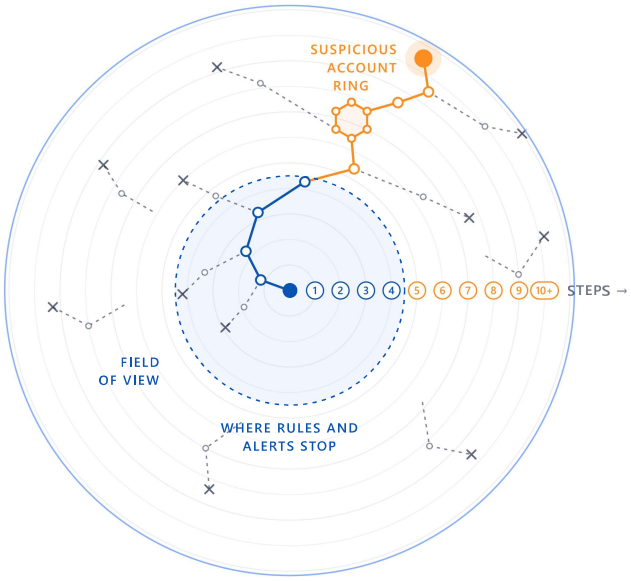
GRAPH TRAVERSAL BENCHMARK

More transaction history. Same hardware.



What happens when an investigation can't follow the evidence far enough? Connections remain hidden, patterns go undetected, and the full picture never emerges. The question isn't just how deep you can go, but how much you can uncover while the question still matters.

0.57 s

12 TRANSFERS DEEP

Follow payments from one bank account, twelve transfers deep, within a defined date window. On a 20 TB transaction graph, DataWalk built and counted 5.37 million connected accounts in 0.57 seconds.

WHY THIS MATTERS

The evidence may be beyond your usual field of view. Graph benchmarks often report how many hops a query can cover in a given time. That tells you little about how long it takes to find a pattern in real data. The pattern may emerge several connections from where you started, and whether you can keep exploring shapes what you'll discover.

THREE CORE BENCHMARK FINDINGS

SUB-SECOND MULTI-HOP SPEED

12 hops deep across 5.37 million connected accounts in 0.57 seconds on 20 TB.

NO MEMORY-WALL EXPANSION

The dataset grew 13× (1.5 TB to 20 TB). Hardware remained identical: 6 servers, about 2.5 TB of total RAM (roughly one-eighth of the dataset size).

WHOLE-NETWORK SWEEP RESILIENCE

Evaluated 37.15 billion relationship checks across all 50 million internal accounts in 61.60 seconds, or about 600 million checks per second.

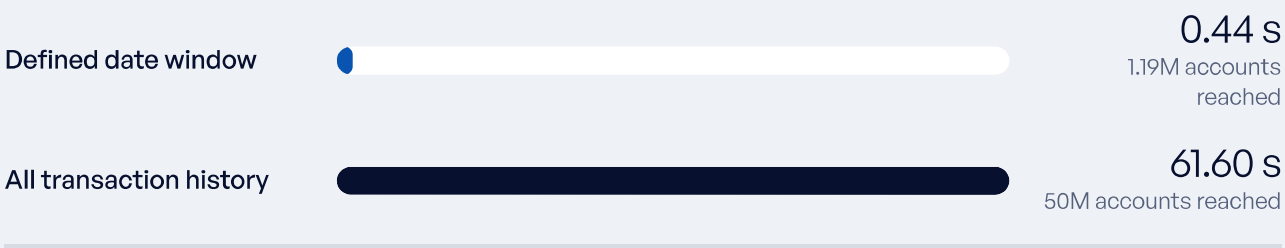
01 / THE CHALLENGE

Depth isn't the measure.

Graph benchmarks usually boast “N hops in X seconds.” But hop count hides the real cost. Two questions can go equally deep while exploring vastly different numbers of connections.

The challenge is handling the explosion of connections at each step and the resulting performance hit, without forcing analysts to restrict their investigation.

SAME 20 TB GRAPH / SEVEN STEPS / BOTH DIRECTIONS



The only change was the date window. The starting account, direction and step count stayed constant, but the number of qualifying transactions and accounts changed. Different scopes produce different answers and very different workloads.

02 / WHAT THE BENCHMARK SHOWS

Three findings. Three practical gains.

These results reveal three different advantages: faster follow-up analysis, control over the work each question demands, and resilience as transaction history grows.

01

Speed where investigators work.

On the 20 TB graph, a directed, date-filtered traversal reached 5.37 million connected accounts in 0.57 seconds at 12 steps, and 45.62 million in 5.34 seconds at 15 steps.

Why it matters: When an unexpected connection appears, the investigator can test another question immediately, instead of making every follow-up a batch job.

02

Depth isn't the measure.

The same seven-step traversal took 0.44 seconds over a defined date window and 61.60 seconds across all history because the work and result population changed.

Why it matters: Analysts choose the direction, dates and populations worth pursuing. Depth alone cannot tell buyers what a benchmark actually processed.

03

Twice the data. Similar answer times.

DataWalk ran the same investigations against 10 TB and 20 TB of transaction history. Both returned the same accounts, with little difference in response times. Even at 15 steps, the results took 4.59 and 5.34 seconds respectively.

Why it matters: The same questions reached the same accounts in roughly the same time, even with twice the transaction history.

CONTROLLED COMPARISON / SAME SELECTED WORK

DIRECTED / SELECTED DATES	10 TB	20 TB	ACCOUNTS
11 steps	0.37 s	**0.31 s**	1.59M
13 steps	1.23 s	**1.18 s**	15.94M
15 steps	4.59 s	**5.34 s**	45.62M

WHAT HAPPENS WHEN YOU WIDEN THE SEARCH?

The whole network, in about a minute.

Starting from one account and following connections in both directions across the full transaction history, DataWalk reached all 50 million internal accounts. The sweep evaluated 37.15 billion relationship rows in 61.60 seconds, about 600 million checks per second.

50M	37.15B	61.60 s
internal accounts reached	relationship checks	full-history, bidirectional sweep

These are relationship checks, not unique stored relationships: the sweep evaluates connections from both directions. This is the other side of the performance story: focused questions can return in fractions of a second; sweeping every connection is a much larger task that DataWalk also completed.

03 / WHY DATAWALK WORKS DIFFERENTLY

Built for the next question.

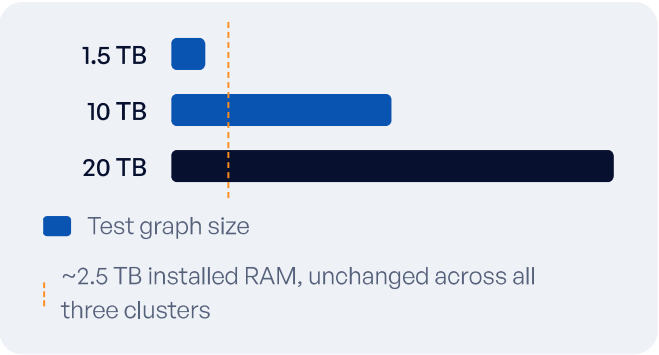
The benchmark results demonstrate the strengths of DataWalk's underlying architecture: relationships prepared during ingestion, exploration across entire sets of connected records, and compact indexes rather than a requirement to keep the whole graph in RAM. The result is an architecture suited to questions that change as new evidence appears.

	THE ALTERNATIVE	DATAWALK	THE GAIN
Relationships ready at runtime	Query-time joins can turn each follow-up question into more reconstruction work.	Binds ingested data to a business ontology and creates defined relationships as it loads.	**Analysts can follow another lead without rebuilding connections.**
Populations, not just paths	Path-by-path approaches can make it difficult to carry an investigation forward.	Traverses sets, carrying selections, filters and analytical context between steps and views.	**Change direction or narrow the population without starting over.**
Large graphs, smaller memory dependence	Memory-dependent engines may require RAM to expand as the graph grows.	Uses compact indexes, set operations and graph data accessed from local storage.	**The test graph grew from 1.5 TB to 20 TB with no increase in installed RAM.**

04 / WHY THE FOOTPRINT MATTERS

The graph outgrew memory. The queries still ran.

All three test sizes used separately deployed, identically specified six-server clusters, each with approximately 2.25 TiB (2.5 TB) of installed RAM. The graph grew from 1.5 TB to 10 TB and 20 TB without an increase in provisioned memory. Every measured request completed within a 12 GB per-query allowance (not a measured peak-use figure).



HOW TO READ THE EVIDENCE

DataWalk-run synthetic banking model: 50 million traversable internal accounts; 1.5, 10 and 20 TB; same DataWalk build, server specification, storage class and workload settings across sizes. One simulated user on warmed node-local storage. Ten runs per request; server-side median for building and counting connected account sets, excluding record retrieval and screen rendering.

DataWalk ran the benchmark with no benchmark-specific tuning and the same configuration at every size. The published results cover a single user on warmed storage. The benchmark did not include an independent audit or a comparison with other products.

Examine the full technical benchmark.

[Full technical report and method: datawalk.com/graph-traversal-benchmark/ →]