TECHNOLOGY BUSINESS MANAGEMENT · CASE STUDY

Bringing transparency to Meridian's infrastructure cost base

Attribution, scenario modelling, and an executive view of a $40.2M run-rate.

$40.2M
Q1 2025 infrastructure base

8 pools → 12 apps
attributed to product & program

2 scenarios
parameterised what-if

Case study on a fictional company and synthetic data

1 The cost picture today is opaque

01

Captured by cost pool

The GL records labour, hardware, cloud and vendor spend, not the product or business that actually consumes it.

02

One number, no drivers

Stakeholders see a single fully-loaded figure, but cannot see what drives it or who consumes it.

03

No way to test change

There is no model to quantify what happens if a workload moves to cloud or a vendor raises its rates.

2 Our approach, in four moves

1

Attribute

Push GL cost down to product and application using consumption drivers.

2

Handle overhead

Treat the shared 8% with an explicit, defensible policy.

3

Model scenarios

Parameterised what-if for migration and vendor rates.

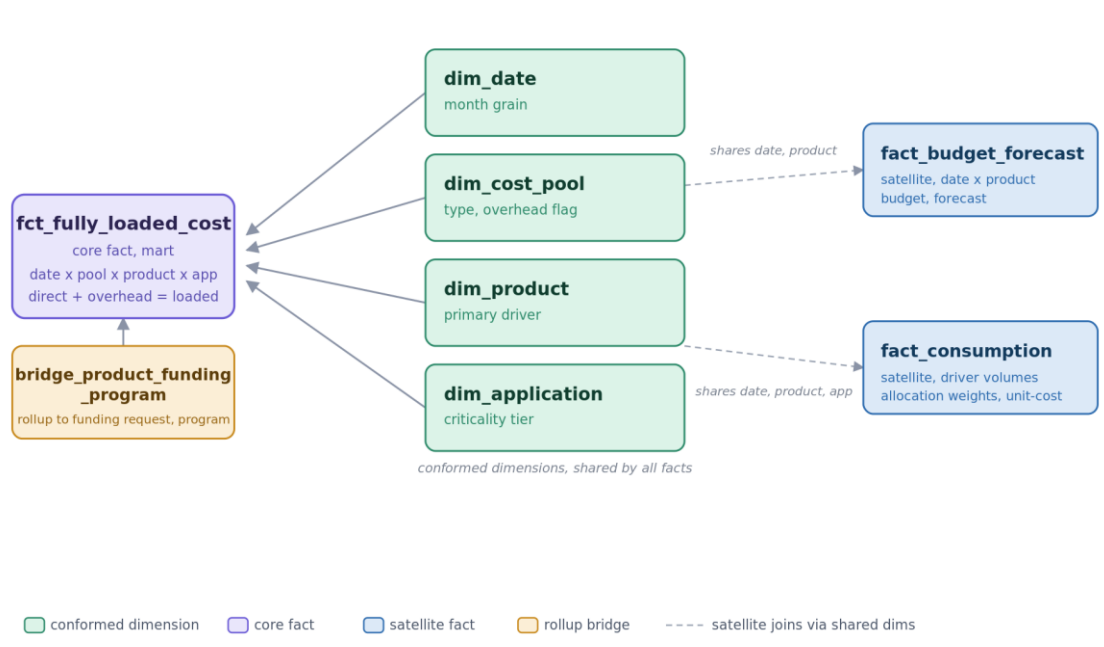
4

Tell the story

An executive view and a clear recommendation.

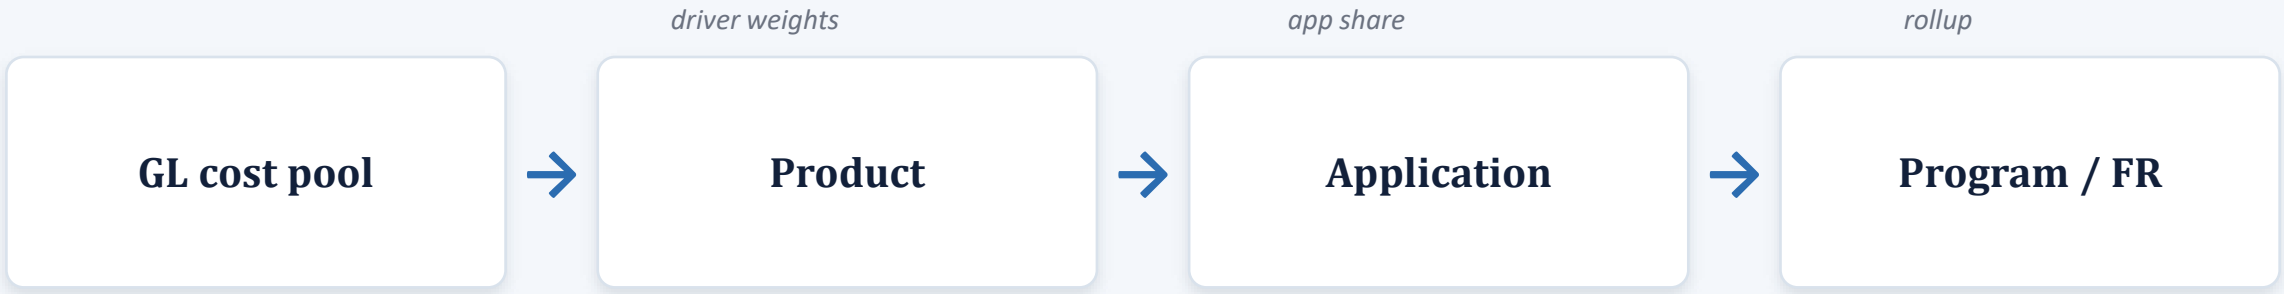
Every figure in this deck is computed in Snowflake and reconciles to the General Ledger.

3 A model built for attribution



- Grain: month · cost pool · product · application
- Cost pool is kept in the grain so scenarios can target one pool
- Overhead carried as its own column, never hidden in product cost
- Program and funding request reached via a product bridge
- Reconciles to the GL within rounding

4 How cost flows: two hops



Judgment call Mixed-driver pools (Labour, Data Centre) split equally across each meter, then by volume within it.

Data resilience A missing monthly usage feed is carried forward from the last known month, so no cost leaks.

5 Worked example: $4.37M of January Labour

Product	Basis	Attributed
Compute	CPU, 66% of meter group	$1.45M
Database	CPU, 34% of meter group	$0.73M
End-User	Headcount, 56% of group	$1.23M
Security	Headcount, 44% of group	$0.96M

1 GL row

lands on 7 applications

The $4.37M splits 50/50: half to the CPU meter (Compute, Database), half to Headcount (End-User, Security). Within each half, a product takes its share of the driver quantity.

This single row is why End-User and Security carry over a million dollars each in January.

6 Worked example: Labour split to 7 applications

Application	Basis (share of meter group)	Attributed
Trade Capture	Compute · 28.2% of CPU group	$615,906
Risk Engine	Compute · 20.8% of CPU group	$454,597
Market Data Hub	Database · 19.5% of CPU group	$425,268
Payments Gateway	Compute · 17.5% of CPU group	$381,275
Settlement Platform	Database · 14.1% of CPU group	$307,953
Collaboration Suite	End-User · 56.2% of HC group	$1,229,062
Identity & Access	Security · 43.8% of HC group	$955,938

The second hop: product to application

Each product's cost from the prior slide now splits to its applications by that application's share of the same driver. Compute's $1.45M spreads across Trade Capture, Risk Engine and Payments Gateway by their vCPU-hours.

7 Shared overhead: the credible call

$3.2M

8% of the base

no driver, no natural home

Policy: spread pro-rata to each application's cost, but report it as a separate line.

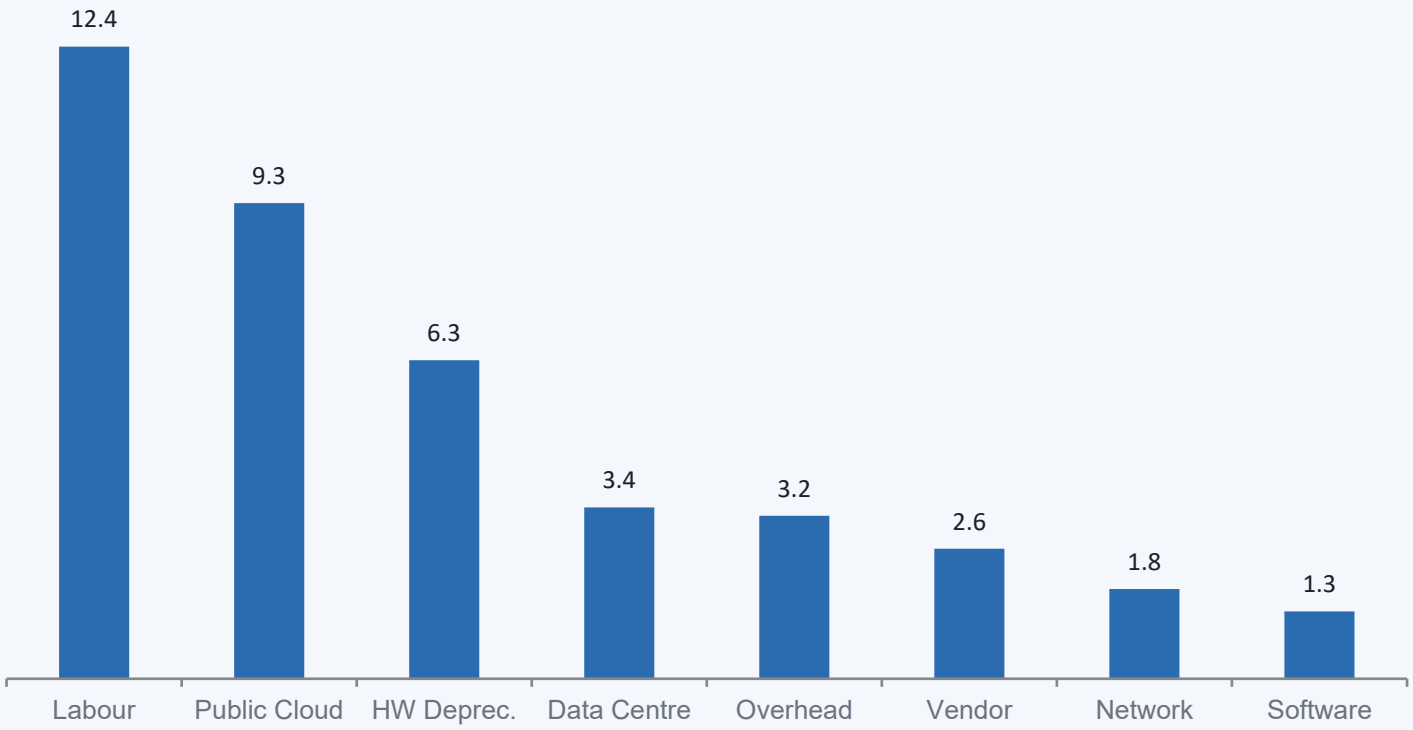
- Why pro-rata: shared overhead scales with the size of the operation it supports, and cost is the best proxy when there is no meter.
- Why a separate column: the assumption stays visible and can be toggled off; nothing is buried in product cost.
- Trade-off: pro-rata makes large products absorb more. Headcount or an even split are defensible alternatives to test.

8 Overhead in practice: January distribution

Each application's share of the month's $16.62M total cost, times the $1.59M pool, is its overhead. Bigger cost base, bigger slice; the pool reconciles exactly.

Application	Direct + allocated	Share	Overhead	Fully-loaded
Trade Capture	$3,098,782	18.6%	$296,454	$3,395,236
Risk Engine	$2,287,196	13.8%	$218,811	$2,506,008
Payments Gateway	$1,918,294	11.5%	$183,519	$2,101,813
Market Data Hub	$1,631,109	9.8%	$156,045	$1,787,154
Collaboration Suite	$1,229,062	7.4%	$117,582	$1,346,644
Settlement Platform	$1,181,148	7.1%	$112,998	$1,294,146
Data Lake	$1,128,572	6.8%	$107,968	$1,236,540
Network Backbone	$1,090,000	6.6%	$104,278	$1,194,278
Identity & Access	$955,938	5.8%	$91,453	$1,047,390
Document Store	$846,428	5.1%	$80,976	$927,404
Client Portal	$683,708	4.1%	$65,409	$749,117
Wealth CRM	$569,756	3.4%	$54,507	$624,263
Total	**$16,619,993**	**100%**	**$1,590,000**	**$18,209,993**

9 Where the money goes



50%

of all cost is Compute, a single product

70%

of the base is Labour, Public Cloud and HW depreciation combined

10 Actual vs budget: watch composition, not the total

- Compute is under budget every month, and the gap is widening (-10% to -21%).
- Database shows the same pattern, worsening to -40%.
- In total, actual lands within about 1% of budget.
- So the total looks fine while the composition is well off plan.

Read this one carefully

End-User and Security appear +5,000% to +9,500% over budget.

That is not overspend. It is how shared labour is attributed to headcount-heavy products against token budgets.

Action: reset those budgets, do not cut those teams.

11 Actual vs Budget vs Forecast, by product (Q1)

Product	Actual	Budget	Forecast	vs Budget	vs Forecast
Compute	$18.49M	$22.10M	$19.98M	-16.3%	-7.4%
Database	$4.42M	$6.10M	$6.21M	-27.5%	-28.8%
End-User	$3.49M	$0.04M	$0.04M	+8,019%	+8,657%
Cloud Platform	$3.14M	$2.90M	$2.82M	+8.2%	+11.3%
Storage	$2.96M	$3.50M	$2.60M	-15.4%	+13.9%
Security	$2.71M	$0.01M	$0.01M	+26,989%	+23,979%
Network	$1.76M	$2.10M	$2.80M	-16.3%	-37.2%
Total	**$36.97M**	**$36.75M**	**$34.45M**	**+0.6%**	**+7.3%**

Real underspend: Compute and Database are below both budget and forecast, and the gap widens through the quarter (Compute -10% to -21%). Delivery is tracking under plan.

Not overspend: End-User and Security dwarf their budgets because shared labour lands on headcount-heavy products given token budgets. Reset the budgets, do not cut the teams.

Budget vs forecast: Forecast ($34.5M) was set below budget ($36.8M); actual ($37.0M) lands about on budget but +7% over forecast, the same attribution effect.

SCENARIO SETUP

Scenario design & assumptions

One parameter table (SCENARIO_PARAM): change one row, re-run. Baseline: fully-loaded Q1 2025 = $40.17M.

Scenario 1: Cloud migration

+$50k to +$555k · +0.13% to +1.38%

LEVER
Move 15% of Data Centre (CP04) workload to public cloud.

MECHANISM
Each CP04 line splits: 85% stays CP04, 15% re-tagged as CP05, same product and application, repriced by the factor.

CLOUD COST FACTOR
1.10 assumed: cloud 10% dearer than owned DC. The data has no cloud rate card, so this is assumed, not derived.

DC COST (FIXED)
CP04 is a Fixed pool. Removing 15% of workload does not remove 15% of DC cost; only the avoidable share is saved.

TAKEAWAY
Avoidable DC: +$50k. Fully-fixed DC: +$555k. Migration pays only if the cloud rate beats the DC cost you can actually shed.

Scenario 2: Vendor rate rise

+$204,079 · +0.51%

LEVER
+8% rate rise on Vendor / Managed Service (CP06).

MECHANISM
Flat +8% on CP06 lines' direct+allocated cost; no pool change; overhead then re-spread.

ASSUMPTION
Full pass-through: no volume or scope change, no renegotiation offset.

LANDS ON
Managed-service consumers only: Database and Cloud Platform.

TAKEAWAY
Larger than S1 and certain in direction. Arrives unless the vendor is pre-empted commercially.

Overhead is re-spread onto each scenario's new base. Pool total unchanged, so deltas are pure economic effect.

12 Scenario 1 — migrate 15% of Data Centre to cloud

+$50,474 fully-loaded · +0.13% of base. Lands on the Data Centre consumers (Storage and Compute); concentrates in Data Lake.

By product

Product	S1 delta
Storage	+$26,991
Compute	+$25,440
Database	-$688
Cloud Platform	-$359
End-User	-$347
Network	-$293
Security	-$270
Total	**+$50,474**

By application

Application	S1 delta
Data Lake	+$15,468
Document Store	+$11,524
Trade Capture	+$10,801
Risk Engine	+$7,955
Payments Gateway	+$6,684
Wealth CRM	-$162
Client Portal	-$196
Identity & Access	-$270
Settlement Platform	-$289
Network Backbone	-$293
Collaboration Suite	-$347
Market Data Hub	-$400
Total	**+$50,474**

13 Scenario 2 — vendor rates rise 8%

+$204,079 fully-loaded · +0.51% of base. Lands only on the managed-service consumers: Database and Cloud Platform.

By product

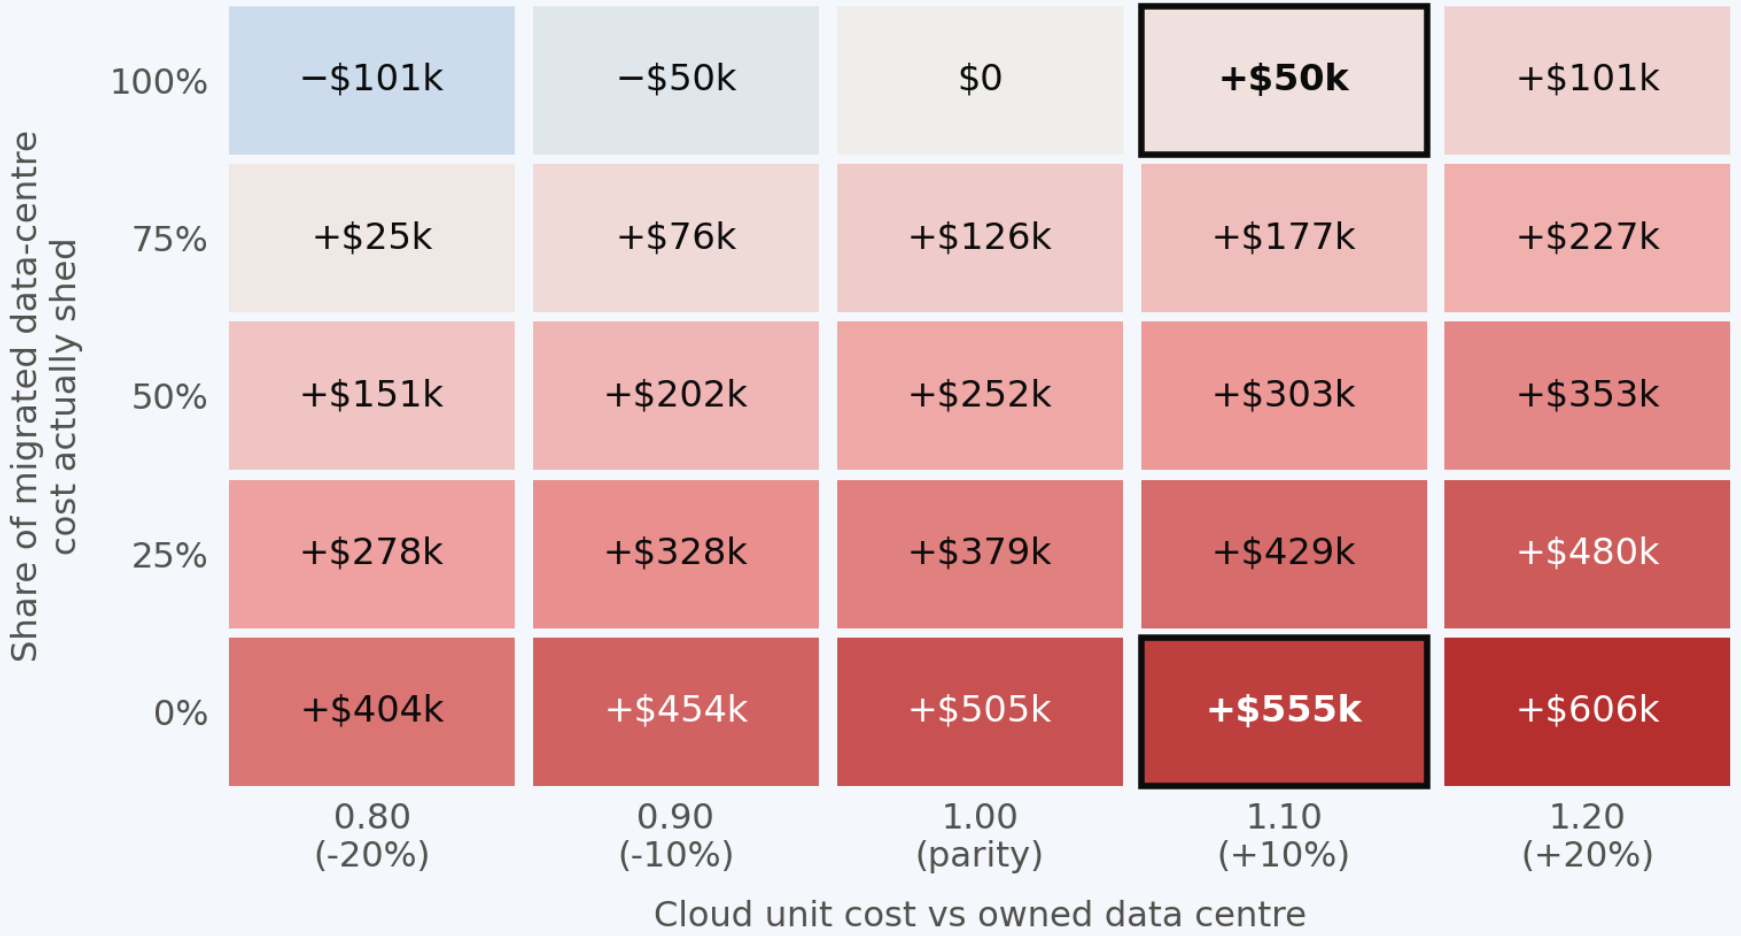
Product	S2 delta
Database	+$125,320
Cloud Platform	+$92,945
Network	-$808
Security	-$1,296
Storage	-$1,586
End-User	-$1,669
Compute	-$8,827
Total	**+$204,079**

By application

Application	S2 delta
Market Data Hub	+$73,092
Settlement Platform	+$52,229
Client Portal	+$51,169
Wealth CRM	+$41,776
Document Store	-$676
Network Backbone	-$808
Data Lake	-$910
Identity & Access	-$1,296
Collaboration Suite	-$1,669
Payments Gateway	-$2,323
Risk Engine	-$2,760
Trade Capture	-$3,744
Total	**+$204,079**

12b Scenario 1 sensitivity: it hinges on stranded cost

Delta = DC cost x 15% x (cloud factor - share of DC cost shed). It only saves money if cloud is cheaper and the data-centre cost really goes away.



Recommendation

Cloud migration is a bet, not a saving. Its impact swings from a +$50k cost to about +$555k, driven by two unknowns: the true cloud unit rate, and how much fixed Data Centre cost you can actually shed. CP04 is a fixed pool, so little sheds in the short term, pushing the likely outcome toward the high end. Prove both numbers before committing a dollar.

Reforecast & reset budgets

Release the under-spending Compute and Database budgets. Reset the token End-User and Security budgets: their extreme variances are an allocation artifact, not overspend. Keep the teams.

Pre-empt the vendor

Open the rate conversation now to blunt the certain +$204k increase.

Focus management attention on Compute: at half the cost base, it is where transparency creates the most value.

14 Assumptions I'd defend, and what's next

Key assumptions

- Equal weight per driver group for mixed-driver pools
- Carry-forward for missing monthly usage feeds
- Overhead spread pro-rata, shown as a separate line
- Variance compared on a pre-overhead basis

With more time

- Sensitivity-test the labour split under alternative bases
- Model product-to-funding-request as a weighted many-to-many
- Extend attribution to region and legal entity
- Add more periods to firm up the trend statements
- Sensitivity-test how fixed the Data Centre pool is, to narrow the Scenario 1 range
- Pin down the cloud cost factor from real cloud rates, replacing the assumed 1.10

Thank you

Questions and discussion welcome.