

DATA STRATEGY TRACKS

Cloud data strategy


Moving data and analytics to the cloud changes cost, speed and control at once. Plan it as a sequence of deliberate choices about what moves, in which order, on which ecosystem, and how spend is governed.

THE SO WHAT OF THIS TRACK

Sequence the move by data gravity and value, choose the ecosystem style deliberately, and run financial governance from day one with price-performance as the measure.

Part of stage 5: Sequence

Cloud decisions fix the platform, the cost model and much of the governance for years. Made as an infrastructure project, they move workloads without moving value. Made inside the data strategy, they follow the priorities on the roadmap, respect data gravity and come with financial governance from the first month.



QUESTIONS THIS TRACK ANSWERS

- Which workloads should move first, and which belong together?
- Which cloud ecosystem style fits our commitments, skills and data spread?
- How will we keep spend predictable when pricing is consumption-based?
- Which on-premises habits must change, and who will change them?
- What do regulation, sovereignty and latency rule in or out?

1

The method

Plan, migrate and operate data and analytics in the cloud

Cloud is different

Habits from the data centre carry over badly. Three differences matter most.

The implementation is hidden

You choose services, not hardware.
Tuning means understanding how each service consumes resources.

Pricing is synthetic

Units, credits and tiers stand between your workload and its cost. The same query can cost very different amounts.

Old best practices must change

Sizing for the peak, keeping everything running and copying data freely all cost money every hour.

Benefits and concerns, weighed together

What draws organisations in

- Rapid provisioning and fast access to new capability
- Elastic scale and paying for what is used
- No capital outlay and fewer infrastructure specialists
- Faster, more agile development

What holds them back

- Regulation, data sovereignty and security
- Governance across many services and accounts
- Latency and compatibility with systems that stay on site
- Loss of control and fear of lock-in

SO WHAT

New financial models and the cultural shift appear on both sides: they are the benefit if managed and the risk if not.

Four hybrid and multicloud patterns

Lifecycle split

Development, test or recovery run in the cloud while production stays where it is, or the reverse.

Use-case specific

Particular workloads, such as data science or a new data product, run in the cloud.

Spanning architecture


One architecture spans on-site and cloud, with data and processing placed where each fits best.

Intercloud

Workloads and data are spread across more than one cloud provider by design.

Let data gravity set the migration order

Systems that exchange data constantly form clusters. Move one and the rest of its cluster follows.



SO WHAT

Migrating one system out of a tight cluster creates latency and transfer cost that erode the case for moving at all.

Cloud data ecosystem styles

Style	Strengths	Watch-outs
Provider-native	Deepest out-of-the-box integration; fastest potential time to value	Lock-in to one provider; harder to reach data in other clouds
Mixed native and independent	Weaker native services can be swapped for better ones; a common path to multicloud	More integration effort up front
Self-contained independent software	Easiest lift and shift; suits an existing software commitment; some portability	More work to use surrounding provider services; lock-in to the software supplier
Intercloud	Uses each provider's unique strengths; fits data already spread across clouds	Hardest to integrate and manage; complex cost governance; transfer fees and latency

SO WHAT

Every style still rests on the provider's identity, billing and monitoring services, so integration with them is never optional.

Pricing models by workload

Match the pricing model to how predictable each workload is. Good fit, consider, or poor fit.

Workload	Resource-based	Consumption-based	Bounded or burst	Blended
Operational and transactional	Good	Poor	Consider	Good
Dashboards	Good	Consider	Good	Good
Standard reports	Good	Consider	Good	Good
Development and test	Consider	Good	Good	Consider
Exploratory	Poor	Good	Consider	Consider
Ad hoc	Poor	Good	Good	Consider

SO WHAT

Steady workloads favour reserved resources; unpredictable ones favour consumption with spending limits.

Operate for scarcity in a world of abundance

Price-performance is the measure

Judge services by cost per unit of useful work, not by speed or list price alone.

Prove it before you commit

Run proofs of concept on real workloads to learn how each service behaves and bills.

Govern spend from day one

Budgets, tags, alerts and owners before the first workload, not after the first overrun.

Expect early overruns

Plan for learning costs in the first quarters and review spend weekly until it settles.

2

Tools and evidence

Score it, measure it, see it in six industries

Which cloud ecosystem style fits you?

Answer six questions. The style with the strongest fit is recommended, with what it implies for the strategy.

Question	Options
Where does your data live today?	Mostly on site; Mostly with one cloud provider; Spread across several clouds
How committed are you to an existing analytics or data platform?	Strongly, and it works well; Partly; some parts we would replace; Not at all; we are starting fresh
How much does avoiding lock-in matter?	It is a board-level concern; Important but negotiable; Speed matters more
How strong are your platform engineering skills?	Strong and growing; Moderate; Limited
What matters most in the first year?	Quick value from new use cases; Moving what we have with least disruption; Best-of-breed capability per workload
Do regulation or sovereignty require data in specific places?	Yes, across several jurisdictions; Yes, one jurisdiction; No strong constraints

What each answer points to

Provider-native

Build on one provider's own data and analytics services end to end. Fastest route to value for new data products
Negotiate exit terms and keep data in open formats to limit lock-in

Mixed native and independent

Use the provider's core services and replace weaker ones with independent tools. Choose a small set of independent tools with clear integration standards
Budget integration effort in the first wave

Self-contained independent software

Move an existing analytics or data platform to the cloud largely as it is. Quickest lift and shift with least retraining
Plan how to reach surrounding provider services over time

Intercloud

Deliberately spread data and workloads across providers. Make governance and cost control cross-cloud from the start
Use federation to query in place and limit transfer fees

SELF-CHECK

Is this in place today?

Score each statement 1 (not at all) to 4 (fully, evidenced).

#	Statement	Score
1	Our cloud moves are sequenced by the data strategy roadmap, not by infrastructure renewal dates.	
2	We have mapped data flows and know which systems must move together.	
3	We chose our cloud ecosystem style deliberately, with exit terms understood.	
4	Every cloud workload has a budget, an owner and spend alerts.	
5	We compare services on price-performance from proofs of concept on real workloads.	

How it looks in six industries

Industry	Situation	The move	Measure
Manufacturer	Plant systems stay on site; analytics on field and quality data needs scale.	A spanning architecture: plant data stays local, quality and installed-base analytics move to the cloud.	Cost per quality-analytics run against the on-site baseline
Retail bank	Regulators expect control over where risk and customer data sit.	Use-case specific moves for analytics, with risk data aggregation kept under one governed design.	Regulatory data held in approved locations, evidenced
Retailer	Seasonal peaks swamp fixed capacity for pricing and stock analytics.	Consumption pricing with spending limits for peak analytics; reserved capacity for daily dashboards.	Analytics spend per trading week against budget
Hospital group	A privacy incident makes patient data placement a board concern.	Move non-identifiable planning and patient-flow analytics first; keep clinical records on site until controls are proven.	Workloads moved with privacy review passed
Telecom operator	Separate billing and CRM stacks are tightly entangled with network data.	Map data gravity first and move billing, CRM and churn analytics as one wave.	Cross-environment transfer cost per month
Insurer	Pricing models need burst compute; reporting runs on steady schedules.	Burst pricing for pricing-model runs, reserved capacity for actuarial and regulatory reporting.	Price-performance of pricing model runs

Measures for the data scorecard

Measure	Definition	Lead or lag
Spend against budget	Monthly cloud data and analytics spend compared with plan, by owner	Lag
Price-performance	Cost per unit of useful work for each major workload, against baseline	Lag
Tagged spend	Share of cloud spend tagged to an owner and a data product	Lead
Workloads moved by plan	Migration waves completed on the sequenced roadmap	Lead
Cross-environment transfer cost	Monthly cost of moving data between clouds and on site	Lag
Time to provision	Days from approved request to a working environment	Lead
Idle resource share	Share of provisioned capacity left unused	Lead

3

Making it happen

Traps, the first 90 days and where SCIKIQ fits

Traps to avoid

Lift and shift everything

Workloads move as they are and cost more than they did on site. Counter-move: Re-plan each wave for cloud pricing and switch off what runs idle.

Breaking a cluster

One system moves and its entangled neighbours start paying latency and transfer fees. Counter-move: Map data gravity and move tightly coupled systems in the same wave.

Governance after the bill

Spend control starts only after the first budget breach. Counter-move: Set budgets, tags, alerts and owners before the first workload goes live.

Multicloud by accident

Teams pick providers independently and the estate fragments. Counter-move: Choose the ecosystem style centrally and approve exceptions explicitly.

Speed as the only test

Services are chosen on benchmark speed and cost overruns follow. Counter-move: Choose on price-performance from proofs of concept on real workloads.

The first 90 days

Days 1-30: Plan

- Map data flows and entangled clusters
- Agree regulatory, sovereignty and latency constraints
- Set up financial governance: budgets, tags, alerts and owners

Days 31-60: Choose and prove

- Choose the ecosystem style using the roadmap and constraints
- Run proofs of concept on two real workloads
- Pick pricing models per workload from the fit table

Days 61-90: Move the first wave

- Migrate one low-entanglement, high-value cluster
- Review spend weekly and tune
- Report price-performance against the baseline to the sponsor

Where SCIKIQ fits

Virtual Data Fabric

Connects and federates data wherever it lives, on site or across clouds, so not everything has to move.

Enterprise Sources

Connects in place to enterprise applications, databases and cloud stores, keeping data where it already sits.

Enterprise Discovery Engine

Discovers metadata bottom-up, showing which systems and data are connected before migration waves are set.

Governance Gate

Applies the same access and compliance checks to every request, whichever environment holds the data.

Take it further

- The SCIKIQ Data Strategy Framework: the cascade this track plugs into
- Strategy Cascade and Scorecard Workbook: add this track's objectives and measures
- Run the interactive tool and self-check at scikiq.org/data-strategy#track-cloud

scikiq.org/data-strategy