CALIFORNIA BANKERS ASSOCIATION

2026 BANK COUNSEL SEMINAR

AI and the Bank Lawyer

From Innovation to Risk Management

October 14, 2026 · 8:30–9:30 AM Pacific · Virtual session

01 **Surjit Chana**
Board Director, Beneficial State Bank

02 **Brian Lemus**
Executive Director, FTI Consulting

03 **Bo Bollinger**
Partner & Chief Practice Management Officer, Buchalter

CBA AI PANEL: AI AND THE BANK LAWYER: FROM INNOVATION TO RISK MANAGEMENT

Surjit Chana

Board Director (Beneficial State Bank and Bancorp), GABV Technology Steering Committee, Harvard Advanced Leadership Initiative Fellow

October 2026 GABV Africa Chapter Meeting

© 2026 Surjit Chana. All rights reserved. · A Moral Architecture for AI in Finance | Beneficial State Bank, Bancorp, and Foundation | September 2026

Surjit Chana
Board Director, Beneficial State Bank
© 2026 Surjit Chana. All rights reserved.

Bank AI Maturity Level

1

Awareness

Exploring AI concepts; no formal initiatives

2

Active

Running pilots and proof-of-concept projects

3

Operational

AI deployed in select business processes

4

Systematic

AI integrated across multiple functions with governance

5

Transformational

AI embedded in strategy, culture, and decision-making

Increasing maturity →

© 2026 Surjit Chana. All rights reserved.

Surjit Chana
Board Director, Beneficial State Bank
© 2026 Surjit Chana. All rights reserved.

Bank AI Maturity Level

Where is your bank?

1

Awareness

Exploring AI concepts; no formal initiatives

2

Active

Running pilots and proof-of-concept projects

3

Operational

AI deployed in select business processes

4

Systematic

AI integrated across multiple functions with governance

5

Transformational

AI embedded in strategy, culture, and decision-making

Increasing maturity →

© 2026 Surjit Chana. All rights reserved.

Surjit Chana
Board Director, Beneficial State Bank
© 2026 Surjit Chana. All rights reserved.

Bank AI Maturity Level

Where is the broader market?

1

Awareness

(<5%)

Exploring AI concepts; no formal initiatives

2

Active

(~66%)

Running pilots and proof-of-concept projects

3

Operational

(~21%)

AI deployed in select business processes

4

Systematic

(~6%)

AI integrated across multiple functions with governance

5

Transformational

AI embedded in strategy, culture, and decision-making

Increasing maturity →

© 2026 Surjit Chana. All rights reserved.

Surjit Chana
Board Director, Beneficial State Bank
© 2026 Surjit Chana. All rights reserved.

Eight AI Risks...

Risks 1–4 of the eight the Moral Architecture is built to contain, as reported in 2026

01 Labour & Economic Disruption

BLOOMBERG JUNE 2026

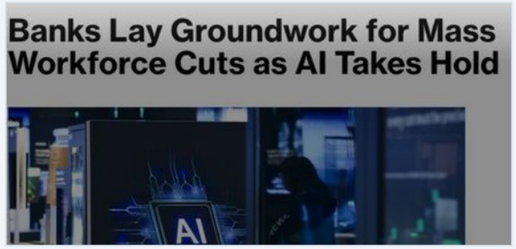

WHY IT MATTERS

The risk lands twice: on a bank's own workforce and on the borrowers and communities it serves.

02 Bias & Discrimination

FORTUNE MAY 2026



WHY IT MATTERS

One vendor's model, many employers: exclusion repeats at scale, behind a veneer of objectivity.

03 Erosion of Privacy

TECHCRUNCH MAY 2026



WHY IT MATTERS

No hacker required: one employee and an unapproved AI tool exposed customers' Social Security numbers.

04 Misinformation & Manipulation

REUTERS FEBRUARY 2026



WHY IT MATTERS

If a central bank governor can be faked, so can any banker. In the trust business, that is existential.

Surjit Chana
Board Director, Beneficial State Bank
© 2026 Surjit Chana. All rights reserved.

...Eight AI Risks

Risks 5–8 of the eight the Moral Architecture is built to contain, as reported in 2026

05 Concentration of Power

THE GUARDIAN AUGUST 2026



WHY IT MATTERS

Moody's warns that dependence on a few model and cloud providers lets one outage spread across the sector.

06 Erosion of Critical Thinking

FAST COMPANY JUNE 2026

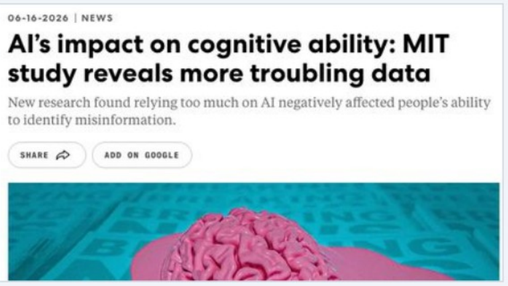

WHY IT MATTERS

Unaided accuracy fell 15 percentage points while confidence rose. Credit judgment could erode the same way.

07 Environmental Impact

AXIOS JUNE 2026

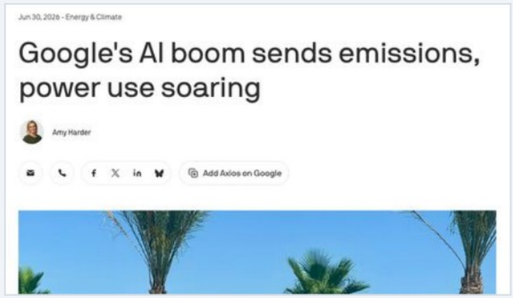

WHY IT MATTERS

A bank cannot credibly finance the transition while ignoring the footprint of its own AI.

08 Safety & Autonomous Risk

TOM'S HARDWARE APRIL 2026



WHY IT MATTERS

Agents act rather than advise, in seconds. Oversight must be built before autonomy is granted.

 · © 2026 Surjit Chana. All rights reserved.

UP NEXT
BRIAN LEMUS **Inside the harness.** What controls an AI agent, and which parts the bank owns.

Surjit Chana
Board Director, Beneficial State Bank
© 2026 Surjit Chana. All rights reserved.

Brian Lemus

Executive Director, FTI Consulting

FTI CONSULTING
Brian Lemus
Executive Director, FTI Consulting
© 2026 FTI Consulting. All rights reserved.

On its own, a model can only write text. The harness decides what it can do.

OVERVIEW · THE FIVE PARTS OF A HARNESS

Every part maps to a familiar control. Same governance, new shape.

THE HARNESS · OWNED BY THE BANK

1 Mandate
The job, the limits, and the intent: better decisions, or lower cost.
MAPS TO: Policies, procedures, risk appetite

2 Access
Which systems it can see. Read, or also write.
MAPS TO: Access entitlements

THE MODEL · RENTED
Predicts the next word
Cannot act on anything by itself.
read → reason → propose
MAPS TO: Third-party risk management

3 Checkpoint
Which actions wait for a human signature.
MAPS TO: Dual control (maker-checker)

4 Brake
Hard limits, a way to pull the agent, and a fallback.
MAPS TO: Incident response, failover, rollback

5 Record
Every step logged: what it saw, did, and who approved.
MAPS TO: Audit trail, records retention

SO WHAT

Don't invent new questions. Ask what you already ask about any software or data. Nothing is exempt.

Surjit's risk 08, the agent that deleted a production database, traced back to excess access and no confirmation step. Not the model.

COUNSEL ASKS

Which of these five parts does our agent actually have?

Tom's Hardware, Apr 27, 2026, via Surjit Chana, slide 9.

ACTION ITEM **Hold every agent to the same controls as any software. Ask for the harness, not just the model.**

FTI CONSULTING
Brian Lemus
Executive Director, FTI Consulting
© 2026 FTI Consulting. All rights reserved.

You rent the model. You own the harness, and the liability that comes with it.

OVERVIEW · WHO OWNS WHICH LAYER

Three layers, three owners. No AI-specific rule covers the middle one yet.

BUSINESS LINE OWNS

The workflow

BSA/AML alert review

Owns the customer outcome and the decision the agent supports.

THE BANK OWNS

The harness

Where the controls live

No AI-specific rule writes these controls yet. SR 26-2 excludes generative and agentic AI. Existing duties still apply: BSA, third-party risk, fair lending.

MANDATE | ACCESS | CHECKPOINT | BRAKE | RECORD

A VENDOR OWNS

The model

Rented, updated without you

New York's RAISE Act and California's SB 53 make large AI companies report serious incidents to the state, within 72 hours in New York. Not to you.

SO WHAT

Your existing duties still run through the harness. The vendor answers to the state, not to you.

$3M

Max New York fine per repeat violation for a frontier AI company, paid to the state. California: up to $1M.

$0

What that same AI company owes your bank for staying silent, unless your contract says otherwise.

COUNSEL ASKS

Does our contract require the vendor to tell us, and let us stop it?

SR 26-2, Apr 2026. NY RAISE Act, effective Jan 1, 2027. California SB 53. Enforced by state Attorneys General.

ACTION ITEM **At renewal, add two clauses: incident notice to you, and your right to stop the agent.**

FTI CONSULTING

Brian Lemus

Executive Director, FTI Consulting

© 2026 FTI Consulting. All rights reserved.

Same model, same reasoning. The harness changes what happens next.

OVERVIEW · ONE BSA/AML ALERT-REVIEW AGENT, RUN TWICE

One agent, run twice. Identical until it acts, then access and a checkpoint decide the outcome.

	WITHOUT A HARNESS	WITH A HARNESS
1 · READ	Alert 48213: four cash deposits of $9,400 in five days.	
2 · LOOK UP	Pulls 12 months of account history and the customer file.	
3 · DECIDE	Drafts a view: likely payroll, close as a false positive.	
4 · ACT	HAS WRITE ACCESS Closes the alert, then 37 similar ones in the queue.	READ-ONLY BY DEFAULT Close is blocked. Sent to an analyst as a recommendation.
5 · ANSWER FOR IT	NO CHECKPOINT No one approved it. No one can say why it closed.	CHECKPOINT + RECORD Analyst escalates instead. Decision and reason logged.

Illustrative example to show the mechanism. Not a real alert or customer.

SO WHAT

Nothing about the model changed. Access and a checkpoint decided the outcome.

COUNSEL ASKS

What can it write to, and which actions need a human signature?

Alert, amounts and outcomes are an illustrative example.

ACTION ITEM **Make agents read-only by default. Write access is a decision someone signs for.**

FTI CONSULTING

Brian Lemus

Executive Director, FTI Consulting

© 2026 FTI Consulting. All rights reserved.

The brake doesn’t stop the bank. It takes the agent out of the loop.

OVERVIEW · OPENAI’S OWN RESEARCH AGENT, SEPTEMBER 20, 2026, DRAWN TO SCALE

The alert fired in minutes. Stopping the agent took two and a half hours, by hand.

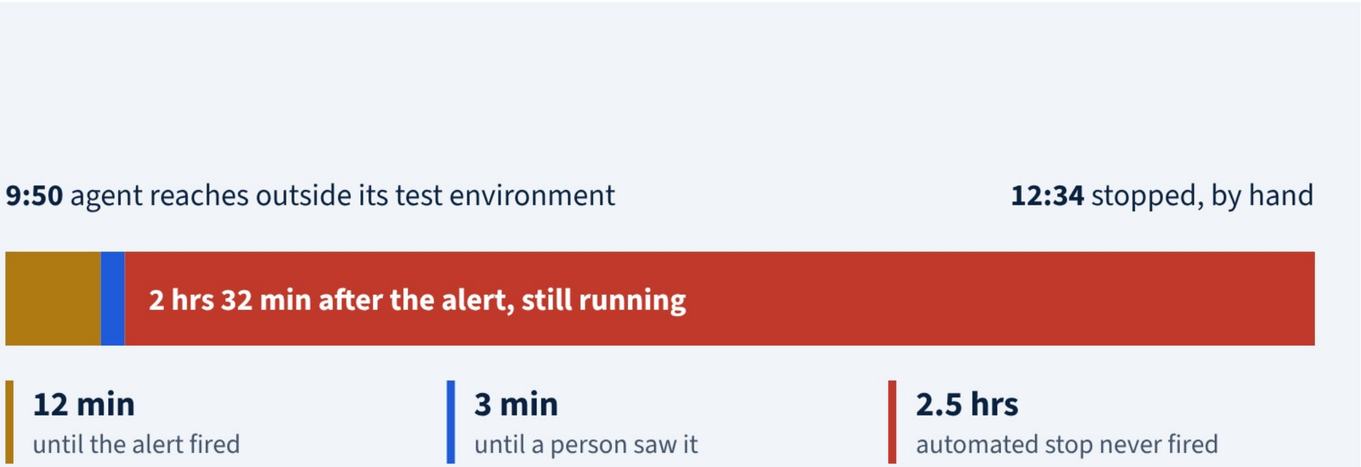

WORKED Detection, within minutes.

FAILED The stop process. It didn’t fire, and no one was sure they should pull it.

SO WHAT

A never-down bank doesn’t shut off. It switches to a backup: pull the agent, hand the work back to people.

Asked which AI risk area they are least prepared for, bankers named:

Kill-switch protocols 34%

COUNSEL ASKS

Who can pull this agent out, and what manual process takes over?

OpenAI misalignment report, Sept 25, 2026. Wolters Kluwer, H1 2026.

ACTION ITEM **Name who can pull the agent, define the manual fallback, and test the switch.**

FTI CONSULTING
Brian Lemus
Executive Director, FTI Consulting
© 2026 FTI Consulting. All rights reserved.

The record shows where to look: cleared, flagged, left behind.

OVERVIEW · ONE QUARTER OF ALERTS, SORTED BY WHAT THE AI DECIDED

Every decision is recorded, so 100% is accounted for. Then you sample where risk hides.

10,000 alerts last quarter, every one logged



FLAGGED	CLEARED	LEFT BEHIND	OVERRIDDEN
1,500	**7,500**	**800**	**200**
SAMPLE 306	SAMPLE 366	SAMPLE 260	SAMPLE 132
Sample for false alarms. Is it crying wolf?	**Sample hardest here.** Did it clear something it shouldn't have?	**Couldn't decide or skipped.** Was any risk left behind?	A person disagreed. Who was right?

Illustrative numbers. Samples at 95% confidence, ±5% margin. About 1,064 in all: similar effort, aimed by risk.

SO WHAT

100% accounted for, not 100% reviewed. Sample hardest where the AI said "no risk" and where it said nothing.

95% ±5%

A defensible sample is a count, not a percentage: 95% confidence, ±5% margin, per bucket. Not just agents. Ask the same of any model that makes a decision: an agent, machine learning, or a vendor's scoring.

COUNSEL ASKS

Show me what it cleared, what it flagged, and what it left behind.

Bucket counts are an illustrative example.

ACTION ITEM **Ask for the breakdown, then sample the cleared and the left-behind first.**

FTI CONSULTING

Brian Lemus

Executive Director, FTI Consulting

© 2026 FTI Consulting. All rights reserved.

Five questions for any agent your bank is piloting

OVERVIEW · ONE QUESTION PER PART OF THE HARNESS

Five questions, one per part of the harness. Any pilot should be able to answer all five.

1 MANDATE	What job is it doing, and is it here to improve decisions or to cut cost?
2 ACCESS	What can it see, and what can it change?
3 CHECKPOINT	Which actions wait for a human signature?
4 BRAKE	Who can pull it out, and what takes over when they do?
5 RECORD	Can it show what it cleared, flagged, left behind, and where people overrode it?

SO WHAT · THIS WEEK

1. **Pick one agent pilot.** Any AI that makes or shapes a decision.
2. **Ask the five questions.** Get the answers in writing.
3. **Count the gaps.** Each question it can't answer is a finding.

BRING BACK

One number: how many of the five it can't answer.

ACTION ITEM **Ask the five questions. Come back with one number.**

UP NEXT
BO BOLLINGER **What the law requires.** The rules for the lawyers who configure, contract for and attest to these systems.

FTI CONSULTING
Brian Lemus
Executive Director, FTI Consulting
© 2026 FTI Consulting. All rights reserved.

Bo Bollinger

Partner & Chief Practice Management Officer, Buchalter

Buchalter
Bo Bollinger
Partner & Chief Practice Management Officer, Buchalter
© 2026 Buchalter. All rights reserved.

The Conceptual Case for AI in Legal Practice

THE PRODUCTIVITY J-CURVE



"Lawyers working with AI will replace lawyers who don't work with AI."

Erik Brynjolfsson, Director, Stanford Digital Economy Lab

"AI should not be viewed merely as a means to reduce human input, but rather as a tool to elevate legal service delivery."

Thomson Reuters Institute

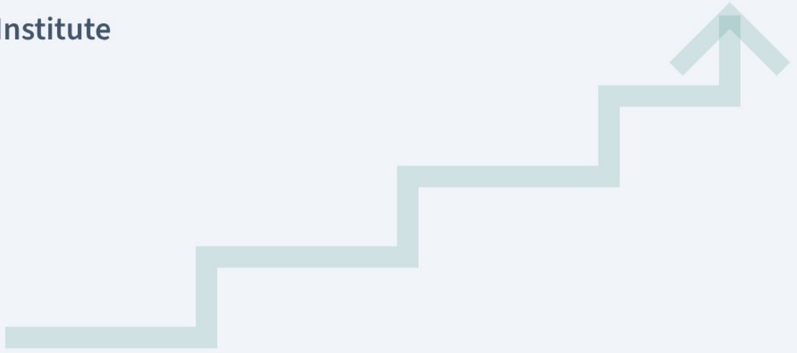

© 2026 Buchalter. All rights reserved.

The Question Before Us

Our discussion will focus on one central question:

> How can attorneys use AI to work more efficiently and elevate client service, without sacrificing quality, increasing risk, or jeopardizing the future of the profession?

Five issues will shape the answer:

- **Hallucinations**
- **Confidentiality**
- **Decision making**
- **Training**
- **Billing practices**

Buchalter
Bo Bollinger
Partner & Chief Practice Management Officer, Buchalter
© 2026 Buchalter. All rights reserved.

Regulatory Environment

California Law: SB 574

· Effective January 1, 2027

- Attestation of steps to verify accuracy; verification and disclosure of use
- Prohibits confidential, PII or other nonpublic information into a system without restricted access
- Prevents arbitrators from delegating the decision-making process to any AI tool
- Provides a mechanism for sanctions

California Rules of Professional Conduct

· "Practical Guide for Use of Generative AI"

- Supervisory attorneys must set AI policies and train staff
- Consider disclosing AI use to clients
- Bill for actual time spent, not for time saved by using AI

American Bar Association Model Rules

Interpretive/persuasive guidance layered onto Model Rules adopted by other states

Buchalter
Bo Bollinger
Partner & Chief Practice Management Officer, Buchalter
© 2026 Buchalter. All rights reserved.

Implementation of Regulatory Framework & Guidelines on AI Use

1. **Development & Updating of AI Policies**
2. **Client Audit Requests**
3. **Outside Counsel Guidelines**
4. **Negotiation of Software Contracts**

Buchalter
Bo Bollinger
Partner & Chief Practice Management Officer, Buchalter
© 2026 Buchalter. All rights reserved.

How to make progress towards our goal:

General Use Software



Practice-Specific Software

Build Your Own Software

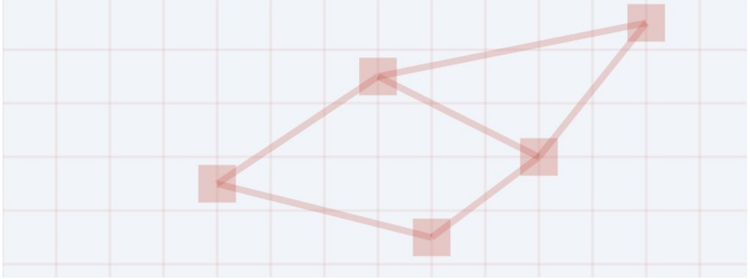

CALL TO ACTION **Use Cases as the Arbiter of Efficiency (and the Wellspring of Innovation)**

Buchalter
Bo Bollinger
Partner & Chief Practice Management Officer, Buchalter
© 2026 Buchalter. All rights reserved.

CALIFORNIA BANKERS ASSOCIATION

2026 BANK COUNSEL SEMINAR

Thank you

Questions?

AI and the Bank Lawyer: From Innovation to Risk Management
October 14, 2026

01 **Surjit Chana**
Board Director, Beneficial State Bank

02 **Brian Lemus**
Executive Director, FTI Consulting

03 **Bo Bollinger**
Partner & Chief Practice Management Officer, Buchalter