analysys mason

Perspective

Using AIOps solutions in the telecoms industry: a market assessment

August 2021

Justin van der Lande

Sponsored by Red Hat

RESEARCH

Contents

List of figures

1. Executive summary

This paper examines the key challenges that telecoms operators face when using artificial intelligence (AI) for IT operations (AIOps) to support their key operational processes. AIOps is a term first used by Gartner to summarise the three technology areas originally devised to support IT operational processes: big data, machine learning (ML) and automated operations. The term AIOps is now being applied to a wider set of processes that includes most aspects of OSS/BSS.

The telecoms industry has been a strong supporter of AIOps-based approaches. For example, TM Forum set up the 'AIOps - redesigning operations processes' initiative to help operators to manage their AI deployments.[1] The use of AI is also being incorporated into standards for 3GPP 5G, and the European Telecommunications Standards Institute (ETSI) has defined the network data analytics function (NWDAF) to support automation and the use of ML technologies. In addition, ETSI has defined five categories for AI usage: infrastructure management; network assurance; network operation; service orchestration and management; and network security.

AIOps is helping to shape a larger set of automations and transformations for use across the telecoms sector for which big data is combined with ML to support the automation of key operations. AIOps has reached an advanced stage of development in network assurance functions, including event correlation and analysis, anomaly detection, root cause analysis (RCA), alert prioritisation and using natural language processing to augment report creation and searches. The industry is discussing how the same systems, approaches and tools can be applied to additional operational areas.

This paper defines AIOps and discusses its associated use cases. It provides insights into key technology components, the offerings by major vendors and examples of implementations within CSPs' businesses. our analysis of these implementations includes a discussion of the challenges faced by CSPs, details of the technologies that they used and a summary of the results.

2. What is AIOps?

2.1 Definition

AIOps, or artificial intelligence for IT operations, is the application of advanced analytics in the form of ML and AI and is used to help operations teams to achieve increased efficiency and to better achieve their business goals by enabling operational requests to be processed quickly and accurately.

[1] For more information, see TM Forum's *IG1190 AIOps Service Management Deployment Phase v2.0.1.* Available at: https://www.tmforum.org/resources/standard/ig1190-aiops-service-management-v5-1-0/.

2.2 Which use cases does AIOps cover?

AIOps platforms are increasingly being used to support IT operations such as event correlation and analysis, anomaly detection, RCA and natural language processing (NLP), as well as a wider set of functions such as IT service management (ITSM). Attempts are also being made to use AIOps in conjunction with DevOps solutions, but these efforts are progressing at a slower pace because solutions need to work with complex data sets and must support a large number of operational outcomes. However, such use cases have the potential to have a greater impact on telecoms operators' businesses in the long term.

Figure 2.1: Use cases supported by AIOps

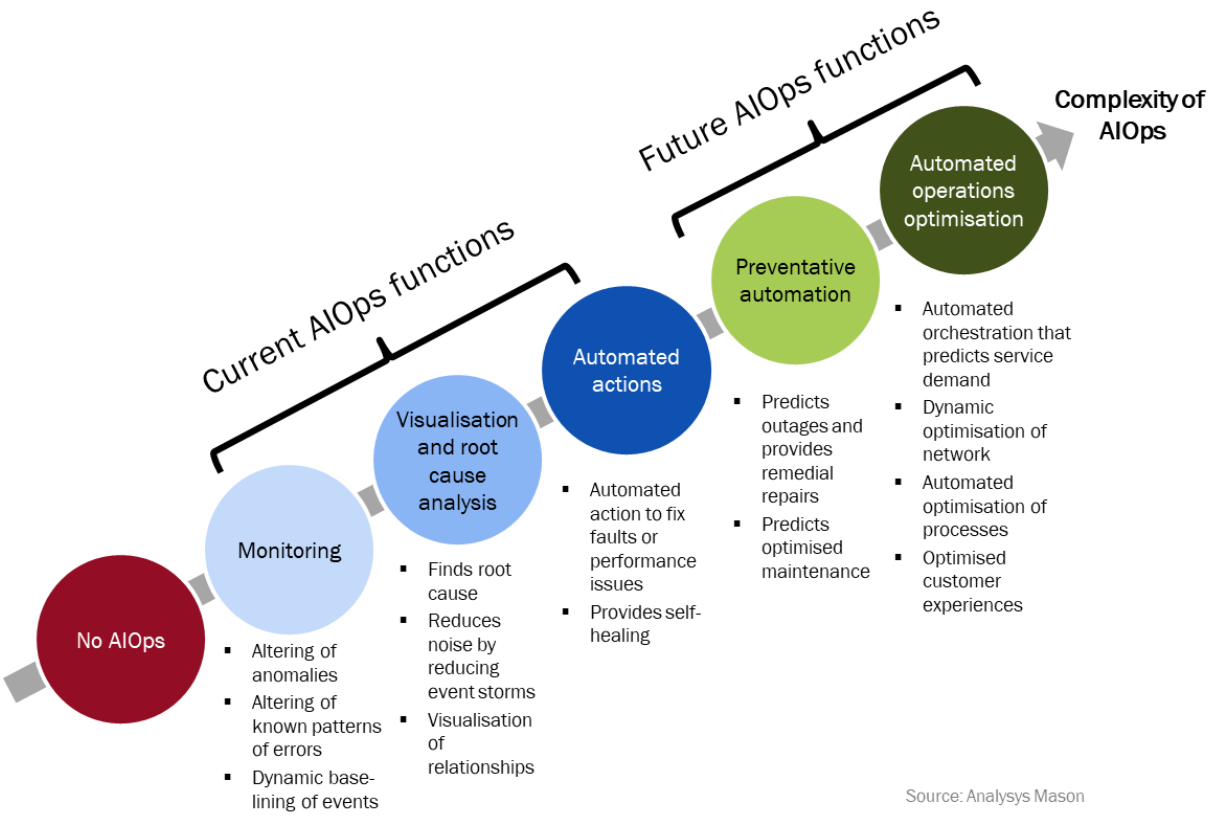

Current AIOps functions in the telecoms industry include the following.

- **Intelligent alerting**. By ingesting data from any part of the IT environment, AIOps filters and correlates meaningful data into 'incidents'. This prevents alert storms where, for example, a failure in system A triggers an alert that affects system B, which also triggers an alert, and is then repeated with each recurring event. Intelligent alerting helps to reduce the number of meaningful events and helps operators to prioritise the issues that are affecting users and the business. Typically, this functionality creates multiple 'rules' that require manual adjustment over time, which limits the ability to swiftly cope with changing situations.

- **Cohort analysis.** The analysis of large data sets becomes difficult for manual or human-based processes beyond five or six data sets. With modern, highly distributed architecture where hundreds of data sets are running at the same time, identifying outliers in configurations or deployed application versions is an insurmountable task for humans. However, AI can collect baseline data and look for outliers to augment human analysis, automatically for performing tasks such as automated maintenance.

- **Cross-domain situational understanding.** AIOps aggregates data and creates causality/relationships, which provides operations teams with an overview of what is at stake and enables them to slice the information as needed to better understand the situation. For example, complex and weak data signals can

be found between different data sources, such as those from networks, IT, applications, weather and outages and services, and once detected, operators are able to act.

- **Automated identification of probable root causes.** Root cause analysis (RCA) identifies the cause of an issue. It builds on the activities of cross-domain situational understanding by drawing on knowledge-based data from previous incidents to provide the most likely cause of an incident.

- **Automated remediation.** AIOps helps to automate closed-loop remediation for known issues. Once problems are identified using RCA analysis of historical issues, AIOps can then suggest the best approach to accelerate remediation. These fixes can then be automatically monitored to check if the issues are resolved.

- **Anomaly detection.** Telecoms data patterns are generally highly predictable which means that anomalies that diverge from 'normal' behaviours can be modelled so that CSPs can flag potential issues, even if they have not previously been encountered. This approach is used in cases such as identity theft, data breaches, fraudulent transactions and to monitor network performance, network latency and network events. AIOps can also create new models using advanced ML algorithms based on big data, having learnt to model 'normal' data patterns with a high degree of accuracy. The predicted anomalies are then prioritised for operations staff, who assess what actions are required, if any.

- **Proactive performance monitoring.** Not all faults cause severe service issues, but the level of severity can increase over time. Performance is one such example; changes in performance over time do not result in service failure but could cause problems in the future. AIOps provides the modelling capability that can help CSPs to predict future issues.

- **Improving customer experience.** The telecoms industry is customer-centric. Long waits in queues for technical support, resolving billing disputes or updating credentials are examples of activities that can lead to customer churn. AIOps can help to provide better network services, reduce outages and improve service performance. In addition, it can help CSPs to understand the complex relationship between network services, customer services and subscriber satisfaction, often expressed as Net Promoter Score (NPS). AIOps provides a framework to enable this highly complex calculation to be made at scale and monitored, which can help to prevent churn or can assist with prioritising service fixes. AIOps is also associated with improving aspects of customer service. For example, it has been used with the introduction of chatbots to automatically resolve customer queries and with ticket routing to reduce customer wait times.

- **Customer segmentation.** AIOps is largely focused on network and IT operations, but marketing operations are a future extension. AIOps can support micro-segmentation for mass-personalisation of subscribers based on their service information, billing information, payment type, revenue generation and location. AIOps identifies thousands of clusters and can then offer real-time suggestions to change buying behaviour.

2.3 Key technology components and core functions

An AIOps framework must include core functions such as data collection, data storage and manipulation, understanding the data context, building and managing data algorithms and the delivery of automations.

Figure 2.2: Key components of AIOps platforms

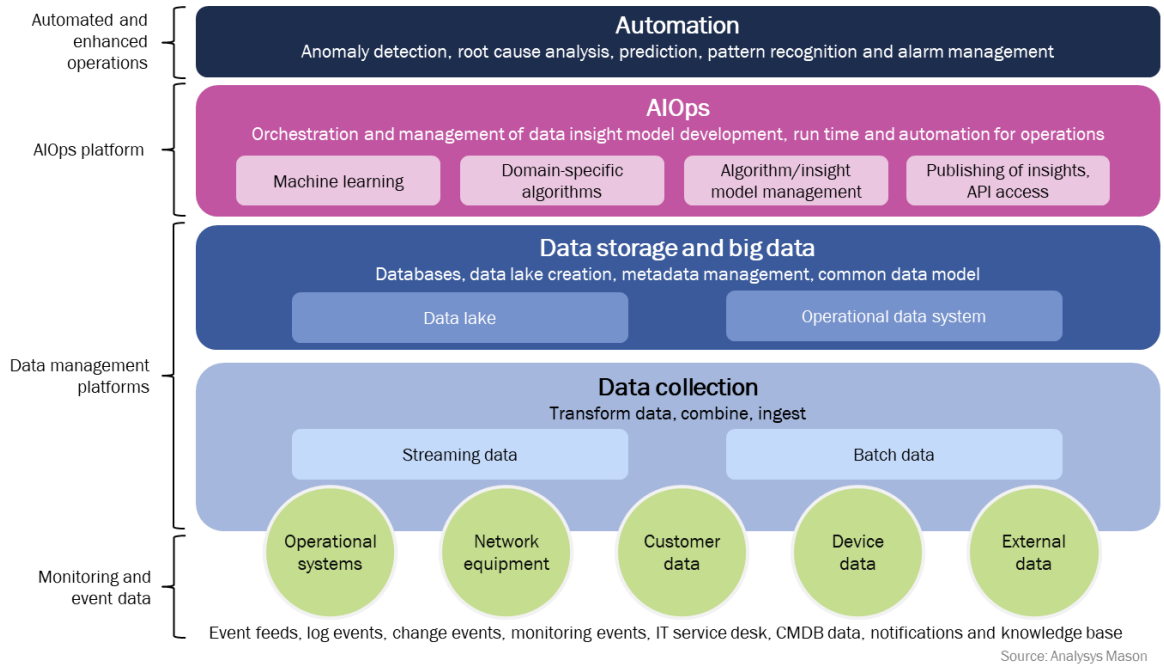

An AIOps platform must be developed in a way that it can be integrated with the components of each operator's larger AI and data framework. A flexible and open approach must therefore be adopted when creating the AIOps platform (where possible), and the use of open-source tools should be considered. Open-source tools help to facilitate collaboration between different vendors because they are widely adopted, represent a neutral method for integration and are usually cheaper than commercial alternatives.

Figure 2.3: Relationship of OSS/BSS applications and processes to the AIOps platform

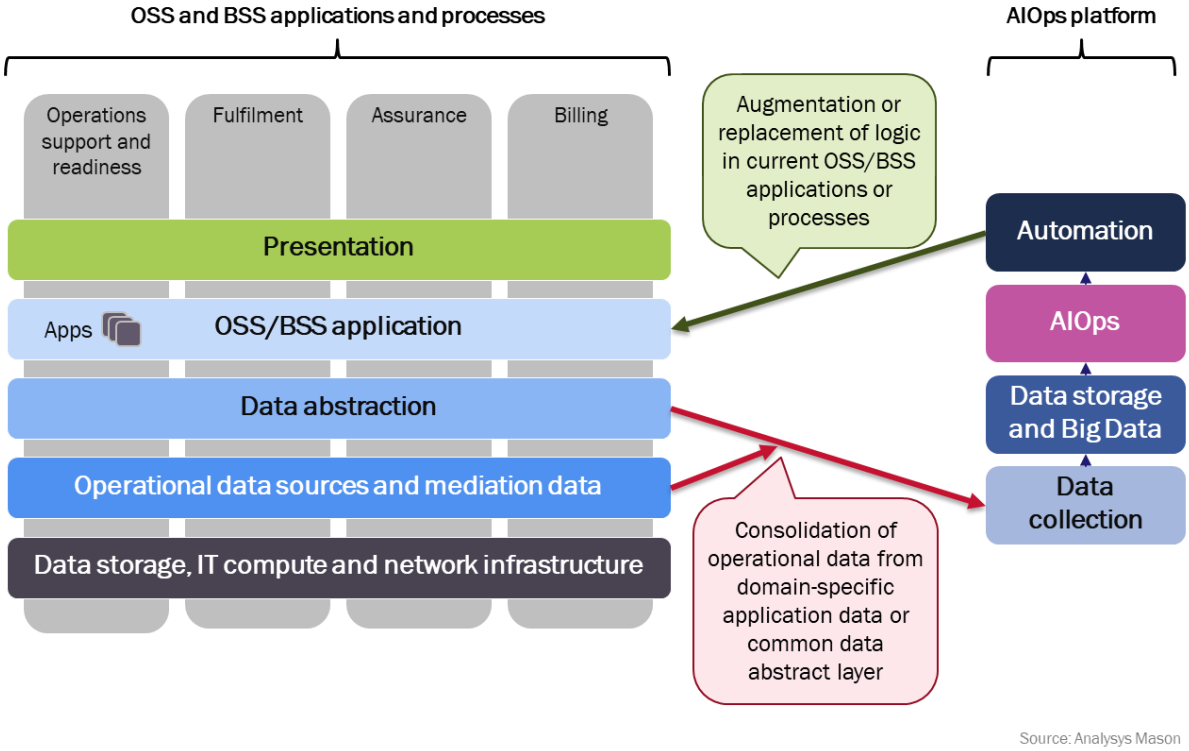

The AIOps platform needs to work in conjunction with existing OSS/BSS applications and processes. This may require the consolidation of operational data from multiple applications to form a common data store, or for the operational data to be accessible for the development of AI tools. The resulting automations and insights can then be fed back to the current applications for action, or the AIOps system can replace current application or process logic.

An off-the-shelf solution for delivering automations rapidly to operations is highly attractive, but it is important to consider the extensibility of this type of solution and its ability to support automations over time. An AIOps platform must also be evaluated in terms of its ability to support and deploy new automations, as well as its ability to update automations that have already been deployed.

2.4 Key characteristics of AIOps solutions

Using AI/ML in operational systems goes beyond building data models (see figure below) with ML technologies; it requires the entire ML modelling lifecycle to be operationalised to work reliably at scale. Building data models is often easier than deploying them, which is why only 25% of modelling projects make it into production.

Building an AIOps solution for the development and deployment of data models requires multiple systems and processes to be created. In addition, the underlying infrastructure must support the rigorous requirements needed to support AI/ML data models in production and during their development. Delays in AI projects are often caused by the preparation of software tools and by needing to provide suitable computer and storage infrastructure.

In a rapidly evolving telecoms market, it is a complex task to establish an integrated set of tools that enable data scientists to rapidly train, develop, test, deploy and optimise their data models. Adopting a suite of software or services from a single cloud provider can resolve many integration and complexity issues but can be costly and reduces the amount of flexibility that independent software vendors (ISVs) or operators have when choosing a technology partner. The openness of a software solution is critical for supporting flexibility: the ability to run on different cloud provider environments or on premises, regardless of the environment that an application is deployed on.

Figure 2.4: Lifecycle of the machine learning data model

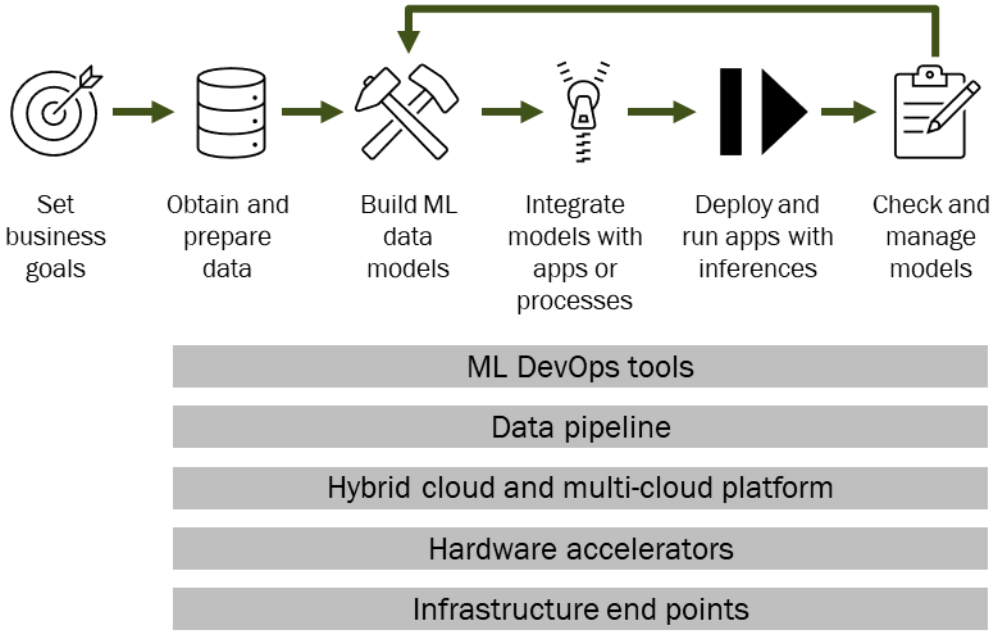

Source: Analysys Mason

AIOps systems requirements

AIOps systems tools need to support the entire ML data model lifecycle after the business goals have been firmly established. The initial gathering of data required for model training relies on the ability to collate all types of the data at scale and from multiple systems. The data is then used to build data models through a number of iterative cycles, and once the data sources are established for each run-time model, data pipelines manage data flows to ensure that the quality of the data is consistent and reliable at scale in production.

ML data modelling software provides the libraries on which data models are created, as well as their management and optimisation. Each model potentially goes through several iterations using different ML or statistical modelling techniques before being released into production.

Each model should be managed by a DevOps process before being integrated into an application or deployed as an automation within a process. Processes may be part of a continuous integration development pipeline and integrated with a deployment process (CI/CD). Highly automated processes enable new data models to be continually deployed, which improves processes and allows data scientists to concentrate on their primary task of modelling.

AIOps infrastructure requirements

Data scientists rely on the underlying IT infrastructure to deliver the resources needed for model training and run-time data model deployment. The use of containers enables resources to be delivered in a consistent and flexible way over any infrastructure. In addition, data scientists, who may require temporary access to additional storage and computer resources, should have an ability to self-serve their needs without having to procure new hardware or services through their IT departments. This would potentially be provided in-house or as a service offering. Having an integrated set of tools that support automated operations for common tasks such as provisioning, configuration and deployment of applications, ensures ML staff are able to support more data models into production.

3. Examples of AIOps in practice

3.1 Elisa Automate

Challenge

Elisa is unusual in that it has developed its own AI-based tools and owns the vendor Polystar OSIX AB Group, a software and probe company that focuses on the telecoms market. Elisa is using its solutions internally and is actively reselling these solutions to other operators, notably T-Mobile (Czech Republic) and Slovak Telekom. Its AIOps tools have been developed to support operations, save energy, plan 5G networks, predict network alarms, automate first-line network operation centre (NOC) responses, automate capacity planning and network configuration. Elisa's AIOps tools help to contain operational costs and to support more network elements, while supporting the delivery of increasingly complex services and customer service improvements alongside an average 44GB per month data for each subscriber – the highest in the world.

Solution

Elisa Automate supports network processes with its AI-built automated solutions, Virtual NOC and Elisa SON. Its Virtual NOC solution provides configuration management for fulfilment processes, as well as fault management, network and service monitoring, and performance management for assurance processes. Elisa SON provides automations to help improve customer experiences through self-organising networks (SON) and it is built on Red Hat® OpenShift® scalable container technology to support closed-loop processes, which make up to 2000 daily network changes. A further 120 use cases using ML and AI are deployed for network optimisation and operation as part of the Elisa solution.

Outcome

The Elisa NOC has evolved since 2010 towards full automation. Automations are delivered using ML-based functions that require minimal manual intervention by operational staff in their NOC and provide dynamically updated logic and rules as requirements change. Network incident alarms are captured, resolved by automations and new resolutions are updated.

In 2020, 90% of network incidents were resolved automatically and as many as 86% of these were proactively corrected before they became visible to customers. Elisa is working to increase the percentage of automated resolutions to reach a goal of 100%.[2]

3.2 SK Telecom

Challenge

SK Telecom is the leading mobile operator in South Korea. It wanted to maintain its leading market position by improving its customer services and rapidly rolling out 5G, but without increasing its operational costs. It wanted to use AI solutions in its pursuit of more-efficient operations and improved customer experiences.

[2] For more information, see Analysys Mason's *Elisa's journey to CI/CD*, available at: https://www.analysysmason.com/research/content/case-studies/elisa-ci-cd-rma14. Please also see Vanilla Plus's *Elisa offers network automation solution for operators*, available at: https://www.vanillaplus.com/2018/02/15/35810-elisa-offers-network-automation-solution-operators/.

Solution

AI and big data are central to the operational efficiency of SK Telecom. New capabilities have been included in its Telco Advanced Next-Generation OSS (TANGO), an in-house-developed AI-based OSS that covers inventory, planning, construction and operation. SK Telecom integrated TANGO with its Network Analyser (TiNA) solution, which is a unified network monitoring and operations tool used for data collection, storage and visualisation. Both TiNA and TANGO are based on the same big data and analytics tools and built on an in-house-developed Metatron platform and other open-source components.

Figure 3.1: SK Telecom's AI platform

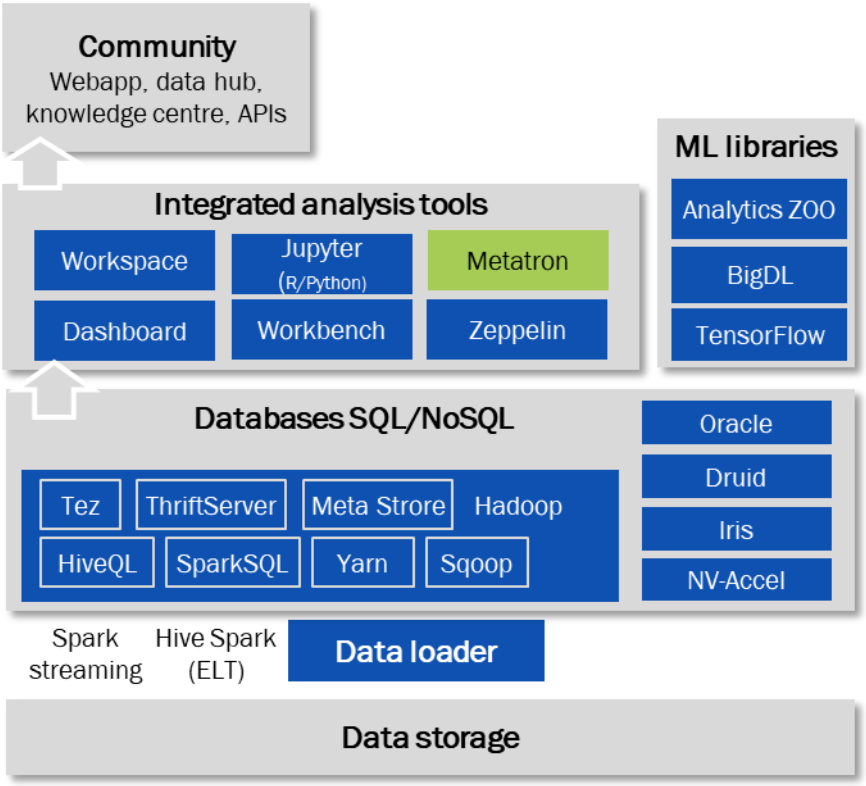

Source: SK Telecom

Outcome

Automation has helped SK Telecom to reduce operational costs, provide better customer experiences and optimise operational processes.

Conversational AI support for automated human interactions and NLP for text-based interactions has helped to reduce customer care costs and has improved SK Telecom's NPS. The TANGO development has improved planning, which has reduced capex, and its automations have reduced opex. The TANGO development resulted in a 40% cost reduction, as compared to the total cost of operating legacy OSS.[3]

[3] For more information, see Analysys Mason's *SK Telecom: big data and analytics tools and products*. Available at: https://www.analysysmason.com/research/content/case-studies/sk-telecom-analytics-rma14/.

3.3 EE (BT)

Challenge

EE is the UK's largest mobile network and is owned by BT. EE needed to manage and support key digital channel assets, which include the EE home page, customer portals for specific customer access to bills and services configurations, the community site and help pages and coverage checking applications. EE requires these services to run 24 hours a day, 7 days a week because subscribers are increasingly using these channels to access new services. These assets account for GBP10 billion of revenue for EE.

EE has shifted its IT infrastructure to the AWS cloud but still needs access to established on-premise systems. In times of acute usage, such as the launch of a new phone, data traffic peaks caused pages to slow or freeze. EE was unable to trace root causes fast enough to prevent poor customer experiences and loss of revenue as potential customers tried alternative providers. A reduction of 1 second in page-loading speed is widely reported to deliver a 7% impact on conversion rate within the industry, resulting in a significant loss of revenue for EE.

Solution

EE adopted an AIOps approach to resolve its performance and outage issues. It used a software tool that scans through the log data associated with all servers to support its digital channel and to add trouble ticket data. It deployed AI/ML-based AIOps tools from Dynatrace to correlate outages and performance issues with log file information, enabling it to predict performance issues before customers experienced them.

The Dynatrace system works by reducing delays to page loading, which is considered to be the critical KPI. It automatically works through large log data sets to build a predictive model to understand when the KPI is affected and by what. The use of ML means that algorithms are able to work effectively out-of-the-box and can also be refined over time because they are exposed to more training data.

The AIOps approach detects components and dependencies for each service to create an application architecture. Automatic baselining of data to understand 'normal' behaviour reduces set-up times and effort and provides a clear indication of abnormal behaviour for monitoring purposes. By combining the application architecture and monitoring, RCA can be used to resolve problems swiftly. This self-learning approach enables services and their underlying infrastructure to change, yet still be supported and managed without having to rewrite hard-coded tools.

Outcome

The AIOps solution reduced EE's digital incidents by 50% and reduced full outages of web pages by 50% compared with the previous year. This also represents a 50% year-on-year outage reduction. EE was able to be more proactive in finding the root cause of issues, leaving staff to focus on optimising customer experiences.[4]

3.4 T-Mobile (USA)

Challenge

T-Mobile (USA) needed a way to scale operations to accommodate the growing volume of network events and outages that it was managing. Cell tower outages were having a major impact on its customers, so T-Mobile needed a way to validate customer service complaints, perform diagnostics to determine problem areas and fix these issues quickly. With a daily rate of 500 000 alarms, the Tier-1 operator's analysts could not keep pace

[4] Dynatrace. *BT calls on Dynatrace to optimize digital experiences and conversions*. Available at: https://www.dynatrace.com/news/customer-stories/ee/.

when each alarm was taking 15 minutes to address. T-Mobile also wanted to automate a set of complex manual testing procedures on thousands of cell towers. By running these procedures consistently, T-Mobile could verify mobile connectivity for its millions of subscribers.

Solution

T-Mobile now polls the entire radio network every hour with its Resolve base AIOps system and compares the outages that it finds with dedicated tickets and events. It then automatically opens these tickets and generates validated events to the NOC for follow-up. T-Mobile also sends complete coverage and outage reports by geography to customer services, which gives T-Mobile's front-line staff visibility into real-time network health and performance.

Outcome

T-Mobile has saved 40 000 person hours annually with automated hourly testing of an entire radio network, automated incident resolution and production of daily network health reports.

T-Mobile has also improved visibility and control. Outages are detected quickly and fixed before most customers even notice; standard operational activities are automatically executed, and events are removed from the event stream after they are resolved.

Service delivery is faster. In 6 months of use, the T-Mobile team used automated systems to deliver 1100 port activations and execute 4500+ run-book automations per month without manual intervention. These automated systems work with existing systems and provide access to insights via APIs.[5]

3.5 Türk Telekom

Challenge

Türk Telekom is supporting growth in mobile gaming, as well as facilitating increased streaming of HD and 4K video and the emergence of 5G, all of which cause delays in services while demand is rising. Excessive use of network resources can affect other services, so careful monitoring of network traffic is needed. Established systems that rely on rules and thresholds being established manually were unable to scale as new services were rolled out.

Preventing excessive use of its resources, whether malicious or not, calls for vigilant awareness of network traffic. Network performance monitoring requires the ability to recognise anomalies as early as possible and to act to prevent more severe outages.

Solution

Türk Telekom used an in-house approach to collect resource performance, service performance and fault management data for analysis in near-real-time and then used it to detect anomalies in service quality to provide support for teams in resolving the issues. This required the building of a new data platform, converting the stream of data and reducing the 'noise' in the data to allow for key signals to be seen by operations staff.

[5] Resolve. *T-Mobile Saves 40K Man Hours Annually by Automating and Accelerating Network Testing*. Available at: https://resolve.io/customers/t-mobile-network-automation-case-study.

Türk Telekom implemented anomaly detection scenarios that fall into three main categories: round trip time (RTT), network traffic and alarm duration. All three categories provide enhanced visibility into network performance and opportunities for proactive network maintenance.

- **RTT.** Türk Telekom analysed approximately 1GB of RTT data over an 8-month period, sharing a total of 1105 detected anomalies with network operations, which led to improvements of up to 9% in fault reductions.

- **Network traffic.** The lead department analysed approximately 700 million rows of network traffic data, detecting 1800 anomalies, which it shared with respective operations groups to improve customer experience.

- **Alarm duration.** Türk Telekom focused on the incident resolution time for its MPLS network. At the start of the project, anomalies were reported on 88 different devices. With constant monitoring and proactive measures taken by net ops teams, this fell to 34 devices.

Outcome

Türk Telekom's data collection and sharing has driven operational improvements. The AI/ML-based solutions have given net ops teams more consistent and accurate network visibility. Türk Telekom plans to extend its AI/ML approach to include predictive maintenance and to assist in other AI use cases.

- The company realised reductions in RTT delays of 8%, 5% and 9%, respectively, in peering, international and mobile links.

- 1800 anomalies in network traffic were detected, leading to improved customer experience and monetisation opportunities.

- The number of MPLS network devices reporting service tunnel convergence anomalies was reduced by over 60%.[6]

3.6 Turkcell

Challenge

In 2018, Turkcell identified four areas of innovation as having a significant positive impact on their markets: data science, IoT, blockchain and AI. To deliver these services, it was essential to remove the constraints associated with Turkcell's (then) monolithic systems and processes had on the development of new solutions. The company's use of batch-based application software releases based on dedicated hardware for each project was considered to be a significant barrier to innovation. ML and analytics processes were being run on older, non-GPU-based hardware without optimised ML hardware acceleration. This resulted in slow data analysis and greater costs because more servers were needed.

Solution

Turkcell[7] made the strategic decision to build its own AI/ML solution based on the Red Hat OpenShift container platform. The use of Red Hat OpenShift enabled an integrated DevOps process to be established that incorporated AI data models into the development process, supporting more-frequent updates that were typically

[6] TM Forum (May 2020), *Turk Telekom deploys AI and ML in network operations*. Available at: https://inform.tmforum.org/casestudy/turk-telekom-deploys-ai-and-ml-in-network-operations/.

[7] Turkcell Democratizes Data Science and Drives Artificial Intelligence Innovation with Red Hat OpenShift

https://www.redhat.com/en/about/press-releases/turkcell-democratizes-data-science-and-drives-artificial-intelligence-innovation-red-hat-openshift?sc_cid=701f2000000txokAAA&utm_source=bambu&utm_medium=social&utm_campaign=abm

being used for Java-based projects. By using open-source-based components, Turkcell was able to reduce costs when compared to using alternative commercial software and was able to draw on the large community of developers for support. Nvidia, the global computer systems design services company, worked closely with Red Hat to help provide hardware acceleration on the NVIDA GPU chipsets to speed up analytics modelling and inference in run-time applications.

Turkcell's use of the Red Hat OpenShift platform accelerated development work on a number of its projects, including:

- the creation of the Foyer platform for the validation of documents to help reduce fraud
- chatbots, for customer interactions using speech-to-text in Turkish
- text-to-speech conversion, to deliver commercial audio and smart voice assistant services, as well as to support a visual application that helps children with autism to better communicate by tapping images that are then converted into speech
- social media to support tools that recommend images for users that get the most positive responses
- a community charity initiative that used AI tools to measure smiles in a "spread the smile campaign," which triggered food donations to an animal shelter.

Outcome

Turkcell reports that it is able to get AI-based applications to market twice as fast with its new solution, cutting development costs for AI by 70%. The combination of using containers running on Red Hat OpenShift and NVIDA GPUs has reduced the number of servers needed for AI-based workloads, and this has led to a saving of 50–70% on hardware costs.

Turkcell has built new services based on its data science platform. Employees can start a server with a GPU-enabled notebook, without direct support from the IT group. Self-service approaches also enable resources to be scaled up and down to support each AI project.

4. Summary

4.1 Strategic planning assumptions

By 2023, 40% of DevOps teams in telecoms will augment application and infrastructure monitoring tools using AIOps capabilities. AIOps platforms will address the need for operational transformation by combining big data, ML and DevOps functionality to analyse the data generated by networks, OSS/BSS and devices.

Several classes of vendor have emerged to meet operators' differing requirements.

- Vendors of general-purpose AI tools that have applied their solutions across industries include Alteryx, Microsoft Power BI and Splunk. These tools have the advantage of bringing a broad set of technology-based functions. However, they lack the out-of-the-box capabilities that telecoms-focused tools can provide. Consequently, the use of such tools requires telecoms operators to have access to resources that bring telecoms-specific expertise in data analytics and management.

- Vendors of major network hardware that have industry-specific solutions are somewhat compromised in their approach because they have to protect business logic held in their solutions and applications. Examples include Huawei SmartCare and AUTIN solutions. In many cases, these vendors have opted to

enrich their products by embedding AI capabilities for ease of use and to improve the time to market for their customers.

- A growing number of domain-specific ISVs (not network vendors) have developed applications that support AIOps use cases that have been optimised for the telecoms sector. Examples include Guavus's Guavus-IQ tool and Moogsoft's AIOps. These vendors provide many new choices for applications that focus specifically on the needs of telecoms operators.

The most-prevalent use cases focus on network operational issues around fault and performance management and associated actions to prevent or reduce their impact. These include functions based on event management, incident management and problem management, which also support customer services operations.

Future use cases include adding automations to extend service assurance to support applications, as well as the automated resolution of issues to create full lights-out operations. In addition, 5G roll-outs will lead to a strong demand for AIOps for network provisioning and to support new service types and capabilities.

5. Red Hat

Red Hat® solutions are built on open-source software that benefits from the collective skills and technology of developer communities, but also provides commercial enterprises and digital service providers with software solutions that offer the assurance and quality needed to run their businesses.

Red Hat delivers foundational open-source data and container-based technologies that provide a hybrid cloud foundation for developing, deploying, and lifecycle managing AI-powered intelligent apps that have usage for AIOps and also MLOps, supporting workflows, components that are used in the efficient and reliable delivery of AI/ML solutions. Red Hat supports AIOps solutions indirectly through open-source communities as a highly active participant, as well as directly through its commercial portfolio of solutions that can be purchased by enterprise customers, including telecoms operators or partners that supply solutions to operators. The Red Hat portfolio supports operators' various AI deployment strategies, from integrating third-party AI applications to developing AI-based applications in-house or leveraging managed AI services delivered on public cloud infrastructure.

Service providers that are interested in developing a company-wide, comprehensive, scalable AI strategy can join a new open-source project, the Enterprise Neurosystem Framework. Red Hat and industry leaders created an open-source community to launch several open-source projects that define the Enterprise Neurosystem Framework, which connects AI/ML functions and services across the whole enterprise. This overarching intelligence framework aims to collect and correlate data and data insights from all parts of a company's operations, autonomously providing load balancing for lower-level functions. and adjusting the corporate 'system' in real time where needed. The long-established OpenInfra Labs AIOps initiative is used as the governance framework and is already used in the OpenStack Foundation. Partners for the Enterprise Neurosystem Framework already include América Móvil, Equinix, Ericsson AI, EY AI Engineering, Intel, Kove, Lambda Labs, Nvidia, PerceptiLabs, Red Hat and Verizon Media.

The initial work for the Enterprise Neurosystem Framework is focused on monitoring hardware and software applications, with work to be done at a later stage on using analytics to predict failures in high-availability systems, anomaly detections and in monitoring the impact on systems, devices and the network of software upgrades.

In most AI projects, the immediate requirements come from data scientists, who are looking for an easy-to-use platform and environment to develop and test their AI applications and algorithms. Red Hat OpenShift® Data Science is a managed cloud service for data scientists and developers. The service provides an environment in which new AI models can be created, trained, tested and deployed in the public cloud.

Red Hat OpenShift Data Science offers an on-demand subset of pre-integrated software from the Open Data Hub functional architecture. This architecture is fully supported by Red Hat and currently runs on the AWS cloud. The platform-as-a-service (PaaS) approach simplifies the procurement of software and provides IT infrastructure without impacting current on-premise or cloud-based operational processes. This allows data scientists to focus on building data models rather than on IT administration. The on-demand nature of the service lends itself to scaling computationally intensive temporary workflows that are used during model development and training, which can then be shifted to potentially more-efficient infrastructure once the models are in production and their computational and storage demands are predictable.

For service providers and ISVs that are developing AI-based solutions, deploying AI/ML in production is an iterative process that extends beyond simply creating AI/ML models.

Red Hat's technology portfolio, expertise, and strategic partnerships with key ISV and hardware acceleration partners enables it to support its customers' AI/ML projects. This includes building production-ready AI/ML environments, as well as services and training.

Red Hat OpenShift is a Kubernetes container platform that uses automated operations to manage hybrid cloud deployments. It provides support for on-demand compute resources, hardware acceleration and enables consistency across on-premise and cloud providers' processes. Red Hat's provisioning functions help AI/ML teams to access computer resources efficiently, and its collaboration features enable different teams to share their results. The Red Hat Application Services solution helps to integrate application development, delivery, integration, and automation functions and data integration services can be used to build data pipelines.

Red Hat's platform and storage products (including Red Hat Enterprise Linux®, Red Hat Virtualization, Red Hat OpenStack® Platform, Red Hat OpenShift Data Foundation and Red Hat Ceph® Storage) provide the core software-defined infrastructure for cloud based, on premises and hybrid operations on which the AIOps structure is delivered. The Red Hat certified partner ecosystem allows a choice of AI/ML and application development tools to be used in this architecture, which enables simple, automated deployment and life-cycle management.

A rising number of companies offer new AI-based applications to support service providers' business and network operations, so it is important for service providers to have an open container platform. Red Hat's overall portfolio utilizes open-source software that is built on community developed solutions and partnerships from the Red Hat Marketplace to deliver a growing set of applications. These include applications such as Avanseus Cognitive Assistant, which predicts potential network failures and provides fixes.

Figure 4.1: Red Hat's open hybrid cloud platform for AI/ML projects

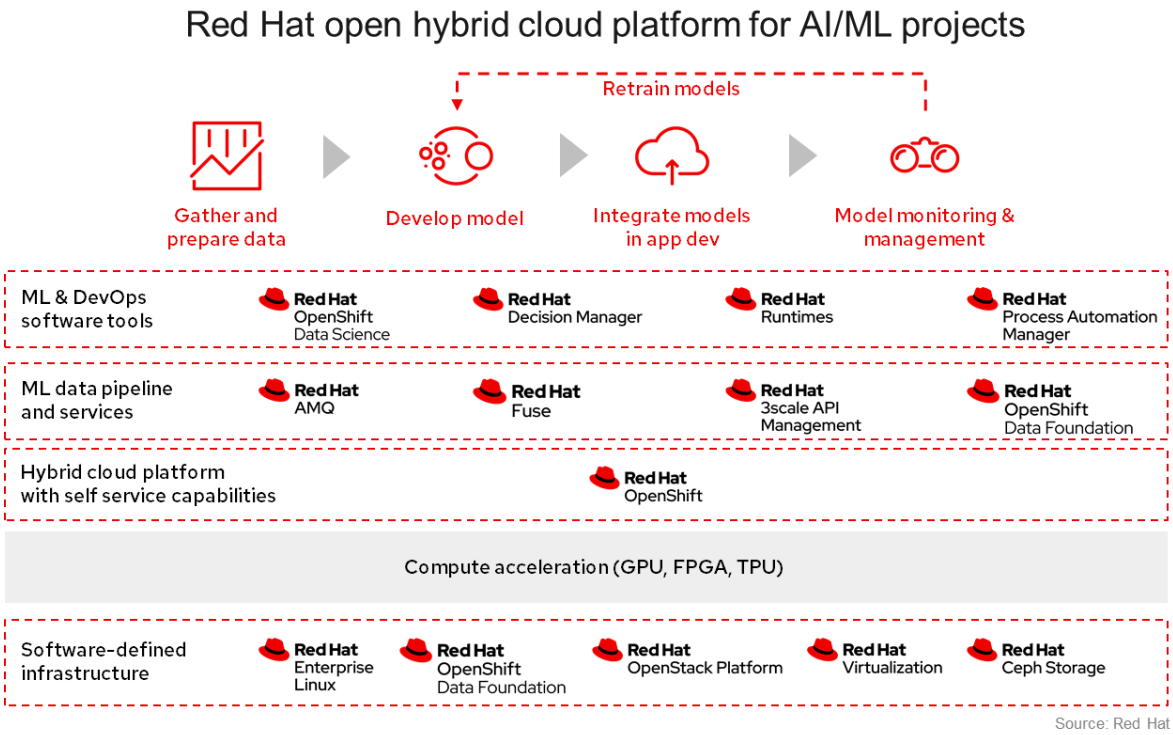

In order to implement AIOps, support is required from multiple vendors and solution creators. This includes support for ML libraries, prebuilt workflows to support telecoms processes, prebuilt interfaces for major data sources to support data collection and the support of APIs within telecoms operators' businesses to facilitate actions based on insights. An open framework is important for providing a neutral environment in which a common set of tools, processes and applications can be shared to allow AIOps solutions to be built and delivered on public, private and hybrid infrastructures.

6. Conclusions

The AIOps market is not well-defined, but the core assets of AI, big data and automation are familiar to CTOs and CIOs within telecom companies and are of high interest to them. AIOps is more likely to be adopted by operators that want to control their own processes and applications. They will have adopted DevOps solutions early and will be seeking automations based on AI and big data to help resolve their complex automation challenges.

The primary users of AIOps' industry-specific tools will be the same customers of point-based solutions, which use logic based on static and simple rules. However, AIOps is considered to be an enterprise-wide approach that incorporates and potentially retires established tools. Operators need to consider their abilities to build and support solutions based on AIOps and how they will continue to maintain solutions. This will also influence their choice of vendor and systems integrator to support their approach.

7. About the author

Justin van der Lande (Principal Analyst) leads the Data, AI and Development Platforms research programme, which is part of Analysys Mason's Telecoms Software and Networks research stream. He specialises in business intelligence and analytics tools, which are used in all telecoms business processes and systems. In addition, Justin provides technical expertise for Analysys Mason in consultancy and bespoke large-scale custom research projects. He has more than 20 years' experience in the communications industry in software development, marketing, and research. He has held senior positions at NCR/AT&T, Micromuse (IBM), Granite Systems (Telcordia) and at the TM Forum. Justin holds a BSc in Management Science and Computer Studies from the University of Wales.

This perspective was commissioned by Red Hat. Analysys Mason does not endorse any of the vendor's products or services.