White paper
Cisco public

Autonomous Networks for Service Providers

Contents

Executive summary

As service providers strive to enhance network performance and customer satisfaction, they face increasing challenges from the dynamic telecommunications environment. This white paper presents a comprehensive exploration of the transition from traditional network management to autonomous networks, emphasizing the integration of Artificial Intelligence (AI) and Machine Learning (ML). A critical aspect of this transition is the concept of the "long tail" in automation, which focuses on addressing a multitude of niche tasks that are often overlooked in favor of more frequent, high-impact activities. By tapping into this long tail, service providers can unlock significant opportunities for innovation and efficiency, ultimately leading to better service delivery.

This paper is based on customer proposals focused on the autonomous network journey, integrating recent insights and innovations discussed within industry forums such as the Telemanagement Forum (TM Forum). Our work aims to enhance the development of autonomous network solutions that are not only responsive to current market demands but also leverage cutting-edge technology and best practices from across the industry. It highlights collaborative approaches between service providers and technology vendors, the importance of a structured Proof of Concept (POC), and the implementation of a clean room strategy to protect intellectual property and customer data. Within the framework of the structured POC, a strategic approach is taken to create domain-based AI agents, which will eventually evolve into autonomous AI agents. These agents are designed to communicate with each other through an agentic framework, enhancing coordination and collaboration across network functions.

By adopting an autonomous network framework based on standards like the TM Forum model, service providers can achieve greater operational efficiency, reduced response times, and enhanced customer experiences. This paper covers validated AI and ML use cases to automate and optimize infrequent yet critical tasks, driving substantial improvements in network management and operational resilience.

Introduction

In today's rapidly evolving telecommunications landscape, the need for efficiency and adaptability has never been more critical. Large service providers are increasingly recognizing the imperative to automate various aspects of their operations. Automation not only encompasses management tasks, service demands, and operational activities but also extends to more advanced use cases that facilitate the development of a closed-loop automation system. This transition is vital for ensuring responsiveness and agility in an increasingly competitive market.

Moreover, we are witnessing a significant evolution in the application of AI within network management. The journey from reactive AI, where responses are generated post-incident, to proactive and preventive AI, which anticipates issues before they arise, marks a pivotal shift in operational strategy. This progression further advances to prescriptive AI, where intelligent systems provide actionable recommendations based on data-driven insights, enabling operators to make informed decisions that enhance service delivery.

In addition to embracing automation and AI, traditional network companies are undergoing a transformation to become service-oriented companies. This shift entails not only a change in operational focus but also a fundamental reimagining of how services are delivered and consumed. By prioritizing customer experience and leveraging advanced technologies, these organizations can create value-added services that cater to the evolving needs of their clientele.

As we explore the journey toward an autonomous network, it is essential to consider how these trends integrate with the broader business objectives of the service provider. By aligning automation and AI initiatives with strategic goals, large service providers can position themselves at the forefront of innovation, ultimately enhancing their competitive advantage in the marketplace.

The need for autonomous networks

Challenges with traditional network management

Service providers encounter several significant challenges with conventional network management, including:

- **Scalability issues:** Traditional systems struggle to dynamically scale resources in response to changing demands.
- **High operational costs:** Increased complexity leads to elevated costs associated with maintenance and operations.
- **Delayed response times:** Service disruptions can result in prolonged downtimes and diminished customer satisfaction.

Benefits of autonomous networks

Autonomous networks present a viable solution to these challenges by:

- **Enhancing operational efficiency:** Automation streamlines network operations, diminishing the need for manual intervention.
- **Improving response times:** AI-driven analytics facilitate real-time monitoring and proactive issue management.
- **Boosting customer satisfaction:** Reliable service delivery leads to improved customer experiences.

TM Forum for autonomous networks

The TM Forum provides a framework for assessing the maturity of service providers in their journey toward autonomous networks. This maturity model is structured across different levels, as shown in Figure 1:

- **Level 1:** Basic automation with manual intervention; service providers rely primarily on traditional management systems.
- **Level 2:** Intermediate automation; some automated processes are in place, but significant manual oversight is still required.
- **Level 3:** Advanced automation; service providers use AI agents to manage network operations with minimal manual intervention inside an autonomous domain, improving efficiency and reducing response times.
- **Level 4:** Proactive automation; AI and ML are extensively integrated to predict network demands and issues across multiple domains or cross-domains, enabling self-healing capabilities.
- **Level 5:** Fully autonomous operations; comprehensive end-to-end automation across all network domains, with seamless integration of business processes that includes applications.

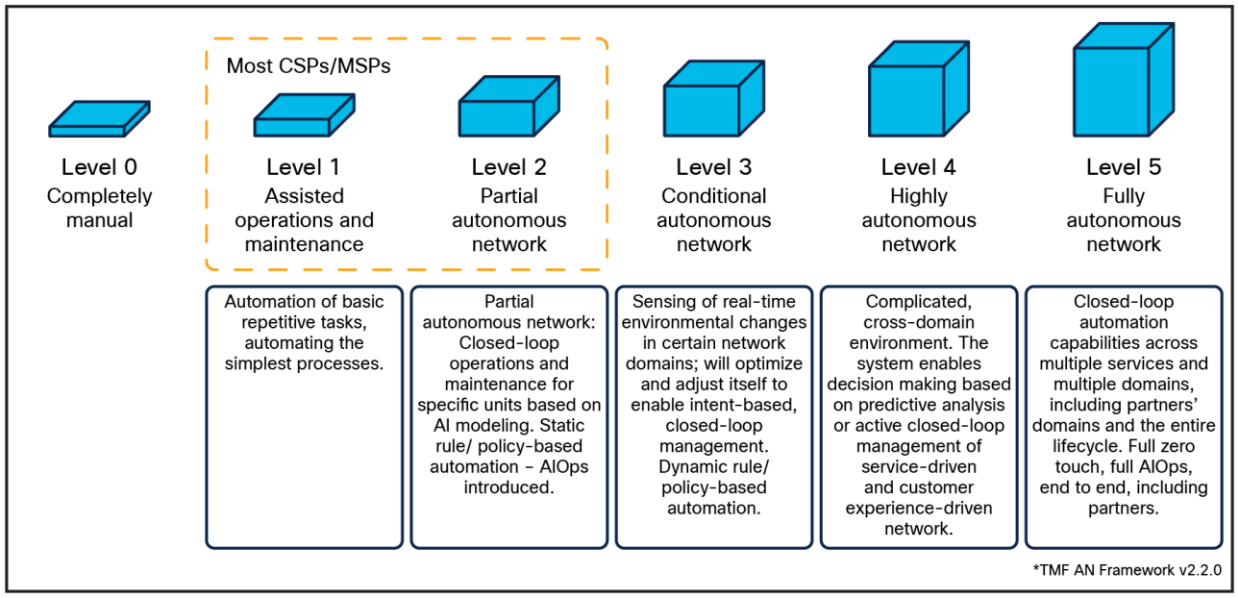

Figure 1.
Autonomous network transition to levels 3 through 5

Most service providers currently operate at **level 2:** they have begun implementing some automation but still rely heavily on manual oversight. To advance to **level 3, 4, or even 5**, the adoption of AI through AI agents based on Large Language Models (LLM) or Small Language Models (SLM), or ML, is crucial.

TM Forum maturity model

Service providers should evaluate their current maturity levels within the TM Forum framework and outline a strategic path toward attaining advanced automation. The primary objective should be to implement closed-loop automation within autonomous domains such as transport, Radio Access Network (RAN), access, or core networks, aligning with TM Forum Level 3 capabilities. Ultimately, the aim is to develop service closed-loop automations that integrate multiple autonomous domains, progressing toward a comprehensive business closed loop that interlinks various service automations across the organization, reflecting the aspirations of TM Forum level 4.

An assessment of the service provider's existing automation level using the TM Forum level evaluation criteria can be conducted using a manual assessment sheet like the one shown in Table 1.

Table 1. Autonomous network assessment tool

Autonomous Levels	L0: Manual Operation and Maintenance	L1: Assisted Operation and Maintenance	L2: Partial Autonomous Networks	L3: Conditional Autonomous Networks	L4: High Autonomous Networks	L5: Full Autonomous Networks
Execution	P	P/S	**S**	**S**	**S**	**S**
Awareness	P	P/S	P/S	**S**	**S**	**S**
Analysis	P	P	P/S	P/S	**S**	**S**
Decision	P	P	P	P/S	**S**	**S**

Autonomous Levels	L0: Manual Operation and Maintenance	L1: Assisted Operation and Maintenance	L2: Partial Autonomous Networks	L3: Conditional Autonomous Networks	L4: High Autonomous Networks	L5: Full Autonomous Networks
Intent/Experience	P	P	P	P	P/S	**S**
Applicability	N/A	Select scenarios				All scenarios

P : People (manual); S : Systems (autonomous)

Autonomous Levels	L0: Manual Operation and Maintenance	L1: Assisted Operation and Maintenance	L2: Partial Autonomous Networks	L3: Conditional Autonomous Networks	L4: High Autonomous Networks	L5: Full Autonomous Networks
Execution	0	0.5	1	1	1	1
Awareness	0	0.5	0.5	1	1	1
Analysis	0	0	0.5	0.5	1	1
Decision	0	0	0	0.5	1	1
Intent/Experience	0	0	0	0	0.5	1
Total	**0**	**1**	**2**	**3**	**4.5**	**5**

Numerical model P = 0 S = 1

Describe operational flow of the evaluation object

Foundation blocks of an autonomous network

The foundation of an autonomous network is built upon three critical blocks, as shown in Figure 2, that ensure the effective functioning of AI-driven systems:

1. Insight/observability

- **High-quality data:** Autonomous networks rely on both real-time and non-real-time data to function effectively. High-quality data is essential for accurate analysis, enabling better decision making and operational efficiency.

2. Decision making using AI

- **AI-driven insights:** Using AI, autonomous networks determine the "what," "why," "when," "where," "how," and "who" related to network operations. This structured decision-making process enables the network to react intelligently to various situations.

3. Automation

- **Intent activation:** This component focuses on activating network intents related to configuration and operational state. Automation can occur in two modes:
  - **Scheduled:** Predefined tasks executed at specific intervals.
  - **On demand:** Tasks triggered based on immediate network requirements.

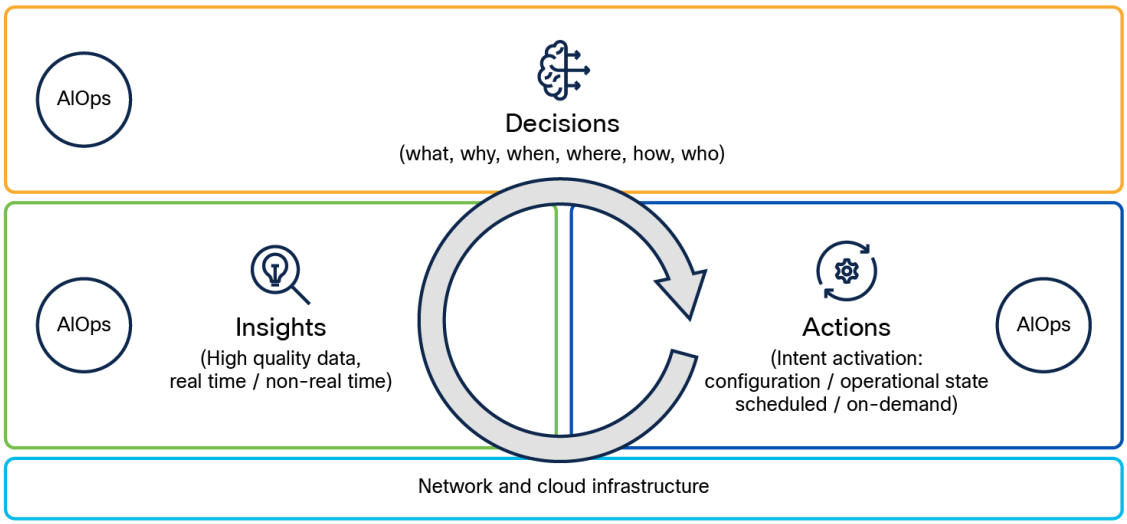

Figure 2.
Automation architecture – foundation blocks

Goal-driven automation

Historically (and this is not specific just to networks but to processes in many industries), automation has focused on tasks. A set of tasks, as shown in Figure 3, is defined as part of a higher-level abstraction such as an intent or, even further, a goal. Traditionally, it is the humans that transform a goal into a set of intents or tasks to be executed. The autonomous network will require us to elevate the scope of automation not just to tasks but also to goals. There is now technology capable of scoping and planning the execution of a goal in ways that were not possible only two years ago. This revolution from task- or rule-centric automation to goal-centric automation is one of the key shifts that will make the autonomous network an achievable goal, enabled, in a big part, by LLM-powered AI autonomous agents.

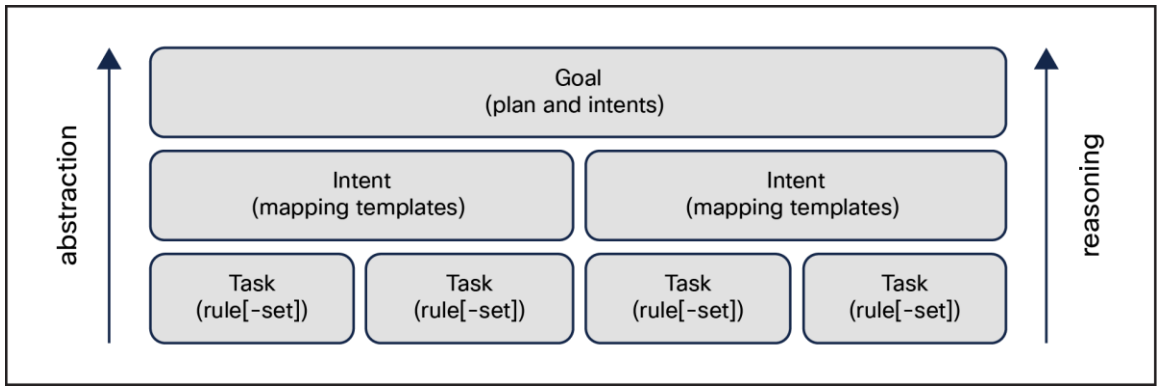

Figure 3.
Goal-driven automation

Simplified network architecture

A simplified network architecture supports these foundational blocks, allowing for seamless integration and efficient operations. This architecture focuses on four key areas: network optimization, layer convergence, service convergence, and edge intelligence.

Network optimization

Network optimization involves creating a converged network architecture that supports both IP-based services and traditional telecommunications services. By reducing the number of routing protocols in use, network operators can streamline operations and enhance performance. The goal is to establish standardized, modular, and resilient designs that facilitate easy scalability and maintenance. This includes:

- **Reduction of multiple routing protocols:** By minimizing the variety of routing protocols, the network can operate more efficiently, reducing complexity and potential points of failure.
- **Standardized modular resilient designs:** Implementing modular designs enables quick adaptation to changing service requirements while maintaining high levels of redundancy and reliability.

Layer convergence

Layer convergence refers to the integration of multiple network layers into a single routed optical network. This approach simplifies the network architecture by merging different layers, such as data link, network, and transport layers, into one cohesive structure. Benefits of layer convergence include:

- **Increased efficiency:** Reducing the layers involved in service delivery enables network management that is more efficient and less prone to errors.
- **Optimized resource utilization:** Converging layers allows for better use of available resources, leading to improved performance and reduced operational costs.

Service convergence

Service convergence uses advanced technologies such as segment routing to facilitate the handling of various types of services while ensuring that Quality of Service (QoS) requirements and Service-Level Agreements (SLAs) are met. This includes:

- **Flexible service management:** Segment routing enables the efficient management of diverse services over a unified infrastructure, allowing for dynamic adjustments based on demand.
- **QoS and SLA compliance:** Implementing service convergence strategies can enable networks to ensure that all services operate within the defined QoS parameters and adhere to agreed SLAs, resulting in improved customer satisfaction.

Edge intelligence

Edge intelligence involves deploying distributed edge data centers or multiaccess edge computing solutions to address the increasing demand for low-latency services. This approach enhances network performance and reliability through:

- **Low latency:** By processing data closer to the source, edge intelligence significantly reduces latency, which is crucial for real-time applications and services.
- **Reduced blast radius during outages:** In the event of a network outage, edge intelligence minimizes the impact by ensuring that critical services can continue to function independently, thereby enhancing overall network resilience.

Through the integration of these elements—network optimization, layer convergence, service convergence, and edge intelligence—autonomous networks can achieve a higher level of efficiency and adaptability, ultimately leading to improved operational performance and customer experiences.

Customer case study

We aim to provide an approach grounded in our extensive experience with large service providers. In existing environments, these organizations must effectively address their "run the business" issues while simultaneously undertaking significant transformations. These transformations are crucial, as they not only lead to substantial efficiency gains but also contribute to improvements in other critical Key Performance Indicators (KPIs). Specifically, organizations can expect enhanced time to market, increased availability, better compliance rates, and improved customer satisfaction.

This section covers the proposed use cases based on key customer organizations such as network operations, engineering, testing/quality assurance, and regulatory compliance. By focusing on these functional areas, the autonomous journey aims to optimize these business structures, driving efficiency and effectiveness throughout the organization. We will outline initiatives that can be prioritized for POC implementation within these domains. By aligning operational improvements with strategic objectives, large service providers can successfully navigate the complexities of transformation while meeting both current and future market demands.

Exploring the long tail in network automation

In the ever-evolving landscape of technology, significant business breakthroughs often emerge by tapping into the "long tail." This concept, popularized by Chris Anderson, refers to the multitude of niche markets or tasks that are frequently overlooked because they may not seem as immediately profitable or prominent as the "head" of the distribution curve. Industry leaders like Google and Amazon have successfully transformed their sectors by automating processes that address these neglected areas, unlocking substantial opportunities for innovation and efficiency.

Network automation, much like other domains, has traditionally focused on automating high-frequency or critical tasks—the "head" of the curve. Examples of such tasks include:

- Creating Threshold Crossing Alert (TCA) rules for frequently monitored metrics.
- Developing parsing rules for common syslog messages.
- Setting up event correlation rules for typical event combinations.

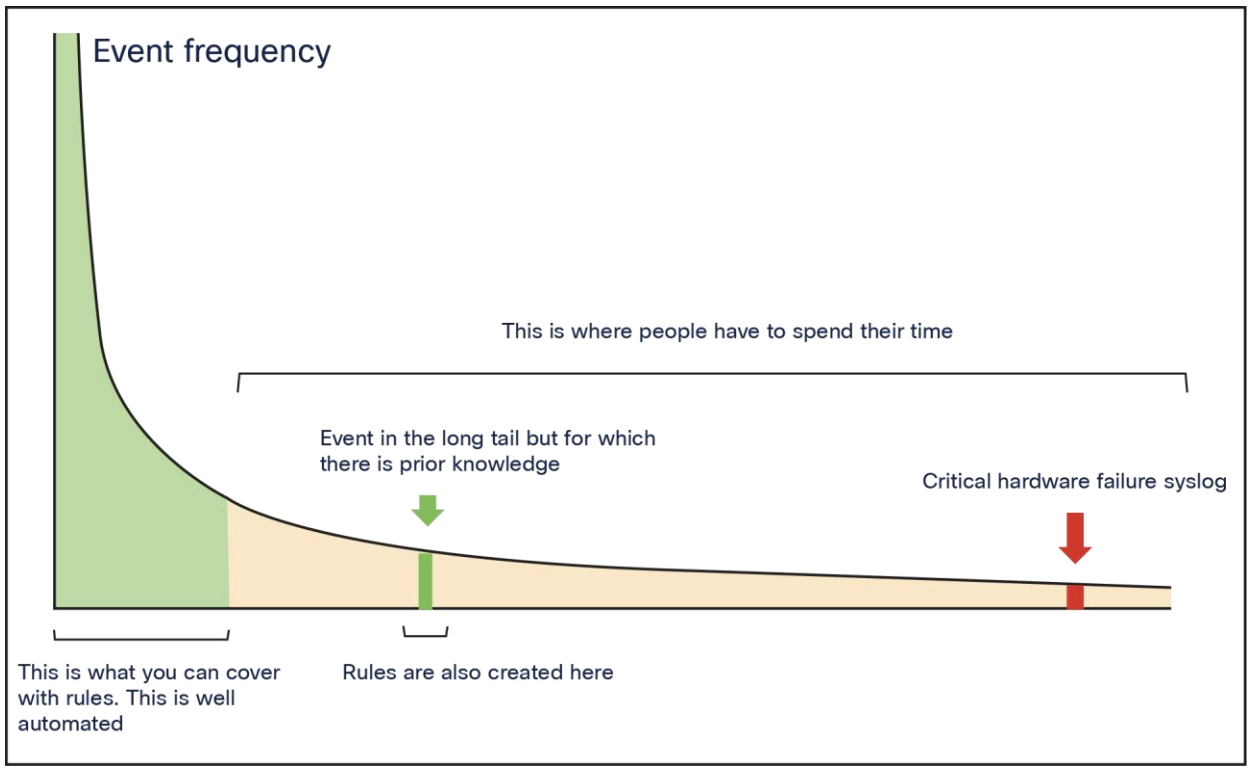

Figure 4.
The long tail in network automation

However, this narrow focus leaves a significant number of less frequent yet potentially critical tasks unaddressed—the long tail, as shown in Figure 4. The next network outage or an essential syslog message may arise from this overlooked segment, leading to disruptions that could have been prevented with proactive automation.

Leveraging AI and ML for transformation

To effectively tackle the challenges presented by the long tail, a fundamentally different approach is required—harnessing the capabilities of AI and ML. These advanced technologies provide powerful tools for automating and optimizing tasks without necessitating explicit rules for every conceivable situation.

Examples of AI and ML applications in network automation

- **Dynamic thresholds:** Unlike static TCA rules, ML algorithms can dynamically adjust thresholds based on real-time patterns and trends, ensuring more responsive monitoring.
- **Comprehensive monitoring:** ML can simultaneously track a wide array of quality factors, identifying anomalies that traditional static monitoring might overlook and enhancing overall network visibility.
- **Intelligent log analysis:** AI can learn from historical data to infer and classify new syslog messages, enabling it to recognize critical events without the need for predefined parsing rules.
- **Automated event correlation:** ML algorithms can detect correlations between events that might not be immediately obvious, which significantly improves incident response times.
- **Adaptive reporting:** AI can generate custom reports in real time, tailored to current network conditions and requirements, ensuring that relevant insights are readily available.

- **Dynamic intent creation:** Leveraging ML, network intents and templates can be created dynamically, allowing them to adapt to evolving network policies and configurations seamlessly.
- **Intelligent workflows:** AI agents can automate a diverse range of tasks, learning from past actions to enhance efficiency and streamline operations.
- **Automated compliance:** Intelligent agents can continuously verify compliance across all devices and configurations, eliminating the need for manually crafted rules and ensuring adherence to organizational standards.

Based on this innovative approach to addressing the long tail of network automation, we have developed a series of use cases that aim to enhance the efficiency, reliability, and adaptability of service provider networks.

LLM or SLM or ML – _when and what use cases?

AI Agents built using LLMs will result in a low-code/no-code capabilities and minimal maintenance requirements. Most of AI development capabilities needed to implement these AI agents can be provided by Multi AI enabler agent Platform SDKs. We will discuss about Multiagentic Platform in the next section.

An example of LLM based AI agents is captured very well in TMF Catalyst Project on Incident Co-Pilot, which focuses on using GenAI-powered solutions to enhance proactive incident management, reduce repeat incidents, and shorten resolution times. This is the basis for NOCLess or MANless NOC or Dark NOC, where lot of process automation can be achieved.

With LLM technique, we do not have to train or fine-tune the model, as this process is highly expensive and does not leverage real-time data as effectively as the Retrieval-Augmented Generation (RAG) technique. LLMs are best suited for complex tasks requiring a deep understanding of language and context, making them ideal for applications like natural language processing to convert a request to an API call, content generation, and advanced analytics. In contrast, SLMs are more efficient for specific, narrowly defined tasks such as configuration auditing and anomaly detection, where they can leverage their lightweight architecture for faster processing and lower resource consumption. Organizations should choose LLMs when broader contextual knowledge and nuanced understanding are necessary, while SLMs are preferable for straightforward, targeted operations that prioritize speed and efficiency. We will cover the use cases in subsequent sections.

Autonomous AI agents

The evolution toward autonomous agents in network automation marks a significant shift from traditional rule-based approaches to a more intuitive, goal-driven paradigm. Over the past decade, network automation has progressed from a task-centric methodology to an intent-based management, which, while beneficial, still relies on predefined rules. Now, advancements in technology allow for the development of systems that autonomously understand and act upon goals, enhancing decision-making capabilities. This new era of goal-driven network automation emphasizes rule independence and natural language accessibility, enabling users to define goals intuitively. With high autonomy, these systems can independently devise plans that account for context and constraints, leading to improved operational efficiency and self-improvement through continuous learning. This transition aligns with the directions set by organizations like the TM Forum, advocating for a more abstracted and ambitious automation framework. The evolution of AI Agents are shown in Figure 5. The AI Agents covered in this document are autonomous agents in the “soon” stage. These AI agents finally will be able to call other AI agents in the “Later” stage using the Agentic platform. An agentic platform is a system that enables communication, collaboration, and coordination between autonomous AI agents. These platforms provide the infrastructure and tools needed for AI agents to interact with each other, external systems, and human users to achieve complex tasks with minimal human intervention.

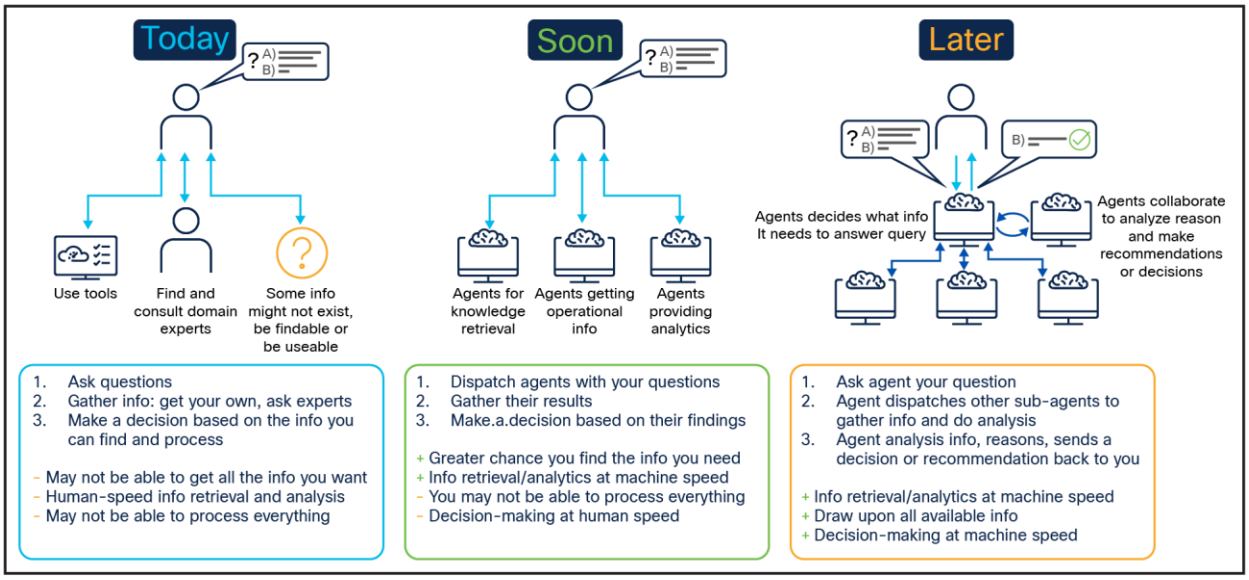

Figure 5.
AI Agents Evolution

Autonomous agents to enable closed-loop operations

Autonomous agents play a crucial role in enabling closed-loop operations within networks by automating various processes across different domains, as shown in Figure 6. These closed loops include:

- **Network element closed loop:** This loop focuses on the individual network elements, automating their operations and maintenance. Autonomous agents can monitor performance, detect faults, and initiate self-healing actions, ensuring that network elements operate efficiently without human intervention.
- **Resource closed loop:** This loop is concerned with managing network resources at a domain level. Autonomous agents optimize resource allocation and utilization, enabling real-time adjustments based on service demands. They ensure that resources are dynamically configured and reconfigured as needed, supporting functionalities like self-optimization and self-healing.
- **Service closed loop:** This loop encompasses end-to-end service management across multiple domains. Autonomous agents manage service requests, fulfillment, and lifecycle management, ensuring consistent service quality. They translate service intents into resource intents, coordinating the necessary actions across different resource domains.
- **Business closed loop:** This loop connects business objectives with service and resource operations. Autonomous agents facilitate communication between these layers, translating business intents into actionable service and resource requests. They continuously monitor and assure that business goals are met, adapting to changes in user demands or market conditions.

These autonomous loops promote efficiency, reduce operational costs, and enhance user experience by ensuring that the entire lifecycle of network services—from resource allocation to service delivery—is automated and optimized dynamically. This approach leverages the latest technologies such as AI, 5G, and edge computing to create a more responsive and self-sufficient network environment.

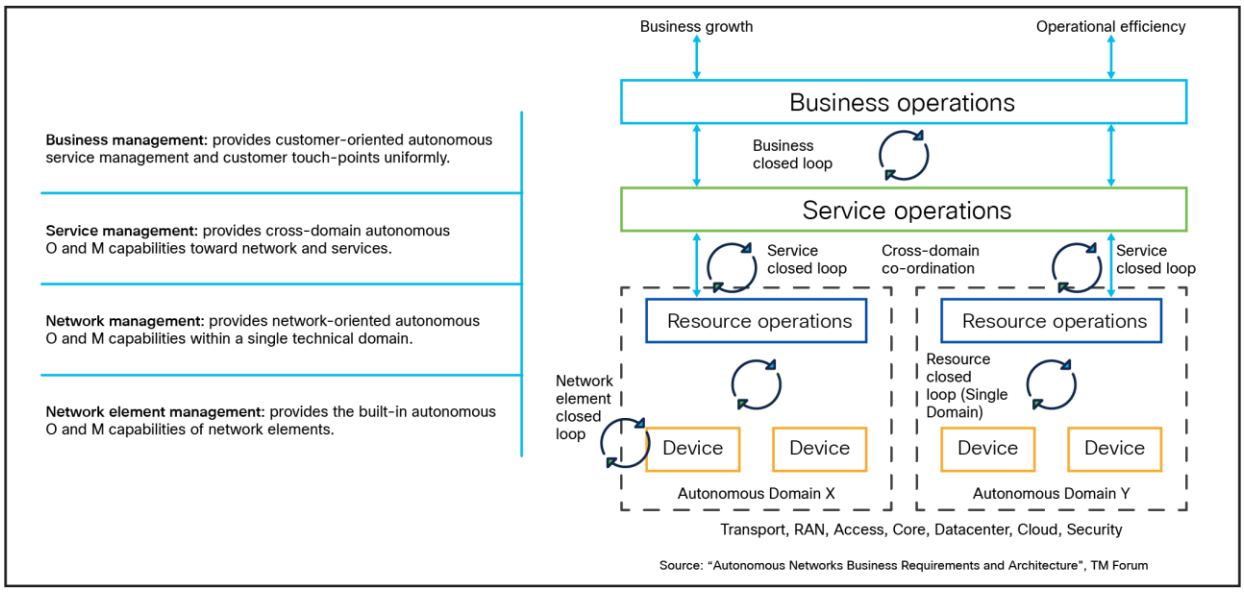

Figure 6.
Autonomous network closed loops as defined by the TM Forum

Autonomous network proposed use cases

We have proposed use cases for service provider customers based on multiple discovery workshops and interviews. These use cases, when implemented, will have an immediate impact on moving to levels 3 and 4 in various autonomous domains.

To identify use cases for levels 3 and 4 in an autonomous network context, organizations can leverage the TM Forum's proposed method for defining level 4 objectives based on "value outcomes and autonomous capabilities." This approach emphasizes that value outcomes should align with enterprise strategies and business goals, thereby guiding the development of autonomous capabilities and validating the results of these constructions.

Value outcome goals can be established from four key aspects:

- **Customer experience**, which focuses on enhancing user satisfaction and engagement.
- **Operational efficiency**, aimed at optimizing processes to reduce costs and improve resource utilization.
- **Revenue growth**, which seeks to identify opportunities for new services and market expansion.
- **Risk management**, which involves minimizing operational risks and ensuring compliance with regulatory standards.

By systematically assessing these four dimensions, organizations can effectively define and prioritize use cases that not only advance their autonomous capabilities but also deliver tangible business value.

When selecting use cases, avoid focusing on simply replacing the old with the new. Instead, aim for solutions that offer improvements over the existing ones and yield multiple KPI gains.

Proposed use cases

- **NOC-less system (incident copilot)**

  Technical details: A multiagent system based on a web application and API enables network engineers to perform basic tasks using natural language processing without requiring prior knowledge. It includes agents for incident management, query and response, audit creation, troubleshooting, optimization, and healing. By leveraging modern software practices such as APIs, microservices, and containers, the inferences from these agents can be used in process automation to achieve a closed-loop outcome.

  Architecture: Multiagent approach, LLM, Retrieval-Augmented Generation (RAG).

  Outcome: Improved ease of use and efficiency, with real-time access to information.

- **Digital twin**

  Technical details: A system that creates a virtual representation of physical systems. In some cases, a smaller modular physical representation can be used for performance analysis. By automating the Continuous Integration/Continuous Delivery/Continuous Testing (CI/CD/CT) pipeline, we can fully automate various tasks such as software upgrades and configuration changes. This approach enables us to achieve a high level of closed-loop automation.

  Architecture: CI/CD tools, test manager, pre-production environment and Cisco® Modeling Labs.

  Outcome: Reduced operational costs and improved setup availability (5 nines availability).

- **Layer 2/3 healing agent**

  Technical details: Anomaly detection coupled with automated remedial actions and rollback options, supported by an optional approval workflow and a situation room providing a 360-degree view of anomalies. This leverages automation to close the resource closed loops.

  Architecture: SLM, telemetry data for optical and network layers, automation tools.

  Outcome: Self-healing networks that minimize Mean Time To Detect (MTTD) and Mean Time To Repair (MTTR), enhancing customer experience.

- **Single visualization dashboard for incident agent**

  Technical details: A unified dashboard that provides operational, planning, and executive views in a single interface. Service-level monitoring and Root Cause Analysis (RCA). The intent here is to leverage data beyond network infrastructure for a service-level closed loop. This is a foundational approach for levels 4 and 5, leveraging goal-driven automation.

  Architecture: RCA based on correlation techniques on syslogs and other structured and unstructured data using ML.

  Outcome: Improved MTTD and MTTR, facilitating effective cross-application launches.

- **Configuration drift detection**

  Technical details: The solution leverages machine learning models to analyze network configurations across devices and flag the deviations detected on the device configurations under test. These models learn normal configuration patterns and detect deviations in real time, effectively addressing the limitations of static rule-based systems. By automating the detection process, this solution ensures that configurations remain consistent with the intended settings, reducing the risk of performance issues, security vulnerabilities, and compliance failures. By integrating the inference from this model into an automation tool like an end-to-end service activation platform, we can achieve a high degree of closed-loop automation.

  Architecture: SLM and/or traditional ML models.

  Outcome: Early detection of configuration anomalies, leading to minimized outages and a cost-effective, simplified approach.

- **LLM-based configuration audit**

  Technical details: A simplified auditing tool that leverages best practice rules for real-time detection of configuration anomalies in production environments. Focus is on network resiliency. The audit findings should be leveraged to proactively remediate issues. The findings can be leveraged by an end-to-end automated service activation platform to achieve a high degree of closed-loop automation. This also crosses the autonomous domains and can help achieve levels 3 and 4.

  Architecture: LLM, RAG with best practice rules.

  Outcome: High availability and reduced outages with lower operational costs.

- **End-to-end automated service activation**

  Technical details: An open orchestration platform that enables automation across multiple domains, including a service catalog and comprehensive capabilities for pre- and post-checks. AI agents dedicated to configuration change management can integrate with a multiagent framework and take actions based on triggers from other agents. This will be an important enabler for achieving levels 3, 4, and 5 of the TM Forum autonomous network model.

  Architecture: LLM with RAG providing API mapping with natural language requests.

  Outcome: Automated configuration changes leading to lower costs, faster provisioning, and a high level of service availability (5 nines availability). This also allows closed-loop automation in service closed loops and business closed loops.

- **Toxic factor detection**

  Technical details: A monitoring system that uses ML algorithms to identify and flag toxic factors within network traffic or application performance metrics. This proactive system helps detect issues by analyzing historical trends and creating a health and diagnostic framework. Issues such as soft failures, often caused by a combination of software versions and hardware types, can be categorized with a high degree of toxic factors. This proactive detection can be integrated into an automation platform to achieve a closed-loop outcome.

  Architecture: ML models, historical data.

  Outcome: Enhanced security and performance management through proactive addressing of harmful factors.

- **Application experience assurance**

  Technical details: Solutions designed to monitor and assure end-user experience across applications, leveraging real-time analytics and user feedback. This proactive detection can be integrated into an automation platform to achieve a closed-loop outcome.

  Architecture: SLM on real-time telemetry data using network probes on services.

  Outcome: Improved application performance and user satisfaction through proactive issue resolution.

- **Metric forecasting**

  Technical details: A generic functionality that may be leveraged for multiple use cases, including congestion mitigation, path optimization, capacity planning, automated maintenance, failure prediction, and sustainability optimization. This proactive detection can be integrated into an automation platform to achieve a closed-loop outcome.

  Architecture: SLM and statistical techniques.

  Outcome: Maintenance of SLA, improved customer satisfaction, and reduced cost by minimizing outages.

- **Autonomous network evaluation**

  Technical details: Assessment framework for evaluating the maturity and performance of autonomous network systems, focusing on self-management capabilities. Feed it the details on the processes involved in a task and whether they are automated or manual.

  Architecture: LLM on various AI standards using RAG or a manual spreadsheet.

  Outcome: Improved understanding of network operational efficiencies and areas for enhancement.

- **Responsible AI**

  Technical details: Implementation of ethical standards and best practices for AI usage within network management to ensure fairness, accountability, and transparency.

  Architecture: SLM or ML approach or manual assessments to secure the data in play.

  Outcome: Enhanced trust and compliance in AI applications, leading to responsible network operations.

Use Case Summary

Some of these proposed use cases focus on leveraging advanced technologies and methodologies to enhance network performance, improve user experience, and ensure responsible AI practices in network management. The integration of these use cases aims to create a more efficient, resilient, and ethical network environment.

Most of the use cases will result in closed-loop automation. This closed-loop automation will leverage inferences from multiple agents to close the loop through the automation of remediation steps. These remediation steps may include configuration changes, upgrades of line cards, or modifications in software. This approach is intended to enhance operational efficiency and minimize manual intervention, ultimately leading to a more self-sustaining network infrastructure.

Use cases from the proof of concept

As service providers embark on their journey toward autonomous networks, the implementation of a detailed POC is critical. This POC should be conducted in a co-development manner between the service provider and network vendors to ensure seamless integration of technologies and processes. A collaborative approach allows both parties to leverage their expertise, fostering innovation and accelerating the path to automation.

The importance of a clean room

To protect the intellectual property of vendors and safeguard customer data during the POC process, the concept of a clean room can be applied. A clean room is a controlled environment where sensitive information can be analyzed and processed without the risk of exposing proprietary technologies or customer data. This approach ensures:

- **Intellectual property protection:** Vendors can share their technologies and methodologies without fear of their proprietary information being misused.
- **Data security:** Customer data remains confidential, preventing any potential breaches that could arise during model training and learning processes.

By implementing a clean room strategy, both the service provider and network vendors can focus on developing effective solutions while maintaining the integrity of their respective assets.

Objectives of the POC

The POC serves several essential purposes in the journey toward an autonomous network:

- **Validation of concepts:** The POC allows the service provider to validate the feasibility and effectiveness of proposed solutions before full-scale implementation.
- **Risk mitigation:** Testing concepts in a controlled environment can identify potential risks and challenges so that they can be addressed early in the development process.
- **Performance benchmarking:** The POC enables the measurement of performance metrics against predefined benchmarks, ensuring that solutions meet desired outcomes.
- **Collaboration and learning:** The POC fosters collaboration between service providers and vendors, facilitating knowledge transfer and shared learning experiences.

Examples of POC use cases and expected outcomes

The examples here cover a subset of POCs for use cases proposed to customers and details on expected outcomes.

- **NOC-less system (incident copilot)**

  Description: The incident copilot is designed to assist network operations centers (NOCs) in efficiently diagnosing and resolving network issues.

  Reason for POC: The POC aims to validate the effectiveness of the system in real-world operational scenarios, ensuring that it can process data from various sources and provide actionable insights.

Acceptance and success criteria:

- **Faster resolution times:** The POC is expected to demonstrate an increase in the speed at which network issues are identified and resolved, reducing downtime and improving overall network reliability.
- **Enhanced troubleshooting accuracy:** By leveraging AI-driven analytics, the system should improve the accuracy of troubleshooting efforts, minimizing the risk of misdiagnosis.

- **Configuration drift detection**

  Description: This use case involves monitoring configuration changes in the network to ensure compliance and integrity, allowing for quick identification of unauthorized changes. It is fully based on SLM, and there is no need for rules.

  Reason for POC: The POC will assess the capability of AI agents to detect configuration drifts in real time, providing insights into compliance with regulatory standards and internal policies. If the customer's automation tool failed in implementing desired configurations, this use case will catch the anomalies.

  Acceptance and success criteria:

  - **Increased compliance assurance:** The POC will validate the system's ability to maintain compliance with SLAs by providing timely alerts for unauthorized configuration changes.
  - **Reduced risk of network vulnerabilities:** By proactively identifying configuration drifts, the service provider can mitigate risks associated with security vulnerabilities and operational inefficiencies.

- **LLM-based configuration audit**

  Description: Using LLMs, this use case conducts comprehensive audits of network configurations and policies to ensure compliance with industry standards.

  Reason for POC: The POC is intended to evaluate the effectiveness of LLMs in analyzing complex configuration data and generating actionable audit reports. Network resiliency with many devices is the prime example for this use case.

  Acceptance and success criteria:

  - **Enhanced audit accuracy:** The POC will demonstrate improved accuracy in identifying noncompliant configurations, leading to more reliable audit results.
  - **Time savings in audit processes:** By automating the auditing process, the POC is expected to reduce the time and resources required to conduct audits, allowing for more frequent evaluations.

- **Single visualization dashboard (Splunk)**

  Description: This use case involves creating a unified dashboard that provides real-time insights into network performance metrics, enhancing visibility for operators.

  Reason for POC: The POC will test the dashboard's capability to aggregate data from multiple sources and present it intuitively. By using SLMs and LLMs, it can target correlation and RCA.

  Acceptance and success criteria:

  - **Improved decision-making:** The POC will validate the dashboard's ability to deliver actionable insights quickly, enabling operators to make informed decisions based on current network conditions.
  - **Enhanced operational visibility:** By consolidating performance metrics into a single view, the POC is expected to enhance operators' understanding of network health and performance, facilitating proactive management.

POC Summary

In conclusion, the journey toward an autonomous network necessitates a comprehensive and collaborative POC that embraces the clean room concept to protect both intellectual property and customer data. By validating concepts, mitigating risks, and benchmarking performance, service providers can ensure successful implementation of autonomous solutions. The four highlighted use cases—incident copilot, configuration drift detection, LLM-based configuration audit, and the single visualization dashboard—demonstrate the potential outcomes that can be achieved through this process. Each use case offers unique insights and enhancements that will ultimately contribute to the operational efficiency and reliability of the service provider's network.

By focusing on these critical areas during the POC, service providers can chart a clear path toward achieving higher levels of automation and realizing the full benefits of autonomous networks.

Conclusion

Transitioning to autonomous networks is a critical step for service providers aiming to enhance operational efficiency and customer satisfaction. By integrating AI and ML into their operations, service providers can not only automate processes but also proactively manage network challenges. Implementing an autonomous network framework will be essential for service providers to maintain competitiveness in an increasingly complex telecommunications landscape.

Learn more

Service providers are encouraged to explore the integration of autonomous network solutions into their operational strategies. By embracing these advanced technologies, they can successfully meet the demands of the future telecommunications landscape. Cisco can help in this journey

Reference

Javier Antich Romaguera, Network Automation. Automating [in] the Long Tail

Authors:

Vijay Raghavendran – Distinguished Engineer, Cisco Systems

Javier Antich – Principal Product Management Engineer, Cisco Systems

Reviewers:

Adrian Quaife – Principal Architect, Cisco Systems

Velimir Vujnovic – Principal Architect, Cisco Systems

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **https://www.cisco.com/go/offices**.

Printed in USA C11-4887932-00 03/25