

Broadcom Charts Platform Engineering 2.0: Extending the IDP for the AI-Native Enterprise

July 06, 2026

By: [Jim Mercer](#), [Matthew Flug](#)

IDC'S QUICK TAKE

Broadcom's VMware Cloud Foundation (VCF) business recently introduced [Platform Engineering 2.0](#), a five-pillar framework that extends the discipline's foundation, platform as a product, golden paths, and self-service internal developer platforms (IDPs) into the AI era. The framework adds AI-native infrastructure, a multi-persona experience model, embedded FinOps, security that shifts down into the platform substrate, and a composable, API-first architecture to the existing foundation of platform engineering principles.

IDC research finds that 93% of organizations are piloting, using, expanding, or planning to use an IDP within the next 12 months (see *Driving Platform Engineering Adoption, Capabilities, and Practices, 2025*, IDC #US53710125, August 2025), indicating that platform adoption is no longer the question. In the age of AI, the question will be whether the underlying platform can absorb the complexities of AI workloads, additional personas, increased operational costs, and security pressures. Broadcom's framework, anchored in VCF, is intended to address that gap and provide platform leaders with a vocabulary and road map for the next 3 to 4 years.

NEWS HIGHLIGHTS

With the introduction of Platform Engineering 2.0, Broadcom contends that IDPs have achieved broad adoption but were originally designed for a single class of users before agentic development and GPU-intensive workloads reshaped platform requirements. Broadcom organizes its response to these developments into five pillars and maps each to capabilities already available or planned within VCF. The five-pillar framework for Platform Engineering 2.0 includes: AI-native platform, multi-persona experience, embedded FinOps, security shifts down, and more composable architecture, all built on top of the existing platform as product operating model. Details of each pillar include:

- An **AI-native platform** pillar anchored on the concept of an agentic development platform (ADP), the structural evolution of the IDP that supports GPU/TPU provisioning, model serving and registries, Model Context Protocol (MCP) server infrastructure, and bounded-autonomy guardrails for AI agents acting as a new, nonhuman platform persona
- A **multi-persona experience** model that expands the platform's addressable user base from the application developer to six personas: app developers, platform engineers, engineering and business leaders, security and compliance teams, data scientists/ML engineers, and AI agents themselves
- **Embedded FinOps capabilities**, including self-service cost transparency, real-time cost attribution, pre-deployment cost gates, token/GPU-specific "tokenomics" cost attribution, and autonomous resource reclamation agents, to move cost governance from retrospective reporting to provisioning time decisioning

- A security **shifts down** model that complements (rather than replaces) shift-left practices by embedding policy as code, secure-by-default configuration, and runtime anomaly detection into the platform and runtime layers, with explicit controls for new AI attack surfaces such as shadow AI sprawl, prompt injection, model poisoning, and inference data leakage
- A **composable, five-layer reference architecture** (experience, orchestration, capabilities, integration, and infrastructure layers) designed around API-first contracts, CNCF-conformant building blocks, and progressive migration, positioned explicitly as "compose, not build versus buy"

VCF-specific capabilities are mapped to each pillar, including GPU as a service, VKS (vSphere Kubernetes Service) for FinOps-aware Kubernetes cost attribution, a curated model store with RBAC for AI model governance, defense-in-depth security spanning hardware to workload layers, and support for an open AI ecosystem of frameworks, models, hardware, and protocols. By embedding these enterprise-grade capabilities directly into the infrastructure, Platform Engineering 2.0 shifts the operational burden from developers to a streamlined, self-service platform. Ultimately, this alignment between VCF and modern platform principles accelerates application delivery, allowing organizations to rapidly deploy secure, cost-optimized, and AI-ready workloads with unprecedented agility.

IDC'S POINT OF VIEW

Platform engineering has crossed the initial adoption threshold and is no longer a niche practice but a mainstream operating model for enterprise software delivery. Broadcom's announcement suggests that adoption alone is no longer the differentiator because the operating model that drove that adoption was built for a containerized, human-paced, single-persona world that no longer reflects how enterprises build software. IDC survey data shows that 81% say agentic AI tools have increased their software development productivity, with over half reporting savings of 5 or more hours per developer per week (see *IDC Survey: Goals, Benefits, and Barriers to Agentic AI Adoption in Software Development — Insights from IDC's Agentic Application Development and DevOps Survey*, IDC #US54283526, March 2026). That upstream gain does not fully remove inefficiencies in the software development life cycle. Instead, it shifts the bottleneck downstream to the platform, impacting its capacity to deliver, validate, and govern the additional output. Organizations that have not re-architected their IDPs to absorb this throughput may find their golden paths becoming, in Broadcom's own words, "golden cages."

An important distinction surrounding this announcement is positioning Platform Engineering 2.0 as an extension rather than a replacement. Some vendors tend to declare prior investments obsolete to justify new platform purchases. Broadcom instead preserves the platform engineering substrate, while layering AI-native, multi-persona, and cost- and security-aware capabilities. So, the underlying goals of an IDP have not changed. IDPs should enable persona-based application development while providing the governance and compliance guardrails that platform teams require. What has changed is the resource mix, the complexity of AI applications and agents, the size of the attack surface, and the workflows themselves. This evolution of platform engineering is a realistic approach because ripping out a working IDP to chase an AI-native rebuild introduces unnecessary risk and delay at a time when organizations need platform stability to tackle AI-driven change. It behooves platform leaders

evaluating this framework, or any competitor's equivalent, to consider auditing their current platform against the five pillars rather than treating the framework as a forklift upgrade.

The multi-persona experience pillar reframes platform engineering's value proposition for the C-suite. IDC's research shows that platform engineering maturity correlates with how strategically an organization treats DevOps and how many roles actively rely on the platform: Architects, platform engineers, and DevOps engineers report the highest maturity because they are the platform's primary builders and users. At the same time, organizations that confine the platform to application developers alone set a ceiling on their maturity.

Extending the platform to security, data science, FinOps, and business leader personas, and treating AI agents as a first-class, API-consuming persona, is not optional scope creep. Instead, it is how the maturity ceiling gets raised. However, organizations should recognize that serving six personas with a single underlying platform team may require additional investments in persona-specific experience layers (portals, CLIs, dashboards) built on a shared API surface. Organizations that attempt this without recognizing the needs of each persona risk introducing fragmentation that platform engineering was created to eliminate, just with a longer list of roles.

Embedded FinOps and security are the two most urgent and least mature pillars across the industry. IDC research reveals that organizations with more years of DevOps adoption prioritize measuring cloud cost management and security alongside developer productivity. A sign that cost and security become first-class platform signals after an organization has moved well past initial adoption. This relationship appears to be the gap Broadcom is targeting with embedded FinOps.

The emphasis on pre-deployment cost gates and real-time attribution is directionally correct. However, emerging challenges such as "tokenomics" cost attribution, tracking GPU time, inference cost, and token usage per team and per model, remain nascent across the vendor landscape, not just within VCF. Organizations may need to build internal tooling to supplement platform-native capabilities in the near term. However, emerging players such as Larridin, Exceeds AI, and others are promoting solutions in this space.

Similarly, the AI-specific security frontier Broadcom describes, including shadow AI sprawl, prompt injection, model poisoning, and inference data leaks, is real, larger than the traditional attack surface, and is generally not protected by existing shift-left tooling. Because these vulnerabilities exist at the data and runtime layers rather than in static source code, legacy vulnerability scanners remain fundamentally blind to them.

At the end of the day, the platform, not the pipeline, is the correct enforcement point for these new attack surfaces, since no SAST or DAST tool can inspect a live inference stream. Centralized platforms can help to make AI governance tractable at scale. As AI workloads, agents, and their associated costs, integrations, and tool calls proliferate across teams, the platform becomes the single control point for defining and applying cost and security policies once and uniformly. As a result, organizations gain consistent visibility and enforcement across all AI workloads.

Broadcom's differentiation is architectural because VCF's composable layering and its position beneath both containers and VMs give it a well-founded claim to unify traditional, modern, and AI workloads on a single operational framework, a claim that pure-play IDP vendors operating

solely at the experience and orchestration layers cannot make. Platform engineering technology evaluations shouldn't be limited to vendors' head-to-head comparisons in isolation. Instead, organizations must weigh infrastructure control against user-experience flexibility. So, for enterprises with large private clouds already standardized on VMware, extending their current setup can potentially offer a much faster path to value than building an internal developer platform from scratch with a new vendor.

Key takeaways

- Platform Engineering 2.0 is best understood as an extension of, not a replacement for, the existing foundation of platform as product, golden paths, and self-service IDPs.
- The shift from a developer-only platform to an expanded-persona model, including AI agents as a first-class, API-consuming persona, may require an increased product management investment to avoid recreating siloed fragmentation.
- Embedded FinOps and AI-specific security controls are critical to ensure the success of platform engineering 2.0.
- Platform leaders should consider auditing their current platform against the five pillars rather than treating the framework as a forklift upgrade when evaluating this framework or any competitor's equivalent.
- Organizations evaluating this framework should also weigh it against adjacent and competitive offerings. Platform-as-a-service and IDP vendors are advancing multi-persona and FinOps capabilities from a software-only, cloud-agnostic posture.
- Broadcom's true differentiator lies in its architecture. By positioning VCF's composable layers beneath both containers and VMs, Broadcom can unify traditional, modern, and AI workloads under a single operational framework. A feat that pure-play IDP vendors cannot match with experience and orchestration alone.
- It may be more beneficial for organizations with existing VMware and private cloud investments to extend those estates rather than adopt a parallel, cloud-native-only IDP stack.

Subscriptions Covered:

[AI Assurance, ALM, Quality and Portfolio Strategies](#), [Intelligent Application Modernization and Deployment Platforms](#)

Please contact the IDC Hotline at 800.343.4952, ext.7988 (or +1.508.988.7988) or sales@idc.com for information on applying the price of this document toward the purchase of an IDC or Industry Insights service or for information on additional copies or Web rights. Visit us on the Web at www.idc.com. To view a list of IDC offices worldwide, visit www.idc.com/offices. Copyright 2026 IDC. Reproduction is forbidden unless authorized. All rights reserved.