Clear Technologies™

WHITE PAPER

The Practitioner's Guide to Hybrid Infrastructure Management: IBM Automation, FinOps, and the Path to Unified Visibility

HOW ENTERPRISE IT TEAMS ARE BREAKING DOWN SILOS, REDUCING WASTE, AND GAINING CONTROL OF HYBRID INFRASTRUCTURE COSTS

March 2026
www.cleartechnologies.net

Table of Contents

The Hybrid Infrastructure Challenge 3

The Rise of Automation and Observability 6

Automation Solutions Enabling Traditional FinOps 8

Hybrid FinOps 12

Adopting the FinOps Practice for Your Environment 17

Conclusion: Advancing Towards Hybrid FinOps 19

ClearTechnologies™

Modern infrastructure is no longer confined to a single environment.

Applications that once ran on a dedicated server or inside a centralized data center now operate across a complex ecosystem of public cloud platforms, Kubernetes clusters, virtual machines, containers, and legacy on-premises systems. Most organizations now manage workloads across multiple providers, internal infrastructure, and rapidly scaling application architectures.

In fact, **93% of organizations now operate a multicloud strategy**,[1] with global spending on public cloud services continuing to accelerate year over year. What began as a way to increase flexibility and innovation has quickly evolved into a new operational challenge: managing a rapidly expanding and increasingly fragmented infrastructure environment.

This complexity is only expected to grow. By 2027, Gartner predicts that **more than 75% of organizations will be running containerized applications in production**,[2] further increasing the scale and distribution of modern infrastructure environments.

As cloud adoption grows, so does complexity.

Engineering teams can provision resources instantly using infrastructure-as-code. Kubernetes clusters can scale workloads dynamically across nodes and regions. Cloud providers offer an almost limitless array of services and pricing models designed to support innovation. While this flexibility enables organizations to move faster, it also breaks traditional models of financial oversight and infrastructure management.

Unlike traditional IT procurement, cloud infrastructure introduces variable operational spend, decentralized purchasing decisions, and resource provisioning that often happens outside of financial governance processes. Engineering teams can spin up resources with a few lines of code, while finance and procurement teams struggle to maintain visibility into where spending is occurring.

The result is an environment where infrastructure is easier than ever to deploy—but significantly harder to understand.

Organizations today face three simultaneous challenges.

Rising infrastructure costs

Cloud adoption has accelerated the demand for compute, storage, and networking resources. The growing use of AI, data analytics, and distributed applications has further increased resource consumption, driving cloud spending upward across industries.[3]

Increasing operational complexity

Applications now span microservices, containers, APIs, and hybrid cloud environments. Monitoring tools often operate in silos, forcing IT teams to manually correlate data across metrics, logs, and traces in order to identify issues. As organizations adopt containerized architectures and microservices-based development, the volume of infrastructure signals and operational data continues to grow rapidly.[4]

Excessive infrastructure waste

Industry research suggests that up to 30% of cloud spending is wasted, often due to overprovisioned resources, idle workloads, or a lack of visibility into how infrastructure is actually being utilized.[5] In Kubernetes environments, costs can become particularly opaque, with teams struggling to attribute cloud spend to specific workloads, teams, or business units.

The consequences extend beyond infrastructure budgets.

Application performance issues can directly impact customer experience and revenue. The average cost of downtime for mid-size and large enterprises can exceed $300,000 per hour, depending on industry and scale.[6] Poor digital performance continues to frustrate customers and erode brand trust.

At the same time, IT teams are increasingly forced into reactive operating models. Without unified visibility across applications and infrastructure, troubleshooting often requires "war room" scenarios where multiple teams manually investigate performance issues across disparate monitoring systems.

What organizations need is not simply more infrastructure—but greater intelligence about how that infrastructure is used.

Understanding where IT spend originates, how workloads consume resources, and how applications perform across hybrid environments has become critical to maintaining both operational efficiency and financial control

Solving this challenge requires a new approach to infrastructure management—one that combines financial transparency, workload-level visibility, automated optimization, and real-time application observability.

Sources

1. Flexera — State of the Cloud Report (multicloud adoption statistics).
2. Gartner — Containerized applications adoption forecast.
3. IDC — Worldwide Public Cloud Spending Guide.
4. Gartner — Observability and monitoring trends in distributed architectures.
5. FinOps Foundation — State of FinOps Report (cloud waste estimates).
6. Uptime Institute — Annual Data Center Outage Analysis.

The Rise of Automation and Observability

As hybrid infrastructure environments have grown more complex, traditional approaches to infrastructure monitoring and management have struggled to keep pace.

Historically, IT operations teams relied on static monitoring tools designed for centralized infrastructure environments. These tools were built to track the health of individual servers, databases, or applications within a relatively predictable architecture. Alerts were generated when predefined thresholds were exceeded, allowing operations teams to respond to performance issues as they occurred.

In modern hybrid environments, however, this approach quickly breaks down.

Applications are no longer monolithic systems running on a small number of servers. Instead, they are composed of distributed microservices, containerized workloads, and dynamic infrastructure that can scale up or down automatically in response to demand.

A single user transaction may now traverse dozens of services across multiple cloud environments before completing.

As a result, infrastructure operations teams are no longer managing a static environment—they are managing a constantly evolving system.

This shift has dramatically increased the volume and complexity of operational data. Every application request generates telemetry across metrics, logs, traces, and infrastructure signals. Container orchestration platforms such as Kubernetes dynamically schedule workloads across clusters, creating environments where infrastructure components can appear and disappear in seconds.

Traditional monitoring tools were not designed for this level of dynamism.

Instead of providing clear insights into system behavior, legacy monitoring solutions often generate overwhelming volumes of alerts without meaningful context. Operations teams may receive hundreds or even thousands of alerts during a single incident, forcing engineers to manually correlate signals across multiple monitoring systems in order to identify the root cause of an issue.

The result is an operational model that is **reactive** rather than **proactive**. Engineers spend valuable time responding to incidents, investigating alerts, and attempting to trace performance problems across distributed systems. In large organizations, these investigations often escalate into cross-team troubleshooting efforts, sometimes referred to as “war room” scenarios, where multiple teams collaborate to isolate the source of a problem.

At the same time, financial complexity has increased alongside operational complexity.

Cloud computing introduced a new consumption-based model for infrastructure spending. Resources can be provisioned instantly and scaled dynamically, which enables innovation but also creates significant challenges for cost governance. Engineering teams can deploy new services rapidly, while finance teams struggle to understand how infrastructure spending maps to business outcomes.

This convergence of operational complexity and financial opacity has driven the emergence of two closely related disciplines: **automation and observability**.

Observability provides the visibility required to understand how modern applications and infrastructure behave in real time. By collecting and correlating telemetry data across distributed systems, observability platforms enable teams to trace application performance, identify anomalies, and diagnose the root cause of issues quickly.

Automation, meanwhile, enables organizations to act on those insights.

Instead of relying solely on manual intervention, automation platforms continuously analyze infrastructure utilization and application demand, making real-time adjustments to resource allocations in order to maintain performance while minimizing waste. These automated optimization capabilities are increasingly essential in environments where infrastructure can change faster than human operators can respond.

Together, observability and automation represent a fundamental shift in how organizations manage modern infrastructure.

Rather than reacting to infrastructure problems after they occur, organizations can now continuously monitor system behavior, detect emerging issues, and automatically optimize resources in response to changing conditions.

This shift is critical for operating hybrid infrastructure environments at scale. Without automation and observability, organizations struggle to maintain both application performance and financial efficiency. With the right visibility and optimization capabilities in place, however, hybrid infrastructure can become not just manageable—but strategically advantageous.

Categories of Solutions

As hybrid infrastructure environments have expanded in scale and complexity, a corresponding ecosystem of operational tools has emerged to help organizations manage them. These tools are not interchangeable. Each addresses a different layer of the problem created by distributed infrastructure: understanding how systems behave, how resources are consumed, and how technology investments translate into operational outcomes.

While vendors often present these capabilities as isolated products, they can be more clearly understood as belonging to several distinct categories. Together, these categories form the operational intelligence layer required to manage modern infrastructure.

Financial Visibility Platforms

The first challenge many organizations encounter when adopting hybrid infrastructure is **financial opacity**.

Traditional IT procurement models relied on predictable capital expenditures and centralized purchasing processes. Cloud infrastructure fundamentally alters this model. Resources can be provisioned instantly, scaled dynamically, and consumed on a usage basis. While this flexibility enables faster development and experimentation, it also introduces a level of financial complexity that traditional accounting frameworks struggle to capture.

Financial visibility platforms emerged to address this gap. Rather than focusing solely on infrastructure metrics, these systems focus on the relationship between technology spending and business outcomes. By integrating financial data with infrastructure consumption metrics, financial visibility platforms allow organizations to allocate technology costs to business units, projects, and operational initiatives.

This approach forms the foundation of **Technology Business Management (TBM)**, a framework designed to help enterprises manage technology spending with the same discipline applied to other financial investments. Platforms in this category enable organizations to model infrastructure costs, forecast future spending, and connect operational consumption with financial accountability.

In large enterprises where infrastructure budgets may span multiple cloud providers, internal data centers, and third-party services, financial visibility becomes essential for maintaining strategic control over technology investments.

Workload-Level Cost Intelligence

While financial visibility platforms provide a macro-level view of technology spending, they often lack the granularity required to understand how individual applications consume infrastructure resources.

This limitation becomes particularly pronounced in containerized environments. Kubernetes, which has become the dominant platform for container orchestration, distributes workloads dynamically across clusters of compute resources. While this architecture improves scalability and resilience, it also obscures the relationship between infrastructure usage and application behavior.

Workload-level cost intelligence tools emerged to address this challenge by attributing infrastructure consumption to specific services, namespaces, or workloads within a container environment. By analyzing resource utilization across clusters, these systems allow engineering teams to understand how architectural decisions translate into operational costs.

This level of visibility enables organizations to identify inefficiencies such as over-provisioned containers, idle workloads, or poorly configured scaling policies. It also enables engineering teams to evaluate the financial impact of design choices at the application level, creating a direct connection between software architecture and infrastructure economics.

As containerized environments continue to grow in scale, workload-level cost intelligence has become a critical component of modern FinOps practices.

Infrastructure Optimization Platforms

Understanding how infrastructure resources are consumed is only one part of the challenge. Organizations must also ensure that those resources are allocated efficiently.

Hybrid environments often contain significant amounts of unused or misallocated infrastructure capacity. Virtual machines may be provisioned with more resources than required, containers may request more compute than they actually consume, and infrastructure scaling policies may fail to adapt to changing workloads.

Infrastructure optimization platforms address this problem by continuously analyzing application demand and adjusting resource allocations to maintain performance while minimizing waste.

Unlike traditional capacity planning tools, which rely on static analysis or manual intervention, modern optimization platforms operate dynamically. They ingest telemetry from applications and infrastructure layers, evaluate performance requirements in real time, and recommend or automatically implement adjustments to resource allocation.

These adjustments may include resizing virtual machines, reallocating container resources, or shifting workloads across infrastructure environments. By aligning infrastructure supply with application demand, optimization platforms ensure that applications receive the resources they require without unnecessary overprovisioning.

This approach not only improves operational efficiency but also reduces the financial impact of infrastructure waste.

Application Observability Platforms

While financial visibility and infrastructure optimization address the economics of infrastructure, organizations must also understand how applications behave within these environments.

Modern applications are no longer monolithic systems running on a single server. They are distributed architectures composed of microservices, APIs, message queues, and external services that may operate across multiple infrastructure environments.

In such systems, diagnosing performance issues becomes significantly more complex. A single user transaction may traverse dozens of services before completing, making it difficult to identify the source of latency or failure.

Application observability platforms emerged to provide the visibility required to manage these distributed systems. By collecting telemetry data across metrics, logs, traces, and service dependencies, observability platforms allow organizations to reconstruct how applications behave in real time.

This capability enables engineering teams to trace transactions across service boundaries, identify performance bottlenecks, and diagnose failures more quickly. In complex hybrid environments where infrastructure and applications interact across multiple layers, observability becomes essential for maintaining operational reliability.

Why These Categories Matter

These categories represent different perspectives on the same infrastructure environment.

Financial visibility platforms focus on technology spending.
Workload cost intelligence tools focus on resource consumption.
Infrastructure optimization platforms focus on resource efficiency.
Observability platforms focus on application behavior.

Each addresses a different dimension of the hybrid infrastructure challenge.

Taken together, they form the operational intelligence layer required to manage modern distributed systems.

The next section examines how IBM's portfolio of infrastructure management platforms aligns with these categories and how they can be used together to provide a unified view of hybrid infrastructure operations.

Automation Solutions Enabling Traditional FinOps

The categories outlined in the previous section represent different dimensions of the hybrid infrastructure challenge: financial visibility, workload-level cost intelligence, infrastructure optimization, and application observability. In practice, organizations rarely deploy these capabilities as isolated tools. Instead, they form a layered operational model that allows technology teams to understand how infrastructure resources are consumed, how applications behave, and how technology investments translate into business outcomes.

IBM's infrastructure management portfolio provides solutions that align closely with each of these operational layers.

Financial Visibility: Apptio

Financial visibility begins with understanding how technology investments translate into business value. IBM Apptio provides the foundation for this layer through its implementation of the Technology Business Management (TBM) framework.

Apptio enables organizations to model and track technology spending across cloud providers, internal infrastructure, and third-party services. By translating infrastructure consumption into financial metrics aligned with business units and operational initiatives, Apptio allows finance and technology leaders to understand the economic impact of their infrastructure decisions.

For organizations operating hybrid environments, this financial layer is critical. Without a clear understanding of how infrastructure spending maps to applications and business services, it becomes difficult to evaluate technology investments or manage costs strategically.

Workload Cost Intelligence: Kubecost

While financial management platforms provide macro-level visibility into technology spending, engineering teams often require a more granular understanding of how infrastructure resources are consumed at the application level.

IBM Kubecost addresses this challenge in containerized environments by attributing cloud and infrastructure costs to specific workloads within Kubernetes clusters. By mapping infrastructure consumption to namespaces, services, and applications, Kubecost allows engineering teams to understand how software architecture decisions influence operational spending.

This visibility enables organizations to identify inefficient workloads, optimize container resource allocation, and align engineering decisions with financial outcomes. In organizations where Kubernetes has become the foundation of application deployment, workload-level cost intelligence is a key component of effective FinOps practices.

Infrastructure Optimization: Turbonomic

Even with strong financial visibility and cost attribution, organizations must still ensure that infrastructure resources are allocated efficiently.

IBM Turbonomic addresses this challenge through application resource management, continuously analyzing application demand and infrastructure capacity to ensure that workloads receive the resources they require without unnecessary overprovisioning.

Unlike traditional capacity planning tools that rely on static analysis, Turbonomic operates dynamically. By ingesting telemetry from applications and infrastructure systems, the platform identifies opportunities to resize resources, rebalance workloads, or adjust infrastructure allocations in real time.

The result is an infrastructure environment where performance and cost efficiency are managed simultaneously rather than treated as competing objectives.

Application Observability: Instana

INSTANA
an IBM Company

The final layer of the hybrid infrastructure intelligence stack focuses on application behavior. As modern applications become increasingly distributed, diagnosing performance issues requires visibility into how services interact across infrastructure environments.

IBM Instana provides this capability through automated observability. By collecting telemetry across metrics, traces, logs, and service dependencies, Instana allows organizations to understand how applications behave across complex hybrid architectures.

This visibility allows engineering teams to trace transactions across microservices, identify bottlenecks in real time, and resolve performance issues more quickly. In distributed environments where application performance is directly tied to infrastructure behavior, observability becomes a critical component of operational reliability.

Visual One Intelligence: Operationalizing Hybrid FinOps

While these platforms each address important aspects of hybrid infrastructure management, organizations often struggle to integrate their insights into a single operational model. Financial systems, infrastructure telemetry, and application monitoring tools frequently operate in separate domains, requiring teams to manually reconcile data in order to understand the full economic picture of their infrastructure environment.

This challenge becomes especially pronounced in hybrid environments where workloads span both public cloud platforms and on-premises infrastructure.

Many FinOps tools were designed primarily for cloud-native environments, relying on billing APIs and native metadata tagging to track resource consumption. However, enterprise infrastructure rarely exists exclusively in the cloud. Data centers, virtualized environments, and legacy systems often represent a substantial portion of an organization's infrastructure footprint, yet they lack the tagging and billing structures required by traditional FinOps platforms.

Visual One Intelligence addresses this gap through a **"Hybrid FinOps"** platform designed specifically for hybrid infrastructure environments. The platform aggregates technical and financial data across cloud and on-prem systems, creating a unified operational view of infrastructure consumption and cost.

Key capabilities distinguish this approach.

Unified economic modeling

Visual One Intelligence converts on-premises capital expenditures—such as hardware depreciation, power, maintenance, and operational labor—into normalized daily costs comparable to cloud operational expenses. This enables organizations to evaluate infrastructure spending across environments using a single financial framework.

Automated tag normalization.

Rather than requiring perfect tagging compliance across every infrastructure platform, the platform applies translation and normalization at the data layer. This allows inconsistent tags—such as variations of production or environment labels—to be resolved automatically into consistent business dimensions.

Hybrid infrastructure visibility

By integrating telemetry and cost data across cloud platforms, virtualization layers, and physical infrastructure, VisualOne Intelligence enables organizations to compare the true cost of workloads across environments. This visibility allows infrastructure teams to evaluate migration decisions, optimize resource allocation, and understand the full cost of running applications across hybrid environments.

Taken together, these capabilities allow organizations to move beyond isolated infrastructure monitoring or cost reporting toward a fully operational **Hybrid FinOps model**—one in which financial accountability, infrastructure efficiency, and application performance are managed as part of a single system.

Hybrid infrastructure environments vary widely in scale and complexity. Some organizations operate primarily in public cloud environments, while others maintain significant on-premises infrastructure alongside cloud workloads. As a result, the tools and operational capabilities required to manage these environments effectively will differ from one organization to another.

Determining which capabilities are most relevant begins with understanding how infrastructure is currently managed across several key operational dimensions.

Assessing Infrastructure Visibility

The first step in evaluating hybrid infrastructure management capabilities is determining whether the organization has sufficient visibility into application behavior and infrastructure performance.

In distributed environments, a single application transaction may pass through dozens of services, infrastructure layers, and external dependencies before completing. Without observability capabilities capable of tracing these interactions, diagnosing performance issues becomes increasingly difficult as architectures grow more complex.

Organizations that lack clear insight into how applications behave across hybrid infrastructure environments often benefit from implementing observability platforms capable of automatically mapping service dependencies and collecting telemetry across metrics, logs, and distributed traces.

Evaluating Infrastructure Efficiency

This layer of visibility provides the foundation required for understanding how infrastructure decisions ultimately affect application performance and reliability.

This layer of visibility provides the foundation required for understanding how infrastructure decisions ultimately affect application performance and reliability.

Once visibility into application behavior is established, organizations must determine whether their infrastructure resources are being used efficiently.

Hybrid environments frequently contain over-provisioned infrastructure capacity, particularly in virtualized environments where workloads may be allocated more compute, memory, or storage resources than they actually require. Over time, these inefficiencies accumulate and can result in substantial infrastructure waste.

Infrastructure optimization platforms address this challenge by continuously analyzing application demand and adjusting resource allocations dynamically. By aligning infrastructure supply with application demand, these systems help organizations maintain performance while minimizing unnecessary resource consumption.

Understanding Workload-Level Costs

For organizations operating containerized environments, understanding how individual workloads consume infrastructure resources becomes increasingly important.

Container orchestration platforms such as Kubernetes distribute workloads dynamically across clusters, making it difficult to attribute infrastructure spending to specific applications or teams without specialized cost intelligence tools.

Workload-level cost intelligence platforms provide this capability by mapping infrastructure consumption to namespaces, services, and workloads within Kubernetes environments. This visibility allows engineering teams to understand the financial implications of their architecture decisions and optimize container resource usage accordingly.

Establishing Financial Governance

Infrastructure visibility and optimization capabilities must ultimately connect to financial governance frameworks that allow organizations to manage technology spending strategically.

Technology Business Management platforms enable organizations to translate infrastructure consumption into financial metrics aligned with business units and operational initiatives. By integrating infrastructure telemetry with financial models, these systems provide leadership teams with a clearer understanding of how technology investments contribute to business outcomes.

This financial perspective is essential for organizations seeking to manage infrastructure spending as a strategic business function rather than a purely technical concern.

Advancing Toward Hybrid FinOps

While the capabilities described above address important aspects of infrastructure management, many organizations encounter an additional challenge when operating hybrid environments: reconciling the economics of cloud infrastructure with the economics of on-premises systems.

Most FinOps platforms were originally designed for cloud environments, where billing data and infrastructure usage can be tracked directly through provider APIs. However, enterprise environments often include substantial on-premises infrastructure investments whose costs are distributed across capital expenditures, operational expenses, maintenance contracts, and facility overhead.

Without a mechanism for translating these costs into a comparable operational model, organizations struggle to evaluate the true economics of hybrid infrastructure environments. Decisions about workload placement, cloud migration, and infrastructure investment may therefore be based on incomplete financial information.

Hybrid FinOps platforms address this challenge by normalizing infrastructure economics across both cloud and on-premises environments. By converting capital expenditures and operational costs associated with data center infrastructure into comparable daily or usage-based cost models, organizations can evaluate infrastructure decisions using a consistent financial framework.

VisualOne Intelligence was designed specifically to support this Hybrid FinOps model.
By integrating infrastructure telemetry, cost attribution data, and normalized financial models across hybrid environments, the platform enables organizations to understand the full economic impact of their infrastructure decisions.

Conclusion

Cloud workloads, virtualized environments, and physical infrastructure can be evaluated within the same operational framework, allowing technology leaders to make more informed decisions about resource allocation, workload placement, and long-term infrastructure strategy.

In practice, Hybrid FinOps allows organizations to move beyond isolated infrastructure monitoring or cost reporting toward a unified operational model in which application performance, infrastructure efficiency, and financial accountability are managed together.