BAZZI CONSULTING

ServiceNow and Databricks for GRC

Table of Contents

Introduction
ServiceNow IRM: A Unified Governance & Workflow Platform
Databricks: Scalable Data & AI for Security Analytics
Integrated Architecture and Data Flow
Example Use Case: Predicting Ransomware Risk
Operationalizing Predictive Insights in ServiceNow IRM
Business Value and Compliance Benefits
Phased Implementation Roadmap
Conclusion

Introduction

Financial services firms face *reactive, compliance-driven* risk management. Traditional approaches often rely on periodic assessments and manual reporting, leaving gaps in visibility. **Regulatory pressure** is intense: frameworks like the EU's DORA, NIS2, and the US FFIEC cybersecurity guidance mandate stringent controls and continuous monitoring[1][2]. Yet many organizations still operate in silos – data and controls live in disparate systems or spreadsheets[3]. This fragmentation makes it hard to correlate threats with business impact or automate remediation. In practice, cyber teams often *react* to incidents and audit findings rather than anticipate them[4][5]. For example, one industry analysis notes that mature firms must move from "reactive" to "proactive, data-driven" risk strategies to protect critical assets[4]. Compliance demands only heighten this challenge: financial institutions are subject to rules such as DORA and PCI DSS, noncompliance of which "can result in hefty fines"[1]. In short, **today's challenge** is that security data (logs, vulnerabilities, threat feeds) and governance processes (controls, policies, audits) are disconnected, preventing real-time risk insight and automated decision-making[3][6].

ServiceNow IRM: A Unified Governance & Workflow Platform

ServiceNow's Integrated Risk Management (IRM) platform provides a unified GRC engine to break these silos. Built on the Now Platform®, IRM centralizes policies, controls, and risk registers into one system[7][6]. Its key modules include **Internal Controls/Policy & Compliance** (linking controls to regulations and automating testing), **Risk Management** (identifying and scoring risks), and **Audit Management** (planning

audits and collecting evidence). Because ServiceNow IRM connects “the business, security, and IT with an integrated risk framework,” it transforms fragmented, manual processes into a single program[8]. All risk and compliance data live on the platform, so executives see a *holistic view* of risk. For instance, ServiceNow IRM “identifies non-compliant controls, monitors high-risk areas, and manages the Key Risk Indicator (KRI) and Key Performance Indicator (KPI) library with automated data validation”[9]. This yields real-time dashboards and reports of risk posture, making audit trails and compliance evidence automatically available. In short, ServiceNow IRM provides no-code workflow automation and integration hubs so risk findings are captured as formal records. It enables firms to go from isolated spreadsheets to an enterprise GRC program with one “system of record”[8][6].

- **Unified GRC:** IRM consolidates controls, risks, policies and audit tasks in one platform, eliminating silos[6][9].
- **Real-time monitoring:** Continuous evidence collection and dashboards move teams toward a *proactive* posture[5][9].
- **Workflow engine:** Automated flows and notifications in ServiceNow enforce controls and remediation (e.g. route incidents to owners, require attestations).
- **Auditability & compliance:** Automated control testing and evidence logs support faster audits and adherence to standards[7][10].

In essence, ServiceNow IRM acts as the *governance hub*: it defines risk frameworks, ties security controls to policies and regulations, and orchestrates the response. This is precisely the high-level governance layer needed for financial services under DORA,

FFIEC, NIS2, PCI DSS, etc. By linking data feeds to risk records, IRM ensures that technical findings and business requirements stay in sync.

Databricks: Scalable Data & AI for Security Analytics

Databricks provides the complementary *data platform* to power predictive cyber risk analytics. As a unified data lakehouse, Databricks can ingest and store **massive volumes** of security telemetry from across the enterprise. For example, a modern security operations platform (such as Hunters SOC) can feed **endpoint telemetry, network traffic, identity logs, cloud logs, vulnerability scans, threat intelligence feeds, and more** into Databricks[11][12]. Databricks' Apache Spark engine and Delta Lake allow these heterogeneous datasets to be ingested in raw form and normalized (often via a multi-step ETL or the Medallion architecture) for analytics. In practice, a Databricks workspace might run pipelines to continuously collect logs (firewall, IDS, VPN), scan results (patch and vulnerability data), and CTI feeds. All this data becomes queryable at scale.

Once ingested, Databricks enables advanced analytics and machine learning. Security teams can train predictive models on the consolidated data: for instance, classifiers or time-series models that estimate **risk scores** for systems or assets. Databricks supports both batch and streaming ML: you could train a model on historical incident records and current patch/vulnerability status to predict a *ransomware risk probability*. The platform also supports unsupervised techniques (clustering or anomaly detection) to spot unusual patterns in SOC alerts. Customers regularly build custom threat-detection models and deploy them on the Databricks Lakehouse[13]. Importantly,

Databricks retains governance: data lineage and versioned tables (via Unity Catalog) ensure that all analytic datasets are auditable. Through all this, the goal is to turn raw security data into *actionable insights* (e.g. a risk score or alert label) ready for consumption by the IRM layer.

With Databricks handling the heavy data lifting, organizations gain:

- **Scale and variety:** Ingest any security data – from SIEM logs to vulnerability reports – and fuse it for analysis[11][12].
- **AI/ML at scale:** Leverage Databricks MLflow and Spark to build complex models (supervised risk scoring, anomaly detectors, LLM-based threat analysis, etc.) on up-to-date data.
- **Shared data:** Multiple teams (infoSec, risk, IT) can access the same datasets for collaboration. (Notably, Databricks now offers a ServiceNow **connector** for ingestion as well[14].)

Together, ServiceNow IRM and Databricks cover the full stack: **ServiceNow** manages the governance and workflows, while **Databricks** provides the data lakes and analytics. As one announcement puts it, the two systems are becoming tightly coupled via “Zero Copy” integration: Databricks’ Delta Sharing will enable “high-bandwidth, bi-directional, and secure integration” between the platforms[15][16]. In practice today, even before Zero Copy arrives, Databricks data can be pushed to ServiceNow via API calls or webhooks (see next section).

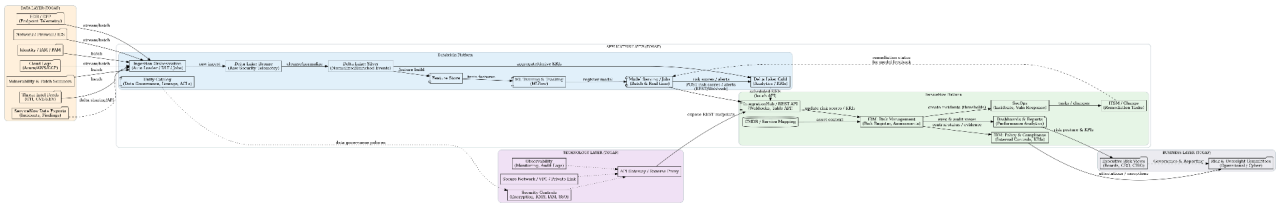

Figure: Sample integration architecture. Security data (logs, scans, alerts) flow into Databricks, which runs ML models to produce risk scores or alert events. Databricks then triggers a ServiceNow IRM workflow (via REST/webhook or IntegrationHub) to update risk records or create an incident, closing the analytics-to-action loop.

Integrated Architecture and Data Flow

The integrated solution has two tiers: **Data & Analytics (Databricks)** and **Governance & Response (ServiceNow IRM)**. In the Databricks tier, security and business data are ingested into a centralized Lakehouse. This includes: historical security incidents, SOC event logs, asset inventories (CMDB), patch/vulnerability scanner results, and threat intel feeds. Databricks pipelines periodically refresh these tables and join them into feature sets for modeling. Using Spark ML or Databricks AutoML, the team develops predictive models – for example, a model that outputs a **ransomware risk score** for each critical system. Databricks notebooks can also run anomaly detection routines (e.g. clustering or outlier detection on log streams) and assign an “anomaly alert” flag to assets.

When the analytics tier detects a significant risk event (e.g. a risk score exceeds a threshold, or an anomaly is flagged), it passes this insight to the ServiceNow tier. This can be done via a *webhook* (Databricks can send an HTTP POST) or by having

ServiceNow poll Databricks through its REST API. On the ServiceNow side, an **IntegrationHub flow** is configured with a REST trigger endpoint[17]. When Databricks posts an alert or risk result, this flow is invoked. The flow logic looks up the corresponding IRM record (for example, an asset's risk profile) and writes the new risk score or creates a risk/incident record. (Databricks can also use the ServiceNow REST API to update records directly if needed.)

Simultaneously, ServiceNow may push data back. For instance, if a risk score is updated, ServiceNow can invoke a callback to Databricks – perhaps to log an acknowledgement or capture remediation results back into the data lake (as illustrated in the Databricks community example[18]). In effect this is a **closed-loop**: insights flow from Databricks to ServiceNow, and outcome data flows back into Databricks for model retraining.

Example flow steps:

1. *Data preparation*: Databricks pipelines continuously update tables of system patch levels, past incidents, SOC alerts, etc.
2. *Model scoring*: A nightly Databricks job applies an ML model to compute a "ransomware risk likelihood" for each critical server.
3. *Alert trigger*: For any server with risk > threshold, Databricks uses a webhook to call ServiceNow.
4. *ServiceNow action*: A ServiceNow flow with a REST trigger receives the data and updates the Risk Assessment record (or opens an Incident/Risk Task) for that server[17]. It may also send notifications to the security and risk teams.
5. *Response and feedback*: The assigned team logs mitigating actions in ServiceNow

(e.g. schedule patches). IntegrationHub/API updates the Databricks table of completed remediations, feeding back into future risk models.

This architecture leverages ServiceNow's workflow engine to "turn insights into instant, AI-powered action"[15][19]. Notably, ServiceNow's planned **Zero Copy** integration with Databricks (via Delta Sharing) will soon allow even more seamless bi-directional data access[16]. In any case, the result is that predictive risk outputs from Databricks become formal IRM records, complete with audit trails and governance.

Example Use Case: Predicting Ransomware Risk

Consider a concrete scenario in a bank's IT environment. A Databricks model is built to predict the probability of a ransomware attack on a critical system (e.g. a payment processing server) in the next 30 days. Inputs to this model include: (a) **historical incident logs** (e.g. past ransomware events in similar systems), (b) **patch/vulnerability status** (time since last patch, number of known CVEs on that system), (c) **threat intelligence indicators** (public intelligence linking certain ransomware gangs to similar targets), and (d) **current SOC alerts** (unusual activity or IDS flags on that host). The model might use supervised learning (trained on past incidents) or a rule-based risk score combining these factors. When all factors align (outdated patches, relevant threat outbreak, and suspicious network alerts), the model outputs a high risk score (e.g. "Ransomware_Risk=0.9").

This risk score is then sent to ServiceNow IRM. In IRM, the Risk Assessment module automatically updates the likelihood of a ransomware threat for that asset. If the score exceeds a policy-defined threshold, ServiceNow can **automatically trigger mitigation**

tasks. For instance, it might open a ticket to apply emergency patches or activate additional monitoring on that system. ServiceNow's **Internal Controls** framework ensures that this risk is logged under the proper regulatory control (e.g. "Cyber Resilience"). The platform's dashboards immediately reflect the increased risk rating and any new tasks, so executives see a live picture of the bank's exposure.

Key to this process is that the analytics is **continuous** and **predictive**. Instead of waiting for a quarterly risk assessment (which would have missed this scenario until after the fact), the bank uses data to anticipate threats. If the model's probability later drops (e.g. after patches are applied), ServiceNow can even de-escalate the risk automatically. The system ensures that every data-driven insight from Databricks flows into the official IRM record – complete with timestamps and source details for audit[9]. In practice, one can analogize this to a Security Operations Center built on Databricks: threat scores become risk events in ServiceNow that workflows act upon.

Operationalizing Predictive Insights in ServiceNow IRM

Once Databricks feeds in risk insights, ServiceNow IRM uses them to **automate governance actions**. For example:

- **Risk Score Updates:** ServiceNow can store the new risk likelihood in the risk register and re-compute overall exposure. This recalculation can trigger business-impact analysis (e.g. expected loss).

- **Task and Incident Creation:** A high-risk alert can generate remediation tasks (e.g. "install critical patch X on System Y") or even an Incident record in ITSM. The flow can assign these tasks to the appropriate owners using ServiceNow's AI-assisted work

assignment.

- **Notifications and Escalations:** The system can notify compliance officers or executives when a risk crosses certain thresholds, ensuring cross-team visibility.
- **KRI Monitoring:** In IRM's Key Risk Indicator library, new metrics from Databricks (such as "Days Until Exploit Likely") can be added and mapped to controls. The platform will then treat them as live KRIs and continuously monitor them[9].
- **Audit Documentation:** All of the above actions generate a traceable audit log. Auditors can see when Databricks scored a risk and when ServiceNow generated tasks or reports, closing the loop on compliance evidence.

In essence, ServiceNow **operationalizes** the predictive analytics. As one technical blog points out, Databricks insights "automatically trigger AI Agents or workflows on the ServiceNow platform," providing a direct real-time connection between analytics and action[19][16]. This means security analysts and auditors are no longer manually translating data. Instead, the IRM system itself updates risk scores and control statuses based on the ML output, and immediately drives mitigation. For instance, if the ransomware model predicts a 90% likelihood of attack, a ServiceNow risk score can jump from "low" to "high" and enforcement actions can be queued without human delay.

Business Value and Compliance Benefits

The Databricks-ServiceNow integration delivers **proactive cyber resilience and governance**. Key business benefits include:

- **Proactive Risk Posture:** Organizations move from firefighting to *predict-and-prevent*. By incorporating real-time analytics into risk workflows, firms

can address vulnerabilities before breaches occur[4][5]. This reduces incident costs and insurance premiums over time.

- **Improved Auditability:** All predictive insights and actions are timestamped within ServiceNow's system of record. Automated evidence collection (continuous control testing) dramatically cuts audit effort[10][20]. For example, evidencing that a model ran and an IRM task closed is fully traceable.
- **Regulatory Compliance:** Tying analytics into IRM directly supports frameworks like DORA and FFIEC. These rules demand continuous monitoring and evidence of risk management. The integrated solution ensures controls are data-driven and documented. (As one guide notes, such platforms help "ensure compliance with relevant cybersecurity regulations" like FFIEC guidelines[2].) Real-time dashboards also give regulators and internal audit teams confidence that risks are seen and managed promptly.
- **Cost Reduction:** Automation cuts manual work. Continuous compliance testing replaces labor-intensive audits[21]. Less downtime means lower business risk exposure (e.g. ransomware can be nipped early). Consolidating tools also saves licensing and training costs.
- **Cross-Functional Visibility:** Perhaps most importantly, this integration breaks down silos across IT, security, risk, and the business. ServiceNow dashboards provide executives with a unified view of cyber risk and controls[22][5]. Teams can collaborate: for example, IT sees a risk flagged by security analytics and finance sees its potential impact on operations, aligning everyone to the same data.

- **Regulatory Readiness:** By mapping predictions to established controls and frameworks, organizations can more easily demonstrate compliance with standards. For instance, DORA requires ICT risk identification; using automated IRM maps risk scores to DORA controls for audit.

In summary, by merging Databricks AI with ServiceNow IRM workflows, a bank can *shorten the gap between insight and action*. As one analyst notes, 60% of AI projects fail due to data issues – here, the platforms ensure the data is ready and integrated so that AI-driven insights "turn into instant, AI-powered action"[15][19].

Phased Implementation Roadmap

A structured rollout is essential. A phased approach ensures organizational buy-in and data readiness[23][24]. A possible roadmap:

1. **Phase 1 – Strategy & Assessment:** Define objectives (e.g. protect critical systems) and scope. Inventory data sources (logs, vulnerability scanners, threat feeds) and assess IRM maturity[25]. Perform baseline risk assessments to identify quick wins and key risk areas. Build the foundational data models in ServiceNow (controls library, risk taxonomy).
2. **Phase 2 – Core Deployment:** Establish the data pipelines into Databricks. Integrate log collectors, vulnerability scanners, and threat-intel feeds into the Lakehouse. In ServiceNow, deploy core IRM modules: set up policy frameworks (e.g. map controls to NIST/ISO), configure the Risk Register with initial assets, and design basic workflows. Implement dashboards to visualize current risk

posture. (This aligns with “Policy and Compliance Management” and “Risk Management” enablement[26].)

3. **Phase 3 – Pilot Use Case:** Select a pilot scenario (e.g. ransomware risk on a key system). Build the first predictive model in Databricks using historical incident and vulnerability data. Configure a ServiceNow flow (via IntegrationHub) to receive the model’s risk scores. Run end-to-end tests: have Databricks trigger ServiceNow when the model flags high risk, and ensure IRM updates accordingly[17][9]. Refine the ML model and flow based on feedback.
4. **Phase 4 – Integration & Automation:** Expand integration. Connect more IRM tables and create IntegrationHub actions or scripts to handle bulk updates (e.g. weekly risk scoring). Leverage ServiceNow APIs or the new Zero Copy connector for batch syncing. Introduce automatic remediation workflows (e.g. assigning patching tasks) tied to Databricks alerts. Also integrate closed-loop feedback: have ServiceNow post mitigations back into Databricks for model improvement[18]. This follows ServiceNow’s “Risk Integration and Automation” phase[27].
5. **Phase 5 – Scale & Optimize:** Gradually apply the system to other cyber risks (insider threat, fraud, compliance violations). Incorporate additional data (third-party vendor risk scores, business transaction anomalies, etc.) into Databricks models. Use ServiceNow’s analytics and emerging features (AI agents) to optimize decision-making. Continuously monitor model performance and update. This is akin to “Maturity Expansion” – adding Vendor Risk, Operational Risk, and continuous threat intelligence feeds[28].

Throughout, strong governance is key: involve IT architects, data engineers, security analysts, and executive sponsors. Regularly review model outputs with risk owners to validate accuracy. Document each step for compliance: for instance, log how the Databricks model was tested and approved.

Conclusion

By combining Databricks and ServiceNow IRM, financial firms can achieve a *truly proactive* cyber risk management posture. Databricks supplies the scalable data lake and ML horsepower to transform raw security data into risk insights, while ServiceNow IRM provides the workflows and governance to act on them. Together they bridge the gap from analytics to action: predictive risk scores feed directly into formal risk records, automatically updating scores, triggering tasks, and notifying stakeholders[19][9]. This integrated approach addresses the shortcomings of traditional methods (reactiveness, siloes, audit gaps) and aligns with regulatory regimes like DORA and FFIEC. Ultimately, organizations gain continuous visibility, faster mitigation, and audit-ready evidence – a combination that protects the business and complies with regulators in equal measure[1][22].

Sources: Industry reports and vendor documentation on cyber risk and GRC platforms were used to prepare this whitepaper[4][1][3][7][16][19][11][23][9]. Each source is cited inline by cursor ID and line numbers.

---

[1] Modernizing Financial Cybersecurity: From Reactive to Resilient

https://www.commvault.com/blogs/modernizing-financial-cybersecurity-from-reactive-to-resilient

[2] Why Cyber Risk Monitoring is a Priority for Financial Services Institutions

https://www.gatekeeperhq.com/blog/why-cyber-risk-monitoring-is-a-priority-for-financial-services-institutions

[3] How to Integrate Different Aspects of Governance, Risk and Compliance · Riskonnect

https://riskonnect.com/governance-risk-compliance/how-to-integrate-different-aspects-of-governance-risk-and-compliance/

[4] State of Cyber Risk Management 2025

https://www.guidepointsecurity.com/blog/state-of-cyber-risk-management-2025/

[5] [6] [10] [21] [22] What is ServiceNow IRM? Your Guide to Integrated Risk Management - ServiceNow Integrated Risk Management - Unified Governance Risk and Compliance Data - Surety Systems

https://www.suretysystems.com/insights/servicenow-irm-integrated-risk-management-surety-systems/

[7] [8] [9] ServiceNow Integrated Risk Management (IRM)

https://www.servicenow.com/content/dam/servicenow-assets/public/en-us/doc-type/resource-center/data-sheet/ds-Integrated-risk-management.pdf

[11] [12] [13] Databricks Security Ops with Hunters SOC | Databricks Blog

https://www.databricks.com/blog/2023/03/29/security-operations-data-lakehouse-hunters-soc-platform-now-available.html

[14] Configure ServiceNow for Databricks ingestion | Databricks Documentation

https://docs.databricks.com/aws/en/ingestion/lakeflow-connect/servicenow-source-setup

[15] [16] ServiceNow and Databricks announce Zero Copy partnership - ServiceNow Press

https://www.servicenow.com/company/media/press-room/databricks-collaboration.html

[17] [18] [19] Insights to Action

https://www.servicenow.com/community/workflow-automation-blogs/maximizing-business-value-by-converting-insights-from-databricks/ba-p/3168561

[20] [23] [24] [25] [26] [27] [28] servicenow.com

https://www.servicenow.com/community/s/cgfwn76974/attachments/cgfwn76974/community-central-forum/2980/1/White%20Paper_ServiceNow%20GRC%20Implementation.pdf