thatDot Streaming Graph

Answer the tough, deep questions of graph analysis immediately.

Stay out of the papers with real-time cybersecurity and fraud prevention that reduces loss windows to near zero.

WHAT IT'S FOR

Cybersecurity

Tough challenges like Advanced Persistent Threats (APT) and insider threats can attack low and slow with no time window limitations. Streaming Graph was developed by DARPA and funded by Crowdstrike to do the powerful graph analysis that catches threats while eliminating the time window limitations that hamstring other streaming data processors.

Financial

Spot patterns from POS devices, ATMs, etc. Turn fraud detection into prevention. Adjust risk exposure in seconds, not hours or days by switching from overnight batch to real-time.

And a Lot MORE...

Entity resolution, digital twins, AIOps, data pipelines, observability...

Combine the real-time speed of an event stream processor such as Flink, with the power of graph analyses such as in Neo4J, and you get ...
thatDot Streaming Graph.

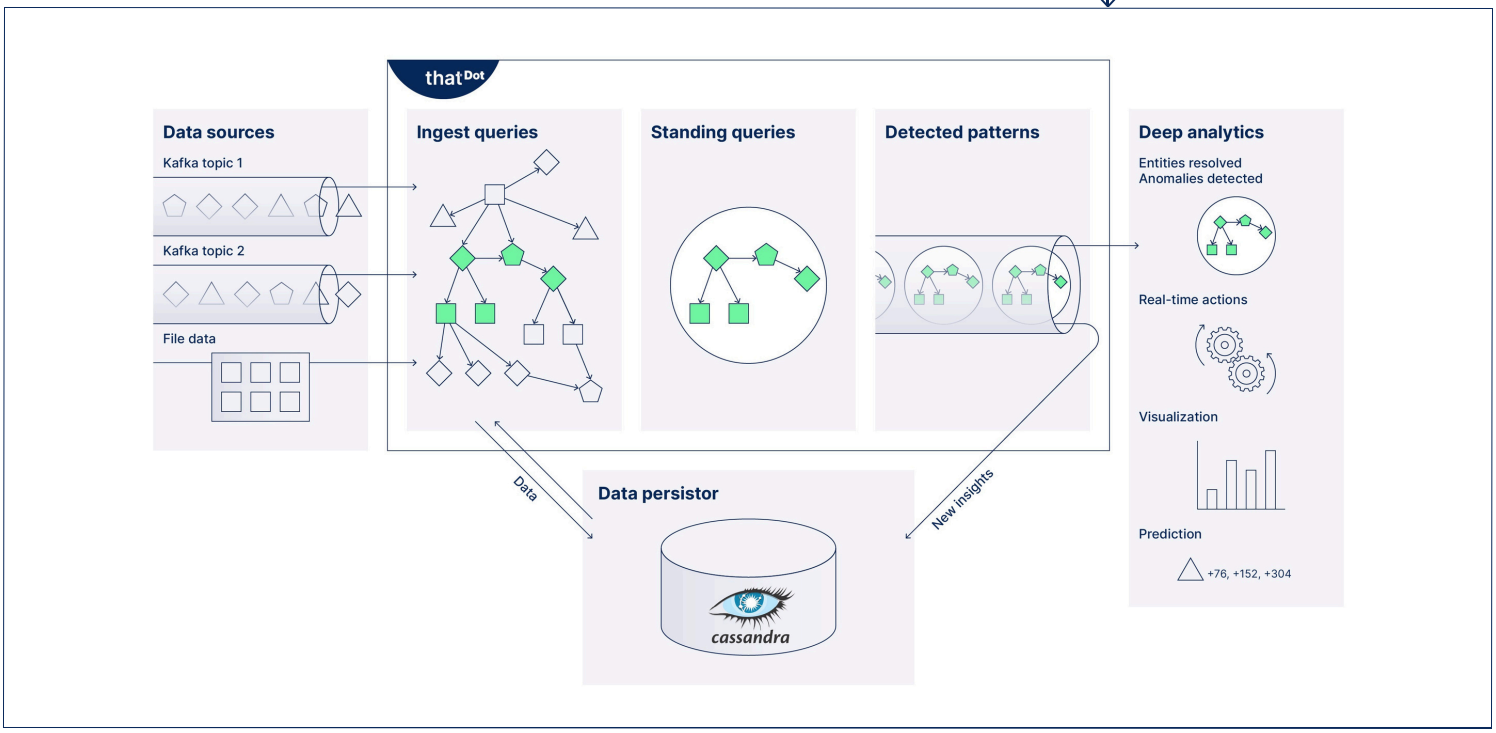

Real-Time Deep Analysis at Scale

ADVANTAGES

Mountains of Data are No Problem

Supernodes, entities with many relationships, have long been the weak point of graph analysis at scale. Streaming Graph has been tested with over 1 mill events/sec with 20,000 standing query responses, even on supernodes with 100's of thousands of edges. No scaling upper limit has been found.

Integrate and Develop Quickly

Use Streaming Graph as both a source and sink for Apache Kafka, AWS Kinesis or SQS, Apache Pulsar, etc. Drop it right into your existing stack and develop with normal developer skills, robust APIs and the standard Cypher query language.

React Fast Enough to Prevent Loss

Graph databases are powerful but slow. Mean time to answer (MTTA) for Streaming Graph can be measured in milliseconds, not hours or days, reducing loss windows to near zero.

HOW IT WORKS

Unify Your Data As It Flows In

Ingest queries turn multiple streams of incoming data into a dynamic graph with no scale limitations. Rather than waiting hours while important categorical data like user and entity behavior (IP addresses, people, accounts, file paths, etc.) is changed to numbers and summaries in a database, pull in real data immediately from multiple data streams, plus batch or file data for context. Find and resolve duplicates, discover relationships, and build a digital twin that genuinely represents your current systems and changes as they do.

Define the Pattern You're Looking For

Standing queries define the patterns you need to find. Use the standard graph analysis Cypher query language. Standing queries don't depend on snapshots or time windows. Low and slow attack patterns no longer defeat the software. Historical data combines with brand new data to match the pattern you started looking for last week.

Drive Real-Time Workflows

Robust APIs make Streaming Graph ideal for embedding in applications, workflows, or data pipelines. Interactively explore or monitor the graph of your systems as it is revealed. Each time a bit of data completes the pattern, the results are immediately sent into a new Kafka pipeline, or wherever you need them. Trigger actions, alert subject matter experts, or push key information to monitoring software. And within milliseconds, you've caught a bad actor and can act to stop data exfiltration, financial fraud, or other damage.

Persist Key Findings and Data

To keep data long term, plug in one of many common databases, Cassandra, Clickhouse, or any Cassandra compatible database like RocksDB. Keep your new insights, and also keep the original data to regenerate the graph in case of power outages, for audit purposes, or to train graph machine learning models such as graph neural networks (GNN).