meridian IT

CLOUDSMART BY MERIDIAN IT

Your Cloud. Your Way.

A workload-first strategy for hybrid IT, and why placement, not platform, now determines cost, resilience, and control.

Meridian IT Inc. | 2026

Author: James Moore

EXECUTIVE SUMMARY

Cloud strategy is no longer a platform decision

For a decade, cloud strategy was framed as a destination: get out of the data center and into the public cloud. That framing has quietly collapsed. Rising and unpredictable public cloud bills, sweeping changes to VMware licensing, new AI workloads that demand proximity to data, and board-level scrutiny of cyber resilience have forced a more disciplined question: not "which platform?" but "where does each workload actually belong?"

This paper examines the four forces reshaping enterprise infrastructure decisions, makes the case for a workload-first placement strategy across on-premises, private cloud, and public cloud environments, and describes how organizations are using data-driven assessment, rather than assumption, to decide what to keep, move, protect, and modernize before their next renewal, refresh, or migration.

83%
of enterprises plan to repatriate at least some workloads from public cloud (Barclays CIO Survey)

$311.8B
projected hybrid cloud market by 2030, up from $172.8B in 2025 (Mordor Intelligence)

>50%
of virtual machines in a typical estate sit inactive, consuming capacity and licenses

The Cloud Conversation Has Changed

Cloud adoption followed a predictable arc. Organizations moved from traditional on-premises infrastructure to public cloud platforms in pursuit of scalability and speed, and for many workloads, they found both. But as estates grew, so did the consequences of treating the cloud as a single destination. Security, compliance, cost control, and data governance climbed the priority list, and private cloud adoption rose with them.

Today, the correction is visible in both directions. Organizations that went all-in on public cloud are repatriating select workloads, moving them back to private or hybrid environments because the economics, performance, or compliance posture no longer hold at scale. Meanwhile, organizations that never left the data center are discovering that some of their workloads would run better and cheaper in the cloud. Neither group got the platform wrong. They got the placement wrong.

The lesson of the last five years is that cloud strategy is no longer a platform decision. It is a workload decision. Elastic, bursty, and cloud-native workloads still favor public cloud. Regulated, latency-sensitive, and data-heavy workloads favor dedicated infrastructure. Most enterprises need both, plus the connective tissue to govern and protect the whole. This is why hybrid has emerged as the dominant model: it is the only architecture that lets each workload run where it makes the most sense.

> **"90% of organizations will adopt a hybrid cloud approach by 2027."**
>
> Gartner Forecasts, 2024

Four Forces Reshaping Infrastructure Decisions

Hybrid adoption is not being driven by fashion. Four concrete, largely unavoidable pressures are forcing enterprises to revisit workload placement, often on someone else's timeline.

01

VMware licensing upheaval

Broadcom's acquisition of VMware ended perpetual licensing, restructured bundles, and pushed many customers toward subscription models they did not plan for, with material cost increases arriving at renewal. Organizations now face three broad paths: absorb and renegotiate, migrate to an alternative hypervisor, or rearchitect onto a private cloud platform. For many IT leaders, the next VMware renewal date is the single hardest deadline on the infrastructure calendar.

02

Public cloud cost scrutiny

FinOps has matured from a buzzword into a board-level discipline, and the analysis it produces is uncomfortable for some workloads: what made economic sense at one scale does not at another. Repatriation is not an admission of failure. It is evidence that the analysis is finally being done. The strategic question is which workloads to move back, and to what.

03 AI and data gravity

New AI and analytics workloads have specific infrastructure requirements: GPU access, high-throughput storage, and proximity to the data they process. Data gravity is a real constraint: moving compute to the data is usually cheaper and faster than moving the data to the compute. These workloads often land best on dedicated infrastructure, on-premises or in a managed private cloud.

04 Cyber resilience as a board expectation

Ransomware has moved recovery from an IT metric to a governance question. Boards now ask about recovery time objectives, air-gapped backups, and what happens in the first 24 hours after production goes dark. Backup, disaster recovery, and high availability are design criteria that shape the placement decision itself.

FOUR FORCES, ONE COMMON PRESSURE POINT

01
VMware licensing
Forcing question:
Absorb, migrate, or rearchitect before renewal?

02
Cloud cost scrutiny
Forcing question:
Which workloads move back, and to what?

03
AI and data gravity
Forcing question:
Where does the data already live?

04
Cyber resilience
Forcing question:
Can we recover in the first 24 hours?

Every force converges on the same decision: workload placement ↓

A Workload-First Placement Strategy

A workload-first strategy starts from a simple premise: don't assume where workloads belong. Each application, database, and platform has its own profile: performance requirements, compliance constraints, cost behavior, integration dependencies, and recovery needs. Placement should follow that profile, and in practice, workloads sort into three broad tiers.

Environment	Best suited for	Trade-offs to manage
On-premises	Latency-sensitive, tightly integrated, or high-data-gravity workloads; systems with strict sovereignty requirements.	Hardware lifecycle and capital planning; specialized skills to operate; capacity fixed between refresh cycles.
Private cloud (dedicated/managed)	Regulated workloads, predictable-performance systems, and repatriated applications needing cloud economics with dedicated control.	Requires the right partner or in-house depth; design decisions carry longer commitments than public cloud defaults.
Public cloud	Elastic, bursty, seasonal, and cloud-native workloads; development and testing; rapid go-to-market services (SaaS, PaaS).	Cost growth without governance; compliance complexity; egress charges and vendor lock-in.

The tiers are not silos. A workload-first architecture treats them as a single estate, governed consistently, deliberately protected, and rebalanced as economics and requirements change. Four questions determine placement for any given workload:

- **Cost and lifecycle.** What decisions are coming: refresh, renewal, expansion, and what is each option genuinely worth?
- **Control and security.** Which workloads require dedicated performance, compliance controls, or data-location guarantees?
- **Resiliency needs.** Which systems need enhanced backup, disaster recovery, or high availability, and what RTO/RPO does the business actually require?
- **Operating model.** What should internal teams run, and what is better managed by a partner?

Resilience Is a Design Decision, Not an Afterthought

Placement determines where a workload runs; resilience determines whether the business survives when it doesn't. The most common failure pattern is a blanket policy: one backup product, one DR approach applied uniformly to workloads with very different consequences of downtime. Modern resilience is tiered to the workload, delivered as managed services.

BACKUP AS A SERVICE

Cloud-based backup with continuous monitoring, daily deduplication to reduce network load, and data vaulted across geographically diverse facilities, with retention schedules built for regulatory requirements.

DISASTER RECOVERY AS A SERVICE

Recovery of mission-critical applications across mixed environments (Microsoft, VMware, AIX, Red Hat), tiered from contractually guaranteed recovery windows to continuous replication with near-zero recovery objectives.

The tiering matters because resilience is a cost like any other. Continuous replication for a development server is waste; nightly backup for a revenue system is negligence. A workload-first strategy assigns each system the protection its business impact justifies, no more, no less.

The Economics: Consumption Models and Operational Relief

One reason "repatriation" no longer means "going backward" is that the commercial model of dedicated infrastructure has changed. Consumption-based platforms such as HPE GreenLake deliver a cloud operating model: pay for what you use, buffer capacity installed and ready to grow into, self-service, and automation on infrastructure that remains dedicated and under the organization's control. The CAPEX-to-OPEX shift that once required a hyperscaler is now available on-premises or in a partner's data center, without the compliance exposure or egress economics of shared platforms.

The second economic lever is operational. Infrastructure teams are stretched; turnover, retirements, and 24x7 coverage gaps are among the most commonly cited pressures in the industry. Managed operations (monitoring, patching, performance optimization, lifecycle management) let organizations modernize without adding headcount, while governance and compliance posture remain in-house.

When to Act, and Who Should

Workload placement is rarely revisited on its own initiative; it gets revisited when an event forces the question. These are the moments when a placement review creates the most value, because the architecture conversation is already open and a commitment has not yet been made:

- A VMware renewal, licensing cost increase, or migration decision
- Public cloud spend that is rising or hard to predict
- A server, storage, or networking refresh due in the next 12–24 months
- A data center exit, consolidation, or colocation move
- New security, compliance, or data sovereignty requirements
- IT staffing gaps, turnover, or 24x7 coverage challenges
- A mandate to modernize without adding headcount
- Pressure to shift infrastructure spend from CAPEX to OPEX
- New AI or analytics initiatives that need GPU capacity and proximity to data
- A ransomware event, audit finding, or failed DR test

The approach described in this paper fits medium and large enterprises running mission-critical application systems where downtime carries real business or financial consequence, and where high availability, disaster recovery, and 24x7 operational support are requirements rather than options. It is especially relevant in regulated and security-sensitive sectors: manufacturing and industrial, banking and financial services, insurance, healthcare, distribution and supply chain, education, transportation and logistics, and retail.

Before You Decide: Put Data Behind the Placement

Every section of this paper ends in the same place: the right answer depends on the estate. Assumptions about workload behavior, utilization, licensing exposure, capacity headroom, and recovery readiness are exactly the inputs that turn out to be wrong most often. In estates we analyze, more than half of virtual machines typically sit inactive, consuming capacity and license spend that could be reclaimed. That is not a detail; it can change the entire business case for a renewal or refresh.

This is why the workload-first approach begins with measurement. The CloudSmart Assessment, built on HPE CloudPhysics, turns an organization's own system data into a placement roadmap:

1. DEPLOY

An agentless collector installs in about five minutes, with read-only access and no firewall changes.

2. COLLECT

Runs passively for roughly 14 days to build a statistically meaningful picture; first trends visible within 24 hours.

3. REVIEW

Scenario analysis across 200+ metrics is distilled into findings and reviewed with your team.

4. DECIDE

A prioritized, data-backed roadmap: current-state documentation, placement plan, resiliency roadmap, costed options.

The natural moment to run an assessment is before a major commitment, such as a VMware renewal, a hardware refresh, a data center exit, or a DR redesign. Running afterward can only tell you what you should have done.

Cloud, on Your Terms

The organizations navigating this era well are not the ones that picked the right platform; they are the ones that stopped treating the platform as the decision. They place each workload deliberately, protect it proportionately, buy it on commercial terms that match how it behaves, and revisit those choices with data as conditions change. That is the whole of the CloudSmart philosophy: your cloud, your way, decided by evidence, not momentum.

Sources: Gartner Forecasts (2024); Barclays CIO Survey; Mordor Intelligence, Hybrid Cloud Market (2025–2030); Meridian IT assessment data.