Threat Modeling and Risk Assessment Report

Nexa Ledger — Multi-Tenant Reconciliation Platform

Nexa Systems • Security Engineering • STRIDE Methodology

Document ID	NEXA-SEC-TMRA-2026-014	**Version**	2.1
Classification	Confidential — Restricted Distribution	**Status**	Final — Approved
Owner	Priya Nandakumar, Director of Security Engineering	**Approved By**	Marcus Ferreira, Chief Information Security Officer
Effective Date	August 2, 2026	**Next Review**	February 2, 2027

Document Purpose

This report presents the STRIDE-based threat model and qualitative risk assessment for **Nexa Ledger**, the multi-tenant payment reconciliation and ledger API platform operated by Nexa Systems. It enumerates in-scope assets and trust boundaries, profiles credible threat actors, catalogues threats per system component, scores each threat on a likelihood/impact matrix, and records the mitigations and residual risk carried forward. It is the authoritative reference for the Q3 2026 security architecture review and for any engineering change that touches authentication, the reconciliation engine, or the public webhook surface.

1 Executive Summary

Nexa Ledger is the system of record for transaction reconciliation across Nexa Systems' fintech client base, settling entries between merchant acquirers, the Vertex Payments processing network, and the Synapse Bank Connect account-aggregation service. Because the platform holds live financial data for 340+ tenant organizations, Security Engineering commissioned this threat model ahead of the Q3 2026 architecture review to establish a documented, repeatable baseline for how the system can be attacked and where controls currently stand.

The assessment applied the STRIDE framework [6] against a decomposed data-flow diagram of the platform (Section 3), profiled five threat actors with credible motive and access to the system (Section 4), and scored every identified threat on a five-by-five likelihood/impact matrix consistent with NIST SP 800-30 [3] (Section 6). In total, **14 distinct threats** were catalogued across seven in-scope components.

Severity Band	Count	Share of Findings
CRITICAL	3	21%
HIGH	5	36%
MEDIUM	4	29%
LOW	2	14%

The three Critical findings concentrate in a single area: trust assumptions on the public webhook ingestion path used by Vertex Payments to deliver asynchronous settlement events. An unauthenticated or replayed webhook can currently inject fabricated settlement records into the reconciliation engine before signature verification is enforced consistently across all environments. This is tracked as **VULN-01** in Section 7 and is the single highest-priority remediation in this report.

Key recommendations, detailed fully in Section 8, are:

1. Enforce HMAC signature verification and replay-window checks on *every* environment for the Vertex webhook endpoint, not only production (VULN-01, VULN-04).
2. Retire the Admin Console's direct database query path in favor of calls that route through the Reconciliation Engine, and move console sessions onto short-lived, scoped tokens (VULN-02).
3. Add row-level audit triggers on the Ledger Database so tampering by a credentialed insider is detectable within one reconciliation cycle rather than at month-end close (VULN-03).
4. Rotate and scope Vertex Payments API credentials per tenant instead of a single shared integration key, limiting blast radius if a partner-side credential leaks (VULN-06).

Twelve of the fourteen findings have an assigned owner and target remediation date in Section 7; residual risk after planned mitigations lands at LOW to MEDIUM for all components except the webhook ingestion path, which remains HIGH until VULN-01 closes.

2 Scope and Methodology

2.1 In-Scope Components

This assessment covers the production deployment of Nexa Ledger on Meridian Cloud, spanning the API gateway, authentication service, reconciliation engine, webhook ingestion endpoint, admin console, ledger database, and audit log store. The following are explicitly **out of scope**: the Nexa Editor product (a separate codebase and threat surface), physical security of Meridian Cloud data centers (covered by Meridian's own SOC 2 attestation), and client-side security of tenant integrations built against the Nexa Ledger SDK, which is addressed in a separate partner security guideline.

2.2 Framework: STRIDE

Threats were elicited per component using **STRIDE** [6], the threat categorization model that decomposes adversarial behavior into six classes:

Category	Security Property Violated
Spoofing	Authentication — an actor claims an identity that is not theirs
Tampering	Integrity — data or code is modified without authorization
Repudiation	Non-repudiation — an actor denies performing an action, and the system cannot prove otherwise
Information Disclosure	Confidentiality — data is exposed to an actor not entitled to see it
Denial of Service	Availability — a legitimate actor is denied service
Elevation of Privilege	Authorization — an actor gains capabilities beyond their entitlement

Each in-scope component from Section 3 was evaluated against all six categories; categories with no credible threat path for a given component are omitted from Section 5 rather than padded with placeholder entries.

2.3 Risk Scoring

Every identified threat is scored qualitatively on **Likelihood** and **Impact**, each rated 1 (lowest) to 5 (highest), following the qualitative risk-rating approach in NIST SP 800-30 [3]. The product of the two scores maps to a severity band:

Severity Band	Score Range	Response Expectation
CRITICAL	15–25	Remediate before next release; escalate to CISO
HIGH	8–14	Remediate within current quarter
MEDIUM	4–7	Remediate within two quarters, or accept with sign-off
LOW	1–3	Track; remediate opportunistically

Likelihood accounts for attacker motivation and required access (informed by the threat actor profiles in Section 4), and Impact accounts for financial loss, regulatory exposure, and reconciliation integrity across tenants. Where a component processes card-network-adjacent data, CVSS v3.1 base scores [1] are additionally recorded for individual vulnerabilities in Section 7 to align with Vertex Payments' partner disclosure requirements.

2.4 Sources

Component boundaries and data flows were reconstructed from the current production Terraform state and the Nexa Ledger service catalog; asset owners were confirmed against the internal architecture register [4]. STRIDE categorization was cross-checked against relevant technique clusters in the MITRE ATT&CK knowledge base [2] and the current OWASP Top 10 [5] to avoid omitting well-documented attack classes.

3 System Architecture and Asset Inventory

Nexa Ledger ingests two categories of external input: interactive requests from tenant client applications through the API Gateway, and asynchronous settlement events pushed by Vertex Payments to a public webhook endpoint. Both paths converge on the Reconciliation Engine, which is the only component permitted to write ledger entries. A separate pull path retrieves account-balance snapshots from Synapse Bank Connect on a scheduled basis. The Admin Console exists for tenant support staff to investigate reconciliation discrepancies and, notably, retains a direct query path into the Ledger Database that bypasses the Reconciliation Engine's write-path validation — this shortcut is the subject of finding VULN-02 in Section 7.

3.1 Data Flow Diagram

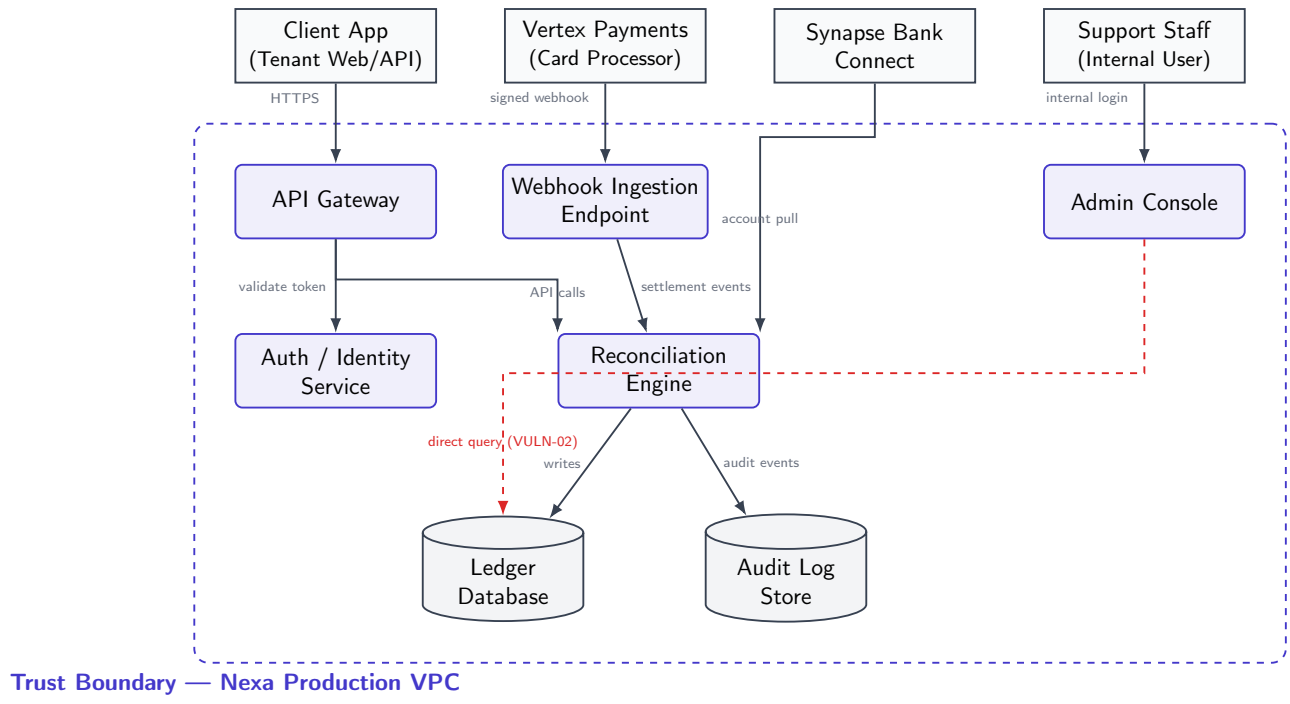

Figure 1: Data flow diagram for Nexa Ledger. External entities sit above the dashed trust boundary; the dashed red edge marks the Admin Console's direct database path flagged in the vulnerability findings.

External Entity Process Data Store Trust Boundary Flagged Risk Path

3.2 Asset Inventory

Table 1 lists every component inside the trust boundary along with the data it handles and its assigned owner. Data classification follows the Nexa internal scheme: *Restricted* (financial or regulated data, encrypted

at rest and in transit, access logged), *Confidential* (business-sensitive, access-controlled), and *Internal* (routing/operational metadata with no direct financial content).

Asset	Type	Classification	Trust Zone	Owner
API Gateway	Process, public-facing	Internal	Perimeter (DMZ)	Platform Infrastructure
Webhook Ingestion Endpoint	Process, public-facing	Confidential	Perimeter (DMZ)	Ledger Platform Team
Auth / Identity Service	Process	Confidential	Internal, app tier	Identity Team
Reconciliation Engine	Process	Restricted	Internal, app tier	Ledger Platform Team
Admin Console	Process, internal tool	Restricted	Internal, app tier	Support Engineering
Ledger Database	Data store (Postgres)	Restricted	Internal, data tier	Data Platform Team
Audit Log Store	Data store (object storage)	Restricted	Internal, data tier	Security Engineering

Table 1: In-scope assets, their data classification, and ownership.

4 Threat Actor Profiles

Likelihood ratings in Section 6 are anchored to the five actors below, ordered by the risk they currently pose to Nexa Ledger rather than by sophistication alone. Capability describes technical resources; access describes the starting foothold each actor is assumed to already have or can plausibly acquire.

Actor	Motivation	Capability	Primary Target
Organized fraud ring	Direct financial gain via fabricated or altered settlement entries	High	Webhook ingestion, reconciliation engine
Malicious insider	Financial gain, coercion, or grievance; abuses legitimate access	Medium	Admin console, ledger database
Compromised third-party vendor	Unwitting supply-chain foothold via a trusted integration	Medium	Vertex/Synapse integration credentials
Opportunistic credential stuffer	Low-effort account takeover for resale or fraud	Low	Client-facing API gateway, auth service
Nation-state-adjacent actor	Espionage or destabilization targeting high-value fintech clients	High	Ledger database, audit log store

Actor Spotlight: Organized Fraud Ring ("Ledger Skimmers")

This is the actor most likely to act against Nexa Ledger in the next twelve months, and it is the actor behind the scenario in VULN-01. Groups in this category have previously been observed against other Vertex Payments merchants injecting forged settlement webhooks to create phantom refunds, then withdrawing the resulting balance before month-end reconciliation catches the discrepancy. They do not need to compromise Nexa infrastructure directly: any endpoint that accepts a webhook without verifying its signature and a fresh timestamp is a viable target. Their operational tempo is bursts of automated traffic timed around known reconciliation windows, which makes them detectable if request-rate and signature-failure telemetry on the webhook endpoint is monitored — currently it is not (see VULN-01).

Malicious insider scenarios in scope assume a support engineer or contractor with legitimate Admin Console credentials, not a compromised account (that path is covered under the fraud-ring and credential-stuffer profiles). The primary concern is the console's direct database query path (Section 3, VULN-02), which lets a credentialed insider alter ledger rows without generating the audit events the Reconciliation Engine would normally emit.

Compromised third-party vendor covers a scenario where a Vertex Payments or Synapse Bank Connect credential leaks outside Nexa's control — for example through a breach at the partner rather than at Nexa. Because Nexa currently issues one shared integration key per partner rather than per tenant (VULN-06), a single leaked credential grants broader access than any individual tenant relationship would justify.

Opportunistic credential stuffers run automated login attempts against the API Gateway using credential lists purchased from unrelated breaches. Their capability is low and their success depends entirely on tenant password reuse, but their volume is high and constant, making them a persistent background threat rather than a one-time event.

The **nation-state-adjacent actor** is included for completeness given several Nexa Ledger tenants operate in regulated cross-border payment corridors; likelihood is rated low in the absence of any specific intelligence, but impact is rated at the maximum band because this actor's objective (bulk exfiltration of the ledger and audit store) would compromise every tenant simultaneously.

5 STRIDE Threat Analysis

Table 2 enumerates every threat identified by walking each STRIDE category against every component in the data flow diagram (Figure 1). Threats T-01 through T-04, T-06, T-07, T-11, and T-13 carry a **CRITICAL** or **HIGH** rating and are tracked as formal findings with assigned CVSS scores and mitigation owners in Section 7; the remainder are tracked in the risk register and revisited each quarter (Section 8).

ID	Component	STRIDE	Threat Scenario	Risk
T-01	Webhook Ingestion	Spoofing	Forged Vertex settlement webhook accepted without a valid HMAC signature in non-production environments, letting fabricated events reach the reconciliation queue.	CRITICAL
T-02	Webhook Ingestion	Tampering	A previously valid signed webhook payload is replayed with no timestamp or nonce check, re-applying an already-settled transaction a second time.	HIGH
T-03	Reconciliation Engine	Tampering	A leaked, partner-shared integration key lets an actor submit settlement records for a merchant the credential was never scoped to.	HIGH

ID	Component	STRIDE	Threat Scenario	Risk
T-04	Admin Console	Elevation of Privilege	A support session can query and modify the Ledger Database directly, bypassing the Reconciliation Engine's write-path validation and audit-event emission.	CRITICAL
T-05	Auth Service	Spoofing	Credential-stuffing bot succeeds against tenant accounts without MFA, using password lists sourced from unrelated breaches.	MEDIUM
T-06	Ledger Database	Tampering	An insider with data-tier access executes a direct SQL update on a settled ledger row outside any application code path, leaving no application-level trail.	CRITICAL
T-07	Audit Log Store	Repudiation	Audit entries are mutable under the shared data-tier service credential, letting an actor alter or delete evidence of a prior tampering event.	HIGH
T-08	API Gateway	Spoofing	Differential error responses on login allow enumeration of valid tenant email addresses ahead of a targeted credential-stuffing run.	LOW
T-09	API Gateway	Denial of Service	Absence of per-tenant rate quotas lets one noisy or compromised tenant integration degrade gateway latency for every tenant.	MEDIUM
T-10	Reconciliation Engine	Information Disclosure	Verbose error responses on failed reconciliation return internal transaction and merchant identifiers to the calling client.	LOW
T-11	Auth Service	Elevation of Privilege	A downgraded user's JWT role claim is not re-validated on token refresh, so the prior, higher-privilege role persists until natural expiry.	HIGH
T-12	Admin Console	Repudiation	Support staff share a small pool of console logins for on-call coverage, so individual actions cannot be attributed to a specific person.	MEDIUM
T-13	Ledger Database	Information Disclosure	Nightly backups are encrypted at the volume level only; card-adjacent metadata fields are not separately encrypted, widening exposure if a backup is mishandled.	HIGH
T-14	Audit Log Store	Denial of Service	No per-tenant write-throughput isolation; a log flood from one tenant's activity can cause audit events from other tenants to be dropped.	MEDIUM

Table 2: STRIDE threat catalogue across all in-scope components.

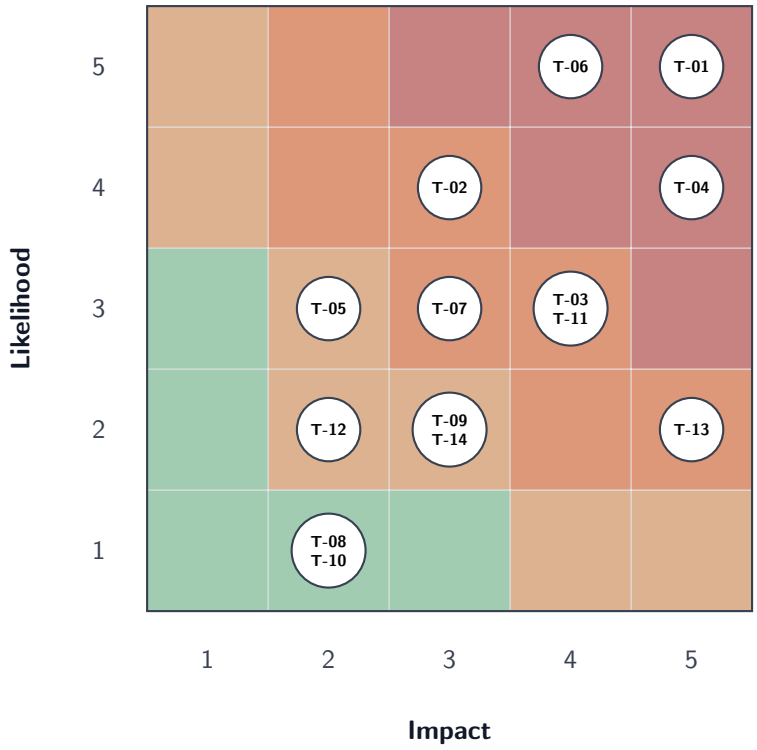

Figure 2: Likelihood–Impact matrix for all fourteen catalogued threats. Cell shading marks the severity band; markers show where each threat currently sits.

6 Risk Assessment Matrix

Each threat in Table 2 was scored on Likelihood (informed by the actor profiles in Section 4) and Impact (informed by the asset classification in Section 3), both on a 1–5 scale, following the banding defined in Section 2. Figure 2 plots all fourteen threats on the resulting grid; cell shading marks the severity band a Likelihood–Impact combination falls into, independent of whether a threat currently occupies that cell.

Band	Count	Threat IDs
CRITICAL	3	T-01, T-04, T-06
HIGH	5	T-02, T-03, T-07, T-11, T-13
MEDIUM	4	T-05, T-09, T-12, T-14
LOW	2	T-08, T-10

The concentration of Critical and High findings in the upper-right quadrant of the matrix is driven almost entirely by the webhook ingestion and admin console trust decisions identified in Section 3 — both are paths where a single missing control (signature verification, write-path enforcement) converts a Medium-likelihood actor into a High- or Critical-impact outcome. This concentration is the basis for prioritizing VULN-01 through VULN-03 ahead of the remaining findings in Section 7.

7 Vulnerability Findings and Mitigations

The eight threats rated CRITICAL or HIGH in Section 5 are tracked below as formal findings with an assigned CVSS v3.1 base score [1], mitigation, owner, and remediation status. Medium and Low threats

remain in the risk register referenced in Section 8 and are revisited each quarterly review rather than tracked individually here.

VULN-01 — Webhook Signature Verification Not Enforced in All Environments (CVSS 9.1, Critical)

Production correctly rejects an unsigned or badly-signed Vertex webhook. Staging and the shared development environment do not: signature verification is gated behind a feature flag that defaults to `off` outside production, a leftover from early integration testing that was never removed. Because the Reconciliation Engine's ingestion logic is identical across environments, any actor who can reach the staging endpoint — which is on a public hostname, only unlisted — can inject settlement events that reconcile against real tenant balances if staging and production ever share a database snapshot during a support investigation. This is the highest-priority item in this report.

ID	Finding	CVSS	Mitigation	Owner	Status
VULN-01	Webhook signature verification disabled by default outside production	9.1	Enforce HMAC-SHA256 verification and a 5-minute replay tolerance in every environment; alert on signature-failure rate.	Ledger Platform Team	IN PROGRESS
VULN-02	Admin Console queries the Ledger Database directly, bypassing write-path validation and audit events	8.6	Route all console mutations through the Reconciliation Engine API; replace session cookies with short-lived scoped tokens.	Support Eng. / Ledger Platform	OPEN
VULN-03	Ledger Database reachable via direct SQL with no row-level audit trigger	8.9	Add PostgreSQL row-level audit triggers writing independently to the Audit Log Store; alert on writes outside the service role.	Data Platform Team	IN PROGRESS
VULN-04	Webhook replay window and event nonce not enforced	7.4	Add a nonce cache and reject duplicate event IDs alongside the VULN-01 signature fix.	Ledger Platform Team	IN PROGRESS
VULN-05	Audit Log Store writable under the shared data-tier service credential	7.1	Move the store to append-only object storage with retention lock; scope the write credential to an append-only policy.	Security Engineering	OPEN
VULN-06	Vertex integration key shared across tenants rather than scoped per tenant	6.8	Issue per-tenant integration credentials; deprecate the shared key.	Ledger Platform Team	OPEN
VULN-07	JWT role claim not re-validated on token refresh	6.5	Re-fetch the current role from the Auth Service on every refresh instead of carrying the original claim forward.	Identity Team	REMEDIATED
VULN-08	Ledger Database backups lack field-level encryption on card-adjacent metadata	6.2	Apply column-level encryption to card-adjacent metadata columns ahead of backup snapshotting.	Data Platform Team	OPEN

Table 3: Formal vulnerability findings, CVSS base scores, and remediation tracking.

VULN-07 closed on July 18, 2026 and is retained here for traceability through the next audit cycle; it will drop from this table at the following revision. The remaining seven findings are reviewed bi-weekly in the Security Engineering standup until each reaches REMEDIATED.

8 Residual Risk and Recommendations

8.1 Residual Risk

Assuming the mitigations in Section 7 land on their target dates, residual risk for Nexa Ledger falls to LOW or MEDIUM for every component except the webhook ingestion path, which remains HIGH until VULN-01 and VULN-04 both close — signature enforcement alone does not fully retire the risk without the accompanying replay protection. No mitigation in this report is expected to bring any component to zero residual risk; the Medium and Low threats in Table 2 (T-05, T-08, T-09, T-10, T-12, T-14) are accepted at their current rating pending the Q4 2026 review, since none individually crosses the threshold that would require an out-of-cycle fix under the response expectations in Section 2.

8.2 Recommended Roadmap

1. **Q3 2026 (by September 30):** Close VULN-01, VULN-02, VULN-03, and VULN-04. These four findings account for all three Critical threats and two of five High threats, and collectively address the webhook trust boundary and the admin console bypass identified as the two structural weaknesses in this assessment.
2. **Q4 2026 (by December 31):** Close VULN-05, VULN-06, and VULN-08. Per-tenant credential scoping (VULN-06) should be sequenced with the Vertex partner team, since it requires a coordinated credential rotation rather than a unilateral Nexa change.
3. **Ongoing:** Re-run this threat model whenever a new external integration is added (a new payment rail or aggregator introduces new external entities and data flows that this document does not yet cover), and at minimum every two quarters regardless of architectural change.
4. **Process:** Add signature-failure rate and admin-console direct-query volume as standing metrics on the Security Engineering dashboard, so that both flagged paths in Figure 1 are observable before the underlying findings close, not only after.

8.3 Sign-Off

Role	Name	Date
Author	Priya Nandakumar, Director of Security Engineering	July 28, 2026
Technical Reviewer	Owen Falk, Principal Engineer, Ledger Platform	July 30, 2026
Approver	Marcus Ferreira, Chief Information Security Officer	August 2, 2026

8.4 Revision History

Version	Date	Summary
1.0	March 12, 2026	Initial threat model covering the API Gateway, Auth Service, and Ledger Database only.
2.0	June 9, 2026	Added Webhook Ingestion Endpoint, Reconciliation Engine, and Synapse Bank Connect data flow following the account-aggregation launch.
2.1	August 2, 2026	Added Admin Console direct-query path (VULN-02) after discovery during the Q3 access review; refreshed all CVSS scores against current CVSS v3.1 guidance.

References

[1] FIRST.org, Inc. Common vulnerability scoring system version 3.1: Specification document. Technical report, Forum of Incident Response and Security Teams, 2019.

[2] MITRE Corporation. MITRE ATT&CK: Adversarial tactics, techniques, and common knowledge. MITRE Corporation, 2024. Enterprise matrix.

[3] National Institute of Standards and Technology. Guide for conducting risk assessments. NIST Special Publication 800-30 Rev. 1, U.S. Department of Commerce, September 2012.

[4] Nexa Ledger Platform Team. *Nexa Ledger Service Catalog and Architecture Register*. Nexa Systems, v2.1 edition, 2026.

[5] OWASP Foundation. Owasp top ten: The ten most critical web application security risks. OWASP Foundation, 2021. Community-maintained project reference.

[6] Adam Shostack. *Threat Modeling: Designing for Security*. Wiley, 2014.