

The Journey to Self-Driving Networks: Agentic AI as a Catalyst for Operational Efficiency and Cost Reduction

Peter Fetterolf PhD



Executive Summary

This paper examines the evolution of HPE Networking's transition into a new era of Agentic AI, highlighting both the technological milestones and the economic benefits of transitioning to self-driving networks.

Hewlett Packard Enterprise (HPE), building on Aruba Central, Juniper Mist, and GreenLake Intelligence, has pioneered AI-driven operations for more than a decade. Components of what HPE calls the self-driving network, these mature solutions reduce operational complexity, improve user experience, and enable measurable labor savings. The introduction of Agentic AI marks a step-change from AI-assisted functions toward fully autonomous, self-driving networks.

Agentic AI brings together intelligent agents, digital twins, large experience models, and conversational AI to orchestrate network operations with unprecedented precision and adaptability. Rather than reacting to issues, networks can proactively prevent failures, self-heal, and optimize service delivery.

The economic analysis demonstrates compelling benefits. In enterprise networks, the transition to self-driving capabilities yields up to **52% TCO reduction over five years**, driven primarily by labor efficiency gains and dramatic OpEx savings. WAN optimization adds further impact: through intelligent QoE-based traffic management using HPE SASE solutions and digital twinning, enterprises achieve a **24% TCO reduction** and payback in less than a month.

These findings confirm that the adoption of Agentic AI not only enhances operational performance but also provides a strong financial case for enterprises and service providers. The combination of reduced downtime, faster mean-time-to-resolution, and more efficient bandwidth utilization delivers rapid ROI and positions organizations for scalable, future-ready network operations.

Background and Evolution of HPE AIOps

For over a decade, Hewlett Packard Enterprise (HPE) has maintained a leadership role in pioneering network-centric AI operations (AIOps). Core to HPE Networking's Secure, AI-Native Network platform delivering 'AI for Networks and Networks for AI Workloads', are industry leading AIOps technologies from HPE Aruba, HPE Juniper Mist, and GreenLake Intelligence. To deliver HPE Networking's vision of the Self-Driving Network, these technologies collectively form the most mature platforms for enabling autonomic, self-driving network functions that target both operational efficiency and end-user quality of service.

Founding, acquisitions, and AI technologies

- **Aruba Networks** was founded in February 2002, quickly gaining recognition for its innovative approach to secure mobility and enterprise Wi-Fi. Aruba was acquired by HPE in 2015, forming the foundation of its Intelligent Edge business. HPE Aruba's AI technology delivers intelligent, security-first networking through a cloud-native platform that leverages advanced analytics and automation. At the core is Aruba Networking Central, which uses AI to optimize performance, enhance security, and simplify operations across wired, wireless, and WAN environments.
- **Mist Systems** was founded in 2014 to overcome the limitations of traditional network management, focusing on proactive issue resolution and enhanced user experience visibility. Acquired by Juniper Networks in early 2019, Mist - powered by Marvis AI - became the foundation for integrating advanced telemetry, AI and machine-learning, and cloud-native architectures into a unified operations platform. HPE Networking announced the acquisition of Juniper Networks in July of 2025.
- **HPE GreenLake Intelligence** was officially introduced on June 24, 2025, at HPE's Discover conference as a transformative evolution of the GreenLake platform, designed to bring agentic AI to hybrid IT operations. GreenLake Intelligence builds on HPE's existing cloud-native infrastructure and integrates technologies from recent acquisitions and innovations in networking, storage, and observability. By leveraging these capabilities, the platform delivers a unified, AI-driven operating model that proactively manages, optimizes, and secures workloads across multi-cloud and on-premises environments.

¹ <https://blog.telegeography.com/used-international-bandwidth-reaches-new-heights>

A foundational, cloud-native microservices architecture is central to HPE Networking's strategy for seamlessly integrating and cross-pollinating innovations across a unified AIOps platform. Early in the journey, HPE Networking is doubling down on a "develop once, deploy everywhere" model with enhancements they will showcase at HPE Discover in Barcelona, Spain in December.

Deployment Scale and Customer Adoption

- **Large enterprise and diverse verticals:** HPE Networking's Self-Driving Networks have been adopted by a wide array of organizations in retail, education, healthcare, financial services, and other distributed enterprise environments. Customers such as PetSmart (U.S.), IKEA, Stanford University, Major League Baseball, and Williams-Sonoma have been cited among early adopters¹.
- **Managed networks and MSPs:** HPE Networking's Self-Driving Networks have also been widely adopted by managed service providers who benefit from operational simplicity.

Focus on Service Quality and Application Performance

A distinct differentiator HPE Networking's Self-Driving Network is the orientation toward user-centric service metrics (often expressed as Service Level Expectations, or SLEs) rather than purely device-availability or raw throughput. The platform tracks metrics such as connection success rate, latency, application responsiveness, and end-to-end user experience. Issues are not just detected but diagnosed and remediated—either automatically or via guided workflows—across wireless, wired, data center and WAN routing domains. In multiple deployments (e.g. ServiceNow), end users' visibility into performance (from device through cloud) has improved markedly, helping not only in issue resolution but in proactive prevention of service degradations².

¹ https://www.futuriom.com/articles/news/juniper-acquisition-of-mist-networks-accelerates-ai-networking-arms-race/2019/03?utm_source=chatgpt.com

² <https://www.juniper.net/content/dam/www/assets/case-studies/us/en/2023/servicenow-case-study.pdf>

Agentic AI and Self-Driving Networks

From AI-Assisted to Fully Autonomous Networks

Over the last decade, HPE has demonstrated measurable benefits in operational efficiency and user experience assurance. The next phase in this journey is the integration of **Agentic AI**—a paradigm shift that transforms automation into autonomous, reasoning-driven operations. Where today's networks are "AI-assisted," the addition of agentic capabilities accelerates the transition toward the **Self-Driving Network**.

Unlike deterministic automation or even traditional AIOps, agentic AI introduces intelligent agents capable of reasoning, collaboration, and self-learning across diverse domains such as switching, routing, wireless, and WAN. These agents function as domain-specific experts, loosely coupled but coordinated through graph-based workflows. This architecture enables dynamic orchestration: decomposing complex operational challenges into subtasks, delegating them to the right agents, and synthesizing results into precise, real-time actions.

Core Enablers of Agentic AI in Networking

HPE Networking's approach emphasizes that agentic AI is not a standalone bolt-on technology but a convergence of key enablers:

- **High-quality, real-time telemetry:** A cloud-native platform provides deep visibility into user and application experience, forming the raw material for intelligent decision-making.
- **Digital twins:** Always-on simulations of user behavior proactively identify performance degradations, feeding insights into AI-driven root cause analysis before issues impact end users.
- **Large Experience Models (LEM):** Purpose-built models, trained on billions of data points, forecast collaboration quality in platforms like Zoom and Teams, with transparent explanations (e.g., SHAP) that link network factors to user outcomes.
- **Conversational AI:** Powered by large language models (LLMs), NLP, and generative AI, provides natural language access to advanced troubleshooting and multi-step workflows.
- **Open, extensible ecosystem:** 100% API-driven architecture and support for emerging protocols such as the Model Context Protocol (MCP) ensure seamless integration across multivendor environments and external systems.

Together, these enablers position HPE Networking to deliver a distributed system of autonomous network agents that can proactively prevent outages, self-heal misconfigurations, and continuously optimize service performance.

Benefits of Agentic AI for Network Operations

The economic and operational implications are significant:

- **Speed:** Rapid detection, analysis, and remediation of complex problems reduce mean time to resolution from hours to minutes.
- **Productivity:** By offloading routine troubleshooting, IT teams can redeploy resources toward innovation and business transformation.
- **Extensibility:** New AI agents can be added incrementally, allowing organizations to adopt autonomy gradually without disrupting workflows.
- **Innovation:** With continuous simulation, prediction, and automated remediation, enterprises can adopt a proactive, experience-first operating model.

Building Trust on the Path to Self-Driving Networks

Much like the adoption of autonomous vehicles, confidence in agentic AI will be built progressively. HPE Networking AI already provides human-in-the-loop (HITL) controls, where IT staff validate AI-driven recommendations before autonomy is fully delegated. Over time, as agents consistently deliver reliable outcomes, enterprises will advance along a five-stage maturity curve (Figure 1) —from assisted-driving functions to fully autonomous self-driving operations.

Journey to Self Driving

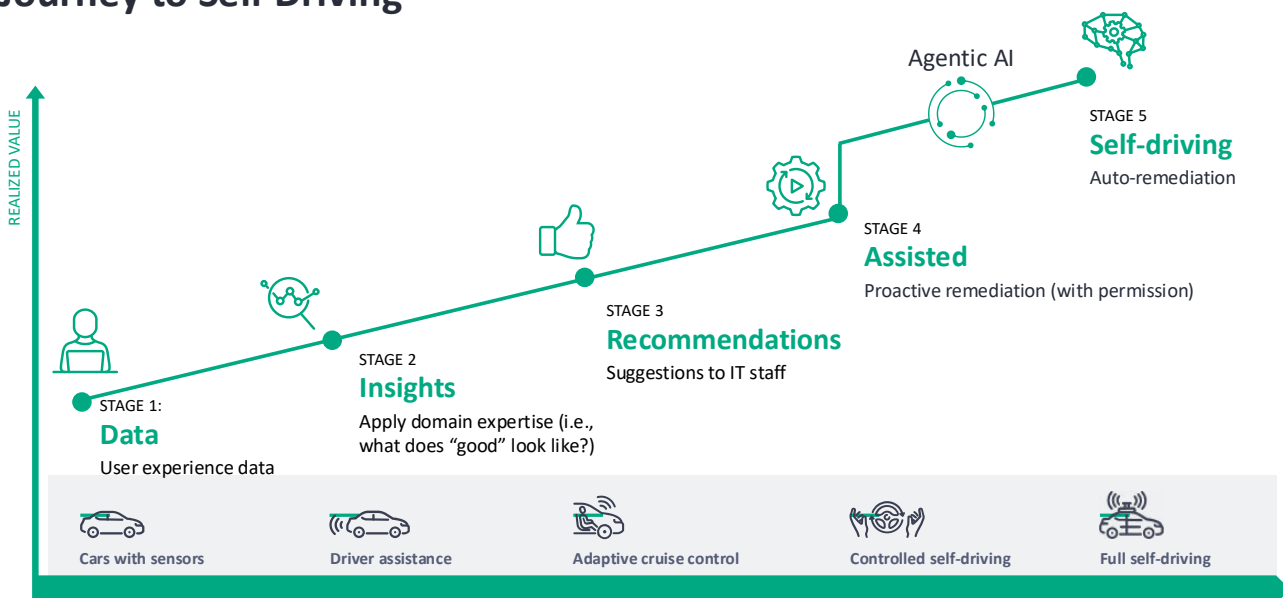


Figure 1: Five-stage journey to self-driving networks

Strategic Outlook

By uniting AI-native infrastructure, agentic intelligence, and an open ecosystem, HPE Networking is uniquely positioned to lead the industry into the era of self-driving networks. The integration of agentic AI into HPE Networking AI is not only a technological milestone but also an economic catalyst, enabling enterprises and service providers to scale operations without proportional increases in staff or cost. The result is a new model for networking: one that is proactive rather than reactive, user-experience driven rather than device-centric, and autonomous rather than manual.

Economic Benefits of Self-Driving Network

Self-driving networks enabled by Agentic AI will provide significant economic benefits. In this paper we consider two primary areas of cost savings:

- Cost savings due to further reductions in labor expenses to manage and support networks
- Cost savings due to WAN network optimization enabled by Agentic AI

We have developed TCO models that quantify the expense savings in each of these areas that are presented in the following sections of this paper.

Expense Savings due to Labor Efficiencies

HPE Networking AI has proven to reduce many network operations expenses primarily in the areas of fault and performance management. With Agentic AI and self-driving the scope and quantity of labor savings expand. Table 1 specifies that categories of labor savings, the current savings with assisted driving AIOps, and the future savings with Agentic AI and self-driving.

Labor Categories	Function/Benefit	Assisted Driving AIOps	Agentic Ai Savings (Self-Driving)
System Deployment & Configuration	Automation speeds up setup and reduces error	50%	90%
Help Desk	AI reduces user-reported issues and automates ticket handling	70%	90%
Fault & Performance Management	Proactive issue identification and remediation, root cause analysis	70%	90%
Change Management	Automated and validated network changes	20%	60%
Software Upgrades	Automation speeds up upgrades and reduces errors leading to network outages.	30%	60%
On-Site Visits & Hardware Replacement	Remote troubleshooting and AI-driven diagnostics reduce the need to dispatch	70%	90%

Table 1: Key areas of labor savings with traditional Mist-AI and Agentic Mist-AI with self-driving

We use these labor savings assumptions to compare three scenarios:

1. Networks without AIOps
2. Networks with Assisted Driving AIOps
3. Networks with enhanced Agentic AI with self-driving

In these scenarios we show the incremental TCO benefits in the journey to self-driving networks.


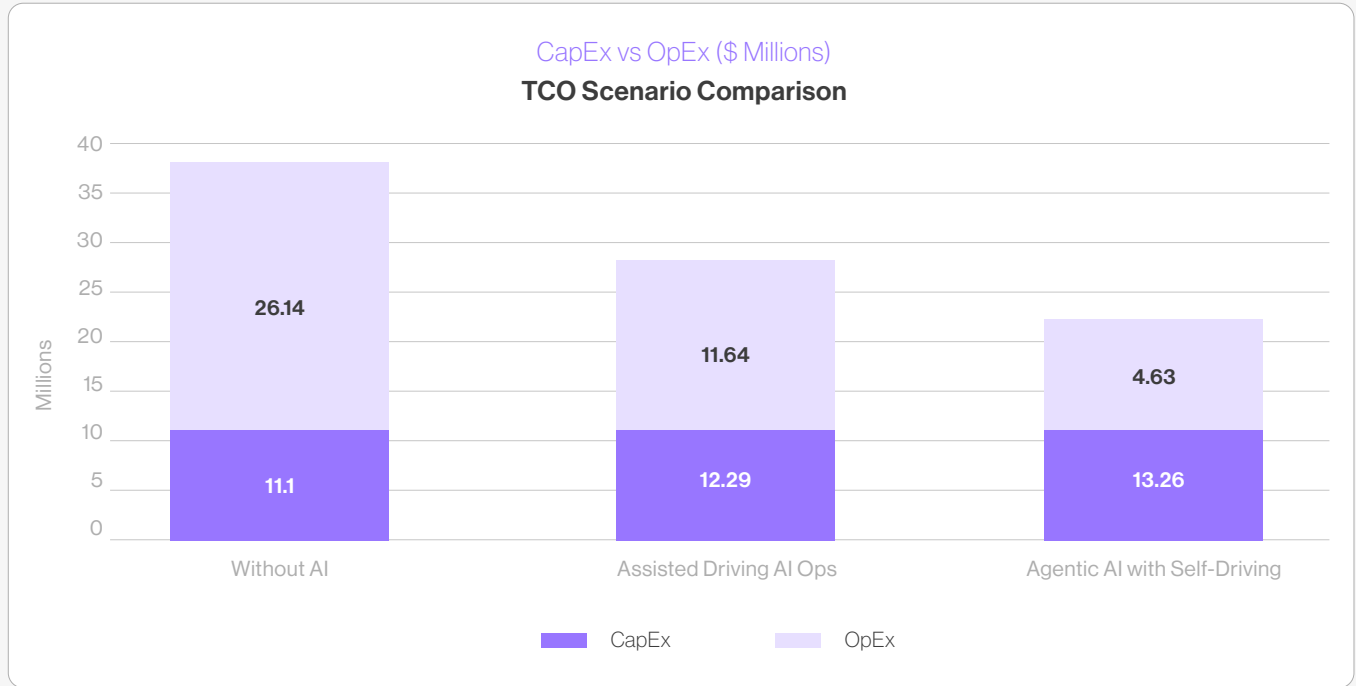
The TCO analysis uses a hypothetical network that would resemble a large enterprise network. The number of offices, types of offices, and number of access points and switches is provided in Table 2.

Network Dimensions	Number of sites	Access Points per site	Switches per site	Total APs	Total Switches
Small Branch Office	500	5	1	2500	500
Medium Branch Office	100	30	4	3000	400
Corporate HQ	3	250	20	750	60
Campus	1	500	50	500	50
Totals				6750	1010

Table 2: Network Dimensions for the TCO comparison

The results of the TCO analysis are presented in Figure 2, and Table 3.

Five-Year Cumulative CapEx vs OpEx Breakdown



Scenario	CapEx (5 Years)	OpEx (5 Years)	Total TCO	OpEx Savings
Without AI	\$11.10M	\$26.14M	\$37.24M	Baseline
Assisted Driving AI Ops	\$12.29M	\$11.64M	\$23.93M	55.5%
Agentic AI Self-Driving Networks	\$13.26M	\$4.63M	\$17.89M	82.3%

OpEx reduction is the main source of savings; CapEx slightly increases but is outweighed by total cost reduction.

Self-driving networks reduce operational costs by 82.3% while only increasing capital expenses by 19.5%

Figure 2: Five-year cumulative CapEx and OpEx comparison of three scenarios

Five-Year Cumulative Savings & Impact Summary

Metrics	Without AI	Assisted Driving AIOps	Agentic AI with Self-Driving Networks
Total Cost of Ownership (5-Year)			
Total TCO	\$37M	\$24M (35.7% savings)	\$18M (52.0% savings)
CapEx	\$11M	\$12M (10.7% increase)	\$13M (19.5% increase)
OpEx	\$26M	\$12M (55.5% savings)	\$5M (82.3% savings)
Cost Savings vs Assisted Driving AIOps			
TCO Savings (5-Year)	Baseline	\$13M	\$19M
OpEx Savings (5-Year)	Baseline	\$15M	\$22M
CapEx Difference	Baseline	-\$1M	-\$2M
Payback Period	N / A	~2 months	~2 months
Business Impact			
Network Downtime Reduction	Baseline	60-70%	80-90%
IT Staff Efficiency Gain	Baseline	40-50%	70-80%
Mean Time to Resolution	Baseline	50% faster	75% faster

All figures based on real enterprise scenario model over 5-year period.

Self-driving networks deliver the highest ROI with 82.3% OpEx reduction and significant business impact improvements.

Table 3: Five-year cumulative savings and impact summary

Summary of Results

- Over five years, self-driving networks yield TCO savings of \$19.35M (52% reduction), while assisted driving approaches yield \$13.31M (35.7% reduction).
- Major savings come from an 82.3% reduction in OpEx with self-driving networks; higher CapEx is offset by rapid ROI (18-24 months).
- Key business impacts include IT staff efficiency, reduced network downtime, faster issue resolution, and improved end-user experience.

Expense Savings due to WAN Optimization

A key expense for most enterprises is on-going monthly WAN expenses. These expenses can include internet, MPLS, and L2/L3 VPN expenses. WAN expenses are a direct function of network traffic and rise with the growth of traffic. WAN bandwidth has direct impact on user quality of experience (QoE). Fundamentally there are two approaches to improving user QoE:

1. **Brute Force:** increase network capacity across the board so utilization is 50% or lower - this is a very expensive solution.
2. **Smart QoE:** Use Agentic AI with Marvis Minis and various performance metrics to optimize the network and QoS policies to improve QoE at a much lower cost than brute force.

HPE Networking's full-stack Campus & Branch solution including wired and wireless access, connected security, network access control (NAC), and indoor location services features intelligent SD-WAN routing. Using smart QoE, HPE SD-WAN routers can be used to provide adequate bandwidth to applications that need high throughput and low latency and jitter (like Zoom) while providing best effort to applications that do not affect QoE (email, file downloads, etc). While this can be done manually, in practice it is difficult to monitor application performance and set the correct parameters across a large network. Agentic AI has the advantage of having access to a wealth of network performance data using digital twin technology (e.g., Marvis Minis) and the ability to automatically adjust the parameters in SD-WAN routers to achieve smart QoE.

To quantify the Agentic AI savings in WAN networks we have created a TCO model of an international MPLS network with the dimensions specified in Table 4. The WAN uses HPE SD-WAN routers powered by Mist to optimize WAN bandwidth while maintaining adequate user QoE. Using Agentic AI, Marvis Minis, and HPE SD-WAN routers we can increase average WAN utilization from 50% to 80% based on estimates from ACG research.

Country	Number of Sites	MPLS Bandwidth
USA	10	800 Mbps
China	5	500 Mbps
Germany	2	150 Mbps
France	3	150 Mbps
Brazil	1	150 Mbps

Table 4: Location and MPLS bandwidth of enterprise WAN sites

The results of the TCO and ROI analysis is presented in Figure 3 and Figure 4.

WAN TCO Analysis: Executive Summary

Traditional WAN vs HPE Juniper Agentic AI Optimization

Key Findings

- ✓ HPE SD-WAN routers with agentic AI optimization deliver **24% reduction** in 5-year TCO
- ✓ CapEx savings of **11%** and OpEx savings of **26%** over 5 years
- ✓ Smart QoE prioritization optimizes bandwidth while maintaining application performance
- ✓ **872% ROI** with 1 month payback period through intelligent network optimization
- ✓ Monthly recurring cost reduction through agentic AI-driven traffic management
- ✓ Improved user QoE without costly bandwidth expansion through intelligent traffic prioritization

Bottom Line Impact:

Total 5-Year Traditional WAN Cost: \$1.69M

Total 5-Year HPE Juniper Optimized WAN Cost: \$1.28M


Total 5-Year Savings: \$404.82K

WAN TCO Analysis | September 2025

Figure 3: WAN TCO analysis of Agentic AI and self-driving networks benefits

5-Year TCO Comparison

Traditional WAN vs HPE Juniper Agentic AI Optimization



Traditional WAN

CapEx	\$215.86K
OpEx	\$1.47M
Total 5-Year Cost:	\$1.69M



HPE Juniper Agentic AI Optimization

CapEx	\$192.35K
OpEx	\$1.09M
Total 5-Year Cost:	\$1.28M

Total 5-Year Savings: \$404.82K (23.99% reduction)

Smart QoE with agentic AI optimization improves application performance with optimized bandwidth allocation

CapEx Savings:

\$23.51K (10.89%)

OpEx Savings:

\$381.31K (25.91%)

WAN TCO Analysis | September 2025

Figure 4: Five-year cumulative CapEx and OpEx savings due to Agentic AI and self-driving networks

Agentic AI and self-driving networks deliver substantial total cost of ownership (TCO) savings and strong return on investment (ROI) for WAN environments, as evidenced by the attached analysis.

TCO Benefits

- Deploying HPE full-stack Campus & Branch SD-WAN routers with agentic AI optimization yields a 24% overall cost reduction—equivalent to \$404.82K saved over 5 years compared to traditional WAN solutions.
- CapEx and OpEx both decrease, especially monthly recurring costs, due to intelligent traffic management and efficient bandwidth utilization.
- OpEx savings are achieved through smarter Quality of Experience (QoE) prioritization, optimizing bandwidth while maintaining or improving application performance without expanding costly infrastructure.
- The largest TCO impact is from a 26% reduction in operating expenses driven by agentic AI.

ROI Benefits

- The solution offers fast ROI with payback achieved within the first month of deployment.
- Agentic AI enables improved application performance, enhanced user experience, and automatic optimizations that reduce IT management overhead by up to 26%.
- The system delivers consistent monthly savings while outperforming the traditional brute force approach, which requires much higher spend to match performance.
- Enhanced scalability and network security posture are realized through efficient bandwidth use and AI-driven anomaly detection.

In summary, agentic AI and self-driving networks unlock robust financial benefits, operational efficiencies, and scalability for modern WAN deployments, with clear evidence supporting rapid migration to these architectures

Conclusion

Agentic AI represents the next logical phase in the journey to self-driving networks. By extending HPE Networking's secure, AI-native network beyond automation into distributed, reasoning-driven intelligence, enterprises gain a platform that is simultaneously proactive, scalable, and economically transformative.

The quantified benefits—82.3% reduction in OpEx and over 50% TCO savings in enterprise environments and 24% cost reductions in WAN deployments—illustrate that this transition is not just a technological upgrade but a business imperative. Beyond financial impact, agentic AI enables IT teams to focus on innovation, reduces dependence on manual troubleshooting, and ensures consistently higher levels of end-user experience.

For organizations navigating increasing complexity and rising costs, immediate adoption of AI-driven operations is recommended, with a roadmap toward full self-driving capabilities. HPE Networking, through agentic AI extensions, provides a proven, extensible framework to achieve these outcomes.

In short, self-driving networks are no longer a future aspiration—they are an attainable reality that delivers both operational excellence and economic advantage.

About ACG Research

ACG Research delivers information and communication technology market share/forecast reports, consulting services, and business case analysis services. Copyright © 2025 ACG Research. The copyright in this publication or the material on this website (including without limitation the text, computer code, artwork, photographs, images, music, audio material, video material and audio-visual material on this website) is owned by ACG Research. All Rights Reserved.

Author: Peter Fetterolf PhD

Source: ACG Research