AI/ML CONFERENCE 2026

Trustworthy AI for Regulatory Data

Why Governance Isn't Optional

AI+IM GLOBAL SUMMIT

CASE STUDY

The $2.3M Question: When AI Gets It Wrong

The Scenario

A retailer's AI orders **4,000 winter coats**

The Algorithm Saw

"Increased searches" for winter gear but missed critical context: hurricane relief donations

The Core Problem

Data without context is meaningless. The AI saw search spikes but lacked semantic understanding of why people searched for winter coats in July.

Financial Services Stakes

- **Regulatory sanctions** and fines
- **Market volatility** and losses
- **Reputational damage** lost credibility
- **Customer remediation** costs

THE COST

$2.3M

excess inventory

> In financial services, these failures don't just cost money—they trigger regulatory action, market losses, or worse .

GLOBAL SUMMIT

THE CHALLENGE

Why Trustworthy AI Is Hard

Complex Data

Regulatory data includes transactions, counterparty details, and risk calculations across multiple systems and formats.

REAL EXAMPLE

A credit risk model pulling from **47 different data sources** across 6 systems.

Fragmented Ownership

Data ownership spread across IT, Risk, Compliance, and Business units with no unified view.

THE FAILURE

One field mislabeled. A single data mapping error in a complex web of dependencies.

Machine Speed

AI amplifies failures at scale—small errors become enterprise-wide problems in seconds.

THE COST

$15M

in miscalculated exposure limits

In complex enterprise environments, trust isn't a feature —it's an architectural requirement . Governance must be engineered, not assumed. You can't delegate this to data scientists and hope for the best.

AI+IM GLOBAL SUMMIT

REGULATORY REQUIREMENTS

The Four Pillars of AI Accountability

What regulators demand—and why "we don't know" isn't acceptable

01 Traceability

Full audit trail from source data to AI output—every step documented and reversible.

REGULATOR ASKS:

"Show us every data point that led to this credit decision"

02 Explainability

Clear logic for AI decisions—no black boxes. Humans must understand why decisions were made.

REGULATOR ASKS:

"Why did the model flag this transaction as suspicious?"

03 Reproducibility

Same inputs must always equal same outputs. Results must be consistent and verifiable over time.

REGULATOR ASKS:

"Re-run last quarter's risk report and prove it matches"

04 Accountability

Clear ownership when failures occur. Someone must be responsible and consequences must be defined.

REGULATOR ASKS:

"Who approved this model for production use?"

In regulated industries, "we don't know" isn't acceptable.

TECHNICAL FRAMEWORK

Enterprise Architecture for Trustworthy AI

The system that enforces trust through engineered controls

Governed Ingestion Layer

All data enters through controlled checkpoints with automated validation before entry. Rejected data is logged and flagged for investigation.

- Accuracy
- Completeness
- Timeliness
- Reconciliation

Centralized Lineage Tracking

Every transformation is documented with version control for all changes, creating an unbroken chain of custody from source to output.

- Documentation
- Version Control
- Chain of Custody

Controlled Analytics Environment

Role-based access controls with comprehensive query logging and mandatory output review before distribution.

- Access Control
- Query Logging
- Review Process

Data Flow

Standardized quality controls ensure data integrity throughout the pipeline

Validation

Comprehensive validation including accuracy, completeness, timeliness, and reconciliation

Trusted AI Output

Ready for regulatory use

CASE STUDY

Architecture in Action: From Chaos to Control

How it looks in Real World

The Situation

Large financial institution, 200+ AI models, regulatory exam pending.

73 models
Couldn't explain data sources

45 models
Used deprecated data feeds

12 models
Conflicting prod/test versions

The 18-Month Transformation

Phase 1 (M 1-6): Discovery

- Created central model registry
- Mapped all data lineage

Phase 2 (M 7-12): Control

- Built governed ingestion layer
- Implemented quality checks

Phase 3 (M 13-18): Validation

- Tested reproducibility
- Trained 200+ staff

The Results

Regulatory Exam
Passed with **zero findings**

Deployment Speed
Reduced time by **40%**

Quality Issues
Caught **15 issues** pre-production

Key Learning

Big transformations need **executive sponsorship** AND **boots-on-the-ground champions**. This is your credibility moment—what actually happened, including the messy parts.

AI+IM GLOBAL SUMMIT

IMPACT METRICS

Data Quality as Your First Line of Defense

Why quality control matters—measured in business outcomes

1 Completeness

Flag missing required fields immediately before they enter the pipeline. If critical data is absent, the system stops.

System halts on missing critical data

2 Accuracy

Cross-check against authoritative sources in real-time. Validate that data values match their reference systems.

Real-time validation against sources

3 Timeliness

Reject stale data automatically. Set time-based thresholds to ensure information is current enough for decision-making.

Auto-reject stale data

4 Consistency

Validate that calculations match across systems. Ensure aggregations and derivations produce consistent results everywhere.

Cross-system calculation checks

Real-World Impact

In financial services, we implemented these controls after a regulatory finding:

Error Rate

12% → 0.8%

Before → After

Detection Time

6 days → 2 hours

Before → After

Audit Findings

3 in 18mo → 0 in 24mo

Quality controls aren't overhead—they're risk reduction .

AI+IM GLOBAL SUMMIT

INTERACTIVE EXERCISE

Spot the Quality Failure

Test Your Governance Intuition

The Scenario

You're reviewing a credit risk AI model before production deployment. Here's what you find:

Dataset Details

- **Source:** Customer transaction database
- **Records:** 2.4 million transactions
- **Time period:** Jan 2023 - Dec 2024
- **Completeness:** 97.8%
- **Accuracy validation:** Passed

Question to Audience

"Would you approve this for production? What questions would you ask?"

The Hidden Problem

That 2.2% missing data (53,000 transactions) was **NOT random**:

- **All missing transactions were >$50,000**
- **All were from Q4 2024**
- **Came from a system migration that failed partway**

What the AI Would Learn

- Large transactions don't exist in Q4 (completely wrong)

This is a critical error that would lead to under-estimating risk for high-value transactions.

What This Teaches

"97.8% complete" sounds good but masks critical gaps

Quality checks must look at **patterns** of missingness

Completeness metrics need **context**

TRACEABILITY

Lineage and Provenance

Traceability in practice—answering the three critical questions

1 What data was used?

Source identification

Complete inventory of all data sources with timestamps, versions, and origin details

2 How was it transformed?

Versioned processing logic

Every calculation, mapping, and business rule preserved with change history

3 When did changes occur?

Complete change history

Timestamped log of every change with user ID, approval, and impact assessment

Real Example: Regulatory Inquiry

"Why did your capital calculation change 3% last quarter?"

Without Lineage

Duration:	2-week investigation
Method:	Manual reconciliation
Result:	Uncertain conclusion
Cost:	Resource intensive

With Lineage

Duration:	2-hour trace
Method:	Automated lineage
Result:	Complete audit trail
Cost:	Minimal effort

REAL INCIDENT

The Cost of Poor Lineage

When You Can't Trace Your Data | Financial Services, 2023

The Trigger

Regulator asks: "Why did your Liquidity Coverage Ratio jump 8% in one day?"

The Question Seems Simple...

It wasn't. Without lineage, it became a 3-week crisis.

3-Week Investigation

Week 1: Confusion

- Data team: no logic changes
- IT: no system issues

Week 2: Chaos

- Manual rec of 15 systems
- Conflicting explanations

Week 3: Discovery

- Vendor changed file format
- Field mapping broke silently

The Damage

- 120 person-hours wasted
- Regulatory credibility damaged
- $500K emergency consulting fees

Delayed strategic initiatives

With Proper Lineage: 2-Hour Resolution

Hour 1: Root Cause ID

- Lineage tool shows source change
- Automated diff shows field mismatch

Hour 2: Fix & Validate

- Corrected mapping deployed
- Historical data reprocessed

ROI Calculation

Lineage system cost: **$200K**

This ONE incident would have paid for it **twice over**

The Broader Point

Lineage isn't about compliance theater . It's about operational efficiency and career preservation .

GLOBAL SUMMIT

IMPLEMENTATION

Governance by Design

Engineering accountability into systems—because policies alone don't work

Clear Data Ownership

Every dataset has a named steward with accountability documented in the system.

- Named Stewards
- Documented
- Escalation

Role-Based Access

You see only what you need. All changes logged with user ID, and privileged access requires justification.

- Need-to-Know
- Logged
- Justified

Approval Workflows

Model changes require sign-off, production deployments are controlled, and emergency overrides are tracked.

- Sign-Off
- Controlled
- Tracked

Separation of Duties

- Developers can't approve their own code
- Data owners can't modify their own access
- Auditors have read-only access

Governance can't be a policy document people ignore . It must be enforced by architecture .

LEARNING FROM FAILURES

Governance Antipatterns

What NOT to Do - Common Failures with Real Examples

1 "Policy as Governance"

❌ What it looks like:

- Beautiful 50-page governance policy
- Annual training everyone clicks through
- No system enforcement
- "Trust people to follow rules"

✅ Why it fails:

- People forget
- Shortcuts happen under pressure
- New hires don't know rules
- No way to verify compliance

Real example: Company had policy requiring dual approval. No system enforcement. Auditor found 40% of deployments had no documented approval .

2 "The Governance Theater"

❌ What it looks like:

- Governance committee meets monthly
- Fancy dashboards with green checkmarks
- Everyone reports "all good"
- Never digs into details

✅ Why it fails:

- Metrics game-able
- No one wants bad news
- Surface-level review misses issues

Real example: Dashboard showed "98% data quality" . Team was only checking data they knew was clean. Production data? Never validated.

3 "The IT Silo"

❌ What it looks like:

- AI governance owned by IT

✅ Why it fails:

- IT doesn't understand business context

ORGANIZATIONAL RESPONSIBILITY

Why Business Leadership Must Own This

This is not an IT problem—organizational structure determines governance reality

AI Risk = Enterprise Risk

When AI fails, it's not the data scientist who faces regulators—**it's the CEO**.

> Technology teams build tools. Leadership decides how tools are used and what risk is acceptable.

Technology Enables

They build tools. Business leaders decide how tools are used and what risk is acceptable.

Who gets promoted when AI succeeds? Who gets fired when it fails? That tells you where accountability really lives.

Leadership Defines

Leadership defines accountability. The organizational chart determines whether governance is real or theater.

If accountability isn't clear in the org chart, it doesn't exist.

In Financial Services

1 Model Risk Management

Reports to **Chief Risk Officer**, not CTO. Independence ensures risk assessment isn't influenced by delivery pressures.

2 Business Unit Sign-Off

Business unit heads approve AI deployments in their domains. They understand the regulatory environment and own the outcomes.

3 Board-Level Visibility

Board receives quarterly AI risk reports, making AI governance a **fiduciary responsibility**, not just a technical concern.

REAL-WORLD IMPLEMENTATION

Operational Reality

What the textbooks don't tell you about building resilient systems

The Real World

Systems fail. Humans make mistakes. Overrides happen. The question isn't "Can we prevent all failures?"

The Real Question:

"Do our controls expect failure and handle it responsibly?"

Good Governance Means

1. **Logged Overrides**
   Every override is logged and automatically reviewed by a human.
2. **Documented Exceptions**
   Exceptions require documented justification, not just an override button.
3. **Pattern Recognition**
   Patterns of workarounds trigger system fixes, not punishment.
4. **Learning Culture**
   Failures become learning moments, not cover-ups.

What Actually Happens

- → Manual overrides when systems freeze
- → Incomplete data when vendors miss delivery
- → Judgment calls when AI flags are wrong

Student Perspective

"I'm learning this in my supply chain classes—the difference between theoretical perfect systems and actual operational reality. This is about building resilience, not perfection."

When the presenter talked about that $15M miscalculation, my first thought was—I want to be the person who **prevents** those conversations.

CAREER DEVELOPMENT

What This Means for Future Leaders

Why you should care about AI governance as you enter this field

1 Trust Is Built Through Discipline

Shortsighted:

"Move fast and break things" doesn't work when regulators are watching.

Strategic:

Shortcuts today become crises tomorrow. The boring work of documentation and controls? That's what separates professionals from amateurs.

2 Good Questions Reduce Risk

- "Where does this data come from?"
  Understanding data sources and transformations
- "How do we know this is accurate?"
  Questioning validation and quality checks
- "What happens if this fails?"
  Considering consequences and mitigation

Asking these questions early = **career insurance**

3 Governance Skills Scale

Learn data quality now

Understand enterprise risk later

Understand lineage in school

Lead audit responses as manager

Practice documentation today

Build trustworthy systems tomorrow

The Opportunity

Most of my peers think AI governance is boring. I think it's where the important jobs will be. When the presenter talked about that $15M miscalculation, my first thought was—someone had to explain that to the CEO. I want to be the person who prevents those conversations, not the one having them.

CAREER DEVELOPMENT

Career Paths in AI Governance

Where These Skills Actually Lead | Student Presenter

When first explained this field to me, I had the same question you might: "Is this actually a career, or just a compliance checkbox?"

The Answer: It's One of the Fastest-Growing Career Tracks

Entry Level (0-3 years)

- **Data Quality Analyst**
- **AI Governance Analyst**
- **Model Validation Associate**

Mid-Career (4-8 years)

- **AI Risk Manager**
- **Model Governance Lead**
- **Data Lineage Architect**

Senior (8+ years)

- **Head of AI Governance**
- **Chief AI Risk Officer**
- **Model Risk Director**

Why Demand Is Exploding

- ✓ EU AI Act now in force
- ✓ US state-level AI regulations increasing
- ✓ Every major company needs these roles
- ✓ Supply of qualified people way below demand

The Educational Path

You don't need a PhD. You need:

1. Data fundamentals (DBs, ETL, quality)
2. Basic statistics/ML concepts
3. Regulatory awareness
4. Communication skills

REGULATORY LANDSCAPE

The Regulatory Landscape

Why This Matters More Next Year Than Today

Current State (2026)

United States

- No comprehensive federal AI law (yet)
- Sector-specific regulation (finance, healthcare)
- State-level laws emerging (CA, NY, CO)

European Union

- EU AI Act in force (phased through 2027)
- High-risk AI systems face strict requirements
- Penalties up to €35M or 7% global revenue

United Kingdom

- Principles-based approach
- Sector regulators issuing guidance
- Focus on financial services AI governance

What's Coming (Next 2-3 Years)

Expect:

- Federal AI accountability law (US)
- Mandatory AI impact assessments
- Third-party AI audits required
- Personal liability for AI failures

For Financial Services:

- Model risk management to all AI
- Real-time monitoring requirements
- Enhanced documentation standards
- Board-level AI risk reporting

What This Means

The systems we're describing aren't "nice to have." They're becoming legally required.

- Can't deploy AI in regulated contexts
- Competitive disadvantage
- Regulatory sanctions
- Difficulty hiring talent

SUMMARY

Key Takeaways

What we hope you remember about trustworthy AI

1

Trustworthy AI is a system outcome, not a feature

You can't add "trustworthiness" at the end. It must be designed in from the start through architecture, controls, and processes. Trust is earned through discipline, not declared in a press release.

2

Governance must be intentional

Hoping people do the right thing doesn't work at scale. Architecture must enforce good behavior through controls, logging, and workflows. If you can bypass it, it's not really governance.

3

Responsibility lives in the organizational chart

Technology enables AI. Leadership owns AI outcomes. If accountability isn't clear in the org chart, it doesn't exist . When AI fails, the CEO faces regulators, not the data scientist.

4

The future belongs to builders who understand limits

The most valuable skill isn't building the most powerful AI. It's building AI that works when it matters , fails safely when it doesn't, and can be explained either way.

DISCUSSION & QUESTIONS

Let's Talk

How these principles apply to your organization and career path

Topics We Can Dive Deeper On

- Technical implementation details
- Change management for governance
- Vendor selection for governance tools
- Career pathways and skill development
- Industry-specific governance challenges

Questions We Get Asked Often

- "How do I convince leadership to invest in this?"
- "Where do we start if we have nothing?"
- "What tools do you actually use?"
- "How do you handle AI governance in cloud environments?"

Time: 10-12 minutes for robust Q&A

Primary Speaker : Bharat Chaturvedi

Financial Institution Expert

Co-Speaker : Dhruv

Student Presenter

Supply Chain Management – TCNJ

Thank You for Your Attention