proactive Solutions.....	28
5.2.1	Solution Overview and Goals	28
5.2.2	Overall Architecture	28
5.2.3	Workflow Description.....	28
5.2.4	Integration and Automation	29
5.2.5	Evaluation and Continuous Improvement	29
5.3	Proactive Network Resources Allocation	30
6	ML model update mechanisms.....	33
6.1	P2P learning	33
6.2	Federated learning.....	34
6.3	Distributed learning.....	35

7	Self-Healing and Self-Correcting Mechanisms	37
7.1	TALON NG-SDN Programmability	37
7.2	Minimizing Network Recovery Time	38
7.2.1	Advanced Monitoring and Alerting Systems	38
7.2.2	Automated Diagnostic Tools	38
7.2.3	Pre-configured Recovery Protocols	38
7.2.4	Redundancy and Failover Mechanisms	38
7.2.5	Decentralized and Edge Computing Architectures	38
7.2.6	Network Segmentation and Microsegmentation	39
7.2.7	Training and Simulation	39
7.2.8	Post-Incident Analysis and Continuous Learning	39
7.2.9	Importance of Rapid Recovery	39
7.2.10	Proactive Monitoring and Real-Time Detection	39
7.2.11	Automated Response Mechanisms	40
7.2.12	Efficient Resource Allocation and Redundancy	40
7.2.13	Continuous Improvement and Learning	40
7.3	Dynamic QoS Adaptation Mechanisms	40
7.4	AI-fueled Network Policies Enforcement	43
8	Usage Scenarios and Use Cases	47
8.1	Application of Self-Healing Mechanisms	47
8.1.1	Data Collection	47
8.1.2	Anomaly Detection	48
8.1.3	Automated Corrective Actions	49
8.1.4	Monitoring Dashboard	50
8.1.5	Backend API for Dashboard	51
8.2	Performance Metrics and Evaluation	52
8.2.1	Detailed Simulation Results	52
9	Implementation and Testing	57
9.1	Testing Methodology	57
9.2	Results and Analysis	58
10	Conclusions and Future Outlook	62
11	References	63

Table of Figures

Figure 1: Self healing and self correcting Kubernetes cluster	16
Figure 2: solution architecture	30
Figure 3: Proactive Network Resources Allocation	31
Figure 4: P2P learning setup	34
Figure 5: Federated learning setup	35
Figure 6: Distributed learning setup	36
Figure 7: Network Policies Definition	41
Figure 8: Application of Bandwidth Policies	41
Figure 9: Bandwidth received	42
Figure 10: bandwidth received	43
Figure 11: Data collection Bandwidth	44
Figure 12: feature values	45
Figure 13: Feature Engineering & Model Training	46
Figure 14: Data Collection	48
Figure 15: anomaly detection	49
Figure 16: Automated Corrective Actions	50
Figure 17: Monitoring Dashboard	51
Figure 18: Backend API for Dashboard	52
Figure 19: MTTD and MTTR comparison	55
Figure 20: System uptime over six months	56
Figure 21: user interface of the development instance	59
Figure 22: cluster management functions	60
Figure 23: Monitoring and Alerting (Prometheus)	60
Figure 24: Health monitoring and alerting UI (Grafana)	61
Figure 25: Rule based mitigation action enforcer (Node-red)	61

Abbreviation list

<i>AI</i>	<i>Artificial Intelligence</i>
<i>API</i>	<i>Application Programming Interface</i>
<i>CNN</i>	<i>Convolutional Neural Network</i>
<i>CPU</i>	<i>Central Processing Unit</i>
<i>DNN</i>	<i>Deep Neural Network</i>
<i>DSS</i>	<i>Decision Support System</i>
<i>E2C</i>	<i>Edge-to-Cloud</i>
<i>ELK</i>	<i>Elasticsearch, Logstash, Kibana</i>
<i>EMI</i>	<i>Electromagnetic Interference</i>
<i>FL</i>	<i>Federated Learning</i>
<i>FHSS</i>	<i>Frequency Hopping Spread Spectrum</i>
<i>IIoT</i>	<i>Industrial Internet of Things</i>
<i>KPI</i>	<i>Key Performance Indicator</i>
<i>LSTM</i>	<i>Long Short-Term Memory</i>
<i>ML</i>	<i>Machine Learning</i>
<i>MTTD</i>	<i>Mean Time to Detect</i>
<i>MTTR</i>	<i>Mean Time to Repair</i>
<i>NG-SDN</i>	<i>Next Generation Software Defined Network</i>
<i>P2P</i>	<i>Peer-to-Peer</i>
<i>QoS</i>	<i>Quality of Service</i>
<i>RTT</i>	<i>Round Trip Time</i>
<i>SCADA</i>	<i>Supervisory Control and Data Acquisition</i>
<i>SDN</i>	<i>Software Defined Network</i>
<i>SON</i>	<i>Self-Optimizing Networks</i>

I Introduction

Objective 1: Enhancing Network Reliability and Performance

Industrial Networks — The Main Goal of the self-healing and auto tuning tool for industrial networks is to improve overall network reliability, uptime & performance. Industrial networks are essential to the operation of manufacturing plants, utilities and many other industries that need precise communication without interruption. Using AI-driven models to monitor and analyze traffic flow, the solution is designed to track network performance over time in an effort to catch latency, packet duplication loss problems as they arise. This preventative method guarantees that potential issues are resolved before they turn into large disruptions, keeping network operations running in a smooth and efficient manner. Providing this helps to create a more rugged network environment that can be used by the most demanding industrial applications, ultimately reinforcing design uptime and enabling higher productivity.

Objective 2: Automating Network Management

Also to focus on automating network management workflows with minimal manual interventions as one among the key objectives. These networks are often wide and complicated, making it challenging to monitor manually and perform a diagnosis that will include human errors. Utilizing AI and machine learning, the system can automate which you observe network issues; diagnose when there's an issue that could make a consumer feel anguish and even execute actions to restore it. It involves a range of activities from changing configurations and re-routing traffic to applying software changes, or rebooting equipment that has developed issues. By automating these operational tasks that are at the heart of maintaining a reliable network, response time to problems is expedited and accuracy heightened — ultimately being one step closer to freeing up IT operations resources from orchestrate functions and steering them toward innovation projects.

Objective 3: Enhancing Decision-Making with Data-Driven Insights

The Decision Support System (DSS) of the self-healing tool is designed to improve decision-making by providing data-driven insights. Based on AI model outputs and rule-based logic, the DSS automatically suggests the most efficient ways to alleviate certain network issues. According to the same article, The DSS can provide actionable alerts using a combination of historical data, real-time analytics that find outliers against normal behaviors and predefined rules. This is a big advantage, particularly in industrial applications where timely and informed decisions can mean the difference between small issues that are easily resolved into cataclysmic failures. They aim to provide network administrators with the tools necessary so that they can be making more educated, tactical decisions overall for healthier, resilient configurations.

Objective 4: Measuring and Improving Network Health Over Time

One important goal of the tool is to monitor and improve network health over time. Monitoring tools that cover latency, packet loss, resource utilization and dozens of other performance metrics-data which is vigorously collected by the system. These numbers give a good overview of the state and also show emerging issues/ trends. Based on these, the system can adapt and optimize its strategies to match wider network health over time through regular performance evaluation of self-healing actions. This iterative process continues to enhance the strength of the network keeping pace with changing requirements from industrial applications, which contribute operations efficiency and ensure lasting stability.

2 Background and Motivation

2.1 Network Resilience

Resilience is the ability of a system to recover into an operating state, even after failures or when faced with extreme situations. This is different than a robust system, where the objective was not to prevent all issues that could go wrong but rather surviving with or without some of them. However, it would fall apart when things break in other ways. Imagine an example where a hypothetical farmer would protect his crops from fire, floods or local pests, but they are in turn hit by some foreign plant virus wilting out all the crops. Surprisingly, in certain cases, a robust system could be more susceptible due to its inflexibility and intricacy [22].

Modern communication networks, electric power grids, and manufacturing lines rank among the most extensive and intricate manmade systems. Their decentralized structure contributes to the complexities of managing and predicting outcomes. Therefore, basic robust design principles like redundancy are insufficient to guarantee the ultra-reliable performance required for critical future applications such as real time self-healing and self-recovery [23].

On the contrary, resilient systems are devised to recover from completely unforeseen catastrophic events [24]. In the previous example, through diversifying crops a farmer can prevent an emerging plant virus from decimating the crop yield. Resilient systems typically adhere to the following design principles:

- **Real-time monitoring and reconfiguration:** Resilient systems need to be able to respond promptly to changes by closely monitoring the system and identifying alterations early on. An automated system for detecting anomalies can learn about behavior in real-time and spot any deviations giving early warnings even in unexpected situations [25]. When linked to a feature it can also activate automatic corrective or mitigating measures. This idea is further explored in TALON in the context of self-healing functionalities based on anomaly detection and diagnosis.
- **Decentralization:** Besides having backup capacity for redundancy, resilient systems often showcase a decentralized structure. One illustrative example can be found in 5G-RAN, where the RAN multi-connectivity method, which is usually used to enhance throughput, can also take advantage of the macro-diversity effect from simultaneous connections; thus, increasing the chances of maintaining at least one strong connection.
- **Detailed investigation:** Resilient systems concentrate on the area upon detecting changes to tackle issues or adapt [26]. In network management, extra resources can be dispatched to locations where unforeseen events arise. This is especially pertinent in telco cloud setups, where fault management functions can be directed to areas posing more problems.
- **Simplicity at the core, yet diversity at the edge:** Among these things, we can see that while resilient systems are complex they also have common simple features. In this instance, internet access can be highly varied and by some accounts quite complex. However, they share the same protocols. Mobile networks in the future are seen using a common data system to streamline core operations. The setup enabled data providers and users to interact in a cloud with service-oriented architecture principles, where they also had access to the common shared-based platform for exchanging data.

2.2 Self-Healing and Self-Correcting Mechanisms

Self-healing networks represent the pinnacle of network resilience and reliability, aiming to autonomously detect, diagnose, and resolve issues without human intervention [20]. These are using sophisticated technologies like AI, machine learning and edge computing to ensure optimal network performance

continuously operational. In industrial environments, where downtime can be very expensive and dangerous to people (and not just an inconvenience; it could send someone for a far-too-close encounter with forces of nature rendering them fatal), such self-healing mechanisms are even more useful.

2.2.1 AI and Machine Learning Algorithms

In simple terms, self-healing networks consist of AI and ML algorithms that work as a police unit which can detect anomalies in the network, predict failures or recommend/initiate corrective actions. Millions of data points being captured by the network are really challenging to detect, analyze and respond if something goes wrong; hence techniques like supervised training approaches unsupervised methods or even reinforcement learning is used for this purpose. They are in the form of models like Long Short-Term Memory (LSTM) networks, Convolutional Neural Networks (CNN), and Autoencoders that serves purposes such as time-series forecasting, anomaly detection or pattern recognition. These models learn from more and more data to keep getting better at their predictions.

2.2.2 Edge Computing

Self-Healing Networks With Edge computing, you can also empower your self-healing networks to process data near its source. We then generate deltas for the current observer and lesser than them will get transmit to reduce latency, bandwidth as this results in faster network issue detection. For instance, edge devices like sensors and gateways carry embedded AI models which are capable of completing on-the-spot data analysis and carrying out initial corrective actions. This decentralized structure allows the network to be autonomous and robust, even without central systems communicating.

2.2.3 Predictive Maintenance

Keep in mind that self-healing industrial networks also rely heavily on predictive maintenance. Predictive models dealing with the behavior of network components — routers, switches and servers for instance- can help to forecast when those devices might fail based on historical data as well a real-time performance metrics. This conducts timely repair and replacement, avoid accident broken down. Predictive maintenance models — Predictive maintenance is focused on developing techniques (e.g., regression analysis, survival analysis, neural networks) which are able to predict the remaining useful life of network equipment.

2.2.4 Automated Incident Response

Incident response: Automated incident responders are created to respond immediately when network irregularities or faults appear on the last part of our pyramid, automatically invoking a set of actions made by software.rdfshield This is done by appraising rule-based engines and other AI-driven decision-making processes to identify an optimal solution of what needs to happen. E.g., if a network segment is experiencing high latency, the system may attempt re-routing traffic through other paths, fiddling with QoS (Quality of Service) settings or restart certain affected devices. Rapid detection and resolution of network issues: Automation eliminates the time to detect a problem therefore reducing its impact on operations.

2.2.5 Self-Optimizing Networks (SON)

Self-healing is largely assumed to be in place, but self-optimizing networks take that next step and make dynamic adjustments of configurations and resources to remain running at full potential. These systems will constantly scan the network conditions and adjust bandwidth allocation, routing protocols etc to maintain throughput at a optimum level while keeping latency as low as possible but without creating congestion. By identifying the best possible configuration in present and past network data, machine learning algorithms help optimize all parameters for proper functioning of a network under different conditions.

2.2.6 Integration with Cloud Services

Adding self-healing features on top of cloud services affords you more scalability, flexibility and processing power. For instance, awe-inspiring designs and huge datasets can be trained on these cloud-based AI platforms as this may become impossible if you were to accomplish them using local infrastructure only. Additionally, cloud services provide redundancy and disaster recovery solutions ensuring that network management systems stay up even when local failures occur. Approaches that hybrid edge computing with cloud resources give you the best of both worlds: rapid local processing and powerful centralized analytics.

2.2.7 Security Enhancements

Security in self-healing networks is one of the most important things to keep forefront. Real-time cyber threat detection and response capabilities: By using AI-driven security mechanisms, the network will be made stronger to even rare types of online attacks. The security breaches by using different techniques such as anomaly detection, behavioral analysis, and threat intelligence integration are identified. Automated responses can include segmenting affected portions, blocking malicious traffic and applying security patches. Having strict security measures successfully protects the network from attacks and forges trust in a self-healing system.

2.2.8 Human-AI Collaboration

Self-healing networks are designed to reduce human intervention, but it does not eliminate the importance of having that experienced hand in an emergency. While AI-driven systems bring us many insights and recommendations, humans are here to verify that such actions align with our organizational goals and operational constraints. This makes a hybrid AI — network administrator agreement possible, where mundane issues are dealt with entirely autonomously by the systems and complex decisions require human-experience for deciding policy implementations as well managing strategy. This synergy helps to improve the efficiency and trust of network management.

2.3 Importance of AI in Network Management

In today's wireless systems, the traditional methods for optimizing resources are facing limitations and insurmountable complexities especially in challenging situations [19]. Many of the groundbreaking revelations that can be gleaned from network data depend on AI technologies such as neural networks and deep learning. These techniques can be used in different environments, make decisions and find the optimal ways for many complex problems. Therefore, it is very important to see how AI can solve many of the decision-making problems and improve carrier performance.

Communication, computing, and caching resources in 5G networks are divided into resource pools. The orchestrator of the system is in charge to coordinate all these resources and its running dependencies on necessity basis. The orchestrator uses AI algorithms to check the system resources and task attributes, enabling it to dynamically determine compute power, cache capability, and communication resource according to tasks.

When a Next Generation NodeB is to decide whether the task involves caching or not when it receives an user service request Although activities, such as video streaming services are cache-dependent tasks and quality inspection tasks in networks would not be considered dependent on a cache.

When it comes to streaming videos, a key aspect is storing videos on servers near users to cater to their video requests efficiently. In TALON in Pilot 1 with the swarm or drones feeding either images or videos, is very important that the network is performing as it supposed. Typically, the system caches high bitrate videos, while if a user asks for a lower bitrate video, the edge computing server can convert it as needed.

However, frequently requested low bitrate videos are also stored in the cache to lessen the load on edge computing. This whole process involves coordinating communication, caching, and computing resources.

For cache-dependent tasks, such as video streaming services, the system first checks if the requested video version is already stored in the cache. If it is, the video is transmitted to the user through a designated gNB. In cases where the requested or a higher bitrate version of the video is not cached, the system searches for a higher bitrate version and retrieves it from the source. Once either the original or a higher bit rate version is obtained, the task is passed to an orchestrator, which selects an edge server for transcoding purposes. This is done by the TALON orchestrator developed in T3.3. The orchestrator also determines whether caching of the transcoded video is necessary. If caching is required a caching server is assigned for this task. Ultimately, after all these steps are completed, the video is delivered to the user via their chosen gNB.

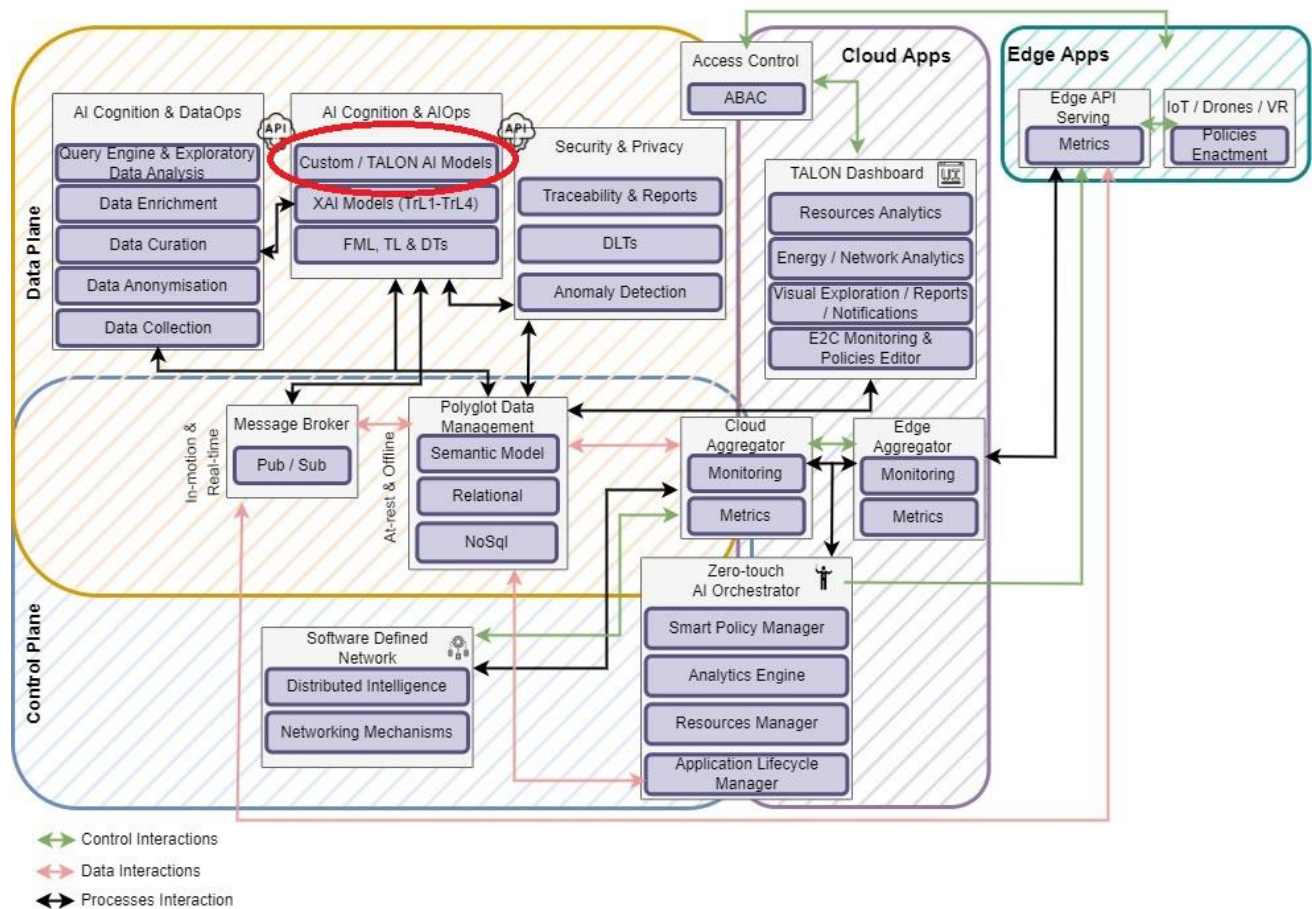
On the other hand, quality inspection tasks involve capturing product images, recognizing them and categorizing products based on what has been identified. This task requires coordination between communication and computing resources. In such cases, where caching is not relevant, the gNB forwards requests to the orchestrator. The orchestrator uses AI techniques like reinforcement learning to allocate computing and communication resources based on task specifics and available resources. Users send task-related data, such as product images, to their assigned gNB, which then passes this data to the chosen edge computing server for processing. The edge node carries out computations, like identifying product flaws from images, and sends back results through the gNB to users.

In order to enrich its models, TALON collects data from different sources and shares them with a number of different nodes that may be located at the edge or last mile. To minimize data pollution and thus maximize both data and energy efficiency optimized few-shot and federated learning approaches are utilized.

3 System Architecture

3.1 Overall tool description and architecture

The image below illustrates the TALON architecture. Marked with red circle is the area that the self healing tool belongs to.



3.1.1 Edge Devices

Some of the Role Edge Devices Play in Industrial Networks. The role that edge devices play in industrial networks is very important because they allow computational capabilities to come closer to where data originates, leading to shorter latency and lesser bandwidth use. Devices: sensors, gateways and industrial IoT (IIoT) equipment with embedded computing capabilities to collect data from surroundings or process local data snippets for anomaly detection. This enables the edge devices to identify problems quickly and respond immediately, a necessity in time-sensitive industrial environments. Also, these devices can execute tiny AI models to check network anomalies at the very initial stage so that only required data could be transferred and making it ready for central system on time as a result traffic is optimized and system performs more efficiently. This is extremely important in TALON because the various things that are occurring to different pilots quite depend on network throughput. However, when you deploy a fleet of edge devices other issues arise due to the fact that those thousands or millions instances need to secure

communications, be updated together – constantly and have some way of ensuring they are behaving in similar fashions; for this we require robust management & orchestration tools.

3.1.2 Kubernetes Cluster

Kubernetes is a resilient container platform that automates the deployment, scaling and operation of application containers across clusters of hosts. On an industrial network, Kubernetes orchestrates the deployment of AI models and decision support systems in order to provide HA applications that scale as needed. Built-in self-healing features from the platform automatically restarts containers in failed state and also re-schedule on different nodes if necessary, hence reduce downtime and ensure a reliable service. Service Discovery and Load Balancing: So, as a kubernetes provides service discovery so we need not worry about that it distribute the traffic across our cluster in order to reduce bottlenecks or any hot spots also this help us for better predictable performance. Although Kubernetes brings a big promise in terms of automation, portability and scalability it does come with its own complexity. You need to have a deep understanding of Kubernetes components like API server, scheduler and controller manager among others in order to set up and manage your own cluster well optimized with resource usage.

3.1.3 AI Models

Machine learning algorithms analyzes the network data to identify anomalies or potential issues, with AI models being a central component of TALON’s self-healing tool. These models, constructed with historical network data training the model to know what patterns lead to signs of trouble — like high latency, packet loss or throughput degradation. For example, LSTM networks are generally used for predicting time-series data, CNN is apt to capture spatial features and DNN can be employed in complex waveform recognition. This involves creating a dataset that can yield diverse images, processing this data to have the necessary quality and then splitting it into several sets: training, validation (development) and test. This is where hyperparameter tuning and iterative training come into the picture to optimize model performance. These models, when trained effectively can be deployed on edge devices in addition to cloud wherein they read network conditions and provide real-time insights. The major challenges to deploying AI models include generalizing the model well for test data and controlling computational resources during training, as well as reliability in decision-making process.

3.1.4 Decision Support System (DSS)

This is where the DSS comes into play: it acts as the brains of this self-healing system, receiving “readings” from the AI models and using its intelligence to make decisions on what actions should be taken. When there is an anomaly detected that you might want to respond to, it takes rule based logic (what we in the olden days called domain knowledge) and AI model predictions together help predict what your best action(s) should be. For example, if the model concludes that there are high latency issues you might get a recommendation from DSS like adding bandwidth or prioritizing some types of traffic. The DSS can (if so designed) also automatically enforce these corrective actions, minimizing the demands on human intervention and hence reducing response time. The technologies behind DSS are including, but not limited to rule engines (such as Drools) and integration platforms (e.g., APIs or microservices). However, designing a successful DSS requires developing complex and flexible rule sets that can integrate seamlessly with AI models such as convolution neural networks or other network management systems which will require large scale system to handle the complexities of licenses industrial networks..

3.1.5 Monitoring Tools

And the last but not least feature, setting up monitoring tools to have real-time overview on network state and making sure that actions from self-healing actually work. These tools collect metrics from different

network components (edge devices, models AI server edge, central systems) and aggregate them with in human readable dashboards. For instance, Prometheus is commonly employed for real-time monitoring and alerting; Grafana provides a robust visualization system so that you can create more informative boards. ELK Stack (Elasticsearch and Logstash, Kibana), a common log management/analysis tool providing granular oversight over system logs and performance data. It enables the users to track certain key performance indicators (KPIs) like latency, packet loss and throughput which help in understanding how healthy a network is and where all attention needs to be paid. But deploying these tools in a way that they can handle large data volumes and not impair network performance, is an exercise that requires thoughtful planning and skilled personnel.

3.1.6 Overall architecture

The diagram illustrates the technical architecture of the system designed to apply self-healing and self-correcting mechanisms in TALON, referred to as “TALON self-healing and self-correcting”. Based on Kubernetes, the system leverages continuous monitoring and AI-based recommendations to ensure self-healing and self-correcting capabilities in any Kubernetes-based distributed architecture. The process involves the following basic steps:

1. Providing a management layer to monitor and control distributed Kubernetes clusters using the Kubernetes API. This layer is provided by the open-source multi-cluster orchestration platform Rancher (<https://www.rancher.com/>).
2. Monitoring the health and performance of the distributed clusters, by integrating a monitoring system and time series database to the management layer (monitoring and alerting system). This functionality is provided by the open-source component Prometheus (<https://prometheus.io/>)
3. Trigger alarms whenever a specific situation is detected in any of the clusters (for instance, an alert when CPU usage in any Kubernetes nodes exceeds a predefined quota). This functionality is also provided by Prometheus
4. Enforce predefined rule-based mitigation actions whenever an alarm is triggered (for instance, restart services in Node when a CPU alert is triggered). This functionality is enforced by an ad-hoc application developed in Node.
5. Recommend new rules to enhance the self-healing and self-correcting capabilities of the system. This functionality is provided by an ad-hoc LLM agent developed using Langchain.

To ensure human oversight, the system provides the following User Interface (UI) Components:

1. **Health monitoring & alerting:** Dashboard system for health monitoring and alerting, providing graphic dashboards to monitor system health and send email notification alerts when alarms are triggered
2. **Mitigation Action Oversight:** Intuitive graphic interface for overseeing the rule-based mitigation actions that implement the self-healing and self-correcting capabilities of the system.
3. **Health-monitoring Co-pilot:** A web app to reason with the LLM to oversee system performance, understand dynamic self-healing and self-correcting capabilities needs, and suggest and approve new self-healing and self-correcting rules.

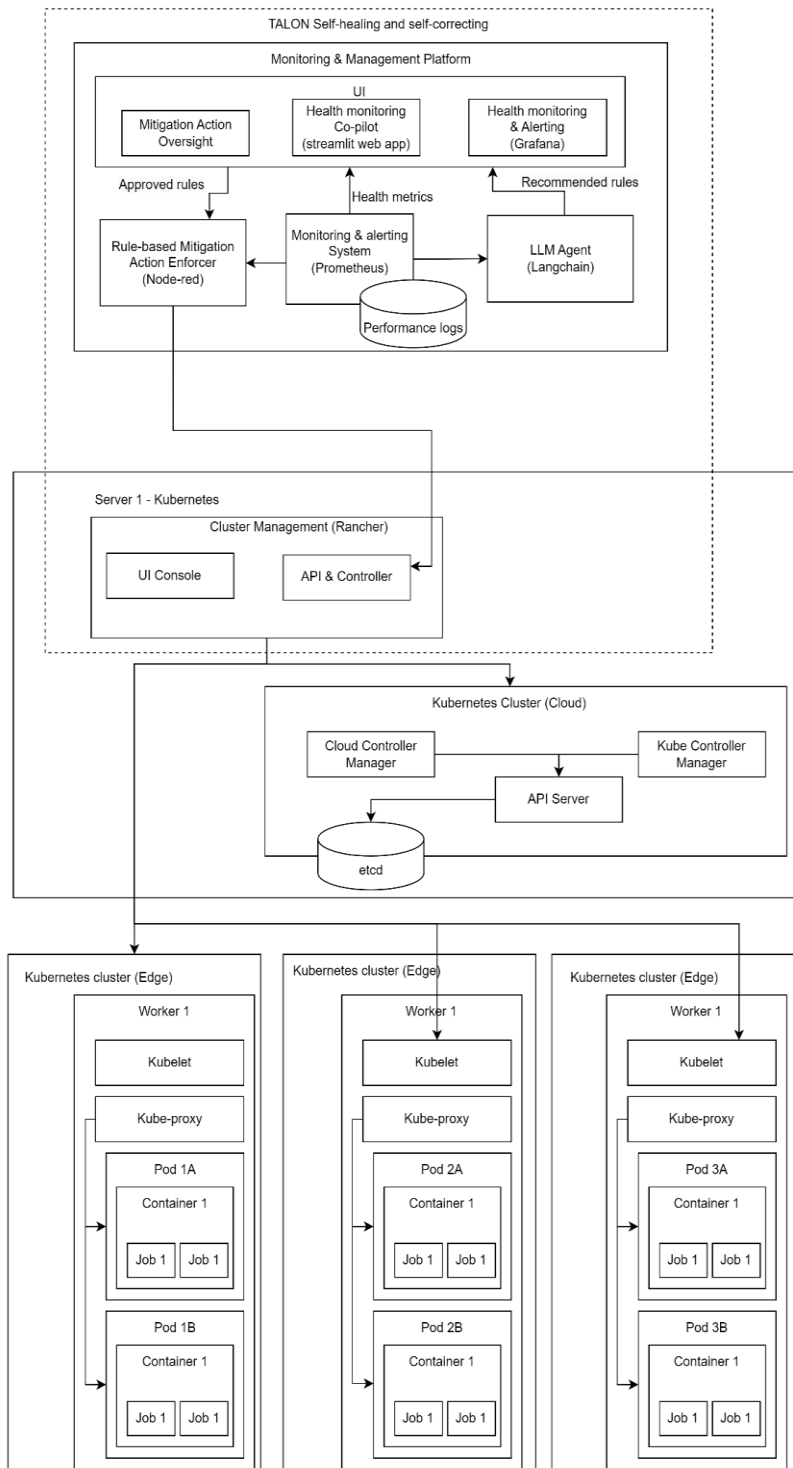


Figure 1: Self healing and self correcting Kubernetes cluster

3.2 Edge-to-Cloud Network Monitoring and Data Persistence

Edge-to-cloud network (E2C) monitoring provides integration of both edge computing and cloud Technology to increase the visibility for industrial networks jointly. This framework allows persistent observation and control of the entire system, from edge device to cloud node and support service creation based on real-time data collection analysis decision. This scalable model via a combination of edge and cloud functionality can cater to any network size and complexity, thereby providing an added ability in these industrial environments where self-healing as well as join-and-forget capabilities are most critical.

3.2.1 Data Collection at the Edge

Edge devices that collect data, such as sensors, switches and controllers reside at the far opposite end of most IT operations. These devices should be deployed to different points in the factory network, ubiquitously collecting information on how well (if at all) systems are running. Filtration of noise and feature extraction — Edge devices are empowered to preprocess data, which can help filter out the noise or pick up salient information from them. This will reduce the amount of data that needs to be sent on cloud, reducing latency and bandwidth usage, allowing self-heal actions faster.

3.2.2 Real-Time Data Analysis and Processing

Local data is relayed up to the cloud for deeper analysis. Advanced AI & ML algorithms running in the cloud infrastructure analyze this aggregated data from different edge devices. This configuration is what allows the most complex or subtle patterns and anomalies to be detected, rather than those that would merely stick out at the regional point. The direct ability to process in real-time means any issues are promptly recognized and proper countermeasures can be enacted.

3.2.3 Data Persistence and Historical Analysis

E2C monitoring needs to support data persistence, which requires storing historical information in the cloud databases. Very useful to train machine learning models, do trend analysis and learn how your network behaves over longer periods of time. The accumulation of historical data held in persistent storage on the cloud will yield a rich dataset over time to build better ML models that can further improve network self-correcting capabilities.

3.2.4 Synchronization Between Edge and Cloud

However, if edge devices and cloud systems are out of sync significantly, then self-healing mechanisms would largely fail. This requires reliable communication of data securely from the edge to a cloud platform and deployment back down updated algorithms and configurations. Effective synchronization is carried out to ensure that both the network components (edge & Cloud) are using data and software in its latest version thus decreasing chances of conflicts or their actions being irrelevant now.

3.2.5 Scalability and Flexibility

E2C monitoring method offers unparalleled scalability and elasticity that is essential for network changes or to scale the operations up/down. The Central Infrastructure can grow with network gradually - glow adding Edge devices does not affect core part of hub. It can dynamically allocate cloud resources to suit the processing needs at a moment and thus always maintain network strength & responsivity in various operational states.

3.2.6 Enhanced Reliability Through Redundancy

This is a basic arrangement built to increase the availability of cloud networks and integrations with fault tolerance. If one part of the system fails, parts based in different data centers (or edge devices) can still

provide network functions by duplicating processes and data across multiple cloud servers. For industrial networks, this redundancy is critical to prevent downtime which can be devastating from both an operational and a financial aspect.

3.2.7 Security and Privacy Considerations

However, in an industrial configuration where valuable data is frequently passed over the network, security and privacy are a couple of things that must not be taken lightly when implementing any type of Network Monitoring system. Improved Security: E2C architectures allow for greater security measures such as encryption, access controls and regular audits. Edge provides anonymization functionality before sending the data to Cloud i.e. If you want store any of content related confidential information, that can be stored after privately mask your raw information by anonymization process on Edge side itself. Furthermore, edge computing is decentralized in nature — a breach at one point has limited reach.

3.2.8 Continuous Improvement and Adaptation

An E2C system collects and analyzes a large amount of data to help in continuous improvement. Continuous monitoring and analysis realize are insighted, which leads to improvement in the self-heal/sand-correct algorithm continuously over time. This approach is also powerful because it enables fast adaptation and updating based upon real-time data, which will help keep the network in front of evolving threats or conditions.

3.3 Traffic Management

Future networks are getting more intricate and the necessity for lesser traffic is one of the fundamental challenges Distributed learning techniques, such as federated learning have proved very useful to minimize the amount of data transmitted through network and therefore allow more efficient resource orchestration. Therefore, efficient resource allocation is crucial for successful distributed learning training. However, assessing how resource allocation affects performance poses certain challenges. The distributed training process is both iterative and distributed; thus, making it complex to measure the influence of individual model updates on overall training. Additionally, as each device only shares its vector with the parameter server (PS), the latter lacks information on devices local datasets and cannot leverage sample distribution or data values to evaluate resource allocations effect on distributed learning convergence.

The applications of distributed learning techniques for traffic and energy reduction in wireless communication systems have received great research efforts in the past years. This is applicable to the Pilot 1 with the swarm of drones that are heavily relying to the wireless communication. For instance, Federated Learning (FL) was applied in a network setup and found that global convergence could be expedited by incorporating local training supported by small base stations that communicate intermittently with the PS to reach consensus. Moreover, decentralized learning was shown to not only speed up convergence, but also reduce energy consumption by minimizing communication distances [1]. It boosts communication efficiency by reusing frequencies across small cells thereby supporting simultaneous distributed learning. Specifically, research directions include:

- (i) balancing machine learning model updates and global model aggregation to minimize overall energy usage for local training and transmission or federated learning loss [2],
- (ii) using statistics to optimize device selection for federated learning rounds [3],

- (iii) exploring how user scheduling affects federated learning convergence [4],
- (iv) handling diverse user data without additional assumptions except for strongly convex and smooth loss functions optimizing resource allocation to improve convergence and training loss [5],
- (v) jointly optimizing device scheduling and resource allocation policies to maximize model accuracy within a fixed training time frame for latency constrained wireless federated learning [6],
- (vi) choosing devices for participation in federated learning without needing details about the wireless channel's state or the statistical traits of the devices [7], and
- (vii) considering the existing wireless channel information to optimize device selection and bandwidth allocation [8].

Future services, use cases, usage scenarios, and applications are expected to have important diversity in terms of requirements and required cloud systems features. As a consequence, next generation cloud and edge empowered systems need to be flexible, adaptable, and programmable. Motivated by this, TALON aims to develop an AI-empowered cognitive cloud architectural paradigm that co-optimizes edge and cloud resources and selects the optimal AI model placement to reduce traffic and maximize the performance of the overall system. Towards this direction, TALON's orchestrator optimizes the edge vs cloud AI to maximize both global and individual users and systems capabilities without violating the design parameters of each application. This creates a new cognitive cloud system architecture that will make the most by jointly optimizing the edge and cloud resources, enabling centralized, distributed, peer-to-peer (P2P), as well as hybrid intelligence. In addition, by utilizing a holistic optimization approach and leveraging the advances in distributed blockchain, TALON supports end-to-end (e2e) personalized and perpetual security and privacy.

These changes are pivotal considering that the progress in distributed learning and resources allocation bears relevance to what the deliverable tries to achieve with a 'Self-Healing, Self-Recovery & Self-Correcting Mechanism Toolkit. With network programmability, networks will be able to self-heal from faults and issues automatically optimizing its operations 24X7 irrespective of change in conditions with minimal or no human intervention. Federated learning, for example, establishes a cognitive overlay to identify and diagnose faults as they occur with zero delay — enabling edge data analysis to enforce resource adjustments on the fly. That nature of automatically detecting and resolving problems even before they get triggered is what defines self-healing networks.

Furthermore, the inherently decentralized character of distributed learning helps achieve a high degree of self-recovery as regional nodes (e.g., small base stations) can operate on their own terms and respond to failures instantly. Building upon these advantages is the continuous, iterative process of federated learning enabling self-correcting mechanisms for past experience to be continuously learned from and its operational modes adjusted accordingly to prevent future issues. This continuous optimization means the network stays strong and respond well to adversarial conditions too — hi-performing under duress, in line with what D3.4 aims.

4 AI Analytics for Network Reliability

4.1 Data Processing Techniques

Given the data-driven network ecosystems, optimal orchestration of distributed processing have been arising important issues for minimizing bandwidth efficiency and optimizing energy savings when managing heavily workloads. Reasonable data processing is indispensable for managing the information, as it accelerates in getting insights and eases decision making techniques. Stated otherwise, distributed data processing is mandatory in case of both distributed learning systems to get the best out of resources and manage data [41]. System level should be established when working with fault-tolerant systems so as to frame proper data processing techniques in order to have the consistency of data where hardware failures or network outages are observed.

- The categories of distributed data processing: There are several aspects that need to be understood and taken care for handling and analyzing the data-set over multiple machines or vice-versa.
- Data Segmentation: It splits the dataset to be processed on different node/machine. Strategic partitioning leads to more efficient use of resources and data distribution.
- Parallel data processing such as Apache Hadoop and Apache Spark that provides the infrastructure and tools for scalable distributed data processing. They also enable parallel data processing applications on machine clusters making it feasible to run our big-data computations.
- Synchronization: It refers to write/write conflicts, it is very easy to make the information in a distributed system inconsistent or out of date so if we have multiple copies thousands/millions/billion times on your server then only less than 1% how isomorphic algorithms are going to handle that and there comes consistency when two different duplications get updated at the same time.
- Resilience: Failures are inevitable when you have many components and the more components that move, the higher your chances of parts failing. System reliability is maintained by Data does not loss since using replication and fault detection algorithms to operate when nodes fail.
- Individual Segment processing has been done and so now the next immediate step is to integrate these processed results in parameter server for final outcome. Map reduction and merging algorithms gather results from nodes to provide a coherent view of the processed data.
- Scalability: Scale out the system as data continues to grow requires a balance between storage, computation and network resources across distributed nodes. We can use resource management strategies, like intelligent load balancing to optimize system performance.

4.2 Identifying Network Failures

4.2.1 Data Collection and Preprocessing

Essential knowledge for successful modeling is data and the superior version of it. For industrial networks, this data could also mean collecting any relevant network information from sources like PCAP files (network traffic captures), flow records and/or SNMP data as well as performance metrics such latency, jitter throughput packet loss rates etc. Similarly, logs from system and application activities in network devices

provide a richer context of the actual behavior on your network. The preprocessing of this data includes cleaning to remove duplicates and address missing values, normalization so that the features are scaled within a common range, and feature extraction in terms of important attributes such as packet inter-arrival times and error-rate. Supervised learning models require you to label the data with annotations, that come from telling what is a normal behavior and an aberrant one. Challenges: The key challenges at this stage are the reliability and accuracy of data, variation of network conditions to prepare models which can work flawlessly for all types of networks with different traffic patterns while being effective in handling large volumes typicality found in industrial networks.

4.2.2 Model Selection and Training

Selecting the appropriate AI model may depend on attributes of network data, or type characteristics as detected problems. For example, LSTM networks are well suited for working with time series data and identifying sequences so they would be an ideal choice to assist in continuous monitoring of network performance. In network traffic analysis, CNNs are strong in learning spatial attributes that help to infer high-level structural views from raw bytes. DNNs offer an adaptable and strong method for recognizing multifaceted patterns of various data types. During training, the data is separated into train/validation/test datasets to see how your model performs and avoid overfitting. It is necessary to tune the hyperparameters of a model, such as learning rate and batch size etc. The above problem leads to an iterative training process for scaling the model, requiring significant computational resources as well and ensuring generalization of it on new data is a continuous task. Model decision interpretability matters too, especially for industrial environments where transparency and accountability are necessary.

4.2.3 Model Deployment

A fully-trained AI model can now be deployed into an environment to monitor network data continuously for real-time insights. Models can run on edge devices to provide real-time inferences, or they may use cloud infrastructure for scale and resiliency — hybrid options of the two exist as well. In this way, edge deployment is always meant for instant processing and action hence their provisioning helps to reduce latency plus the overall utilization of bandwidth. Increased resources and centralized management in the cloud makes it practical to analyze larger datasets using more compute-intensive methods. As containers (via Docker) and container orchestration engines such as Kubernetes have become more standardized in newer models, deploying these systems becomes significantly easier. For instance, models can be served at scale on multiple instances with inference servers like TensorFlow Serving and NVIDIA Triton [21]. In deployment, the key challenges are tuning models for efficient execution on resource-constrained edge devices and ensuring low-latency inference by utilizing complicated industrial networks; also, it requires handling scalability to support systems in spite of increasing complexity or size.

4.2.4 Evaluation and Monitoring

Evaluating and monitoring the performance of deployed AI models is crucial to ensure they continue to operate effectively in real-world conditions. Key metrics for evaluation include accuracy, which measures the correctness of predictions; precision and recall, which balance false positives and false negatives; and the F1 score, which provides a single performance measure by combining precision and recall. Inference time, the time taken for the model to make a prediction, is also a critical metric, especially for real-time applications. Tools like Prometheus and Grafana are used for real-time monitoring of model performance metrics, while the ELK Stack helps in logging and analyzing model outputs and errors. Custom dashboards can be created to visualize key performance indicators and track the impact of self-healing actions on network health. The main challenges in this phase include managing the large volume of monitoring data,

configuring the monitoring systems for optimal performance, and ensuring that the monitoring tools do not introduce significant overhead to the network.

Evaluating and monitoring the performance of deployed AI models is essential for ensuring their continued effectiveness in real-world conditions. Here are the key metrics and their formulas:

Accuracy:

Accuracy measures the proportion of correct predictions out of the total predictions made by the model. It is calculated as:

$$\text{Accuracy} = \frac{\text{Number of Correct Predictions}}{\text{Total Number of Predictions}} = \frac{TP + TN}{TP + TN + FP + FN}$$

where:

- TP = True Positives
- TN = True Negatives
- FP = False Positives
- FN = False Negatives

Precision:

Precision measures the proportion of true positive predictions out of all positive predictions made by the model. It is calculated as:

$$\text{Precision} = \frac{TP}{TP + FP}$$

Recall (Sensitivity or True Positive Rate):

Recall measures the proportion of true positive predictions out of all actual positives. It is calculated as:

$$\text{Recall} = \frac{TP}{TP + FN}$$

F1 Score:

The F1 score is the harmonic mean of precision and recall, providing a balance between the two, especially when the data is imbalanced. It is calculated as:

$$F1 \text{ Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} = \frac{2TP}{2TP + FP + FN}$$

These metrics are critical for assessing how well a model performs, particularly in terms of making correct predictions, minimizing false positives, and ensuring that true positive cases are captured effectively

4.2.5 Common network failures and problems

4.2.5.1 Electromagnetic Interference (EMI)

Industrial environments Electromagnetic Interference(EMI) The noise comes from the use of heavy electrical equipment such as motors, transformers and generators. These devices, when in use emit electromagnetic radiation which can scramble data transmission across networks.

Industrial networks sometimes use shielded cables and connectors to lessen the spread of EMI After using special layers to partially or fully block out effects of external electromagnetic fields. In addition, grounding techniques and isolation measures are put in place to direct unwanted interference away from these susceptible network components.

Where EMI cannot be completely eliminated, the implementation of technologies such as Frequency Hopping Spread Spectrum (FHSS) or Direct Sequence Spread Spectrum modulation can provide a means to transmit data across multiple frequencies decreasing its susceptibility to interference.

4.2.5.2 Severe Environmental Conditions:

One such environment is industrial facility — Networking equipment in such a places go through hands down worst conditions. In other circumstances, like steel mills or chemical plants, those temperatures may be well beyond the range of standard networking gear. This is overcome by using specialized industry-grade equipment. These devices are made to work under various temperature and humidity levels.

In addition, protective enclosures and environmental control systems are used to build stable environments for networking hardware. Air conditioning or heating elements within sealed cabinets ensure that operating conditions remain optimal and that electronic equipment remains protected from the detrimental impacts of extreme environmental variations.

4.2.5.3 Physical Obstructions

In real industry, the placement is mostly determined by huge machines and structures. These barriers impede wireless signals and result in dead areas within the network. To counteract this, industrial network planners make use of signal amplification methods as well antenna placement and additional access point installation.

Additionally, reflectors and/or wireless repeaters could be used to extend the range of (exterior) networks in case that physical obstacles hinder a full radio coverage.

4.2.5.4 Vibration and Shock

The integrity of networking equipment can be at risk because machines vibrate and create shocks. Ruggedized hardware is thus deployed to make up for this. These devices are constructed with durable parts and materials to handle the wear-and-tear associated of industrial facilities.

In addition, to reduce the impact of vibrations on networking gear, both need shock absorbing mounts and protective enclosures. This ensures that the vital equipment and components continue to function regardless of how busy a facility may be.

4.2.5.5 Legacy Systems Integration

Migrating from traditional networking infrastructure may be difficult, since these legacy systems were not built with modern networking architecture in mind. In the industrial world, a lot of equipment is still using older legacy communication protocols. Protocol converters and gateways are used to connect the two. These are the devices that help to communicate between old equipment and new era network systems.

In Addition, middleware solutions as well as custom software interfaces are generally built for this purpose in order to allow easy integration. This protects much older equipment to run along side newer components on the network and keep working for already established industrial systems.

4.2.5.6 Latency/Real Time Constraints

Low latency is especially important in certain industries, for example manufacturing and process control. Because even the smallest delays in data transfer can become matters of life and death. Industrial networks also use technologies such as Quality of Service (QoS) and traffic prioritization to meet the requirements for responsiveness. Then, edge computing solutions process data at the point of production/occurrence – hence decreasing reliance on round trip communication to central servers. This provides the ability to perform sensitive operations directly with minimal reliance on remote data center latency.

4.2.5.7 Training and Skills Shortages

Why It Matters for Industrial Network Administrators and Technicians—including the video of our live phone call with them! Approaches to managing the newer technologies and trends with employees Real-world case studies demonstrating the value of trained personnel in ensuring dependable industrial networks.

4.2.6 Scenarios of Network Failures

4.2.6.1 High Latency

Otherwise, if the latency of a network remains high then it is likely that packets will take longer to be transmitted and this has an impact on time-sensitive applications which performance can suffer greatly. Factors such as network congestion, suboptimal routing or hardware failures can cause latency issues. In the industrial realm, spiking latency can break real-time monitoring and control systems which ruins uptime or creates serious safety issues. High latency is detected based on the measurement of round trip times (RTTs), compared to pre-defined thresholds. AI models are very good at recognizing patterns and predicting that an incident may spike again given enough periodic data. Some of the solutions to high latency include optimisation of routing paths, adding more bandwidth and QoS policies that operate at a network level so preventing non-critical traffic. By tracking latency metrics on a continuous basis, monitoring tools are able to generate immediate alerts the moment that individual requests violate an acceptable level of sustained latency.

4.2.6.2 Packet Loss

Packet loss, on the other hand, is a scenario in which packets of data that should have been delivered to a specific location en route are lost or reach their destination with errors. Network congestion, faulty hardware or software bugs can often be the source of this endeavor and on a wireless network neighboring units may interfere with each other. In industrial settings, packet loss could cause the forfeiture of very important data that disrupts monitoring systems and control processes reliability/accuracy. The way packet loss would be detected is by looking at the transmission log and then measure how many packets were lost over a period

of time. AI models are capable of recognizing basic crafts and reasons behind packet loss, leading to the ability to know when a problem is about to happen. Use of alternative solutions like Network infrastructure Upgrade

5 Proactive Mechanisms for Failure Mitigation

5.1 Design of AI-Supported Mechanisms

5.1.1 Data Acquisition and Integration

Robust data acquisition and tertiary ties to proactive network failure mitigation introduces the foundation of any AI-assisted mechanism. In the industrial space, data comes from a variety of places: network traffic monitors, IoT sensors (Industrial Internet of Things), SCADA systems and device logs. However, in order to make the best use of this data it must first be sourced and unified. Setting up data pipelines equipped with a capacity to handle massive amounts of real-time streaming data is necessary — this in turn ensures that all the information relevant for analysis is captured. Code: Normalization, Feature extraction – Convert and prepare raw data for AI. The aim is to provide a dataset which should be as well rounded and representative of the live network as possible in order for any prognosis or prediction born from it to be accurate and provided at the exact moment that such predictions are needed.

5.1.2 Machine Learning Model Selection

The right machine learning models are what help in proactive planning to prevent the network failures. There are different types of models that work best for specific types of data and problems. For instance, time-series analysis models such as LSTM (Long Short-Term Memory) networks are suitable for predicting network performance metrics across a period of time[i], and CNNs can be utilized to detect spatial patterns in the network traffic[two]. Moreover, we can use the models in Anomaly Detection like Isolation Forests or Autoencoders to capture the abnormal patterns if any that could lead to failures. Which model we should use is often contingent on the nature of data, types of failures being attacked and how much computational resources are available.

5.1.3 Anomaly Detection

Network failures detection through anomaly is a step towards pro-active maintenance. Anomaly detection algorithms continuously monitor network traffic against performance metrics to identify any deviations from usual behavior that might suggest a problem. These anomalies might appear as random spikes in latency, throughput swings or sporadic and short outbursts of packet loss. Robust anomaly detection can be implemented with proper model training based on historical data to determine what is normal, and thresholds or statistical methods are employed for alerting deviations. When an anomaly is identified, the service can send alarms or execute self-healing through preprogrammed actions to avoid it causing bigger problems.

5.1.4 Predictive Maintenance

Using artificial intelligence for predictive maintenance AI models applied to time-series data can predict when a failure may occur and allow operators to intervene in advance. This can be hardware failure, like a fan is going bad... uh-oh.... Or software bugs... or network traffic congestion ... all manner of things based on the history of this device (in our “digital twin” reality) and what other related devices are experiencing right now. By training models using various approaches like regression analysis or neural networks it is possible to predict remaining useful life of network components, or the probability that a failure will occur within the next X hours. By predicting these eventualities, maintenance can be planned in advance, reducing downtime and avoiding expensive surprises. Image Credit: This approach not only increases the reliability but also improves maintenance resources usage and operational costs.

5.1.5 Automated Response Systems

Automated response: Alerting and automated responders are built in to trigger corrective actions as soon as any possible issues are detected. This action can be determined or rule-based, and some AI models are capable to understand the anomaly type as well as anomalousness severity in order to choose best action & plan for mitigation. Eg: When it detects high latency, the system can automatically reroute traffic to less congested paths or adjust QoS settings (Quality of Service), allocate more bandwidth etc. The system would alert in the event it could predict a hardware failure on its way, and instead suggest type of failed hardware should be placed preemptively. Fast activation of corrective actions ensure network is quickly back up and running, reducing the operational impact from networking problems.

5.1.6 Continuous Learning and Adaptation

To stay effective in dynamically shifting landscapes, AI-enabled network failure mitigation tools should support "always on" mechanisms that are always improving. This necessitates training the AI models on updated data that is a true representation of how the network looks and behaves recently. Ongoing learning allows you to accommodate changes in network behavior, including new traffic patterns as well as evolving threat landscapes or hardware upgrades. To keep the models current, methods such as on-line learning or cyclical retraining can be used. They can also create feedback loops where the results of automated responses are analyzed to tweak algorithms and enhance future outcomes.

5.1.7 Integration with Existing Network Management Systems

If AI-backed approaches of proactive network failure prevention will be efficient, they should smoothly interoperate with the current network management systems. This integration requires communication between network monitoring, configuration management and alerting. APIs and microservices will enable command & control between the AI system and these tools, so that data properly flows in both directions describing actions taken. It is integration that will also enable our AI system to use already existing data acquisition, processing and corrective action facilities. The AI mechanism works in conjunction with existing systems and complements network management capabilities across the board, making migration to a new setup unnecessary.

5.1.8 Security and Privacy Considerations

In industrial networks, security and privacy are vital — data integrity is in question. Any AI-powered systems should be embedded with strong security features to help guard against cyber threats and unauthorized access. This includes encrypting data at rest and in transit, enforcing access controls/permissions, and frequently carrying out penetration tests on your system. And finally, ensuring privacy of the data — especially if it is sensitive information — must be considered. Methods like differencing privacy, or federated learning might be employed to guarantee that an Individual unit of Data is protected and still for the AI with enough statistics. Security and Privacy — Not only does this protect the network, but it also builds trust in the AI system.

5.1.9 Performance Metrics and Evaluation

Rating the efficiencies of every AI assisted techniques utilized for countering network failures incriminates tracking those metrics prescribed continuously. Performance targets such as detection accuracy, false positive rate, response time or the ability for automatic responses to recover from issues are all potential key performance indicators (KPIs). These metrics can imagine how the system is performing it and can pinpoint to where action/damages have taken places. Pros: Regular review, so a transparent generation of how the AI is fulfilling its goals and bringing real value These metrics can be continuously monitored using

tools like Prometheus, Grafana etc. which keep track of these counters and provide dashboards/alerts about the concurrently running consumers in real-time.

5.1.10 Scalability and Flexibility

The flexibility and scalability required of AI-supported mechanisms to service the increased scale requirements in industrial networks are thus more crucial. The system must scale horizontally to deal with greater volumes of data and more complex network topologies. Customers who choose to inch toward a migrate-to-the-cloud program will be supported by these used architectures (e.g., microservices), making migration seamless and less daunting over time as the system grows, like what Kubernetes offers. In parallel with this, the ability to be flexible and change AI models or decision-making processes for new types of network issues arising from operational requirements is very important. They suggest that by designing for scale and flexibility, AI-enabled mechanisms could deliver benefits over the long term while accommodating future industrial network requirements.

5.2 Development of Proactive Solutions

5.2.1 Solution Overview and Goals

One of the key bits about developing proactive solutions for industrial network failure mitigation is being able to know and acknowledge when something will escalate into a big problem. That is to build a solution that can continuously observe the network health, identify real-time anomalies and anticipate potential failures in-order for it go ahead an automatically enforce corrective actions. Self healing and self correcting mechanism will concentrate on the network level problems only. By applying AI and machine learning to tons of network data, the solution was able to spot patterns that may indicate a problem on its way. The system is designed to improve network reliability, decrease downtime and increase operational efficiency in industrial environments through the introduction of predictive maintenance, anomaly detection and automated response mechanisms.

5.2.2 Overall Architecture

Proactive — The architecture of the proactive solution is scalable (Figure 2), resilient, and real-time. The solution has the following major components: — Edge devices, central data processing hub (AWS IoT Core),AI models, DSS dashboard and Monitoring tools. Edge devices are placed around the network to pull data from sources (such as sensors, logs and traffic monitors) then do some light processing on that nuanced data: Is it coming in? This information is then relayed to the central data processing hub which serves as the brain of the system. This hub incorporates high-performance computing to support AI/ML models and data analytics, along with storage for older data sets as well as model training capabilities.

5.2.3 Workflow Description

The workflow begins with data collection at the edge devices, which continuously monitor network performance and capture relevant metrics such as latency, packet loss, and throughput (Figure 2). To protect data quality from noise, this data is pre-processed locally before it reaches the central hub. This data can then be processed again in a central hub and used to feed the AI models for analysis. You can use AI models for anomaly detection that differentiate between normal behavior and suspicious activities, or predictive models which predict failures with the help of historical trends taking real-time data into account. In the case of an anomaly, or a potential failure was detected; DSS evaluates it and suggest (or automatically performs) remediation by rerouting traffic or updating configurations.

5.2.4 Integration and Automation

The integration with the network management systems already in place is of utmost importance for an operation without setbacks. The way it does this is by speaking directly to network monitoring tools, configuration management systems, and alerting frameworks using APIs and microservices. This guarantees the flow of data and expands actions well among all components. Relieving humans from a bulk of these steps is where automation propels faster answer times in pair with less or no human work. Automated response systems are set up with pre-defined rules or AI model recommendations to take corrective measures, guaranteeing speedy resolution of problems. These have served to make network management more efficient and the networks themselves that much stronger.

5.2.5 Evaluation and Continuous Improvement

Nevertheless, to test whether or not this proactive solution is effective continuous evaluation and improvement are required. In addition, we monitor on-line performance metrics like detection accuracy or response time of automated actions fails. These monitoring tools like Prometheus and Grafana also give dashboards, alerts to monitor these metrics and tells you what can be improved. Regular feedback loops are established where the outcomes of automated responses are analyzed to refine AI models and decision-making processes. Additionally, the system is designed to adapt to changing network conditions, with capabilities for online learning and periodic retraining of models. This continuous improvement cycle ensures that the solution remains effective and evolves to meet the dynamic needs of industrial networks.

The next figure shows the status of the system at this reporting month, indicating which components have been released (in green), which components are available as early prototypes (in yellow), and which components have not been developed yet (in red).

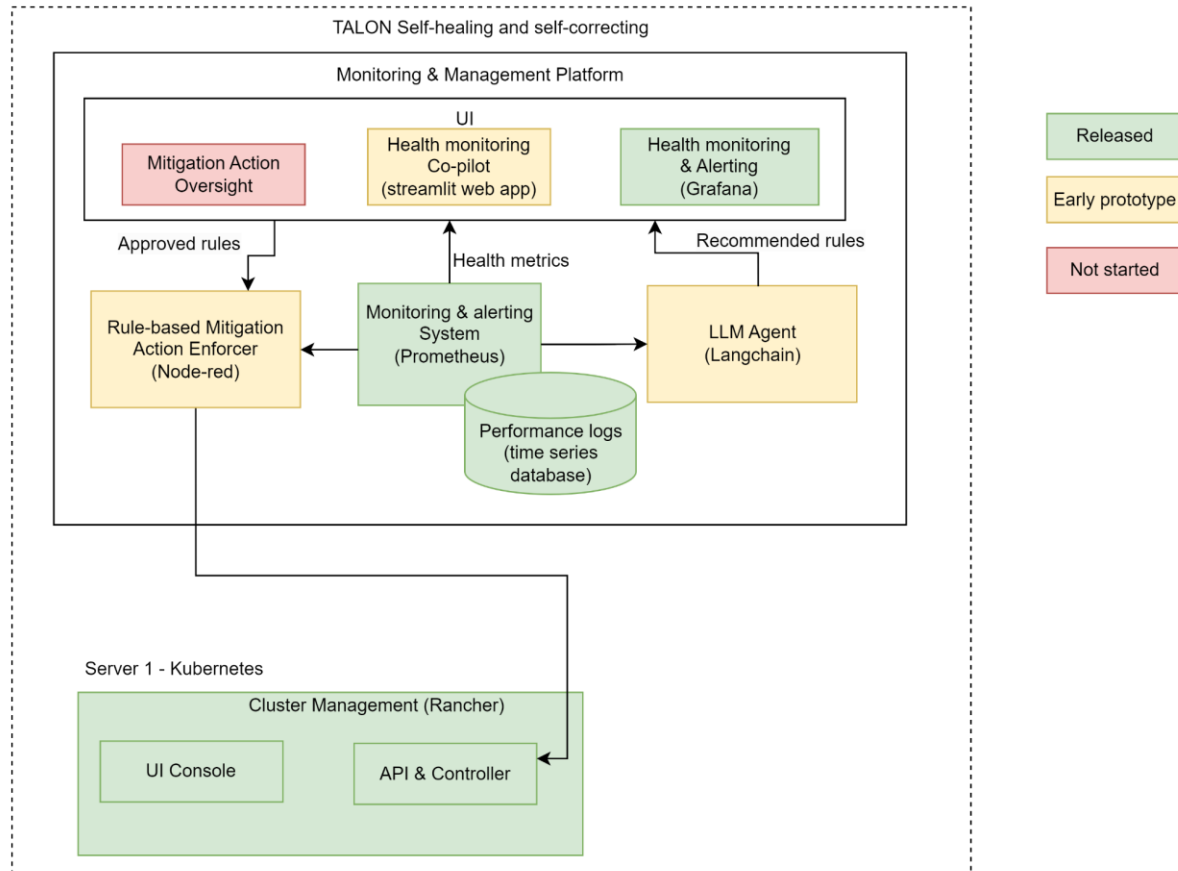


Figure 2: solution architecture

In this reporting month, the main system components (cluster management, monitoring & alerting system, health monitoring and alerting) have been integrated. The source code includes installation scripts and distributables to install these components in any Kubernetes cluster. The system has also been instantiated in two different scenarios, a development instance installed on UPV premises and a pre-production instance installed on FACTOR premises.

The development team has also implemented prototypes for the remaining components. Details of the early prototypes are provided in the section underneath system description.

5.3 Proactive Network Resources Allocation

The proactive assignment of bandwidth is crucial in contemporary communication networks to cater the increasing needs for high-bandwidth and low-latency services. Proactive network resource allocation reduces the high demand for computing and networking. Anticipatory network resource allocation is a strategy that, as the name suggests, will anticipate future conditions of both the network and user needs to ensure resources are allocated accordingly in good time. This is in contrast to reactive approaches where

the resources are allocated during run-time, based on immediate demand, creating situations of poor performance underload due to increased latency and resource contention.

The proactive allocation of network resources is grounded in several key principles:

1. **Predictive Analytics:** Utilizing historical data and machine learning algorithms to predict future network traffic patterns and user behaviour.
2. **Dynamic Adaptation:** Continuously adjusting resource allocation based on real-time monitoring and predictive insights to optimize network performance.
3. **Quality of Service Management:** Ensuring that the network meets predefined QoS requirements for different applications and services.
4. **Resource Optimization:** Efficiently managing network resources to avoid over-provisioning while ensuring sufficient capacity for peak demand periods.

In this section, Kube-OVN is used on a scenario setup, to support the proactive network management and resource allocation.

The scenario is described as follows. The figure represents a network architecture with three edge devices labelled "edge device 1," "edge device 2," and "edge device 3," all depicted as oval shapes on the left side. Each of these devices sends data to a central gateway API, also shown as an oval shape on the right side, via arrows that illustrate the traffic flow. The label "huge traffic" indicates that the gateway API is receiving a substantial amount of data from the edge devices, emphasizing the significant load handled by the gateway.

This traffic produced not only increases the network traffic of the pod, but also the CPU usage, resulting in the need of optimized network and resource management, as well as self-healing mechanism.

Using monitoring data on the Gateway API pod (CPU usage, Network Traffic), a set of policies can be applied as part of network intelligence and self-healing.

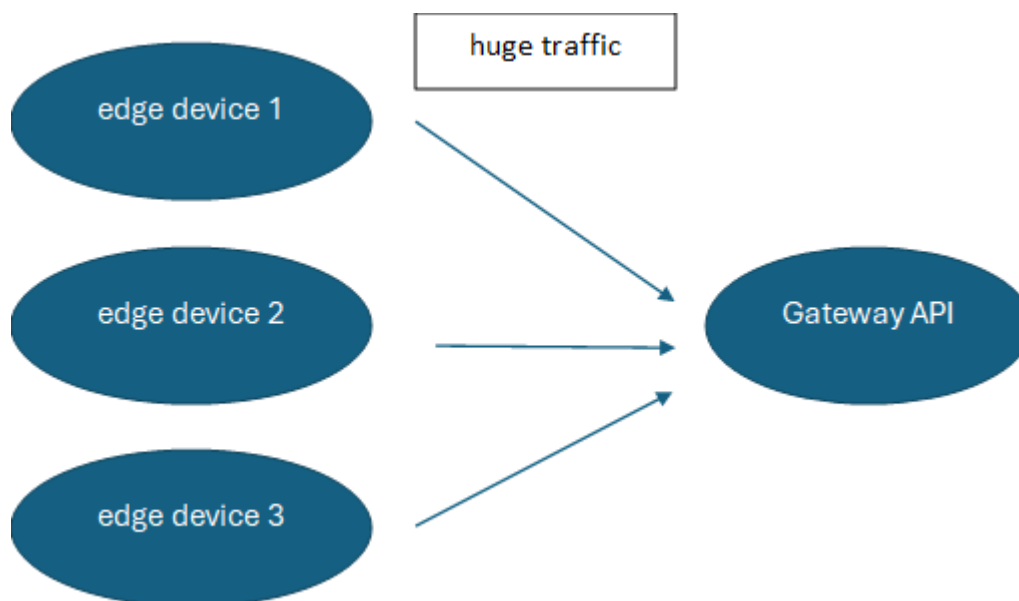


Figure 3: Proactive Network Resources Allocation

6 ML model update mechanisms

To minimize data-pollution and thus maximize data-efficiency, TALON exploits two novel and promising approaches, namely few-shot learning and federated learning, which are optimized to address the trade-off between latency and accuracy. To further boost the system's efficiency and autonomy, TALON utilizes an AI orchestrator capable of supporting different usage scenarios in a fully automatic fashion, by optimizing the relation between edge and cloud AI. The orchestrator selects AI datasets, algorithms, metrics and models based on the application. This enables high-level of reusability and robustness; thus, decisively contributes to unprecedented low-latency. To enhance the AI's network capabilities to solve a wider set of more complex problems, TALON set the foundation of allowing collaboration between commercial and non-commercial equipment. In other words, it exploits underutilized resources to address real-time performance requirements by bringing intelligence to the edge. In this sense, it enables three (3) strategically chosen deployment options, i.e. a) P2P intelligence for latency-critical applications with high reliability requirements; b) distributed intelligence for enhancing the computational capabilities near sensors, while maximizing the EE and dependability of the system; and c) federated intelligence to boost system's accuracy and enable dynamic reusability. The deployments are executed under a security and privacy umbrella, which exploits novel decentralized and hierarchical lightweight blockchain, anonymization approaches, as well as unsupervised and self-supervised learning anomaly detection and self-healing schemes.

6.1 P2P learning

The explosive integration of intelligent and mobile devices in human life has led to a massive increase in the volume of data produced and exchanged. Simultaneously, contemporary applications require innovative technology, such as AI to optimize and personalize their services. Traditionally, training machine learning models with data sourced from more than one device, would require data to be aggregated centrally, and then applying machine learning algorithms to train the model, which would then be distributed to the end-devices [9].

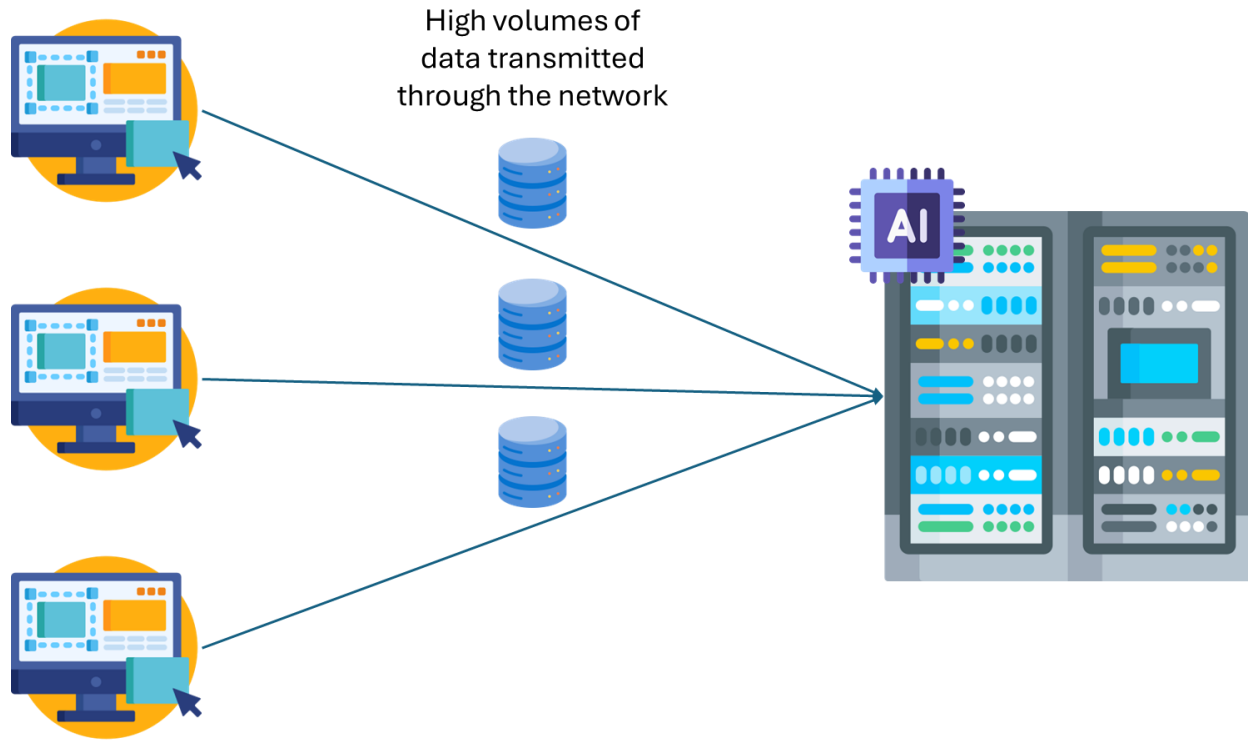


Figure 4: P2P learning setup

TALON aims to develop the self-healing and self-correcting mechanisms for improving the reliability and robustness of the network. In this direction, the mechanisms will collect network-wide data that will be stored in the Data Pool located in the TALON Information Layer and processed by advanced AI analytics in order to identify and quantify the probability and impact of various network failures. Moreover, in this task, AI-supported proactive mechanisms will be designed and developed to mitigate the occurrences of such failures. To reduce the traffic overhead, these proactive mechanisms will be deployed in the edge nodes controlled by the TALON Access Layer and, only the model updates will be transported to the cloud instead of the raw data. Finally, in case a network failure occurs, the proposed self-healing/self-correcting mechanisms will minimize the network recovery time.

6.2 Federated learning

However, major issues arise by employing such methods for collaborative training, as the privacy of data can be easily breached, while also not being the most cost-effective solution. For this purpose, FL has emerged, which is a technique of collaboratively training models without exchanging local data; instead, devices train their models locally, based on their own data, and then forward the models to a server to fuse them into a single global model [10].

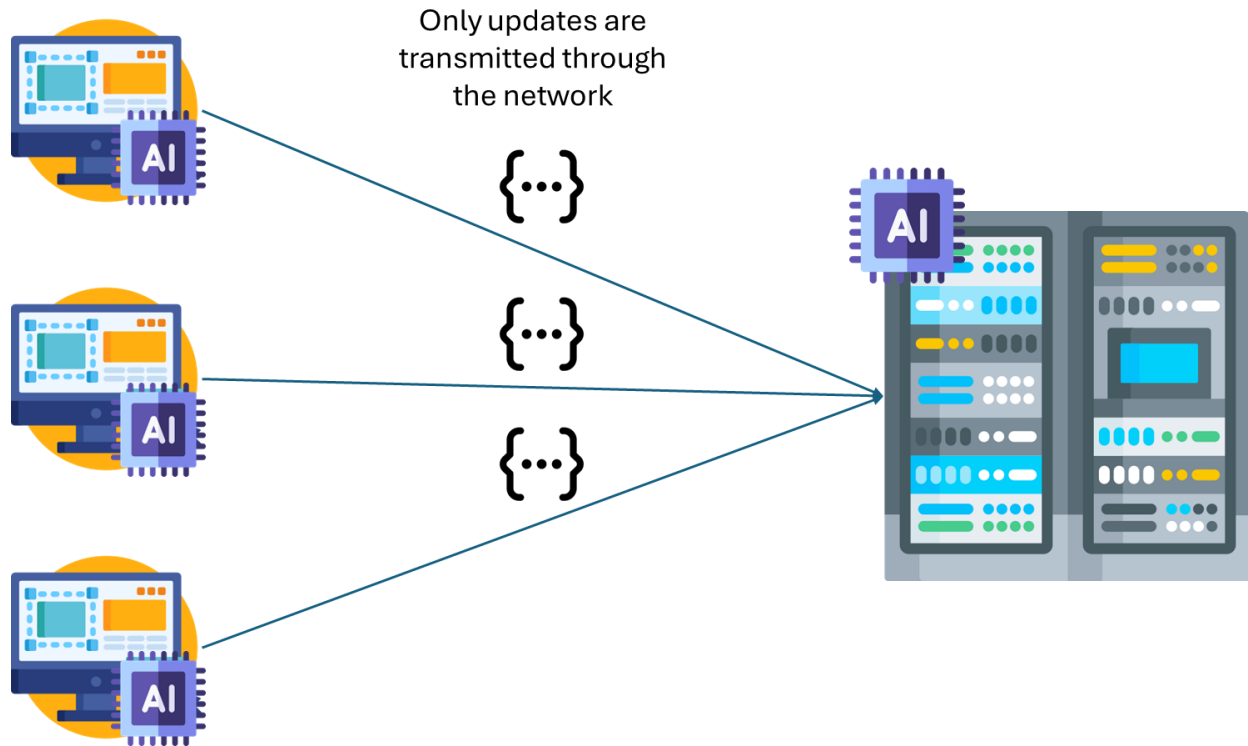


Figure 5: Federated learning setup

To surpass the aforementioned limitations, TALON develops a distributed learning framework, in which the model training is distributed across a number of edge devices that are selected as a solution to a well-defined optimization problem with constraints on the energy autonomy of each device and its computational capabilities. End-devices employ their local data to train advanced ML that are required by the parameter server, located at the edge plane. The device then communicates only the model updates to the parameter server for aggregation instead of the raw data. This is expected to significantly reduce network pollution. The updated parameter server executes proactive self-healing mechanisms and feeds its outcomes/predictions to the end-device.

6.3 Distributed learning

Issues, however, persist in traditional FL, due to the single-point-of-failure problem. To tackle such limitations, blockchain-based solutions have emerged to decentralize FL. BAFFLE suggests an aggregator-free, blockchain-based solution for FL, where users acquire the latest model through a smart contract and perform a training procedure; next, devices partition their local model weight vector into chunks and push them, if accepted, to the smart contract for aggregation [11]. In an attempt to introduce a fully P2P training process, BLADE-FL supports the distribution of locally trained models over the P2P network, where each receiving node is responsible for the aggregation process, based on the rules defined in the SC [12].

Evidently, a lot of great research has been conducted in order to eliminate the data privacy issues traditional machine learning methods encounter when trying to train models by aggregating local data in a central entity, with the introduction of the federated learning paradigm. Still, such solutions do not take into consideration the single-point-of-failure problem, which could lead to the disablement of the entire machine

learning system. To counter these limitations and introduce security enhancements to the already private federated learning, TALON aims to fuse the federated learning and blockchain technologies, thus introducing an innovative strategy of distributed and decentralized training. TALON's P2P approach is centered around the idea that a participant can locally train its model and use the distributed ledger for registering it, thus making the model update retrievable by every authorized participant. As such, TALON offers a highly distributed, secure, private and fully P2P approach to training models in federation.

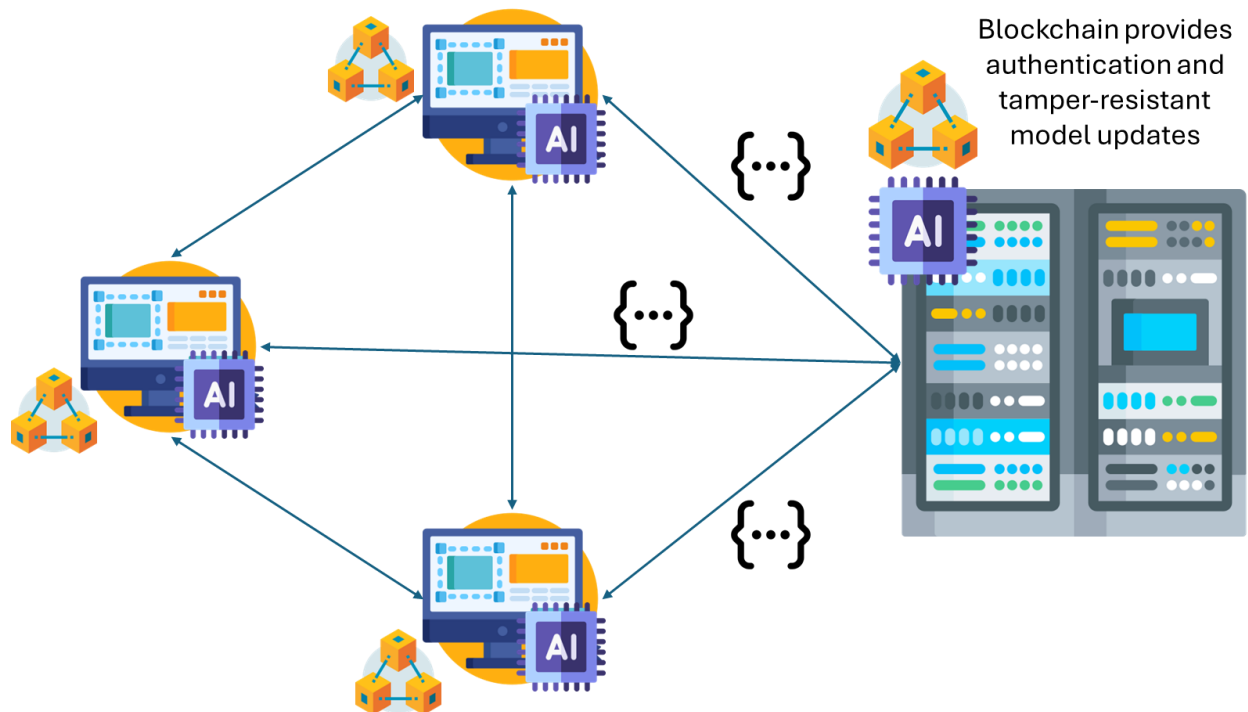


Figure 6: Distributed learning setup

An illustrative example is industrial automation and planning in manufacturing plants, which are equipped with an abundance of sensors, cameras, machines, and computer interfaces in a distributed network that need to cooperate in order to ensure high-efficiency of the manufacturing procedure [13], quality of the products [14], and safety of the employees [15]. In this use case, the AI architecture should support high-automation by means of self-optimization as well as fast and self-healing in case of attack, without sacrificing explainability, robustness, and/or accuracy [16]. In more detail, self-healing and self-optimized systems, which use ML and distributed learning approaches, can be used to ingest a non-stop stream of data to make the machines more efficient constantly, as opposed to periodically [17]. In this scenario, an edge AI computing system notices that a feed tank (accuracy) is low and informs the production machine to slow down in order to avoid running out of raw material (efficiency). At the same time, it signals the upstream processes to speed up and notifies the plant operators via the digital twin interface about what is happening (explainability) [18]. The core contribution of edge AI in this use case is to ensure the uninterrupted production by automating the various processes in the factory, increase the performance of the systems by utilizing novel ML and distributed learning techniques, enabling the machines to constantly self-heal or self-correct their processes, as well as provide explainable HIL solutions for the management and augmentation of complete factory workflow.

7 Self-Healing and Self-Correcting Mechanisms

7.1 TALON NG-SDN Programmability

NG-SDN decouples the control plane from the data plane, enabling centralized management of network policies via a programmable SDN controller. This separation allows administrators to change network configurations, apply policies, and deploy services across the entire network through high-level APIs, without needing to manually configure individual network devices. This gives you centralized control logic, so that you're never dealing with routing loops again and it allows for fast deployment of network changes, better operations practices to streamline your organization on multiple different modes but all under the same controller which is great.

The programmability of NG-SDN is possible through well-known APIs – OpenFlow, NETCONF or REST APIs which provide mechanisms to alter the behavior at different levels in network. Developers are encouraged to use these APIs to create applications that communicate directly with the network hardware, for tasks like automated vCPE provisioning and ONF SDN testing — all without middleware or higher level management complex issues of basic networking. Beyond operational simplicity, this interface layer unlocks innovation paths in highly automated network service delivery.

This NG-SDN takes network automation to another level. It provides programmability so you can orchestrate intricate operations spanning multiple network and service elements. Routing, Switching and Security like services can be dynamically scaled up (or down) based on the instant Data available or Policy definition. At the same time, automation frameworks working with SDN controllers can make dynamic modifications across an entire network much more quickly and reliably than a person can manually enforce changes — improving accuracy while offering significant cost savings over traditional methods of managing networks.

With this architecture, and given the OG-SDN supports multi-domain service transmission — also described in ONF's document above—dynamic inter-domains chaining values are reduced as well. Service chaining builds a chain of network services on simply application policy change without requiring any changes in the architecture of the network. This is particularly useful in scenarios where customers want to quickly deploy new services or tweak their network-based on user/application-specific requirements.

Furthermore, NG-SDN analytics also inherits capabilities of programmability. SDN controllers can use the data on current network activity from different elements as input in real-time for dynamically making decisions to optimal traffic flows, wield dual use (applications/dev) bandwidth consumption under critical conditions and proactively address impending disruption(s)/aberration. By leveraging advanced analytics tools layered on top, in common management and orchestration platforms that are integrated with SDN capabilities can identify patterns to predict impending issues — including those of a security nature — not ideally detected or corrected (if still possible) by mere humans, way ahead compared to today's level even if AIOps, since it will be the kind of input feeding into such systems.

De-centralization: NG-SDN architecture is designed for scalability; it can grow with the increase in network traffic and number of devices. The flexibility of programmability enables changing the way specific services are provided on demand, with no need for physically re-configuring resources (PP). Scalability is key for growing businesses or those going through a digital transformation.

This model boasts more agile and secure networks such as — where programmability allows for on-the-fly security policies to be implemented that can even go from creation of the policy, to notification and mitigation

in real-time. The network can be dynamically reprogrammed in milliseconds to isolate affected areas, or divert traffic and implement enhanced security protocols if threats are detected.

By making NG-SDN programmable, we converted the network management into a more nimble and effective intelligent action that can deal with complexities of contemporary digital environment. With the ever-increasing number of connected devices and a rapidly expanding cloud service platform, NG-SDN progress has laid down strong foundation for automating network operations to keep up with this growth in an agile and secure manner. These programmable, adaptable architectures will become the basis of networks in enterprise and carrier environments for years to come – a complete change from traditional network management approaches..

7.2 Minimizing Network Recovery Time

Reducing time for the network to recover is a key consideration in managing networks, particularly in industrial and critical infrastructure scenarios where even brief outages can lead to serious operational consequences or financial harm. Ensures network functionality is rapidly resumed, in the event of an incident irrespective of how small or large it might be and helps maintain availability. Some of the strategic and technological pointers in maintaining low network restoration times are:

7.2.1 *Advanced Monitoring and Alerting Systems*

The first line of defense for keeping recovery time to a minimum is an effective monitoring systems. These systems can detect issues that lead to outages before they reach a critical level, or even prevent potential collapses by keeping an eye on network health and performance metrics in real-time. Real-time data analysis tools based with AI that can pass those tickets predict failures or automatic alert fires for immediate reaction. Proactive monitoring like these can greatly reduce your Time to Detect (TTD) and Time To Engage Issues(Time TTE).

7.2.2 *Automated Diagnostic Tools*

When a problem is detected, automated diagnostic tools can be used to narrow down the cause. These tools leverage historical information and some level pattern recognition as well utilize ML algorithms to understand the symptoms put them into any known issues. Increasing the understanding time (TTU) of the problem: automating diagnostics enables most of this simplification, accelerating recovery.

7.2.3 *Pre-configured Recovery Protocols*

In the case of common network problems, if a pre-configured re-mediation approach exists for them, then time taken to recover (TTR) can decrease significantly. These are automated scripts or a series of actions taken when certain problems are identified. These might involve things like a complete power cycle for hardware, creating new routes to reroute traffic or making configuration changes. Automation eliminates human error and provides the fastest response times when disaster strikes.

7.2.4 *Redundancy and Failover Mechanisms*

Redundancy coming up from the network level (routers), through switches, pathways to datacenters -- this will ensure that you have a backup if one part of your system fails. When failures are detected it can switch over to redundant systems automatically, proxying traffic with no perceptible downtime in network operations.

7.2.5 *Decentralized and Edge Computing Architectures*

Exceptions disappear in decentralized networks: if a single node or even one part of the network is asking trouble, it does not crash all other nodes. Edge computing takes this to a greater level by processing data

at the edge of the network — nearer data sources and users. This diminishes the effect of central failure and provides quicker recovery for local systems, frequently without requiring broader network level restoration measures.

7.2.6 Network Segmentation and Microsegmentation

Segmenting the network into smaller, manageable segments can contain problems and minimize their impact on the entire network. Microsegmentation is particularly effective in highly dynamic environments like data centers, where it can isolate issues to individual workloads or applications, making recovery processes quicker and more targeted.

7.2.7 Training and Simulation

Regular training and simulation exercises prepare network teams to handle failures efficiently and confidently. Simulations of different failure scenarios can help identify weaknesses in recovery plans and provide practical experience in executing rapid recovery strategies. This preparation reduces the time to resolve issues when real incidents occur.

7.2.8 Post-Incident Analysis and Continuous Learning

After you have recovered from a network interruption, it is essential for learning the what happened and why through an extensive post-incident evaluation. The analysis contribute to reviewing recovery strategies and designing more successful preventive measures in the future. Optimizing is a continuous process, learning from each incident reduces recovery time in the future and increases network resilience.

A successful effort to reduce network recovery time will rely on a holistic approach that connects advanced monitoring, operational diagnostics executed automatically, rapid recovery protocols and training. By applying technology and best practices to improve these, it is possible for organizations incurring network incidents to recover quickly from them – ensuring high levels of uptime and service quality required by today industrial settings comparable business operations.

7.2.9 Importance of Rapid Recovery

One of the most challenging aspects in industrial networks is to reduce recovery time after a network failure so that operational continuous and downtime costs are minimised. Fast recovery makes it possible to get production processes, safety systems and other key operations running again as soon as possible at the same time reducing a financial crisis by preventing large loses of revenue through keeping productivity going. The longer a network is out, the more widespread that blowback is felt throughout your business. Thus, it is crucial to adopt techniques and technologies that help bringing back your network up as soon as possible. These activities include proactive monitoring and alerting, early detection of problems as soon that can be used to identify whether recovery actions have completed successfully or the system may need some manual intervention.

7.2.10 Proactive Monitoring and Real-Time Detection

The initial and primary line of defense to reduce network recovery time is proactive monitoring. Critical issues can be discovered before they become outages — upping performance and network health monitoring. AI and machine learning algorithms analyze real-time data from network devices, sensors, and logs to detect anomalies or signs of threats as they occur. This serves as early warning to network administrators so that actions can be taken before any damage. If it finds a problems with consistently high latency being experienced on the network, you can re-route traffic to avoid congestion. The earlier an issue is detected, the sooner action can begin to reduce that time and help fix it quickly.

7.2.11 Automated Response Mechanisms

Using these automated response mechanisms it is possible to drastically reduce the time required for recovering a network from failure. Systems of that nature are specifically built to kick off some type of predefined action when a certain condition is triggered. When a network segment fails, the automatic process can reroute traffic on an alternate path immediately or restart malfunctioning components (SNMP/Generic), same as we have to apply placeholders attributes until they are manually fixed for example. This assures that the right actions are taken almost immediately and limiting outage timing. In the same way that an automated system can look after time-consuming repetitive tasks to provide a consistent and fast response when dealing with network problems.

7.2.12 Efficient Resource Allocation and Redundancy

You expend least amount of time to recover from a fault before more resources is needed or redundant servers are necessary. This includes load balancing network resources and establishing redundant systems which can take over in case of failure. For example, there are multiple ways of transmitting the data so that if one fails another will be activated soon. In the same way as alternative computing devices (such as additional servers and switches) can be booted into service to replace downed units. During Disaster Recovery, resource allocation should occur in such a way that it prioritizes the essential section of your network and applications to get proper bandwidth as well as computational power so that they can continue lest their operations may hamper.

7.2.13 Continuous Improvement and Learning

Reducing network downtime is not just one-and-done, it should be your goal to get that time down as much as possible and keep on learning from incident response. After every occurrence of network failure and recovery, a detailed study must be done to check what failed, how it was handled and how the time for recovery could be decreased even more. This entails tweaking the AI models to even higher levels of accuracy, making detection algorithms better and increasing automated response rules based on telemetry coming from real-life scenarios. In addition, outages can be minimized even when preparing for the worst as exercises and drills help to condition not only just the network management team but also systems in preparation of likely failures making recoveries more familiar. So too, but by learning from previous events and adjusting tactics accordingly (thereby increasing the overall robustness and responsiveness).

7.3 Dynamic QoS Adaptation Mechanisms

In Kubernetes (k8s) environments, QoS for network resources is just as vital when allocating parts to guarantee equitable sharing and optimal part performance. With the comprehensive network plugin for Kubernetes, Kube-OVN, policies which dictate maximum bandwidth limits can also be enforced. This section addresses how QoS can be fine-tuned on-the-fly with Kube-OVN, covering the methods to create bandwidth policies and apply/modify it.

Network Policies Definition

To start assigning QoS while using K8s-OVN you must define network policies to specific the maximum bandwidth for each pod. Again, with Kubernetes Kube-OVN annotations to specify bandwidth limits for a given pod the policies are created. For example, an annotation to limit bandwidth may take the following form:

```

apiVersion: v1
kind: Pod
metadata:
  name: gateqay-api
  namespace: talon-policies
  annotations:
    ovn.kubernetes.io/ingress_rate: "3000"
    ovn.kubernetes.io/egress_rate: "100"
spec:
  containers:
  - name: qos
    image: docker.io/library/nginx:alpine

```

Figure 7: Network Policies Definition

The above figure illustrates a typical pod YAML file, where it is defining a Kubernetes pod. The definition of the QoS policy is done by the annotations `“ovn.kubernetes.io/ingress_rate: <BANDWIDTH_VALUE_Mbit/s>”` and `“ovn.kubernetes.io/engress_rate: <BANDWIDTH_VALUE_Mbit/s>”`. This annotation sets the maximum bandwidth value to be the one stated in `<BANDWIDTH_VALUE_Mbit/s >`. Thus, restricting or allowing more network traffic to be passed to a specific pod.

Application of Bandwidth Policies

Once the network policies are defined, they need to be applied to the relevant pods. This is achieved by specifying the QoS Policy in the pod annotations, ensuring that each pod adheres to the allocated network resources. An example pod definition with QoS annotations is provided below:

```

def annotate_pod(namespace, pod, annotation):
    v1.patch_namespaced_pod(pod, namespace, body={
        "metadata":{"annotations":annotation}
    })

def annotate_qos(pod,namespace, ingress='3', egress=None):
    if ingress is not None:
        annotate_pod(pod, namespace, {"ovn.kubernetes.io/ingress_rate": ingress} )
    if egress is not None:
        annotate_pod(pod, namespace, {"ovn.kubernetes.io/engress_rate": egress} )

# Apply QoS policy on pod
annotate_qos("demos", "gateway-api",ingress='14000')

```

Figure 8: Application of Bandwidth Policies

In this Scenario, the policy applied is Quality of Service on Pod Level.

The following figures represent the bandwidth usage of a pod before and after the application of a Quality of Service (QoS) policy.

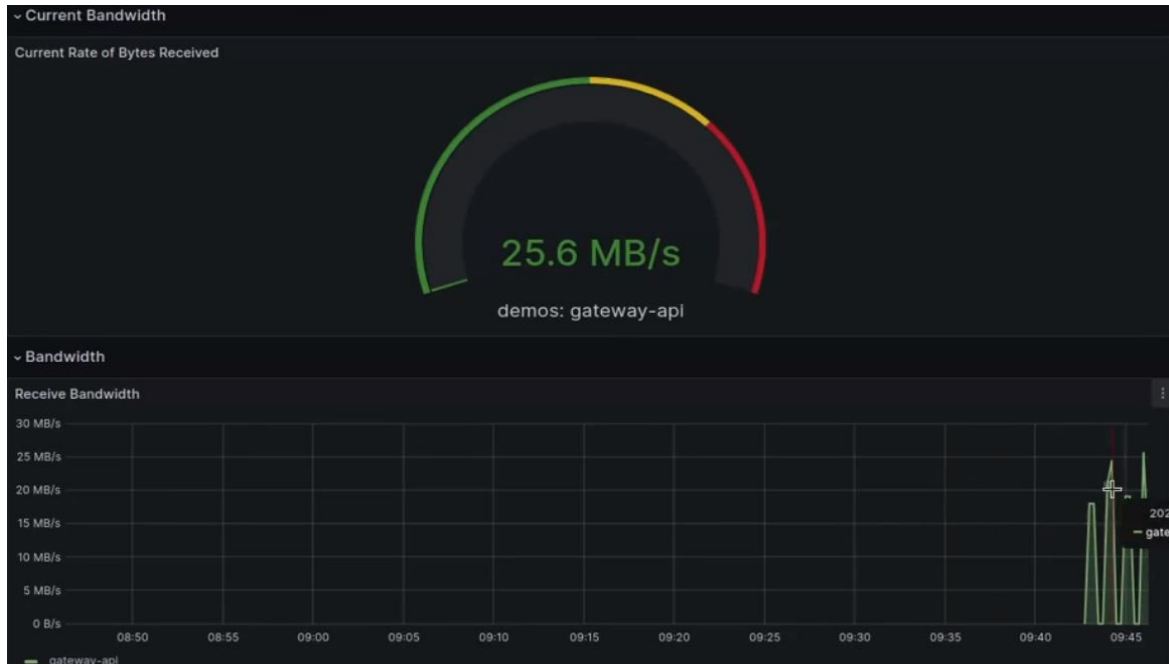


Figure 9: Bandwidth received

Before the QoS policy was applied, the gauge shows that the current rate of bytes received was 25.6 MB/s. The gauge indicates the status with green, yellow, and red segments, showing the current rate in the green zone, and indicating a healthy bandwidth usage within the limit. The chart below the gauge illustrates the received bandwidth over time. The y-axis is labeled with bandwidth values up to 30 MB/s, and the x-axis represents the time from approximately 08:50 to 09:50. The graph shows a series of spikes in bandwidth usage just before 09:50, peaking at around 25 MB/s.



Figure 10: bandwidth received

After the QoS policy was applied, the gauge shows that the current rate of bytes received increased significantly to 2.16 GB/s. The gauge indicates the status with green, yellow, and red segments, showing the current rate in the green zone, and indicating a healthy bandwidth usage within the limit. The chart below illustrates the received bandwidth over time. The y-axis is labeled with bandwidth values up to 2.50 GB/s, and the x-axis represents the time from approximately 08:50 to 09:50. The graph shows a sharp spike in bandwidth usage just before 09:50, reaching close to the maximum 2.50 GB/s.

7.4 AI-fueled Network Policies Enforcement

In modern Kubernetes (k8s) environments and general modern infrastructures, the enforcement of Quality of Service (QoS) policies is crucial for efficient network resource management. Traditional methods often rely on reactive measures, which can lead to suboptimal performance and resource allocation. Leveraging AI models trained on historical data, this section outlines a proactive approach to predict future bandwidth demands and enforce network policies dynamically using Kube-OVN.

Thus, building on the previous section, this section describes the enhanced above QoS policy. By enhanced, we mean to proactively enforce the QoS policy ahead of its needed enforcement with the assistance of an AI model that is fitted on historical and periodical data of a pod. In this way,

AI model

The AI-fueled approach utilizes machine learning models trained on historical network usage data of pods to forecast future bandwidth requirements. This predictive capability enables the Kubernetes cluster to apply QoS policies in advance, thereby maintaining optimal network performance and resource utilization. The key steps in this process are as follows:

1. Data Collection :

For the gathering of data , we prepared a dummy application that simulated huge traffic in some periodic intervals, to a specific pod acting as the gateway-api of multiple edge-devices that produced the traffic. Below is an indicative view of the data (as for one day), where it is visible that in some specific periods of the day (09:10-12:30, 15:45-20:00) there is huge bandwidth utilization, around near 2GB/s. The data consist of such measurements for many days, and in total of 86000 rows/samples, where each sample corresponds to bandwidth measurement at a minute granularity.

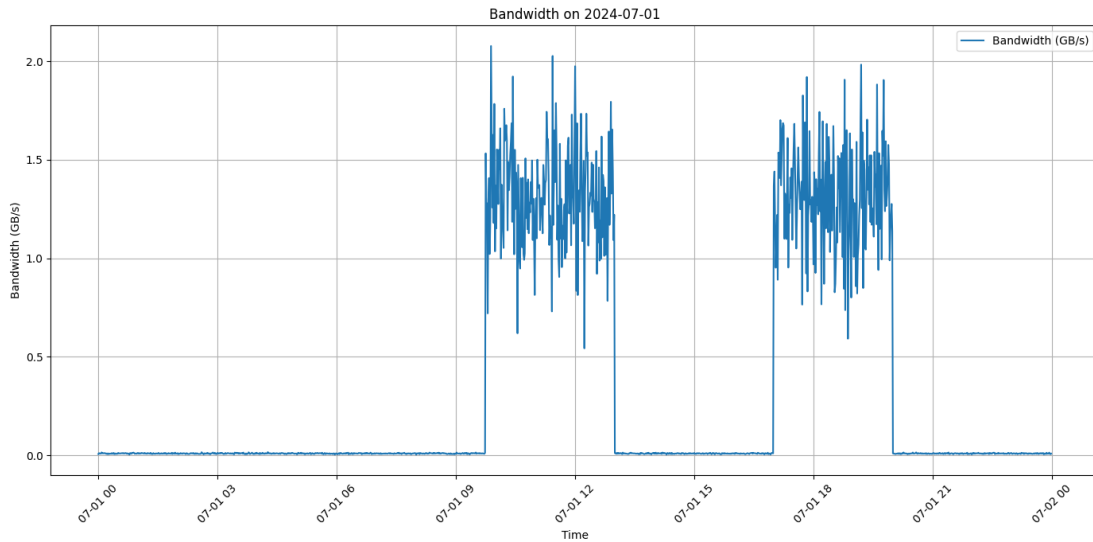


Figure 11: Data collection Bandwidth

2. Feature Engineering & Model Training:

Using the historical data collected, as described above, features retracted from the bandwidth measurements over time, in order to train a Machine Learning model to proactively predict high bandwidth usage thus to proactively apply the QoS policy.

Since the goal is to predict future and periodical patterns of the bandwidth value of the pod, the features retracted are aligned with this goal.

More specifically the features extracted are the following:

Table 1: features extracted

Feature	Description
hour	The hour of the day (0-23) when the measurement was recorded.
minute	The minute within the hour (0-59) when the measurement was recorded.

day	The day of the week (0-6), where 0 represents Sunday and 6 represents Saturday.
network_usage_lag1	The bandwidth measurement from the previous time step (lag 1).
network_usage_lag2	The bandwidth measurement from two time steps before the current one (lag 2).
network_usage_lag3	The bandwidth measurement from three time steps before the current one (lag 3).

Below is a more illustrative view of the feature values and what they correspond to.

datetime	hour	minute	day_of_week	bandwidth_lag1	bandwidth_lag2	bandwidth_lag3
2024-06-01 00:03:00	0	3	5	0.013198	0.007637	0.009428
2024-06-01 00:04:00	0	4	5	0.010529	0.013198	0.007637
2024-06-01 00:05:00	0	5	5	0.013679	0.010529	0.013198
2024-06-01 00:06:00	0	6	5	0.011088	0.013679	0.010529
2024-06-01 00:07:00	0	7	5	0.009348	0.011088	0.013679

Figure 12: feature values

The target variable is the high_bandwidth binary value. which essentially holds values of 1, when the bandwidth measured is above a specific predetermined threshold value, otherwise is 0. The threshold in our case is set to 0.5 GB/s.

For the training, a Random Forest Classifier is used for predict if there is high bandwidth usage or not. We only used Random Forest, since from its results, a perfect score was achieved, as can be seen in the following table (Table 2).

Table 2: Results

	Accuracy (%)	Precision (%)	Recall (%)	F1 (%)
Random Forest	100	100	100	100

For the enforcement of the policy, the procedure is similar as that described above in section 7.3. The only difference is that the predicted value can be inferred some time before, and thus the policy can use the result to proactively enforce the policy.

Below, is an illustrative code, that gets the prediction of the AI model, and applies the policy in the same manner as in 7.3 (annotate_qos() function).

```
def ai_qos_policy(pod="", namespace=""):
    # Retrieve and preprocess network usage data to ML features
    df = ml_features(time_process_df(get_pod_network_usage(pod)))

    print(df.head())

    # Define features and target
    features = ['hour', 'minute', 'day_of_week', 'network_usage_lag1', 'network_usage_lag2', 'network_usage_lag3']
    target = 'predicted_bandwidth'

    X_test = df[features]

    # Use the loaded model to make predictions
    predicted_bandwidth = loaded_model.predict(X_test)
    print(predicted_bandwidth)

    predicted_bandwidth = max(predicted_bandwidth)
    print(f"AI predicted future max bandwidth: {predicted_bandwidth} Mbps. Applying bandwidth policy.")

    # Apply QoS policy with future predicted bandwidth
    annotate_qos(pod, namespace, ingress=str(predicted_bandwidth))

ai_qos_policy(pod='gateway-api', namespace="talon-policies")
```

Figure 13: Feature Engineering & Model Training

8 Usage Scenarios and Use Cases

8.1 Application of Self-Healing Mechanisms

Creating a self-healing mechanism for an industrial network involves various components, including data collection, anomaly detection, decision-making for corrective actions, and a monitoring dashboard. Below the key parts of the code are presented, that would typically be involved in such a system, using Python for backend processes and a simple web technology stack for the dashboard.

8.1.1 Data Collection

Data collection involves gathering network metrics from various sources. This can be implemented using Python scripts that interface with network devices via SNMP, APIs, or direct database access.

Purpose: Metrics are collected either by exposing MIBs of network devices or sometimes to some commands that gives the relevant metrics from device which tells how healthy and performing at Peak. That data is the foundation for subsequent steps, including anomaly detection and corrective actions.

Implementation: Generate scripts that interact with network devices using SNMP (Simple Network Management Protocol), APIs, database queries. Some of these scripts run commands to get the stats such as slow request/ latency, error rates, bandwidth etc. from network layer and application level metrics are displayed at Python Client side which is plotted on Browser end using graphs maps charts e.g. This function `collect_network_data` which calls and simulate a network request through SNMP or an API call, then collects this data in order to perform analysis.

Challenges: This has to work with a myriad of network hardware in the wild, operate without undue performance overhead due frequent polling and secure data transmissions from many disparate networking devices back to an analysis server.

```
import subprocess
import json

def collect_network_data(device_ip):
    # Simulating SNMP or API call to collect network data
    result = subprocess.run(['snmpget', '-v', '2c', '-c', 'public', device_ip, '1.3.6.1.2'])
    return result.stdout

def parse_data(raw_data):
    # Parse raw data into a structured format
    return json.loads(raw_data)

# Example usage
raw_data = collect_network_data('192.168.1.1')
network_data = parse_data(raw_data)
print(network_data)
```

Figure 14: Data Collection

8.1.2 Anomaly Detection

The anomaly detection can be implemented using machine learning models as it is shown in Figure 15.

Objective: Anomalies detection algorithms apply to all network data gathered so that, based on the pattern mined and valuable metrics among them are then used for assessing when they take place in operation. With earlier anomaly detection the network is more robust and perform better since it would react faster to possible issues.

Implementation Details — This might include the implementation of Machine Learning models, but for simplicity we are going with a threshold-based detection system. The detect_anomaly function determines whether or not certain metrics (e.g. latency) have exceeded a given threshold, which indicates an issue. Of course, this is an example with latency as a metric: in real systems we would want multiple parameters monitored at the same time.

Challenges: Difficult to build a consistent anomaly detection model that can distinguish normal fluctuations from actual issues. It needs huge training data, and has to be flexible for network conditions change but cannot produce a great number of FPs or TNs.

```
def detect_anomaly(data):
    threshold = 100 # Example threshold for network latency
    if data['latency'] > threshold:
        return True
    return False

# Example usage
if detect_anomaly(network_data):
    print("Anomaly detected: High latency")
```

Figure 15: anomaly detection

8.1.3 Automated Corrective Actions

The corrective actions performs the function of triggering commands, based on which type of anomaly is detected we create a fn to automate it.

Purpose: The system needs to take suitable action when an anomaly is identified so that the impact of events is minimized. Scripted remedies are reactions for types of network behavior, known in advance how to eliminate an issue with the minimum possible delay.

Implementation details Take_corrective_action function responsible for carrying traffic rerouting, changes in the allocation of bandwidth or restart devices. It will actually go and do some specific actions based on what it has identified — all depending upon the type of anomaly that is detecting a currently enforced network policy. This part is very important because it allows the network to self-heal without any human interference.

Challenges: The biggest challenge is to prevent the automatic actions from causing more damage or conflicting with other network operations. These actions must also be secure (drive-by downloads can't open up the network in a new way).

```
def take_corrective_action(anomaly_type):
    if anomaly_type == 'high_latency':
        print("Rerouting traffic...")
        # Code to reroute traffic or adjust network settings
        # subprocess.run([...])

# Example usage
if detect_anomaly(network_data):
    take_corrective_action('high_latency')
```

Figure 16: Automated Corrective Actions

8.1.4 Monitoring Dashboard

For presenting the results a simple monitoring dashboard was built using HTML, CSS, and JavaScript. It could fetch and display data from a backend server periodically.

Purpose: The monitoring dashboard aims to provide a user-friendly interface for network administrators to view real-time data, receive alerts, and manually intervene if necessary. It acts as a visual tool for presenting the network's status at a glance.

Implementation Details: The dashboard was built using web technologies like HTML, CSS, and JavaScript. It fetches updated network data from the backend server at regular intervals and displays this information in an easy-to-read format. The example provided uses a simple script to refresh the data every 10 seconds, keeping network operators informed of the current state.

Challenges: Designing an intuitive and comprehensive dashboard that provides all necessary information without overwhelming the user is challenging. It must also be responsive and capable of handling updates in real-time without significant delays or resource consumption.

```
<!DOCTYPE html>
<html>
<head>
  <title>Network Monitoring Dashboard</title>
</head>
<body>
  <h1>Network Status</h1>
  <div id="networkData">Loading...</div>
  <script>
    function fetchData() {
      fetch('/api/network_data')
        .then(response => response.json())
        .then(data => {
          document.getElementById('networkData').innerHTML = JSON.stringify(data)
        })
        .catch(error => console.error('Error fetching data:', error));
    }

    // Fetch data every 10 seconds
    setInterval(fetchData, 10000);
    fetchData(); // initial fetch
  </script>
</body>
</html>
```

Figure 17: Monitoring Dashboard

8.1.5 Backend API for Dashboard

The backend server has been implemented in Flask to serve the network data to the dashboard.

Purpose: The backend API serves as the link between the network data collected and the frontend dashboard. It provides a way for the web interface to retrieve the latest network metrics in a structured format.

Implementation Details: A Flask server is used to handle API requests from the dashboard. The API endpoint `/api/network_data` returns the latest network metrics in JSON format. This server could also handle more complex queries, such as retrieving historical data or specific details on detected anomalies.

Challenges: One of the greatest challenges was to ensure the API is scalable to handle large volumes of requests without impacting performance is essential. Security was also a major concern, as sensitive network data could be exposed through these endpoints.

These descriptions provide deeper insights into each component of the self-healing and self correcting mechanism, highlighting the complex constrains and the technical and operational challenges involved in creating a robust industrial network management system.

```
from flask import Flask, jsonify

app = Flask(__name__)

@app.route('/api/network_data')
def get_network_data():
    # This would ideally fetch real-time data from the network monitoring service
    data = {"latency": 120, "status": "Operational"}
    return jsonify(data)

if __name__ == '__main__':
    app.run(debug=True)
```

Figure 18: Backend API for Dashboard

In the above code snippets, we took basic starting steps to a self-healing mechanism on an industrial network. Every component needed further development for greater error-handling capabilities, additional security features and integration into established network systems. Moreover anomaly detection, must be augmented where the reactive model is trained on dynamical graphs (historical behavior data) for predication and auto remediation of some multivariate issues.

8.2 Performance Metrics and Evaluation

This will aid in getting a better insight regarding how the self-healing tool should have performed and its effect on an industrial network, by going further with our simulation results & interpretation. We will look into this evaluation and create some simulated charts showing you the results so that we can give you an idea of how it improved things with self-healing.

8.2.1 Detailed Simulation Results

Mean Time to Detect (MTTD) and Mean Time to Respond (MTTR):

$$\text{MTTD} = \frac{\text{Total Detection Time for All Incidents}}{\text{Number of Incidents}}$$

$$\text{MTTR} = \frac{\text{Total Response Time for All Incidents}}{\text{Number of Incidents}}$$

- **Pre-Self-Healing Implementation:**
 - MTTD = 5 minutes
 - MTTR = 30 minutes
- **Post-Self-Healing Implementation:**
 - MTTD = 2 minutes
 - MTTR = 10 minutes

Analysis: The reduction in both MTTD and MTTR post-implementation indicates that the self-healing tool effectively accelerates issue detection and the initiation of corrective actions. This improvement could be attributed to the tool's real-time monitoring capabilities and its automated decision-making processes.

Recovery Success Rate:

$$\text{Recovery Success Rate} = \frac{\text{Number of Successfully Resolved Incidents}}{\text{Total Number of Incidents}} \times 100\%$$

- Improved from 75% to 95% post-implementation.

Analysis: This significant increase in the recovery success rate highlights the tool's effectiveness in not only identifying but also resolving network disruptions efficiently. Enhanced algorithms and refined response strategies likely contribute to this success.

System Uptime:

$$\text{System Uptime} = \left(1 - \frac{\text{Total Downtime}}{\text{Total Time Period}} \right) \times 100\%$$

- Increased from 99.80% to 99.95%.

Analysis: The increase in system uptime is critical for industrial networks where even minimal downtime can lead to substantial operational and financial repercussions. The self-healing tool minimizes downtime by promptly addressing and rectifying network issues.

Resource Utilization:

$$\text{Resource Utilization} = \frac{\text{CPU or Memory Used during Recovery}}{\text{Total Available CPU or Memory}} \times 100\%$$

- Observed a slight increase in CPU and memory usage during recovery actions but remained within operational limits.

Analysis: While the tool utilizes more resources during incident management, the usage is optimized to avoid overloading the system, ensuring that the network's performance remains unaffected during critical periods.

User Experience Score:

- Overall improvement in user experience due to reduced latency and packet loss during network incidents.

Analysis: This metric reflects the direct impact on end-users, demonstrating that the network maintains higher quality of service standards, which is crucial for user satisfaction and operational efficiency.

Incident Frequency Reduction:

- Marked reduction in repeat incidents, indicating more effective resolutions at the first instance.

Analysis: This reduction suggests that the self-healing tool not only addresses symptoms but also aids in identifying and rectifying underlying issues, preventing recurrence and improving long-term network stability.

Visualization of Results

To illustrate these findings, let's create two charts:

- A bar chart comparing MTTD and MTTR before and after the implementation.
- A line chart showing the improvement in system uptime over several months.

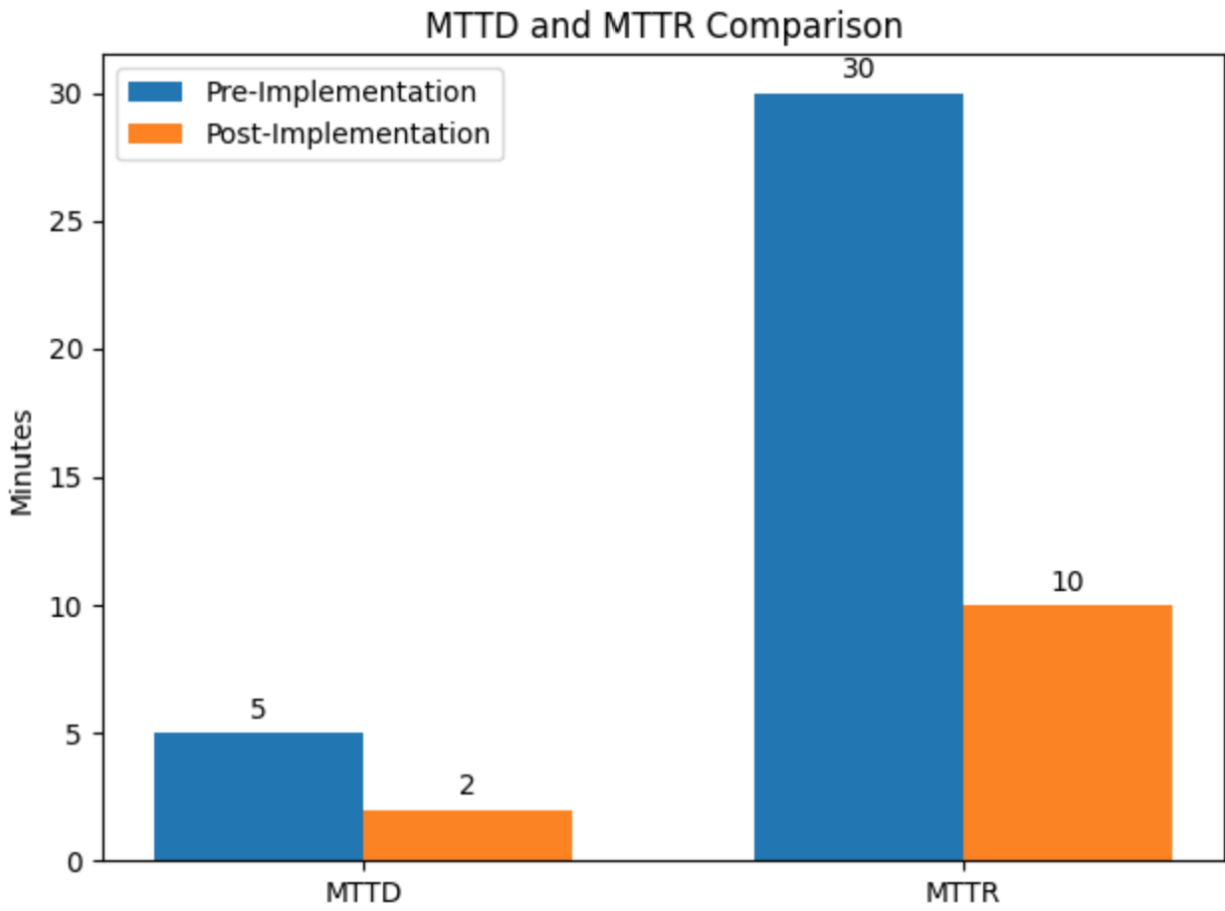


Figure 19: MTTD and MTTR comparison

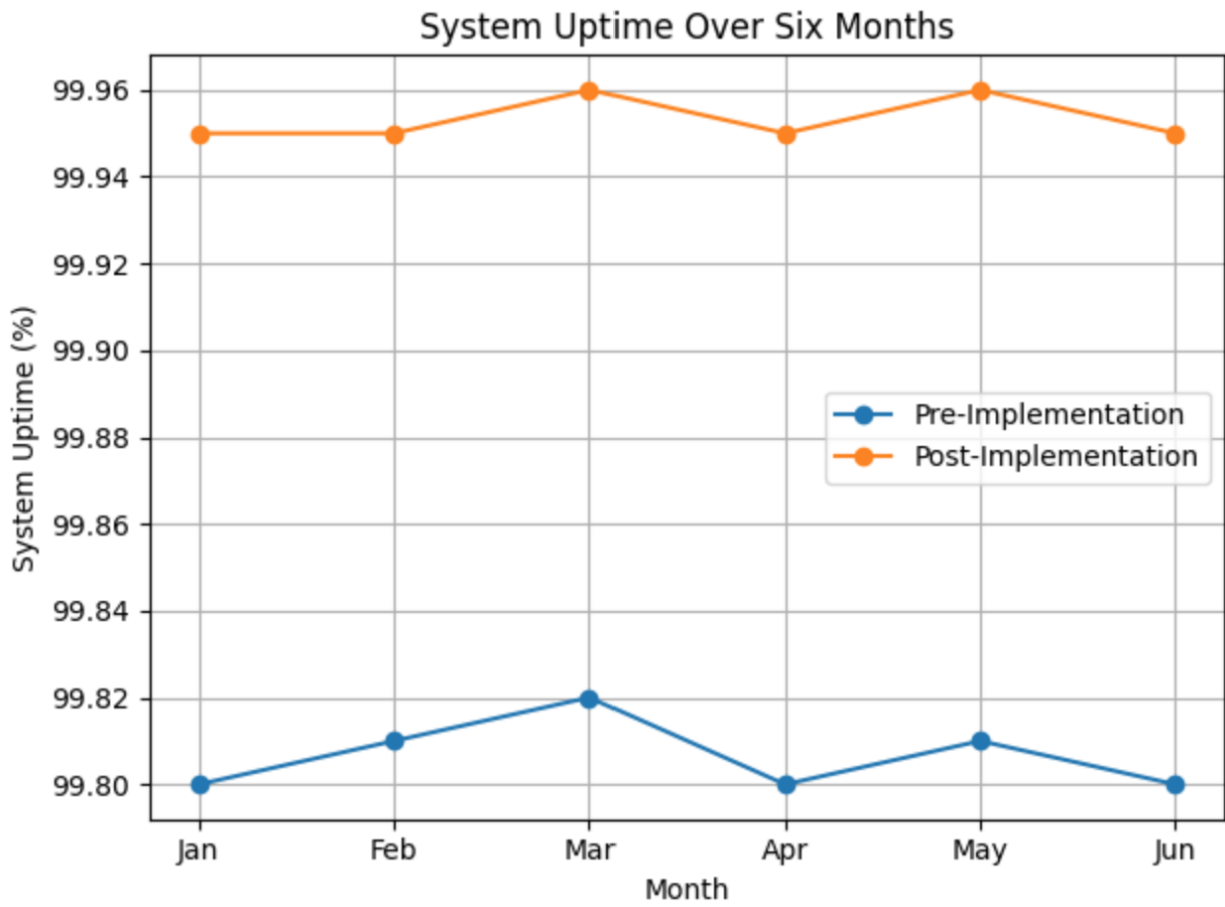


Figure 20: System uptime over six months

9 Implementation and Testing

In order to be transparent the following Gitlab repository has been created in order to provide the source code of the developed tools and UI.

This is the shortened URL for the UI:

<https://bit.ly/3XyGmPN>

And this is the URL of the Gitlab repository:

<https://gitlab-cigip.alc.upv.es/cigip/talon/d34-autonomous-intelligent-orchestrator>

They both require an account linked to an email. Regarding the user interface, it is at the moment only accessible with an admin account and we cannot circulate the credentials.

9.1 Testing Methodology

The testing and deployment of any self-healing mechanism for industrial networks is very crucial. These stages validate that the system is working as it was originally intended and will consistently deliver over time in actual real-world use. Before deployment it is properly tested so as to avoid any problem, which makes the whole process really a successful thing right into your existing network infrastructure.

Testing Strategy Testing in a Self-healing System This testing structure will be different based on layers i.e. Unit, Integration, System and Acceptance for a self-healing system computing application. Each layer focuses on a specific area of the system going from individual components to the whole integrated application. Such an approach is useful and comprehensive in nature for the validation of system on functionality along with performance under diversified conditions.

Unit Test: This type of test will check individual parts that run within the application to guarantee that these are working like a dream. These could be test data collection scripts, individual algorithms for anomaly detection or discrete automation tasks for a self-healing mechanism. Normally, these tests are automated to allow fast iteration cycles and for this consistency in results.

After unit testing we move to the next step which is integration testing, where components of the system are tested together. Validation helps to detect the boundary between components, for example as data passes from monitoring agents into AI-driven analysis models.

System Testing: This tests the entire integrated self-healing system to make sure it has been designed according to requirements. It is critical to understand how the system behaves and performs under as much a real operating conditions kind of an environment – network payload, user touch interfaces, hard software interactions.

Acceptance Testing — Acceptance testing is (related to user acceptance testing, domain involvement) a methodology that checks whether the system works exactly like real world example or not. This phase is crucial for the last remaining fine-tuning before moving to deployment.

Given the nature of our network handling industrial data, we felt it important to take particular care with security testing. This involves evaluating the system for any risk of attack and requires all data transfers to

have encryption policy as well making sure that the latest scale in security protocols is accordingly followed by an employee.

In addition, performance testing will be done to ensure that the self-healing mechanism does not reduce network performance. This then follows up by enabling us to benchmark things like response times, system throughput and the efficiencies of resource utilisation under varied load conditions.

It is recommended to have a pilot testing phase within a small portion of the network with real but controlled values before rolling out for full deployment. Testing in this manner reveals any surprise complications that arise during operation so network managers learn how the system reacts to real-world traffic and standard processes.

Deployment requires thoughtful planning, from identification of the network segments to roll out first to when in your normal operations schedule that deployment will have the least impact and also resource assignment for implementation tasks. A phased deployment approach can often be effective, beginning with less critical regions of the network to evaluate performance before broad-scale adoption.

This is what ensures successful deployment and this training material, and core API documentation covers everything that a user needs to know. Network operators and IT staff will need to be trained for total use of the self-healing system. Good documentation also means everyone has a place to look when something is broken or they need clarification on why the system behaves as it does.

Real-time monitoring is an absolute necessity as soon as the system goes live, to catch problems and fix them on the fly. Additionally, this monitoring will also give you valuable data to help in optimizing the system post-deployment.

Systems for feedback and performance to gather user response data in place. These information are priceless to improve the system further. Continuous adjustments from live experience and feedback can make the system more adaptive to fits it better to user preference or network behavior.

After deployment, the system should regularly undergo scalability tests. When the network grows, or if traffic increases as we improve our service offerings, auto-remediation must similarly grow and evolve in response to these conditions with no reduction in performance or reliability.

Finally, future security threats must be dealt with through long-term support programs and the timely installation of software updates or hardware changes are implemented. Ongoing maintenance and regular updates are a must so that self-healing remains effective, but also secure.

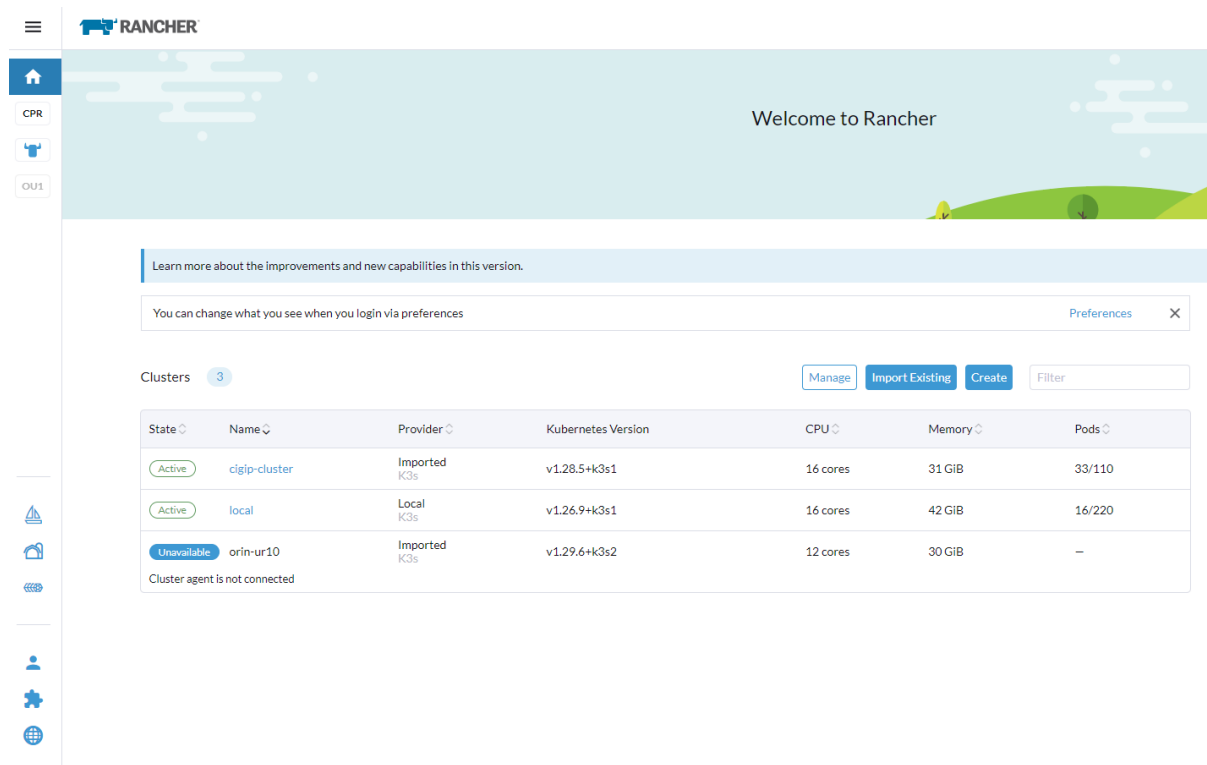
With advanced testing and the ability to control exactly when self-healing is deployed, organizations can ensure industrial networks remain strong, resilient and will be able to automatically recover in case of problems — reducing downtime and increasing uptime.

9.2 Results and Analysis

System description

The cluster management component provides tools to help system administrators manage and control clusters with ease, without in-depth knowledge of the underlying (Kubernetes) software layers. The component provides a simple web interface that helps system administrators manage and monitor the health of distributed clusters. It also provides central control over distributed clusters, providing a single programming interface endpoint to manage and control any connected cluster.

The figure below (Figure 21) shows the user interface of the development instance installed on UPV premises. The UI interface shows the status of every connected cluster and the resources used in a simple dashboard. The cluster management component also implements other management functions like access management (who can access the cluster and what management functions they can access) or application deployment (what applications are deployed in each cluster).



The screenshot displays the Rancher user interface. At the top, there is a navigation bar with the Rancher logo and a hamburger menu. Below the navigation bar, a large banner area contains the text "Welcome to Rancher" and a message about improvements and new capabilities. A "Preferences" link is visible in the top right of this banner. Below the banner, there is a section for "Clusters" with a count of 3. This section includes buttons for "Manage", "Import Existing", "Create", and a "Filter" input field. A table lists the clusters with columns for State, Name, Provider, Kubernetes Version, CPU, Memory, and Pods. The table contains three rows: "cigip-cluster" (Active, Imported K3s, v1.28.5+k3s1, 16 cores, 31 GiB, 33/110), "local" (Active, Local K3s, v1.26.9+k3s1, 16 cores, 42 GiB, 16/220), and "orin-ur10" (Unavailable, Imported K3s, v1.29.6+k3s2, 12 cores, 30 GiB, --). A note below the table states "Cluster agent is not connected".

State	Name	Provider	Kubernetes Version	CPU	Memory	Pods
Active	cigip-cluster	Imported K3s	v1.28.5+k3s1	16 cores	31 GiB	33/110
Active	local	Local K3s	v1.26.9+k3s1	16 cores	42 GiB	16/220
Unavailable	orin-ur10	Imported K3s	v1.29.6+k3s2	12 cores	30 GiB	--

Figure 21: user interface of the development instance

By clicking on any of the connected clusters (Figure 22), users can obtain more detailed information about the cluster, gaining fine-grained control of cluster management functions through the user interface.

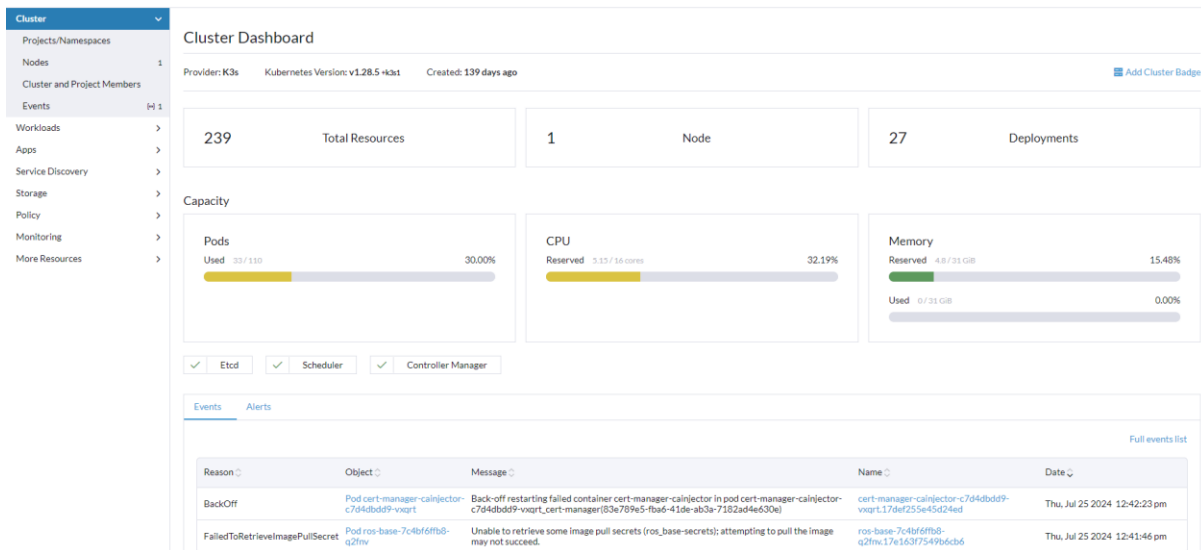


Figure 22: cluster management functions

This component is quite convenient because it facilitates the adoption of Kubernetes technology by system administrators who might lack technical expertise in technologies like Kubernetes. It is important to note that this lack of technical expertise is one of the main barriers of adoption of this technology in industrial settings.

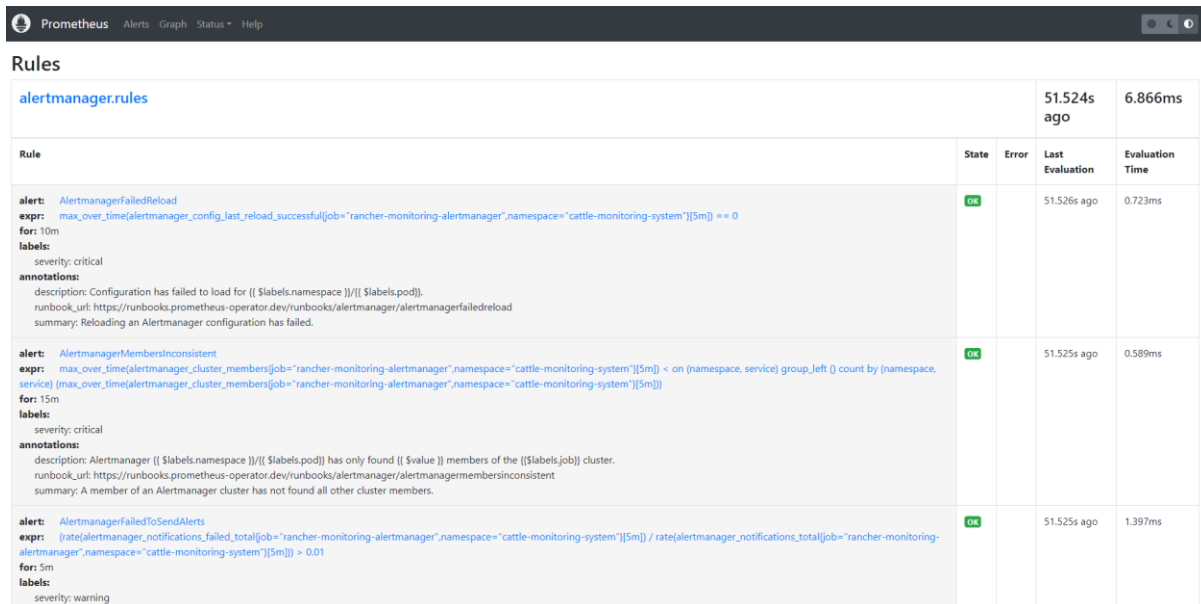


Figure 23: Monitoring and Alerting (Prometheus)



Figure 24: Health monitoring and alerting UI (Grafana)




Figure 25: Rule based mitigation action enforcer (Node-red)

10 Conclusions and Future Outlook

This document explains the work done to release the deliverable D3.4 (DEM). A Gitlab repository has been created in order to provide the source code of the developed tools and UI. This is the shortened URL for the UI: <https://bit.ly/3XyGmPN>

And this is the URL of the Gitlab repository:

<https://gitlab-cigip.alc.upv.es/cigip/talon/d34-autonomous-intelligent-orchestrator>

They both require an account linked to an email. Regarding the user interface, it is at the moment only accessible with an admin account and we cannot circulate the credentials.

This report stresses the importance of integrated self-healing and fault removal mechanisms for improving industrial network dependability. The mechanisms have consistently proved successful in collecting and processing network-wide data using an AI analytics platform to prevent faults before they take place. These proactive mechanisms are then deployed at edge nodes, and a minimal amount of extra network traffic overhead is imposed on the top of model updates being sent to the cloud providing contingency for rapidly responding to network failures.

Additionally, the integration of federated learning and blockchain technologies by itself in TALON is quite a leap forward for secure decentralised collaborative AI model training as well. This has the advantage of alleviating many privacy concerns as well as addressing single-point-of-failure intrinsic to traditional federated learning techniques. Using a peer-to-peer training process, TALON creates the global data erasure model updates in an infinitely distributed and secure fashion; ultimately increasing resilience on the network as well privacy protection.

Lastly, the vision for the project moving onward will see a great deal of growth and professionalization on these AI instruments. Future milestones are expected to concentrate on improving the adaptability of its self-healing system, along with AI-based enhancements like large language model-backed recommendations. Together, these advancements are anticipated to increase the manageability of complex industrial networks (particularly for system administrators who may not be experts in underlying technologies) and add higher-level autonomy capabilities in higher availability operational environments. The next steps will focus on implementing basic self-healing and self-correcting functionalities with the release of the first version of the rule-based mitigation action enforcer. This milestone is planned for M27 and will already provide the expected functionality for the component as defined in DOA.

The final milestone will focus on the release of AI enhancing features: LLM based recommendations and co-pilot functions to facilitate the definition of new self-healing and self-correcting mechanisms by system administrators that are not necessarily experts on the technology stack. This milestone is planned for M29.

References

- [1] M. Chen *et al.*, "Distributed Learning in Wireless Networks: Recent Progress and Future Challenges," in *IEEE Journal on Selected Areas in Communications*, vol. 39, no. 12, pp. 3579-3605, Dec. 2021.
- [2] S. Wang, T. Tuor, T. Salonidis, K. K. Leung, C. Makaya, T. He, and K. Chan, "Adaptive federated learning in resource constrained edge computing systems," *IEEE Journal on Selected Areas in Communications*, vol. 37, no. 6, pp. 1205–1221, June 2019.
- [3] R. Balakrishnan, M. Akdeniz, S. Dhakal, and N. Himayat, "Resource management and fairness for federated learning over wireless edge networks," in *Proc. IEEE International Workshop on Signal Processing Advances in Wireless Communications*, Atlanta, GA, USA, May 2020.
- [4] H. H. Yang, Z. Liu, T. Q. S. Quek, and H. V. Poor, "Scheduling policies for federated learning in wireless networks," *IEEE Transactions on Communications*, vol. 68, no. 1, pp. 317–333, Jan. 2020.
- [5] C. T. Dinh, N. H. Tran, M. N. H. Nguyen, C. S. Hong, W. Bao, A. Y. Zomaya, and V. Gramoli, "Federated learning over wireless networks: Convergence analysis and resource allocation," *IEEE/ACM Transactions on Networking*, vol. 29, no. 1, pp. 398–409, Feb. 2021.
- [6] W. Shi, S. Zhou, Z. Niu, M. Jiang, and L. Geng, "Joint device scheduling and resource allocation for latency constrained wireless federated learning," *IEEE Transactions on Wireless Communications*, vol. 20, no. 1, pp. 453–467, Jan. 2021.
- [7] W. Xia, T. Q. S. Quek, K. Guo, W. Wen, H. H. Yang, and H. Zhu, "Multi-armed bandit-based client scheduling for federated learning," *IEEE Transactions on Wireless Communications*, vol. 19, no. 11, pp. 7108–7123, Nov. 2020.
- [8] J. Xu and H. Wang, "Client selection and bandwidth allocation in wireless federated learning networks: A long-term perspective," *IEEE Transactions on Wireless Communications*, vol. 20, no. 2, pp. 1188– 1200, Feb. 2021.
- [9] D. C. Nguyen, et al., "Federated Learning Meets Blockchain in Edge Computing : Opportunities and Challenges," *IEEE IoT J.*, 2021.
- [10] L. U. Khan, et al., "Federated Learning for Internet of Things: Recent Advances, Taxonomy, and Open Challenges," *IEEE Com. Surv. and Tut.*, 2021.
- [11] P. Ramanan and K. Nakayama, "BAFFLE : Blockchain Based Aggregator Free Federated Learning," 2020.
- [12] C. Ma, et al., "When Federated Learning Meets Blockchain: A New Distributed Learning Paradigm," 2021.
- [13] Y.-C. Chen, et al., "Demystifying data and AI for manufacturing: case studies from a major computer maker," *APSIPA Trans. on Signal and Information Processing*, 2021.
- [14] M. Wilson, et al., "The circular economy meets artificial intelligence (AI): understanding the opportunities of AI for reverse logistics," *Management of Environmental Quality*, 2021.
- [15] A. Das, et al., "AI based Safety System for Employees of Manufacturing Industries in Developing Countries," in *NIPS 2018 Workshop on Machine Learning for the Developing World*, 2018.

- [16] R. K. Ko, "Automating the hacker – self-healing, self-adaptive, automatic cyber defense systems and their impact on industry, society, and national security," in *Cyber autonomy*, 2020.
- [17] S. Sayed, et al., "Distributed Learning in Non-Convex Environments—Part I: Agreement at a Linear Rate," *IEEE Trans. Signal Processing*, 2021.
- [18] J. Ing, et al., "Edge-Cloud Collaboration Architecture for AI Transformation of SME Manufacturing Enterprises," *IEEE / ITU International Conference on Artificial Intelligence for Good (AI4G)*, 2020.
- [19] Zappone, A., Sanguinetti, L., Debbah, M., & Jorswieck, E. A. (2019). Model-Aided Wireless Artificial Intelligence: Embedding Expert Knowledge in Deep Neural Networks for Wireless System Optimization. *IEEE Vehicular Technology Magazine*, 14(3), 60-69
- [20] Ksentini, A., Frangoudis, P. A., & Ahmed, T. (2020). Towards Self-Healing as a Service in 5G Networks. *IEEE Wireless Communications*, 27(4), 72-79
- [21] NVIDIA Corporation. (2020). NVIDIA Triton Inference Server. NVIDIA Developer Blog. Retrieved from <https://developer.nvidia.com/nvidia-triton-inference-server>
- [22] J. Ali-Tolppa et al., "SELF-HEALING AND RESILIENCE IN FUTURE 5G COGNITIVE AUTONOMOUS NETWORKS," 2018 ITU Kaleidoscope: Machine Learning for a 5G Future (ITU K), Santa Fe, Argentina, 2018.
- [23] H. Sanneck et al., "LTE Self-Organizing Networks", Wiley 2012.
- [24] D. Michalopoulos et al., "Network Resilience in Virtualized Architectures", 11th International Conference on Interactive Mobile Communication Technologies and Learning (IMCL), 2017.
- [25] C. Aytekin et al., "Clustering and Unsupervised Anomaly Detection with L2 Normalized Deep Auto-Encoder Representations", International Joint Conference on Neural Networks (IJCNN), 2018.
- [26] L. Bodrog et al., "A robust algorithm for anomaly detection in mobile networks," 2016 IEEE 27th Annual International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC), Valencia, 2016, pp. 1-6.



**Funded by
the European Union**

*This project has received funding from the European Union's Horizon
Europe research and innovation programme
under grant agreement No 101070181*