RISK REGISTER AND ROADMAP

Sample Risk Register and Remediation Roadmap

Anonymized portfolio sample showing prioritization, ownership, closure evidence, and phased engineering work.

CLIENT
Confidential European B2B software company

PLATFORM
Multi-tenant web platform

RISK MODEL
Likelihood x impact with explicit assumptions

PURPOSE
Convert assessment findings into accountable and verifiable remediation

How to read this sample

The register below is intentionally abridged. It shows how findings were translated into risk statements, time horizons, accountable roles, and objective closure evidence. Technical identifiers and implementation details have been generalized.

> **SCORING PRINCIPLE**
> Severity was not derived from labels alone. Likelihood reflected attacker prerequisites and exposure; impact reflected tenant scope, reversibility, recoverability, and business consequence. Assumptions that changed the score were stated explicitly.

Sample risk register

ID	Severity	Condition	Business consequence	Owner	Target
R-01	HIGH	No independent recovery path	Destructive event causes permanent multi-tenant loss.	Platform owner	0-7 days
R-02	HIGH	Production artifact not traceable to reviewed source	Assessment claims, rollback, and incident containment remain uncertain.	Operations	0-7 days
R-03	HIGH*	Untrusted or conditional edge transport	Users cannot establish normal trust; on-path interference remains plausible if public.	Operations	0-7 days
R-04	MEDIUM	Object-storage runtime identity has administrative authority	Application compromise inherits platform-wide storage administration.	Operations	30 days
R-05	MEDIUM	Database runtime identity owns schema objects	Application compromise can modify structure and widen persistence.	Backend / DBA	30 days
R-06	MEDIUM	Alternate listeners bypass intended edge controls	Traffic can avoid TLS, headers, routing policy, or monitoring.	Operations	0-30 days
R-07	MEDIUM	One tenant-scoped write path lacks target authorization	Authenticated user can cause cross-tenant integrity damage.	Backend	0-7 days
R-08	MEDIUM	Privileged renderer input can cross a browser execution boundary	Persistent script execution on a published client site.	Backend / Frontend	0-30 days
R-09	MEDIUM	Security logging and alerting are incomplete	Abuse may remain undetected and difficult to reconstruct.	Operations	30 days
R-10	MEDIUM	Shared internal credential spans all tenants	One component compromise gains platform-wide internal access.	Platform engineering	Strategic

Sample risk register (continued)

ID	Severity	Condition	Business consequence	Owner	Target
R-11	MEDIUM	Unauthenticated telemetry endpoint discloses infrastructure metadata	Aids attacker reconnaissance and targeting of known-version components.	Operations	0-30 days
R-12	MEDIUM	Background job payload lacks schema validation	No confirmed exploitable sink; defense-in-depth gap for future job types.	Backend	30 days
R-13	MEDIUM	Third-party account access and MFA status unconfirmed	A single compromised registrar, DNS, or email account could undermine every other control regardless of application quality.	Platform owner	0-30 days
R-14	LOW	Business-logic workflow allows an out-of-order state transition	Bypasses an internal review step; no additional data access gained.	Backend	30 days
R-15	LOW	Missing baseline HTTP security headers	Reduces defense-in-depth; compounds impact if a render-layer fix is ever incomplete.	Operations	30 days

* Conditional rating: severity depends on whether the affected hosts are public or intentionally private under managed trust.

WORKED SCORING EXAMPLE

R-07 (tenant-scoped write path lacking target authorization) scored likelihood 3 of 5: it requires an authenticated account, which is a real but not high barrier, and a foreign-tenant identifier, which was shown to be obtainable through a legitimately accessible listing endpoint elsewhere in the product, not through guessing. It scored impact 5 of 5: the affected helper was shared code reused across every tenant context, so the consequence was rated at full-platform scope rather than a single customer, even though any one exploit attempt only touches one foreign record at a time. The resulting score of 15 reflects a real, demonstrated path with a moderate barrier to entry and a severe, platform-wide consequence class, not a worst-case assumption stacked on a worst-case assumption.

Closure evidence by risk

ID	Required closure evidence
R-01	Clean-environment restore succeeds; recovery time recorded; retained copies cannot be deleted by production identities.
R-02	Running digest maps to an approved commit, signed build, deployment event, and known-good rollback artifact.
R-03	Normal clients validate the certificate chain; strict transport policy is enabled only after successful deployment.
R-04	Runtime storage identity passes required operations and fails policy, identity-management, bucket-deletion, and backup-deletion tests.
R-05	Runtime role performs application DML but fails schema, role, grant, unrelated-database, and backup-modification tests.
R-06	External and internal scans show only approved ingress; no application or telemetry bypass listener remains.
R-07	Foreign-tenant and global writes are denied before persistence; authorized same-tenant writes remain functional.
R-08	Unsafe values receive deterministic rejection; legacy content cannot execute; restrictive browser policy is verified.
R-09	Sensitive actions produce durable off-box events; priority alerts reach a named owner; a test incident is reconstructed.
R-10	Credentials are tenant or project scoped; rotation, revocation, attribution, and unexpected-source detection are demonstrated.
R-11	Unauthenticated requests to the telemetry endpoint fail from untrusted networks; the authorized collector continues to scrape successfully.
R-12	Malformed job payloads are rejected at enqueue time with no job created; legitimate job types continue to process correctly.
R-13	Written client confirmation of access holders and multi-factor status for registrar, DNS, and email-provider accounts, recorded as client-confirmed evidence.
R-14	Out-of-order state-transition attempts are rejected server-side; the intended step-by-step workflow continues to function.
R-15	Full baseline header set present and correctly scoped across tested public and authenticated routes.

Phased remediation roadmap

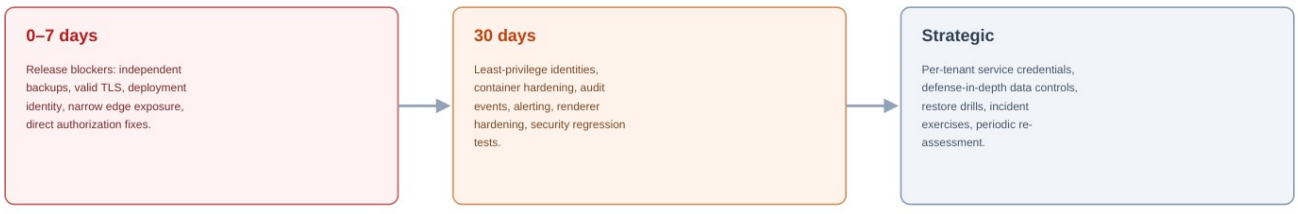

Phase 0: establish facts and contain obvious exposure

1. Record the running artifact digest and effective production configuration.
2. Restrict alternate application, administration, and telemetry listeners.
3. Confirm the real public trust model and intended ingress paths.
4. Create an immediate off-box copy of relational data and object storage using credentials unavailable to the application.
5. Assign named owners for every release-blocking risk.

Phase 1: release blockers

Work item	Why it blocks release	Verification
Independent recovery	A destructive event is otherwise permanent.	Full clean restore and measured recovery time.
Deployment identity	Security findings cannot be tied to production.	Commit, signed digest, deployment event, and runtime metadata agree.
Valid edge trust	Users cannot establish normal trusted transport.	System trust succeeds and approved edge policy is present.
Tenant authorization fix	Confirmed cross-tenant integrity path.	Foreign-target write denied with no side effect.
Renderer input fix	Confirmed persistent browser execution path.	Unsafe values rejected; no execution in public or preview rendering.

Phase 2: containment and operational readiness

- Introduce separate database owner, migration, runtime, worker, read-only, and backup identities.
- Introduce scoped object-storage runtime, administration, backup, and restore identities.
- Run application containers as non-root, drop unnecessary privileges, and reduce writable filesystem paths.
- Move security events off-box; define detections for authorization failures, identity changes, destructive actions, deployment changes, and backup failures.
- Add security regression tests for tenant scope, renderer inputs, storage policy, deployment identity, and failure modes.
- Document incident roles, containment actions, credential rotations, evidence preservation, and communication ownership.

Phase 3: strategic controls

- Replace platform-wide internal tokens with project-scoped credentials and independent rotation.
- Evaluate database row-level security as defense in depth for the highest-value tenant tables.
- Schedule periodic restore drills and destructive-incident tabletop exercises.
- Re-run the threat model after new public routes, privileged roles, storage systems, queues, authentication providers, or trust-boundary changes.
- Create reusable secure-engineering standards from the engagement: authorization, container hardening, service identities, deployment provenance, audit events, and recovery.

Ownership and release gate

Control	Accountable	Responsible	Release gate	Independent verification
Recovery system	Platform owner	Operations	Yes	Security assessor / restore witness
Deployment provenance	Platform owner	Operations	Yes	Security assessor
Edge trust and exposure	Operations owner	Operations	Yes, if public	Security assessor
Tenant authorization	Engineering owner	Backend engineer	Yes	Regression tests + assessor
Renderer input safety	Engineering owner	Backend / Frontend	Yes, for root-cause fix	Security assessor
Least-privilege data identities	Engineering owner	Backend / DBA / Operations	Yes	Negative privilege tests
Telemetry and incident readiness	Platform owner	Operations	Recommended before launch	Tabletop / detection test

Risk acceptance

An unresolved finding was not considered “accepted” merely because the team intended to address it later. A formal acceptance would require a named owner, business justification, compensating controls, residual likelihood and impact, an approval date, an expiry date, and review triggers.

> **RELEASE PRINCIPLE**
>
> No release-blocking risk should be marked closed because a patch exists in source. Closure requires deployment identity, positive and negative tests, adjacent-regression review, and recorded verification evidence.

The roadmap defines required security behaviour, recommended implementation direction, acceptance criteria, and verification procedures. It does not represent completed engineering work. Each owner remains responsible for adapting the proposal to the effective production environment and producing the evidence required for closure.