Modernizing Federal Networks from the Campus and Beyond

Networks and Telecommunications Community of Interest

Date Released: July 17, 2026

Synopsis

The U.S. Federal Government's networking infrastructure faces mounting pressure from cloud adoption, AI workloads, distributed workforces, and an increasingly sophisticated threat landscape that has rendered legacy WAN architectures built on MPLS, leased lines, and TDM fundamentally inadequate. Meeting modern mission demands requires simultaneous compliance with FISMA, FedRAMP, and TIC 3.0 while embedding Zero Trust principles at every layer, a combination that perimeter-based security models cannot support.

This white paper introduces the Future-Agile Network (FAN), an architectural framework that replaces rigidity with intentional adaptability. By blending fiber, 5G/6G, and LEO satellite connectivity with AI-driven network management and continuous Zero Trust validation, FAN enables federal agencies to absorb dynamic workloads, sustain operations through disruption, and treat compliance as a design outcome rather than an afterthought. The FAN positions the network not as a constraint, but as the intelligent, resilient backbone of government modernization.

American Council for Technology-Industry Advisory Council (ACT-IAC)

The American Council for Technology-Industry Advisory Council (ACT-IAC) is a non-profit educational organization established to accelerate government mission outcomes through collaboration, leadership and education. ACT-IAC provides a unique, objective, and trusted forum where government and industry executives are working together to improve public services and agency operations through the use of technology. ACT-IAC contributes to better communication between government and industry, collaborative and innovative problem solving, and a more professional and qualified workforce.

The information, conclusions, and recommendations contained in this publication were produced by volunteers from government and industry who share the ACT-IAC vision of a more effective and innovative government. ACT-IAC volunteers represent a wide diversity of organizations (public and private) and functions. These volunteers use the ACT-IAC collaborative process, refined over forty years of experience, to produce outcomes that are consensus-based.

To maintain the objectivity and integrity of its collaborative process, ACT-IAC welcomes the participation of all public and private organizations committed to improving the delivery of public services through the effective and efficient use of technology. For additional information, visit the ACT-IAC website at www.actiac.org.

Networks and Telecommunications Community of Interest

The Networks and Telecommunications Community of Interest provides clear, impartial feedback and recommendations on networks and telecommunication issues identified in collaboration with the government project sponsors and advisors.

Disclaimer

This document has been prepared to contribute to a more effective, efficient, and innovative government. The information contained in this report is the result of a collaborative process in which several individuals participated. This document does not – nor is it intended to – endorse or recommend any specific technology, product, or vendor. Moreover, the views expressed in this document do not necessarily represent the official views of the individuals and organizations that participated in its development. Every effort has been made to present accurate and reliable information in this report. However, neither ACT-IAC nor its contributors assume any responsibility for consequences resulting from the use of the information herein.

Copyright

For further information, contact the American Council for Technology-Industry Advisory Council at (703) 208-4800 or www.actiac.org.

MODERNIZING U.S. FEDERAL AGENCY NETWORKS

Security & Access Layer
Zero Trust

SASE
SWG | CASB | ZTNA | FWaaS

Remote / Satellite Access (VSAT)

Remote / Satellite Access (VSAT)

Autonomous Networking

SON (Future St.) Self-Organizing Networks

IBN (Future St.) Intent Based Networking

SD-WAN
CORE ORCHESTRATION LAYER

MPLS

Core & Data Networks

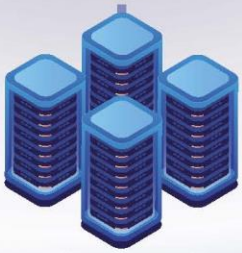

Federal Data Center

5G / NTN

Public Internet

Proprietary Terminals

Table of Contents

Introduction

The U.S. Federal Government is at an inflection point in the evolution of its networking infrastructure. For decades, Wide Area Network (WAN) architectures rooted in Multiprotocol Label Switching (MPLS), leased lines, and even legacy time division multiplexing (TDM) delivered reliable connectivity across agencies and mission environments. These traditional designs served their purpose in a different era, one where traffic patterns were linear, applications were concentrated in centralized data centers, and collaboration rarely extended beyond email, voice, and siloed internal systems. Predictability and stability were their defining strengths, but flexibility was never the design goal.

That landscape has shifted irreversibly. Federal missions today depend on cloud-hosted services, mobile-first workforces, and data-intensive platforms such as artificial intelligence, geospatial analytics, and real-time sensor fusion. The shift from centralized to distributed applications has created unpredictable traffic patterns that can no longer be satisfied by rigid backbones. Agencies are expected to deliver secure integration with private sector partners, support interagency collaboration on national security and emergency response, and sustain continuous operations even in moments of crisis, all while ensuring compliance with a complex regulatory framework that includes but is not limited to: Federal Information Security Modernization Act (FISMA), Federal Risk and Authorization Management Program (FedRAMP), Criminal Justice Information System (CJIS), and Trusted Internet Connections (TIC) 3.0 guidelines.

Overlaying these security controls and operational demands is the reality of a transformed threat environment. The perimeter-only security model is obsolete. Nation-state adversaries, insider risks, and the exponential growth of endpoints, from Internet of Things (IoT) sensors to tactical edge devices, mean every access point is a potential attack vector. Federal networks must embed Zero Trust principles into their very fabric, continuously validating users, devices, and applications. Security can no longer be bolted on; it must become intrinsic to the design. Continuous monitoring, real-time compliance validation, and automated remediation powered by artificial intelligence (AI) are the only ways to keep pace with evolving threats.

AI itself is the defining disruptor. It is not simply another application consuming bandwidth. Training clusters, distributed inference workloads, machine-to-machine communication, and real-time decision support systems introduce asymmetric and bursty traffic flows that stress even modernized infrastructures. AI workloads demand ultralow latency, massive east-west data movement within data centers, and elastic north-south traffic between agencies, clouds, and the edge. If federal networks remain locked in legacy patterns, they risk becoming bottlenecks that slow down missions rather than platforms that accelerate them.

This convergence of pressures has led to the articulation of what is called a Future-Agile Network (FAN) model. FAN is not a single product or service but an architectural and operational philosophy

that acknowledges the speed of change and designs for it. It is built to be resilient, surviving disruptions through intelligent routing, redundant paths, and multi-bearer connectivity that blends fiber, cellular, and satellite seamlessly. It is inherently secure, embedding Zero Trust into every transaction and validating continuously through telemetry and automated controls. It is adaptive, scaling dynamically as AI workloads spike, as collaboration surges during crises, and as emerging technologies such as Wi-Fi 7, 5G/6G, and low-earth orbit (LEO) satellite constellations reshape what is possible. Above all, it is mission-aligned, recognizing that government agencies have unique operational demands that require flexibility, vendor neutrality, and cost efficiency without sacrificing speed or security.

The FAN positions the network as a strategic enabler rather than a constraint. It turns compliance into an outcome of design rather than an afterthought, treating FedRAMP and TIC 3.0 as foundations rather than hurdles. It enables real-time collaboration to become a first-class capability, with networks that can dynamically shift traffic between fiber, satellite, and 5G to sustain secure video or command and control applications during a crisis without degradation. And it prepares agencies for the AI era, ensuring that the network is not overwhelmed by new workloads but instead anticipates, absorbs, and optimizes them.

By replacing rigidity with agility, the FAN framework ensures that federal agencies are not merely reacting to change but are proactively shaping the future of secure, mission-driven networking. In this model, security is not a defensive wall but an adaptive immune system. AI is not a burden on the network but a copilot in its management. And the network itself becomes more than a conduit for traffic; it becomes the digital backbone of government modernization and resilience in a world where speed, trust, and intelligence define mission success.

The Recent State of Federal Agency WAN Architectures

In examining the current landscape of federal agency WAN architecture, several consistent themes emerge. Most agencies have relied on established design principles, centralized controls, standardized configurations, and an emphasis on perimeter-based protections which have changed little over the past decade. Although these foundational elements provided stability and compliance, they have not kept pace with rapid advancements seen in commercial network architectures or with the flexibility demanded by cloud and mobile-first strategies. This stagnation can be traced primarily to the influence of the initial TIC mandate and its stringent security directives, which required large, centralized security inspection gateways and shaped the evolution of federal WANs more than any other policy in recent years.

Restrictive Controls Introduce Increased Architectural Costs

The focus on network security at times reduced agility. While the guidance helped protect federal infrastructure, it also created obstacles to modernization, requiring CIOs to balance protection with adaptability.

Ensuring secure traffic was only part of the solution. Federal agencies require high network availability for essential sites such as data centers, headquarters, and internet gateways. The demands for both security and high availability led to costly and inflexible network architectures.

These growing expenses underscore a broader challenge: as agencies pursue resilient architectures to ensure operational continuity, they are also confronted with complex cost-benefit calculations tied to legacy systems and modernization efforts.

TIC 3.0 Provides Agencies the Opportunity to Innovate

Today, federal agencies stand at a pivotal moment with the introduction of TIC 3.0 guidelines, which balance the importance of securing our federal networks with the demands of new data flows (Trusted Internet Connections (TIC) 3.0 Core Guidance Documents 2025). It is a transformative shift that balances robust security with unprecedented architectural flexibility. While agency CIOs and CISOs remain accountable for maintaining centralized security logs and upholding comprehensive control over network security, TIC 3.0 allows decentralization of Internet access points. Instead of routing all traffic through a handful of centralized TIC locations, agencies can now move egress to branch offices and remote sites. This evolution significantly enhances the performance of Internet-based applications, reduces the reliance on high-cost centralized TIC bandwidth, and enables more agile responses to mission demands. The federal WAN of today and tomorrow is thus beginning to resemble the dynamic, distributed infrastructures of leading commercial enterprises and Fortune 500 companies, empowering their users through streamlined connectivity, rapid adoption of emerging technologies, and the flexibility to support a modern, mobile-first workforce.

Far from merely mirroring the dynamic, distributed infrastructures of leading commercial enterprises, the flexibility of TIC 3.0 guidance enables federal CIOs to drive genuine innovation in how they deliver services to their end users, not only in traditional workplaces but also in challenging environments. With more streamlined connectivity, accelerated adoption of emerging technologies, and enhanced flexibility to support a modern, mobile-first workforce, federal agencies are now positioned to actively shape the future of government IT, breaking through barriers of the past to lead in both operational excellence and technological advancement.

Federal Compliance Requirements

The critical nature of the telecommunications infrastructure and the vulnerability of those networks to malicious security attacks have led to a host of federal mandates designed to protect federal networks and sensitive data. Compliance with these standards is a critical component that agency personnel must consider when designing a WAN and/or selecting a WAN service provider.

These range from adherence to FISMA standards for securing authorization to operate (ATO), and alignment with National Institute of Standards and Technology (NIST) Special Publication (SP) 800-53 Revision 5 (NIST SP 800-53 Rev. 5 2020), to specialized handling of criminal justice information to comply with CJIS policies. To support their federal customers and in the interest of national security,

telecom providers are constantly working to address these federal-specific security controls with partitioned system architectures to protect customer systems and data, and complementary cryptographic standards to protect data in motion and at rest. These include but are not limited to post-quantum cryptographic key exchange practices, rigorous access controls, Zero Trust principles, multifactor authentication, continuous monitoring, rapid incident reporting, and remediation.

Enhancing WAN Security

Federal agencies transitioning from traditional perimeter-based security models, such as those defined by the TIC 2.2 approach, which rely on a limited number of Internet-to-enterprise gateways for centralized inspection, are embarking on an important journey toward a zero-trust security model. With the disappearance of traditional network perimeters and security boundaries, agencies now have strong options for embedding security directly into the network in ways that are efficient and aligned with newer federal mandates and guidance. The TIC 3.0 model, as outlined in the ACT-IAC Zero Trust white paper from April 2019 and further defined in NIST SP800-215 (NIST SP800-215 n.d.), embraces the Zero Trust paradigm: securing users, devices, networks, applications, and policy enforcement through Automation and near real-time Analytics. Unlike legacy strategies, Zero Trust recognizes that threats can originate from anywhere, requiring continuous validation and inspection at every layer, including basic protections like those against Denial-of-Service (DoS) and Distributed Denial-of-Service (DDoS) attacks. Agencies now have the flexibility to place Policy Enforcement Points (PEPs) and security controls where they are most effective. This includes at the edge, Internet exchange points, cloud service providers, or core data centers, fully aligning with the TIC 3.0 model. Zero Trust Network Access (ZTNA) is a critical part of this framework, complemented by cloud-based protections enabled through Software-Defined-WAN (SD-WAN) and Secure Access Service Edge (SASE) architectures.

Many agencies have already implemented security measures offering inspection and reinspection across a subset of Zero Trust pillars, but few have achieved a comprehensive, integrated Zero Trust approach. Recognizing that zero trust is a journey, not a one-time project, and no single vendor currently addresses all Zero Trust pillars, agencies should prioritize incremental, interoperable enhancements over disruptive, forklift upgrades. Leveraging service provider lab environments to validate the interoperability and effectiveness of these components can help ensure a smooth and effective transition, strengthening an agency's overall security resilience in a measured and sustainable way.

Modernizing Networks by Enabling an Intelligent Edge

The decentralization of security under TIC 3.0 guidance has been matched by the introduction of technologies, including SD-WAN and SASE, that provide agencies with the ability to modernize their networks, drive down costs, and push intelligence, flexibility, and resilience to the network edge. These technologies offer a cost-effective and superior or modernized complementary service to MPLS networks, providing organizations with enhanced flexibility across equipment, networking, licensing, and engineering resources. By integrating SD-WAN alongside existing MPLS infrastructure, organizations can optimize their network architecture through a hybrid approach that leverages the reliability and performance guarantees of MPLS for critical applications while utilizing SD-WAN's redundancy, scalability, and potential cost savings. A more in-depth security section follows later in this document.

Various industry studies show that moving to SASE architectures, which incorporate SD-WAN, can deliver 30% to 50% savings over traditional WANs. (Thorne 2020) (IDC 2022) However, benefits can vary significantly from agency to agency depending on the existing environment, future requirements, and the ability of agencies to implement changes while maintaining mission effectiveness.

TDM and Technical Debt Considerations

As agencies transition to SD-WAN and IP-centric networks, they must still support legacy TDM-based voice and radio systems that remain mission-critical, despite rapidly rising costs. As IP adoption accelerates, the number of TDM ports continues to decline, driving up the per-port cost of maintaining central office infrastructure. CIOs are therefore faced with a strategic choice: continue investing in aging, increasingly expensive legacy platforms, or accelerate migration to IP-enabled solutions. While both paths carry financial and operational implications, moving to IP radios and voice platforms not only modernizes capabilities for end users but also delivers immediate reductions in network operating costs while improving long-term sustainability.

How FAN Delivers Measurable Business Outcomes

Enhanced Efficiency

SD-WAN enhances network agility and fundamentally transforms network economics by intelligently using a mix of underlays including secure internet via Dedicated Internet Access (DIA), Broadband, cellular, and satellite, while retaining MPLS where it delivers the most value. Policies for applications that are highly sensitive to latency and jitter might prioritize MPLS as the primary underlay to benefit from predictable performance and consistent packet delivery. Conversely, applications that are less sensitive to latency and jitter can be steered over broadband, satellite, and cellular underlays, with the SD-WAN application dynamically selecting the best available path and failing over as conditions change.

While each organization’s costs may vary, our analysis incorporates Enterprise Infrastructure Solutions (EIS) Managed Trusted Internet Protocols Service (MTIPS) list pricing benchmarks and finds that TIC 3.0 compliant architectures employing Branch Office Internet connectivity

will benefit from up to a 50% reduction in cost per Mbps for secure Internet access when leveraging a FedRAMP-authorized SASE vendor. Upon extending Internet connectivity to branch offices to deliver better Internet-based application performance at a lower cost, agencies are best served by electing to use this new Internet circuit for site resiliency using SD-WAN. Beyond circuit costs, SD-WAN reduces technical debt by eliminating complex legacy architectures, obsolete hardware, and outdated software platforms that drain IT budgets and constrain agility.

The solution streamlines workforce economics by cross-leveraging the specialized skills of existing technical staff to design, deploy, and manage network infrastructure. Furthermore, organizations can significantly decrease costs associated with hiring, training, and retaining scarce technical talent while improving operational efficiency through centralized management and automation of the SD-WAN.

Technology Convergence

SD-WAN solutions can also reduce product, maintenance, and operational costs by combining traditional technology layers such as routers, firewalls, and WAN optimization products and operating systems into single platforms.

Strengthened Resiliency

Network downtime directly impacts revenue, productivity, and customer satisfaction. SD-WAN provides intelligent path selection and automatic failover across multiple transport options, ensuring business continuity even when individual circuits fail.

Proven Scalability

As agencies expand into new locations or respond to changing business requirements, SD-WAN enables rapid deployment without the lead times and complexity of traditional WAN provisioning. Zero-touch provisioning and cloud-based management and orchestration allow new sites to be activated in days rather than months, directly supporting business growth objectives while maintaining consistent security and performance standards across the agency's entire network footprint.

Comprehensive Cost Analysis Framework

The return on investment (ROI) framework we are providing examines the complete financial picture across six critical dimensions:

Infrastructure Costs: Quantifies savings from optimized circuit selection, reduced hardware requirements, and elimination of technical debt in system architecture and legacy platforms.

Human Capital: Calculates the financial impact of developing talent, streamlined training requirements, and improved retention through simplified operations.

Security Operations: Evaluates integrated security service costs including patching, vulnerability management, identity and access control, event monitoring and logging, and seamless integration with an agency's existing security stack.

Data Management: Assesses costs associated with data silos, inadequate governance frameworks, and poor data quality that impact network visibility and decision-making.

Business Continuity: Measures the value of improved uptime, reduced mean time to repair, and enhanced application performance.

Strategic Agility: Captures the competitive advantage of faster deployment cycles, improved scalability, and enhanced ability to support digital transformation initiatives.

Figure 1: Cost Analysis Framework

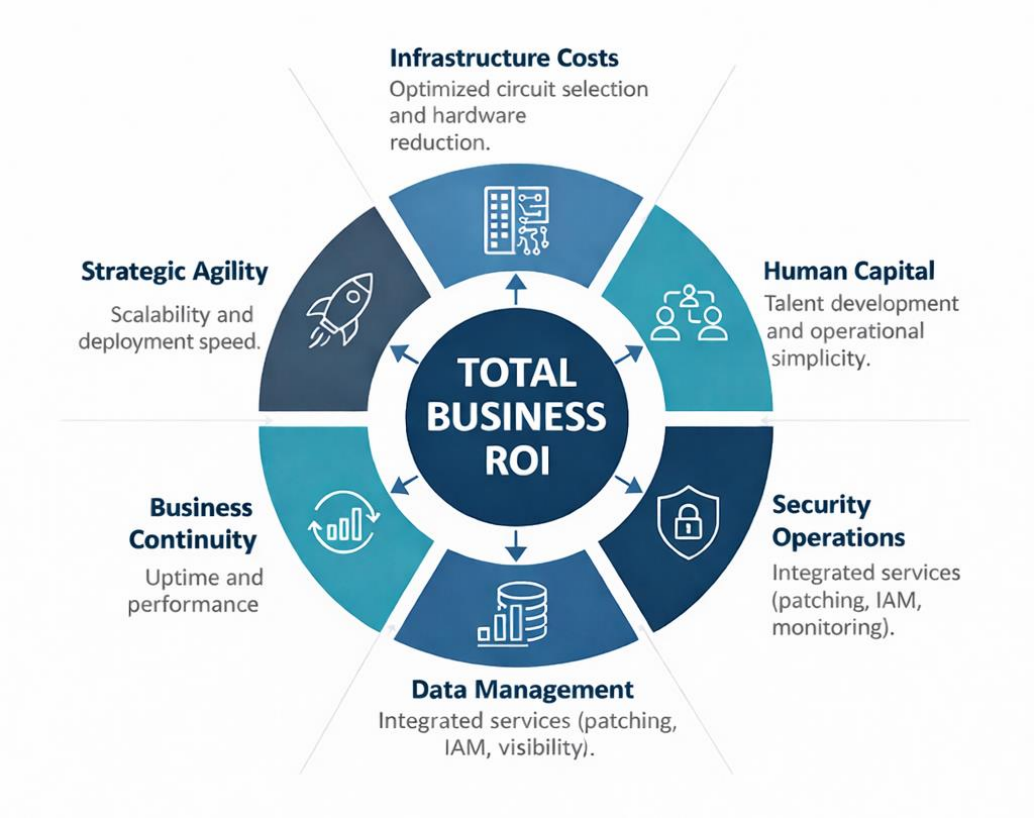

Introducing Networks for Federal Agencies in the AI Age

Federal networks are now entering a new operational phase defined not by incremental modernization, but by systemic transformation. Artificial intelligence, edge computing, multi-cloud architecture, and machine-to-machine collaboration fundamentally change what the network must do, how fast it must respond, and how risk must be managed.

In the AI age, the network is no longer just a transport layer. It becomes:

- A real time control fabric for distributed systems.
- A security enforcement surface embedded everywhere, not centralized anywhere.
- A data movement platform for high volume, asymmetric, and burst driven workloads.
- A policy execution engine that must operate at machine speed.
- A compliance instrument that must provide continuous, provable assurance.

AI workloads do not tolerate static routing, slow provisioning cycles, or human-in-the-loop troubleshooting as the primary operating model. Training clusters, distributed inference, edge analytics, and enhanced AI collaboration create traffic patterns that are highly dynamic, latency sensitive, and operationally unforgiving. In this environment, a network that was designed to be stable becomes a liability. Stability must be replaced with controlled adaptability.

At the same time, the threat model has shifted. AI expands both the value and the vulnerability of data in motion. Distributed inspection, Zero Trust enforcement, continuous telemetry, and automated remediation are no longer advanced features; they are baseline survival requirements. The network must behave less like a static perimeter and more like an adaptive immune system, responding continuously to changing conditions, threats, and mission demands.

For federal agencies, this shift is not optional. It is driven by mission realities:

- AI enabled decision support and analytics
- Edge and tactical data collection
- Multi-cloud and interagency integration
- Continuous compliance expectations
- Rising adversary sophistication
- The coming impact of post-quantum cryptography

These forces collectively redefine what “good” looks like for a federal network.

This is the context in which the FAN framework emerges, not as a technology refresh, but as a new architectural operating model for the AI era. FAN represents the transition from static, perimeter-

oriented WANs to adaptive, intelligence-assisted, mission-aligned network fabrics designed for continuous change.

The sections that follow move beyond the limitations of legacy architectures and define what a federal network must become in the AI age: resilient by design, secure by default, adaptive by nature, and optimized for both human oversight and machine-speed execution.

Networks in the AI Era: AIOps, Transport, and Traffic Growth

Artificial intelligence is transforming not only the applications federal agencies use but also the very networks that support them. What was once a relatively stable foundation optimized for predictable traffic, centralized data centers, and long provisioning cycles is now being pushed into a new era of dynamism. As agencies adopt cloud, edge, and AI-driven capabilities, they must rethink how networks are designed, operated, and secured. The overall impact of AI on a federal network is profound: it redefines operational roles, alters traffic flows, and expands the security and compliance perimeter. Three key themes capture this shift: AIOps (Artificial Intelligence for IT Operations) with guardrails, the reshaping of WAN architecture in the AI era, and the growth of AI-driven traffic patterns.

AIOps Autonomous Operations with Guardrails

The traditional model of network operations, including manual troubleshooting, reactive fault isolation, and static policy enforcement, cannot keep pace with modern demands. AI introduces far more variability, from dynamic workloads shifting between data centers and clouds, to mission-critical inference tasks at the edge that cannot tolerate downtime. Enter AIOps (Artificial Intelligence for IT Operations), which brings machine learning and advanced analytics into daily network management. AIOps can automatically detect anomalies, predict failures before they occur, and dynamically adjust network paths to optimize performance.

But in a federal mission environment, autonomy without accountability is unacceptable. Networks that power national security, healthcare delivery, and citizen services cannot risk “black-box” automation. Guardrails are essential. Predefined policies, compliance checks, and human-in-the-loop oversight ensure that AI-driven operations remain transparent, explainable, and aligned with mission requirements. Here, AIOps is not about replacing human operators but augmenting them, shifting their focus from firefighting to higher-level strategic oversight. This evolution both reduces operational fatigue and raises confidence in meeting stringent regulatory requirements such as FISMA, FedRAMP, CJIS, and TIC 3.0.

WAN Impact Architecting Transport for the AI Era

WANs have always been the backbone of federal connectivity, linking data centers, cloud providers, field offices, and mission sites. But the rise of AI workloads fundamentally reshapes WAN

requirements. Training and inference processes generate heavy east-west traffic between data centers, large scale transfers of unstructured datasets, and latency-sensitive north-south traffic that flows between users, clouds, and edge environments. Unlike legacy applications, these patterns are bursty, data heavy, and often unpredictable.

Architecting transport for the AI era requires building resiliency and adaptability directly into the WAN. High-capacity fiber, redundant routes, and multi path resiliency are no longer optional, they are baseline. Integration with satellite, 5G/6G, and edge connectivity ensures continuity in austere or disconnected environments. Meanwhile, SD-WAN overlays and Secure Access Service Edge (SASE) frameworks enable agile rerouting of traffic while enforcing consistent security and compliance. In effect, WANs must evolve from static highways into intelligent transport fabrics that are self-optimizing, secure, and capable of supporting the immense data gravity of AI.

AI Where the Traffic Will Grow

AI is not just another application; it is a force multiplier for traffic growth. Large language models and machine learning algorithms require massive datasets to be trained. Inference services must deliver answers in real time, whether mission planning, citizen-facing services, or security analysis. Collaboration tools increasingly embed AI for translation, transcription, and knowledge retrieval, amplifying everyday bandwidth demands.

This growth is most visible in three domains:

- **Data Center to Cloud:** Training and retraining models, dataset migration, and large-scale storage synchronization.
- **Edge to Core:** Mission data collected in the field, pre-processed by edge AI, and transmitted back to central systems.
- **User to Cloud Applications:** Routine collaboration enhanced by AI features that generate background traffic agencies never had to plan for before.

Without predictive capacity planning, bottlenecks will erode both performance and mission readiness. Agencies need telemetry-driven forecasting models that simulate AI traffic flows and help design networks that scale over time.

From AI-Driven Scale to Quantum-Era Risk

As agencies modernize WAN architectures to absorb the explosive growth of AI traffic, they are simultaneously accelerating toward a second, less visible inflection point. The same transport fabrics being strengthened to move unprecedented volumes of data, more bandwidth, more paths, more endpoints, more automation also expand the attack surface and longevity of that data in motion. AI forces networks to move faster and farther; quantum computing threatens to look backward in time

and break what was once considered secure. This convergence creates a new design tension for federal WANs: it is no longer sufficient to architect transport solely for performance, resiliency, and adaptability. Transport must now be built with cryptographic survivability in mind.

In practical terms, AI makes WAN more valuable to adversaries. High volume model training data, mission telemetry, and enhanced AI collaboration flows concentrate sensitive information into long-lived, high-throughput links. Quantum computing changes the threat horizon by rendering "secure today" encryption potentially vulnerable tomorrow. Data intercepted now may be unreadable for years until it suddenly is not. This reality reframes WAN modernization as a time-bound security decision rather than a purely architectural one.

The impact is profound: WANs are no longer just conduits for AI-era workloads; they are the frontline where quantum risk materializes first. Edge routers, transport encryption, routing overlays, and key exchange mechanisms become strategic controls, not implementation details. The question shifts from *how fast and flexible the network can be* to *how long the confidentiality and integrity of data in motion will endure*. This is where the discussion must expand beyond AI driven traffic growth and the impact of architecting post-quantum transport, as well as cryptographic integrity in an adversarial future.

Impact of Quantum Computing

Quantum computing (QC) promises transformative benefits for fields ranging from cryptography and logistics to scientific research, thanks to its unparalleled processing power and ability to solve complex problems exponentially faster than classical computers. However, this leap in computational capability also brings significant risks to federal networks, particularly through its potential to break traditional encryption methods and exacerbate the risks associated with man-in-the-middle attacks. Today, adversaries may be harvesting encrypted data with the intent to decrypt it in the future as quantum capabilities mature. This strategy is referred to as "harvest now, decrypt later." This reality underscores the urgent need to adopt quantum resistant key exchange architectures, which are specifically designed to thwart man-in-the-middle attacks and safeguard sensitive communications.

The WAN represents the most exposed segment of the federal enterprise to quantum computing threats, making it a prime target for adversaries seeking to intercept sensitive data in transit. By prioritizing the deployment of Post-Quantum Cryptography (PQC) at agency edge routers, CIOs and CISOs can deliver the greatest impact on their organization's quantum resilience. Leading PQC vendors now support FIPS 203 compliant key exchange architectures, such as those utilizing ML-KE-1024 and ML-DSA-87 GCM/CCM, to negotiate encryption keys out-of-band (OOB) and safeguard IP traffic as it traverses the WAN. Securing the agency's WAN at the edge with PQC encryption is the most effective step agencies can take today to future-proof their cybersecurity posture against quantum era threats.

Today's quantum threat may force CIOs and CISOs to debate some important questions about their plans for modernization.

- Should they weigh the flexibility and performance enhancements offered by TIC 3.0 architecture, which are designed to counter today's threats, against the robust security controls available in more restrictive network designs?
- While SD-WAN delivers modern, adaptable networking solutions that can be strengthened with quantum-resistant encryption, do private network architectures like MPLS, with their additional layers of control and segmentation, remain essential for enterprises and agencies transmitting the most sensitive, mission critical data?
- In an era where quantum threats are rapidly evolving, could these more restrictive architectures provide the enduring protection needed to secure our most valuable information?

Considering these questions, federal agencies may benefit most from a hybrid approach: leveraging the flexibility and performance enhancements of TIC 3.0 by deploying PQC key exchange architectures to secure their SD-WAN, while continuing to use MPLS as the underlay for their most sensitive data flows. This strategy allows agencies to adapt to evolving network demands and quantum era threats, all while maintaining the highest levels of control and security for mission critical communications.

Beyond SD-WAN: The Path to an Autonomous Network

SD-WAN is an important tool giving agencies centralized control, visibility, and agility over their WAN. But SD-WAN is not the end goal; it is the critical foundation for the next true leap in networking: the autonomous, self-driving network.

This future is defined by two closely related concepts: Intent-Based Networking (IBN) and Self Organizing Networks (SON). These technologies move beyond the automation of tasks to the autonomy of outcomes, and agencies with an SD-WAN footprint are uniquely positioned to lead this transition.

Intent-Based Networking (IBN)

Intent-Based Networking represents a structural change in how humans relate to digital infrastructure. It is not simply an evolution of routing protocols or a refinement of policy management. It is a shift in abstraction. It changes the level at which humans operate.

In traditional networking, including modern SD-WAN deployments, the human remains responsible for defining how the network should behave. Engineers configure routing tables, build quality-of-service policies, define access control lists, tune bandwidth allocations, and explicitly determine

failover paths. Even when automation tools are layered on top, the paradigm remains procedural. Humans instruct the network on the mechanics.

For example, in a traditional model, engineers might define a policy that routes all Zoom traffic from the executive office over a high-bandwidth MPLS circuit while sending all other traffic over a broadband link. They specify application signatures and bind the signatures to interfaces. The engineer configures thresholds, builds fallback rules, and monitors the performance dashboards. If jitter increases or latency spikes, an engineer investigates the configuration and adjusts the thresholds to improve performance. The system behaves exactly as instructed, but only within the boundaries of what was explicitly defined.

IBN changes the locus of control. Instead of describing the mechanism, the engineer describes the desired outcome. Rather than instructing the network how to prioritize packets, the engineer defines the business and operations objective such as ensuring that executives always experience high quality, jitter-free video conferencing. The difference appears subtle, but it is foundational. One model is command-driven. The other is outcome-driven.

In an IBN architecture, artificial intelligence and machine learning systems interpret this declared intent and convert it into network behavior. Humans no longer specify dozens of discrete configurations. Instead, the system performs a continuous closed loop that includes translation, activation, and assurance.

Translation is the first stage. The system ingests business language and contextual data. It understands application performance requirements, user roles, topology constraints, historical performance baselines, and security policies. From that understanding, it generates the complex set of device-level configurations required to satisfy the stated outcome. This may include routing adjustments, QoS markings, segmentation policies, firewall rules, and path selection logic across routers, switches, wireless controllers, and security appliances. The human expresses intent in business terms. The system renders it into operational reality.

Activation follows translation. The policies generated by the system are automatically deployed across the network fabric. This is not a static push of configuration. It is policy orchestration at scale. The deployment respects topology dependencies, sequencing requirements, and validation checkpoints. What previously required coordinated change windows and manual verification becomes software-driven execution.

Assurance is where IBN becomes transformative. Traditional networks are largely reactive. Monitoring tools surface alerts. Engineers interpret them. Adjustments are made manually or through predefined automation scripts. In contrast, IBN systems continuously validate whether the declared intent is being met. Telemetry is collected in real time. Performance metrics such as latency, jitter, packet loss, and throughput are evaluated against the original business objective.

If the system detects degradation, perhaps the MPLS link begins to experience jitter that threatens video quality, it does not simply raise an alarm. It recalculates. It determines alternative paths. It dynamically reroutes traffic to a better-performing circuit. It may adjust queue weights, reallocate bandwidth, or apply temporary prioritization policies. All of this occurs autonomously, within guardrails defined by governance policies. The original intent remains the north star. The system self-corrects to preserve it.

This closed-loop process operates continuously. Intent is declared once, but assurance never stops. The network becomes a self-optimizing system aligned to business outcomes rather than static configurations.

The deeper implication is that IBN shifts networking from device-centric management to intent-centric governance. Engineers move from configuring syntax to defining strategy. Their role evolves upward in abstraction. Instead of tuning individual links, they define performance objectives, risk tolerances, and segmentation models. The AI-driven control plane handles the translation and enforcement.

Intent-Based Networking therefore represents more than automation. It represents cognitive infrastructure. It embeds feedback loops into the network itself. It treats configuration not as a one-time event but as a continuously validated hypothesis: is the business outcome being achieved right now?

In this model, networking becomes less about telling the system how to behave and more about declaring what success looks like. The intelligence within the network assumes responsibility for preserving that success, adapting in real time as conditions change. That is the fundamental shift.

Self-Organizing Networks (SON)

While the term Self-Organizing Network, or SON, is most frequently associated with 5G and mobile carrier environments, its foundational principles extend far beyond radio access networks. SON provides the operational blueprint for any infrastructure that aspires to autonomy. The same logic that allows cellular networks to adapt dynamically to shifting user density and interference patterns is the logic that enables enterprise networks to become self-driving. When applied to wired, wireless, and hybrid cloud environments, SON becomes the engine that makes Intent-Based Networking operationally viable.

At its core, SON rests on three interdependent capabilities: self-configuration, self-optimization, and self-healing. Together, these pillars transform a network from a manually maintained system into an adaptive, continuously responsive platform.

Self-configuration addresses the moment of introduction: when a new device, location, or service joins the environment. In traditional networking models, deploying a new router or switch requires staging, manual configuration, policy alignment, and validation by experienced engineers. Each step introduces potential inconsistency and delay. In a SON-enabled architecture, the process is

fundamentally different. A new device powers on, authenticates itself securely, and automatically retrieves the appropriate configuration and policy set from a centralized control plane. It integrates into the network fabric without command-line intervention. This concept is familiar to many through Zero Touch Provisioning in SD-WAN deployments, but within a SON framework it becomes part of a broader automation philosophy. The device is not merely configured; it is absorbed into a living system that understands its role, applies segmentation policies, enforces security parameters, and aligns performance expectations from the moment it comes online. The human defines standards and intent once, and the system ensures consistent implementation everywhere.

Self-optimization extends beyond initial configuration into continuous performance alignment. Traditional networks operate reactively. Monitoring tools report metrics. Thresholds trigger alerts. Engineers investigate and adjust policies when degradation becomes visible. In contrast, a SON-enabled system treats optimization as an ongoing, predictive activity. It constantly analyzes telemetry streams that include traffic distribution, application response times, link utilization, latency variation, packet loss, and congestion patterns. Machine learning models detect subtle deviations from normal baselines before they become service-impacting events. Instead of waiting for a link to saturate or fail, the system proactively redistributes traffic, recalculates path selection, adjusts queue weights, or modifies bandwidth allocations to maintain optimal conditions. Optimization is no longer an episodic event tied to troubleshooting; it becomes a continuous function embedded within the control plane. Users experience stability not because the network never changes, but because it is always adjusting.

Self-healing represents the resilience layer. Failures are inevitable in distributed systems. Hardware components degrade. Circuits experience outages. Software processes encounter faults. In conventional environments, detection depends on monitoring systems that escalate alerts to operations teams, who then diagnose and remediate the issue. Time to detection and time to repair introduce service disruption. In a SON architecture, fault detection is immediate and autonomous. When a device becomes unreachable or a link drops below performance thresholds, the system identifies the anomaly in real time. It recalculates routing tables, reroutes traffic across alternate paths, enforces updated policies, and preserves session continuity wherever possible. The corrective action occurs in milliseconds, often before users perceive degradation and well before a support ticket can be created. The network effectively repairs itself within the boundaries of defined intent and governance.

When these three principles operate together, the network transitions from a static configuration domain into a dynamic organism. Self-configuration ensures consistent and rapid onboarding. Self-optimization maintains alignment with performance objectives. Self-healing preserves continuity during failure events. None of these capabilities function in isolation. They rely on shared telemetry, centralized policy models, distributed enforcement, and intelligent feedback loops.

Applying SON principles to the enterprise environment unlocks the assurance and autonomy associated with Intent-Based Networking. Assurance, in this context, is not simply monitoring. It is the active verification that declared business objectives are being met continuously. The system

measures its own state against intent, detects divergence, and course corrects automatically. The network becomes self-aware within operational boundaries.

This is where the connection between SON and IBN becomes clear. Intent defines the destination. SON provides an adaptive engine that keeps the system aligned with that destination despite changing conditions. Without self-configuration, intent could not scale. Without self-optimization, intent could not adapt. Without self-healing, intent could not persist through failure.

The convergence of these capabilities marks the evolution from programmable infrastructure to autonomous infrastructure. In such an environment, human operators move from tactical responders to strategic supervisors. They define outcomes, establish governance guardrails, and analyze systemic trends, while the network handles real-time adjustments. The enterprise network, once dependent on manual precision and reactive troubleshooting, begins to behave more like a living system that is responsive, resilient, and continuously aligned with business objectives.

The SD-WAN Advantage: Why Agencies Have a Head Start

Agencies that have already invested in SD-WAN have not just built a better network, they have laid the essential groundwork for IBN and SON. They are not starting from scratch; they are simply upgrading the "brain" of a system they already own.

Here is how an agency's existing SD-WAN components provide a direct path forward:

Existing SD-WAN Component	**How It Evolves to IBN/SON**
Centralized Orchestrator	This is the "single pane of glass" for management. In IBN, this orchestrator is infused with AI and ML, transforming it from a management dashboard into an autonomous "intent engine."
Policy-Based Engine	SD-WAN allows you to build policies for applications. IBN abstracts this one step further. The IBN system *auto generates and manages* these policies based on your high-level intent.
End-to-End Visibility	SD-WAN provides visibility into app performance and underlay/overlay paths. This visibility becomes the raw data feed for IBN's **Assurance** loop. The system uses this data to learn, predict, and act.

Automation (e.g., ZTP)	SD-WAN's automation of provisioning and path selection is the foundation for the closed-loop automation of IBN, where the network configures, heals, and optimizes itself.

The Next Frontier: Beyond 5G and Predictive Networking

While IBN and SON represent the autonomous future, the evolution doesn't stop there. As we look "beyond 5G" toward 6G and AI-native architectures, the network's role will fundamentally change again from *autonomous* to *cognitive*.

This next frontier is built on **predictive networking**.

- **Today's Autonomous Network (IBN/SON):** The network *reacts* in real time. It detects jitter and *then* moves the call. This is a *proactive* and *self-healing* model.
- **Tomorrow's Cognitive Network (Beyond 5G):** The network *anticipates* needs before they happen.

Powered by massive-scale AI and "Digital Twin" models (real time simulations of the entire network), the system will not just wait for a problem. It will run continuous simulations to predict network behavior.

For example, it knows that a large data center backup is scheduled to begin in ten minutes. It will *predict* that this transfer will cause congestion on a specific link, which would *in the future* impact an executive's video call.

Based on this prediction, the network will preemptively and transparently move the video call to an optimal path *before* a single packet is ever dropped and before any jitter could possibly occur.

In this future, dynamic policy and real-time configuration evolve into prescriptive, predictive policy. The network stops just responding to intent and starts anticipating it, truly becoming a cognitive partner that ensures outcomes are met before the user even knows there could be a problem.

Balancing Autonomy with Cybersecurity Rigor in AI Driven Networks

As the next generation of SD-WAN evolves toward Intent-Based Networking (IBN) and Self-Organizing Networks (SON), network architectures are set to benefit from remarkable levels of automation and intelligence. However, increased autonomy must be accompanied by stringent adherence to foundational security principles such as least privilege, zero trust, defense in depth, and continuous monitoring. While AI-driven systems can rapidly optimize network performance and enforce high-level directives like “ensure all executives have a high-quality, jitter-free video conferencing experience,” skilled network professionals must remain in the loop. Their expertise in understanding ports, protocols, application prioritization, and anomaly detection is essential to validate that the

system's actions truly align with business objectives and have not been compromised by malicious actors or misconfigurations.

When AI systems prioritize applications, they must rely exclusively on trusted, up-to-date data sources drawing from inventories that have been validated and are actively maintained. Failure to do so could result in prioritizing outdated tools, missing new business applications, or even inadvertently elevating malicious software. This highlights the principle of secure data provenance: the requirement that all data influencing network decisions must be authenticated and integrity checked. Similarly, automated assurance and remediation features must only act upon current, authorized lists of business applications; otherwise, automated policy changes could disrupt vital operations.

To manage these risks, cybersecurity best practices dictate a layered, proactive approach. This means applying rigorous patch management, comprehensive vulnerability assessments, and secure supply chain management, including lab-based validation of all software and hardware before deployment. Principles like change management and separation of duties become even more important as network autonomy increases, ensuring that no single point of failure or unauthorized change can undermine network integrity. Regular reviews and updates to network policies and the underlying data are essential, as is robust access control to ensure only authorized personnel can modify critical parameters.

Ultimately, as we grant our networks more autonomy to self-organize and self-heal, embedding cybersecurity at every layer through disciplined governance, multi-factor authentication, network segmentation, and continuous monitoring is critical. By maintaining this balance between intelligent automation and cybersecurity rigor, organizations can harness the benefits of next-generation networks while safeguarding their core business operations from evolving threats.

Moving to Future-Agile Networks for Federal Agencies

Taken together, these forces show that AI is not simply layering new requirements on top of old infrastructure; it is driving a fundamental redefinition of the federal network. Operations shift from manual to intelligent. WANs evolve into adaptive fabrics. Traffic growth becomes not a challenge to endure but a design principle to anticipate.

Security, too, is inseparable from this transformation. AI expands the attack surface: malicious models, poisoned data, and adversarial traffic all demand continuous monitoring and automated mitigation. Future agile federal networks must therefore integrate Zero Trust principles, policy-as-code enforcement, and AI-enabled cybersecurity defenses directly into the fabric of the network.

The future is not about managing yesterday's WAN. It is about architecting networks that thrive in the AI era and are resilient, compliant, adaptive, and mission aligned. Federal agencies that embrace this shift will find themselves not just keeping up with technological change but leading it, using AI

not only to power applications but also to strengthen the very foundations of the networks that carry them.

Top Impacts for Federal Agencies in Moving to a Future-Agile Network

Based on the context of current U.S. federal agency modernization efforts such as TIC 3.0, and OMB Zero Trust mandates, here are the top 10 issues agencies face in moving to a Future-Agile Network.

Zero Trust Architecture (ZTA) Implementation

- **The Issue:** Moving from a "castle and moat" perimeter security model to one where no user or device is trusted by default. This requires a massive cultural and technical shift to verify every access request.
- **Security Context:** Agencies must implement granular identity governance and micro-segmentation, ensuring that a breach in one area does not compromise the entire network.

Figure 2: Federal Zero Trust Architecture and NIST SP800-207

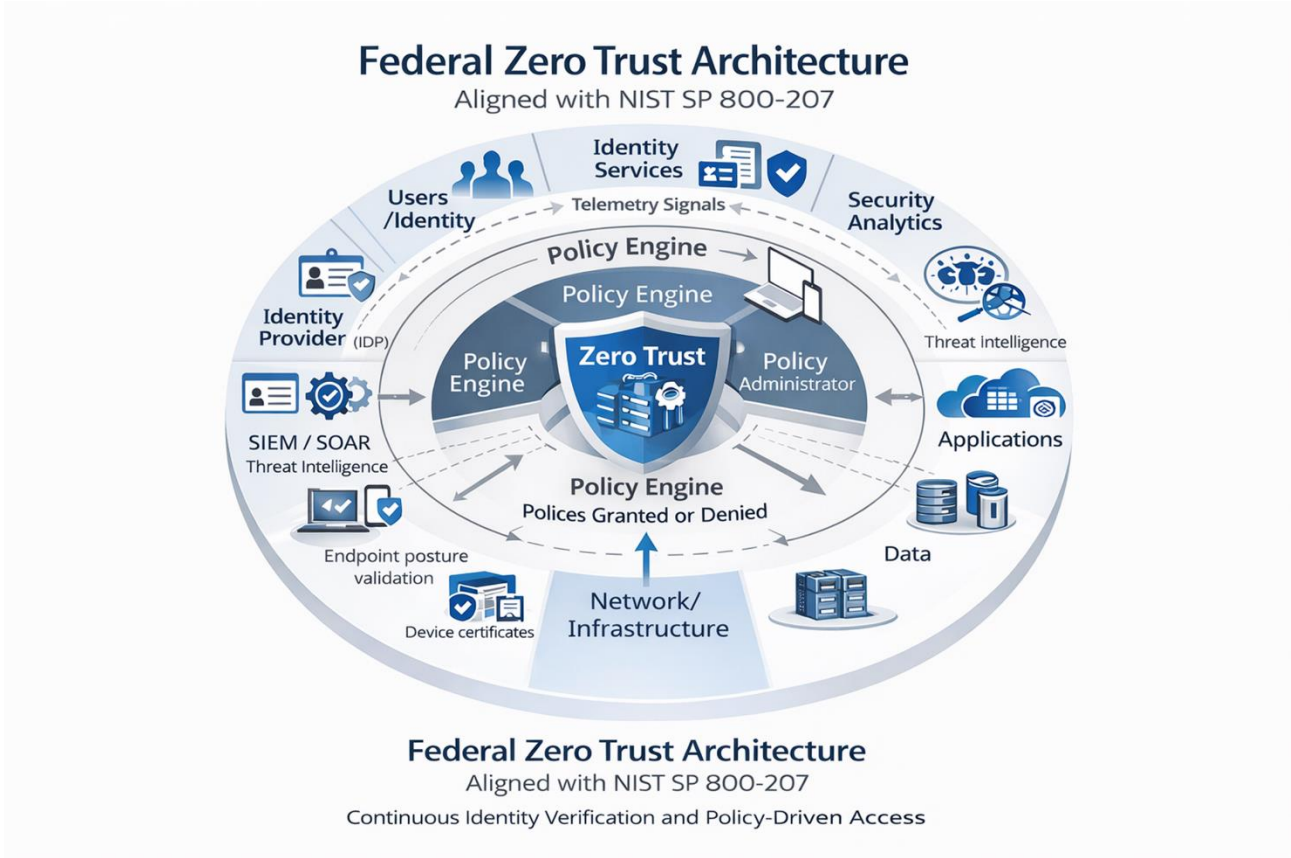

Adopting Secure Access Service Edge (SASE)

- **The Issue:** Converging wide area networking (WAN) with security services (like cloud-access security broker (CASB) and firewall as a service (FWaaS)) into a single cloud-delivered model. Agencies struggle with the complexity of replacing on-premises appliances with cloud-native enforcement points.
- **Security Context:** SASE is critical for "future agile" networks because it moves security inspection to the edge where the user is rather than backhauling traffic to a central data center, which reduces latency but requires robust cloud security policies.

Trusted Internet Connections (TIC) 3.0 Compliance

- **The Issue:** Transitioning from TIC 2.0 (routing all traffic through physical access points) to TIC 3.0 (flexible, distributed security). Agencies often face difficulties interpreting the new "trust zone" guidance and validating that distributed architectures meet CISA standards.
- **Security Context:** This shift eliminates the "bottleneck" of centralized security but requires agencies to deploy sophisticated telemetry to maintain visibility over decentralized traffic flows.

Integration of Non-Terrestrial Networks (LEO Satellites)

- **The Issue:** Incorporating Low Earth Orbit (LEO) satellite connectivity (like Starlink or Amazon LEO) to provide high-speed, low-latency redundancy for remote sites or disaster recovery.
- **Security Context:** While LEO offers agility and resilience, it introduces new attack vectors (signal jamming, interception) and requires ensuring that traffic traversing commercial satellite constellations remains encrypted and isolated from public internet traffic.

Migration to Software-Defined WAN (SD-WAN)

- **The Issue:** Replacing rigid, expensive MPLS circuits with flexible SD-WAN overlays that can route traffic dynamically over broadband, LTE, and 5G.
- **Security Context:** SD-WAN creates a more agile network but increases the attack surface. Security must be "baked in" to the SD-WAN fabric (using encryption and segmentation) rather than bolted on, preventing lateral movement if a branch node is compromised.

Managing Hybrid and Multi-Cloud Complexity

- **The Issue:** Federal data now lives everywhere on premises, in AWS/Azure/Google Cloud, and at the edge. Networking teams struggle to create a unified "network fabric" that connects these disparate environments seamlessly.
- **Security Context:** Inconsistent security policies across different cloud providers create gaps. "Future agile" networks require a unified control plane that enforces the same security rules regardless of where the data resides.

Legacy Technical Debt & Interoperability

- **The Issue:** Integrating modern, agile technologies (like 5G and WiFi 6) with decades old mainframes and legacy cabling infrastructure that cannot support high bandwidth or modern encryption standards.
- **Security Context:** Legacy systems often cannot support modern authentication methods (like MFA or FIDO2), leaving "holes" in the Zero Trust fabric that adversaries can exploit.

Supply Chain Risk Management (SCRM)

- **The Issue:** As agencies buy more commercial off-the-shelf (COTS) software and hardware for agility, they inherit the risks of their vendors.
- **Security Context:** "Future agile" relies on complex software stacks. Agencies must rigorously vet the "Software Bill of Materials" (SBOM) to ensure that the agile network controllers and edge devices do not contain hidden backdoors or vulnerabilities from foreign adversaries.

Workforce Skills Gap

- **The Issue:** The shift to "future agile" networks turns network engineers into software developers. There is a critical shortage of federal staff who understand *both* traditional networking (BGP, OSPF) and modern automation (Python, Ansible, Terraform).
- **Security Context:** Misconfigurations in software-defined networks are a leading cause of breaches. Without a skilled workforce to write secure code for network automation, agility leads to fragility.

Visibility and Telemetry at the Edge

- **The Issue:** As the network's perimeter dissolves, agencies lose the "single pane of glass" visibility they had with centralized firewalls.

- **Security Context:** To maintain security in an agile network, agencies must deploy Extended Detection and Response (XDR) tools that can collect and correlate logs from endpoints, cloud, and network edges in real time to detect anomalies.

Modernization Roadmap for U.S. Federal Agencies and the Maturity Model

Level 1	Legacy WAN, static security, limited cloud
Level 2	Multi-Cloud integration, AIOps-assisted operations, introduction of SASE Principles
Level 3	SD-WAN, Zero Trust adoption, partial automation, mature SASE adoption
Level 4	Fully future-agile intelligent edge, AI-optimized, continuous compliance, resilient convergent access

Level 1 represents the Legacy WAN environment, defined by static architecture and operational constraints. At this stage, agencies depend heavily on traditional wide area networking technologies such as MPLS circuits, leased lines, and in some cases even legacy time division multiplexing infrastructure. Security remains perimeter-centric, built primarily around static firewall deployments that provide limited visibility into east-west traffic or internal segmentation risks. Cloud adoption is minimal and cautious, typically confined to isolated pilot projects or basic SaaS tools rather than mission-critical systems. Network provisioning is manual, often requiring lengthy change management cycles and coordination across multiple teams. While this environment offers predictability for legacy applications, it lacks flexibility. Scalability is limited, resiliency is dependent on fixed circuits, and security posture is vulnerable to modern threats that exploit static defenses. The architecture is stable but rigid, increasingly strained by the demands of artificial intelligence workloads, distributed mobility, and cloud-native application patterns.

Level 2 marks the Transitional WAN phase, characterized by controlled experimentation and incremental modernization. Agencies begin introducing SD-WAN technologies to reduce reliance on fixed circuits and improve routing agility. Traffic steering becomes more dynamic, and cost optimization opportunities emerge as broadband and hybrid transport options supplement traditional MPLS. At the same time, Zero Trust concepts begin replacing implicit trust models, shifting access control toward identity-centric authentication and segmentation. Automation tools are introduced, but their scope remains limited. They may include configuration scripts, basic orchestration tools, or alert-driven workflows, but they do not yet constitute fully autonomous or closed-loop systems. Cloud adoption expands during this phase, although workloads are carefully segmented and migration decisions remain risk-averse. Legacy and modern environments operate in parallel, creating operational complexity. Artificial intelligence plays a limited role, often confined to targeted analytics, reporting dashboards, or isolated proof-of-concept deployments. Agencies at this

level gain flexibility and reduce certain costs, but fragmentation persists across platforms and policy domains.

Level 3 represents the Integrated WAN stage, where modernization becomes cohesive and intelligence begins operating at scale. Agencies leverage scaled SD-WAN deployments to integrate multi-cloud environments, distributing workloads across providers such as AWS, Azure, Google Cloud, and government cloud platforms while maintaining consistent security controls and policy enforcement. Networking and security converge through frameworks like Secure Access Service Edge, allowing unified protection for distributed users, remote sites, and mobile workforces. Artificial Intelligence for IT Operations becomes operational rather than experimental. AIOps platforms ingest telemetry across domains, providing anomaly detection, predictive capacity modeling, and automated remediation workflows. Visibility expands dramatically. Agencies can observe application behavior, user experience metrics, and infrastructure performance across on-premises and cloud environments in near real time. The network evolves from a static transport utility into an adaptive environment that anticipates shifting traffic patterns, including those generated by AI-driven analytics and machine learning workloads. At this level, infrastructure decisions are increasingly aligned with mission objectives rather than constrained by legacy limitations.

Level 4 defines the Future Agile WAN, where networks are designed for continuous adaptation and mission alignment. Edge computing becomes integral, with AI processing occurring closer to data sources to reduce latency and preserve bandwidth. Only essential or aggregated traffic flows back to centralized cores. Artificial intelligence is embedded deeply within both operations and security layers, dynamically optimizing routing paths, enforcing compliance policies in real time, and predicting capacity demands with high precision. Compliance processes evolve from periodic audits to continuous validation models driven by telemetry and automated reporting, ensuring adherence to federal standards such as FISMA and FedRAMP. The access ecosystem becomes fully resilient, blending fiber, satellite, 5G, emerging 6G connectivity, and mobile edge technologies into a unified transport fabric. In this mature state, human operators shift from tactical troubleshooting to strategic oversight. AI systems manage daily optimization, fault response, and performance alignment within governance guardrails. The network is no longer merely prepared for AI workloads; it is architected around them. It becomes a mission-aligned, intelligence-optimized platform capable of adapting to disruption, scaling under pressure, and sustaining operational readiness under any condition.

Conclusion

The transition from traditional WAN architecture to SD-WAN solutions that leverage TIC 3.0 principles offer an excellent opportunity for federal agencies to modernize their network infrastructure. This shift not only offers significant cost savings and operational efficiencies but also lays the groundwork for future advancements in networking, such as autonomous, self-driving networks. By leveraging SD-WAN, organizations can achieve greater agility, centralized control, and enhanced visibility, positioning themselves for growth and innovation in an increasingly digital and cloud-centric world. Furthermore, the adoption of SD-WAN can lead to improved mission effectiveness and citizen

services, particularly in remote or challenging environments, by providing more reliable and responsive connectivity.

As federal agencies navigate the complexities of modernization, the financial clarity and strategic insights provided by this analysis will be crucial in making informed decisions that align with their long-term goals. The comprehensive ROI framework offered in this paper can assist decision-makers with quantifying the tangible business value of transitioning to SD-WAN and cloud-based SASE solutions. This financial analysis framework helps agencies understand the cost-benefit dynamics of modernizing their network infrastructure, ensuring that investments are strategically aligned with operational needs and future growth objectives. By addressing the challenges posed by legacy systems and embracing innovative network solutions, agencies can overcome past barriers and lead in both operational excellence and technological advancement.

Finally, the FAN model articulated in this white paper underscores the importance of designing network architectures that are resilient, secure, and adaptive. By embedding Zero Trust principles and leveraging intelligent routing, redundant paths, and multi-bearer connectivity, the FAN model ensures that government agencies can maintain operational continuity and security in the face of disruptions. This architectural philosophy not only supports the dynamic demands of modern IT environments but also aligns with the unique operational requirements of government agencies. As emerging technologies such as AI, QC, LEO satellite constellations, and future Wi-Fi and cellular generations continue to reshape the networking landscape, the FAN model provides a robust framework for agencies to remain agile, mission-aligned and prepared for the future.

Authors & Affiliations

This white paper was written by a consortium of government and industry. The organizational affiliations of these contributors are included for information purposes only. The views expressed in this document do not necessarily represent the official views of the individuals and organizations that participated in its development.

Stephen Matherne	Spectrum Business
Erik Hand	AT&T
Phil Moser	Fortinet Federal
Christopher Grindstaff	AT&T
Mohamed Hersi	ACI Solutions
Joseph Tahan	Spectrum Business
Scott Andersen	Creative Technology and Innovation

References

ACT-IAC. 2019. *Zero Trust Cybersecurity: Current Trends.* April. https://www.actiac.org/documents/zero-trust-cybersecurity-current-trends.

IDC. 2022. *The Business Value of Cisco SD-WAN.* August. Accessed June 16, 2026. https://www.cisco.com/c/en/us/solutions/enterprise-networks/sd-wan/idc-business-value-white-paper.html.

2020. *NIST SP 800-53 Rev. 5.* December 10. https://csrc.nist.gov/pubs/sp/800/53/r5/upd1/final.

n.d. *NIST SP800-215.* https://csrc.nist.gov/pubs/sp/800/215/final.

Thorne, Elizabeth. 2020. *TeleGeography.* May 14. Accessed June 16, 2026. https://resources.telegeography.com/wan-mythbusters-is-it-true-that-sd-wan-can-cut-your-network-spend-in-half?utm_source=chatgpt.com.

2025. *Trusted Internet Connections (TIC) 3.0 Core Guidance Documents.* July 2. https://www.cisa.gov/resources-tools/resources/trusted-internet-connections-tic-30-core-guidance-documents

Appendix — Acronyms & Abbreviations

5G / 6G: Fifth / Sixth Generation; cellular network standards providing high-bandwidth, low-latency wireless connectivity (6G is emerging/future).

AI: Artificial Intelligence; systems that perform tasks requiring human-like intelligence (e.g., prediction, classification, optimization).

AIOps: Artificial Intelligence for IT Operations; applying AI/ML to operations data to detect anomalies, predict incidents, and automate remediation.

ATO: Authorization to Operate; formal approval (risk acceptance decision) permitting a system to operate in a production environment.

AWS: Amazon Web Services; a major cloud service provider (CSP).

BGP: Border Gateway Protocol; the primary inter-domain routing protocol used on the Internet and in large enterprise cores.

BIS: Broadband Internet Service; typically, lower-cost Internet access (cable/DSL/fixed wireless) used as primary or secondary transport.

CASB: Cloud Access Security Broker; security controls/policy enforcement between cloud consumers and cloud providers.

CIO: Chief Information Officer; senior executive accountable for IT strategy and operations.

CISO: Chief Information Security Officer; senior executive accountable for cybersecurity strategy, risk, and security governance.

CJIS: Criminal Justice Information Services; (often referenced as CJIS policy); requirements for protecting criminal justice information.

CCM: Counter with CBC-MAC; an authenticated encryption mode (commonly paired with AES).

COTS: Commercial Off-The-Shelf; commercially available products/services procured without custom development.

CSP: Cloud Service Provider; entity delivering cloud services (IaaS/PaaS/SaaS).

DAS: Distributed Antenna System; indoor/outdoor cellular coverage enhancement architecture.

DDoS: Distributed Denial of Service; attack that overwhelms a target using traffic from many distributed sources.

DIA: Dedicated Internet Access; business-class Internet service with defined bandwidth/SLA (often symmetrical).

DoS: Denial of Service; attack that degrades or interrupts service availability (not necessarily distributed).

EIS: Enterprise Infrastructure Solutions; a GSA government-wide contract vehicle for telecom/networking services.

FAN: Future-Agile Network; as defined in this document, an architectural/operational model emphasizing resiliency, security-by-design, adaptability, and mission alignment.

FedRAMP: Federal Risk and Authorization Management Program; standard approach for security assessment, authorization, and continuous monitoring for cloud offerings used by U.S. federal agencies.

FIDO2: Fast Identity Online 2; passwordless authentication standards using strong cryptography (e.g., security keys, platform authenticators).

FIPS: Federal Information Processing Standards; U.S. government standards for security/crypto and related IT requirements.

FISMA: Federal Information Security Modernization Act; U.S. law requiring federal information security programs and risk-based protections.

FWaaS: Firewall as a Service; cloud-delivered firewall and policy enforcement (common component within SASE).

GCM: Galois/Counter Mode; authenticated encryption mode (commonly AES-GCM) providing confidentiality + integrity.

IBN: Intent-Based Networking; networking model where operators declare desired outcomes ("intent") and the system translates/enforces/assures those outcomes continuously.

IaaS: Infrastructure as a Service; cloud model delivering compute/network/storage building blocks.

IoT: Internet of Things; connected sensors/devices (often at edge sites) producing data and requiring network access.

IP: Internet Protocol; foundational networking protocol for addressing and routing packets.

ISP: Internet Service Provider; company that provides Internet connectivity services.

LEO: Low Earth Orbit; satellite orbits typically ~160–2,000 km altitude; often used for lower-latency satellite broadband compared to GEO.

LTE: Long-Term Evolution; 4G cellular standard used widely for primary/backup WAN transport.

Mbps: Megabits per second; unit of data rate (bandwidth).

MFA: Multi-Factor Authentication; authentication using two or more factors (something you know/have/are).

ML: Machine Learning; subset of AI using statistical models to learn patterns from data.

ML-DSA: Module-Lattice Digital Signature Algorithm; post-quantum digital signature family referenced in emerging standards discussions.

ML-KE: Module-Lattice Key Encapsulation / Key Exchange; post-quantum key establishment family referenced in emerging standards discussions.

MPLS: Multiprotocol Label Switching; WAN transport technique using label switching for traffic engineering, VPNs, and predictable pathing.

MTIPS: Managed Trusted Internet Protocol Services; a managed secure Internet access service approach associated with federal TIC requirements (commonly ordered under EIS).

NIST: National Institute of Standards and Technology; publishes security standards and guidance widely used by federal agencies.

OMB: Office of Management and Budget; federal office that issues policy memos affecting agency IT/security programs (including cloud and TIC-related guidance).

OOB: Out-of-Band most commonly refers to "out of band" management in IT (using a separate, secure channel for network administration).

OSPF: Open Shortest Path First; common interior gateway routing protocol (IGP) in enterprise networks.

OT: Operational Technology; systems controlling physical processes (industrial/utility environments).

PaaS: Platform as a Service; cloud model delivering managed platforms/runtime services for applications.

PEP: Policy Enforcement Point; where access/security policy is enforced (e.g., gateway, agent, proxy, cloud edge).

PQC: Post-Quantum Cryptography; cryptographic algorithms designed to resist attacks by quantum computers.

QC: Quantum Computing; computing approach leveraging quantum mechanics that may threaten certain classical cryptography at scale.

QoS: Quality of Service; traffic handling mechanisms to prioritize/shape/queue traffic for performance (latency/jitter/packet loss control).

RFP: Request for Proposal; procurement document inviting vendors to submit bids/solutions.

ROI: Return on Investment; financial metric comparing gains/benefits to costs.

SaaS: Software as a Service; cloud model delivering complete applications over the Internet.

SASE: Secure Access Service Edge; architecture converging WAN connectivity and cloud-delivered security services (often includes SD-WAN, SWG, CASB, ZTNA, FWaaS).

SBOM: Software Bill of Materials; list of components/dependencies in software used for supply chain risk management.

SC / SCRM: Supply Chain Risk Management; processes to manage vendor/component risks across hardware/software/services.

SD-WAN: Software-Defined Wide Area Network; overlay architecture enabling centralized control, application-aware routing, and multi-transport path selection.

SON: Self-Organizing Network(s); networks capable of self-configuration, self-optimization, and self-healing (often used in cellular contexts but applied more broadly in your text).

SP: Special Publication; NIST document series (e.g., NIST SP 800-53) providing standards/guidance.

TDM: Time Division Multiplexing; legacy transport technique multiplexing multiple signals by time slots (common in older telecom networks).

TIC: Trusted Internet Connections; federal initiative/guidance for securing external connections and modernizing security perimeters (e.g., TIC 3.0 use cases).

WAN: Wide Area Network; network connecting geographically separated sites (branches, data centers, agencies, cloud regions).

Wi-Fi 6 / Wi-Fi 7: Wireless LAN standards (802.11ax / 802.11be); providing higher throughput, lower latency, and improved efficiency.

XDR: Extended Detection and Response; security capability correlating telemetry across endpoints, network, cloud, and identity to detect/respond to threats.

ZTA: Zero Trust Architecture; security model where no implicit trust is granted; continuous verification of users/devices/apps and least-privilege access.

ZTNA: Zero Trust Network Access; technology pattern to provide application-specific access based on identity, device posture, and policy—without broad network access.

ZTP: Zero Touch Provisioning; automated onboarding/provisioning of devices without manual configuration (common in SD-WAN).