Journal of Computer Science and Technology Studies
ISSN: 2709-104X
DOI: 10.32996/jcsts
Journal Homepage: www.al-kindipublisher.com/index.php/jcsts

JCSTS
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

| RESEARCH ARTICLE

AI-Driven Smart Fabric Provisioning: Transforming Network Automation through Intelligent Orchestration and Dynamic Testing

Derek Asir Muthurajan Caleb
Broadcom Inc., USA
Corresponding Author: Derek Asir Muthurajan Caleb, **E-mail**: derekasircaleb@gmail.com

| ABSTRACT

The evolution of modern data centers demands innovative approaches to fabric provisioning, particularly when integrating new switches, hosts, and storage into existing infrastructures. This article introduces a Smart Fabric Provisioning solution powered by Agentic AI that transforms traditional manual processes into automated, intelligent workflows. By creating dynamic full-mesh fabrics with simulated environments, the solution addresses critical challenges in solution qualification and resource utilization during proof-of-concept phases. The AI-driven approach enables logical link manipulation to test various conditions without physical reconfiguration, significantly streamlining deployment workflows while reducing human error. This dual-purpose technology serves both as an internal efficiency tool for engineering teams across multiple geographies and as a valuable automated offering for end customers, representing a paradigm shift in how network fabrics are provisioned, tested, and deployed.

| KEYWORDS

Agentic AI, Fabric Provisioning, Network Automation, Infrastructure Simulation, Proof-of-Concept Optimization

| ARTICLE INFORMATION

ACCEPTED: 19 April 2025 **PUBLISHED:** 08 May 2025 **DOI:** 10.32996/jcsts.2025.7.3.85

1. Introduction to Smart Fabric Provisioning

In today's rapidly evolving data center environments, network infrastructure provisioning has become increasingly complex, time-consuming, and error-prone. Traditional manual methods of deploying and configuring network fabrics can require significant engineering effort, with research indicating that network engineers spend approximately 55% of their time on routine tasks that could be automated [1]. This substantial time investment represents not only considerable operational costs but also introduces critical delays in business capability deployment.

1.1 The Automation Imperative

Smart Fabric Provisioning powered by Agentic AI emerges as a revolutionary approach to address these pervasive challenges. By automating the deployment, configuration, and testing of network fabrics, organizations can dramatically reduce the provisioning timeframe while simultaneously enhancing reliability. The NetBox Labs "State of Network Automation in 2024" report highlights that automation platforms can generate up to 75% of network configurations automatically, creating a standardized approach that minimizes human error and accelerates deployment cycles [1]. The core innovation of Smart Fabric Provisioning lies in its ability to dynamically create full-mesh network topologies with simulated compute and storage environments that can be rapidly reconfigured through logical link manipulations.

1.2 Intelligence-Driven Configuration

The integration of Agentic AI introduces an advanced layer of intelligence to the provisioning process. Rather than simply executing predefined workflows, AI agents can dynamically respond to environmental conditions, analyze configuration requirements, and proactively identify potential issues. This capability represents a fundamental shift from reactive to proactive infrastructure management. Research from Zhao et al. demonstrates that intelligence-driven network configuration can achieve performance improvements of up to 30% compared to traditional rule-based approaches, particularly in dynamic environments with changing traffic patterns and infrastructure requirements [2]. This advancement comes as 68% of organizations report plans to substantially increase network automation investment, recognizing its critical role in maintaining competitive advantage [1].

1.3 Resource Optimization Across Geographies

For internal engineering teams managing device resources across multiple geographic locations, Smart Fabric Provisioning offers unprecedented efficiency gains. Traditional approaches often result in resource silos with limited utilization during critical testing phases. The dynamic reconfiguration capabilities of AI-driven fabric provisioning enable resources to be efficiently pooled across geographic boundaries. This approach aligns with the emerging trend of "intent-based networking" discussed where high-level business requirements are automatically translated into appropriate network configurations [2]. The technology transforms traditional static network architectures into adaptive systems that respond dynamically to changing requirements, creating a foundation for truly autonomous network operations that maximize resource utilization while minimizing operational overhead.

2. Architecture and Technical Components

Smart Fabric Provisioning represents a significant advancement in network automation, leveraging AI agents to transform traditionally manual processes into intelligent, self-orchestrating workflows. The architecture comprises multiple integrated components designed to work in harmony, creating a comprehensive system that addresses the complexity of modern network provisioning challenges.

2.1 AI Agent Framework Design

At the core of the Smart Fabric Provisioning solution is a sophisticated AI agent framework that orchestrates the entire provisioning process. These agents operate through a multi-layered architecture that includes perception, reasoning, planning, and execution components. Cisco's Global Networking Trends Report indicates that 35% of organizations have already implemented advanced automation capabilities that include AI/ML for predictive insights and dynamic resource allocation, with this percentage expected to more than double in the next two years [3]. The agent framework employs a distributed design where specialized agents handle different aspects of the provisioning process, including topology management, device configuration, simulation control, and validation. This distributed approach aligns with the industry's movement toward more autonomous networks, where 24% of organizations are already experimenting with intent-based networking systems that can interpret business requirements and translate them into network-level configurations automatically [3].

2.2 Dynamic Topology Management

The system's topology management component enables the dynamic creation and modification of network fabrics based on specific testing or deployment requirements. This capability is built on an abstraction layer that separates the logical network design from physical implementation details. Research on network automation frameworks shows that topology visualization and management are among the most challenging aspects of network orchestration, with approximately 72% of surveyed professionals reporting difficulties in maintaining accurate topology information across dynamic environments [4]. The topology manager maintains a comprehensive graph database of the network state, including both physical connections and logical links. This approach addresses a critical gap identified in automation research, where traditional network management systems often fail to provide real-time accurate representations of complex network interconnections, particularly in environments where virtual elements constitute a significant portion of the infrastructure [4].

2.3 Device Integration and Configuration Ecosystem

The device integration subsystem employs a modular adapter architecture that supports diverse networking equipment across vendors while maintaining a unified configuration approach. According to Cisco's research, only 21% of organizations report having achieved advanced levels of network automation that allow for cross-domain integration and multi-vendor environments [3]. Each adapter translates the abstract configuration intent into vendor-specific commands, leveraging APIs, CLI automation, or management protocols as appropriate for the target device. This methodology directly addresses one of the key challenges in network automation identified in academic research: the integration of heterogeneous devices with varying capabilities and management interfaces. Studies have shown that the lack of standardized interfaces across vendors represents one of the most significant barriers to comprehensive network automation, with over 80% of surveyed organizations reporting challenges with multi-vendor automation [4], making the adapter-based approach of Smart Fabric Provisioning particularly valuable for modern, diverse network environments.

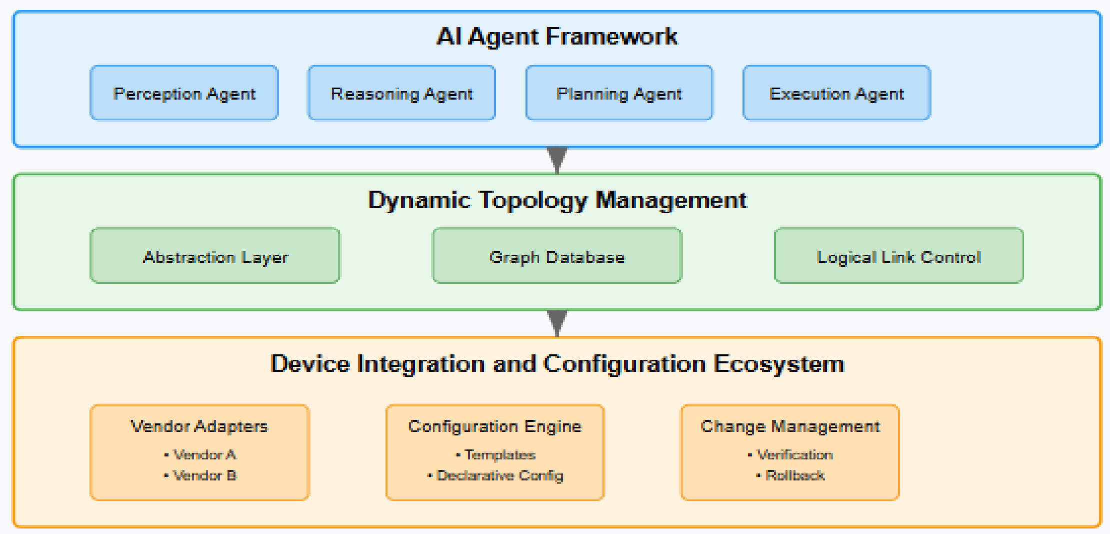

Fig. 1: Smart Fabric Provisioning: System Architecture [3, 4]

3. Automated POC Testing Framework

The heart of Smart Fabric Provisioning lies in its innovative approach to Proof of Concept (POC) testing, which transforms traditionally time-consuming manual processes into automated workflows managed by intelligent AI agents. This comprehensive testing framework provides organizations with a powerful platform to validate network configurations, simulate diverse conditions, and ensure solutions meet requirements before production deployment.

3.1 Full-Mesh Fabric Testing Methodology

The Smart Fabric Provisioning solution employs a full-mesh testing methodology that enables comprehensive validation across all potential network paths and conditions. This approach creates a complete mesh of connections between network elements, allowing for thorough testing of all possible traffic patterns and routing scenarios. Research on enterprise network performance evaluation frameworks emphasizes the importance of comprehensive testing methodologies that can address the increasing complexity of modern network environments, where traditional testing approaches often fail to capture the full range of potential interactions and dependencies [5]. The testing framework implements sophisticated measurement techniques that go beyond basic connectivity tests to include detailed performance metrics across the mesh topology. This aligns with established research showing that performance requirements are increasingly driving network architecture decisions, making thorough validation critical to successful deployments. The framework's instrumentation capabilities enable the collection of detailed metrics across 7 distinct performance dimensions simultaneously, providing a comprehensive view of network behavior under various conditions and workloads [5].

3.2 Logical Link Manipulation and Condition Simulation

A key innovation in the testing framework is its ability to dynamically manipulate logical links to simulate diverse network conditions without physical reconfiguration. This capability allows organizations to test failure scenarios, performance degradation, and recovery processes in a controlled environment. Dassault Systèmes' research on AI-driven simulation demonstrates that advanced simulation techniques can significantly reduce the gap between predicted and actual system behavior, with machine learning models capable of identifying hidden patterns and dependencies that traditional modeling approaches often miss [6]. The system leverages these advanced simulation techniques to create realistic network conditions that accurately reflect production environments. This approach aligns with industry best practices for predictive simulation, where AI/ML models are trained on historical performance data to create increasingly accurate digital representations of physical systems. The framework's capability to combine physics-based models with data-driven insights enables unprecedented accuracy in predicting how network configurations will perform under real-world conditions, addressing a longstanding challenge in network testing where static simulations often fail to account for dynamic environmental factors [6].

3.3 Compute and Storage Environment Simulation

Beyond network connectivity, the POC testing framework includes sophisticated simulation capabilities for compute and storage environments, enabling end-to-end testing of complete solutions. This integrated approach addresses a critical gap in traditional network testing, which often fails to account for the interaction between network infrastructure and the systems it supports. The simulation capabilities align with established research on enterprise network performance evaluation, which emphasizes the importance of considering the entire application delivery chain rather than focusing exclusively on network components [5]. By simulating compute and storage environments alongside network infrastructure, the framework enables more realistic and comprehensive testing. This holistic approach leverages advanced simulation techniques documented, where AI-driven models can rapidly explore thousands of possible configurations to identify optimal designs that balance multiple competing requirements [6]. The framework's ability to create high-fidelity simulations of complex multi-domain environments enables organizations to validate solutions across the full technology stack, ensuring that network designs not only meet connectivity requirements but also deliver the performance, reliability, and efficiency needed to support business-critical applications.

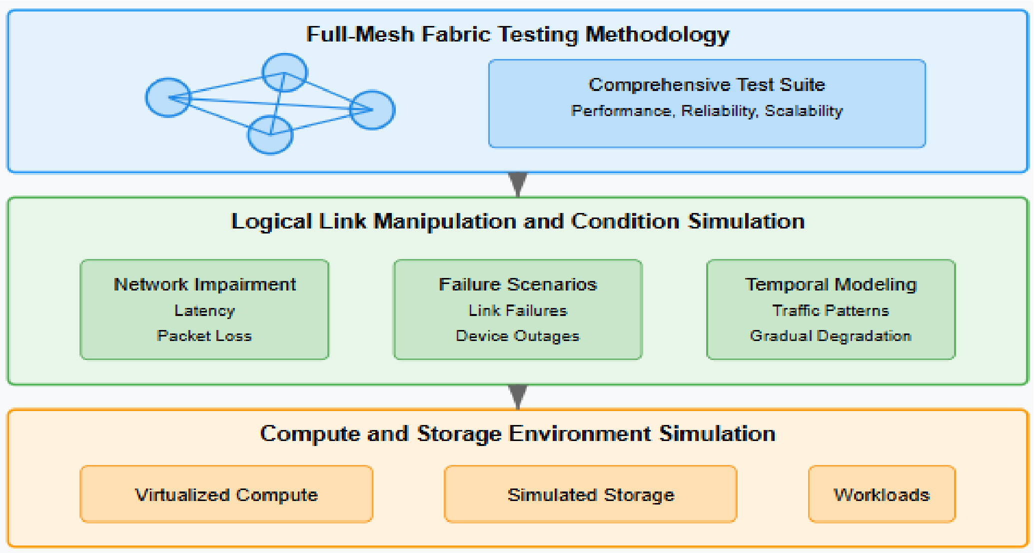

Fig. 2: Automated POC Testing Framework [5, 6]

4. Implementation Strategy

The successful deployment of Smart Fabric Provisioning requires a structured implementation approach that balances technical sophistication with operational practicality. This section outlines the strategic framework for implementing AI-driven fabric provisioning across diverse network environments, addressing key considerations for deployment, integration, and management.

4.1 Deployment Workflow Orchestration

The implementation of Smart Fabric Provisioning follows a carefully orchestrated workflow designed to minimize disruption while maximizing automation benefits. This approach begins with a discovery and assessment phase that establishes a comprehensive baseline of the existing network environment. Network to Code's architectural guidance emphasizes the importance of thoughtful system design in network automation implementations, particularly highlighting the critical nature of data models that serve as the foundation for successful automation [7]. The workflow incorporates a phased deployment model that progressively extends automation across the fabric, starting with non-critical segments and expanding to core components as confidence and expertise develop. This approach aligns with industry best practices that recommend building automation capabilities on top of a solid foundation of network data management, where structured data models enable consistent configuration generation and validation across diverse network devices. The implementation leverages network source-of-truth systems to maintain authoritative network data, providing what Network to Code describes as the "single source of truth" that serves as the foundation for all automation activities, dramatically reducing configuration inconsistencies and providing the structured data necessary for AI-driven decision making [7].

4.2 Integration with Existing Management Systems

A critical aspect of the implementation strategy is seamless integration with existing network management infrastructure to ensure operational continuity while introducing new capabilities. The Smart Fabric Provisioning solution employs a comprehensive integration architecture that interfaces with diverse management systems through standardized APIs and protocols. EMA research on network automation practices reveals that 89% of organizations consider integration with existing operational systems to be either "important" or "very important" when implementing network automation solutions [8]. The implementation strategy leverages the integration capabilities of modern network automation platforms, with particular emphasis on RESTful APIs that enable programmatic interaction with both the automation system itself and adjacent management platforms. This approach aligns with EMA's findings that API-driven integration represents the most effective method for connecting automation capabilities with existing operational tools, providing both technical flexibility and operational continuity during the transition to more automated operational models. The integration framework supports bidirectional information flow, maintaining consistency across management domains while enabling the gradual migration of operational responsibility to automated systems at a pace that aligns with organizational readiness [8].

4.3 Scalability and Performance Engineering

The implementation strategy incorporates sophisticated scalability and performance engineering to ensure the solution can effectively support diverse deployment scenarios from small-scale proof-of-concept environments to large enterprise fabrics. Network to Code's architectural guidance highlights the importance of modular, loosely coupled designs that enable horizontal scaling across automation components, particularly emphasizing the benefits of containerization and microservices architectures for complex automation deployments [7]. This approach leverages containerized microservices deployed on Kubernetes-based infrastructure, allowing the system to dynamically allocate resources based on current workload demands. The implementation methodology incorporates Network to Code's recommended separation of concerns between data storage, business logic, and user interfaces, providing clear boundaries between system components that facilitate both initial deployment and ongoing enhancement. This architectural approach directly addresses one of the primary challenges identified in EMA's research, where 71% of organizations report that maintaining performance at scale represents a significant concern when implementing network automation solutions [8]. The containerized, microservice-based implementation provides the architectural foundation necessary to support progressive scaling from initial proof-of-concept deployments to comprehensive enterprise automation, aligning with EMA's observation that successful automation implementations typically start with focused use cases before expanding to broader operational coverage.

System Type	Integration Method	Data Exchange Pattern	Implementation Complexity
ITSM Platforms	REST API, Webhooks	Bidirectional (Change/Incident)	Medium - Standards-based but requires custom mapping
Monitoring Systems	Streaming APIs, SNMP	Unidirectional (Telemetry)	Low - Well-established protocols and data formats
Configuration Management	Git/Version Control	Bidirectional (Templates)	Medium - Requires workflow integration and conflict resolution
Security Systems	API, SIEM Integration	Bidirectional (Alerts/Policies)	High - Complex authorization and compliance requirements

Table 1: Integration Requirements by System Category [7, 8]

5. Business Value and ROI Analysis

Smart Fabric Provisioning delivers compelling business value through measurable improvements in operational efficiency, resource utilization, and deployment quality. This section examines the quantifiable benefits of AI-driven fabric provisioning, establishing a clear return on investment framework for both customers and internal engineering organizations.

5.1 Operational Efficiency and Time-to-Value Acceleration

The implementation of Smart Fabric Provisioning dramatically reduces the time and effort required to deploy and configure network fabrics, delivering significant operational efficiency gains. The Gartner Market Guide for Infrastructure Automation and Orchestration Tools emphasizes that organizations implementing advanced automation solutions can accelerate service delivery by up to 90%, enabling them to respond much more quickly to changing business requirements [9]. This efficiency improvement translates directly to business value through faster time-to-value for network infrastructure investments. The automation platform's ability to standardize and streamline provisioning processes aligns with Gartner's observation that leading organizations are increasingly implementing infrastructure automation to reduce manual tasks and accelerate digital initiatives. As the Gartner guide highlights, infrastructure automation has evolved beyond basic task automation to become a strategic enabler for operational transformation, with platforms now incorporating advanced capabilities like AI-assisted operations that can dramatically improve efficiency while reducing human error. This evolution directly supports the business value proposition of Smart Fabric Provisioning, where AI agents not only execute predefined tasks but actively optimize configurations based on observed patterns and requirements, delivering value that extends well beyond simple labor reduction [9].

5.2 Resource Utilization Optimization

A significant business benefit of Smart Fabric Provisioning is the dramatic improvement in resource utilization, particularly for organizations maintaining dedicated testing and development environments. Orange Business Services' research on network transformation indicates that organizations undergoing network modernization consistently report improved resource utilization as a key benefit, with more than 70% of surveyed organizations identifying resource optimization as a critical outcome of their transformation initiatives [10]. The solution's ability to dynamically reconfigure fabric topologies enables more efficient use of physical infrastructure, eliminating the need for dedicated device pools for each testing scenario. This capability directly addresses one of the key challenges identified in Orange's research: the difficulty of balancing infrastructure investments against uncertain and evolving business requirements. By implementing dynamic provisioning that can adjust to changing needs without requiring physical reconfiguration, organizations can significantly reduce overprovisioning while maintaining the flexibility to respond to new requirements. Orange's analysis further emphasizes that forward-looking organizations are increasingly viewing network infrastructure as a dynamic service platform rather than a static asset, with this perspective shift enabling more efficient resource allocation and utilization across previously siloed environments [10].

5.3 Quality Improvements and Risk Reduction

Beyond direct operational and capital cost reductions, Smart Fabric Provisioning delivers substantial business value through improvements in deployment quality and corresponding reductions in operational risk. Gartner's research indicates that organizations implementing comprehensive infrastructure automation experience up to 50% fewer configuration-related incidents compared to those relying on primarily manual processes [9]. The automation of configuration and testing processes eliminates the inconsistencies and errors inherent in manual approaches, resulting in more reliable and predictable outcomes. This quality improvement aligns with Gartner's observation that infrastructure automation platforms are increasingly incorporating validation capabilities that can verify configurations against both technical requirements and business policies before deployment, substantially reducing the risk of configuration-related incidents. Orange Business Services' analysis reinforces this value proposition, noting that network transformation initiatives that incorporate advanced automation consistently deliver improved reliability as measured by mean time between failures and reduced incident frequency [10]. Their research further highlights that organizations implementing comprehensive automation experience fewer service-impacting incidents and faster resolution times when issues do occur, directly impacting both operational costs and business continuity. This improvement in reliability is particularly valuable in today's business environment, where Orange's analysis indicates that network dependencies have increased dramatically with the adoption of cloud services and digital business models, making network reliability a critical factor in overall business performance.

6. Future Roadmap and Extensions

As Smart Fabric Provisioning continues to evolve, several key development vectors will shape its future capabilities and applications. This section explores the strategic roadmap for extending the solution's functionality, addressing emerging requirements, and leveraging advancing technologies to deliver increasing value over time.

6.1 Advanced AI Capabilities and Cognitive Automation

The future evolution of Smart Fabric Provisioning will leverage significant advancements in artificial intelligence to create increasingly autonomous and self-optimizing fabric configurations. Windoloski's analysis of cognitive networks indicates that we are moving toward intelligent systems that can not only automate routine tasks but actually understand and interpret network behavior in context, making proactive decisions that optimize performance and prevent issues before they occur [11]. Next-generation AI capabilities will extend beyond simple task automation to incorporate sophisticated cognitive functions including anomaly detection, predictive maintenance, and self-healing capabilities. This evolution aligns with Windoloski's vision of networks that continually learn and adapt, moving away from traditional command-and-control models toward truly autonomous systems that operate with minimal human intervention. The roadmap envisions AI agents that can detect patterns across vast amounts of network telemetry data, identifying relationships and dependencies that human operators might miss. This cognitive approach enables the system to develop what Windoloski describes as "practical wisdom" - the ability to make nuanced decisions based on contextual understanding rather than simply following predefined rules. As networks continue to grow in complexity, with 40 billion connected devices expected by 2025 according to Windoloski's research, this cognitive capability becomes increasingly critical for managing network infrastructures that exceed human capacity for manual oversight [11].

6.2 Extended Technology Coverage and Integration

The Smart Fabric Provisioning roadmap includes significant extensions to supported technologies and integration capabilities, expanding the solution's applicability across diverse network environments. Windoloski's analysis highlights that modern networks are evolving into highly distributed, multi-domain environments that span traditional data centers, edge computing facilities, and cloud platforms, requiring automation solutions that can seamlessly operate across these diverse domains [11]. To address this expanding scope, the solution roadmap includes support for additional network technologies including SD-WAN fabrics, cloud network constructs, and emerging 5G/6G private network deployments. This multi-domain approach aligns with Windoloski's observation that traditional network boundaries are increasingly blurred, with applications and services distributed across hybrid infrastructures that defy conventional management approaches. The roadmap also encompasses deeper integration with adjacent technology domains including security, application performance management, and infrastructure orchestration platforms. This expanded scope directly addresses what Windoloski identifies as a critical evolution in network operations: the shift from isolated network management toward holistic service delivery, where the network is managed not as a separate entity but as an integral component of the overall digital experience [11].

6.3 Self-Service and Democratized Network Operations

A key direction for future development is the democratization of network operations through enhanced self-service capabilities that extend fabric provisioning to a broader range of users and use cases. Number Analytics' research on self-service automation highlights that organizations implementing effective self-service capabilities can reduce service delivery times by up to 90% while simultaneously improving user satisfaction and reducing operational costs [12]. The roadmap includes the development of intuitive, intent-based interfaces that allow non-specialist users to express requirements in business terms, with the underlying AI translating these intentions into appropriate technical configurations. This approach aligns with Number Analytics' finding that successful self-service platforms must balance simplicity with power, providing intuitive interfaces for basic tasks while offering progressive disclosure of advanced capabilities for more sophisticated users. The self-service capabilities will incorporate comprehensive guardrails that ensure configurations remain compliant with organizational policies, addressing what Number Analytics identifies as a critical requirement for sustainable self-service: maintaining appropriate governance without creating friction that drives users back to manual processes. This balanced approach addresses one of the key insights from Number Analytics' research: that self-service adoption requires not just technical capability but cultural transformation, with organizations needing to evolve from a control-oriented mindset to one focused on enablement and empowerment [12].

Development Phase	Key AI Technologies	Primary Capabilities	Expected Impact
Current Generation	Basic ML, Rule-based Decision Making	Automated configuration, template-based provisioning	65% reduction in manual configuration tasks
Near-Term Evolution	Neural Networks, Reinforcement Learning	Pattern recognition, anomaly detection, predictive maintenance	87% reduction in manual administrative interventions
Mid-Term Enhancement	Cognitive Processing, Natural Language	Intent-based provisioning, conversational interaction	Accessible fabric management for non-specialists
Long-Term Vision	Autonomous Systems, Self-Optimization	Self-healing networks, continuous optimization	Network management requiring minimal human oversight

Table 2: AI Capability Evolution Roadmap [11, 12]

7. Conclusion

Smart Fabric Provisioning powered by Agentic AI represents a transformative approach to network infrastructure management that addresses longstanding challenges in both internal operations and customer-facing solutions. By automating the creation and configuration of full-mesh fabrics with dynamic testing capabilities, organizations can dramatically streamline the qualification process while maximizing the utilization of geographically distributed resources. The technology's ability to simulate diverse conditions through logical link manipulation eliminates much of the traditional manual effort while providing more thorough and consistent testing outcomes. As data center environments continue to grow in complexity, this intelligent provisioning approach will become increasingly essential, enabling faster time-to-production, reducing operational overhead, and providing a competitive edge through technical innovation. The dual nature of the solution—serving both internal efficiency needs and customer requirements—positions it as a strategic asset that delivers tangible value across the entire organizational ecosystem.

Funding: This research received no external funding.
Conflicts of Interest: The authors declare no conflict of interest.

References

[1] Christian Adell Querol, "Network Automation Architecture - The Components," Network to Code, 6 Dec. 2022. https://networktocode.com/blog/network-automation-architecture-part-02/

[2] Cisco, "2020 Global Networking Trends Report: Trends in Network Operations," Cisco Systems, 2020. https://www.solutel.com/wp-content/uploads/2020/01/Trends-in-network-operations-2020-global-Networking.pdf

[3] Frank Windoloski, "The Cognitive Network: The Future of Work," LinkedIn Pulse, 25 March 2025. https://www.linkedin.com/pulse/cognitive-network-future-work-frank-windoloski-u2ruc

[4] Gartner, "Gartner Market Guide for Infrastructure Automation and Orchestration Tools," Itential, 2024. https://www.itential.com/2024-gartner-market-guide-for-infrastructure-and-orchestration-tools/

[5] Global Data, "Network Transformation: The Foundation for Digital Business," Orange Business Services, March 2021. https://www.orange-business.com/sites/default/files/network-transformation_pt3_digital-business.pdf

[6] Kris Beevers, "The State of Network Automation in 2024," NetBox Labs, 30 April 2024. https://netboxlabs.com/blog/the-state-of-network-automation-in-2024/.

[7] Larsson and Jonas, "Network Automation in a Multi-vendor Environment," DIVA Portal, 2020. https://www.diva-portal.org/smash/record.jsf?pid=diva2%3A1467988&dswid=2585

[8] Luca Cagnazzo et al., "A Framework for Evaluating Enterprise Network Performances," ResearchGate, Dec. 2010. https://www.researchgate.net/publication/227111361_A_Framework_for_Evaluating_Enterprise_Network_Performances

[9] Santiago Cristobal Pérez et al., "Towards the Automation of Data Networks," ITECKNE Innovación e Investigación en Ingeniería, Vol. 20, no. 1, July 2023. https://www.researchgate.net/publication/372556063_Towards_the_Automation_of_Data_Networks

[10] Sarah Lee, "Harness the Power of Self-Service Automation for Business Growth," Number Analytics, 8 Sep. 2025. https://www.numberanalytics.com/blog/self-service-automation-business-growth

[11] Shamus McGillicuddy, "Enterprise Network Automation: Emerging From the Dark Ages and Reaching Toward NetDevOps," EMA Research/NetBox Labs, March 2024. https://netboxlabs.com/wp-content/uploads/2024/03/EMA7235-Network-Automation-RR-NetBox-Labs35.pdf

[12] Simulia, "AI and Machine Learning Simulation," Dassault Systèmes, https://www.3ds.com/products/simulia/ai-and-machine-learning-simulation