

Measured on a full terabyte of industry-standard data (TPC-DS) on BigQuery, June 2026 — the same answer, raw versus accelerated.

GOVERNED ANALYTICS, AT NEAR-ZERO COST

The same answer, for a fraction of the cost — on a terabyte.

Tessallite answers the questions your business repeats every day — in Excel, Power BI, the API and the agent — from a tiny maintained summary instead of re-reading the whole terabyte. On a full terabyte of industry-standard data, that turns hundreds of dollars per thousand runs into a few cents, with the identical answer.

COST PER 1,000 RUNS

\$330 → 5¢

A question that costs hundreds of dollars to repeat runs for pennies.

DATA SCANNED PER QUESTION

66 GB → a few KB

Read from a maintained summary, not the whole terabyte.

ANSWERS IDENTICAL

4 / 4

Every accelerated answer matched the raw answer, exactly.

~\$116,000 a year → ~\$18 a year

A team running 40 dashboards refreshed hourly asks these questions roughly 350,000 times a year. Today that is about \$116,000 in BigQuery scan; with Tessallite it is about \$18 — the same answers, delivered to every tool. (BigQuery on-demand at ~\$5/TB; per-run figures measured below.)

The proof — four everyday business questions SF1000 MEASURED

We asked each question through the live connection the way Excel or Power BI would — first the normal way (raw BigQuery), then with Tessallite — and the two answers came back identical. "Cost per 1,000 runs" uses BigQuery's on-demand price (~\$5 per terabyte scanned, 10 MB minimum per query).

Business question	Answered from	Data scanned today	Data scanned	Cost per 1,000 runs	Same answer?
-------------------	---------------	--------------------	--------------	---------------------	--------------

		(raw BigQuery)	with Tessallite	(today → with Tessallite)	
Sales revenue by year	maintained summary	66.0 GB	1,480 bytes	\$330 → \$0.05	yes
Sales revenue by product category	maintained summary	67.1 GB	255 bytes	\$335 → \$0.05	yes
Sales revenue by sales region	maintained summary	66.0 GB	41 bytes	\$330 → \$0.05	yes
Sales revenue by customer type	maintained summary	66.1 GB	3,348 bytes	\$330 → \$0.05	yes


Tens of gigabytes the normal way; a few bytes with Tessallite — the same answer, for a fraction of a cent instead of a few hundred dollars per thousand runs. ("Customer type" is the retailer's preferred-customer flag.)

See it at true scale SF1000 MEASURED

The cloud bill and the data each question moves, on a real linear scale — not compressed. Lower is better.

What it costs to ask the same question 1,000 times


Three everyday questions, from a simple one to a heavy one. Red = today (raw BigQuery, each time); green = with Tessallite.



The green bars are a few cents — too small to see beside hundreds of dollars. That is the point.

Data scanned to answer one question

The same three questions, in gigabytes. Red = raw BigQuery; green = with Tessallite.



The green bars are a few kilobytes to a few megabytes — invisible beside 66–155 GB. Same answer, tiny read.

What it is worth to you

- **A smaller cloud bill, now.** The questions your teams repeat every day read kilobytes instead of tens of gigabytes — hundreds of dollars per thousand runs become a few cents, and the saving recurs on every refresh.
- **One answer, everywhere.** The same governed definition and the same accelerated result in Excel, Power BI, the API and the agent — one source of truth, not five.
- **Nothing to hand-build or babysit.** Tessallite suggests and maintains the acceleration from how the business actually queries; nothing silently breaks when the model changes.
- **The same answer, always.** Every accelerated result is the raw-source result — faster and far cheaper, never a different number.

Your whole question set — 17 questions, all accelerated

SF1000 MEASURED

Not one query — the everyday set a commercial, finance or merchandising team actually asks. For each, the data scanned the normal way versus with Tessallite, and the recurring cost at 1,000 runs.

Business question	Data scanned today (raw BigQuery)	Data scanned with Tessallite	Cost per 1,000 runs (today → with Tessallite)
Sales revenue by region	66.0 GB	41 bytes	\$330 → \$0.05
Sales revenue by quarter	66.0 GB	112 bytes	\$330 → \$0.05

Sales revenue by product category	67.1 GB	255 bytes	\$335 → \$0.05
Sales revenue by month	66.0 GB	304 bytes	\$330 → \$0.05
Sales revenue by state	66.0 GB	436 bytes	\$330 → \$0.05
Sales revenue by region, month by month	88.0 GB	725 bytes	\$440 → \$0.05
Sales revenue by year and month	66.0 GB	2.0 KB	\$330 → \$0.05
Sales revenue by product class	67.1 GB	2.6 KB	\$335 → \$0.05
Sales revenue by product category, month by month	89.1 GB	4.4 KB	\$445 → \$0.05
Sales revenue by customer type, month by month	88.1 GB	6.3 KB	\$440 → \$0.05
Sales revenue by state, month by month	88.0 GB	7.8 KB	\$440 → \$0.05
Sales revenue by product category, by year and month	89.1 GB	26.5 KB	\$445 → \$0.05
Sales revenue by brand	67.1 GB	24.4 KB	\$335 → \$0.05
Sales revenue by product category and brand	67.1 GB	47.7 KB	\$335 → \$0.05
Sales revenue by region and state, month by month	88.0 GB	69.4 KB	\$440 → \$0.05
Sales revenue by product category and brand, by quarter	89.1 GB	1.2 MB	\$445 → \$0.05
Promotion impact: revenue by promotion type and category, month by month	155.1 GB	81.8 MB	\$775 → \$0.41

Two reusable governed lists — favourite product categories and key stores by state — are defined on the model so everyone filters on the same definition. The heaviest question in the set still drops from \$775 to 41 cents per thousand runs.

How Tessallite does it

The capabilities behind these results — and what your teams use day to day.

Capability	What it does	What it did here
------------	--------------	------------------

Governed semantic model	Business measures (revenue, margin, returns), dimensions (year, category, region, customer type) and the table joins are defined <i>once</i> , centrally.	"Revenue by category" means the same thing in Excel, the API and the agent — correct by design, not re-derived per tool.
Aggregate acceleration	Tessallite maintains small summaries of the huge fact table at the grains people actually query, and routes matching questions to them automatically.	The headline win: questions answered from byte- to kilobyte-sized summaries instead of scanning 66–155 GB — same answer, a fraction of the cost.
Transparent routing	For every query Tessallite records the route it chose, the bytes it scanned, and the summary it used.	That is how the numbers in this report were captured — straight from Tessallite's own log and the live run.
Multi-protocol gateway	One governed model served over SQL clients, Excel / Power BI, the REST API and the conversational agent.	The same accelerated answer reached every tool; these numbers came through the real connection a BI user uses.
Governed BigQuery source	A read-only, access-controlled connection to BigQuery; every query flows through the gateway.	BigQuery stays the engine and system of record — Tessallite adds governance and acceleration without copying data out.

How we ran it

The industry's standard test, at full scale. TPC-DS is the benchmark the major data-platform vendors use to prove performance: a complete model of a real retailer — sales, returns, customers, promotions, calendar. We ran it at SF1000: a full terabyte, 2.88 billion sales records, 24 tables, loaded clean.

SCALE

~1 TB

SF1000 — about one terabyte of analytics data.

SALES RECORDS

2.88 billion

2,879,987,999 line items — years of a large retailer's activity.

TABLES

24

Stores, products, promotions, customers, dates — 0 load errors.

Same question, both ways, identical answer. We ran each question on the same model through the same connection — once accelerated, once straight to BigQuery. Identical overhead on both sides, so the gap is pure data work, and the answers matched to the digit on every question. Byte

figures come from BigQuery's own dry-run with the result cache off, so each one is the real cost to scan.

Where the win shows up

At a terabyte, the win is your bill — the number that compounds. BigQuery scans ~60 GB in seconds, so raw and accelerated finish in similar wall-clock; but the accelerated question reads a few kilobytes and costs a few cents, every time it runs. That saving recurs on every dashboard refresh, for every team, all year.

Model `tpcds_retail_sf1000` · measured June 2026 · raw run records: `tpcds-sf1000-ab-postfix.json`, `tpcds-sf1000-bytes-matrix.json`, `tpcds-sf1000-superset-matrix.json`. TPC-DS-derived workload; not an official TPC audited result.