SELECTOR

A Maturity Model for Network AIOps

Table of Contents

Executive Summary

AIOps, a term introduced by Gartner[1] to describe the use of AI in IT Operations, is ushering in a new era for operations within the enterprise. It allows CIOs to take definitive steps towards realizing a state in which operations are fully autonomous under the supervision of operations staff who can finally focus on business value creation instead of constantly fighting operational fires. However, getting to this state is a journey in the evolution of current practices and technologies. In this document, we present one such model that recognizes the realities of today's enterprises, laying out a four stage approach, going from the lowest level, 'Passive Ops' (highly reactive, manual stage), to the highest 'NoOps' (fully autonomous operations). This document highlights key network operations (NetOps) specific considerations. Where relevant, we will also call out opportunities where security operations (SecOps) and development operations (DevOps) can benefit from network AIOps outcomes.

The four stages in our model are:

1. **Passive Ops**: When operations are highly manual and reactive, with Ops teams spending the bulk of their time firefighting to prevent systems from melting down.

2. **Active Ops**: The graduation of operations wherein an Observability-based methodology allows for better insights into issues, using telemetry from across the network to make better informed holistic decisions.

3. **AIOps**: A state that can be achieved today, using AIOps solutions. At this maturity level, enterprises have started to realize the benefits of using AI to augment operations staff's analytical and decision-making abilities.

4. **NoOps:** A visionary state that at the highest point of maturity could be considered somewhat utopian. Getting here implies fully autonomous operations with humans only supervising AI to get the job done. Getting to this state will pre-req advancement in many technologies besides core AI itself. However, with advances in today's AI techniques combined with robotic process automation, digital workers and other business process automation/analytics capabilities, the industry is likely going to offer the ability for enterprises to get to an entry-level NoOps state in the near-term.

AIOps is a discipline of robust monitoring and observability, and we will discuss the relationship among the three. To conclude, we will look at a real enterprise AIOps maturity success story with Tracfone, an American pre-paid mobile phone provider.

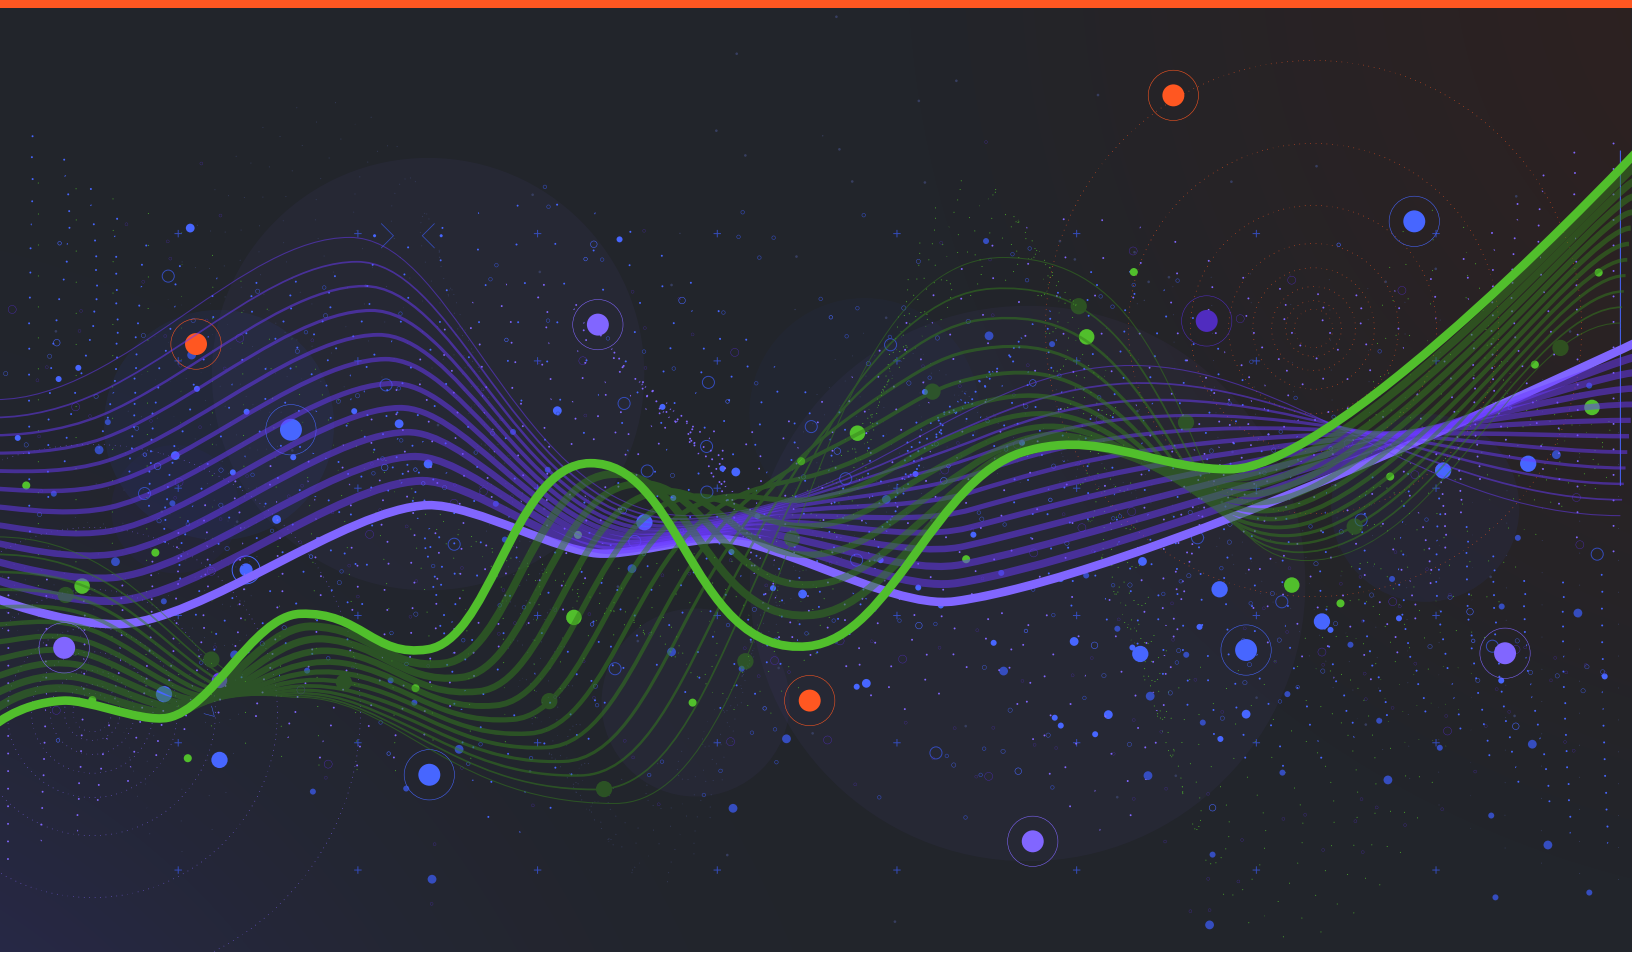

Introduction

Ask an enterprise IT operations (Ops) executive how much IT incidents cost the business annually, and you will hear numbers that could go into the millions of dollars per hour[2]. The current tools, techniques, and people that make up today's Ops functions are struggling to cope with the demands of the modern day enterprise. Blame the rapid advancement in technology across multiple disciplines since it has introduced a reality in which it has become humanly impossible to cope with the volume, variety, and velocity of data (aka 'big data') being produced by today's environments. An estimate by Statista[3] shows data volumes worldwide growing by 60x from 2 Zettabytes (Zettabyte = 1021 bytes, or one trillion Gigabytes) in 2010, to 120 Zettabytes in 2023. And this is only going to get worse, with a 50%+ growth estimate in the next 2 years (181 Zettabytes by 2025). This should come as no surprise though if one were to consider the proliferation of handheld devices, IoT devices, and digital business models that are pumping out data at growing volumes (with higher velocity and even more variety) year-over-year.

Along with the progress in technology, we would expect to see an evolution of our tools, techniques, and workforce that help manage and operate these modern systems. This is not just essential, but also inevitable, and leads to these enterprise IT goals:

1. **Reliability:** Ensure enterprise applications and infrastructure performance meet the expected service levels in complex multi hybrid-cloud environments.
2. **Productivity:** Minimize the noise, toil, and alert fatigue that IT operations teams have to deal with, allowing them to make room for business value creation activities.
3. **Data-driven:** Adopt a discipline of data-driven, actionable insights in all aspects of the business, including IT operations.
4. **Modernized:** Adopt advanced enterprise architectures, analytics, and AI to enhance decision-making, tackle 'tools sprawl' and enhance market competitiveness. Seamlessly integrate all forms of IT operations, including, network operations (NetOps), software development operations (DevOps), security operations (SecOps) and machine learning operations (MLOps).
5. **Automated:** Automate current manual and tedious processes, which will also help mitigate errors inadvertently introduced by human operators.

In this document, we will present a prescriptive approach, in the form of a maturity model, which the enterprise can adopt to evolve their IT operations posture in general. As a large chunk of IT operations, we will zoom into viewpoints around NetOps and highlight considerations for SecOps and DevOps where relevant, since there is opportunity to cross-pollinate across these disciplines. Before we dive deeper into the model, let's establish a common understanding of a few important concepts that will be used within the context of the model and the various stages of maturity therein.

What is Monitoring, Observability, and AIOps?

If one was to survey the dozens of vendors who offer monitoring, observability and AIOps products, whether purpose-built (for e.g., an industry, or domain such as networking) or general-purpose (for e.g., full stack, cross-industry), it would become obvious there are differing views that can easily confuse prospective enterprise buyers. While we will not be doing a deep comparison in this document, it is important we level-set on the context in which the three will be used within the maturity model. The TL;DR is this – monitoring is the foundation on which observability and AIOps are built and observability/AIOps complement each other in ways that make them both as essential to the enterprise as monitoring. Confused? Try this visual below:

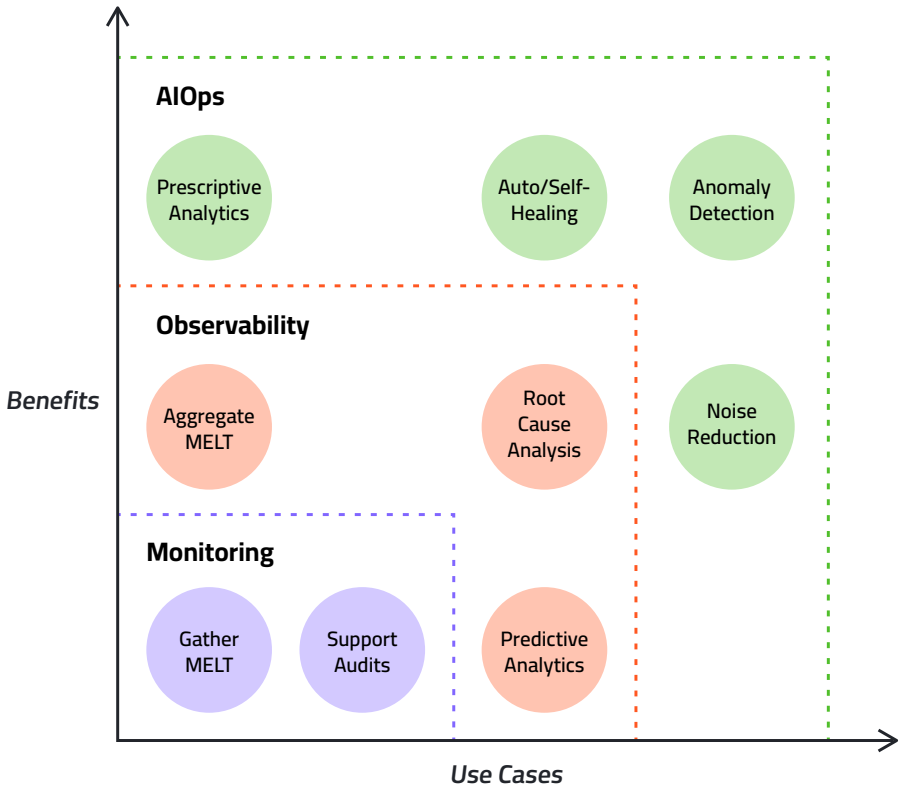

Figure 1: The relationship between monitoring, observability and AIOps

Monitoring

Simply put, the act of 'monitoring' an entity (such as a system, application, equipment, component, process, digital twin, 'bot') is about gathering data that describes a quantitative/qualitative state of the entity at periodic intervals, key actions being taken by the entity (including interactions with other entities), snapshots of its state when something goes wrong, and then offering this data in the form of metrics, events, logs or traces (aka 'MELT' data). Monitoring is therefore the foundation of any 'Ops' discipline. Without it, there is no data to work with, to review, analyze, correlate, and infer insights from. A solid framework of monitoring is a pre-requisite to any initiatives around observability or AIOps.

Observability

Taking it several steps further, observability builds on monitoring by using MELT data to drive analytical insights that help operators understand, describe, and predict behavior of the systems being monitored. Through observability comes the understanding of what's happening, why something has happened (descriptive analytics) and, in many cases, also what could likely happen next (predictive analytics). The success of observability depends on the creation of robust, reliable telemetry pipelines through collection, aggregation, filtering (e.g., de-duplication) and enrichment (e.g., correlation) of telemetry across all systems being monitored.

Observability becomes ever more relevant in today's highly distributed, multi-hybrid cloud world, where cloud native architectures have given rise to highly containerized applications based on the micro-services architectural pattern. Modern-day applications are typically micro-services-based (compared to the older monolithic architecture). Observability solutions are needed to ensure data is being collected and aggregated — not just from every distributed node, host, network component or VM, but also from every micro-service that is typically running within a Kubernetes container platform, which itself produces very large amounts of MELT data.

It would be prudent to note here that one might come across the concept of an observability 'tool' vs. an observability 'platform' defined by several vendors in the market[4], with the latter being considered a broader set of integrated capabilities (satisfying key observability tenets) that allows for holistic insights that are system-wide. This document does not get into the nuances, and we will therefore abstract both, using the term 'solutions' instead.

AIOps

The term was first defined by Gartner[1] to describe the use of AI in IT operations. As a newer concept in the Ops arena, AIOps is about adding the ever-growing promise and power of artificial intelligence, as a broad discipline (not just ML), to enterprise operations use cases. Just like any other AI application, AIOps needs good data and that is where monitoring and observability come in. While one will find many variants in the manifestation of the key AIOps tenets in commercially available solutions, most would agree that AIOps coupled with observability is the way to go. And of course, this is built on top of a strong foundation of monitoring.

There is an important consideration for any AIOps solution to be truly enterprise ready – it must ingest various types of system data (e.g., alerts, events, telemetry, configuration files), work with different data formats (e.g., time series streams) and be readily customizable to capture anything new (e.g., custom formats used by home-grown monitoring tools or legacy systems within the enterprise). This would be needed out of the box or with minimal effort on the part of the enterprise (or vendor) to ensure quick adoption. In some senses, this is a 'turnkey' data ingest and processing capability that would ensure quick time to value for any AIOps solution within the enterprise.

Here are several enterprise use cases where the AI in AIOps kicks in to differentiate itself from what observability can achieve:

- **Filtering out Noise, Event Correlation, Anomaly Detection**: Modern-day enterprises are dealing with a deluge of telemetry, which AI-based automation is well-suited to tackle. The application of AI ranges from data processing (e.g., filtering out noisy data, de-duplicating events) to event correlation and anomaly detection, where it produces superior outcomes.
- **Prescriptive Analytics**: The ability to prescribe solutions to a problem using statistical and machine learning techniques that suggests an ideal course of remediation action based on likely outcomes of following one of many potential actions that could be taken.
- **Closed-loop Automation**: Built on the foundation of prescriptive analytics and ML models, closed-loop automation (also referred to as 'self-healing') is about trusting the AI to make the right choice of remediation actions possible and initiating an automated workflow to execute it.
- **Natural Language Interactions**: Advances in AI with large language models (LLM) have brought in modern, automated approaches to tackle the current, manually intensive interactions that operators have been traditionally having with Ops management solutions. A centralized LLM-based natural language conversational interface alleviates the inherent inefficiencies and inaccuracies by allowing operators to use natural language queries instead. It frees up operators from needing to iteratively gather contextually relevant data through several interfaces, pouring through a very large number of metrics, running several reports, and browsing through many dashboards to find what they need, only to realize they needed more. An AI chatbot will instead accept an operator's queries in a natural language format, offload the manual analysis to an automated backend that uses AI techniques to respond with contextually relevant data, and follow up with the operator in a natural language style.

Today, vendors and enterprises looking to explore LLMs have a myriad of open LLMs to choose from[5].

And many more, which we will discuss later in this document. It should therefore come as no surprise that AIOps solutions are, by definition, designed to integrate seamlessly with well-known observability solutions out in the market, as well as some of the popular open source observability frameworks and enterprise collaboration tools. This gives AIOps solutions the ability to become a single pane of glass and the one source of truth for all Ops tasks within the enterprise.

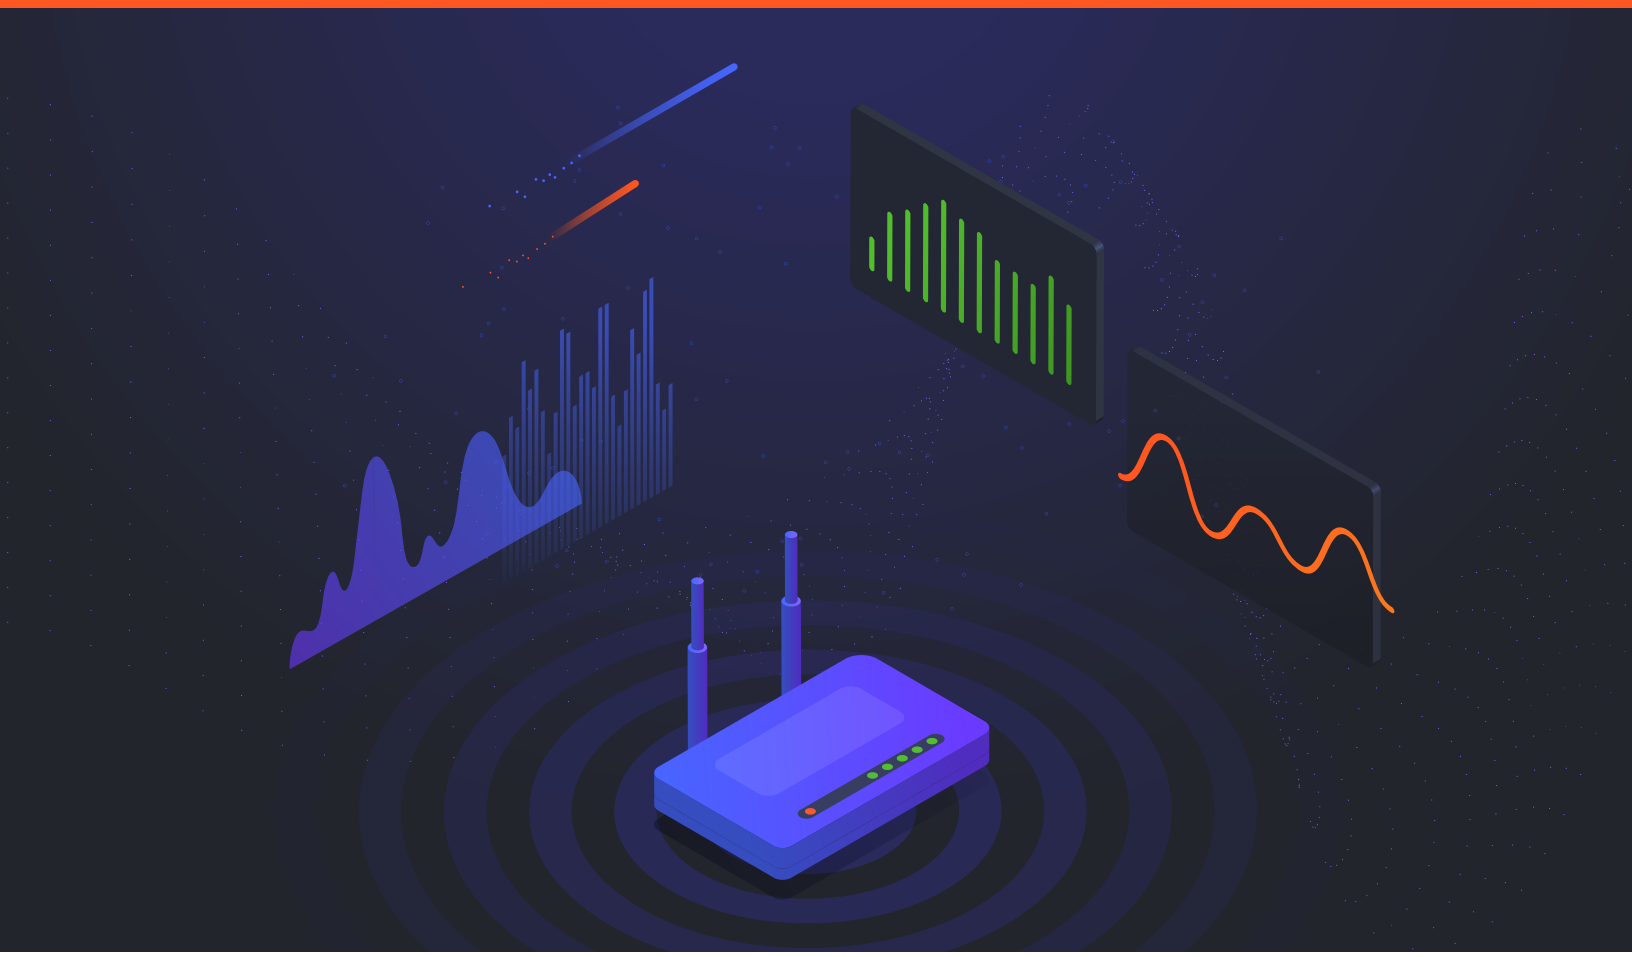

Network Observability and Network AIOps

Observability for networks follows the same fundamental tenets for general observability across other disciplines, be it DevOps, DataOps, MLOps, or SecOps. The fundamentals of telemetry, data ingest and aggregation, data management and analytics, and alerts as well as the end-user views (dashboards, graphs, etc.), among others, remain the same. One key difference is that there are new types of data producers in the form of network specific protocols and equipment, which must now also be supported. This includes BGP, gNMI, HAProxy, Netbox, SNMP, Zabbix, NetFlow, IPFIX, among many others, and those from well-known vendors of networking infrastructure.

So… what then is network AIOps? And how does network AIOps differ from general IT AIOps? The answer lies in the data, the AI, and the domain specificity. The heart of any good network

AIOps solution is the AI, which, just like any other domain-specific AI application, needs to be driven by AI techniques and ML models that are purpose built for network operations use cases. The ML models would be trained with data ('features') that are domain (networking) specific, as discussed above. In addition, there are unique characteristics within a network that need to be taken into consideration such as topology changes (planned or otherwise) which affect underlying behavior and performance. It is also important for network issues to be intelligently correlated with application performance management metrics so that operations staff can either confirm that the network is the reason for application performance degradation or to prove conclusively, with data, that the network isn't to blame.

With this foundational understanding in mind, let's now shift gears and dive into the maturity model for network AIOps. Given the discussion thus far, it should be no surprise that a network AIOps maturity model will align closely with a general AIOps maturity model. There will be network operations specific characteristics which we will point out where appropriate.

Maturity Model for Network AIOps

The maturity model's goal is to help enterprises advance meaningfully in their network operations management modernization journey. That will, in turn, drive better business outcomes. It is designed to help enhance the productivity, efficiency, and effectiveness of Ops staff and achieve better ROI from the investment in Ops management, observability, and AIOps solutions, while minimizing Ops staff 'toil' ('busy work' that takes up an inordinate amount of productivity bandwidth, preventing operations staff from adding greater value to the business and its customers). Along the way, it will tackle many of the enterprise Ops challenges we discussed at the beginning of this document. When reviewing the model, keep in mind that change takes time and dedication, and it is therefore likely that an enterprise could be partly at one level of maturity and partly in the next. The maturity model should be viewed as a journey, staying true to the best practices of AIOps adoption which call for a gradual, phased deployment plan.

At one glance, the figure below will help you understand the 4 levels of network AIOps maturity we will be discussing. The maturity levels are qualified in terms of these key dimensions:

1. **Underlying Theme**: The essence of a particular state of maturity.
2. **Strategic Drivers**: Key enterprise goal/s at the level.
3. **NetOps use cases**: Network operations specific use cases that would be served in the maturity state under consideration. These are often applicable to broader Ops as well.
4. **Dependencies**: Key capabilities that success in a level would depend on (pre-requisites).
5. **Features**: Features typically expected at the particular level of maturity. This includes different types of technologies that are currently in use or would be needed for enterprise operations management.
6. **Data management strategy**: A critical part of any implementation across the maturity model, the data management strategy helps support the intended analytical outcomes. This recognizes the fact that behind any AI implementation is well-managed, trustworthy data.
7. **What's needed to graduate**: Key considerations if one wants to graduate to the next maturity level.

	Passive Ops (In Use Today)	Active Ops (In Use Today)	AIOps (Available Today)	NoOps (Visionary)
Underlying Theme	▪ Reative, manual, very basic automation using available telemetry ▪ Siloed Ops teams; very limited collaboration ▪ Low ops staff efficiency and productivity	▪ Observability helps tackle problems proactively ▪ Some automation (e.g., with ops management tools) ▪ Siloed ops teams; better collaboration ▪ Medium ops staff efficiency and productivity	▪ AI augmented operations using AIOps solutions ▪ Medium to high ops automation ▪ Strong ops team collaboration ▪ High ops staff efficiency and productivity	▪ Autonomous AI driven operations ▪ Fully automated, supervised by humans ▪ Minimal ops team required ▪ Ops staff re-focused to value creation tasks
Strategic Drivers	Keep the lights on and prevent system meltdown	Modernize operations to proactive mode	Lower costs; enhance ops efficiency and effectiveness	Digital self-service, fully automated operations
NetOps Use Cases	▪ Diagnose alerts and failures that have already happened (manually) ▪ Audit log files and alerts (manually) to assess risk of network failure that could happen ▪ Remediation to problem (manual or low-key automation assisted) without due consideration for bigger systemic issues that could be lurking ▪ NetOps collaboration with SecOps, DevOps, MLOps (limited and inefficient)	▪ Proactive network failure alerting ▪ Detect hazardous network conditions ▪ Early exposure to unknown "unknowns" ▪ Collect and analyze network performance data to meet SLOs and SLAs ▪ Reduce noise and alert fatigue ▪ Event correlation across the network for deeper insights ▪ Detect some network topology changes ▪ SecOps: proactive threat detection reusing NetOps telemetry	▪ Noise reduction, event correlation, anomaly and outlier detection (e.g., BGP and port instability, interface errors) with greater accuracy and speed of root cause analysis ▪ Reduce MTTD and MTTR ▪ AI-guided remediation suggestions ▪ Predictive analytics (e.g., forecast network capacity, performance and hardware issues) ▪ SecOps: Threat detection and incident response ▪ DevOps: Standardized environment provisioning ▪ Optimize network resource allocation	▪ Fully self-diagnosing with closed-loop automation ▪ Tightly integrated NetOps with SecOps, DevOps, and other relevant Ops functions within the enterprise
Dependencies	Metrics, events, logs, traces available natively from applications and infrastructure in use	Robust, reliable telemetry pipelines: collect, aggregate, filter, enrich system-wide telemetry	Fully integrated data ingest and processing capabilities for rapid enterprise adoption	Advancements in AI technology and infrastructure to run it
Features	▪ Reliance on telemetry made available natively by applications and infrastructure in use ▪ Custom scripts and runbooks for common operational tasks ▪ Manual and tedious root cause analysis processes ▪ Manually apply necessary remediation (or use custom-built, low-key automation)	▪ Telemetry correlation ▪ Use of manually defined static thresholds, policies, and rules to effect proactive outcomes ▪ Emphasis on end user experience	▪ AI-driven advanced analytics (predict and prescribe) ▪ Integrates readily with popular Ops tools in multi-hybrid cloud ecosystems ▪ Seamless integration with popular team collaboration tools ▪ Customized for the domain (e.g., NetOps) ▪ Single pane of glass and source of truth	▪ 100% AI-led Ops management ▪ Self-learning AI, traps unknowns and adds the learning to the collective intelligence ▪ Human operators supervise the AI, ensuring high-quality outcomes ▪ Standardized protocols of information exchange across all forms of Ops within the enterprise ▪ Automation enablers: workflows, domain specific RPA bots, digital twins, digital workers ▪ AI governance
Data Management Strategy	▪ Fragmented, custom data formats and stores ▪ Designed primarily for log storage and manual auditing	▪ Data silos (custom for each Ops tool) ▪ Archives for historical/audit purposes	▪ Big data management strategy to enable analytics, ML and AI usage patterns ▪ Active metadata management	▪ Autonomous big data management (AI managed)
Graduation Criteria	▪ Comprehensive infrastructure visbility ▪ Executive sponsorship (secure budgets and establish aligned goals) ▪ Revamp or replace existing processes with a tools-based approach	▪ Executive sponsorship (secure budgets and establish business-aligned goals) ▪ Refactor processes and establish new ones around AI automation and augmentation ▪ Small, incremental step adoption plan	▪ Advancement in AI ▪ Acquire the underlying infrastructure for hungry AI solutions ▪ Standardized protocols for information and metadata exchange across all enterprise Ops	▪ N/A

Figure 2: The four stages of enterprise Network AIOps maturity

Level 1: Passive Ops

This is the most basic stage in our model and quite commonly found within enterprises today. At this level:

- Ops teams are tasked primarily with 'keeping the lights on'.
- Ops staff fulfill their duties manually, reacting to system generated alerts using telemetry acquired with a low rate of sampling across applications and infrastructure in use. Key metrics are pulled with insufficient granularity (e.g., SNMP, once every 15 minutes), which handicaps the ability to pinpoint when and where an issue started and ended.
- Telemetry is sometimes inadvertently or intentionally ignored from key sources (e.g., due to lack of proper network topology views and timely change management; in some cases, there is no choice but to ignore telemetry from older equipment due to security policies that ban unencrypted network traffic).
- The absence of end-to-end instrumentation further compromises the depth of telemetry.
- Automation is also very low key, using basic techniques like scripts and runbooks to help with diagnosis and remediation tasks.

In summary, the lack of investment in advanced technology, disciplines, and Ops management tools becomes evident at this level.

Example Use Cases

Sample network operations use cases at this level include:

1. Collecting and diagnosing alerts and failures, using the telemetry available, to determine a root cause (manually).
2. Auditing log files and alerts periodically to assess the exposure risk of network failure that has already happened or is likely to happen (manually, with reliance on operations staff experience to foresee issues based on current conditions).
3. Remediation actions taken by operations staff that are based on available information and fix narrow problems, often missing bigger systemic issues that are lurking out there.
4. NetOps, SecOps, and DevOps collaboration that is tedious and fraught with inefficiencies due to the lack of system-wide holistic views.

The level of toil as well as noise that the Ops teams must deal with puts them into constant firefighting mode. Operations teams are working in silos, with limited collaboration. Data management at this level is mostly about storing the logs generated within the network and supports applications and infrastructure such that manual auditing can be carried out effectively. Data is highly fragmented and there is no easy way to have it all be accessed centrally for holistic insights.

Graduation Criteria

Key considerations to graduate to the next level of maturity include:

1. Adopting a discipline of comprehensive visibility into applications and infrastructure across the IT landscape, most importantly starting with those considered to be critical path. One way to achieve this is to bring in end-to-end observability, which will require a dedicated focus from architects, operations teams, and developers.
2. Executive sponsorship to provide the necessary budgets and establish business aligned goals.
3. The necessity of leaving behind existing manual operations to make way for new processes that work in a tools-led Ops landscape.

Level 2: Active Ops

This second state of maturity comes with some significant benefits in cost reduction and moves enterprise IT Operations towards a proactive strategy. It is characterized by:

- The use of Ops management and observability solutions which provide better insights than the simpler monitoring tools we discussed in level 1.
- Establishment of telemetry pipelines which collect metrics, logs, events, and traces across the entire system. These pipelines also aggregate, filter out the noise, and enrich the data. This is typically done using an observability solution. It would be important to note that the extent of filtering (e.g., de-duplication) and enrichment (e.g., correlation) is basic compared to what one can expect to see in an AIOps solution at level 3 in our maturity model.
- Ops staff leverage the built-in analytics capabilities within the observability solution to produce quicker, better, holistic diagnostic and remediation outcomes (than level 1).
- Use of static threshold-based alerts through manually defined policies and rules that rely on pre-tagged telemetry. This becomes untenable to maintain though without machine learning (as in level 3) due to the sheer volume of telemetry and alerts being generated by a typical enterprise.
- Better organized Ops teams who can collaborate more meaningfully (than level 1) due to a holistic view of the applications and infrastructure being monitored through observability solutions. This is facilitated by emphasizing a superior end-user experience that allows for ease-of-use and cross-team collaboration.

The market seems to be polarized by vendors of AIOps solutions who are born into holistic AI and those who claim observability plus advanced statistical/ML techniques make them ready to serve enterprise AIOps outcomes. As a reminder, ML is one of many components of AI. Enterprises should take the time to understand the core tenets of AIOps and evaluate their solution choices carefully when trying to move from level 2 to level 3. It would be beneficial to look 'under the hood' to fully comprehend what a solution is truly offering – statistical techniques vs. basic/advanced ML models vs. true AI (which encompasses ML). Can the solution truly scale with the needs of the enterprise? Will it allow them to eventually reach level 4 maturity (NoOps) where AI is pervasive across all forms of IT operations? Also, as discussed earlier, AIOps outcomes are always at their best when complementing a discipline of observability within the enterprise.

Example Use Cases

Network Operations use cases at this level of maturity could include:

1. Alerting when conditions are created that would typically lead to network failure.
2. Collecting and analyzing network performance data to ensure SLOs and SLAs are met .
3. Filtering out the noise in telemetry and reducing alert fatigue.
4. Getting alerted to potential changes in network topology.
5. SecOps benefits: proactively detecting threats using network operations telemetry to discover potential network intrusion conditions.

Data management at this level of maturity is driven by silos of data in use across each Ops management or observability solution, with the latter often helping consolidate data into one repository. However, the problem of data silos persists due to the existence of 'tool sprawl' (many tools in use across multiple Ops teams and functions), and enterprises must maintain archives for historical/audit use cases.

Graduation Criteria

Key considerations to graduate to the next level of maturity include:

1. Executive sponsorship to provide the necessary budgets and establish business aligned goals.
2. Adoption of an enterprise AI discipline, which will require a dedicated focus from architects, operations teams, and developers.
3. Refactoring existing Ops processes or establishing new processes that are geared towards AI-based automation and staff augmentation brought on by using AIOps solutions.
4. Making measured, incremental progress to eventually elevate AIOps solutions to the center of operations functions across the enterprise.
5. Ensuring AIOps solutions can work seamlessly with existing monitoring and observability solutions to avoid expensive migrations unless truly necessary.

Level 3: AIOps

At this state, the enterprise has started to adopt AI for its IT operations. This is a journey and getting started is key. As discussed earlier, vendors of AIOps solutions can make it easy to adopt their solutions by providing a well-integrated set of data ingest and pre-processing capabilities that can be plugged into any enterprise ecosystem with any type of data format or type, using features that come out of the box or through a framework that requires only a small amount of customization effort. This key capability will determine the speed of progression of the enterprise through the AIOps maturity state.

At this level, we recognize the following defining characteristics:

- Operations teams would use AI-based augmentation for decision making. One might hear the term 'cognition' used in this context as well. What AIOps does essentially is it augments the cognitive abilities of the human operator by adding deeper and broader insights (through both descriptive and predictive analytics) than what an Ops management and observability solutions would typically provide.
- AI-based augmentation could also come in the form of recommendations for remediations to an issue under investigation (through prescriptive analytics), based on machine learning capabilities. This can help build automated incident response workflows that can be triggered manually, after review by an operator, or automatically for well-understood common problems.
- Telemetry continues to be collected as one would in the previous maturity stage and observability solutions play an important complementary role here, as do various monitoring tools. It is important to recognize that AIOps solutions will typically not replace existing tooling — they would complement them. This is even more relevant in a large, distributed enterprise where microservices-architecture-based services are hosted on Kubernetes container platforms in multi-hybrid cloud environments.
- AIOps solutions are designed to work with very large volumes of data that are being produced across the entire IT ecosystem at real-time streaming speeds, which therefore mandates advanced principles in management of big data and related metadata. This could manifest in the form of modern data lakehouse or traditional data lakes that are geared towards AI and ML usage patterns.

- Collaboration among Ops teams is significantly enhanced due to the centralized nature of AIOps, as the single pane of glass and source of truth for all things operations related across the enterprise. A good user experience is therefore essential in any AIOps solution, as is the ability to integrate seamlessly with existing team collaboration tools.
- As discussed previously, for any AI to be relevant, it must be customized to the domain it serves. For example, in the network AIOps space, the AI must be trained to serve NetOps specific use cases, with network operations specific data, conditions, events, actions and various other 'features' that train ML models supporting the AI. Only then can a network AIOps solution do effective event correlation, data aggregation, data filtering, root cause analysis, remediation recommendations and even natural language user interactions for NetOps use cases.

At this point, it would be prudent to point out that AIOps has become a highly 'overloaded' term in the industry (borrowing the meaning of the term from object-oriented programming). AI forms the foundation of an AIOps solution, and a deeper dive into some of the commercially available AIOps offerings would likely reveal they are nothing but an 'AI façade' around existing Ops management or observability solutions. IT operations leaders who decide to get to this maturity state are looking towards AI to enhance their Ops team efficiency and effectiveness, thereby enhancing overall productivity. They want to give these teams the opportunity to start creating business value, rather than being in constant firefighting mode. It is therefore important for any Ops leader who has chosen to get to this level of maturity to do their due diligence when making their AIOps solution choice.

Example Use Cases

Network operations use cases this state serves could include, among others:

1. Enhancing network operations team productivity, reducing MTTD and MTTR.
2. Reducing the noise in the typical deluge of signals and alerts across the network to enhance the quality of the data being used to make diagnostic and remediation recommendations, leading to faster incident response.
3. Improved outcomes through the correlation of key events across the network.
4. Anomaly and outlier detection (e.g., BGP and port instability, interface errors, etc.).
5. Detecting unknown 'unknowns' using advanced machine learning models.

6. Predicting/forecasting network capacity and performance problems and hardware issues.
7. Optimizing network resource allocation based on suggestions from AI-guided insights into the correlation of application performance to network equipment and infrastructure.
8. Collaborating with SecOps teams: Enabling faster threat detection and incident response using network telemetry and insights generated by network AIOps.
9. Collaborating with DevOps teams: Automating the provisioning of environments with standardized configuration that include the supporting network infrastructure.

Graduation Criteria

Finally, let's look at what's needed to get beyond this stage of maturity. The next level brings in fully autonomous operations, which will need:

1. Advancement in key technologies (industry wide) across AI, digital twins, digital workers, and RPA as well as AI driven autonomous data management.
2. Acquisition of budgets to fund underlying infrastructure that would run resource hungry AI systems.
3. A visionary set of industry leaders who will help standardize protocols and metadata in use across enterprise Ops, whether NetOps, SecOps, MLOps, DevOps or, in general, any Ops.
4. Executive sponsorship and an aligned business integration plan to ensure this huge culture shift for most enterprises will have a chance at succeeding.

Level 4: NoOps

'NoOps', a term coined by Forrester[6], is the most advanced, 'visionary' state of maturity in our model. Businesses recognize the fact that many enterprise application and infrastructure problems are inadvertently human induced. Autonomous operations help tackle this challenge. At this stage of maturity, the IT operations of a business are fully or mostly run by AI systems that monitor, diagnose, and fix problems, under the supervision of human operators. A 'self-service' culture emerges within the enterprise wherein users of IT assets interact with an AI 'bot' to get routine tasks completed.

This level of maturity would rely on 2 key technologies to monitor, diagnose, and fix issues:

- **Robotic Process Automation** (aka 'RPA' – software that automates repetitive business tasks and processes, often referred to as 'RPA bots')[7] and
- **Digital Workers** (an artificially intelligent, software-based virtual employee that integrates various AI technologies with RPA bots to handle routine computing-based business processess)[8]

This is an aspirational state for many enterprises and AIOps vendors — their driving vision for the near future. Many market leaders are making and adopting rapid advances in AI, which can be coupled with automation technologies like RPA and digital workers to make entry-level NoOps a near-term reality. Key characteristics from the Ops point of view in this state include:

- Fully autonomous operations, that are run by an AI system which oversees all aspect of the enterprise Ops, including NetOps.
- Industry standard protocols of information exchange across all forms of Ops within the enterprise to ensure seamless sharing of telemetry, alerts, and metadata.
- Supervision of the AI by human operators, while they focus primarily on value creating activities.
- An elevated level of data governance required to guarantee high quality and trustworthy data, along with accurate and contextually relevant metadata.

Example Use Cases

A few network operations use cases that this stage would satisfy:

1. Self-service, self-diagnosing, and self-healing networks. Since an autonomous system is designed to proactively prevent failures and recover quickly, should one occur, the mean time between failures (MTBF) becomes more relevant to track than the time to repair (MTTR).
2. Tightly integrated and correlated NetOps with SecOps, DevOps, and other relevant Ops functions within the enterprise, given that network operations capture relevant telemetry that other Ops functions would find useful as well.
3. Unknown 'unknowns' that are not just discovered (using advanced ML models) but also autonomously remediated, with the collective learning then added to the AI knowledge base.

Tracfone: A Real-life Customer Success Story with AIOps Maturity

There is nothing better than a real customer case study to prove that the maturity model discussed above can work for a real business with real network operations challenges. We use Selector's Tracfone case study[9] wherein the enterprise successfully used AIOps solutions from Selector to minimize network downtime and reduce MTTR, moving from a Level 2 (Active Ops) to Level 3 (AIOps) as defined in the model above.

Tracfone's 'Before State'

Level 2 (Active Ops) maturity

- Enterprise operations teams were struggling to cope with IT operations issues that had implications across network, infrastructure, and application layers.
- Operations staff worked in silos, doing very well with what they did in their domain, but unable to effectively collaborate cross-domain.
- Various operations management tools, observability solutions, and open source frameworks were already in use that were not truly providing ROI for a cross-domain holistic diagnosis.
- High MTTR, noisy alerts, and high levels of toil were experienced by operations teams when tackling cross-domain challenges.

What did Tracfone do?

Employed an AIOps Solution

- Worked with Selector to devise an umbrella solution that ingested data from 18 different sources including several operations management tools, open source frameworks, and observability solutions that were already in active use. This effectively broke down the operational data silos that previously existed, allowing for one central repository, i.e., a single source of truth.
- Employed AI-driven alert and event filtering (de-duplication) to reduce the noisy data inputs.
- Achieved AI-based correlation of data inputs that resulted in actionable cross-domain analytical insights that led to quick root cause determination.
- Gained the use of graphical aids (dashboards, graphs, etc.) that presented a unified view of the cross-domain data and analytical insights to all operations teams.

Tracfone's 'After State'

Level 3 (AIOps) maturity and beyond

- A single pane of glass for end-to-end visibility, using inputs from 18 different tools/ solutions (Ops management, observability, open source) to derive actionable insights tackling cross-domain operations challenges.
- Huge reduction in MTTR, major reduction of overall incidents, enhanced system availability.
- Enhanced overall experience due to significantly improved service levels that Tracfone could provide its customers.
- Recognition by IT Leaders of the immense potential of AIOps and the ROI it offers.
- A path for Tracfone to enhance AI adoption and automation for its operations, moving it closer towards closed-loop automation (level 4 maturity).

References

1. "AIOps (Artificial Intelligence for IT Operations)" by Gartner
2. "Hourly Cost Of Downtime" by Laura Didio, Principal Analyst, ITIC
3. "Volume of data/information created, captured, copied, and consumed worldwide from 2010 to 2020, with forecasts from 2021 to 2025" by Statista
4. "Understanding Observability Platforms: The Key to Improved Performance and Reliability", a blog by Selector
5. "Open LLM Leaderboard" at Huggingface.io
6. "I Don't Want DevOps. I Want NoOps." by Mike Gualtieri, VP Principal Analyst at Forrester
7. "Robotic Process Automation(RPA)" as defined by Gartner
8. "3 Ways Digital Worker Technology Transforms Future of Work" by Cem Dilmegani, Principal Analyst at AIMultiple
9. "How Selector Analytics Minimized Network Downtime and Reduced MTTR for TracFone" at Selector.ai

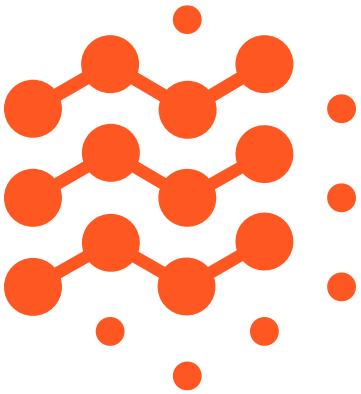

About Selector

As an industry-leading networking AIOps company, Selector provides fully managed AI-based analytical solutions and services that give enterprises instant, actionable insights to manage their multi-hybrid-cloud network and application infrastructures. By reducing the time taken to identify, detect, and resolve network issues, Selector can help improve network performance and efficiency, thereby reducing overall costs.

Selector offers a fully integrated capability stack, with comprehensive data ingest (Collect), which is analyzed for hidden insights (Correlate) and provided back to users as actionable insights that allow operations teams to get straight to the root cause of the problem (Collaborate). Selector is trusted by dozens of leading companies across different industries. Learn more at https://www.selector.ai

About the Contributors

Ron Reuben, *Author*

Ron Reuben is a software product and business leader with over 25 years industry experience in enterprise data and AI. He enjoys writing technical, product and market oriented collateral to share his experience and knowledge with enterprise practitioners and decision-makers.

Ron currently serves as a global executive consultant, helping both startup and large enterprise companies overcome data and AI challenges. He previously held roles in executive leadership, including Head of Product, at technology companies in Silicon Valley.

As a product leader, Ron has brought flagship products to market, growing some of these from the ground up, and has introduced technical innovations to serve both industry specific and cross-industry enterprise technology needs.

Kevin Kamel, *Editor*

Kevin is a multidisciplinary executive with 20 years of experience in product development, software engineering, marketing, and customer success. Kevin has spent the majority of his career in fast-paced startup environments.

Before joining Selector as VP of Product Management, Kevin led product, customer success, and sales engineering at Circonus, helping guide the company's growth from seed through Series B. Previously, he served as the VP of Product Development at MailerMailer LLC where the product portfolio he developed won numerous awards and contributed to the company's acquisition by Ziff-Davis in 2017.

Kevin earned a B.S. in Computer Engineering from the University of Maryland, while working as a system safety and reliability engineer at NASA. After graduation, he accepted a faculty position at UMD's MIND Lab, a public-private partnership which served as a DC-area tech incubator.

Chelsea Rio, *Designer*

Chelsea Rio is a creative professional with nearly 20 years of experience in design and marketing. She is especially passionate about elevating visuals, messaging, and overall user experience to inspire and delight audiences.

As the Marketing Manager at Selector, Chelsea plans and implements strategies that increase brand awareness, attract and retain customers, and drive sales. Prior to this role, she served as the creative lead for several SaaS companies as well as a marketing firm. She has also taught graphic design as an adjunct professor at the University of Maryland, College Park.

Chelsea earned her MFA in graphic design from Savannah College of Art and Design and holds a BA in studio art and BS in zoology from the University of Maryland, College Park.