



FROST & SULLIVAN
BEST PRACTICES



2026

GLOBAL ENTERPRISE
AI CONTROL PLANE
**TRANSFORMATIONAL
INNOVATION LEADERSHIP**

Best Practices Criteria for World-Class Performance	3
The Rise of the Enterprise AI Control Plane Industry	3
Unified AI Control Plane Leadership	4
Frictionless Adoption and Customer Enablement	8
Market Leadership and Brand Strength	9
Conclusion	11
What You Need to Know about the Transformational Innovation Leadership Recognition	12
Frost & Sullivan’s Transformational Innovation Leadership Recognition is its top honor and recognizes the market participant that exemplifies visionary innovation, market-leading performance, and unmatched customer care.	12
Best Practices Recognition Analysis	12
For the Transformational Innovation Leadership Recognition, Frost & Sullivan analysts independently evaluated the criteria listed below.	12
Transformational Innovation	12
Customer Impact	12
Best Practices Recognition Analytics Methodology	13
Inspire the World to Support True Leaders	13
About Frost & Sullivan	14
The Growth Pipeline Generator™	14
The Innovation Generator™	14

Best Practices Criteria for World-Class Performance

Frost & Sullivan applies a rigorous analytical process to evaluate multiple nominees for each recognition category before determining the final recognition recipient. The process involves a detailed evaluation of best practices criteria across two dimensions for each nominated company. TrueFoundry excels in many of the criteria in the enterprise AI control plane space.

RECOGNITION CRITERIA	
Technology Leverage	Customer Impact
Commitment to Innovation	Price/Performance Value
Commitment to Creativity	Customer Purchase Experience
Stage Gate Efficiency	Customer Ownership Experience
Commercialization Success	Customer Service Experience
Application Diversity	Brand Equity

The Rise of the Enterprise AI Control Plane Industry

Enterprise adoption of artificial intelligence (AI) is accelerating globally, but many organizations continue to face structural and operational challenges that limit the transition from experimentation to secure, governed, and production-scale deployment. Organizations across industries grapple with fragmented AI ecosystems, where models, deployment frameworks, governance solutions, observability platforms, and agent orchestration systems often operate in silos.

These challenges have given rise to a new category of enterprise technology: the enterprise AI control plane. Similar to how cloud management platforms unified infrastructure operations and observability platforms centralized monitoring, enterprise AI control planes provide a centralized operational layer through which organizations can govern, secure, observe, optimize, and scale AI across models, agents, applications, and infrastructure environments.

The need for such platforms is especially acute in regulated industries such as banking, healthcare, telecommunications, manufacturing, and semiconductor design, where organizations must balance innovation with strict requirements around security, compliance, auditability, and data sovereignty. Enterprises increasingly require centralized systems capable of enforcing governance policies, managing heterogeneous AI environments, and ensuring operational consistency across cloud, hybrid, on-premise, and air-gapped deployments. At the center of this control plane sits the AI Gateway, the operational layer where governance is enforced against live traffic.

Founded in 2021 and headquartered in San Francisco, California, TrueFoundry has emerged as a category-defining leader in the enterprise AI control plane market, serving enterprise customers across North America, Latin America, Europe, Australia and Asia-Pacific.

Frost & Sullivan believes TrueFoundry has established an early architectural advantage by investing ahead of the market in governed agent infrastructure, including MCP Gateways, Agent Gateways, Model gateway, Skills Registries, and enterprise-wide AI governance. As organizations transition from isolated large language model (LLM) applications toward fleets of autonomous agents, this early investment positions the company to benefit from one of the fastest-growing segments of enterprise AI infrastructure.

Today, TrueFoundry is trusted by leading Fortune 500s running production-ready agentic applications workflows through the AI Gateway. The company's customer portfolio spans telecom, healthcare, banking, semiconductors and more, with enterprises like NetApp, Automation Anywhere, Innovaccer, and Cox Communications running production workflows via the TrueFoundry AI Gateway. An estimated one trillion tokens a day pass through TrueFoundry's AI Gateway, representing a leading use case in a highly regulated industry.

Unified AI Control Plane Leadership

TrueFoundry's key differentiator is its ability to unify the fragmented layers of enterprise AI into a single non-opinionated operational platform. The company's AI Control Plane establishes a common governance layer, enabling organizations to apply consistent security policies, route requests intelligently, maintain operational visibility, enforce compliance requirements, and optimize resource utilization, regardless of the underlying model provider or deployment environment.

TrueFoundry applies this governance through the AI Gateway, the point where the control plane's rules meet live traffic and get enforced on every request. The AI Gateway incorporates three sibling gateway layers: an LLM Gateway governing model access, a Model Context Protocol (MCP) Gateway governing tool access, and an Agent Gateway governing agent operations complemented by an Agent Registry and a reusable Skills Registry.

The platform's architecture is inherently enterprise-ready, built on a Kubernetes-native foundation that supports deployment across multi-cloud, hybrid, on-premise, and air-gapped environments. This flexibility ensures that organizations retain full control over data residency, security and compliance requirements, while maintaining a consistent operational model across infrastructure environments. The ability to deploy consistently across Amazon Web Services (AWS), Microsoft Azure, Google Cloud, and on-premise infrastructure without vendor lock-in directly addresses one of the most pressing concerns for large enterprises adopting AI at scale.

This purpose-built design is what separates TrueFoundry's AI Gateway from a conventional Application Programming Interface (API) Gateway. A general-purpose API gateway retrofitted with an AI plugin has no native concept of token economics, prompt-level tracing, or the identity of an agent invoking a tool on a user's behalf. The company's AI-native architecture replaces that per-model configuration with a single virtual-model abstraction, stitches the LLM, MCP, guardrail, and agent call into one end-to-end trace, and enforces cost and access controls at the gateway layer itself.

TrueFoundry further strengthens its technology leadership with the launch of TrueForge, an open-source, vendor-neutral agent harness that helps enterprises take AI agents from prototype to production. Frost & Sullivan believes agent harnesses will become a foundational layer of enterprise AI infrastructure as organizations move from deploying individual AI agents to operating fleets of autonomous agents that require centralized governance, security, and operational control. By combining production-grade agent capabilities with the centralized governance, security, observability, and cost controls of the TrueFoundry AI Control Plane, TrueForge eliminates the tradeoff between flexibility and enterprise governance.

"A lot of these agent platforms have their own kind of proprietary connectors and then you're vendor locked in. TrueForge as a better way to be a little bit vendor agnostic and be able to be portable."

- Rob Rubin, Senior Director, Developer Platform Engineering, NetApp

TrueFoundry's performance at scale further reinforces its technology leadership. The platform operates in production environments that process trillions of tokens and billions of requests, demonstrating reliability and robustness under real-world enterprise workloads. Customers leverage the system to support high-volume inference services, large-scale generative AI (GenAI) applications, and mission-critical operations, validating its capability to move beyond experimentation into sustained, production-grade deployments.

Critical Architectural Components of an AI Gateway

Frost & Sullivan believes enterprise buyers increasingly evaluate an AI Gateway based on five strategic capabilities that determine whether AI can be deployed securely, governed consistently, and operated efficiently at scale.

AI Security: As AI systems increasingly process sensitive enterprise data and invoke external tools, security requirements extend beyond controlling model access. Organizations require runtime policy enforcement that governs every model request and tool invocation while supporting regulatory frameworks such as Health Insurance Portability and Accountability Act (HIPAA), General Data Protection Regulation (GDPR), and SOC 2.

TrueFoundry addresses this through OAuth2/OIDC authentication, fine-grained role-based access control, and virtual keys that let teams authenticate without ever holding real provider credentials. Input and output guardrails and audit trails are enforced consistently across the AI Gateway, MCP Gateway, and Agent Gateway.

AI Governance: As organizations expand AI across multiple models, agents, and enterprise tools, governance becomes increasingly difficult to manage through disconnected point solutions. Enterprises require centralized oversight that standardizes policy enforcement, routing, auditability, and lifecycle management across heterogeneous AI environments.

Because its gateway already governs models, tools, and agents in one layer, TrueFoundry resolves identity, applies role-based access control, enforces budget, and writes an audit record for every model call, MCP tool invocation, and agent interaction before it reaches a provider, maintaining consistent policy across cloud, hybrid, on-premises, and air-gapped deployments.

AI Observability: Production AI workloads involve multiple models, agents, and enterprise tools, making end-to-end observability essential for operational reliability. Organizations require visibility into model interactions, agent execution, latency, token consumption, and system performance to troubleshoot increasingly complex AI workflows.

TrueFoundry's gateway is OpenTelemetry-compliant, generating a structured trace for every request that captures model, provider, tokens, latency, and tool calls, and exporting it to any opentelemetry protocol-compatible backend without code changes. This gives enterprises a single operational view across production environments processing trillions of tokens and billions of requests.

AI Sovereignty: As enterprises deploy AI in regulated environments, data sovereignty depends not only on where data is stored, but also on how it moves at runtime. Model routing, retries, failover, agent workflows, MCP tool invocations, and observability pipelines can all create unintended cross-region data movement, making compliance with GDPR, HIPAA, and enterprise data residency requirements increasingly complex.

TrueFoundry's split-plane architecture separates orchestration from execution: a control plane manages configuration and policy, while the gateway plane is where models run, requests are proxied, and artifacts are stored. The gateway plane operates on the customer's own infrastructure. This lets the company enforce residency at the AI Gateway as a hard constraint, failing a request closed if no compliant path exists rather than routing around it. This is also why mid-to-large enterprises with multi-regional or global operations gravitate toward the platform: one consistent operational model, without centralizing the data itself.

AI Cost Control: AI spend becomes difficult to manage at scale for a simple reason: without centralized governance, costs get fragmented across model providers, applications, and business units, and no one has a clear view of where the money is going or how to optimize it.

TrueFoundry addresses this through what it calls AI FinOps, built directly into the Enterprise AI Control Plane:

- Token-level visibility into usage across every model, application, and team
- Intelligent model routing that sends requests to the most cost-effective model capable of handling them
- Semantic caching that avoids paying for repeated or near-duplicate queries
- Centralized cost governance that gives teams one place to track and control spend across the organization

Customers using these controls report infrastructure savings of 30% to 60%, along with meaningful reductions in overall AI costs.

"It would not have been possible to take our Generative AI projects to production this fast without the TrueFoundry platform and the effort and brains of the TrueFoundry team." -

-VP, Digital & Technology, Fortune 50 Healthcare Company

TrueFoundry's versatility is evident in the breadth of applications it supports across industries and use cases. These deployments range from agentic interactive voice response systems handling hundreds of millions of customer interactions, to predictive maintenance systems powered by pre-trained domain-specific language models, to enterprise AI for information technology operations platforms and GenAI application suites.

Customer-Driven Innovation and Impact

A defining aspect of TrueFoundry's innovation approach is its rapid product evolution cycle. The company actively incorporates customer feedback into its development roadmap, continuously refining and expanding platform capabilities in response to real-world usage patterns. Weekly feature releases, publicly shared roadmaps, and a strong feedback loop between customers, engineering teams, and product leadership ensure that innovation closely aligns with market needs. This dynamic and responsive innovation model allows TrueFoundry to remain ahead of industry trends while delivering immediate value to its customers.

Through partnerships with cloud providers, integration with leading AI frameworks, and the development of training programs and certification initiatives, TrueFoundry actively builds an ecosystem that supports knowledge sharing, capability development, and accelerated adoption. This ecosystem-centric approach amplifies the platform's impact and ensures that innovation extends beyond the product itself into the broader AI community.

Case Study: \$100 Million in Savings and a 45 times Latency Improvement at a Fortune 50 Health Brand

Challenge

A Fortune 50 health enterprise operating more than 10,000 United States pharmacy locations relied on a legacy menu-tree interactive voice response (IVR) system to handle roughly 500 million patient calls a year on its prescription-refill line. The system understood only rigid, pre-scripted phrasing, so natural requests such as asking to refill a heart medication failed to map to any recognized intent. Most routine calls escalated to a live pharmacist, consuming clinical capacity on tasks that required no clinical judgment, while downtime on the line carried an estimated cost of roughly \$6,000 per second given the scale of the business behind it.

Solution

TrueFoundry was brought in as the control plane and AI Gateway underneath the health enterprises' internal agent workflow platform, while an internal team retained ownership of the business logic and call-routing workflows built on top of it. The AI Gateway routed every request across more than 40 self-hosted models running on a hybrid infrastructure estate, on-premise NVIDIA GPUs for steady, predictable load and cloud burst capacity for spikes, while the control plane handled model deployment, fine-tuning, and one-click model operations without requiring the health enterprises' teams to manage Kubernetes directly.

Each call passed through a three-stage routing system:

- A deterministic rules layer resolved 90% to 95% of routine requests from cache with no model involved

- A lightweight classifier confirmed scope
- Only genuinely complex requests reached a large reasoning model, cutting unnecessary model calls by 10% to 20%

Safety was enforced at the gateway before any reasoning occurred, including a hard rule that the system could inform and transact on prescriptions but never offer medical advice, toxicity screening on every turn, and automatic escalation of emergency keywords straight to a human line.

Results

- Cut response latency from roughly 90 seconds in prototype to under two seconds in production
- Reached full national rollout across the entire store footprint in six months, against an original 18-month estimate
- Improved call containment by 14 percentage points over the legacy system
- Ran more than 40 self-hosted models in production, with zero reputational incidents across more than four years live
- Expanded from a single IVR deployment to more than 15 production use cases across separate business units, each governed centrally while operating its own isolated environment
- Delivered more than \$100 million in combined model and infrastructure savings

Frictionless Adoption and Customer Enablement

TrueFoundry delivers a differentiated customer experience by aligning its purchase journey and post-deployment engagement model with the realities of enterprise AI adoption. Recognizing that organizations often enter the AI journey with fragmented tooling, limited operational maturity, and unclear pathways to production, the company designs a customer-centric approach that prioritizes speed, flexibility, scalability, and hands-on enablement from the initial engagement through ongoing value realization.

It anchors the purchase experience in a low-friction, value-driven onboarding model. Rather than requiring large upfront commitments or complex integration cycles, TrueFoundry typically engages customers through structured proof-of-value programs. These engagements allow enterprises to test high-impact use cases within a short timeframe, often achieving production-ready deployment within weeks. This approach enables customers to directly experience the platform's capabilities before scaling adoption, thereby reducing perceived risks and accelerating decision-making.

The platform's flexibility and interoperability enable seamless integration into existing environments without requiring replacement of prior investments. Applications, model configurations, and governance frameworks can be migrated or layered with minimal effort, lowering switching barriers and eliminating vendor lock-in, allowing enterprises to adopt the platform incrementally as needs evolve.

Once onboarded, TrueFoundry delivers a structured, high-touch customer experience focused on rapid enablement and continuous engagement. Dedicated teams support environment setup, user onboarding,

and capability adoption across platform administrators and developers, while real-time collaboration channels and close engineering alignment enable fast issue resolution, quick feature turnaround, and direct integration of customer feedback, ensuring a responsive, adaptive partnership rather than a static support model.

TrueFoundry also enhances the customer experience through structured adoption and expansion frameworks. After initial deployment, customers typically progress through a phased rollout model, transitioning from pilot use cases to broader enterprise-wide implementations. Regular touchpoints, including quarterly business reviews, allow the company to assess usage, identify optimization

“TrueFoundry's key differentiator is its ability to unify the fragmented layers of enterprise AI into a single operational platform. By bringing together governance, observability, deployment orchestration, optimization, and agent operations, the company delivers a level of control and simplicity that few competitors can match.”

- Arun Prasath
Principal Consultant, Industrial

opportunities, and guide customers toward additional capabilities such as agentic workflows or advanced governance features. This consultative engagement model ensures that adoption continuously expands.

The platform's comprehensive training and education initiatives further reflects its emphasis on enablement. Through its academy programs, hands-on workshops, and engineering-led sessions, TrueFoundry equips customers with the skills required to independently build, deploy, and scale AI applications. These initiatives address the common enterprise challenge of talent gaps in AI and ensure

that organizations fully leverage the platform's capabilities without relying heavily on external support.

The company's demonstrable outcomes in real-world deployments reinforce its successful customer experience. Enterprises across industries have successfully accelerated their AI initiatives, achieving faster time-to-production, improved operational efficiency, and significant cost savings. In several cases, customers highlight that the speed of deployment and the level of support provided by the TrueFoundry team were critical factors in achieving production readiness in record time.

TrueFoundry's agile operating model further enhances the overall experience. As a fast-growing organization with a highly collaborative structure, the company maintains a strong degree of responsiveness and adaptability. It involves cross-functional teams in customer engagements, ensuring that expertise across product, engineering, and strategy is readily accessible. This integrated approach allows the company to deliver technical depth and strategic guidance, strengthening long-term customer relationships.

Market Leadership and Brand Strength

TrueFoundry establishes a rapidly strengthening presence within the global enterprise AI ecosystem by positioning itself as a category-defining enterprise AI control plane rather than a conventional AI infrastructure provider. As organizations increasingly seek centralized governance, observability, security, and operational control across AI initiatives, the company differentiates itself by providing a unified platform that enables enterprises to manage AI as a strategic business capability rather than a collection of disconnected tools.

Unlike point solutions that focus solely on model serving, observability, gateways, or agent orchestration, TrueFoundry consolidates the operational layers of enterprise AI into a single enterprise AI control plane. This differentiated approach enables organizations to establish centralized governance while simultaneously accelerating AI adoption across business functions. By unifying deployment, governance, observability, optimization, security, and agent lifecycle management, the company creates a compelling value proposition that resonates with enterprise customers seeking to scale AI responsibly and efficiently.

The company's growing presence within enterprise environments further strengthens its brand equity. TrueFoundry is trusted by a diverse portfolio of organizations, including Fortune 500 and Global 2000 enterprises, that rely on the platform to operationalize mission-critical AI workloads. This growing adoption reflects increasing market recognition of the company's ability to solve complex enterprise AI challenges through a unified operational framework.

The company's brand development efforts further reinforce its market position. Through industry events, executive forums, technical workshops, thought leadership initiatives, and developer-focused programs, TrueFoundry actively contributes to conversations shaping the future of enterprise AI. These activities increase visibility while positioning the company as a technology innovator and a trusted advisor to organizations navigating AI transformation.

The strength of TrueFoundry's leadership team further enhances market credibility. Drawing upon extensive experience in large-scale software systems, cloud infrastructure, machine learning operations, and enterprise technology, the company combines deep technical expertise with practical operational understanding. This blend enables TrueFoundry to design solutions that address technological complexity and real-world business requirements, strengthening customer trust and long-term engagement. In June 2026, the company acquired Seldon AI, a United Kingdom-based pioneer in production-grade machine learning (ML) serving, extending its control plane to unify traditional ML and agentic AI operations under a single governance layer.

As the enterprise AI control plane category continues to mature, TrueFoundry's blend of technology leadership, customer success, ecosystem partnerships, and global adoption position it as one of the defining organizations shaping the future of enterprise AI operations.

The company's customer base continues to expand geographically and across industries, demonstrating its ability to scale commercially while maintaining relevance in diverse markets. With a strong presence across North America, Europe, and Asia-Pacific, TrueFoundry establishes itself as a globally relevant player, with a majority of its revenue driven by mature enterprise markets. This geographic diversification enhances the resilience of its growth model and reinforces its position as a scalable, internationally recognized brand.

The company's strong leadership team and technical heritage further reinforce its brand credibility. With roots in large-scale AI deployments and experience from leading technology organizations, TrueFoundry brings a high level of expertise to its platform design and execution. This combination of technical pedigree and practical enterprise experience enhances customer confidence and strengthens the company's position as a trusted partner for AI transformation initiatives.

Conclusion

TrueFoundry is a category-defining leader in the enterprise artificial intelligence (AI) control plane market, providing the operational foundation that enables organizations to govern, deploy, secure, observe, optimize, and scale AI at enterprise scale. Its unified platform brings together governance, deployment orchestration, observability, security, multi-model management, and agent lifecycle operations within a single control plane, helping enterprises overcome fragmented AI ecosystems and increasing compliance requirements. By transforming AI from isolated experimentation into a secure, production-ready capability, TrueFoundry enables organizations to innovate faster while maintaining operational control and visibility. Customers benefit from accelerated deployment timelines, improved infrastructure utilization, significant cost efficiencies, enhanced governance, and production-scale reliability. Supported by strong customer engagement, expanding ecosystem partnerships, and growing adoption across North America, Europe, and Asia-Pacific, the company redefines how enterprises operationalize, manage, and scale AI globally.

With its strong overall performance, TrueFoundry earns Frost & Sullivan's 2026 Global Transformational Innovation Leadership Recognition in enterprise AI control plane industry.

What You Need to Know about the Transformational Innovation Leadership Recognition

Frost & Sullivan's Transformational Innovation Leadership Recognition is its top honor and recognizes the market participant that exemplifies visionary innovation, market-leading performance, and unmatched customer care.

Best Practices Recognition Analysis

For the Transformational Innovation Leadership Recognition, Frost & Sullivan analysts independently evaluated the criteria listed below.

Transformational Innovation

Market Disruption: Innovative new solutions have a genuine potential to disrupt the market, render current solutions obsolete, and shake up competition

Competitive Differentiation: Solutions or products articulate and display unique competitive advantages

Market Gaps: Solution satisfies the needs and opportunities that exist between customers' desired outcomes and their current market solutions

Leadership Focus: Companies' focuses are on building a leadership position in core markets and on creating stiff barriers to entry for new competitors

Passionate Persistence: Tenacity enables the pursuit and achievement of seemingly insurmountable industry obstacles

Customer Impact

Price/Performance Value: Products or services offer the best ROI and superior value compared to similar market offerings

Customer Purchase Experience: Purchase experience with minimal friction and high transparency assures customers that they are buying the optimal solution to address both their needs and constraints

Customer Ownership Excellence: Products and solutions evolve continuously in sync with the customers' own growth journeys, engendering pride of ownership and enhanced customer experience

Customer Service Experience: Customer service is readily accessible and stress-free, and delivered with high quality, high availability, and fast response time

Brand Equity: Customers perceive the brand positively and exhibit high brand loyalty, which is regularly measured and confirmed through a high Net Promoter Score®

Best Practices Recognition Analytics Methodology

Inspire the World to Support True Leaders

This long-term process spans 12 months, beginning with the prioritization of the sector. It involves a rigorous approach that includes comprehensive scanning and analytics to identify key best practice trends. A dedicated team of analysts, advisors, coaches, and experts collaborates closely, ensuring thorough review and input. The goal is to maximize the company's long-term value by leveraging unique perspectives to support each Best Practice Recognition and identify meaningful transformation and impact.

VALUE IMPACT			
STEP		WHAT	WHY
1	Opportunity Universe	Identify Sectors with the Greatest Impact on the Global Economy	Value to Economic Development
2	Transformational Model	Analyze Strategic Imperatives That Drive Transformation	Understand and Create a Winning Strategy
3	Ecosystem	Map Critical Value Chains	Comprehensive Community that Shapes the Sector
4	Growth Generator	Data Foundation That Provides Decision Support System	Spark Opportunities and Accelerate Decision-making
5	Growth Opportunities	Identify Opportunities Generated by Companies	Drive the Transformation of the Industry
6	Frost Radar	Benchmark Companies on Future Growth Potential	Identify Most Powerful Companies to Action
7	Best Practices	Identify Companies Achieving Best Practices in All Critical Perspectives	Inspire the World
8	Companies to Action	Tell Your Story to the World (BICEP*)	Ecosystem Community Supporting Future Success

*Board of Directors, Investors, Customers, Employees, Partners

About Frost & Sullivan

Frost & Sullivan is the Growth Pipeline Company™. We power our clients to a future shaped by growth. Our Growth Pipeline as a Service™ provides the CEO and the CEO's growth team with a continuous and rigorous platform of growth opportunities, ensuring long-term success. To achieve positive outcomes, our team leverages over 60 years of experience, coaching organizations of all types and sizes across 6 continents with our proven best practices. To power your Growth Pipeline future, visit Frost & Sullivan at <http://www.frost.com>.

The Growth Pipeline Generator™

Frost & Sullivan's proprietary model to systematically create ongoing growth opportunities and strategies for our clients is fueled by the Innovation Generator™.

Learn more.

Key Impacts:

- **Growth Pipeline:** Continuous Flow of Growth Opportunities
- **Growth Strategies:** Proven Best Practices
- **Innovation Culture:** Optimized Customer Experience
- **ROI & Margin:** Implementation Excellence
- **Transformational Growth:** Industry Leadership



The Innovation Generator™

Our 6 analytical perspectives are crucial in capturing the broadest range of innovative growth opportunities, most of which occur at the points of these perspectives.

Analytical Perspectives:

- Megatrend (MT)
- Business Model (BM)
- Technology (TE)
- Industries (IN)
- Customer (CU)
- Geographies (GE)

