# Stick with Cloud Foundry

Making and Maintaining the Case To Stay on Cloud Foundry

Michael Coté

Tanzu Division, Broadcom

21 September 2026

vmware by Broadcom

Broadcom Proprietary and Confidential. Copyright © 2026 Broadcom.
All Rights Reserved. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries.

# Let’s migrate!

## Or, why does this keep happening?


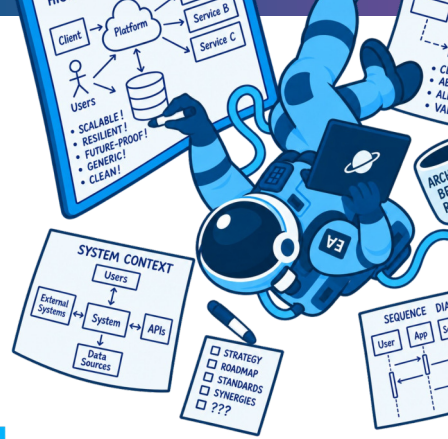


vmware by Broadcom

Broadcom Proprietary and Confidential. Copyright © 2026 Broadcom.
All Rights Reserved. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries.

“They hired me to make changes, so let’s make changes.”


vmware by Broadcom
Broadcom Proprietary and Confidential. Copyright © 2026 Broadcom.
All Rights Reserved. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries.

# A working platform is invisible


vmware by Broadcom
Broadcom Proprietary and Confidential. Copyright © 2026 Broadcom.
All Rights Reserved. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries.

# CoE it from orbit

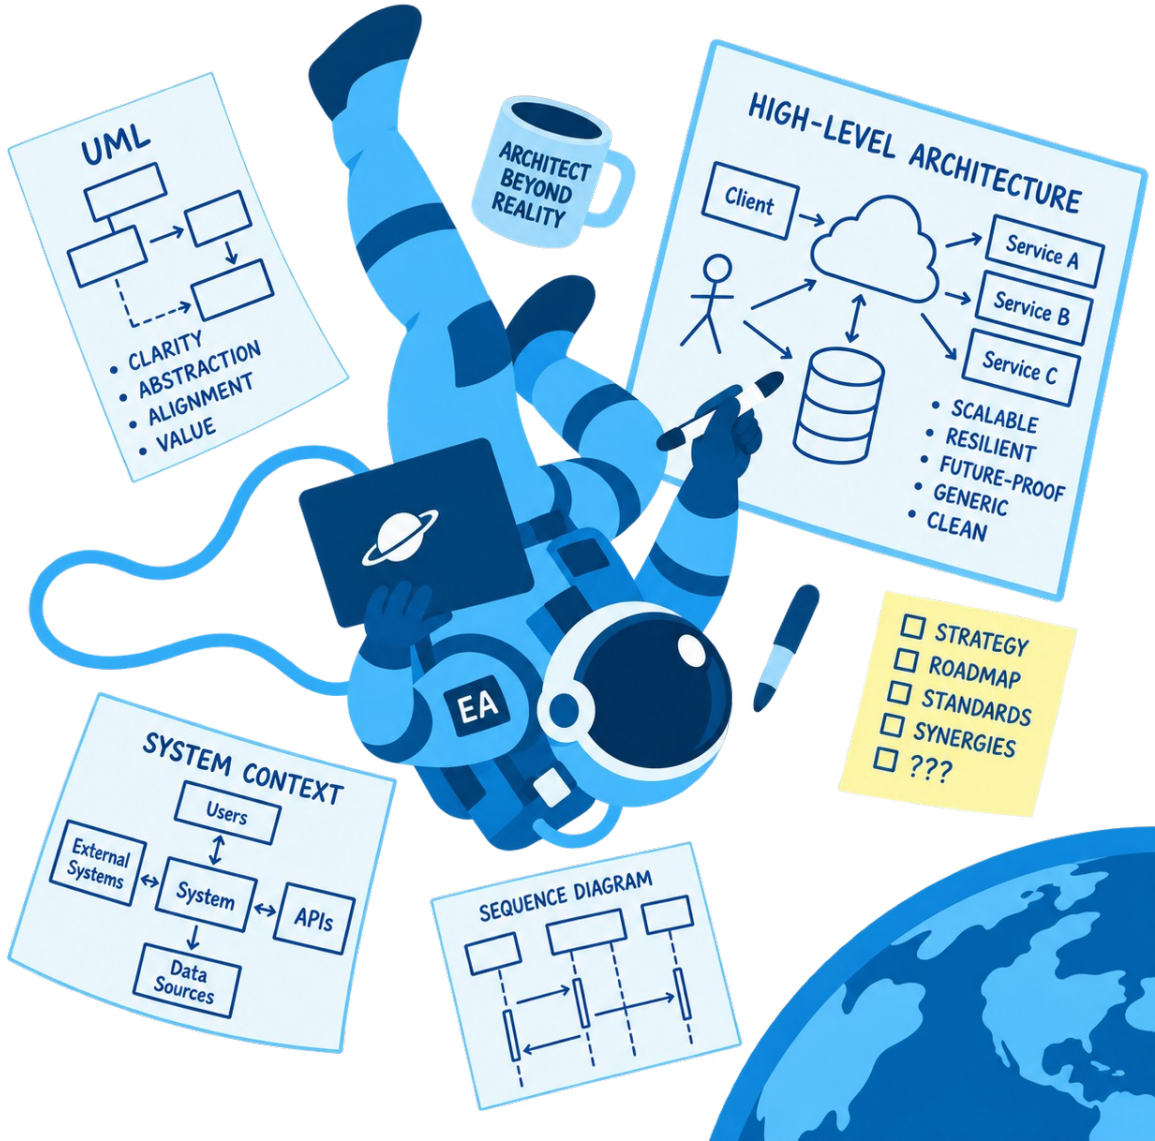


vmware by Broadcom

Broadcom Proprietary and Confidential. Copyright © 2026 Broadcom.
All Rights Reserved. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries.

# The Current New Thing Needs a New Platform


vmware by Broadcom
Broadcom Proprietary and Confidential. Copyright © 2026 Broadcom.
All Rights Reserved. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries.

# Coté

https://www.cote.io/ | cote@broadcom.com

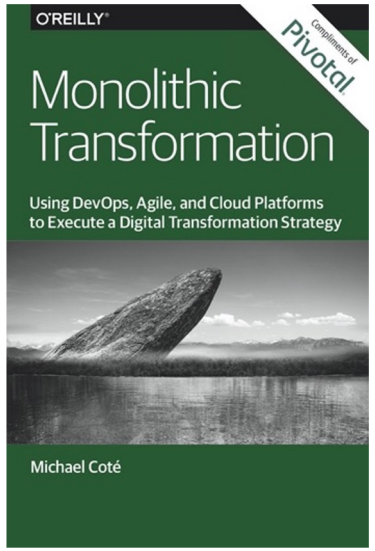


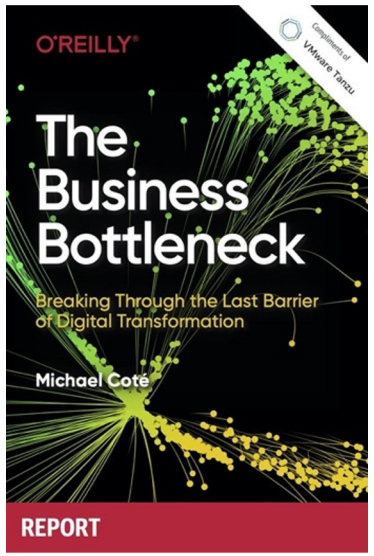


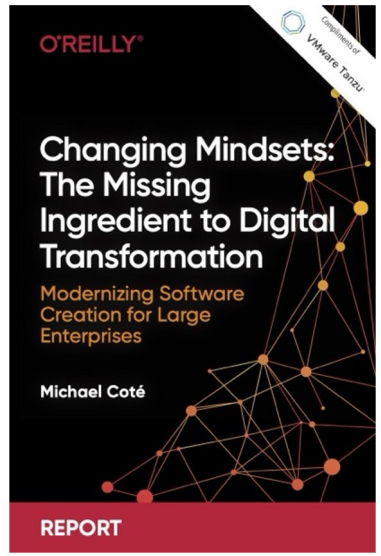



Tanzu Catsup

RedMonk DELL 451 Research S&P Global Market Intelligence Pivotal vmware Tanzu by Broadcom

pyright © 2026 Broadcom.
refers to Broadcom Inc. and/or its subsidiaries.

# It works!

vmware by Broadcom
Broadcom Proprietary and Confidential. Copyright © 2026 Broadcom.
All Rights Reserved. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries.

- 350 apps / 7 ops
- 300 apps / 8 ops
- 1,200 app teams / 4 ops

- 30,000 devs / 50 ops
- 6,500 devs / 16 ops
- 2,500 devs / 5 ops
- 1,200 devs / 6 ops

- 45 app teams / 5 ops

Sources: Kroger, GAIC, Mercedes-Benz, conversations with FSI platform engineers; "Enterprise Grade Platform Engineering at Charles Schwab," Coté, September, 2024, based on Schwab's Explore 2024 panel; Rabobank ops conversations, CF Day EU 2025, Oct 7th, 2025; "3 Cloud Foundry Stories," Coté, CF Day EU, Oct 7th, 2025. "Shaping Our Platform Future: Rabobank's Shift To OSS Cloud Foundry," Jochem Siegel & René Dollevoet, Rabobank, CF Sumit, Sep 21st, 2026.

vmware by Broadcom
Broadcom Proprietary and Confidential. Copyright © 2026 Broadcom
All Rights Reserved. The term "Broadcom" refers to Broadcom Inc. and/or its subsidiaries.

# Four multi-year success stories at large enterprises

## Outcomes 👍

- Trusted to deploy during working hours, heightened awareness periods.
- 99% faster setup - completely self-service developer onboarding.
- 75% reduction in time spent patching.
- ~25% more time coding, over 6,400+ developers.
- Happy developers, happy apps.

As of September, 2025.

16

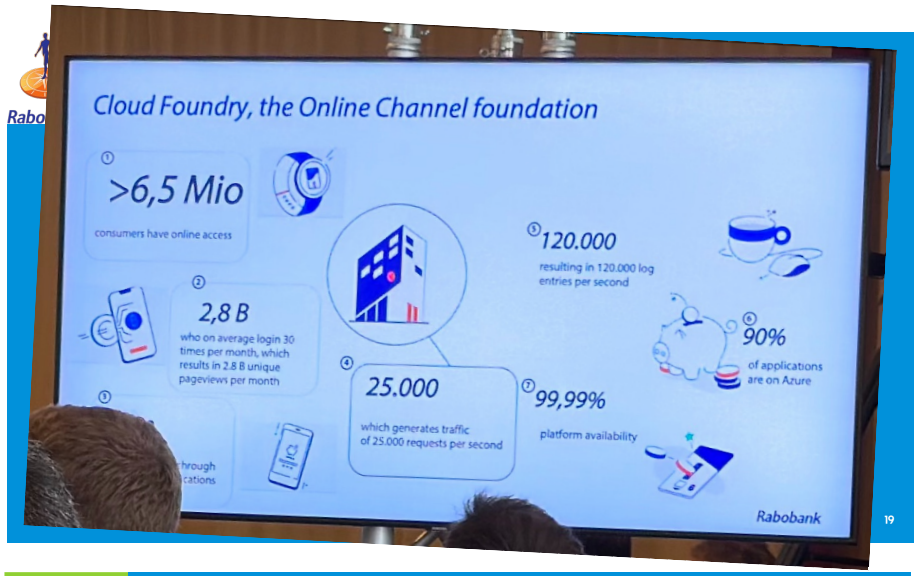


19

charles SCHWAB

## Outcomes 👍

- No-fee trading enabled by keeping operating costs low
- Managing hundreds of thousands of app instances with 3-person team
- Handles daily load, "market-storms."

Easier audits and governance

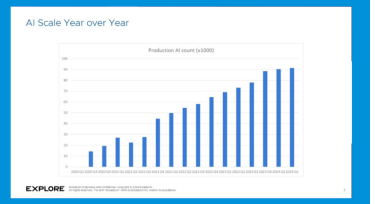


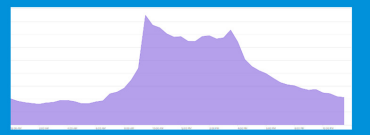

First 15 minutes of market open

Sources: "Enterprise Grade Platform Engineering at Charles Schwab," Coté, September, 2024, based on Schwab's Explore 2024 panel [TANP2089LV].

22

DATEV

## Outcomes 👍

- Modernization and standardization.
- More appdev, less opsdev.
- Patching in 2 to 3 days.
- Integrate with mainframes.

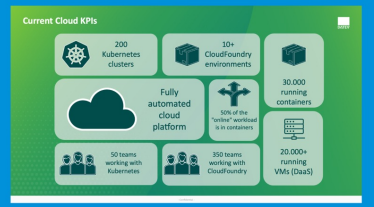



Sources: conversations with Jürgen Sußner, DATEV, October, 2025; DATEV case study and video.

26

vmware by Broadcom

Broadcom Proprietary and Confidential. Copyright © 2026 Broadcom.
All Rights Reserved. The term "Broadcom" refers to Broadcom Inc. and/or its subsidiaries.

Source: "Three Cloud Foundry Stories," Coté, CF Day, October, 2025; "Shaping Our Platform Future: Rabobank's Shift To OSS Cloud Foundry," Jochem Siegel & René Dollevoet, Rabobank, CF Sumit, Sep 21st, 2026.

# Replatform to what?

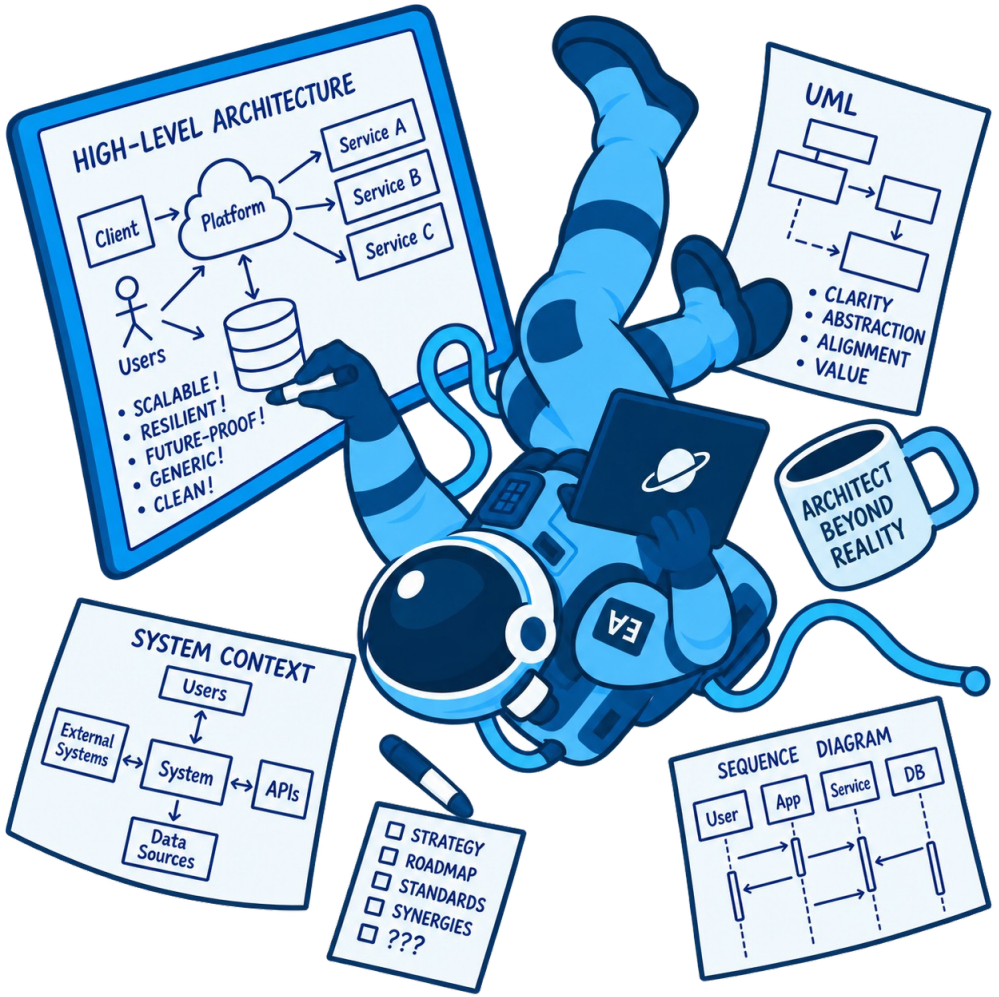


vmware by Broadcom
Broadcom Proprietary and Confidential. Copyright © 2026 Broadcom.
All Rights Reserved. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries.

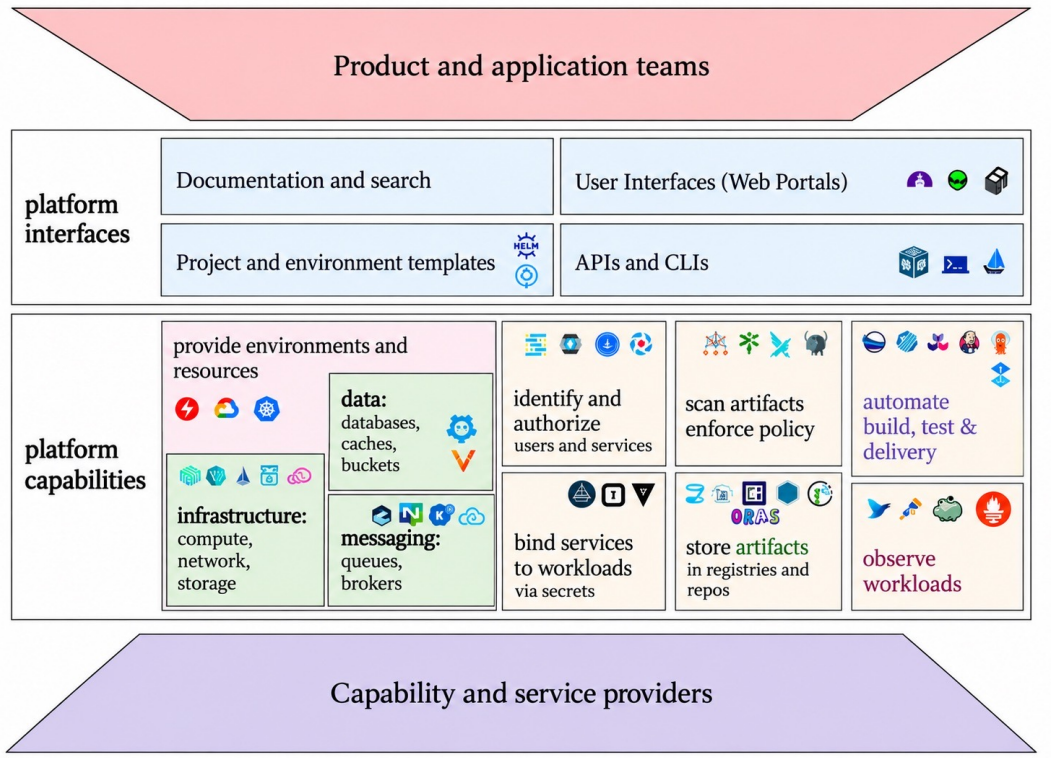


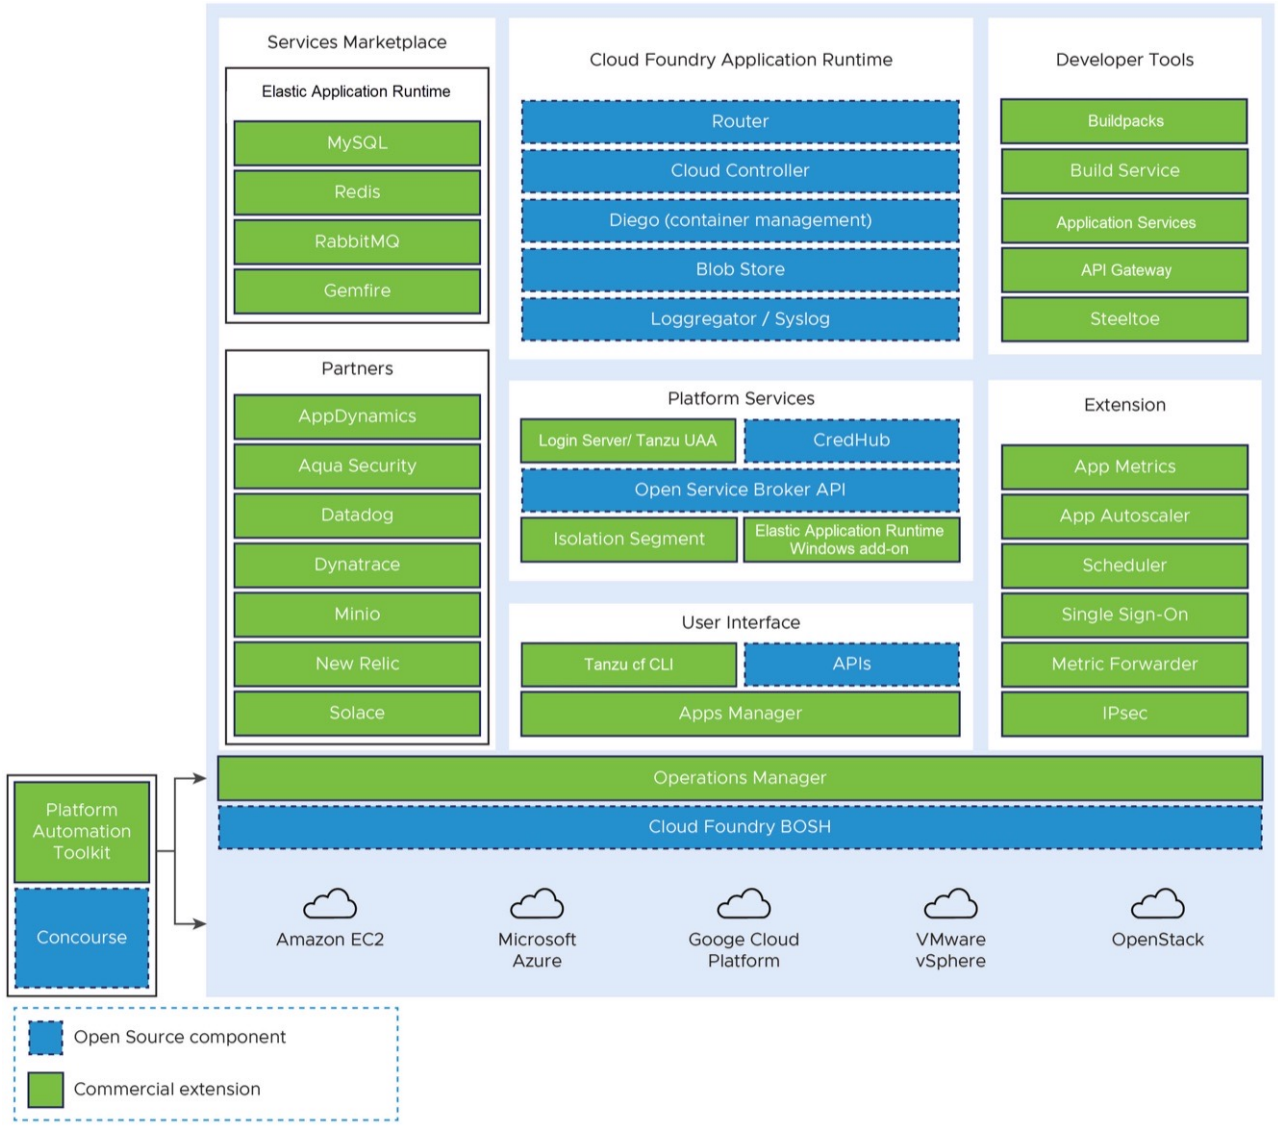


vmware by Broadcom
Broadcom Proprietary and Confidential. Copyright © 2026 Broadcom.
All Rights Reserved. The term "Broadcom" refers to Broadcom Inc. and/or its subsidiaries.

Sources: "CNCF Platforms White Paper," March 2023; Tanzu Platform 10.4 documentation.

# Groups running Cloud Foundry are often mature+

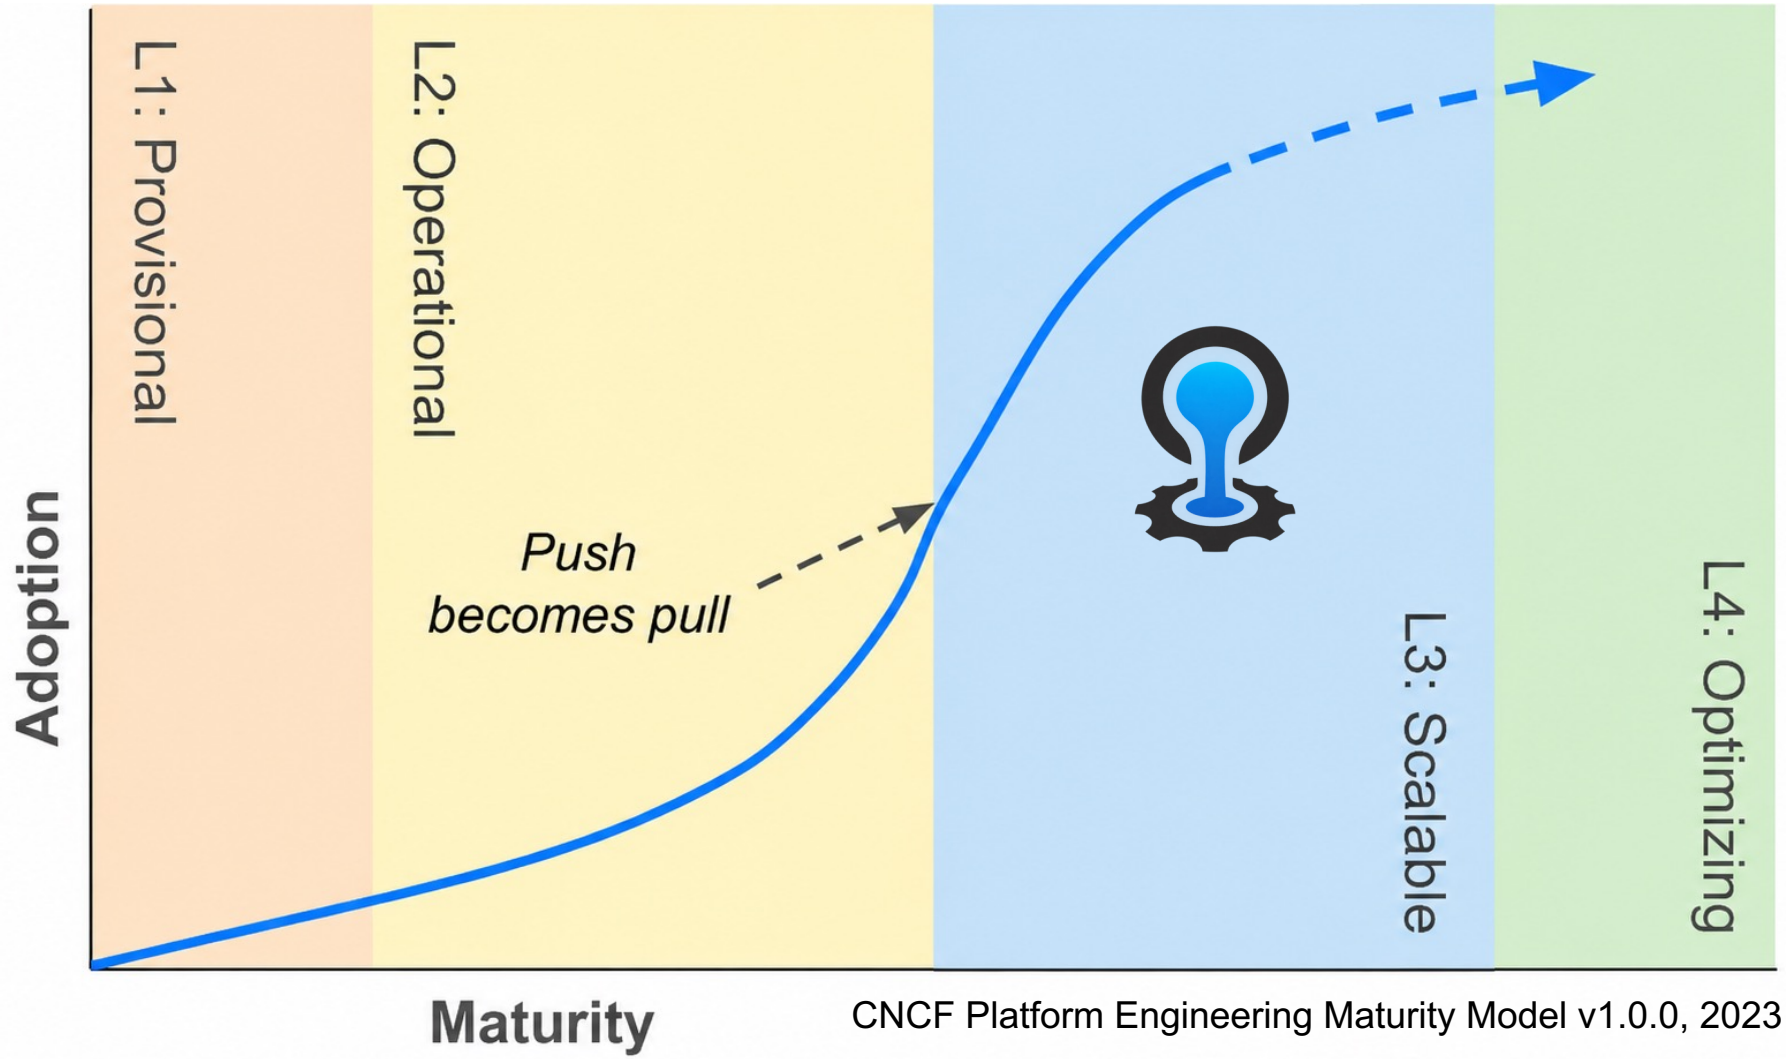


CNCF Platform Engineering Maturity Model v1.0.0, 2023

vmware by Broadcom
Broadcom Proprietary and Confidential. Copyright © 2026 Broadcom.
All Rights Reserved. The term "Broadcom" refers to Broadcom Inc. and/or its subsidiaries.

Sources: CNCF Platform Engineering Maturity Model,2023; *Platform Engineering*, Fournier & Nowland, 2024.

| | Level 1 – Build | Level 2 – Operate | Level 3 – Scale | Level 4 – Improve | Level 5 – Adapt |
|---|---|---|---|---|---|
| **Primary focus** | First cloud native workloads | Production + standardization | Org-wide repeatability | Governance + security by default | Continuous optimization |
| **Platform & infrastructure** | Initial Kubernetes clusters<br>Basic container registry<br>Infrastructure as Code with Terraform | GitOps for apps via Argo CD or Flux<br>Thin internal developer platform (templates, self-service)<br>Centralized environments | Multiple standardized clusters (multi-region / multi-env)<br>Formal Internal Developer Platform (portals, templates, APIs)<br>GitOps extended to platform components | Drift detection + auto-remediation<br>Zero-trust platform access<br>Security posture management integrated with ops | Policy-driven workload placement (cloud/on-prem/edge)<br>Platform continuously reshaped based on usage data |
| **Application delivery** | Helm or raw manifests via Helm<br>Basic CI pipelines<br>1–2 pilot apps | Helm + Kustomize at scale<br>Runtime config via ConfigMaps/Secrets<br>Standard base images + scanning + SBOMs | Artifact signing + automated promotion pipelines<br>Namespaces-as-a-Service<br>Horizontal + event-driven autoscaling from app metrics | End-to-end supply chain security (signed images, enforced SBOMs)<br>Admission controls everywhere | Progressive delivery (canaries, feature flags everywhere<br>)Predictive scaling from production signals |
| **Observability & operations** | Minimal logging/metrics (often cloud defaults + early Prometheus) | Metrics + logs + early tracing using OpenTelemetry<br>Central log aggregation | Distributed tracing as first-class signal<br>Automated backup, DR, cluster lifecycle | Central audit pipelines (SCM, CI, clusters, apps)<br>Runtime policy enforcement | Automated performance/reliability feedback loops<br>Anomaly detection from live telemetry |
| **Networking & security** | Manual secrets<br>Basic perimeter security | Admission controls<br>Container/runtime scanning<br>Early service mesh (often built on Envoy) | Operational service mesh (mTLS, retries, traffic shaping)<br>Multi-tenancy with quotas | Workload identity + automated cert rotation<br>Fine-grained authZ via mesh<br>Zero-trust networking | AI-assisted remediation<br>Continuous security optimization |
| **Cost & optimization** | Mostly manual cost awareness | Initial resource limits<br>Early FinOps signals (namespace quotas, requests/limits) | Chargeback/showback per team or namespace | Integrated FinOps dashboards + budget controls | FinOps directly drives autoscaling and scheduling<br>Continuous cost optimization |
| **AI (where applicable)** | Mostly experimental | First production models<br>Basic observability + access controls | | | |

Sources: "Platform Engineering Maturity Model," CNCF Platforms Working Group, October, 2023. See also "Cloud Native App Platforms: New Research Shows Struggles and Hope," Camille Crowell-Lee and Rita Manachi, June, 2024. See also beta version of Cloud Native Maturity Model: Core content updates for v4.0 (#84) (6e03f6c).

# “Most rearchitectures take three to five years” – all apps even longer

| Scrappy Platform | Run Sc | new platform, migrate apps |
| --- | --- | --- |
| 1 year | 2 to 3 yea | 3 to 5 years |

Assuming it works!!

vmware by Broadcom
Broadcom Proprietary and Confidential. Copyright © 2026 Broadcom.
All Rights Reserved. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries.

Sources: customer interviews; ”over drinks” confessionals; *Platform Engineering*, Fournier & Nowland, 2024; “The Upside-Down Economics of DIY PaaS,” Coté, Jared Ruckle, Bryan Friedman, Matt Walburn, 2026.

# The Easy Apps Are Easy to Move Quick

How organisations use containers, CNCF Annual Survey

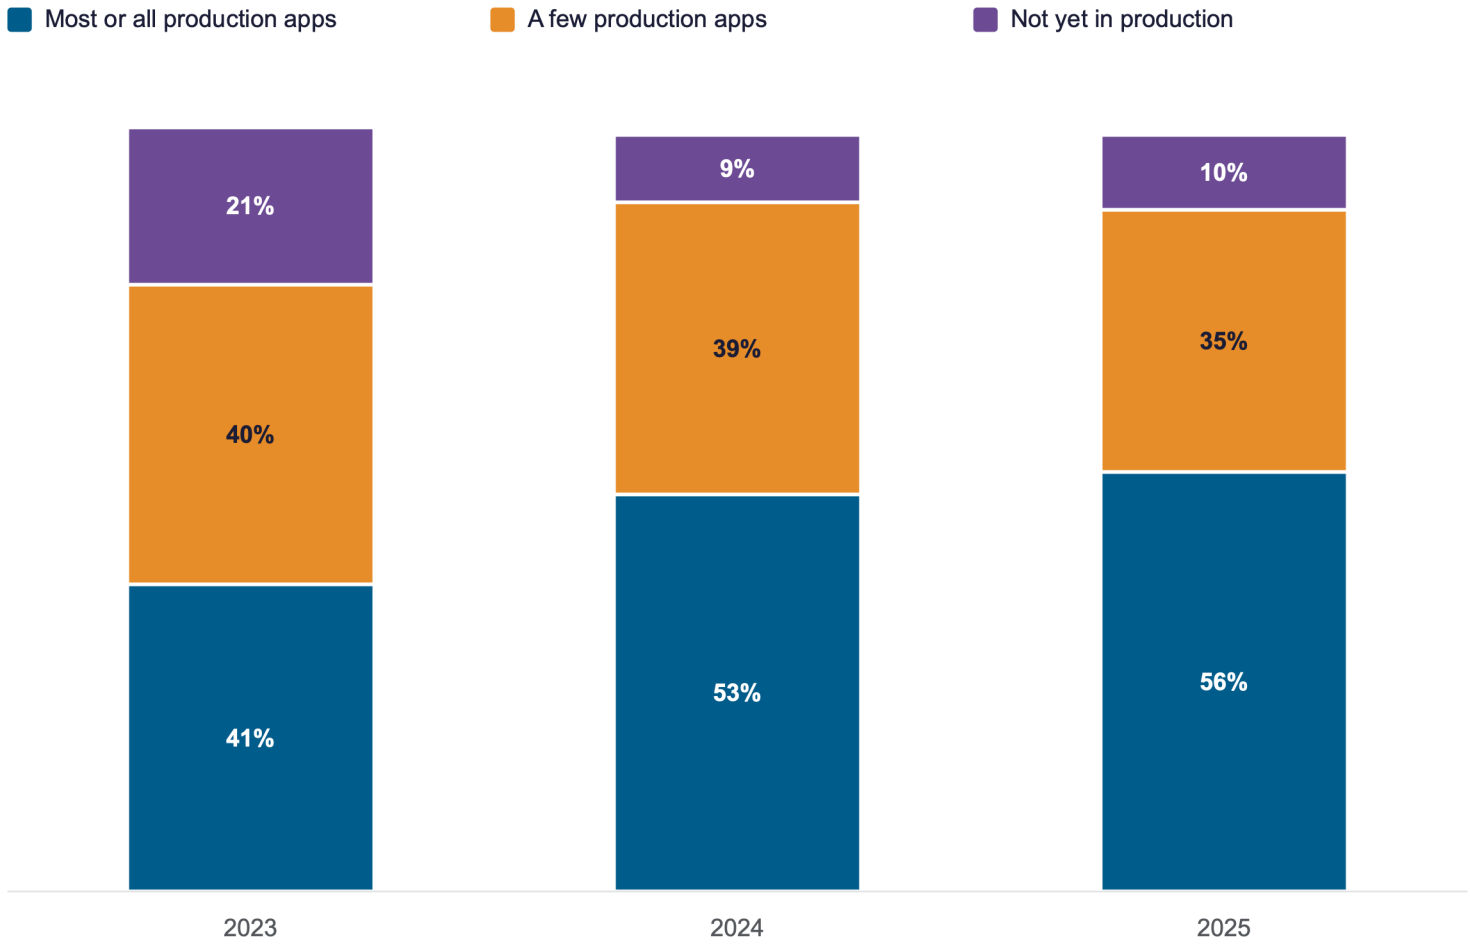


Source: CNCF Annual Cloud Native Survey, Figure 6, 2023-2025. n=522, 408, 365.
Chart generated by AI on 2026-09-18. Likely not checked by a human, so it may have errors.

vmware by Broadcom
Broadcom Proprietary and Confidential. Copyright © 2026 Broadcom.
All Rights Reserved. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries.

Source: CNCF Annual Cloud Native Survey, 2023, 2024, 2025. Figure 6. n=522, 408, 365.

# Seeing Platforms


vmware by Broadcom
Broadcom Proprietary and Confidential. Copyright © 2026 Broadcom.
All Rights Reserved. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries.

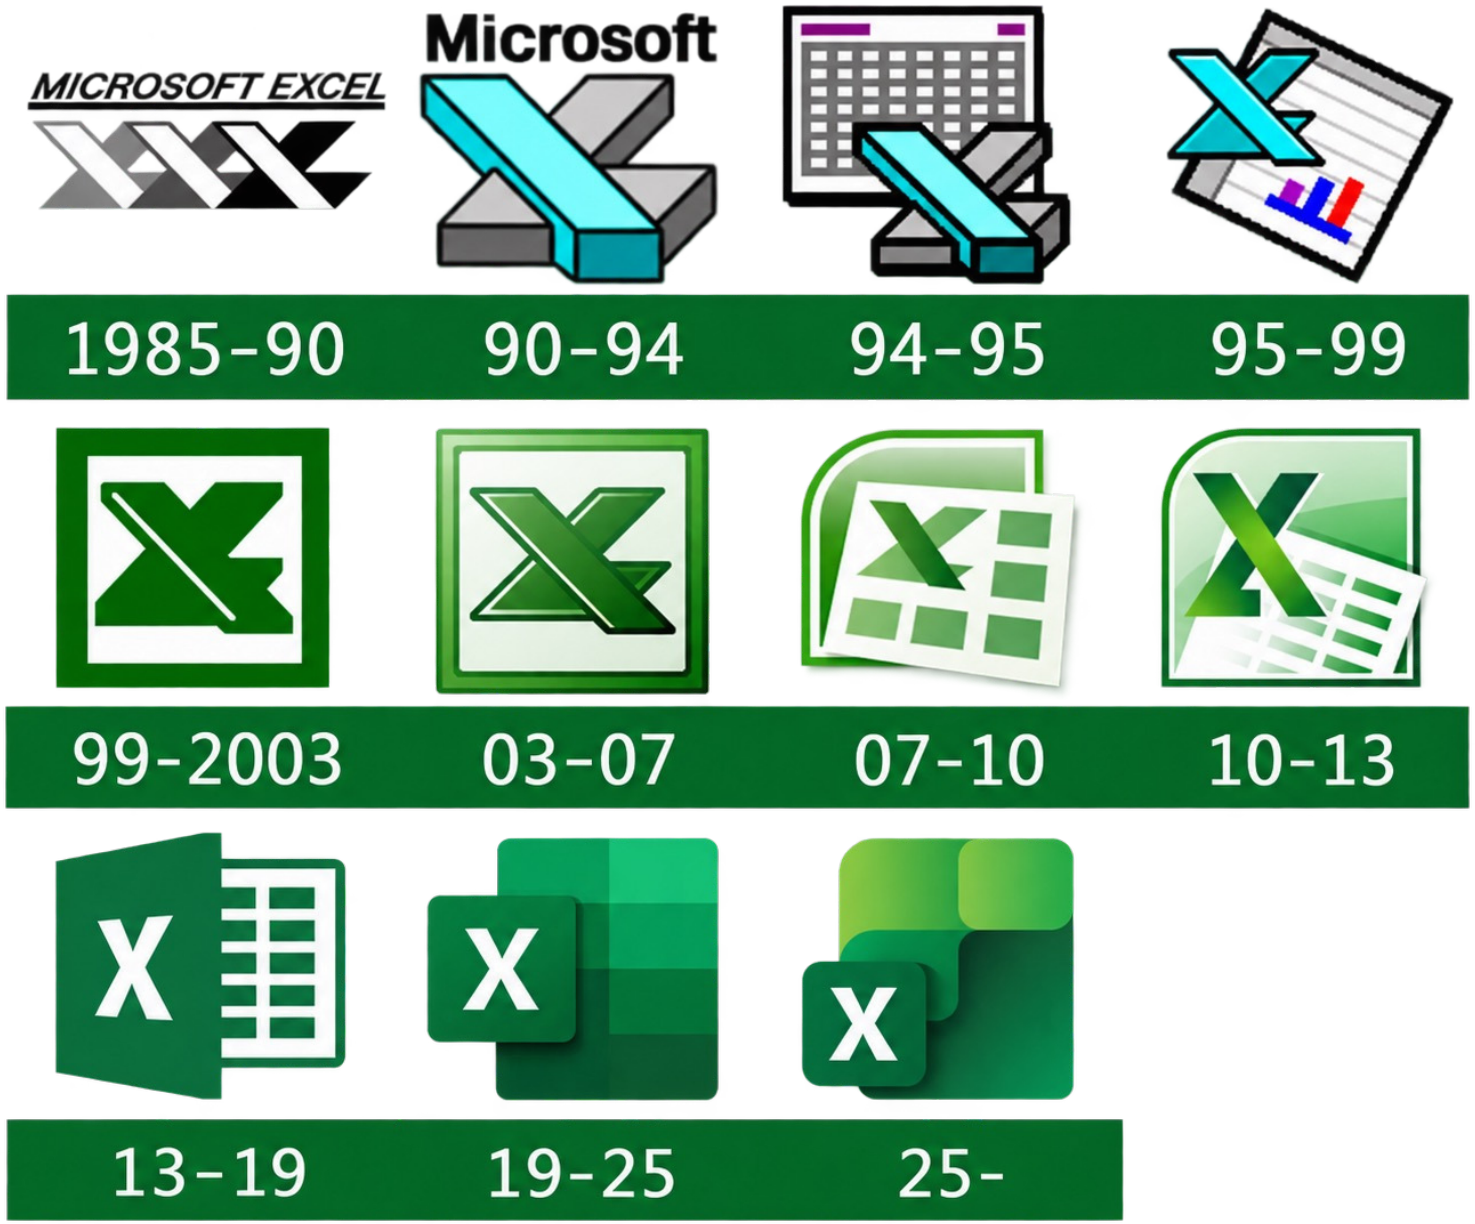


vmware by Broadcom
Broadcom Proprietary and Confidential. Copyright © 2026 Broadcom.
All Rights Reserved. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries.

# Metrics for the LINE OF BUSINESS

Indicators

## Speed

Velocity is a vector comprised of speed and direction.

We bring a raw speed of advantage to the LOBs and also enable them to rapidly and reliably respond to changes in direction in the service of the business based on user feedback loops.

**MEASUREMENTS**

- Time to value (cycle time)
- Frequence of customer feedback
- Time between bug identification and fix
- Time from feedback to deployment of change
- Customer satisfaction (NPS)
- Business satisfaction

## Stability

Reality is a complex landscape of changing priorities, emergent bugs, evolving architectures, and staffing changes.

We help the LOB achieve resiliency and low volatility as they deliver customer value in the face of this complex reality.

**MEASUREMENTS**

Volatility (std dev in velocity / mean velocity)
# of defects generated per developer - year
% of software launches / upgrades delayed due to defects
Employee satisfaction (ENPS)

## Scalability

LOBs need to scale across two dimensions:

People - LOBs strive to attract developers and ramp productivity linearly with personnel.

Apps - LOBs need to rapidly scale their applications and their complexity to handle demand.

**MEASUREMENTS**

- # of products in development
- # of products measuring business success
- Investment ratios: spend developing software vs operating and systems
- Disruption caused by doubling workload
- Ability to attract and retain talent (# of internal referrals)

## Security

To move rapidly the team needs to feel secure in making code changes aggressively. Automated test coverage provides this safety net.

To rapidly search for customer value LOBs must adopt a learning culture that fosters psychological safety necessary to fail and learn from failure.

**MEASUREMENTS**

% teams using CI
% teams doing TDD
Time from commit to deployment

## Savings

Teams must reduce risk and waste through small batch delivery and fast consumer feedback.

This drives significant savings as use of the product grows and is key to maintaining their trust and enabling them to go fast, forever.

**MEASUREMENTS**

- Fraction of developer time spend writing code and delivering value
- Product:dev ratio
- Business satisfaction
- # of go/no-go decisions based on business success

Sources: Tanzu customer interviews and consulting, ~2015 to 2025; "Common platform team mistakes and how to avoid them," Camille Fournier, PlatformCon 2025, July, 2025.

# Metrics for the IT

Indicators

## Speed

IT can efficiently upgrade, patch, and manage the platform.

They rapidly onboard new application teams and provide the necessary services to quickly unblock teams and enable them to deliver consumer value.

**MEASUREMENTS**

- # prod/dev deploys per month
- # platform upgrades per month
- Platform upgrade speed
- # of new apps onboarded/month
- Team distribution of skills

## Stability

Our customers entrust us with their production workloads and their developer productivity.

We must provide adequate SLOs to meet their needs and earn their trust by ensuring compatibility and uptime across platform upgrades.

**MEASUREMENTS**

Minutes of prod outage per year
Minutes of dev outage per year
Mean time to recovery
Mean time between failures
# of upgrade-related failures

## Scalability

IT needs to provide an “at-scale” service on-demand at the whim of the business.

They need to explore all options with minimal friction as they grapple with the mix of workloads on-premise and in the cloud.

**MEASUREMENTS**

- Queries per second
- # of AIs per foundation
- # of SIs per foundation
- # of foundations
- # of teams using the platform
- Does increasing workload on existing

## Security

Security is a paramount concern for our customers. We earn their trust by providing a platform that is secure by default.

We solve for security and reduce security-related friction and toil in order to enable our customers to go fast, forever.

**MEASUREMENTS**

Time between identifying and patching a CVE
Cost in person-hours or dollars of leaked credential
Fraction of operator time spent on security configuration
# of disruptions/suspensions due to security concerns

## Savings

IT must meet the needs of thousands of developers within tight budgetary constraints.

We provide a platform that simultaneously reduces complexity and sprawl and improves the ops:dev ratio.

**MEASUREMENTS**

- Operator:developer ratio
- # of apps per operator
- # of foundations per operator
- Degree of automation for provisioning, build, test, change approval governance, deployment, perf

Sources: Tanzu customer interviews and consulting, ~2015 to 2025; “Common platform team mistakes and how to avoid them," Camille Fournier, PlatformCon 2025, July, 2025.

# Platform marketing & community management






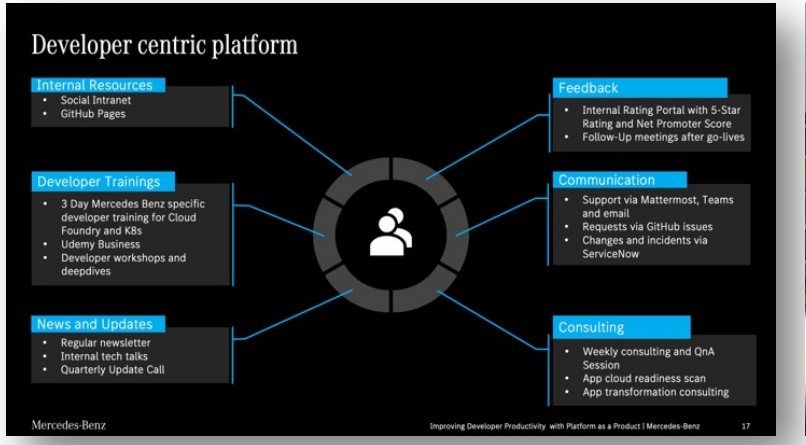



vmware by Broadcom

Broadcom Proprietary and Confidential. Copyright © 2026 Broadcom.
All Rights Reserved. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries.

Sources: ING, 2023; BT Canvas team; MB.io; Duke Energy; Allstate; "Take DevOps to 11 and Sprinkle Cloud on it with Rainbows and Unicorns," Matt Curry, s1p 2017. “Improve Developer Productivity with Platform as a Product,” VMware Explore, Nov. 2022; Kessel Run Wikipedia page (circa Feb 2025); laptop picture cleaned-up in ChatGPT to unblur finger, original available upon request.

“Consulting giant Accenture is trying to figure out how to stop non-technical workers from blowing through companies’ AI token budget on trivial tasks like converting PDFs to presentation slides, according to leaked audio obtained by 404 Media. Across the industry Accenture is seeing ‘soaring token spend,’ according to the audio.”

Sebastian Herrmann, *404 Media*.


vmware by Broadcom

Broadcom Proprietary and Confidential. Copyright © 2026 Broadcom.
All Rights Reserved. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries.

# Already and agentic platform

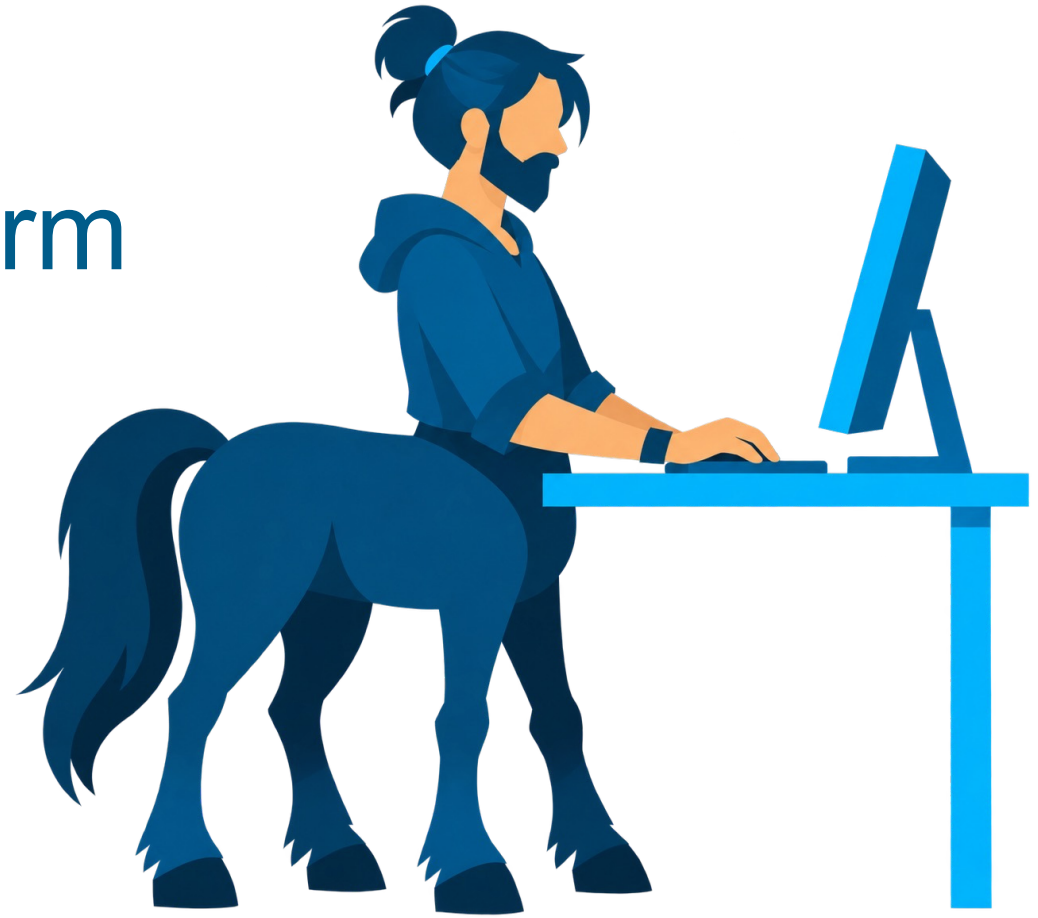

vmware by Broadcom
Broadcom Proprietary and Confidential. Copyright © 2026 Broadcom.
All Rights Reserved. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries.

# Agents in your CaaS

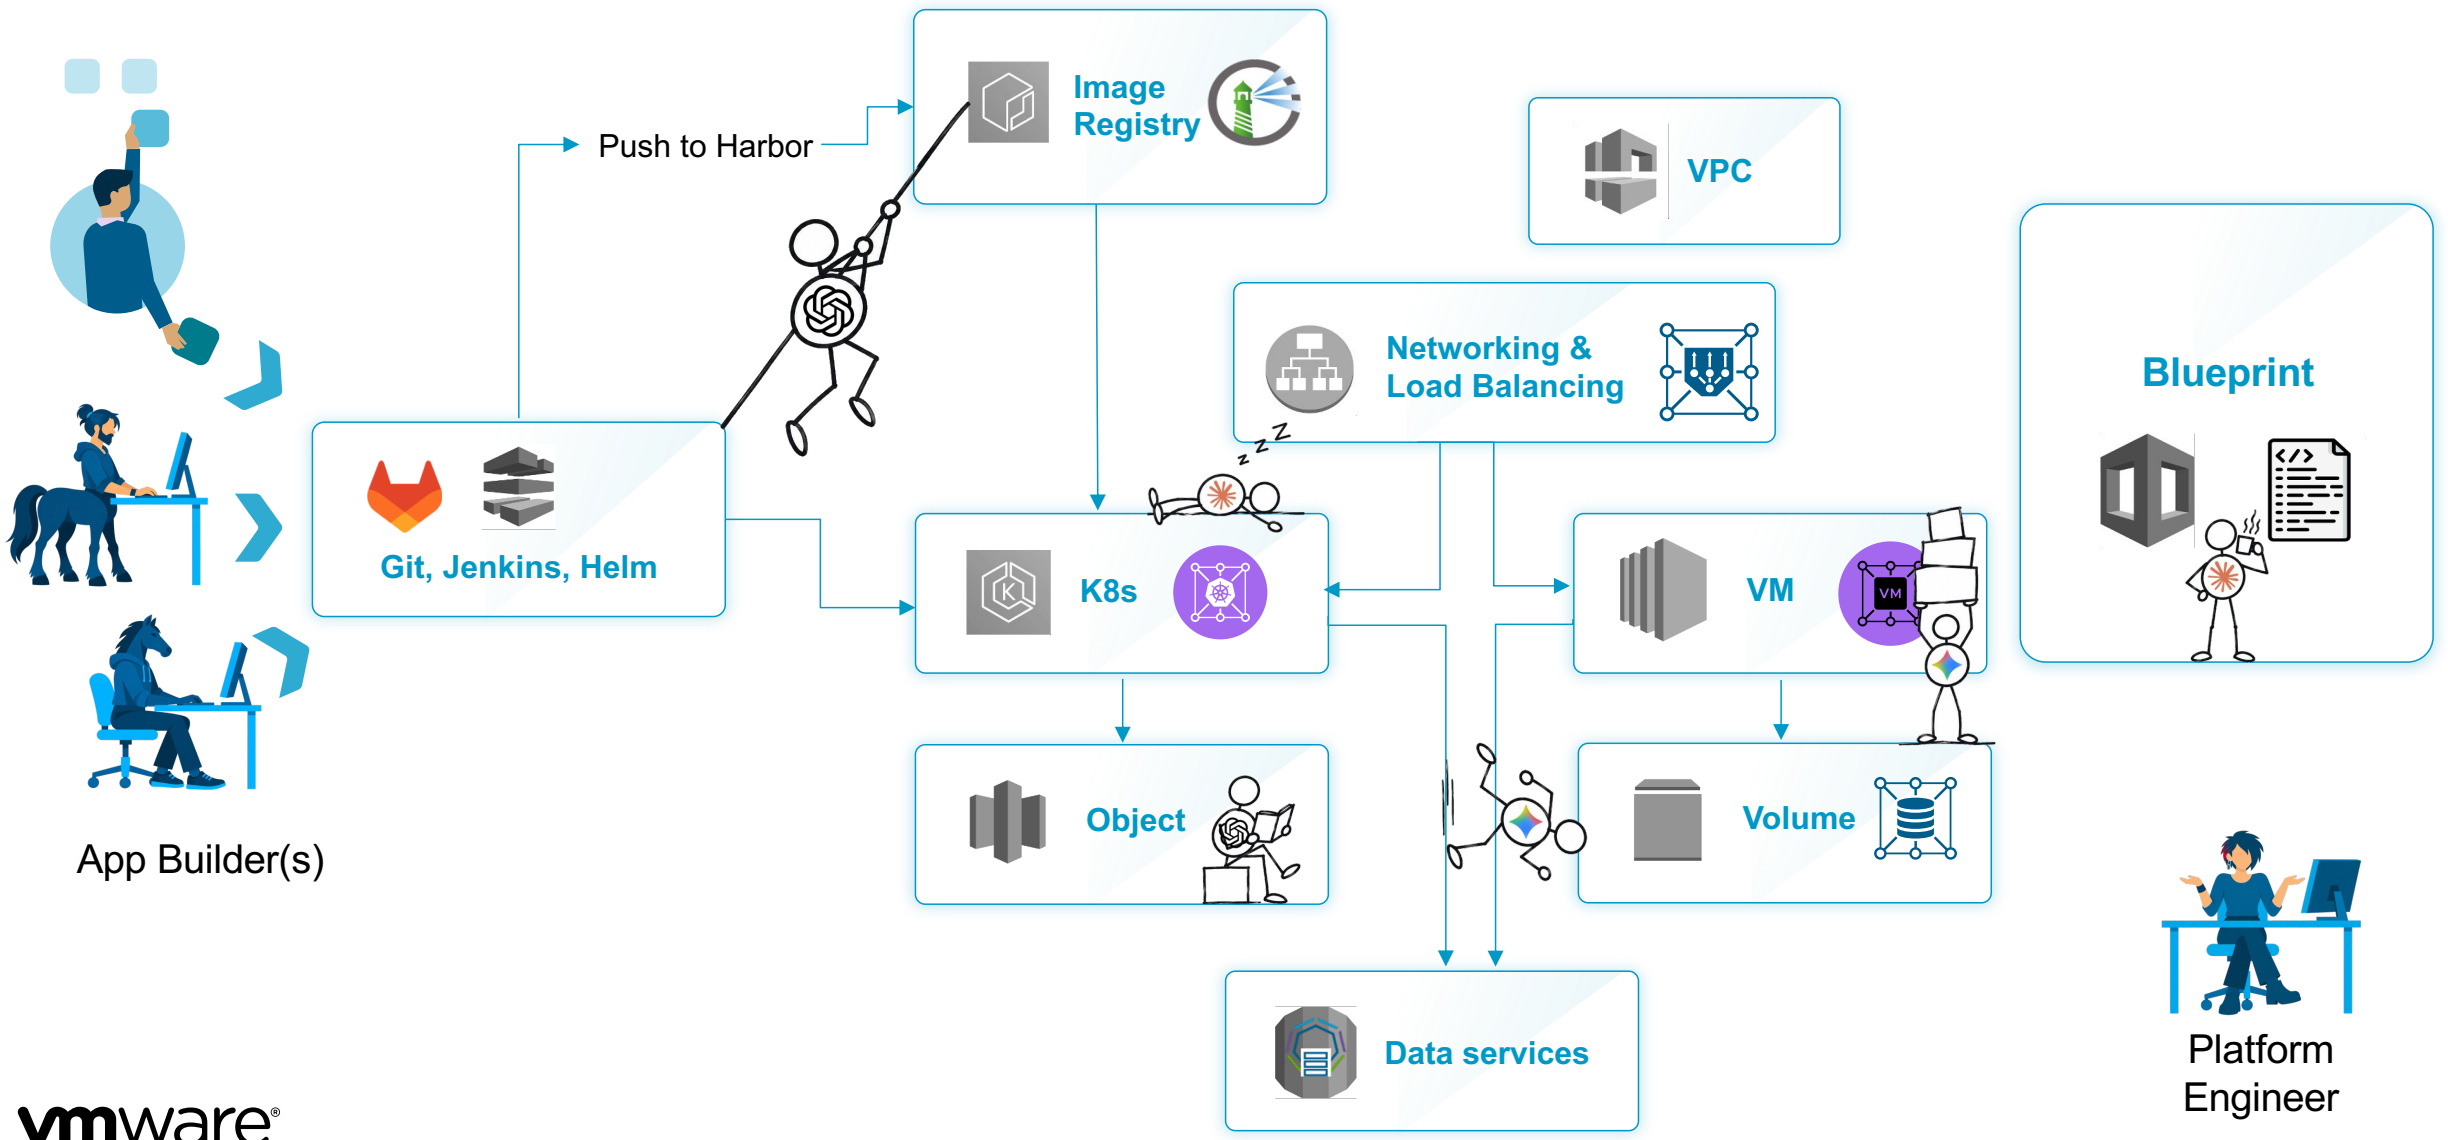


vmware by Broadcom
Broadcom Proprietary and Confidential. Copyright © 2026 Broadcom.
All Rights Reserved. The term "Broadcom" refers to Broadcom Inc. and/or its subsidiaries.

Source: "Inside the Coding Agent-Powered SDLC," Adib Saikali & Frank Denneman, Explore 2026.

# Keep agents out of your PaaS

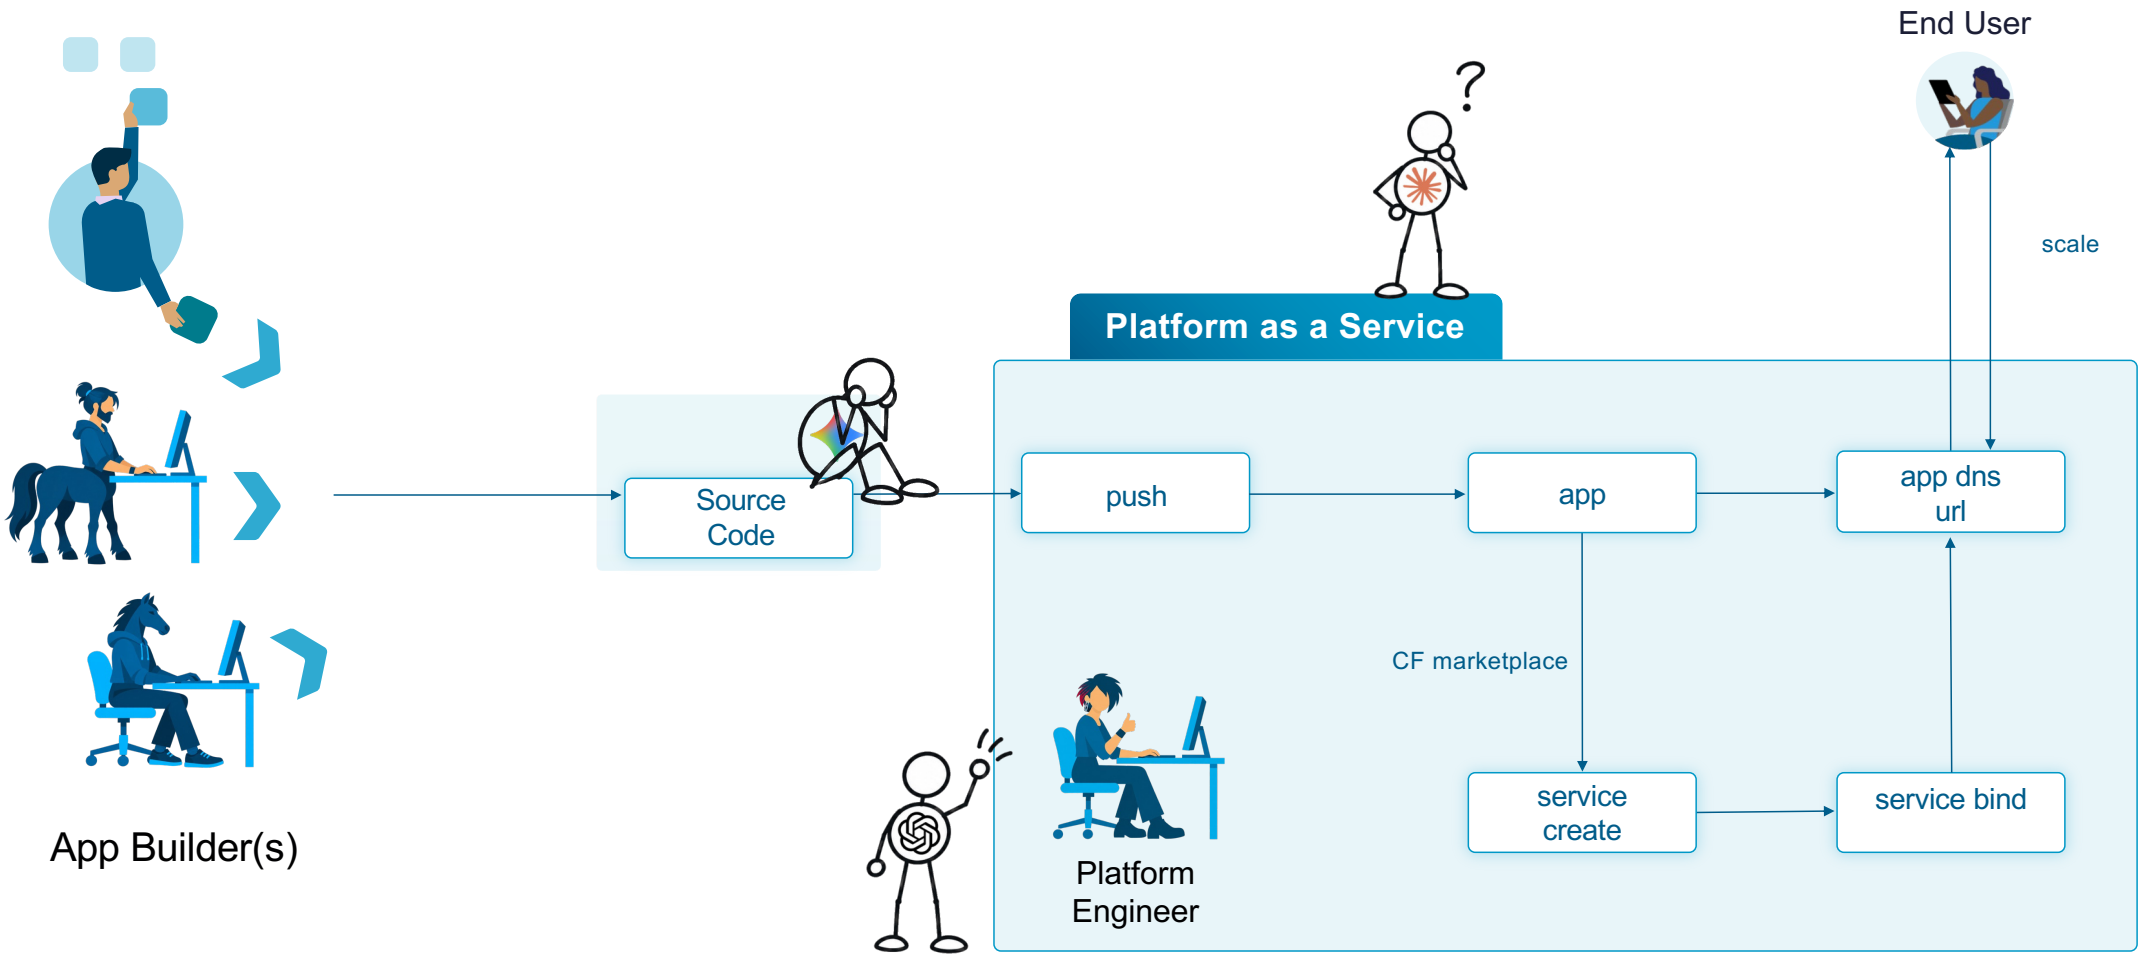


vmware by Broadcom
Broadcom Proprietary and Confidential. Copyright © 2026 Broadcom.
All Rights Reserved. The term "Broadcom" refers to Broadcom Inc. and/or its subsidiaries.

Source: "Inside the Coding Agent-Powered SDLC," Adib Saikali & Frank Denneman, Explore 2026.

# What makes humans productive makes AI even more productive

…and, what makes your enterprise safe from human goofs, helps make it safe from AI goofs

## Agent ***Provisions*** Infrastructure

- Agent must reason about infrastructure topology, sizing, and configuration
- Each agent can make different provisioning decisions
- Dedicated infrastructure can lead to over-provisioning and resource sprawl
- More agent decisions create more opportunities for errors

## Agent ***Consumes*** Services

- Agent selects from platform-defined service plans
- Platform team controls sizing, configuration, policy, and lifecycle
- Services can be backed by shared, right-sized infrastructure
- Every agent gets a repeatable, governed service

vmware by Broadcom
Broadcom Proprietary and Confidential. Copyright © 2026 Broadcom.
All Rights Reserved. The term "Broadcom" refers to Broadcom Inc. and/or its subsidiaries.

Source: "Inside the Coding Agent-Powered SDLC," Adib Saikali & Frank Denneman, Explore 2026.

# Thanks!

cote@broadcom.com

TryTanzu.ai


vmware by Broadcom

Broadcom Proprietary and Confidential. Copyright © 2026 Broadcom.
All Rights Reserved. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries.