Partners Universal International Research Journal (PUIRJ)
Volume: 05 Issue: 01 | January– March 2026 | ISSN: 2583-5602 | www.puirj.com

Self-Driving Networks: AI Automation for Enterprise IT

Dr.A.Shaji George

Independent Researcher, Chennai, Tamil Nadu, India.

Abstract – The pressure to manage enterprise networks has never been greater than it is today. The rise of artificial intelligence workloads, the explosion of IoT devices, multi-cloud computing and a dispersed workforce model have created environments that are too complex to manage. In this article we explore the rise of self-driving networks as an organisational and strategic response to this phenomenon. The article builds on the confluence of developments in AIOps, agentic automation, high-performance network hardware design and built-in security to offer a five-point maturity model for network autonomy, a high-level overview of the key architectural building blocks that support self-driving capability, and recommendations for organisations at various points in the transition. The discussion includes the evolution of network management, the distinct features that distinguish self-driving networks from previous generations of automation, the state of the art, industry applications, and the challenges of transition. The article concludes that the enabling components of autonomy are readily available and tested today, and that organizations that begin the structured journey to autonomy now will stand a much stronger chance of competing, securing and scaling AI-powered operations in the future.

Keywords: Self-Driving Networks, AIOps, Agentic Automation, Zero Trust, SASE, Network Autonomy, AI-Native Infrastructure, Enterprise Networking.

1. INTRODUCTION

1.1 The Network Is Drowning

If you ask any enterprise IT professional what keeps them awake at night, the answer is likely to be variations on a theme. Too many alerts. Too little time. And too much of both on issues that should not need human intervention. Today's enterprise networks are ecosystems that have become incredibly complex. A typical mid-sized enterprise today might have a fleet of thousands of devices across multiple offices, hybrid connectivity to two or three public cloud providers, clusters of AI inference servers running production applications, and a dynamic mix of IoT devices, from office building sensors to factory monitoring tools. Each of these components has telemetry. Each has its own points of failure. And each, in the traditional approach to network management, requires human intervention to be monitored, understood and acted on.

It was an acceptable approach when networks were smaller, more uniform, and less dynamic. A small team of highly capable network engineers could maintain a mental model of the environment to which they were responding, they could respond to a predictable set off alarms, and they could ensure service levels were maintained through their expertise and systematic problem solving. That era is over. Today's enterprise network has more events of greater diversity, arriving at a faster pace than any human team can cope with. The consequence is an all-too-common and expensive predicament IT staff spend most of their time fixing problems rather than avoiding them, the best network engineers are bogged down fixing problems instead of designing the network, and the enterprise debt meter ticks ever higher [1, 4].

Partners Universal International Research Journal (PUIRJ)
Volume: 05 Issue: 01 | January– March 2026 | ISSN: 2583-5602 | www.puirj.com

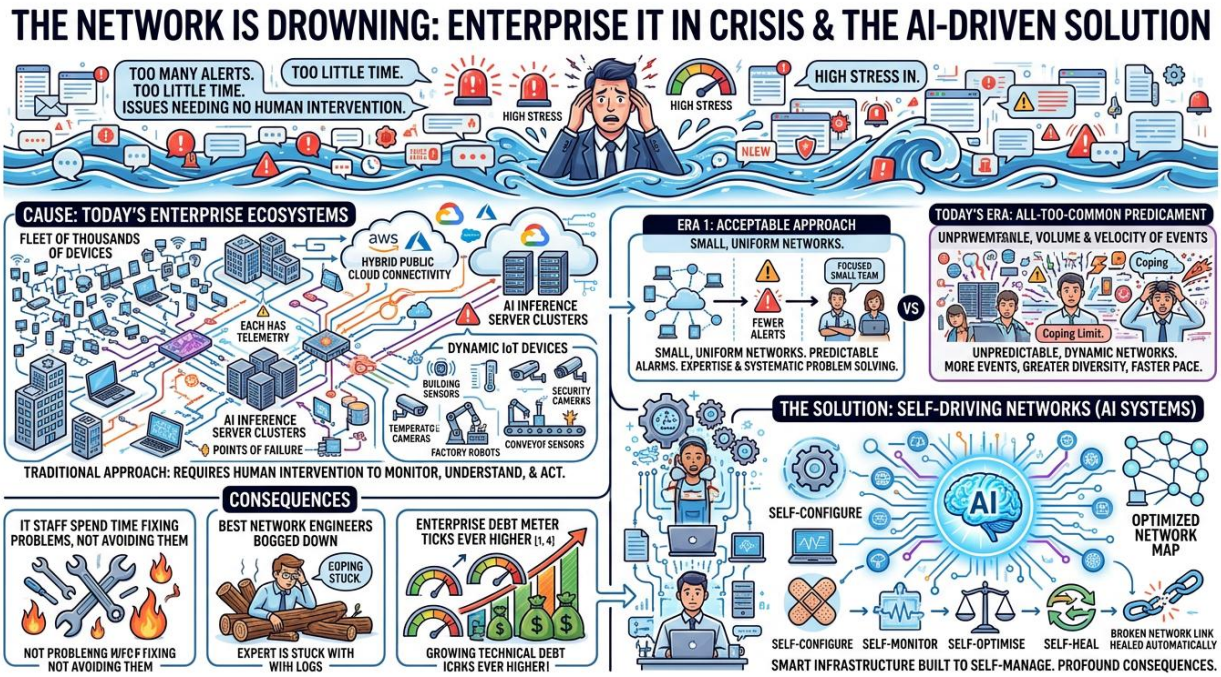

Fig -1: The Network is Drowning Enterprise IT in Crisis & The AI-Driven Solution

The solution to this problem, which is now taking clear shape in the enterprise networking industry, is a new class of smart infrastructure that is increasingly referred to as self-driving networks. These are artificial intelligence (AI) systems that are built to self-configure, self-monitor, self-optimised, and self-heal. The name is apt and the consequences are profound. In this article, we explore the nature of self-driving networks, how they work, their historical origins, and how to transition towards them in a planned and systematic manner [2, 5].

2. OBJECTIVES

This article seeks to achieve the following. First, to give IT leaders, network architects and other decision-makers a conceptually-driven understanding of what self-driving networks are, why they represent an architectural transformation rather than a simple product upgrade, and why they are needed in the current business environment. Second, to chronicle the evolution of network management strategies, and pinpoint where the transition from traditional methods began to falter. Third, to offer a five-phase maturity model for organizations to guide their own autonomy transformation, from collecting telemetry through to becoming a self-driving network. Fourth, to detail the technological building blocks that enable self-driving networks, such as AIOps, agentic automation, AI-ready networking hardware, and network security. Fifth, to examine actual deployment scenarios in retail, healthcare, enterprise campuses and government, to provide evidence that the technology is real and here today. Sixth, to identify the real-world challenges in adopting the technology and provide strategies for overcoming them. Finally, to project the future of the technology, and to identify the key developments that will impact the next generation of enterprise networks.

3. HISTORICAL CONTEXT HOW NETWORK MANAGEMENT EVOLVED INTO ITS CURRENT CRISIS

Partners Universal International Research Journal (PUIRJ)
Volume: 05 Issue: 01 | January– March 2026 | ISSN: 2583-5602 | www.puirj.com

To understand the significance of self-driving networks, it's useful to look at how enterprise networks evolved to their current state of operational dysfunction.

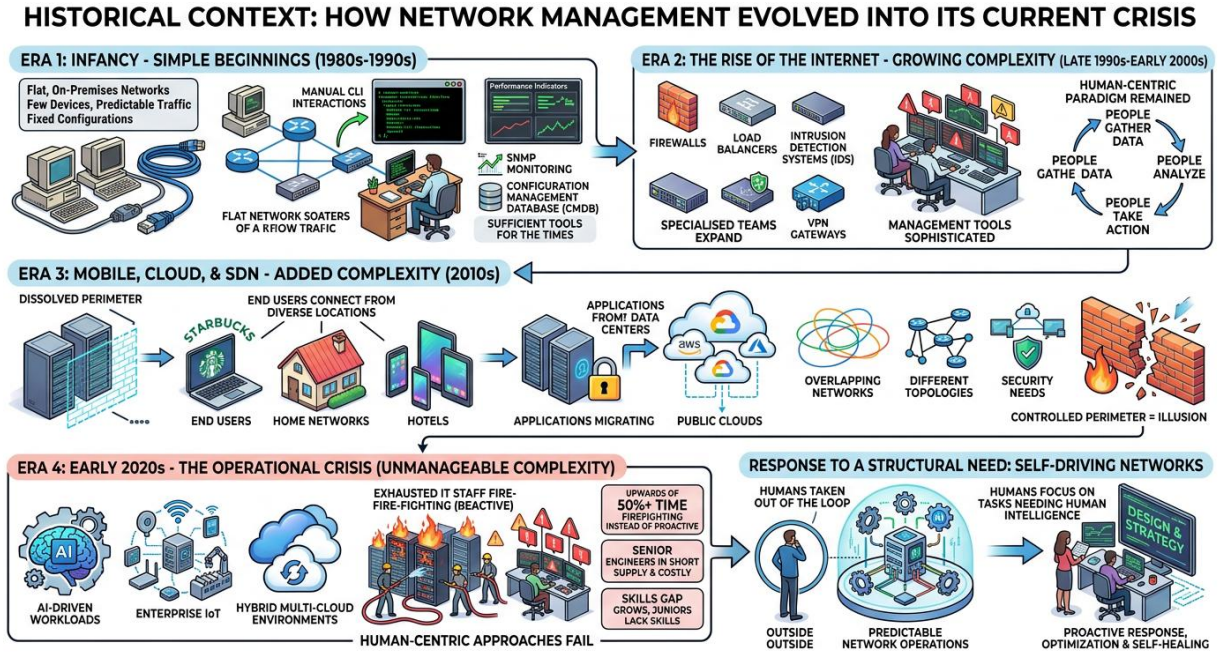

Fig –2: Historical Context How Network Management Evolved into Its Current Crisis

Enterprise networks in their infancy, in the 1980s and 1990s, were simple affairs. Enterprise networks were generally flat, on-premises affairs, with relatively few devices, predictable traffic patterns and fixed configurations. Managing a network involved manually interacting with routers and switches through their command-line interfaces, tracking the health of a limited set of performance indicators, and troubleshooting issues that were usually easy to identify. The network management tools of the day, such as SNMP-based monitoring systems and configuration management databases, were sufficient for the networks they managed.

The rise of the internet in the late 1990s and early 2000s brought complexity. Enterprises began to deploy firewalls, load balancers, intrusion detection systems and VPN gateways. Network teams expanded and became more specialised, and vendors released management tools that were able to merge alarms from different device types onto a single console. But the model remained the same. People gathered information, people analyzed and people took action. The management tools became more sophisticated, but the paradigm remained the same. The transition to mobile devices, cloud computing and software-defined networking in the 2010s added complexity beyond the capabilities of the tools of the day. The perimeter dissolved. End users connected to applications from Starbucks, home networks, and hotels. Applications migrated from the data center to the cloud. The notion of a controlled network perimeter became an illusion. Networks no longer existed in isolation but as a series of overlapping networks with different topologies, security needs, and performance considerations.

By the early 2020s, the introduction of AI-driven workloads, enterprise-scale IoT and hybrid multi-cloud environments meant that the level of complexity was no longer manageable using human-centric approaches in most enterprises. Research indicated that most enterprise IT teams were spending upwards of

Partners Universal International Research Journal (PUIRJ)
Volume: 05 Issue: 01 | January– March 2026 | ISSN: 2583-5602 | www.puirj.com

half their time firefighting instead of working proactively on initiatives [4, 11]. Now add the skills gap Senior network engineers were in short supply and costly, junior engineers lacked the skills to deal with multi-domain problems, and the skill gap in the face of increasing network complexity was growing. Self-driving networks were not a product of marketing requests, but a response to a structural need that recognised the only way forward was to take humans out of the loop for predictable network operations, and to focus human resources on the tasks that needed human intelligence [1, 8].

4. WHAT SELF-DRIVING NETWORKS ACTUALLY ARE

"Self-driving" is a term that comes from the world of automobiles, and the comparison is more apt than it seems. A self-driving car uses real-time processing of sensor data by machine learning algorithms to make thousands of micro-decisions each second that affect speed, steering, lane centring and obstacle avoidance. There's no human in the loop for micro-decisions. The human provides the destination and oversees at the macro level. The system handles execution.

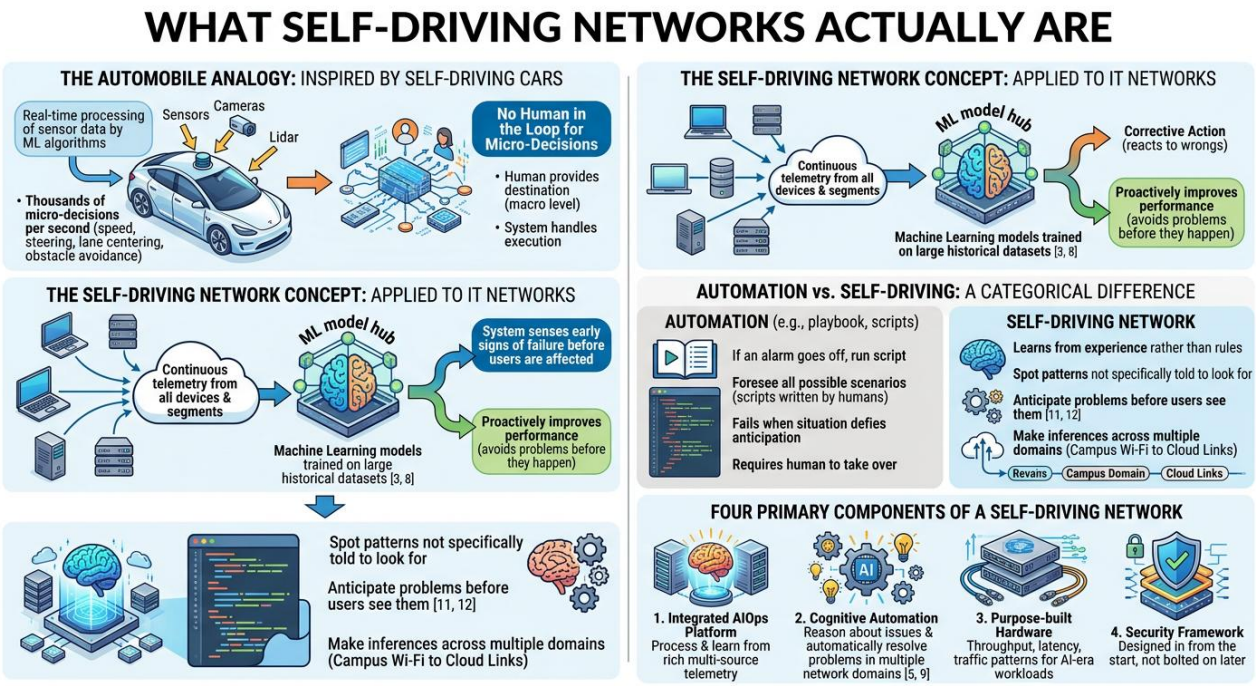

Fig -3: What Self-Driving Networks Actually Are

The same is true of self-driving networks. The network receives continuous telemetry from all devices and segments, feeds it to machine learning models that have been trained on large historical data sets it then detects patterns that mean something is wrong, or going wrong, and takes corrective action, or in more advanced cases, proactively improves performance, rather than waiting for a human to detect the problem and take action [3, 8]. Let's be careful in defining a self-driving network versus an automated one. Automation, as it is typically understood, involves writing a procedure into a playbook or script so that it runs without manual intervention. If an alarm goes off, a script is executed. This is a good thing, and reduces the amount of human work required for common, known circumstances. The downside is that it requires humans to have foreseen all possible scenarios and to have written scripts for those scenarios. When the situation

Partners Universal International Research Journal (PUIRJ)
Volume: 05 Issue: 01 | January– March 2026 | ISSN: 2583-5602 | www.puirj.com

defies anticipation, which happens all the time in complex networks, the automation cannot help and a human has to take over.

Self-driving networks go categorically further. They learn from experience rather than rules. They can spot patterns they haven't been specifically told to look for, because the machine learning algorithms they employ learn from past experience, rather than follow pre-programmed rules. They anticipate problems before they become visible to users, because the AI can sense the early signs of a failure before it affects end users. And they can make inferences across multiple domains, detecting causality that traverses campus Wi-Fi, core switches, WAN and cloud links in a single pass [11, 12]. A self-driving network is comprised of four primary components. The first is an integrated AIOps platform, which stands for artificial intelligence for IT operations, that can process and learn from rich multi-source telemetry. The second is a cognitive automation engine, which can reason about issues and automatically resolve problems in multiple network domains. The third is purpose-built hardware, with the throughput, latency and traffic patterns needed to support the workloads of the AI era. The fourth is a security framework, designed in from the start rather than bolted on as an afterthought [5, 9].

5. CURRENT TRENDS IN AUTONOMOUS NETWORKING

There are a number of trends that are driving implementation of the self-driving network principles in enterprise networks.

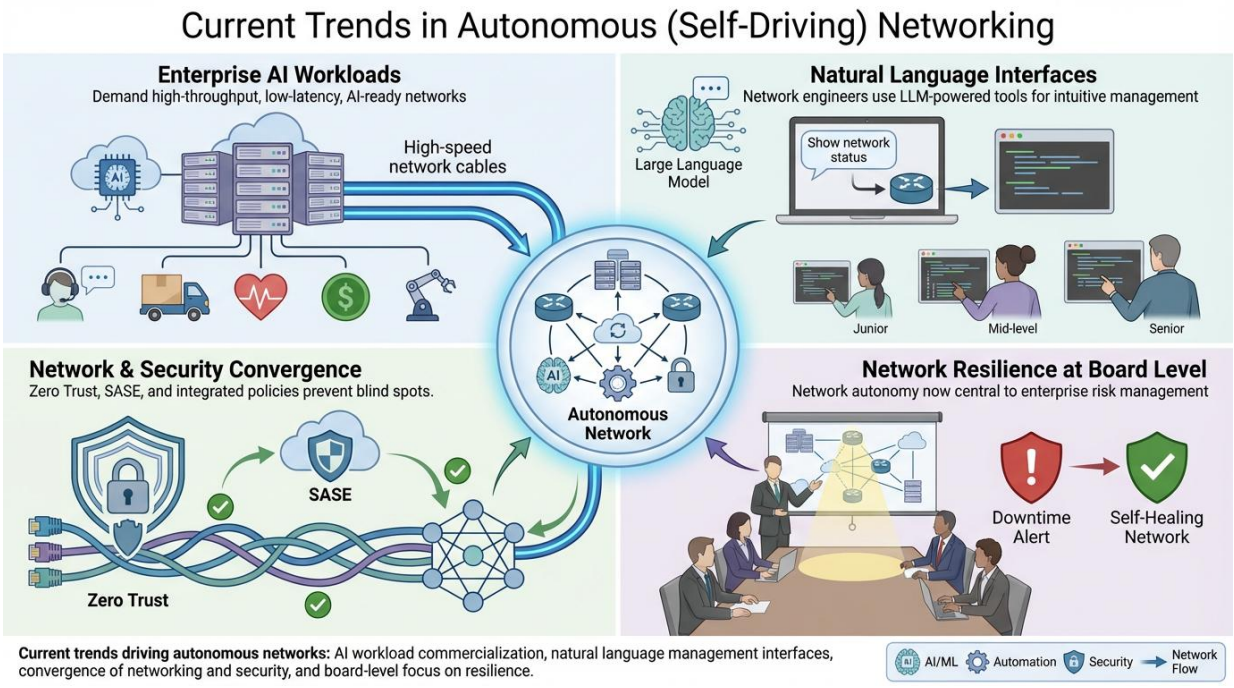

Fig -4: Current Trends in Autonomous (Self Driving) Networking

The most obvious is the commercialisation of AI workloads. Companies are no longer experimenting with AI. They are running AI-powered applications for customer engagement, logistics, medical diagnosis, financial risk prediction and manufacturing process optimisation. These apps demand fast, high-throughput, low-latency network connections and can create traffic patterns that traditional network infrastructure is not built

to handle. The need for AI-ready networking both in terms of network performance and the management visibility and sophistication of the network is driven by the need to run AI applications, not by network teams looking for tools [8, 10].

The second trend is the rapid evolution of large language models as the user interface for network management. Networking tools that were prototypes two years ago now have natural language processing incorporated. IT engineers can ask questions of their networks in natural language and expect meaningful answers [5, 15]. This development is significant not only in terms of user experience but as a new capability. Companies with little networking expertise can now engage with their network in ways that previously demanded senior networking expertise [1, 16].

A third development is the blurring of the lines between networking and security. The adoption of Zero Trust and Secure Access Service Edge (SASE) is progressively blurring the lines between network and security. Security policy is increasingly enforced in the network, and separation of the two disciplines results in blind spots that leave the network vulnerable to attack. Self-driving network platforms with integrated security are increasingly the default architectural model rather than a higher cost premium feature [4, 13].

A fourth trend, connected to the above, is the corporate elevation of network autonomy from being an IT operations issue to one that reaches the board room. As companies rely more on digital services and the impact of network outages and security breaches becomes more visible on the bottom line, network resilience is taking on a more visible role in the board room. Self-driving networks, which promise improved availability and response to security threats, are now topics of enterprise risk management rather than exclusively IT infrastructure conversations [9, 17].

6. THE FIVE STAGES OF NETWORK AUTONOMY

It is crucial for companies to understand the development of self-driving networks as a process rather than an on/off proposition. Autonomy doesn't happen overnight. It unfolds in stages and each stage relies on the capabilities built through the stage before it, and each stage provides operational benefits before autonomy is achieved [1, 11].

Stage 1: Data Collection. The first step on the path to networking autonomy is instrumentation. This is the stage where the network is prepared to collect real-time data from all its elements wireless access points, wired switches, routers, firewalls and end devices. A variety of data is gathered wireless signal strength, wireless channel loading, latency across network segments, packet loss, authentication success and failure, device connection patterns, and traffic volume by type of application and type of user. The effectiveness of the AI that will eventually emerge depends on the effectiveness of this data collection. Companies that use point solutions for different network domains, each with its own monitoring database for wireless, wired, and WAN networks, produce disjoint data that constrains the AI's capacity to discover cross-domain relationships. All-encompassing, integrated, continuous data gathering is not a foundation to be rushed through. It is the foundation for all else [15].

Stage 2: Insights. As enough telemetry is collected and made available to a unified data layer, machine learning algorithms begin to detect patterns and anomalies that are impossible for humans to detect when looking at data through traditional dashboards. At this point, the network is not yet self-managed. But it is delivering intelligence, providing IT teams with a prioritised view of what is going on in the environment and why. The benefits so far are substantial. Rather than acting on a stream of undifferentiated alerts, IT teams are

Partners Universal International Research Journal (PUIRJ)
Volume: 05 Issue: 01 | January– March 2026 | ISSN: 2583-5602 | www.puirj.com

provided with contextual intelligence that explains what is (or is not) really a performance issue, what the likely causes might be, and so on. An engineer who might have taken two hours to identify the cause of a Wi-Fi performance issue in a particular part of the office can now be presented with a second-by-second diagnosis, and the historical pattern that shows exactly what changed and when [12, 13].

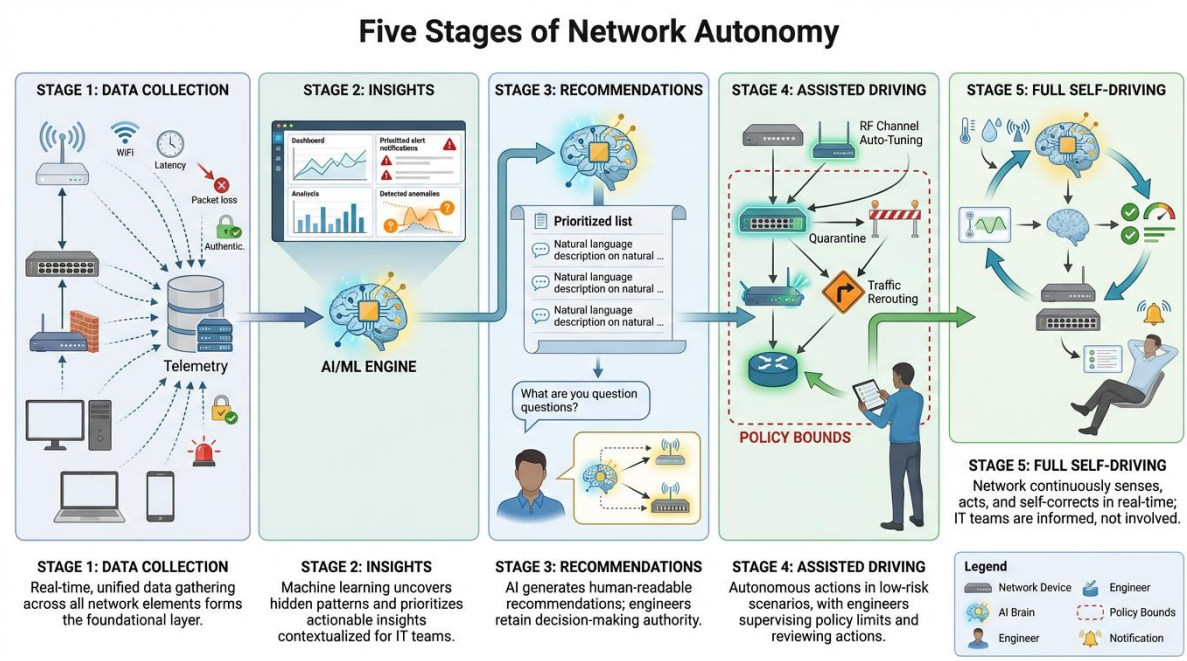

Fig -5: Five Stages of Network Autonomy

Stage 3: Recommendations. The move from diagnosis to recommendation is where the AI takes on an active role in the decision-making process. Having assessed the current conditions of the network and compared those against historical trends and the collective knowledge of the network system via the training data, the system will provide a prioritized list of recommended actions. The suggested actions are presented in natural language rather than configuration language, so that they can be understood by engineers of different skill levels. This is where tools such as conversational interfaces can be used. An IT engineer can ask the system why a particular set of users is seeing a slowdown in response time for a particular application and be presented with an explanation that reveals which factors are at play and what changes to the configuration should be made to resolve each one. The engineer is still in full control, but most of the thinking comes from the AI [14, 16].

Stage 4: Assisted Driving. This is the first stage to include autonomous action in a limited set of known, low-risk situations. Not only does the network recommend action, it takes it, with limits that IT staff specify in advance. Examples of actions that are appropriate at this stage include automatically shifting RF channel settings to eliminate interference, re-routing application traffic away from a congested WAN link, quarantining a device that shows behavior symptoms of a malware infection, or re-applying security policies to a new device type, based on a firmware-induced posture change. It's important to get right the human-in-the-loop model that applies at this stage. It doesn't mean that a human pre-approves all action that would defeat the purpose of speed. It does mean that human judgement sets policy bounds on the actions that an

autonomous system can take, and that the actions of the autonomous system are recorded and reviewable. Engineers go from operators to supervisors [1, 8].

Stage 5: Full Self-Driving. At the top of the maturity scale, the network automatically identifies conditions in its environment, reasons about what action to take, makes changes, tests the results and revises its strategy if the initial course of action is not successful. This is all done in real time, with no human intervention at any level. The IT team is merely notified of what the network has done and can look at its rationale, but they don't have to tell it what to do. This stage is not necessarily deployed across the entire network at any one time. Most companies that add self-driving operational capabilities will achieve full autonomy in isolated use cases, such as wireless coverage in a campus environment, before moving on. What's important is that Stage 5 is not a vision. It is a reality today for some specific use cases in organizations that have systematically evolved through the preceding stages [11, 18].

7. CORE TECHNOLOGICAL COMPONENTS

7.1 Unified AIOps and the Intelligent Data Layer

The analytical heart of any self-driving network architecture is AIOps, i.e. artificial intelligence that is used in IT operations. Its role is to accept the vast amount of telemetry produced by any current enterprise network and turn it into actionable intelligence timely enough to act on it and not to leave it merely in the past.

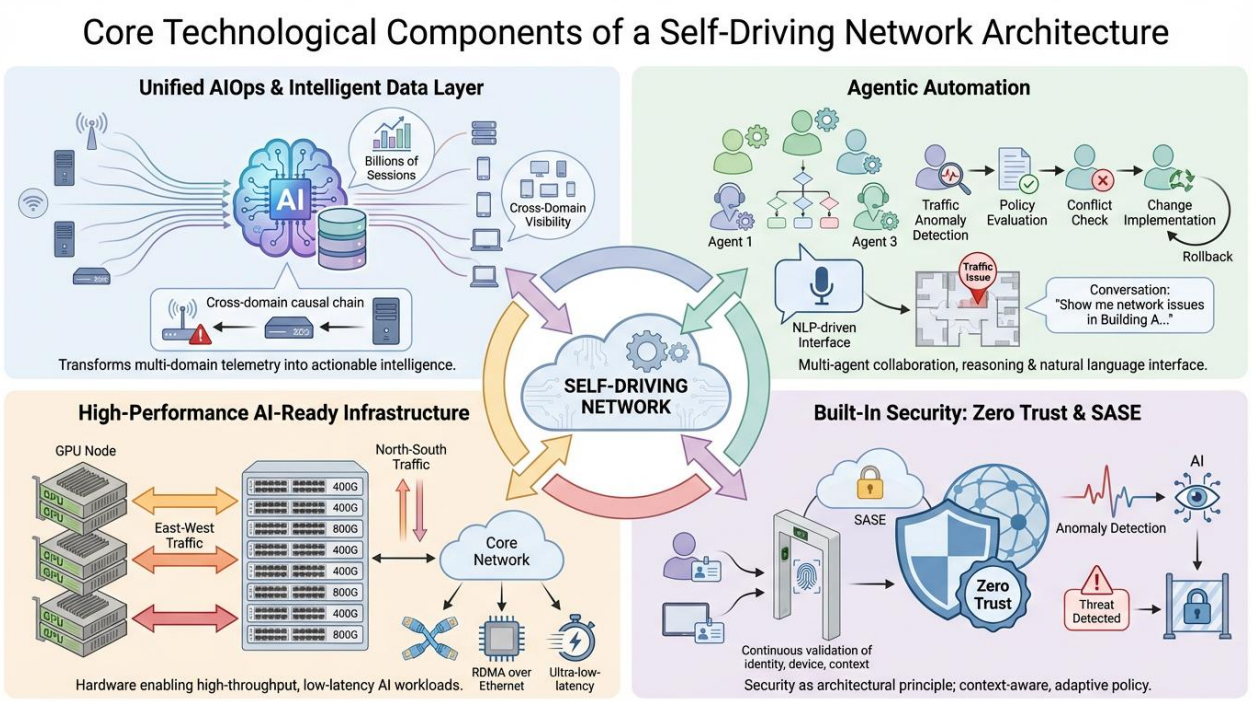

Fig -6: Core Technological Components of a Self-Driving Network Architecture

The difference between the mature AIOps systems and the previous network analytics systems is not necessarily the complexity of the algorithms, although it also counts, but rather the volume and richness of data that the AI has been trained on. The most able platforms have been consuming network telemetry over a decade or longer, and their databases contain billions of client sessions of millions of types of devices in

thousands of network environments. When a pattern that is deemed peculiar is observed on the network of a specific organization nowadays, the AI will be able to match it to a reference set of situations representing a vast array of real-life experiences and find the most probable explanation and the most successful response with a certain level of certainty that a system trained based on the local history data cannot offer [5, 8].

AIOps implementations are also done across network domains instead of being applied as independent solutions to segments of wireless, wired, WAN and data centers. It is hard to overestimate the importance of cross-domain visibility. Cross domain causal chain Many of the most influential and challenging-to-detect network problems consist of cross domain causal chains. An improperly set up core distribution switch can show itself to users as underperforming Wi-Fi, since the wired backhaul to the impacted access points is saturated. A cloud routing problem can manifest itself in the form of application latency on the users who are linked to Wi-Fi at the branch offices. In the absence of cross-domain correlation, these problems are difficult to diagnose without having the expertise of a specialist. They manifest automatically with unified AIOps [11, 13].

7.2 Agentic Automation Beyond Scripts to Reasoning

The most important and practically relevant technology of the last two years in the field of enterprise technology is agentic AI. It is a type of systems where independent software agents, with the ability to sense their surroundings, create goals, and execute series of commands towards those goals, work together to achieve tasks that are complex and multi-step.

An agentic system in the networking context does not merely execute a pre-programmed script when a condition is identified. It puts in place a team of specialized agents that work together. A single agent examines the traffic pattern which raised an alarm. A second estimates the security policy implication of alternative possible responses. The third checks that a proposed change in configuration would not be in conflict with existing policy constraints. Fourth implements the change and measures the result, and can roll back in case of validation failure. Such an organized multi-agent rationality helps the system to deal with the situation of real complexity, which cannot be handled by the single-agent or rule-based automation [4, 17, 19].

The human-facing part of agentic automation is the conversational interface layer, which is driven by natural language processing and large language models. It breaks down what would otherwise be a highly technical dynamism of engagement with complex systems into simple conversation. A non-wireless RF optimization expert who works as an engineer can request, in casual terms, of why a specific floor of a building exhibits poor wireless performance in the morning, and be presented with a structured and correct answer that utilizes multi-domain analysis. This is not a cosmetic aspect. It is a drastic shift towards the ability of anyone to be an effective manager of enterprise networks and how fast they can be managed [1, 3].

7.3 High-Performance AI-Ready Infrastructure

AI workloads are demanding in terms of network infrastructure, and more importantly, self-driving network management software needs a capable hardware to be deployed on. These are connected yet different needs and both are important. The AI workloads are fundamentally different to the traditional enterprise application traffic. Using large models requires enormous, prolonged data transfer between clusters of GPUs, creating so-called east-west traffic flows, that is, server-to-server traffic in a data center instead of the client-to-server north-south flows that most enterprise networks were optimized to. The complication of inference workloads is suggested by latency sensitivity. In cases where an AI model is being used to provide predictions to a production application, any small change in network latency can negatively impact the user experience or cause errors in time-sensitive operations.

Hardware implementation of these needs is in the form of high-density switching platforms with 400 Gigabit Ethernet and 800 Gigabit Ethernet port counts, ultra-low latency forwarding networks that reduce queueing delay, and implementation of protocols that are targeted at the requirements of a GPU cluster network Remote Direct Memory Access over Converged Ethernet, among others. They are architecturally differentiated products that are constructed to operate in the AI factory environment, and are not upgrades to an existing enterprise switching platform [5, 7, 10]. The difference between AI-ready infrastructure and traditional enterprise hardware is vast regarding their performance. Companies trying to execute AI training or inference workloads on the infrastructure optimized to deal with the traffic of traditional applications will experience bottlenecks that cannot be addressed by reconfiguring software. In this context, the hardware layer is a requirement and not an unattainable addition [8, 9].

7.4 Built-In Security Zero Trust and SASE as Architectural Principles

Self-driving networking architecture security is not a product that can be placed alongside networking infrastructure. It is a guideline that is enshrined in any architectural choice. This style has been termed as built in rather than bolted on and it is a valuable and bitter experience taught to us in the history of enterprise security. In the case of enterprise networking, the majority of security was implemented at the borderline. A network edge firewall verified traffic going in and out of the organization. The traffic that passed that check was considered to be safe. This model was clearly limited even at the time when the majority of users were in offices and most applications were on-premise. It proved completely ineffective once people started to be mobile, the apps were transferred to cloud services, and the notion of an outlined perimeter was virtually abolished [1, 13].

The architectural solution to that inadequacy is Zero Trust. With a Zero Trust model, no trust is ever supposed to be given based on location in the network. All access requests, whether initiated within the corporate network or by a remote connection, are considered based on the identity requesting access, health and configuration of the device requesting access, sensitivity of the resource being accessed, and the behavioral context of the session. Access is granted at the lowest level of privilege needed and it is continuously reviewed and not once at the beginning of the session [16, 18].

SASE applies Zero Trust to security services delivered by the cloud, establishing a single policy enforcement system that is consistent in all cases in a corporate office, a branch office, a home network, or a coffee shop. The practical value to the self-driving networks has been that security policy is a coordinated activity of the autonomous functionality of the network, and not an independent control plane, which has to be independently adjusted in response to changes in the configuration of the network. Security policy automatically applies and changes when the network autonomously changes [4, 5].

Behavioral analysis provided by AI will introduce a layer of dynamism in detecting threats. The system has a normal model of expected behaviour of all network devices and user identities. In case behavioral patterns do not match the behavioral patterns based on known attack patterns, compromise indicators, or policy violations, the system determines the anomaly, categorizes its likely cause, and notifies the security operations team, or triggers automated containment, depending on the policy and severity. Such functionality is significantly more efficient than signature-based detection of new threats, as it detects behavioral abnormalities instead of allowing an acquainted attack pattern to be detected in a signature repository [13, 17].

8. PRACTICAL BENEFITS FOR IT ORGANIZATIONS

Partners Universal International Research Journal (PUIRJ)
Volume: 05 Issue: 01 | January– March 2026 | ISSN: 2583-5602 | www.puirj.com

Self-driving network capabilities have operational effects that are tangible and can be measured in various areas.

The first advantage that is most apparent is that mean time to resolution of network incidents improves. In a case where the AI is able to detect the underlying cause of an issue and apply or suggest a remedy within seconds of detecting the precipitating conditions, incidents that once took hours of engineer effort to resolve, are cleared before most of the users are aware that they to begin with. Organizations that have established AIOps implementations have noted declines in the quantity of trouble tickets by up to 90 percent, which is more of a qualitative transformation of the way IT groups can utilize their time instead of an increase in efficiency [8, 19].

Automated onboarding and provisioning saves a lot of time to those whose organization has to operate at scale on all its networks. The addition of a new branch location, the introduction of a new access point or a new type of IoT device into a self-driving network setup entails the definition of a template policy and the automatic application of it by the system when new devices are added. The work that a network engineer might have to spend a day on site configuring devices, one at a time, can now be done remotely and automatically within minutes after the hardware has been connected [2, 9].

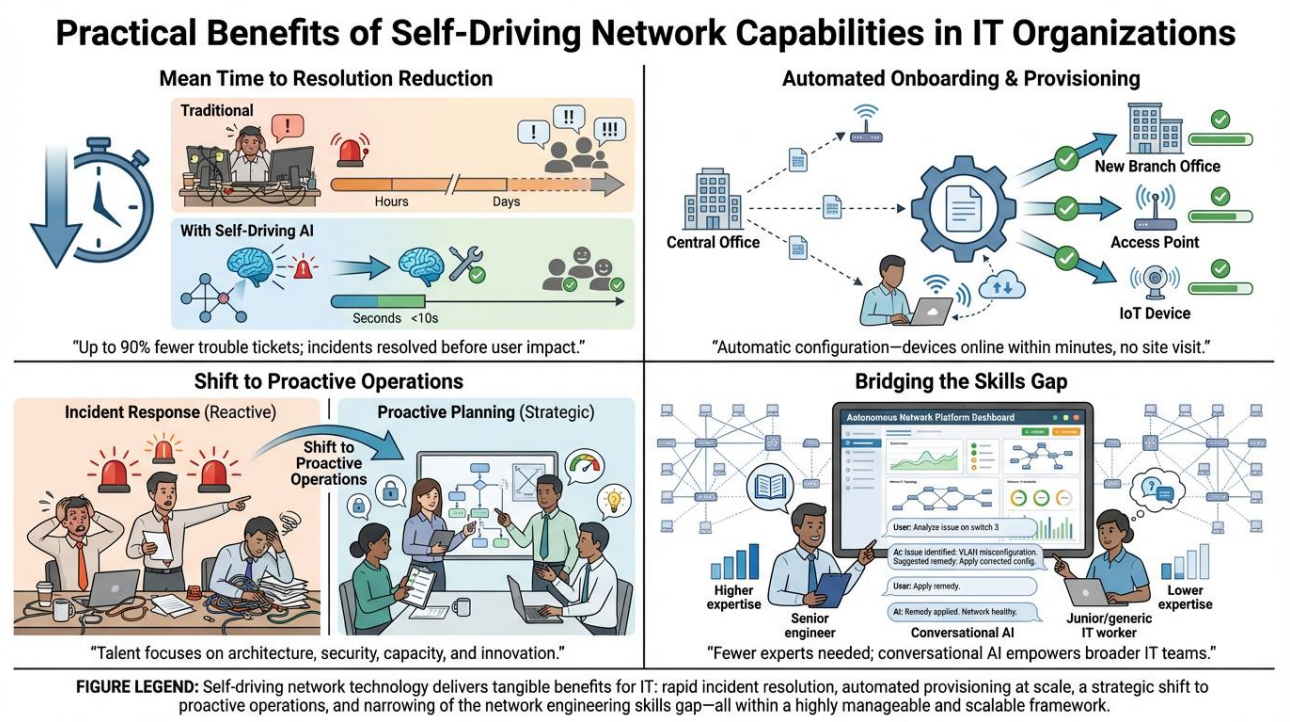

Fig -7: Practical Benefits of Self-Driving Network Capabilities in IT Organizations

Perhaps the most strategically important advantage is the change in operations towards the proactive. A team that is no longer wasting most of its time dealing with incident response is a team that can apply its capability to architecture augmentation, security patches, capacity planning, and to the assimilation of new functions that aid business aspirations. The value of that shift in the distribution of engineering talent cannot be precisely measured, however, and is consistently mentioned by those organizations that did the transition as one of the most quantifiable business effects [11, 18].

Partners Universal International Research Journal (PUIRJ)
Volume: 05 Issue: 01 | January– March 2026 | ISSN: 2583-5602 | www.puirj.com

The skills gap is an advantage that is increasingly gaining relevance with as the shortage of experienced network engineers is escalating. Autonomous network platforms not only lower the degree of expert knowledge needed to manage the daily activities of an organization but also allow organizations to manage environments of high complexity with fewer people. Conversational AI interfaces take this advantage one step further and enable engineers with general knowledge of IT to work on a broader variety of network problems without necessarily specializing in wireless RF theory or advanced routing protocols [4, 12, 17].

9. REAL-WORLD APPLICATIONS ACROSS SECTORS

The real-world application of self-driving network principles in various industries is solid evidence that the technology is a reality and not a dream.

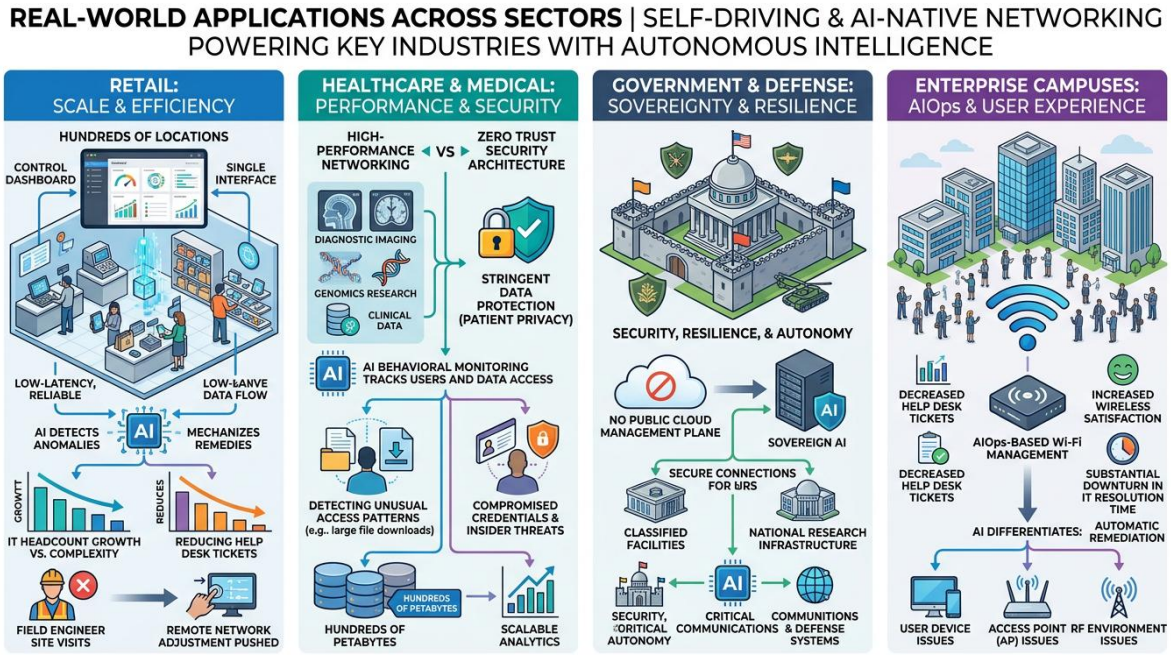

Fig -8: Real-World Applications Across Sectors

AI native networking platforms have also facilitated the centralized control of connectivity and security policies in hundreds of locations through a single interface in large-scale retail settings. The point-of-sale systems, inventory management software, and customer engagement software all require a reliable low latency network connectivity. The use of AI to identify anomalies and mechanize remedies has enabled retail IT teams to have their headcount (growth) rise in disproportion to the complexity of the environment, especially when it is geographically distributed. In the event whereby connectivity at a particular store is interfered with or changed due to configuration, the system detects and corrects this without necessarily having to send a field engineer to the location [8, 13].

In medical environments, high-performance networking combined with Zero Trust security architecture covers two, both important and traditionally non-reconcilable, needs: high performance connectivity to support diagnostic images, genomics research, and clinical data systems and stringent data protection mandated by patient privacy laws. The behavioral monitoring based on AI introduces an additional

protection level since it helps to detect unusual data access patterns that can signify the use of compromised credentials or insider threat actions. There have been reports of institutions utilizing AI-native infrastructure to increase the extent of their analysis without affecting their security posture, institutions that operate large volumes of research data, including those in the hundreds of petabytes range [14].

The particular use case of government and defense agencies, where the concepts of network security, resilience, and autonomy overlap with sovereignty needs, is an independent use case. On-premises operation of AI-native networks not based on public cloud management planes is essential to the classified facilities and national research infrastructure. The implementations of AI sovereignly in large research centers and military departments have shown that the principles of self-driving networks can be applied even in the most sensitive security context, including not only the business world but also the military domain [6].

The environments of enterprise campuses are among the most visible instances of the quantifiable performance enhancement brought about by the AIOps-based Wi-Fi management. Companies that have implemented AI-native wireless management systems record steady decreases in the number of help desk tickets associated with connectivity, as well as wireless satisfaction ratings, and a substantial downturn in the IT departments taking to resolve wireless performance problems. This is because the AI can differentiate between user device issues and access point issues and RF environment issues, and hence automatically direct one type of issue to the correct remediation line [5, 11].

10. CHALLENGES TO ADOPTION

Seeing the potential benefits of a self-driving network and actually adopting it are two different things, and we should be realistic about both.

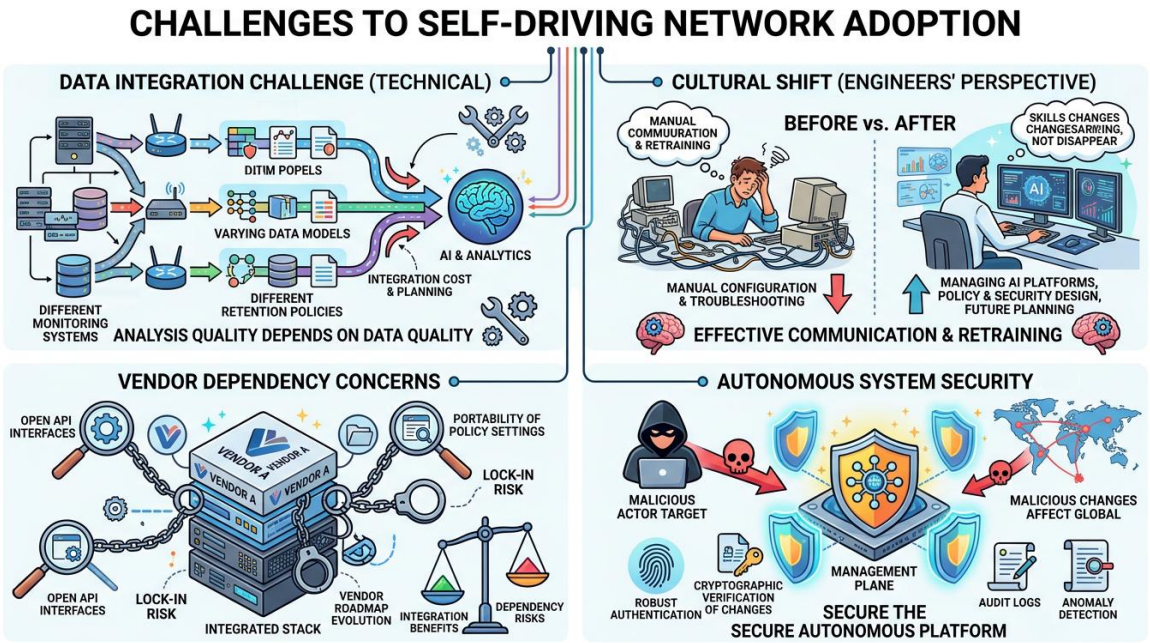

Fig -9: Challenges to Self-Driving Network Adoption

The single biggest technical challenge is data integration. The quality of the analysis that can be performed by a self-driving network platform can only be as good as the quality of the telemetry data available. Companies that maintain monitoring systems for different aspects of their network, each with its own data models and policy for data retention, will find that considerable work is required to integrate the various sources into a common data layer before the AI can be used effectively. This is not to say that it is not worth starting now, but it is worth planning for the integration cost [12, 13].

Culture is another challenge that is at least as important as technology. Professional engineers who have developed expertise in manual configuration of network devices and troubleshooting may perceive self-driving networks as taking away rather than adding to the skills they have developed. This is a legitimate concern and should be addressed. The truth is that self-driving networks do alter the tasks network engineers perform. The need for skills in manual configuration and troubleshooting decreases while skills in managing AI platforms, policy design, security design and planning for the future increase. Companies that effectively communicate the change and invest in retraining their staff as they roll out new technology will be better positioned for success than those that treat the change as purely a technology project [9, 17].

Vendor dependency is a legitimate concern. The most effective self-driving network systems are integrated stacks of hardware, management software, artificial intelligence for network analytics and security services, all from the same vendor ecosystem. Integration is often the key to the capabilities businesses desire. But it results in a level of vendor dependency that should be assessed. The openness of API interfaces, the portability of policy settings and the likely evolution of the vendor's roadmap are all part of due diligence [4, 19]. Security of the autonomous system is not simple. A platform that can change network configuration is a valuable target for malicious actors. If an attacker is able to take control of the management plane of a self-driving network, then they can make changes to the network at scale and in a timely fashion. Robust authentication to access the platform, cryptographic verification of network changes, audit logs and anomaly detection of the management plane are all critically important elements of a secure self-driving network [16, 18].

11. FUTURE PROSPECTS

11.1 Where This Technology Is Going

The future of self-driving network technology is clear, if not the pace and route.

Agentic AI abilities will advance considerably over the next 3-5 years. The range of network problems that can be solved by self-driving systems without human intervention will broaden as reasoning-capable large language models advance and multi-agent orchestration capabilities mature. Issues currently requiring HITL oversight because they involve tradeoffs that are difficult to formalise as policy rules will be solvable by reasoning AI agents. The net result will be another expansion of the range of network operations not requiring human intervention [9, 19].

Networking and security will converge. Zero Trust and SASE architectures will be mainstream rather than cutting-edge in enterprise networking in the next three to five years due to both regulatory and practical pressures of perimeter-based security in a fully distributed enterprise environment and the cost advantage of integrated management and enforcement of network security policies rather than separate toolsets for network management and security. The next generation of self-driving networks will have security policy enforcement baked in and not bolted on [5, 13].

Partners Universal International Research Journal (PUIRJ)
Volume: 05 Issue: 01 | January– March 2026 | ISSN: 2583-5602 | www.puirj.com

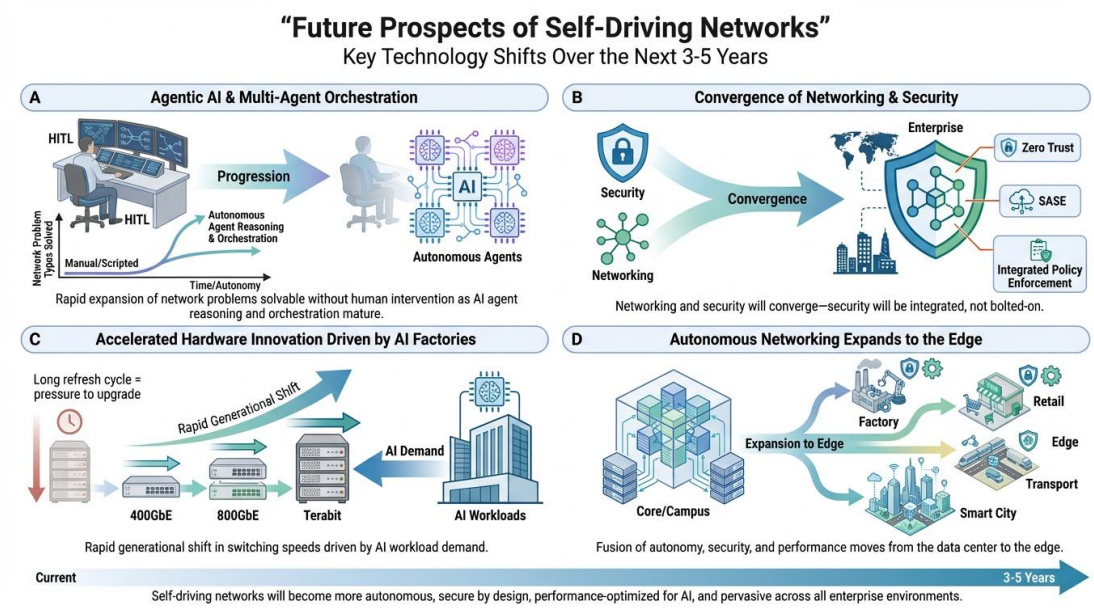

Fig -10: Prospects of Self-Driving Network

Innovation in hardware will be driven at an accelerated pace by the demands of AI factories. The move from 400GbE to 800GbE to new terabit switching technologies will be the most rapid generational shift in enterprise networking history driven by the need for larger training clusters and higher inference throughput for AI workloads. Companies that extend refresh cycles will be unable to support the demands of AI workloads, so they will be under pressure to upgrade even if they have longer refresh cycles [7, 8, 10]. The extension of autonomous networking from the core and campus to the edge is another key near-term trend. As AI inference workloads are pushed closer to the edge of the data center, factories, retail, transport and smart cities will need to adopt the same fusion of performance, autonomy and security as enterprise campus and data center environments are doing today [11, 18].

12. THE COMPETITIVE LANDSCAPE, OPEN STANDARDS, AND VENDOR-NEUTRAL EVALUATION

12.1 Navigating the Market Who Builds Self-Driving Networks and How to Compare Them

For any enterprise considering self-driving network capabilities in 2016, the market is full of several serious vendors with considerable resources and different architectural approaches, capabilities and strategies. We need to understand that landscape in its own right, rather than with a single-vendor lens, to make investment decisions that are in the best interests of the enterprise rather than the vendor.

12.2 The Major Players

Cisco Systems is by far the world's largest enterprise networking vendor. Its AI networking portfolio has evolved considerably. Cisco Catalyst Center (formerly DNA Center) offers AIOps and automation for the campus and branch. ThousandEyes offers end-to-end network visibility across internet, cloud and enterprise networks, and provides the multi-domain visibility needed for self-driving networks. Cisco's AI Network Analytics capability correlates telemetry across wireless, wired and SD-WAN, and its integration with other Cisco security products such as Cisco Secure Access and Duo makes it a viable Zero Trust solution.

Partners Universal International Research Journal (PUIRJ)
Volume: 05 Issue: 01 | January– March 2026 | ISSN: 2583-5602 | www.puirj.com

HPE's 2024 acquisition of Juniper Networks brought together the Mist AI platform and its AI-native wireless and wired management systems with HPE's data center and edge networking systems. Mist AI is well known for its conversational AI interface, Marvis, and the depth of the portfolio makes it a formidable rival across the campus, branch and data center.

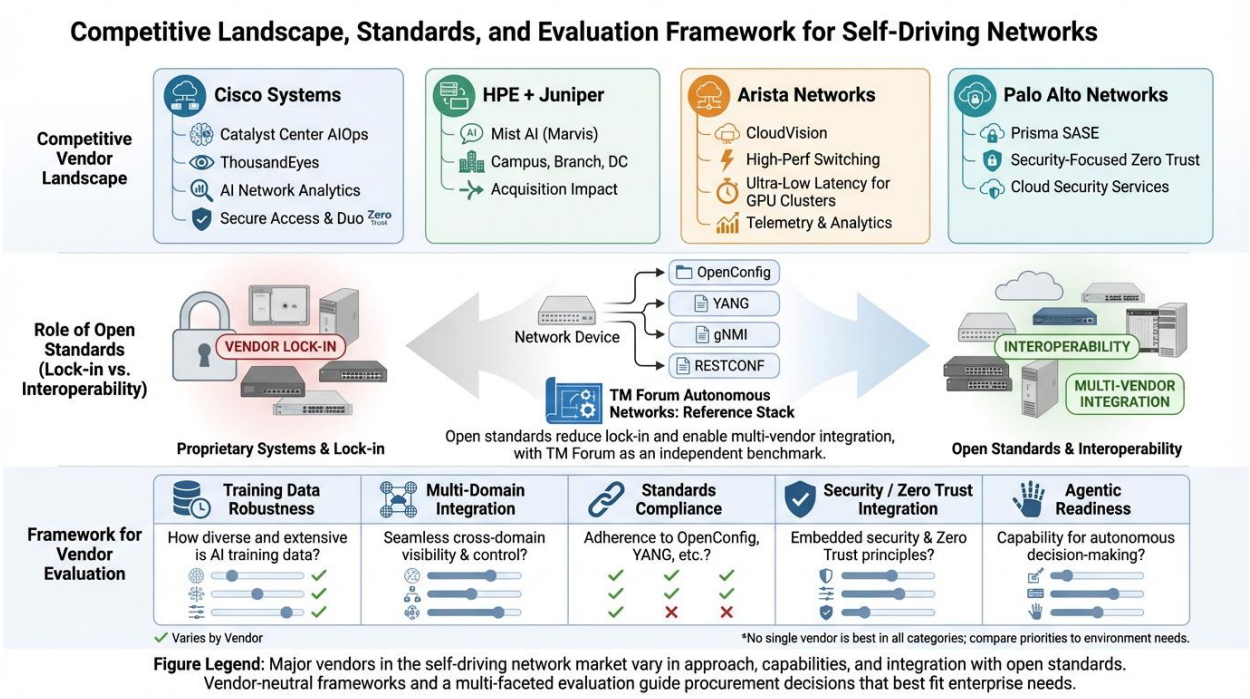

Fig -11: Framework for Self-Driving Networks

Arista Networks is a leading data center and AI factory vendor. Its CloudVision platform offers network-wide automation, telemetry and analytics for high-performance switching networks. For the specific networking needs of GPU clusters, Arista's expertise in ultra-low latency fabric design is a key value proposition for organizations looking to build or scale AI training capacity.

Palo Alto Networks, as a security company, has made a claim for its Prisma SASE as a network security product that absorbs many of the functions of network management. For those with a strong focus on building secure networks, its unified Zero Trust and cloud-based security services approach is a viable alternative.

12.3 The Role of Open Standards

There are concerns about vendor lock-in in this area, which is partially addressed by the maturity of open standards. OpenConfig provides a vendor-neutral network device configuration and state data model. YANG modeling language, gNMI for telemetry and RESTCONF for configuration allow organizations to develop their own monitoring and automation tools that, theoretically, support multiple vendors' devices. How well specific vendors support these standards is variable and should be tested as part of the procurement process. TM Forum's Autonomous Networks framework, a vendor-neutral maturity model created by a global network of network operators and vendors, offers a vendor-independent reference architecture for autonomous

networking. When evaluating self-driving network platforms, companies gain insight by comparing vendor claims against this independent model, rather than vendor-asserted levels of maturity.

12.4 A Framework for Vendor Evaluation

Comparing approaches to self-driving network platforms requires the following considerations.

First, depth and breadth of training data: how long has the vendor's AI been consuming production network telemetry, and how many different types of deployment environments.

Second, multi-domain integration: is there a holistic view and autonomous management across wireless, wired, WAN, and data center, or is it good in some domains and not others.

Third, standards compliance: to what extent is the platform compliant with OpenConfig, gNMI and other non-proprietary interfaces.

Fourth, security integration: is Zero Trust enforcement integrated natively or through third party integration that comes with extra management overhead.

Fifth, agentic readiness: what types of multi-step proactive actions can the system execute today, with customer references, and what is planned for the future with specific timelines.

No vendor will be equally strong across the five areas. Knowing where each is best and where they may have compromises is the first step towards a procurement decision that meets the needs of the specific environment, risk posture and priorities.

13. A PRACTICAL ROADMAP FOR ORGANIZATIONS BEGINNING THE TRANSITION

For those organizations convinced of the benefits of self-driving networks but unsure how to get started, a roadmap for transition offers a way forward.

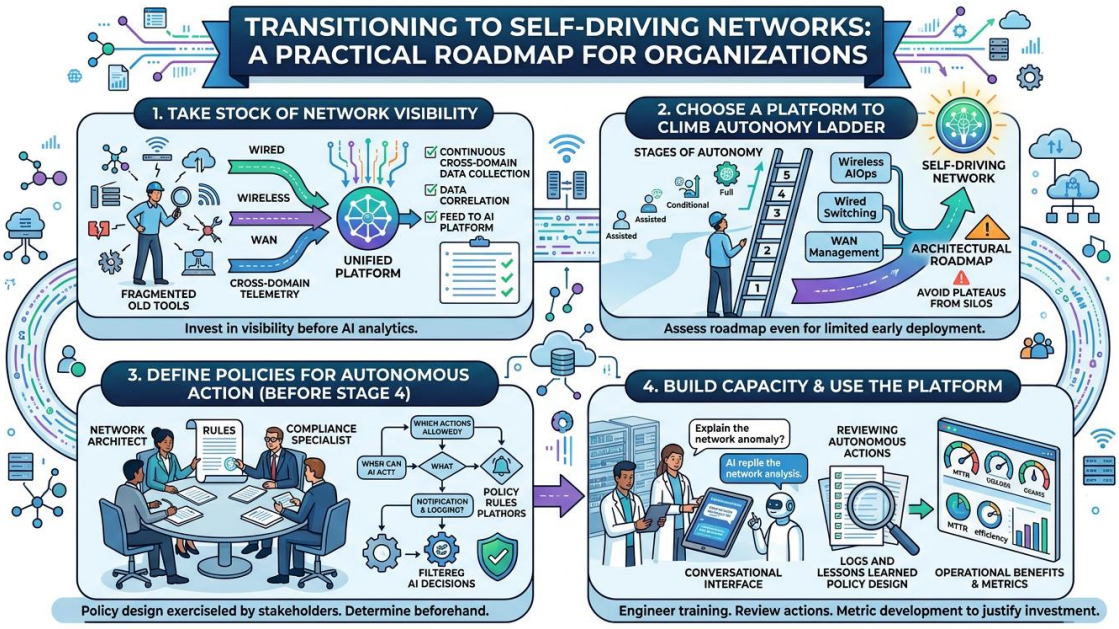

Fig -11: Transitioning To Self-Driving Networks

The first is to take stock of the level of network visibility. It is not whether the company has monitoring tools, it probably does, but whether those tools collect, correlate and provide continuous, cross-domain telemetry to an AI platform that can use it. If not, the first step should be to invest in a unified platform capable of doing so, before investing in AI-based analytics on top of it.

The second step is to choose a platform that can climb the ladder of the five stages of autonomy, rather than focusing on the next immediate need. An enterprise that procures a wireless AIOps solution, for instance, without an expectation of how that solution will sit with wired switching and WAN management analytics will reach a plateau on the benefits it can derive. It is important to assess the architectural roadmap of platform candidates even if the near-term deployment will be limited.

The third step is to define policies for autonomous action before enabling Stage 4. Which actions the AI should be allowed to take autonomously, when, and with what notification and logging should be determined by a policy design exercise led by network architects, security specialists and in some cases, compliance specialists. This is best done before enabling autonomous action rather than having to resolve differences in opinion about what is or is not appropriate after the AI has done something unexpected.

The fourth step is building the capacity to use the platform, not just build it. This includes equipping network engineers with training on how to use the conversational interface and AI analytics capabilities, processes for reviewing autonomous actions and incorporating lessons learned into the design of policies, and developing metrics that capture the operational benefit of the autonomous capabilities over time so that the investment can be assessed and justified into the future.

14. CONCLUSION

14.1 The Strategic Imperative

Self-driving networks are not a technology under development, or a vendor-invented category with no proof of concept. The building blocks, robust AIOps solutions, agentic engines of automation, AI-ready switch and router hardware, and integrated Zero Trust security controls are in use today across multiple industries. They are delivering tangible improvements in incident count, time to resolve incidents, security, and a reallocation of engineering time from operations to strategy [1, 11, 18]. The five stages of autonomy provide a viable roadmap. It is not necessary to reach full autonomy to benefit from the process. Each stage of progress brings its own advantages and moving through the stages in sequence builds technical capability and an organisational comfort with the technology that makes further progress possible.

The strategic business case is clear and growing stronger. The complexity of enterprise networks is outstripping manual management. Manual management skills are in short supply and costly. The likely cost of network incidents, in terms of lost productivity, damaged customer relationships, regulatory risk and loss of market share, is increasing. Given this, the choice to operate networks primarily as human-focused reactive operations is not a passive one. It's a decision to increase operational debt, compounded. Those that start the planned shift to self-driving network management now are going to be better placed to massively scale AI workloads, and to be secure in a hostile network environment while releasing human talent to take on the tasks that can only be done by humans. The longer they wait, the more difficult it will be for them to catch up. The network has traditionally been infrastructure. In the AI age, the smart, autonomous network is a weapon. Those that understand that difference early, and respond accordingly, will have a competitive edge that will be felt beyond their IT organizations [2, 8, 17].

REFERENCES

[1] Hewlett Packard Enterprise. HPE Networking: AI-Native Architecture and Self-Driving Network Operations Overview. Available at: https://www.hpe.com
[2] Hewlett Packard Enterprise. GreenLake Platform and Hybrid Cloud Simplification. Available at: https://www.hpe.com
[3] Hewlett Packard Enterprise. Marvis Virtual Network Assistant: Conversational AI for IT Operations. Available at: https://www.hpe.com
[4] Atlaxion. "The Rise of Self-Driving Networking: How AI is Automating Enterprise IT." Atlaxion, December 4, 2025. Available at: https://atlaxion.com
[5] Juniper Networks, now part of HPE. Mist AI-Native Networking Platform: AIOps, Marvis AI, and AI Data Center Solutions. Available at: https://www.juniper.net
[6] GovEvents. Oracle Federal Forum: AI and Operational Efficiency in Federal Agencies. Available at: https://www.govevents.com
[7] Hewlett Packard Enterprise. HPE Alletra Storage and AI Factory Infrastructure Portfolio. Available at: https://www.hpe.com
[8] theCUBE Research. "Operationalizing AI at the Edge." Bob Laliberte. March 25, 2026; "NVIDIA's GTC Focuses on Full-Stack AI Factories, Inferencing, Agents, and Tokenomics." Bob Laliberte. March 20, 2026. Available at: https://thecuberesearch.com
[9] ITTech Pulse Staff Insight. "Build vs Buy: Should Enterprises Develop or License Domain-Specific Language Models?" ITTech Pulse, March 25, 2026. Available at: https://ittech-pulse.com
[10] Hewlett Packard Enterprise. HPE Private Cloud AI and AI Factory for Sovereign and At-Scale Deployments. Available at: https://www.hpe.com
[11] Hewlett Packard Enterprise. AIOps and the Journey from Reactive to Proactive Network Management. Available at: https://www.hpe.com
[12] TechTarget. "Domo Doubles Down on AI with Latest Platform Additions." TechTarget, March 25, 2026; "U.S. Federal AI Framework Deemed Aspirational, Noncommittal." March 24, 2026. Available at: https://www.techtarget.com
[13] NCS Singapore. Hybrid AI, Data Infrastructure, Cyber Security Managed Services, and Intelligent Transport Solutions. Available at: https://www.ncs.co/en-sg/
[14] IgMin Research. Dubey, S. "From Test Case Design to Test Data Generation: How AI is Transforming End-to-End Quality Assurance in Agile and DevOps Environments." IgMin Research, February 3, 2026. DOI: 10.61927/igmin331. Available at: https://www.igminresearch.com
[15] Hewlett Packard Enterprise. HPE Aruba Networking Central: Real-Time Telemetry and Network Visibility. Available at: https://www.hpe.com
[16] ManageEngine. IT Operations Management and Observability: AIOps, Security Information and Event Management, and Identity and Access Management. Available at: https://www.manageengine.com
[17] Hewlett Packard Enterprise. HPE AI Security Tools: Reducing Risk and Boosting Enterprise Resilience Across Cloud, Core, and Edge. Available at: https://www.hpe.com
[18] Hewlett Packard Enterprise. Zero Trust and SASE Principles in AI-Native Network Security Architecture. Available at: https://www.hpe.com
[19] Juniper Networks, now part of HPE. "HPE Accelerates Self-Driving Network Operations with New Mist Agentic AI-Native Innovations." Juniper Networks Press Release. Available at: https://www.juniper.net