GARTNER DATA & ANALYTICS SUMMIT 2026

Query to Conviction

Ensuring Trust and Verification in AI-Driven Analytics

David Dixon & **Scott Cairncross**

Enterprise Transformation @ TextQL

THE PROBLEM

The Trust Gap in AI Analytics

"80%

of data and analytics governance initiatives will fail by 2027, due to a lack of a real or manufactured crisis."

Gartner Press Release, 28 Feb 2024

Speed Without Confidence

AI delivers answers fast, but leaders hesitate without understanding how results were derived.

Regulatory Pressure

SOX, HIPAA, GDPR demand auditable, explainable analytics processes.

Operational Risk

Unverified insights reaching production cascade into compliance failures.

ROOT CAUSES

01

The Black Box Problem

AI generates answers without showing reasoning. Stakeholders see outputs but not logic.

02

Metric Inconsistency

Different teams define metrics differently. Revenue calculated three ways destroys confidence.

03

No Audit Trail

Without logging who asked what, compliance teams cannot validate or reproduce results.

04

Ungoverned Access

Without role-based controls, sensitive data leaks and regulatory lines are crossed.

FRAMEWORK

The Trust Architecture

Four pillars that transform AI analytics from a black box into a system leaders and regulators can stand behind.

Transparency

Full visibility into queries, reasoning, and data sources. Users see how every insight is produced.

Governance

Role-based access and semantic layers ensure consistent definitions and proper data boundaries.

Accountability

Complete audit trails log every interaction -- an immutable compliance record.

Verification

Human-in-the-loop workflows and version-controlled logic validate AI outputs before action.

PILLAR ONE

Transparency

AI analytics should never be a black box. Every answer must come with a visible chain of evidence -- the SQL executed, the tables joined, and the logic applied.

Query Visibility

Every SQL query shown before and after execution. Natural language reasoning displayed alongside technical output.

Full Data Lineage

Trace any number back to its source tables, columns, joins, and transformations. End-to-end from result to raw data.

Reasoning Chains

Step-by-step breakdown of how the AI interpreted your question, chose its approach, and arrived at the answer.

PILLAR TWO

Governance & the Semantic Layer

Consistent definitions. Controlled access. One version of the truth.

"3.4x

more likely to achieve effectiveness at AI governance practices"

with AI governance platforms (AIGPs)

Gartner, "Build AI Governance Programs That Keep Pace With Regulatory Change," 29 January 2026

Semantic Layer

Centralized metric definitions ensure "revenue" means the same thing for every team, every query, every time.

Role-Based Access

Row-level and column-level permissions enforce data boundaries. Sensitive data stays within authorized scopes.

Version-Controlled Logic

Business rules stored in version control with full change attribution. Every modification tracked and reversible.

PILLAR THREE

Accountability

Complete Audit Trails

Every query, every user, every timestamp -- logged immutably. Regulators can reproduce any historical analysis.

Interaction Logging

Who asked what, when, and what data was accessed. Full chain of custody for every analytical output.

Compliance-Ready

SOX, HIPAA, GDPR audit requirements met automatically. No manual documentation overhead.

PILLAR FOUR

Verification

Cross-Validation

AI verifies its own answers using different methods. Discrepancies are surfaced and explained, never hidden.

Human-in-the-Loop

Critical outputs reviewed by domain experts before action. Humans maintain meaningful control over AI decisions.

Honest Uncertainty

When data does not exist, the system says so. Every assumption documented with caveats and confidence levels.

LIVE DEMONSTRATION

The Glass Box Test

We will show you the query, cite its sources, document its assumptions, and create reproducibility.

Transparency

Can you see the SQL?
The reasoning?
The data sources?

Cross-Validation

Does the system verify its own answers using different methods?

Reproducible

Will it provide explanation and files so that you can independently produce the same answer?

Audit Trail

Are assumptions logged?
Can you reproduce every result?

Live Demo

The Glass Box Test: Prompts that probe transparency, lineage, cross-validation, and provide certainty

WHAT YOU WILL SEE

The Glass Box in Practice

Every query visible

SQL shown before execution. Natural language reasoning alongside technical output. No hidden logic.

Full data lineage

Source tables, columns, joins, and transformations -- traceable from result to source on demand.

Self-verification

Cross-validation using different methods. Discrepancies surfaced and explained, not hidden.

Customers In Production (updated as of 3/12/26): 542k+ tool calls | 59.1% cross-validation rate | Median 12 verification steps per conversation

YOUR EVALUATION FRAMEWORK

The Trust Checklist

Seven questions to ask any AI analytics vendor.

- ☐ Can users see the exact query and logic behind every AI-generated insight?
- ☐ Is every data access logged with user identity, timestamp, and context?
- ☐ Are metric definitions centrally governed with version history?
- ☐ Do role-based controls restrict access at object type levels?
- ☐ Can a 3rd party reproduce any analysis from the audit trail?
- ☐ Do humans review and approve AI outputs before they reach production?
- ☐ Is business logic stored in version control with change attribution?

YOUR EVALUATION FRAMEWORK

The Trust Checklist

Six prompts and 1 configuration mapping to each Trust Checklist question.

- [ ] What's driving total invoice revenue trends over the past 3 years? Show it to me in an interactive dashboard. Separately, I want to see the complete data lineage with an interactive visualization.
- [ ] Generate an audit report for this session that is comprehensive showing all queries executed, tables accessed, and data retrieved. Include user identity, timestamps, and context. Separately, show a visual timeline.
- [ ] Check the context repository for the version history of the 'Total Invoice Revenue' metric definition. Has the calculation changed over time? Who approved the changes? Separately, show a visual timeline.
- [ ] Do role-based controls restrict access at object type levels?
- [ ] Model the revenue impact of a 10% price increase across all products. Show me the methodology, assumptions, and export this as a reproducible package.
- [ ] Build me an approval workflow dashboard for the 10% price increase model...
- [ ] Show me the change history your context memory of our Revenue Recognition Policy. What changed, who changed it, and why? The same for the business rules behind Customer Segmentation. Separately, create an interactive visual timeline.

KEY TAKEAWAYS

Building Trust-First AI Analytics

Transparency is Non-Negotiable

Every AI output must be inspectable. If stakeholders cannot see the reasoning, they will not trust the result.

Governance Enables Speed

Proper controls eliminate the reconciliation work that actually creates bottlenecks.

Audit Trails Are Competitive Advantage

Comprehensive logging builds institutional knowledge and accelerates future analysis.

Gartner: *"Organizations using AI governance platforms (AIGPs) are 3.4 times more likely to achieve effectiveness at AI governance practices."*

Gartner, "Build AI Governance Programs That Keep Pace With Regulatory Change," 29 January 2026. GARTNER is a trademark of Gartner, Inc. and its affiliates.

Questions & Discussion

David Dixon & **Scott Cairncross**

TextQL | textql.com