



Multi-Cloud Visibility & Governance Scorecard

Assess where multi-cloud visibility breaks — and build a defensible internal case for fixing it.

Built for Cloud Operations, Platform Engineering, CMDB/ITSM, FinOps, Security, and GRC teams.

Cloud visibility breaks before cloud governance does

Most enterprises do not lack cloud data. They lack a consistent way to turn data from provider consoles, CMDBs, monitoring systems, billing exports, vulnerability scanners, and ITSM platforms into the same operational answer.

An AWS resource can appear active in the provider console, stale in the CMDB, unallocated in the FinOps report, exposed in a security tool, and assigned to a closed team in ITSM. Each record may be technically valid. Together, they cannot support a reliable decision.

This is why multi-cloud visibility should not be measured by the number of dashboards available.

Visibility should be measured by whether teams can answer four questions consistently.

Use the scorecard to test four operational questions:

01 What is in scope?

Can we establish a reconciled baseline of the infrastructure that should be visible?

02 Can we trust what we see?

Can we reconcile identity, source authority, freshness, and conflicting records?

03 Who can act?

Can a condition be routed to an accountable owner with the authority to remediate, approve, fund, or accept risk?

04 Can we make and defend a decision?

Can teams connect infrastructure context to applications, impact, cost, security, evidence, and the next action?




More data does not create governance. Governance begins when infrastructure state can be resolved to scope, ownership, business context, and action.



Mikhail Malamud
General Manager

Fix the earliest weak dependency, not the loudest symptom

The scorecard evaluates seven connected capabilities. Each domain depends on context established earlier in the model.



Scoring method

Each domain contains two questions. Add both answers for a score from 0 to 4.

- 0 Not established
- 1 Partial
- 2 Operational

Find the first place your operating model breaks

Each domain depends on context established earlier.
A downstream capability can look mature while still
relying on incomplete scope, stale data, or unresolved
ownership.



What infrastructure is actually within the organization’s operational scope?

Cloud and hybrid estates rarely develop through a single, controlled architecture. They expand through new cloud accounts, subscriptions, projects, acquired environments, regional deployments, Kubernetes clusters, and infrastructure created outside standard onboarding processes.

Each team has a plausible inventory, but no one can demonstrate that it represents the complete operational scope. A missing subscription or project may be absent from the CMDB, security coverage, monitoring, cost allocation, and reporting at the same time.

This is not simply an inventory problem. It is a failure to reconcile expected operational scope with observed coverage. Until that comparison is repeatable, every downstream governance decision rests on an unverified baseline.

Expected scope vs. Observed inventory

	AWS accounts Expected: 128 accounts	
	Azure subscriptions Expected: 96 subscriptions	
	GCP projects Expected: 72 projects	
	VMware estates Expected: 18 estates	

Assess your maturity

1. Can the organization produce a current, reconciled list of all in-scope environments?

0

Not established

1

Partial

2

Operational

2. Can teams identify infrastructure that is missing, stale, or outside coverage?

0

Not established

1

Partial

2

Operational

0-1 Unknown scope · 2-3 Partial baseline · 4 Operational baseline

How reliably do systems describe the same asset and its current state?

The same infrastructure asset may appear across cloud platforms, CMDB, monitoring, security, FinOps, and ITSM systems. Each record can be technically valid while using a different identifier, update cycle, ownership value, lifecycle state, or classification.

The issue is not that every value necessarily conflicts. A cloud provider’s **Running** state and a CMDB’s **Active** lifecycle state may both be correct. The problem is that teams often lack a defined rule for determining whether records describe the same asset, which source governs each decision-critical field, whether the value is current enough for the decision, and how disagreements are resolved.

Records may refer to the same asset, but cannot yet support the same decision

Cloud provider	State: Running
CMDB	State: Active · Updated 17 days ago
Security scanner	Owner: Unknown
FinOps mapping	Cost center: Unallocated

Trusted infrastructure data connects stable asset identity with field-level authority, defined freshness expectations, and managed conflict resolution.

Assess your maturity

1. Are critical records reconciled using stable identity rules and field-level source authority?



2. Can teams verify freshness and resolve conflicting values through a defined process?




0-1 Unknown scope · 2-3 Partial baseline · 4 Operational baseline

Who is accountable when an infrastructure condition requires action?

Ownership exists when an infrastructure condition can be routed to the current person or team authorized to act. A resource is not operationally owned simply because an **Owner** field exists.

Operational ownership means the assigned person or team can make the required decision, or the condition follows a defined fallback route when ownership is missing, stale, or disputed.



Visibility without accountable routing does not produce action.

Assess your maturity

1. Can conditions resolve to the correct accountable owner for the required action?



2. Do missing or outdated ownership assignments enter a managed fallback process?



0-1 Unknown scope · 2-3 Partial baseline · 4 Operational baseline

Can you trace every cloud resource to business impact?

An infrastructure alert identifies the affected resource, but not what the business depends on. When `payments-prod-rds-03` degrades, teams still need to determine the application, downstream services, blast radius, recovery owner, and business impact.

Without maintained relationships, this context is reconstructed from tags, diagrams, Terraform, monitoring labels, tickets, and individual knowledge. The resource is visible, but the impact is not.

Application and dependency context connects infrastructure state to applications, shared services, business processes, owners, and recovery priority.

A visible resource is not the same as visible business impact.

Application / service	Owner	Environment	Critical dependencies	Business context
Payments API	Payments Platform	Production	Checkout screening, Fraud	Revenue-critical, PCI DSS
Patient records	Clinical Systems	Staging	Identity, Data Platform	HIPAA
Customer portal	Digital Experience	Production	Identity, Notification service	Customer-facing, GDPR

Assess your maturity

1. Can critical assets be linked to current applications, services, and accountable owners?

0

Not established

1

Partial

2

Operational

2. Can teams determine direct and shared dependencies during incidents or change?

0

Not established

1

Partial

2

Operational


0-1 Unknown scope · 2-3 Partial baseline · 4 Operational baseline

How much reconstruction is required to produce a trusted operational answer?

The question usually sounds simple: [Which EC2 instances belong to this application, and which of them changed before the incident?](#)

But answering it often means opening several systems: AWS Console for the current resource state; CMDB for application mapping; monitoring for the incident timeline; security tooling for findings and exposure; ITSM for recent changes; Slack or Teams to find the person who knows which data is actually current.

By the time the analyst reconciles resource IDs, checks stale CIs, compares timestamps, and validates ownership, the result may already be outdated. Another engineer may repeat the same investigation and reach a different answer because they used different filters, exports, or source precedence.



Assess your maturity

1. Can teams answer recurring cross-system questions through a reusable path?

0

Not established

1

Partial

2

Operational

2. Does the result contain enough current context to support the next decision?

0

Not established

1

Partial

2

Operational

0-1 Unknown scope · 2-3 Partial baseline · 4 Operational baseline

Does cloud spend resolve to accountable owners and actionable decisions?

Idle and underused resources remain active. RI and Savings Plan utilization is difficult to explain. Cost anomalies arrive without an accountable owner. Finance, FinOps, engineering, and product teams produce different numbers because allocation still depends on incomplete tags and spreadsheet correction.

The issue is not access to billing data. It is whether cloud and Kubernetes spend can be mapped consistently to applications, environments, cost centers, and owners, then converted into a decision someone can implement.

Account	Application	Environment	Cost center	Owner	MOM
datastream	Data Platform	Production	897	Data Engineering	+24%
trynex	Payments API	Production	754	Payments Platform	-17%
elyra	Customer Portal	Staging	443	Digital Experience	

Trynex cost context — Trynex / Payments API		
Cost category	Monthly cost	Allocation status
AWS EC2 Instances	\$2,820	Allocated
AWS EBS Volumes	\$1,378	Allocated
AWS RDS Instances	\$129.78	Allocated
AWS VPC Route Tables	\$84.27	Shared platform cost
AWS ELB Load Balancers	\$67.72	Allocated
Total: \$4,080		

Assess your maturity

1. Can cloud and Kubernetes spend be attributed without manual spreadsheet correction?

0

Not established

1

Partial

2

Operational

2. Do cost signals reach owners able to make or implement the decision?

0

Not established

1

Partial

2

Operational

0-1 Unknown scope · 2-3 Partial baseline · 4 Operational baseline

Are security and compliance conditions connected to priority, ownership, action, and evidence?

Configuration drift, benchmark failures, and policy violations are often found through separate provider tools, scheduled scans, or ad hoc API exports. Teams receive findings without consistent asset context, business criticality, ownership, SLA, or exception status.

The same configuration may be evaluated differently across AWS, Azure, GCP, Kubernetes, and hybrid infrastructure. Evidence for ISO, NIST, HIPAA, or internal policies is then reconstructed from screenshots, exports, tickets, and spreadsheets.



Assess your maturity

1. Are assets evaluated against consistent policies, benchmarks, and compliance requirements?

0

Not established

1

Partial

2

Operational

2. Are findings prioritized, routed, and closed with ownership, SLAs, exceptions, and evidence?

0

Not established

1

Partial

2

Operational

0-1 Unknown scope · 2-3 Partial baseline · 4 Operational baseline

Record the score and the evidence in one place

Score each domain from 0 to 4 using the lowest capability level that fully describes the current state across the assessed scope.

Domain	Score	Evidence
1. Baseline coverage	— / 4	Yes / Partial / No
2. Data trust	— / 4	Yes / Partial / No
3. Ownership and accountability	— / 4	Yes / Partial / No
4. Application and dependency context	— / 4	Yes / Partial / No
5. Infrastructure answerability	— / 4	Yes / Partial / No
6. FinOps and cost management	— / 4	Yes / Partial / No
7. Security and compliance context	— / 4	Yes / Partial / No
Total		— / 28

Scoring rule

Do not award a higher score for planned improvements, isolated examples, or tools that are available but not used consistently.

Case study

How Boeing made multi-cloud and AI spend explainable

Clouware connected cost data across AWS, Azure, and Google Cloud with the workloads, accounts, and stakeholders responsible for it.

1 Overview

Boeing runs engineering, analytics, and manufacturing workloads across AWS, Azure, and Google Cloud.

As AI usage became a material cost driver, the FinOps team needed faster, consistent explanations of what caused spend and who could act on it.

34%

Faster cost reporting and chargeback cycles with less manual effort

37%

Drop in cost anomalies thanks to near real-time alerting

\$958,250

Annual savings from S3 cleanup, SP/RI optimization, and rightsizing

2 Challenge

- 1 AI costs were difficult to map to teams and workloads
- 2 SP and RI coverage was not consistently tracked
- 3 Idle and underused S3 storage lacked a clear cleanup owner
- 4 Anomalies, chargeback, and forecasting depended on reactive or manual workflows

3 Clouware operating layer

1 Established a governed cost scope

Clouware connected AWS, Azure, and GCP billing feeds for account-, workload-, and stakeholder-level reporting.

2 Exposed AI cost drivers

AI cost tracking separated Azure OpenAI input and output tokens and monitored Bedrock and SageMaker usage for allocation and forecasting.

3 Turned waste into owned work


SP/RI coverage, S3 optimization, anomaly detection, and rightsizing identified the affected account, financial impact, and required action.

Connect fragmented infrastructure data to accountable decisions

The scorecard identifies where multi-cloud visibility and governance break. Clouware turns fragmented multi-cloud infrastructure into one operational picture.

It brings AWS, Microsoft Azure, Google Cloud, Kubernetes, VMware, and connected tools into a CMDB built for cloud operations. With Clouware, teams can see what exists, who owns it, what it supports, what it costs, and where risk or control gaps require action. The platform connects Cloud Operations, FinOps, Security, and Governance around the same infrastructure context — so decisions do not begin with manual reconstruction.

Clouware partners with AWS, Microsoft Azure, and Google Cloud, is available in their marketplaces, and is a FinOps Certified Platform.



Establish multi-cloud visibility and trusted infrastructure context

Discover, normalize, search, and reconcile infrastructure across cloud, hybrid, and connected operational sources. Connect assets with applications, owners, environments, relationships, cost context, monitoring status, and change history.

Connect cloud spend to accountable decisions

Normalize billing and usage data, enrich spend with CMDB context, and allocate cost to applications, teams, environments, cost centers, and shared-cost buckets. Support forecasting, anomaly investigation, optimization, and showback or chargeback workflows.

Connect controls to context, exceptions, and evidence

Evaluate infrastructure against defined policies and provider-specific controls. Connect findings to affected assets, applicable requirements, accountable owners, exceptions, remediation workflows, and audit evidence.

Turn fragmented risk findings into accountable remediation

Bring vulnerability findings into infrastructure, application, ownership, and change context. Prioritize remediation by affected asset and service impact, then route the work through operational workflows.

Review your score with Cloudaware

See how Cloudaware can help establish a trusted infrastructure baseline and connect cloud inventory with ownership, cost, security, compliance, and operational workflows.

[Schedule demo](#)

Bring your completed scorecard.
We'll review the first broken dependency.

