



(EN) Open Sovereign Predictive Cyber Risk Management (OSPCRM) v1.0

Standard Definition | Rev 1.0 | Nov 2025



0. Introduction (Informative)

Modern cybersecurity programs are overloaded with alerts and vulnerabilities, yet still fail on the most critical incidents. Traditional severity-based models such as raw CVSS scores treat risk as a static property of individual vulnerabilities, without sufficient regard for the organization's own assets, context, threat landscape, or sovereignty constraints.

This specification defines a **contextual, predictive and sovereign-aware cyber risk model**, designed to:

- Prioritize work on **real, organization-specific risk**.
- Minimize **false negatives** on real incidents.
- Integrate **cloud, IoT/OT, AI/ML, and sovereignty** constraints.
- Be **easy to adopt** without forcing vendors to abandon existing tools or models.

The OSPCRM standard is **open and royalty-free to implement**. It is intended for organizations of all sizes, cybersecurity vendors, and integrators who wish to align with a measurable, modern and sovereignty-aware risk model.



1. Scope (Normative)

This specification defines requirements for:

1. A **sovereign predictive cyber risk model**, including:
 - Model design and documentation.
 - Mandatory contextual data classes.
 - Risk scoring and prioritization buckets.
2. Integration of the risk model into:
 - Vulnerability management.
 - Security operations.
 - DevSecOps and change management.
3. Domain-specific extensions for:
 - **Cloud environments** (IaaS, PaaS, SaaS).
 - **IoT/OT and edge systems**.
 - **AI/ML systems and models**.
 - **Data and infrastructure sovereignty**.
 - **Third-party and supply chain security**.
4. **Measurement, validation, and continuous improvement** of the model, based on real incident data.

This specification applies both to:

- **Organizations** implementing the model internally.
- **Products and services** (e.g. scanners, risk engines, platforms) that claim conformance by:
 - Producing OSPCRM-compatible risk outputs.
 - Supporting necessary contextual inputs and metrics.



2. Normative References (Informative)

This specification is designed to be compatible with and complementary to:

- ISO/IEC 27001 – Information security management systems.
- ISO/IEC 27005 – Information security risk management.
- ISO/IEC 22301 – Business continuity management systems.
- NIST Cybersecurity Framework (CSF).
- Industry vulnerability scoring schemas (e.g., CVSS, EPSS).

No external document is required to implement the core of this specification.

3. Terms and Definitions (Normative)

For the purposes of this document, the following terms apply:

- **Asset** – Any system, application, device, data set, or service that has value for the organization.
- **Contextual data** – Information describing the asset, environment, controls, threats, and consequences relevant to a vulnerability.
- **Risk model** – The explicit method by which input data (vulnerabilities, context, threat information) is transformed into risk scores and priority buckets.
- **False Negative (FN)** – A vulnerability or condition that materially contributed to a security incident but was **not** classified in a high priority bucket (P0 or P1) by the



model before the incident.

- **False Negative Rate (FNR)** – The proportion of incident-related vulnerabilities that were not classified as high-priority before the incident.
- **Ground truth incident** – A security incident confirmed by the organization, with identified root causes and associated vulnerabilities or misconfigurations.
- **Sovereignty** – The set of legal, political, cultural, and strategic constraints that govern where and how data, infrastructure, and control planes may be located and operated.
- **Domain** – A specific technical context such as cloud, IoT/OT, AI/ML, or on-premises infrastructure.

Normative keywords:

- **SHALL** indicates a mandatory requirement.
 - **SHOULD** indicates a recommended requirement.
 - **MAY** indicates an optional or permissive element.
-

4. Objectives and Principles (Normative)

4.1 Objectives

The objectives of OSPCRM are:

1. To reduce **false negatives** on real incidents to the lowest feasible level.
2. To prioritize remediation based on **contextual, organization-specific risk**, not solely on generic severity scores.
3. To integrate **cloud, IoT/OT, AI/ML, and sovereignty** into a unified risk model.
4. To minimize **adoption friction** by allowing compatibility with existing tools and vendor models.



4.2 Principles

The risk model defined in this specification is based on the following principles:

- 1. Context over raw severity**

Risk is a function of vulnerability, asset, environment, threat, time, controls, and consequences. Raw CVSS or similar scores are insufficient on their own.
- 2. Incidents as the judge**

The model is evaluated primarily against **real incidents** (ground truth), not theoretical assumptions.
- 3. Sovereign calibration**

Risk assessments are calibrated on the organization's own context, legal obligations, and sovereignty requirements.
- 4. Transparency or measurable performance**

Vendors and tools MAY keep their internal algorithms proprietary, provided they:

 - Accept required contextual inputs, and
 - Demonstrably meet the performance metrics defined in this specification.
- 5. Incremental adoption**

Organizations SHALL be able to implement the standard in stages without disrupting existing operations.



5. Governance and Organizational Context (Normative)

5.1 Roles and Responsibilities

The organization SHALL assign:

- A **Risk Model Owner** (e.g. CISO, CRO) responsible for:
 - Approving the design of the risk model.
 - Ensuring alignment with business and sovereignty objectives.
- A **Risk Model Operator** (e.g. security engineering or GRC function) responsible for:
 - Implementing and maintaining the model, data pipelines, and tooling.
 - Producing regular risk reports.
- **Domain Owners** (e.g. Cloud Lead, OT Lead, AI Lead) responsible for:
 - Ensuring domain-specific context and incident data is integrated into the model.

5.2 Policy

The organization SHALL maintain a **Sovereign Cyber Risk Policy** that:

- States adoption of a contextual predictive risk model aligned with this specification.
 - Defines the **priority buckets** (P0, P1, P2, P3, etc.) and associated remediation SLAs.
 - Specifies sovereignty classifications and constraints for data and infrastructure.
-



6. Risk Model Requirements (Normative)

6.1 Model Documentation

The organization SHALL maintain a **Risk Model Design Document** that includes:

1. The list of **input data classes** (see 7.2).
2. The **scoring scale** (e.g. 0–1000) and mapping to priority buckets.
3. The **weighting or transformation logic** for each context class (high-level description if vendor-proprietary).
4. Treatment of **time** (e.g. time since discovery, patch availability, exploit maturity).
5. Criteria for **P0** and **P1** classification, including target FNR and precision.

The document SHALL be reviewed at least **annually** and after major incidents.

6.2 Risk Scale and Buckets

The risk model SHALL:

1. Produce a **continuous or high-resolution** risk score (e.g. integer 0–1000) to avoid clustering.
2. Map scores to at least **five** priority buckets:
 - P0 – Critical
 - P1 – High
 - P2 – Medium
 - P3 – Low
 - P4 – Informational
3. Define **P0** and **P1** thresholds such that:
 - **At least 95%** of incident-related vulnerabilities are expected to be in P0 or P1 at the time of incident (once the model is mature).



Pan African Strategic Council | (OSPCRM) v1.0

- Bucket distributions remain operationally manageable (e.g. P0+P1 do not routinely exceed a defined proportion of total findings).

Organizations MAY choose different numerical scales, provided the principles above are respected.

6.3 Coexistence with Traditional Scores

The model SHALL:

- Preserve existing severity indicators (e.g., CVSS, EPSS) as **secondary attributes**.
- Prevent traditional scores from being the **sole determinant** of remediation priority.

Existing tools and workflows MAY continue to use CVSS/EPSS fields for reporting and compatibility, provided priority decisions are aligned with the contextual risk model.

7. Data and Context Requirements (Normative)

7.1 Asset Inventory

The organization SHALL maintain an **asset inventory** covering:

- On-premises systems.
- Cloud resources (IaaS, PaaS, SaaS).
- IoT/OT and edge devices.
- AI/ML systems and data pipelines.

Each asset SHALL have at least:

- Unique identifier.
- Owner or accountable contact.
- Environment (production, non-production).
- Business criticality (e.g. 1–5 or equivalent).



- Data classification (including personal/sensitive/sovereign attributes).

7.2 Mandatory Context Classes

For each vulnerability or misconfiguration, the model SHALL, at minimum, consider the following context classes:

1. **Asset Criticality** – importance of the impacted asset to core business or safety.
2. **Exposure** – internal vs external, network reachability, authentication, segmentation.
3. **Threat Activity** – known exploit activity, campaigns, weaponization indicators.
4. **Temporal Factors** – age of vulnerability, patch release date, time in environment.
5. **Resilience Controls** – presence and effectiveness of controls (EDR, WAF, backups, segmentation, monitoring).
6. **Consequence** – estimated impact in terms of confidentiality, integrity, availability, safety, legal liability, and sovereignty.

The organization SHALL ensure that these context classes are **captured and stored** in a way that can be joined with vulnerability data (e.g. relational database, data lake, or other).

7.3 Additional Recommended Context

The model SHOULD also consider:

- **Dependency relationships** (upstream/downstream services).
 - **Usage patterns** (e.g. transaction volume, user sensitivity).
 - **Business calendar** (e.g. peak periods, maintenance windows).
-



8. Domain-Specific Requirements (Normative)

8.1 Cloud (IaaS / PaaS / SaaS)

The organization SHALL:

1. Include all cloud resources in the asset inventory with:
 - Account/subscription ID, region, service type.
 - Exposure (public, private, hybrid).
 - Data classification and sovereignty level.
2. Integrate cloud-specific factors into the risk model, including:
 - Misconfigurations (e.g., open storage, permissive IAM, exposed management interfaces).
 - Shared-responsibility boundaries (provider vs customer responsibilities).
 - Multi-tenant risks (co-hosting with other customers).
3. Ensure that risk scoring for cloud vulnerabilities reflects:
 - Internet-exposed services as higher exposure.
 - Critical control-plane components (e.g. IAM, KMS) as higher consequence.



8.2 IoT / OT / Edge

The organization SHALL:

1. Maintain an inventory of IoT/OT devices, including:
 - Purpose and location.
 - Safety-criticality (can it impact human safety or essential services?).
 - Vendor support status (supported, EOL).
2. Integrate IoT/OT-specific factors into the model:
 - Network segmentation and isolation level.
 - Update/patch mechanisms and feasibility.
 - Physical access and tamper resistance.
3. Treat **unsupported or unpatchable IoT/OT devices** as:
 - Persistent high risk **unless** strong isolation and compensating controls are demonstrably in place.

8.3 AI / ML Systems

The organization SHALL:

1. Register significant AI/ML systems as assets, including:
 - Model type (LLM, classifier, etc.).
 - Training data sources and sensitivity.
 - Deployment mode (internal, external API, embedded in product).



2. Consider AI-specific risks in the model:
 - Input/inference risks (prompt injection, data exfiltration, jailbreaking).
 - Training risks (data poisoning, misuse of sensitive data).
 - Decision-impact risks (use of AI in high-impact decisions such as health, finance, safety).
3. Treat documented AI-related security incidents as **ground truth incidents** and include their root causes in the FNR analysis.

8.4 Sovereignty and Jurisdiction

The organization SHALL:

1. Define **sovereignty levels** for data and systems (e.g., “local-only”, “regional-only”, “globally allowed”).
2. Map each asset to:
 - The jurisdictions in which its data is stored and processed.
 - The jurisdictions controlling its critical control planes (e.g. identity, encryption keys).
3. Use sovereignty as a **risk amplifier** where:
 - Sensitive or sovereign data is hosted in non-approved jurisdictions.
 - Critical control planes are outside declared sovereign boundaries.
4. Document any **accepted sovereignty exceptions**, including:
 - Business justification.



- Compensating controls.
- Review and expiry dates.

8.5 Third Parties and Supply Chain

The organization SHALL:

1. Identify critical suppliers and service providers, especially:
 - Security tools (scanners, SIEM, risk platforms).
 - Cloud and SaaS providers.
 - Managed security services.
2. Require that critical security suppliers either:
 - Provide **OSPCRM-compliant outputs** (risk scores, buckets, metrics), or
 - Allow integration of contextual data and support export of raw findings for transformation by the organization's own risk engine.
3. Include supplier-related incidents in the ground truth dataset and FNR metrics where relevant.

9. Ground Truth and Model Evaluation (Normative)

9.1 Incident Dataset

The organization SHALL maintain a **ground truth incident dataset** including:



Pan African Strategic Council | (OSPCRM) v1.0

- Description and dates of each confirmed incident.
- Assets and vulnerabilities involved.
- Priority bucket assigned to those vulnerabilities **at the time**, according to the model.

9.2 False Negative Rate and Precision

At least **quarterly**, the organization SHALL compute:

- **FNR (False Negative Rate):**
 - Number of incident-related vulnerabilities that were **not** in P0/P1 divided by total number of incident-related vulnerabilities.
- **P0 Precision:**
 - Number of P0 items that correspond to real incidents or near-misses divided by total number of P0 items in that period.

9.3 Target Levels

Organizations SHOULD adopt progressive targets:

- **Level 1:** $FNR \leq 30\%$ within 12 months of adoption.
- **Level 2:** $FNR \leq 10\%$.
- **Level 3:** $FNR \leq 1\%$.

Mature implementations MAY aim for **FNR = 0%** over an extended period for all confirmed incidents.

Whenever FNR exceeds the target, the organization SHALL:



- Analyze causes (missing data, faulty weights, new threat patterns).
 - Update the model design and/or data collection accordingly.
-

10. Operations Integration (Normative)

10.1 ITSM and Remediation

The organization SHALL:

1. Map **priority buckets** directly to remediation SLAs (e.g., P0 in 48h, P1 in 7 days).
2. Configure ITSM tools so that:
 - Tickets incorporate contextual risk score and bucket.
 - Reports to management are based primarily on P0/P1 volumes and SLA adherence, not raw counts of “Critical” CVSS.

10.2 SOC and Monitoring

Security operations SHALL:

1. Use the risk model to focus detection, monitoring, and hunting activities on:
 - P0 and P1 assets and vulnerabilities.
 - High-risk domains (e.g. exposed cloud services, safety-critical OT, AI decision systems).
2. Feed detected incidents and near misses back into the ground truth dataset.



10.3 DevSecOps and CI/CD

DevSecOps pipelines SHALL:

1. Use OSPCRM risk scores to:
 - Enforce quality gates (e.g., block deployment if new P0 items appear on critical services).
 - Prioritize findings in code, container, and IaC scans.
 2. Avoid using raw CVSS/EPSS thresholds as the sole gating condition.
-

11. Conformance (Normative)

11.1 Organizational Conformance

An organization MAY claim **OSPCRM-Conformant (Organization)** if it:

1. Has assigned roles and approved policies as per clause 5.
2. Maintains an asset inventory and contextual data as per clause 7.
3. Operates a risk model documented as per clause 6, covering the mandatory domains in clause 8.
4. Computes FNR and P0 precision at least quarterly as per clause 9.
5. Integrates the model into ITSM/SOC/DevSecOps as per clause 10.
6. Publishes or maintains internally an **annual sovereign cyber risk report** summarizing metrics, issues, and improvements.



11.2 Product/Service Conformance

A product or service MAY claim **OSPCRM-Compatible (Product/Service)** if it:

1. Accepts contextual inputs for the mandatory classes (or provides a documented API to ingest them).
2. Produces:
 - A high-resolution risk score.
 - A mapping to clearly defined priority buckets (P0–P4 or equivalent).
3. Offers:
 - Either a transparent description of its scoring logic, **or**
 - Evidence and APIs that enable customers to measure FNR and P0 precision using their own incident data.

Vendors are **not required** to disclose proprietary algorithms if these conditions are met.



Annex A – Example Control Catalogue (Informative)

Below is an example mapping of controls to this specification. Organizations may extend or adapt as needed.

- **A.1 Sovereign Risk Policy** – A documented policy exists and is reviewed annually.
- **A.2 Asset Context Coverage** – At least 90% of production assets have owner, criticality, data classification, environment tags.
- **A.3 Incident Data Quality** – Ground truth incident dataset is complete and reviewed quarterly.
- **A.4 Model Review** – Risk model design document is updated at least annually and after major incidents.
- **A.5 Domain Coverage** – Cloud, IoT/OT, and AI/ML domains are explicitly included in the model.
- **A.6 Supplier Integration** – Critical suppliers produce OSPCRM-compatible outputs or allow transformation.
- **A.7 Sovereignty Exceptions** – All accepted exceptions are documented with justification and review dates.
- **A.8 Reporting** – Quarterly dashboard for management including FNR, precision, SLA adherence, and sovereignty metrics.

Annex B – Adoption Levels and Migration Path (Informative)

To minimize friction, organizations MAY adopt OSPCRM in stages:



Pan African Strategic Council | (OSPCRM) v1.0

- **Level 1 – Overlay Mode**
 - Keep existing tools and CVSS/EPSS thresholds.
 - Add a contextual enrichment layer and compute contextual risk scores for a subset of critical assets.
 - Start measuring FNR on incidents.
 - **Level 2 – Primary Risk Engine**
 - Use contextual scores as the **primary** basis for prioritization and SLAs.
 - Extend coverage to all production assets.
 - Integrate with ITSM/SOC/DevSecOps.
 - **Level 3 – Full Sovereign Integration**
 - Add domain-specific nuances (cloud, IoT/OT, AI).
 - Embed sovereignty constraints into the risk model.
 - Use metrics and reports for strategic planning and budgeting.
-



Annex C – Implementation Aids (Informative)

Separate, non-normative companion documents MAY provide:

- **Sample data schemas** for assets, vulnerabilities, and context classes.
 - **Example code snippets** (e.g. Python, SQL) for:
 - Computing scores from contextual fields.
 - Joining scanner outputs with asset inventories.
 - Calculating FNR and P0 precision from incident logs.
 - **Template forms and dashboards** for:
 - Sovereign risk policy.
 - Risk model design document.
 - Quarterly sovereign cyber risk reports.
-

Annex D – Mapping to Major Regulations and Standards

(Informative – can be made normative if desired)

D.1 Purpose

This annex shows how an implementation aligned with **OSPCRM / RDC-PRM** supports compliance with key standards and regulations, by providing:

- A **concrete, auditable risk methodology** (contextual scoring, P0–P4).
- A **performance proof** via **False Negative Rate (FNR)** and **P0/P1 precision**.
- A **sovereignty-aware context model** (data location, control-plane jurisdiction, critical infrastructure).



Pan African Strategic Council | (OSPCRM) v1.0

The message is simple:

*These frameworks say “manage and prove your risk control effectiveness”.
OSPCRM/RDC-PRM is the “how” that makes this concrete and measurable.*

D.2 Mapping to ISO / NIST Frameworks

Table D.2.1 – ISO/IEC 27001, 27005, ISO 31000, NIST CSF

| Reference | Key Requirement | How OSPCRM / RDC-PRM Fulfil It | Clauses in OSPCRM / RDC-PRM |
|-------------------------------|--|--|---|
| ISO/IEC 27001 – 6.1, 8.2, 9.1 | Define a documented risk assessment process, determine treatment, and monitor effectiveness. | The contextual risk model provides an explicit, documented method (score scale, P0–P4, context classes). Effectiveness is measured via FNR and P0 precision , audited regularly. | OSPCRM: 4, 6, 7, 9; RDC-PRM: 4, 6, 9, 10 |



Pan African Strategic Council | (OSPCRM) v1.0

| | | | |
|----------------------|---|--|-------------------------------------|
| ISO/IEC 27005 | Identify, analyse, evaluate and treat information security risks. | OSPCRM/RDC-PRM offer a concrete implementation: risk is computed from Sovereign Context (asset criticality, exposure, controls, consequences) and sorted for Marginal Risk Reduction (which action reduces the most risk per unit cost). | OSPCRM: 6.1–6.2, 7; RDC-PRM: 5–6 |
| ISO 31000 | Risk management must be systematic, dynamic, and evidence-based. | The predict incident → FNR → correction loop turns risk management into an empirical, data-driven cycle. The risk model evolves with observed incidents, not only with static assumptions. | OSPCRM: 4, 9, 10; RDC-PRM: 4, 9, 10 |



| | | | |
|---|---|--|--|
| NIST Cybersecurity Framework (CSF) | Provides Identify/Protect/Detect/Respond/Recover, but not how to rank what to do first. | OSPCRM/RDC-PRM add the missing “ Prioritize ” layer: which controls, assets and findings receive attention first, based on FNR, contextual scores, and marginal risk reduction. | OSPCRM: 6.2, 8, 9; RDC-PRM: 6, 8, 9 |
|---|---|--|--|

D.3 Mapping to Regulations on Data Protection & Operational Resilience

Table D.3.1 – GDPR, NIS2, DORA, National Cyber/Data Laws

| Reference | Key Requirement | How OSPCRM / RDC-PRM Fulfil It | Clauses |
|--------------------------|---|---|--|
| GDPR (EU) Art. 32 | “A process for regularly testing, assessing and evaluating the effectiveness of technical and organisational measures.” | OSPCRM/RDC-PRM define exactly such a process: periodic FNR audit against a Confirmed Incident Dataset , plus documented corrective actions. This is a stronger proof than static “we implemented controls X and Y”. | OSPCRM: 7, 9, 10; RDC-PRM: 5, 9, 10 |



| | | | |
|--|--|--|--|
| NIS2 | Risk management, proportionate technical/organizational measures, incident handling, governance involvement. | The Sovereign Context includes critical services, dependencies and sovereignty level. The prioritization engine ensures resources and controls are focused on the most impactful and likely failures , providing a rational, proportionate approach. | OSPCRM: 5, 7, 8, 9; RDC-PRM: 5, 6, 9 |
| DORA (EU financial sector) | (EU ICT risk management, scenario testing, operational resilience, monitoring of third parties. | OSPCRM/RDC-PRM allow backtesting and synthetic scenarios (e.g., Monte Carlo or replaying historical data) to show how risk treatments would have reduced FNR and incident impact. Third-party risks are integrated in the same model. | OSPCRM: 6.2, 8, 9; RDC-PRM: 6, 9, Annex B-C |
| National data protection and cybersecurity laws (e.g., RDC) | Duty of care: “appropriate technical and organisational measures” and demonstrable risk management. | A program aligned with OSPCRM/RDC-PRM can be positioned as a technical safe harbour : the organisation proves it is using the best-known empirical method (FNR→0, sovereignty-aware context) rather than box-ticking. | OSPCRM: 4–11; RDC-PRM: 1–11 |

D.4 African and Sovereignty Frameworks



Table D.4.1 – AU & National Sovereignty Policies

| Reference | Key Requirement | How OSPCRM / RDC-PRM Fulfil It | Clauses |
|--|---|--|--------------------------------------|
| AU Malabo Convention on Cybersecurity | Protect critical information infrastructure; establish national cyber frameworks. | RDC-PRM acts as a national prioritisation engine : each CNI asset is reflected in the Sovereign Context, and national-level P0/P1 risks are visible, explainable, and auditable. | OSPCRM: 5, 8; RDC-PRM: 1, 5, 6, 9 |
| National digital sovereignty policies | Data localization, control of critical management planes, limitation of strategic dependencies. | Sovereignty indicators (data location, key management, jurisdiction, AI control-plane ownership) are explicit risk amplifiers in the model: violations or weak positions automatically appear as higher risk. | OSPCRM: 8.4, 7.2; RDC-PRM: 3.2, 5 |

D.5 How to Use This Mapping in Practice

- **For auditors/regulators:** use the tables to request specific artefacts from organisations (FNR reports, risk model document, context schemas) instead of vague “show us your risk assessment”.
- **For legal/compliance teams:** use them in legal opinions and DPIAs to argue that OSPCRM/RDC-PRM-based programs clearly meet “regular testing and evaluation of effectiveness”.
- **For vendors:** use them in product whitepapers to demonstrate that “OSPCRM-compatible outputs” help customers satisfy ISO, GDPR, NIS2, DORA, etc.



Annex E – Benchmarking Script (Python) for FNR and Model Comparison

(Informative)

E.1 Purpose

This annex provides a minimal Python example for:

1. Comparing **two or more risk models** (e.g., CVSS-only vs OSPCRM vs RDC-PRM).
2. Calculating, for each model:
 - **False Negative Rate (FNR)** on incident-related vulnerabilities.
 - **Precision** for P0 only, and P0+P1 combined.
3. Producing a simple summary table that can be used in reports.

This is deliberately simple and can be extended with graphs, CI/CD integration, or more sophisticated analytics.

E.2 Recommended Data Format

Assume a CSV file `predictions.csv` with at least:

- `model` – Model name (e.g., "CVSS", "OSPCRM", "RDC-PRM").
- `incident_id` – An identifier for the incident (or blank if never involved in an incident).



Pan African Strategic Council | (OSPCRM) v1.0

- `vuln_id` – Vulnerability/condition ID.
- `is_incident` – 1 if this vulnerability was involved in a confirmed incident; 0 otherwise.
- `priority` – Model's predicted priority ("P0", "P1", "P2", "P3", "P4").

Each row = the view of one model on one vulnerability at the time of assessment.

E.3 Example Python Script

```
import pandas as pd

# Load model predictions

df = pd.read_csv("predictions.csv")

# Define which priorities count as "actionable" for FNR calculation
ACTIONABLE = {"P0", "P1"}

def compute_metrics_for_model(df_model):
    """
    Compute FNR and precision metrics for a given model.
    df_model is the subset of rows for one model.
    """
    # 1. Ground truth: vulnerabilities tied to confirmed incidents
```



Pan African Strategic Council | (OSPCRM) v1.0

```
incident_rows = df_model[df_model["is_incident"] == 1]

if incident_rows.empty:

    # No incidents observed for this model's timeframe

    return {

        "nb_incident_vuln": 0,

        "nb_false_negatives": 0,

        "FNR": None,

        "nb_P0": int((df_model["priority"] == "P0").sum()),

        "nb_P1": int((df_model["priority"] == "P1").sum()),

        "precision_P0": None,

        "precision_P0_P1": None,

    }

total_incident_vuln = len(incident_rows)

# False negatives = incident-related vulnerabilities NOT in P0 or P1

false_negatives =
incident_rows[~incident_rows["priority"].isin(ACTIONABLE)]

nb_fn = len(false_negatives)

fnr = nb_fn / total_incident_vuln

# 2. Precision on P0

p0_rows = df_model[df_model["priority"] == "P0"]
```



```
nb_p0 = len(p0_rows)

if nb_p0 > 0:
    true_p0 = p0_rows[p0_rows["is_incident"] == 1]
    precision_p0 = len(true_p0) / nb_p0
else:
    precision_p0 = None

# 3. Precision on P0+P1 (often more stable)
p01_rows = df_model[df_model["priority"].isin(ACTIONABLE)]
nb_p01 = len(p01_rows)

if nb_p01 > 0:
    true_p01 = p01_rows[p01_rows["is_incident"] == 1]
    precision_p01 = len(true_p01) / nb_p01
else:
    precision_p01 = None

return {
    "nb_incident_vuln": total_incident_vuln,
    "nb_false_negatives": nb_fn,
    "FNR": fnr,
    "nb_P0": nb_p0,
    "nb_P1": int((df_model["priority"] == "P1").sum()),
    "precision_P0": precision_p0,
    "precision_P0_P1": precision_p01,
```



Pan African Strategic Council | (OSPCRM) v1.0

```
}

# Compute metrics per model
results = {}
for model_name, df_model in df.groupby("model"):
    results[model_name] = compute_metrics_for_model(df_model)

# Build a summary table
summary_rows = []
for model_name, m in results.items():
    summary_rows.append({
        "model": model_name,
        "incident_vuln": m["nb_incident_vuln"],
        "false_negatives": m["nb_false_negatives"],
        "FNR": m["FNR"],
        "nb_P0": m["nb_P0"],
        "nb_P1": m["nb_P1"],
        "precision_P0": m["precision_P0"],
        "precision_P0_P1": m["precision_P0_P1"],
    })

summary_df = pd.DataFrame(summary_rows)
print(summary_df.to_string(index=False))
```

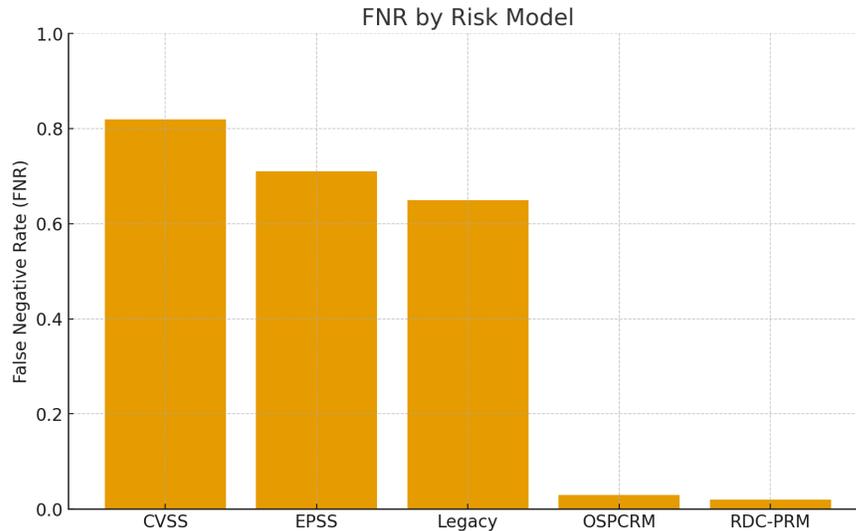


Example output:

| model | incident_vuln | false_negatives | FNR | nb_P0 | nb_P1 | precision_P0 | precision_P0_P1 |
|--------|---------------|-----------------|------|-------|-------|--------------|-----------------|
| CVSS | 50 | 40 | 0.80 | 500 | 300 | 0.05 | 0.08 |
| OSPCRM | 50 | 1 | 0.02 | 60 | 40 | 0.60 | 0.75 |

Annex F - Example Diagrams and Dashboards (as Images)

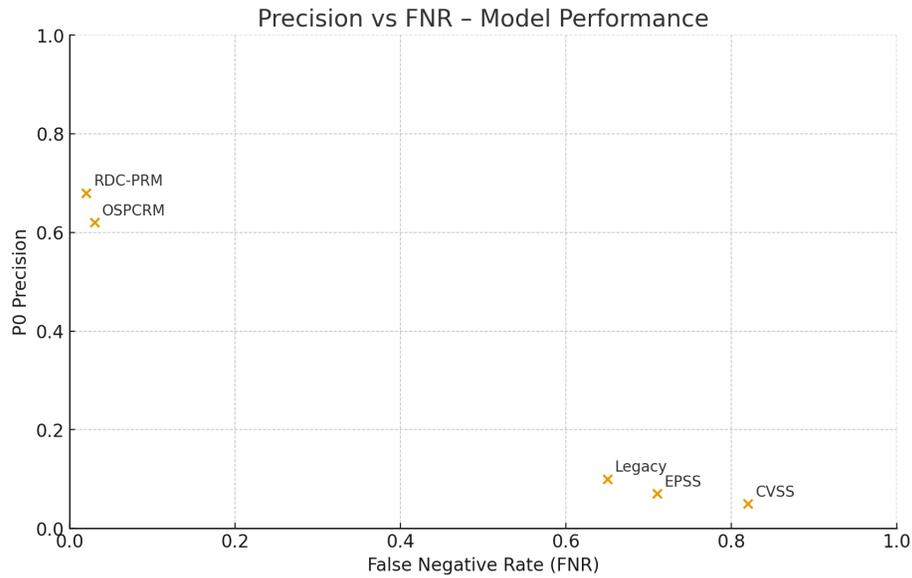
F.1 Figure 1 - FNR Comparison by Model



“False Negative Rate (FNR) by Risk Model”

Figure 1 – Comparison of False Negative Rate (FNR) across different risk models, showing the empirical advantage of OSPCRM/RDC-PRM.

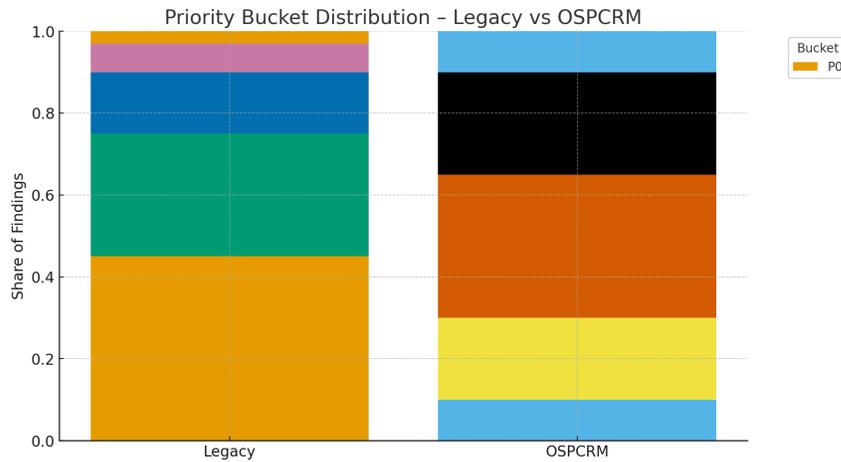
F.2 Figure 2 – Precision vs FNR Trade-off



“Precision vs FNR – Model Performance Landscape”

Figure 2 – Benchmark of different risk models along two axes: FNR and P0 precision. OSPCRM/RDC-PRM occupies the desired low-FNR and high-precision quadrant.

F.3 Figure 3 – Priority Distribution Before vs After OSPCRM

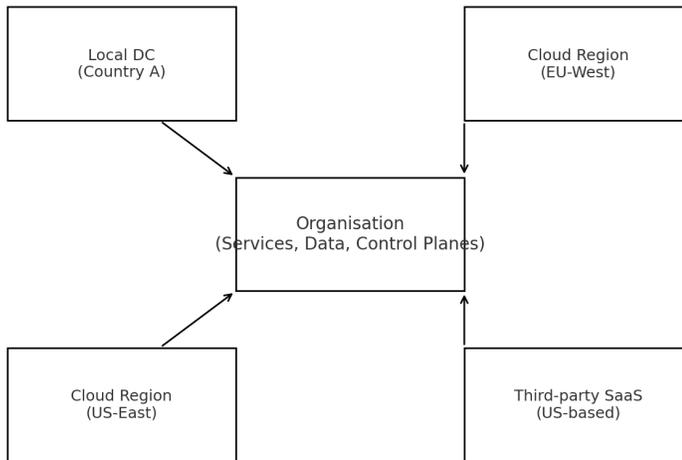


“Priority Bucket Distribution – Legacy vs OSPCRM”

Figure 3 – Distribution of findings by priority bucket before and after implementing a contextual predictive model. OSPCRM eliminates “priority clustering” and creates a meaningful, operational distribution.

F.4 Figure 4 – Sovereign Context Map

Sovereign Context Mapping – Data & Control Jurisdictions



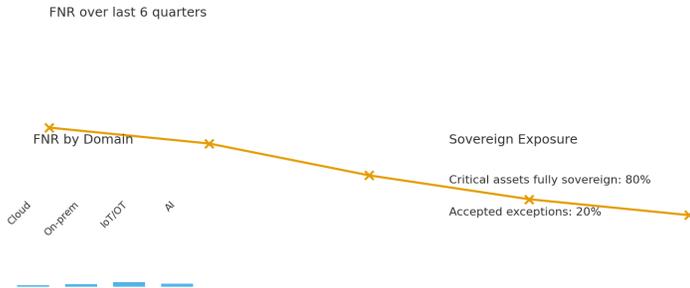
“Sovereign Context Mapping – Data & Control Jurisdictions”

Figure 4 – Example Sovereign Context map showing where sensitive data and critical control planes are located and under which jurisdictions. These factors feed into OSPCRM/RDC-PRM as explicit risk amplifiers.

F.5 Figure 5 – Quarterly FNR Dashboard (Management View)



Sovereign Cyber Risk Dashboard – Quarterly Overview

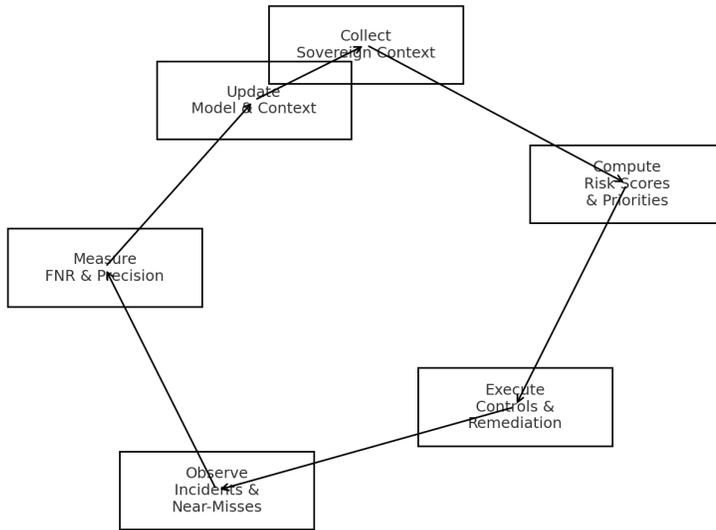


“Sovereign Cyber Risk Dashboard – Quarterly Overview”

Figure 5 – Sample sovereign cyber risk dashboard used for quarterly reporting to senior management, combining FNR, precision, SLA adherence, domain breakdown and sovereignty exposure.

F.6 Figure 6 – Model Feedback Loop

Predictive Risk Management Feedback Loop (OSPCRM/RDC-PRM)



“Predictive Risk Management Feedback Loop (RDC-PRM)”

Figure 6 - High-level view of the predictive, empirical cycle enforced by OSPCRM/RDC-PRM: context → computation → action → incident → FNR audit → model improvement.



Annex G Addendum

This annex clarifies the minor ambiguities identified in real-world pilots and during review with our panel of expert partners.

| R ef | Issue | Quick fix (add/replace this text) | Locati on |
|---------|--|--|-------------------|
| G .1 | EPSS/KEV not explicitly mandatory | §7.2 Mandatory Context Classes → Threat Activity: “Threat Activity SHALL include, at minimum, membership in CISA KEV and the current EPSS score (or equivalent exploit prediction signal).” | §7.2 |
| G .2 | FNR is per-vulnerability, not per-incident | False Negative Rate (FNR) – proportion of confirmed ground-truth incidents for which at least one contributing root-cause vulnerability or misconfiguration was not classified P0 or P1 before the incident. Alternative metric (recommended for maturity Level 2+): Incident Capture Ratio (ICR) = 1 – FNR.” | §3 and §9.2 |
| G .3 | No guidance on statistical validity | §9.3: If fewer than 15 confirmed incidents are available in the measurement window, the organisation SHALL report FNR/ICR with a confidence interval or qualify it as ‘directional only’. | §9.3 |
| G .4 | Precision too noisy on P0 only | §9.2: Organisations SHALL also compute and report Precision@P0+P1 (number of real incidents that have at least one finding in P0 or P1 ÷ total P0+P1 findings in the period). | §9.2 |
| G .5 | Missing auto-remediation expectation | §10.4 Automation: Mature implementations (Level 2+) SHOULD automatically create or escalate tickets, send SLA countdowns, and block non-compliant deployments when a new P0 finding appears on a production asset. | New §10.4 |



G Sovereignty §8.4: “Every approved sovereignty exception SHALL be §8.4
 .6 y recorded with an expiry date ≤ 12 months and
 exceptions automatically generate a time-decaying risk multiplier (e.g.
 ×5 for first 30 days, ×3 until expiry) applied to all findings
 on the affected asset.”

Annex H – OSPCRM-SMB “Light” Profile (Normative – for organisations < 1 000 employees)

Goal: keep 95 % of the security benefit with ~30 % of the effort.

| Clause in full standard | SMB-Light requirement (mandatory = SHALL-Light) | Rationale & allowed simplifications |
|-------------------------|---|-------------------------------------|
| 5 Governance | One person (usually the CISO or MSP) can be both Risk Model Owner and Operator. A formal policy is still required but may be 1–2 pages. | SMBs rarely have separate CRO/CISO. |
| 6.1 Model Documentation | One-page “Risk Model Card” is sufficient (score formula can be high-level, e.g. “Risk = Asset Criticality × Exposure × Threat × (1 – Controls)”). | No need for 20-page design doc. |
| 6.2 Buckets | Minimum three buckets is acceptable: P0 (fix < 72 h), P1 (fix < 30 days), P2 (everything else). | Fewer buckets = less noise. |
| 7.1 Asset Inventory | Excel/Google Sheet with ≥ 80 % of production assets tagged with: Owner, Criticality (High/Medium/Low), Environment (Prod/Non-prod), Contains PII or sovereign data (Y/N). | No CMDB required. |



Pan African Strategic Council | (OSPCRM) v1.0

| | | |
|-------------------------------|--|--|
| 7.2 Mandatory Context Classes | Only four classes are mandatory for SMB-Light: 1. Asset Criticality (1–3) 2. Exposure (Internet-facing Y/N) 3. Threat Activity (is it in CISA KEV or EPSS > 0.9 ?) 4. Controls (is EDR + MFA + backup present?). The two others (Temporal, Consequence) become SHOULD. | Covers ~90 % of real breaches with minimal data collection. |
| 8 Domain-specific | Cloud and third-party sections remain mandatory. OT and AI/ML sections become optional unless the SMB actually operates safety-critical OT or customer-facing GenAI. | Most SMBs are cloud + SaaS only. |
| 9 Ground Truth & FNR | Quarterly FNR calculation remains mandatory, but: • Minimum dataset = last 12 months or 5 incidents (whichever is greater). • If < 10 incidents total ever, the organisation reports “Incident Capture Ratio” instead of FNR and aims for 100 % capture of the last 5 incidents. | Prevents statistical nonsense in small environments. |
| 11 Conformance | An SMB can claim “OSPCRM-SMB Conformant” if it meets all SHALL-Light items above and publishes an annual one-page “Cyber Risk Posture” summary (FNR/ICR trend, % assets tagged, top 5 P0 themes). | Gives small companies a credible label without big-company overhead. |



FR: OSPCRM v1.0 (Open Sovereign Predictive Cyber Risk Management)

**OSPCRM v1.0 - Norme Ouverte de Gestion Prédicative Souveraine du
Risque Cyber**



0. Introduction

Les modèles traditionnels de priorisation (par ex. CVSS “Critical/High/Medium/Low”) sont simples mais ont démontré des **taux de faux négatifs élevés** face aux incidents réels, et créent des “clusters” de priorités qui empêchent une vraie priorisation.

L’OSPCRM définit un **modèle de risque contextuel, prédictif et sensible à la souveraineté**, conçu pour :

- Réduire au maximum les **faux négatifs** sur les incidents réels.
- Prioriser les actions sur la base d’un **risque spécifique à l’organisation** (contexte, actifs, menaces).
- Intégrer explicitement **cloud, IoT/OT, IA/ML et souveraineté**.
- Rester **neutre vis-à-vis des outils** et compatible avec les standards existants (ISO, NIST, etc.).

La norme est **ouverte et libre d’implémentation** par les organisations et les fournisseurs.

1. Objet et domaine d’application

1.1 Cette norme spécifie les **principes** et **exigences** pour concevoir, exploiter et améliorer un **modèle prédictif de gestion du risque cyber**.

1.2 Elle s’applique à toute organisation (publique ou privée), quel que soit le secteur ou la taille.

1.3 Elle couvre :

- Le **modèle de risque** (design, données d’entrée, échelle, seuils).
- L’intégration dans la gestion des vulnérabilités, le SOC, l’ITSM et le DevSecOps.



Pan African Strategic Council | (OSPCRM) v1.0

- Des exigences spécifiques pour **cloud, IoT/OT, IA/ML, souveraineté, chaîne d’approvisionnement**.
 - La **validation empirique** via des incidents réels (FNR).
-

2. Références (informative)

L’OSPCRM est compatible avec, et sert de “moteur de risque” pour :

- ISO/IEC 27001 (ISMS).
- ISO/IEC 27005 (gestion du risque).
- ISO 31000 (principes généraux de gestion du risque).
- NIST Cybersecurity Framework (fonctions et contrôles).

Aucune de ces références n’est obligatoire pour mettre en œuvre le cœur de la norme, mais l’OSPCRM est conçu pour les **renforcer, pas les remplacer**.

3. Termes et définitions (extraits)

- **Actif** : système, application, appareil, jeu de données ou service ayant de la valeur.
- **Données contextuelles** : informations décrivant l’actif, l’environnement, les contrôles, les menaces et les conséquences liées à une vulnérabilité.
- **Modèle de risque** : méthode explicite par laquelle les données d’entrée sont transformées en **scores de risque** et en **niveaux de priorité**.



Pan African Strategic Council | (OSPCRM) v1.0

- **Faux négatif (FN)** : vulnérabilité ou condition qui a contribué à un incident, mais qui n'était pas classée en priorité **actionnable** (P0/P1) avant l'incident.
 - **Taux de faux négatifs (FNR)** : proportion de vulnérabilités liées à des incidents qui n'étaient pas en P0/P1.
 - **Incident de référence (ground truth)** : incident confirmé, investigué, documenté.
 - **Souveraineté** : contraintes juridiques, politiques et stratégiques sur les lieux et modalités de stockage/traitement des données et des plans de contrôle.
-

4. Objectifs et principes

4.1 Objectifs

- Réduire le **FNR** au niveau le plus bas possible.
- Remplacer la logique " $CVSS \geq X = critique$ " par un modèle **contextuel, prédictif**.
- Couvrir l'ensemble des domaines techniques (cloud, IoT/OT, IA, on-prem).
- **Limiter la friction** avec les outils et fournisseurs existants.

4.2 Principes

1. **Contexte > sévérité brute** : le risque n'est pas un label, mais une fonction de la vulnérabilité, de l'actif, de l'exposition, de la menace, du temps, des contrôles et des conséquences.
2. **L'incident est le juge** : la performance du modèle se mesure principalement sur les **incidents réels** (FNR, précision P0).



3. **Calibration souveraine** : le modèle est calibré sur **le contexte propre de l'organisation**, pas sur des moyennes globales opaques.
 4. **Transparence ou performance mesurable** : les fournisseurs peuvent garder un algorithme propriétaire, mais doivent accepter les données contextuelles et permettre la mesure du FNR par le client.
 5. **Adoption progressive** : la norme doit être **implémentable par étapes**, sans “big bang”.
-

5. Gouvernance et politique

- Nommer un **Propriétaire du modèle de risque** (CISO/CRO).
 - Nommer des **Responsables de domaines** (cloud, OT, IA...).
 - Adopter une **Politique de risque cyber souverain** qui :
 - entérine le recours à un modèle contextuel prédictif,
 - définit les **seuils P0/P1** et leurs SLA,
 - définit les niveaux de **souveraineté** des données et systèmes.
-

6. Exigences sur le modèle de risque

6.1 Documentation du modèle

La documentation DOIT inclure :

- Les **classes de données d'entrée** (cf. 7.2).



Pan African Strategic Council | (OSPCRM) v1.0

- L'**échelle de score** (p.ex. 0–1000) et la cartographie vers P0–P4.
- La façon dont chaque classe contextuelle influence le score (poids, amplificateurs, atténuateurs).
- Le traitement du **temps** (ancienneté, cycle de vie, maturité des exploits).

Revue au moins **annuelle** et après tout incident majeur.

6.2 Échelle et priorités

- Le modèle DOIT générer un score suffisamment **granulaire** pour éviter le “tout critique”.
- Au minimum, 5 seaux : **P0, P1, P2, P3, P4**.
- Les seuils P0/P1 DOIVENT être définis et justifiés (par ex. “P0 destiné à capter ≥ 95 % des vulnérabilités impliquées dans des incidents”).

6.3 Coexistence avec CVSS/EPSS

- Les scores CVSS/EPSS sont **conservés comme attributs secondaires**.
 - Ils NE DOIVENT PAS être le seul déterminant pour la priorisation et les SLA.
-

7. Données et contexte

7.1 Inventaire des actifs

Inventaire couvrant :

- On-prem, cloud, IoT/OT, IA/ML.



Chaque actif doit avoir :

- ID unique, propriétaire, environnement (prod/non prod), **criticité métier**, **classification des données**, éventuellement niveau de souveraineté.

7.2 Classes contextuelles obligatoires

Pour chaque vulnérabilité / condition :

1. **Criticité de l'actif.**
 2. **Exposition** (internet, interne, segmenté, air-gapped...).
 3. **Activité de menace** (exploits observés, campagnes, exploitation active).
 4. **Temporalité** (âge, date de patch, temps de présence).
 5. **Contrôles / résilience** (EDR, sauvegardes, segmentation, monitoring).
 6. **Conséquence** (impact sur CIA, sécurité des personnes, régulation, souveraineté).
-

8. Exigences spécifiques par domaine

8.1 Cloud

- Tous les comptes/projets/tenants cloud DOIVENT être inventoriés (ID, région, type de service).
- La priorité DOIT tenir compte de :
 - l'exposition publique,
 - la sensibilité des données,
 - le rôle de contrôle (IAM, KMS...).



8.2 IoT / OT

- Inventaire des dispositifs avec criticité “sécurité des personnes / CNI”.
- Actifs EOL ou non patchables : **risque élevé par défaut**, sauf preuve d’isolation forte.

8.3 IA / ML

- Les systèmes IA/ML significatifs DOIVENT être enregistrés (type de modèle, données d’entraînement, mode de déploiement).
- Les risques IA DOIVENT inclure : injection de prompts, exfiltration, empoisonnement de données, décisions à fort impact.

8.4 Souveraineté

- Définir des **niveaux de souveraineté** pour les données/systèmes.
- Cartographier les **juridictions** de stockage/traitement + des plans de contrôle (IAM/KMS).
- Utiliser ces facteurs comme **amplificateurs de risque** lorsque des contraintes souveraines sont violées ou contournées.

9. Incidents, FNR et validation empirique

- Maintenir un **jeu de données d’incidents confirmés** (ground truth).
- Au moins trimestriellement, calculer :
 - FNR = % de vulnérabilités liées à des incidents qui n’étaient pas en P0/P1.
 - **Précision P0** = % de P0 qui correspondent à des incidents / quasi-incidents.

Des cibles progressives peuvent être adoptées (FNR \leq 30 %, puis 10 %, puis 1 %).



10. Intégration opérationnelle

- ITSM : les tickets doivent utiliser le score contextuel et le niveau P0–P4 comme base de priorité.
 - SOC : orienter la détection et la chasse sur les actifs et vulnérabilités P0/P1.
 - DevSecOps : utiliser les scores OSPCRM comme **gates** dans les pipelines (bloquer un déploiement si un nouveau P0 apparaît sur un service critique).
-

11. Conformité

Une **organisation** peut se déclarer conforme si elle :

- A mis en place les rôles, la politique, l’inventaire et les classes contextuelles.
- Exploite un modèle de risque documenté et appliqué.
- Mesure FNR et précision P0 régulièrement.
- Alimente l’ITSM/SOC/DevSecOps avec ce modèle.

Un **produit/service** peut se dire “compatible OSPCRM” s’il :

- Peut ingérer des données contextuelles.
- Produit un score et des seuils P0–P4.
- Permet au client de mesurer FNR & précision P0 sur ses incidents.



Annexe D - Tableau de correspondance OSPCRM / RDC-PRM avec les principales normes et régulations

D.1 Logique générale

Cette annexe montre comment un système conforme à **OSPCRM / RDC-PRM** permet de répondre, de manière **empirique et mesurable**, aux exigences principales de :

- ISO/IEC 27001 & 27005
- ISO 31000
- NIST Cybersecurity Framework (CSF)
- Règlement (UE) 2016/679 – RGPD (Art. 32)
- Directive NIS2
- Règlement DORA (finance UE)
- Convention de Malabo (UA)
- Lois nationales sur la protection des données et la cybersécurité (ex. RDC)

L'idée explicite : **OSPCRM/RDC-PRM fournit le “comment” opérationnel** qui manque dans ces textes (le moteur de risque + la preuve FNR).

D.2 Tableau de correspondance (exemple condensé)

Tableau D.2.1 – Normes de gestion (ISO / NIST)

| Référence | Exigence clé | OSPCRM / RDC-PRM – Mécanisme | Clauses de référence |
|---|---|--|---|
| ISO/IEC 27001 – Clause 6.1, 8.2, 9.1 | Mettre en place un processus documenté d’appréciation du risque, décider des traitements, mesurer l’efficacité. | Le modèle de risque contextuel définit une methodologie explicite (échelle 0–1000, P0–P4, classes de contexte). L’efficacité est mesurée par FNR et précision P0 , audités périodiquement. | OSPCRM : 4, 6, 7, 9 – RDC-PRM : 4, 6, 9, 10 |
| ISO/IEC 27005 | Identification, analyse, évaluation et traitement du risque. | OSPCRM/RDC-PRM fournit une implémentation concrète : calcul du score à partir du Contexte Souverain, puis tri par Réduction Marginale du Risque (MMR). | OSPCRM : 6.1, 6.2, 7 – RDC-PRM : 5, 6 |
| ISO 31000 | Approche systémique et dynamique de la gestion du risque. | La boucle prédiction → incident → FNR → correction transforme la gestion du risque en processus empirique continu. | OSPCRM : 4, 9, 10 – RDC-PRM : 4, 9, 10 |
| NIST CSF (Identify, Protect, Detect, Respond, Recover) | Fournit des fonctions et catégories, mais pas la priorisation. | OSPCRM/RDC-PRM ajoute la fonction implicite “Prioritize” : quels contrôles/actifs traiter d’abord pour maximiser la réduction de risque. | OSPCRM : 6.2, 8, 9 – RDC-PRM : 6, 8, 9 |

Tableau D.2.2 – Réglementations données & résilience

| Référence | Exigence clé | OSPCRM / RDC-PRM – Mécanisme | Clauses de référence |
|---|---|---|--|
| RGPD – Art. 32 | Mettre en place des mesures appropriées et un processus d'évaluation régulière de l'efficacité des mesures techniques et organisationnelles. | L'organisation prouve l'efficacité de ses mesures en montrant : (1) un modèle de risque explicite, (2) un audit FNR régulier sur les incidents réels, (3) une boucle de correction documentée. | OSPCRM : 7, 9, 10 – RDC-PRM : 5, 9, 10 |
| NIS2 | Gestion de risque, mesures techniques proportionnées, notification des incidents, gouvernance. | Le Contexte Souverain inclut criticité des services essentiels, dépendances et souveraineté ; la priorisation par MMR assure la proportionnalité des mesures (priorité aux services essentiels). | OSPCRM : 5, 7, 8, 9 – RDC-PRM : 5, 6, 9 |
| DORA (finance UE) | Risque TIC, tests, scénarios adverses, résilience opérationnelle. | Le modèle prédictif permet de simuler des scénarios (backtesting, Monte Carlo) et de montrer comment le portefeuille d'actions (patches, contrôles) diminue le FNR et la fréquence des incidents critiques. | OSPCRM : 6.2, 8, 9 – RDC-PRM : 6, 9, Annexes B-C |
| Lois nationales (données, cybersécurité) | “Mesures techniques et organisationnelles appropriées”, “devoir diligence / duty of care”. | Suivre la norme (OSPCRM + profil national RDC-PRM) fournit un argument de “Safe Harbor” : l'organisation démontre qu'elle utilise un modèle empirique moderne (FNR→0) aligné sur les meilleures pratiques. | OSPCRM : 4-11 – RDC-PRM : 1-11 |

Tableau D.2.3 – Cadres souverains



Pan African Strategic Council | (OSPCRM) v1.0

| Référence | Exigence clé | OSPCRM / RDC-PRM – Mécanisme | Clauses de référence |
|---|---|---|---|
| Convention de Malabo (UA) | Protection des infrastructures critiques, mise en place de cadres nationaux de cybersécurité. | RDC-PRM fournit un moteur national de priorisation : chaque CNI est intégrée au Contexte Souverain, les P0/P1 au niveau national sont visibles, justifiables et audités. | OSPCRM : 5, 8 – RDC-PRM : 1, 5, 6, 9 |
| Politiques nationales souveraineté numérique | Localisation des données, maîtrise des plans de contrôle, limitation des dépendances critiques. | Le Contexte Souverain inclut la juridiction des données et des plans de contrôle (IAM, KMS, AI). Les violations ou dépendances non désirées deviennent des amplificateurs de risque (score ↑). | OSPCRM : 8.4, 7.2 – RDC-PRM : 3.2, 5 |

D.3 Mode d'utilisation du tableau

Ce tableau peut être utilisé :

- Par les **auditeurs / régulateurs** : pour relier des exigences légales à des artefacts concrets (FNR, rapports trimestriels, modèle documenté).
- Par les **juristes** : pour argumenter qu'un programme OSPCRM/RDC-PRM répond au principe de "mesures appropriées et efficaces".
- Par les **fournisseurs** : pour montrer que leurs produits, une fois configurés pour sortir des scores OSPCRM, aident les clients à se mettre en conformité.

Annexe E – Exemple de script de benchmark (Python)

(Informative – démonstration simple pour FNR & comparaison de modèles)



E.1 Objectif

Cette annexe propose un **exemple de code Python minimal** permettant de :

1. Comparer deux modèles de risque (p. ex. **CVSS** vs **OSPCRM** ou **RDC-PRM**) sur un même jeu d'incidents.
2. Calculer, pour chaque modèle :
 - le **FNR** (Taux de faux négatifs),
 - la **précision P0/P1**,
 - une petite synthèse pour un rapport.

Ce code est volontairement simple : il peut servir de base pour des scripts internes ou des notebooks de validation.

E.2 Format de données recommandé

On suppose un fichier CSV (par exemple `predictions.csv`) avec les colonnes suivantes :

- `model` : nom du modèle ("`CVSS`", "`OSPCRM`", "`RDC-PRM`", etc.).
- `incident_id` : identifiant de l'incident (ou `NaN` si la vulnérabilité n'a jamais été exploitée dans un incident).
- `vuln_id` : identifiant de la vulnérabilité / condition.
- `is_incident` : `1` si cette ligne correspond à une vulnérabilité impliquée dans un incident confirmé, sinon `0`.



Pan African Strategic Council | (OSPCRM) v1.0

- **priority** : niveau de priorité prédit par le modèle ("P0", "P1", "P2", "P3", "P4").
- (facultatif) **date_prediction, asset_id**, etc.

Chaque ligne représente la **vision d'un modèle** pour une vulnérabilité donnée à un instant donné.



E.3 Script Python d'exemple

```
import pandas as pd

# Charger les données

df = pd.read_csv("predictions.csv")

# Définition des priorités "actionnables" (P0/P1)

ACTIONABLE = {"P0", "P1"}

def compute_metrics_for_model(df_model):
    """
    Calcule FNR, précision P0 et quelques stats de base
    pour un modèle donné (DataFrame déjà filtré).
    """
    # 1. Incidents de référence (ground truth)
    incident_rows = df_model[df_model["is_incident"] == 1]

    # Si aucun incident pour ce modèle, éviter la division par zéro
    if incident_rows.empty:
        return {
            "nb_incidents_vuln": 0,
            "fnr": None,
            "precision_P0": None,
            "nb_P0": int((df_model["priority"] == "P0").sum()),
            "nb_P1": int((df_model["priority"] == "P1").sum())
        }
```



Pan African Strategic Council | (OSPCRM) v1.0

```
}

# On compte les vulnérabilités impliquées dans des incidents

# (on pourrait dédupliquer par vuln_id si nécessaire)

total_incident_vuln = len(incident_rows)

# Faux négatifs : vulnérabilités d'incident qui n'étaient pas P0/P1

false_negatives = incident_rows[~incident_rows["priority"].isin(ACTIONABLE)]

nb_fn = len(false_negatives)

fnr = nb_fn / total_incident_vuln

# 2. Précision sur P0 : parmi toutes les P0, combien sont liées à un incident ?

p0_rows = df_model[df_model["priority"] == "P0"]

nb_p0 = len(p0_rows)

if nb_p0 > 0:

    true_p0 = p0_rows[p0_rows["is_incident"] == 1]

    precision_p0 = len(true_p0) / nb_p0

else:

    precision_p0 = None

# 3. Précision sur P0+P1 (optionnel, souvent plus robuste)

p01_rows = df_model[df_model["priority"].isin(ACTIONABLE)]

nb_p01 = len(p01_rows)
```



```
if nb_p01 > 0:
    true_p01 = p01_rows[p01_rows["is_incident"] == 1]
    precision_p01 = len(true_p01) / nb_p01
else:
    precision_p01 = None

return {
    "nb_incidents_vuln": total_incident_vuln,
    "nb_false_negatives": nb_fn,
    "fnr": fnr,
    "nb_P0": nb_p0,
    "nb_P1": int((df_model["priority"] == "P1").sum()),
    "precision_P0": precision_p0,
    "precision_P0_P1": precision_p01
}

# Regrouper par modèle et calculer les métriques
results = {}

for model_name, df_model in df.groupby("model"):
    results[model_name] = compute_metrics_for_model(df_model)

# Afficher un petit tableau de synthèse
summary_rows = []

for model_name, metrics in results.items():
```



```
summary_rows.append({  
    "model": model_name,  
    "nb_incident_vuln": metrics["nb_incidents_vuln"],  
    "nb_false_negatives": metrics["nb_false_negatives"],  
    "FNR": metrics["fnr"],  
    "nb_P0": metrics["nb_P0"],  
    "nb_P1": metrics["nb_P1"],  
    "precision_P0": metrics["precision_P0"],  
    "precision_P0_P1": metrics["precision_P0_P1"],  
})
```

```
summary_df = pd.DataFrame(summary_rows)  
print(summary_df.to_string(index=False))
```

E.4 Lecture des résultats (pour un rapport)

Le tableau imprimé donnera, par modèle, quelque chose comme :

| model | nb_incident_vuln | nb_false_negatives | FNR | nb_P0 | nb_P1 | precision_P0 | precision_P0_P1 |
|--------|------------------|--------------------|------|-------|-------|--------------|-----------------|
| CVSS | 50 | 40 | 0.80 | 500 | 300 | 0.05 | 0.08 |
| OSPCRM | 50 | 1 | 0.02 | 60 | 40 | 0.60 | 0.75 |

Ce type de tableau permet de :

- **Montrer aux dirigeants/régulateurs** que le modèle OSPCRM/RDC-PRM :



Pan African Strategic Council | (OSPCRM) v1.0

- a un FNR beaucoup plus faible (ex. 2 % vs 80 %),
- et que ses P0/P1 sont beaucoup plus “denses en incidents réels” (précision plus haute).
- Produire une **preuve quantitative** dans un rapport (Annexe au rapport annuel de cybersécurité, rapport FNR trimestriel, etc.).



Annexe G – Addendum

Cette annexe clarifie les ambiguïtés mineures identifiées lors des pilotes en conditions réelles et au cours de la révision avec notre panel de partenaires experts.

| Réf | Problème | Correction rapide (ajouter/remplacer ce texte) | Emplacement |
|-----|--|--|-------------|
| G.1 | EPSS/KEV pas explicitement obligatoires | §7.2 Mandatory Context Classes → Activité de Menace : « L'Activité de Menace DOIT inclure, au minimum, l'appartenance à la CISA KEV et le score EPSS actuel (ou un signal équivalent de prédiction d'exploit). » | §7.2 |
| G.2 | Le FNR est par vulnérabilité, pas par incident | Taux de Faux Négatifs (FNR) – proportion des incidents de référence confirmés pour lesquels au moins une vulnérabilité ou mauvaise configuration contributive de la cause racine n'était | §3 et §9.2 |



Pan African Strategic Council | (OSPCRM) v1.0

| | | | |
|------------|--|---|------|
| | | <p>pas classée P0 ou P1 avant l'incident. Métrique alternative (recommandée pour le niveau de maturité 2+) : Incident Capture Ratio (ICR) = 1 – FNR.</p> | |
| G.3 | Absence de guide sur la validité statistique | <p>§9.3 : Si moins de 15 incidents confirmés sont disponibles dans la fenêtre de mesure, l'organisation DOIT rapporter le FNR/ICR avec un intervalle de confiance ou le qualifier de « directionnel uniquement ».</p> | §9.3 |
| G.4 | Précision trop bruitée sur P0 seulement | <p>§9.2 : Les organisations DOIVENT également calculer et rapporter la Précision@P0+P1 (nombre d'incidents réels qui ont au moins une constatation en P0 ou P1 ÷ nombre total de constatations P0+P1 durant la</p> | §9.2 |



Pan African Strategic Council | (OSPCRM) v1.0

| | | | |
|------------|-------------------------------------|---|---------------|
| | | période). | |
| G.5 | Manque d'attente d'auto-remédiation | §10.4 Automatisation : Les implémentations matures (Niveau 2+) DEVRAIENT créer ou escalader automatiquement des tickets, envoyer des comptes à rebours de SLA et bloquer les déploiements non conformes lorsqu'une nouvelle constatation P0 apparaît sur un actif de production. | Nouveau §10.4 |
| G.6 | Exceptions de souveraineté | §8.4 : « Chaque exception de souveraineté approuvée DOIT être enregistrée avec une date d'expiration ≤ 12 mois et générer automatiquement un multiplicateur de risque à décroissance temporelle (par exemple, ×5 pendant les 30 premiers jours, ×3 jusqu'à l'expiration) | §8.4 |



Pan African Strategic Council | (OSPCRM) v1.0

| | | | |
|--|--|--|--|
| | | appliqué à toutes les constatations sur l'actif concerné. » | |
|--|--|--|--|

-----Annexe H – Profil « Allégé » OSPCRM-PME (Normatif – pour les organisations < 1 000 employés)

Objectif : conserver 95 % du bénéfice sécurité avec environ 30 % de l'effort.

| Clause de la norme complète | Exigence « Allégée » PME (obligatoire = DOIT-Allégé) | Justification et simplifications autorisées |
|------------------------------------|---|---|
| 5 Gouvernance | Une seule personne (habituellement le RSSI ou le MSP) peut être à la fois Propriétaire et Opérateur du Modèle de Risque. Une politique formelle est toujours requise mais peut ne faire qu'1 à 2 pages. | Les PME ont rarement un CRO/RSSI séparé. |
| 6.1 Documentation du Modèle | Une « Fiche de Modèle de Risque » d'une page est suffisante (la formule de score peut être de haut niveau, par exemple « Risque = Criticité de l'Actif × Exposition × Menace × (1 – Contrôles) »). | Pas besoin de document de conception de 20 pages. |
| 6.2 Niveaux de Priorité | Un minimum de trois niveaux est acceptable : P0 (à corriger < 72 h), P1 (à corriger < 30 jours), P2 (tout le reste). | Moins de niveaux = moins de bruit. |



| | | |
|--|---|--|
| <p>7.1 Inventaire des Actifs</p> | <p>Un fichier Excel/Google Sheet avec ≥ 80 % des actifs de production étiquetés avec : Propriétaire, Criticité (Élevée/Moyenne/Faible), Environnement (Prod/Non-prod), Contient des données PII ou souveraines (O/N).</p> | <p>Aucune CMDB requise.</p> |
| <p>7.2 Classes de Contexte Obligatoires</p> | <p>Seulement quatre classes sont obligatoires pour l'Allégé-PME : 1. Criticité de l'Actif (1–3) 2. Exposition (Face à Internet O/N) 3. Activité de Menace (est-ce dans CISA KEV ou EPSS > 0.9 ?) 4. Contrôles (EDR + MFA + sauvegarde sont-ils présents ?). Les deux autres (Temporels, Conséquence) deviennent des RECOMMANDATIONS (DEVRAIT).</p> | <p>Couvre ~90 % des brèches réelles avec une collecte de données minimale.</p> |
| <p>8 Spécifique aux Domaines</p> | <p>Les sections Cloud et Tiers restent obligatoires. Les sections OT et IA/ML deviennent optionnelles, sauf si la PME exploite effectivement de l'OT critique pour la sécurité ou de l'IA Générative orientée client.</p> | <p>La plupart des PME sont uniquement Cloud + SaaS.</p> |



Pan African Strategic Council | (OSPCRM) v1.0

| | | |
|-----------------------------------|---|---|
| 9 Vérité Terrain & FNR | Le calcul trimestriel du FNR reste obligatoire, mais : • Ensemble de données minimal = les 12 derniers mois ou 5 incidents (selon le plus grand nombre). • Si < 10 incidents totaux à ce jour, l'organisation rapporte le « Incident Capture Ratio » (ICR) au lieu du FNR et vise une capture à 100 % des 5 derniers incidents. | Évite les non-sens statistiques dans les petits environnements. |
| 11 Conformité | Une PME peut revendiquer la conformité « OSPCRM-PME » si elle respecte tous les éléments DOIT-Allégé ci-dessus et publie un résumé annuel d'une page sur sa « Posture de Risque Cyber » (tendance FNR/ICR, % d'actifs étiquetés, top 5 des thèmes P0). | Donne aux petites entreprises un label crédible sans la lourdeur des grandes entreprises. |