ASSESSED INTELLIGENCE

CASE STUDY · SHADOW AI IN A REGULATED ENVIRONMENT

Shadow AI at the Edge of Governance

The CB Financial Services Disclosure and the **ARISE Framework** ™ Control Gaps That Preceded It

REF · AI-CASE-2026-002
ANALYSIS DATE · 2026-06-30
CLASSIFICATION · PUBLIC

Prepared by Assessed Intelligence. This case study analyzes a publicly disclosed cybersecurity incident using only information in the company's SEC filing and reputable public reporting. References to specific **ARISE Framework** controls describe the controls that the public record indicates were absent or untested; they are reasoned inferences from disclosed facts, not findings from an engagement with the company.

01 EXECUTIVE SUMMARY

On May 11, 2026, CB Financial Services, Inc., the holding company for the Pennsylvania-based Community Bank, filed a Form 8-K under Item 1.05 to disclose a material cybersecurity incident.[1] The cause was not an external attacker. An employee handled non-public customer information, including customer names, Social Security numbers, and dates of birth, by entering it into an unauthorized artificial intelligence (AI) application while preparing a presentation.[2]

Legal analysis of the filing identifies it as the first Form 8-K to attribute a material cybersecurity incident to *shadow AI*, the practice of employees using AI tools without organizational approval or security review.[3] A review of the SEC full-text filing database found that the phrase "unauthorized artificial intelligence" appeared in exactly one 8-K on record: this one.[4]

The incident is instructive precisely because it was ordinary. No system was breached, no service was interrupted, and the data exposure resulted from a well-meaning shortcut rather than malice. The harm was governance harm: customer identity data left organizational control, and the event triggered disclosure duties and regulatory exposure. This analysis maps the incident to the **ARISE Framework™**, identifies the specific controls whose absence allowed the exposure to occur, and orders the remediation by where the work most directly belongs. One observation should be stated at the outset: the company's response was disciplined. The gaps were upstream of the response, in prevention and detection.

[1]CB Financial Services, Inc., Form 8-K (Item 1.05), filed with the SEC, May 11, 2026.
[2]Form 8-K, May 11, 2026; data elements (name, Social Security number, date of birth) and the presentation-preparation context as reported by American Banker and the Community Development Bankers Association, 2026.
[3]Wilson Sonsini Goodrich & Rosati, "'Shadow AI' Triggers First SEC Form 8-K for Unauthorized AI Use," May 29, 2026.
[4]American Banker, "A bank breaks its silence on its shadow-AI breach," 2026 (review of the SEC full-text filing database).

02 WHAT HAPPENED

Community Bank discovered the incident on May 5, 2026. Two days later, on May 7, the company determined the incident to be material. It filed the Form 8-K on May 11, signed by John H. Montgomery, the company's president and chief executive officer.[5] The four-business-day disclosure clock under Item 1.05 runs from the materiality determination, not from the date the incident is detected.[6]

By the company's account, the employee submitted customer data to an external AI application to complete a routine task. The company reported that it reached the application's vendor before the data could be used to train a model.[7] Notably, CB Financial determined the incident to be material even though it did not disrupt operations, customer account access, payment systems, or core infrastructure, and was not expected to have a material effect on the company's financial condition.[8] Data sensitivity alone supported the materiality conclusion. One commentator characterized the combination of name, Social Security number, and date of birth as a complete identity-theft starter set.[9]

> *"With shadow AI, the call is coming from inside the house. The threat originates not with a malicious outsider but with an employee looking for a shortcut."* — Public commentary on the CB Financial disclosure

03 WHY THIS PATTERN RECURS

Shadow AI is not an exotic risk; it is a common one. IBM's 2025 Cost of a Data Breach analysis recorded shadow AI as a distinct breach category and found that organizations with high shadow-AI involvement incurred approximately **$670,000** in additional breach cost and took a median of 247 days to identify the incident.[10] In the financial sector, unmanaged AI adoption accounted for roughly 20 percent of AI-related breaches in 2025, and among organizations that experienced an AI-related security incident, 97 percent lacked adequate AI access controls.[11]

The behavioral data explains the exposure. The generative-AI channel has become the single largest path for data movement from corporate to personal control inside the enterprise browser.[12] Survey evidence shows that leadership confidence outpaces actual visibility: a majority of knowledge workers report using unsanctioned AI tools at work, while most executives believe their employees use AI responsibly and that the organization has clear visibility into AI use.[13] That gap between assumed control and actual control is the condition in which a single shortcut becomes a disclosure event.

[5]Wilson Sonsini, May 29, 2026; Intelligize, "Shadow AI Lands on the SEC's Radar," 2026 (incident timeline and CEO signatory).
[6]Wilson Sonsini, May 29, 2026 (Item 1.05 four-business-day clock runs from materiality determination).
[7]American Banker, 2026 (vendor reached before model training).
[8]Wilson Sonsini, May 29, 2026 (materiality determined absent operational, customer-access, or financial-condition impact).
[9]Intelligize, "Shadow AI Lands on the SEC's Radar," 2026, quoting Andrew Hoog, Board Cybersecurity.
[10]IBM, Cost of a Data Breach Report 2025, via Cloud Security Alliance research note, 2026.
[11]Help Net Security, financial-sector cyber-threat reporting, 2026.
[12]LayerX, Browser Security Report 2025.
[13]Survey reporting via Diginomica, 2026.

Assessed Intelligence research places this in context: 76 percent of critical-sector organizations have AI risk-governance gaps.[14] CB Financial is one disclosed instance of a condition that is widely distributed and largely unmeasured.

[14]Assessed Intelligence Research, updated Q1 2026.

04 ARISE CONTROL GAP ANALYSIS

The **ARISE Framework** organizes assurance into seven domains: Govern, Identify, Protect, Detect, Respond, Manage, and Validate. The analysis below assigns each evident gap to the domain and the specific control where the corresponding work belongs. The control identifiers are drawn from the framework's control set. Read together, the gaps describe a prevention and detection layer that the public facts indicate was not in place, and a response layer that was.

GOVERN Policy, accountability, and AI risk ownership

The most direct control gap sits in acceptable use. The **ARISE Framework** requires an acceptable-use standard that explicitly prohibits the submission of non-public data into unapproved or public AI tools, paired with user acknowledgment and defined sanctions. A regulated institution operating that control, with an annual acknowledgment on file, establishes a clear, enforceable rule against the exact act that occurred. Policy alone does not stop a determined actor; it does establish the standard against which prevention, training, and enforcement are built.

ARISE CONTROL	EVIDENT GAP	PREVENTIVE OR MITIGATING EFFECT
`G.GV.C-01` **Acceptable Use**	No enforced acceptable-use rule explicitly prohibiting submission of non-public data into unapproved or public AI tools, with acknowledgment and sanctions.	Establishes the bright-line prohibition the employee crossed; creates the basis for enforcement, training content, and disciplinary consequence.
`G.GV.P-04` **AI Policies**	No codified AI policy defining approved tools, ownership, and conditions of use for AI across the workforce.	Replaces ad hoc individual judgment with an organizational standard for which AI tools may touch which data classes.
`G.RM-03` **AI Risk-Management Program**	Shadow AI not represented in an AI risk register with an assigned owner, tier, and mitigation plan.	Surfaces shadow AI as a named, owned, and tracked risk rather than an unmanaged condition discovered only after exposure.

MANAGE AI risk operations and human resources

Governance intent becomes durable only when it is operated. Operational AI risk management maintains the register, lifecycle assessments, and escalation paths that keep shadow AI in view between policy refresh cycles. The human-resources control set carries acknowledgment, onboarding, and sanction into the employment relationship, so that the acceptable-use rule is acknowledged, understood, and enforceable.

ARISE CONTROL	EVIDENT GAP	PREVENTIVE OR MITIGATING EFFECT
`M.RM-02` **AI Risk Management**	No operating cadence maintaining an AI risk register and lifecycle assessments through which shadow AI would be tracked and mitigated.	Keeps the risk owned and current between governance cycles rather than rediscovered through an incident.

ARISE CONTROL	EVIDENT GAP	PREVENTIVE OR MITIGATING EFFECT
M.HR **Human Resources**	Acceptable-use acknowledgment, onboarding, and sanction not integrated into the employment lifecycle for AI-specific conduct.	Makes the prohibition acknowledged and enforceable, and connects conduct expectations to consequence.

IDENTIFY Inventory, data mapping, and insider risk

An organization cannot govern AI use it cannot see. The framework requires an authoritative inventory of AI systems with registration before use, and a current map of where sensitive data flows. Customer identity data that is mapped is data to which handling controls can be attached. The incident also reflects the unintentional-insider pattern: a trusted employee, acting in good faith, moved regulated data outside organizational control.

ARISE CONTROL	EVIDENT GAP	PREVENTIVE OR MITIGATING EFFECT
I.AM-01 **AI System Inventory**	No living inventory of approved AI systems and no registration-before-use gate, so unsanctioned tools operate unseen.	Distinguishes approved from unapproved tools and gives enforcement and monitoring a defined baseline to act against.
I.DG-01 **Data Mapping**	Customer identity data flows not mapped end to end, so no control hooks were tied to the destinations that data could reach.	Ties classification to data-loss-prevention, encryption, and access controls at the points where regulated data could egress.
I.RM-02 **Insider Threat**	No program correlating human-resources, behavioral, and technical signals for unintentional mishandling of sensitive data.	Treats the well-meaning insider as a managed risk category with monitoring and escalation, not solely a training problem.

PROTECT Data handling, technical enforcement, and awareness

Prevention is where policy becomes mechanism. Data classification with handling rules, endpoint and egress controls that constrain transfers to unsanctioned destinations, and role-based AI training with in-application awareness prompts together convert a written prohibition into an enforced one. The framework specifies just-in-time awareness prompts that warn a user at the moment of a risky action; such a prompt at the point of upload addresses the precise sequence that produced this exposure.

ARISE CONTROL	EVIDENT GAP	PREVENTIVE OR MITIGATING EFFECT
P.DS-01 **Classification & Handling**	Customer data not labeled and handled under a classification scheme enforced by data-loss-prevention tags.	Makes regulated data recognizable to controls so that transfer to an external AI tool can be flagged or blocked by handling rule.
P.ES **Endpoint Security**	No endpoint data-loss-prevention or application control to detect or block uploads of sensitive data to unsanctioned AI applications.	Provides the technical enforcement point that stops or quarantines the action the acceptable-use rule prohibits.

ARISE CONTROL	EVIDENT GAP	PREVENTIVE OR MITIGATING EFFECT
P.AT-01 **AI Training**	No role-based AI-risk training and no just-in-time, in-application awareness prompts for risky data actions.	Equips staff to recognize the risk and warns at the moment of action, interrupting the shortcut before data leaves control.

DETECT Monitoring and anomaly detection

Detection determines how quickly an exposure becomes a managed event. The company learned of the incident and acted, but the public timeline indicates discovery after the fact rather than at the moment of upload. Continuous monitoring of data egress and anomaly detection tuned to bulk movement of sensitive fields to external destinations shorten the interval between action and awareness, and that interval governs both harm and disclosure posture.

ARISE CONTROL	EVIDENT GAP	PREVENTIVE OR MITIGATING EFFECT
D.CM **Continuous Monitoring**	No real-time telemetry or alerting on sensitive-data egress to external AI services; awareness followed the event.	Compresses time-to-detection so an exposure is caught at or near the moment of transfer rather than discovered later.
D.AD **Anomaly Detection**	No behavioral baseline to flag an out-of-pattern transfer of identity data to an unapproved external destination.	Raises an alert on the anomalous pattern and routes it into incident playbooks before the window for mitigation closes.

RESPOND Incident response and regulatory reporting

This domain is where the public record shows controls functioning rather than failing. The company moved from discovery on May 5 to a materiality determination on May 7 to an 8-K filing on May 11, within the four-business-day window, and reported reaching the vendor before the data could train a model. That sequence is consistent with a privacy incident-response process and a regulatory-reporting capability operating as intended. The lesson is not that response failed; it is that mature response cannot substitute for the prevention and detection layer that would have kept the data inside the organization in the first place.

ARISE CONTROL	EVIDENT GAP	PREVENTIVE OR MITIGATING EFFECT
R.IR-03 **Incident Response: Privacy**	Observed as functioning. Containment and assessment of a personal-data exposure proceeded on a defined timeline.	Limited propagation by reaching the vendor before model training; demonstrates the control class operating under real conditions.
R.IC.E-01 **Regulatory Reporting**	Observed as functioning. Materiality was assessed and the 8-K filed within the statutory window.	Met the SEC Item 1.05 obligation on time; confirms a working deadline-tracking and reporting capability.

VALIDATE Independent assurance

Controls that are assumed but never tested tend to be the controls that are missing. Independent internal audit of the AI-use control environment verifies that prohibition, inventory, classification,

enforcement, and monitoring exist and operate. Validation is the domain that would have surfaced the prevention and detection gap as a finding before it surfaced as a disclosure.

ARISE CONTROL	EVIDENT GAP	PREVENTIVE OR MITIGATING EFFECT
V.CI-03 **Internal Audit: Security**	No independent assurance that controls governing AI use existed and operated; the gap persisted unverified until the incident.	Surfaces the missing prevention and detection controls as an audit finding, with time to remediate before exposure occurs.

05 THE REMEDIATION PATH

The controls above are not equally urgent. The sequence below orders them by leverage: each action makes the next one enforceable. The path moves from establishing the rule, to making AI use visible, to enforcing the rule technically, to detecting failure, and finally to verifying that the layer holds.

| Prohibit and define

Stand up the acceptable-use standard (**G.GV.C-01**) and the AI policy (**G.GV.P-04**). State the prohibition on entering non-public data into unapproved AI tools, name the approved tools, and require annual acknowledgment with defined sanctions.

| Inventory and gate

Establish the AI system inventory with registration before use (**I.AM-01**) and map the flows of regulated customer data (**I.DG-01**). Visibility into approved tools and sensitive-data paths is the precondition for every enforcement control that follows.

| Classify and enforce at the endpoint

Apply the classification and handling scheme (**P.DS-01**) and the endpoint and egress controls (**P.ES**) that detect or block transfers of labeled data to unsanctioned destinations. This is the point at which the written prohibition becomes a technical one.

| Detect the pattern

Tune continuous monitoring (**D.CM**) and anomaly detection (**D.AD**) to sensitive-data egress, so that an exposure is caught at the moment of transfer and routed into the response process that the company has already shown it can run.

| Train with context

Deliver role-based AI training with just-in-time, in-application awareness prompts (**P.AT-01**). A warning at the moment of upload addresses the precise sequence that produced this exposure, and reinforces the rule where work actually happens.

| Own the risk and validate it

Carry shadow AI in the AI risk register with an assigned owner (**G.RM-03** / **M.RM-02**), and verify the full control set through independent internal audit (**V.CI-03**). Validation converts an assumed control environment into an evidenced one.

06 IMPLICATIONS

Two conclusions follow from this case. First, shadow AI is a governance and disclosure issue before it is a technical one; the CB Financial filing establishes that data sensitivity alone can render an AI-misuse incident material under SEC rules, independent of operational disruption.[15] Organizations that treat

[15] Wilson Sonsini, May 29, 2026.

unsanctioned AI use as an inconvenience rather than a reportable risk are exposed to the same sequence.

Second, prevention and detection must carry the weight that response cannot. CB Financial responded well; the exposure still occurred and still required disclosure. The controls that would have changed the outcome are upstream, in the prohibition, the inventory, the classification, the endpoint enforcement, and the monitoring. The **ARISE Framework** names those controls and orders the work. Organizations operating in regulated sectors must assume that employees are already using AI tools that the organization has not approved, and must build the control layer on that assumption rather than on the confidence that they are not.

Governance that holds up when it matters most. Assessed Intelligence applies the ARISE Framework™ to make AI use visible, governed, and verifiable before an incident forces the question.

References

1. CB Financial Services, Inc., Form 8-K (Item 1.05), filed with the U.S. Securities and Exchange Commission, May 11, 2026.
2. American Banker, "A bank breaks its silence on its shadow-AI breach," 2026.
3. Wilson Sonsini Goodrich & Rosati, "'Shadow AI' Triggers First SEC Form 8-K for Unauthorized AI Use," client alert, May 29, 2026.
4. American Banker review of the SEC full-text filing database, as reported in "A bank breaks its silence on its shadow-AI breach," 2026.
5. CB Financial Services, Inc., Form 8-K, May 11, 2026; timeline (discovery May 5, materiality determination May 7, filing May 11) as reported by Wilson Sonsini and Intelligize, 2026.
6. Wilson Sonsini Goodrich & Rosati, client alert on SEC Item 1.05 disclosure timing, May 29, 2026.
7. American Banker, "A bank breaks its silence on its shadow-AI breach," 2026 (vendor reached before model training).
8. Wilson Sonsini Goodrich & Rosati, client alert, May 29, 2026 (materiality determined despite no operational or financial disruption).
9. Andrew Hoog, Board Cybersecurity, as quoted in Intelligize, "Shadow AI Lands on the SEC's Radar," 2026.
10. IBM, Cost of a Data Breach Report 2025 (shadow AI category: approximately $670,000 in additional cost; 247-day median time to identify shadow AI Breach lifecycle), as cited by Cloud Security Alliance research note, 2026.
11. Help Net Security, financial-sector cyber-threat reporting, 2026 (shadow AI approximately 20 percent of AI-related breaches; 97 percent lacked adequate AI access controls).
12. LayerX, Browser Security Report 2025 (generative AI as the largest corporate-to-personal data-movement channel in the enterprise browser).
13. Nutanix and related 2026 survey reporting on shadow-AI prevalence and the gap between executive confidence and worker behavior, via Diginomica, 2026.
14. Assessed Intelligence Research, AI risk-governance gap analysis, updated Q1 2026.
15. Wilson Sonsini Goodrich & Rosati, client alert, May 29, 2026 (data sensitivity sufficient to support materiality independent of operational impact).